Another Look at Inflation and Expectations during the U.S. Great Depression

Patrick J. Coe   Shaun P. Vahey
Carleton University   Warwick Business School

October 2014

This is very preliminary and incomplete - please do not quote

Abstract

We look at the extent to which price and output declines could have been predicted during the Great Depression. Rather than focus on point forecasts, we estimate the forecast density and derive probabilities for events of different magnitudes. In this draft we show that the means of the forecast densities from simple autoregressive models do not match the declines in prices and output. However, by late 1929 (output) or early 1930 (prices) the forecast densities imply an increase in the probability of falls in output and prices as large as those that actually occurred. In subsequent drafts we plan to expand our model space to use a richer set of models to generate forecast densities.

1 Introduction

The US Great Depression saw declines in output and levels of unemployment that have not been seen since. Between July 1929 and March 1933 industrial production in the United States fell by an annual average rate of approximately 18%.[1] At the same the Consumer Price Index fell at an annual average rate of approximately 8.3%. The validity or otherwise of a number of candidate explanations for the U.S. Great Depression depend on the extent to which this decline in prices was anticipated by agents. For example, if monetary factors were important the standard channel requires an increase in ex ante real interest rates. However, in general the period from late 1929 to 1933 saw nominal interest rates follow a downward trend. As a result, any increase in the ex ante real interest rate required that to at least some extent the fall in aggregate prices was expected. Romer and Romer (2013) go further and argue that for the monetary contraction channel to have been important any expectations of deflation must be linked to Fed actions. Alternatively the debt-deflation and financial intermediation channels identified by Fisher (1933) and Bernanke (1983) stress the high ex post real interest rates and debt burdens implied by unanticipated deflation. In this case deflation was not anticipated at the time nominal loan contracts were written.

Motivated by this debate, number of studies have looked at the extent to which the deflation of the early 1930s was either predicted by economic agents at the time, or predictable given information available to agents at that time. Hamilton (1987, 1992) uses data from commodities futures markets to extract expectations of future price movements and concludes that these markets were expecting price rises rather than declines. Cecchetti (1992) takes an alternative approach and uses ARIMA models to forecast inflation at the three and six month horizons for 1930-32. He finds that the persistence in the inflation process means that once it had started deflation could have

---

[1]The Industrial Production Index used is the Jan 22nd, 1946 release of the Board of Governors of the Federal Reserve System's series obtained from the St Louis Federal Reserve Economic Database (FRED). The CPI that follows is that of the Bureau of Labor Statistics, and is also obtained from FRED.

be anticipated. In particular he concludes that certainly by late 1930, and maybe as early as late 1929, deflation was forecastable. Evans and Wachtel (1993) cast doubt on Cecchetti's conclusions by arguing that there would have been no guarantee that prices would have continued to follow the same process in the future as they had in the past. Using a Markov switching model to allow agents to incorporate the possibility of a change in regime for the inflation process they conclude that it is unlikely the extent of the deflation of the early 1930s would have been predicted at the time.

Klug et al (2005) also express doubt as to whether forecasts from time series models are representative of agents expectations at the time. They compare the forecast errors of railway freight car loadings published by contemporary railroad shippers to forecasts from time series models that they construct for the same object. Klug et al find that the forecast errors made by the contemporary shippers are much greater than those from the time series models suggesting the forecasters at the time used much more than just the past history of the series. They document that the contemporary forecasts of freight car loadings consistently exceeded the actual numbers for 1929-33, suggesting that it was not just the price declines that surprised agents, but also output declines. This result is consistent with the results of Dominguez at al (1988) who use VAR models to forecast industrial production as well as prices. These VARs include interest rates, money, stock prices and business indicators, but fail to predict the large declines in either output or prices following the stock market crash. Dominguez et al buttress these results with the forecasts of the Harvard Economic Service and Irving Fisher, both of which remained optimistic until late 1931. Romer and Romer (2013) also document that *Business Week* magazine did not expect prolonged deflation until late 1930 and argue that these expectations were driven by perceptions of the Federal Reserve policy.

A common feature of the studies that use econometric models to undertake pseudo real-time forecasting exercises in order to ask whether or not deflation could have been predicted during the early 1930s is that they focus on the point forecast. They ask

whether or not these point forecasts are close to the realized value of the series. When this point forecast implies a decline in prices much less than what actually happened then the interpretation is that the deflation could not have been anticipated. In this paper we take a different approach and look at the whole of the forecast density and not just its mean. In expected utility theory households make decisions based on the probabilities that they assign to a number of different states of the world and the outcomes in those states, not just the outcome in the most likely state of the world as determined by the point forecast. By looking at the whole of the forecast density we are able to generate probability forecasts for deflations of various magnitudes at horizons of one, three and six months throughout the Great Depression. Therefore, rather than judge whether or not the deflation was predictable on the basis of how close the point forecast is to zero, we judge the extent to which an event was predictable based on the probability assigned to it by the forecast density.

Similarly policy makers make decisions that take into account the fact that there are a number of possibilities for the future state of the economy. Friedman and Schwartz (1963) criticize the Federal Reserve for its inaction following the stock market crash and the banking panics of the early 1930s. One interpretation of the literature, which for the most part concludes that output and price declines were not anticipated, is that while the case for more aggressive monetary policy might be apparent with hindsight, the depth of the depression could not have been predicted at the time. As such, a failure to act by the Fed is in some sense understandable. However, as with private agents, policymakers also make decisions with a number of different states of the world possible. Therefore, while a depression of the magnitude of the Great Depression might not have been the most likely outcome immediately after the 1929 stock market crash (for example) our approach allows us to explore whether it is reasonable to think that the probability of such an event increased around this time. If it does, then perhaps the inaction of the Fed is harder to explain.

4

Finally, while Cecchetti shows that some univariate time series models can predict inflation relatively during some stages of the Great Depression, as Klug et al point out it is not necessarily the case that agents would have used such models. It is also possible that they would have used a number of different models and that the weight they put on the predictions of each model would change over time as suggested by Evans and Wachtel. Futhermore, model uncertainty can mean that focusing attention on a single model can understate the true forecast uncertainty. Therefore, rather than confine our attention to a single model, we generate predictions from a number of different models. We then use a linear mixture of experts framework (sometimes referred to as the Linear Opinion Pool, LOP) to combine these into a single forecast. see (among others) Timmermann (2006) or Coe and Vahey (2014). We then use the forecast densities of the individual models and the LOP are then used to calculate probabilities of declines in output and prices at forecast horizons of one, three and six months starting in January 1927.

The remainder of this paper is organized as follows. In section two we discuss the data, the models using to generate forecasts and forecast combination. In section three we present preliminary results based on time series models consistent with agents forming expectations using simple rules of thumb. In section four we discuss these results and outline next steps.

2 Data and Methods

Figure 1 shows the US consumer price index from 1919 to 1941. Two episodes stand out, the sharp post-WW1 inflation and deflation of the early 1920s and the sharp deflation during the Great Depression. In between those two episodes we see a period in which there was a general upward trend in prices until 1926 and then a slight downward trend until 1929. However, as panel (b) shows it was not uncommon for the CPI to fall in

one month, having risen in the previous month and vice-versa. This is in contrast to the period from late 1929 to mid-1933 in which there was only one month in which the CPI did not fall. Figure 2 shows the US industrial production index. In addition to the Great Depression of 1929-33 it also shows the milder downturns in 1920/1 and 1937/8. Details on the data used in this paper are available in the appendix.

2.1   Models and Forecasts

We plan to use a variety of models to generate forecasts for inflation and industrial production growth. For this draft of the paper we restrict our attention to univariate auto-regressive models similar to those employed by Cecchetti (1992). These are consistent with agents forming expectations using simple rules of thumb based on current and lagged values of series of interest. Let x_t denote either the inflation rate (Δp_t) or the growth rate of industrial production (Δy_t) and assume that in period t we are interested in forecasting this object h periods ahead, that is x_{t+h}. We first estimate an AR(p) model for x_t using data up to and including period t and given by:

$$x_t = \alpha + \sum_{j=h}^{h+p-1} \beta_j x_{t-j} + u_t \tag{1}$$

where α is an intercept and the parameters $\beta_1....\beta_p$ capture persistence. The error term $u_{h,t}$ is assumed to be mean zero, serially uncorrelated and with variance σ^2. We consider values of p from one to four.[2]

The predictive density for x_{t+h} from model i is given by $g(x_{t+h} \mid I_{i,t})$. Note, the information set, $I_{i,t}$, will differ across models. For the AR(p) models considered in this draft, the information set contains current and lagged values of x, given by $x_t, x_{t-1}, ...x_{t-p+1}$ and the estimated parameters from (1). For richer models it will

---

[2]Using our full sample of data from 1919M2 to 1941M12 we estimated models with $p = 1, 2, ..., 12$ and calculated the Akaike and Bayesian Information Criteria. There was very little support for values of $p > 4$ either for univariate autoregressive models or a variety of VAR models. Therefore we restrict our attention to values of $p \leq 4$.

also include current and lagged values of other variables and the parameters describing the correlations between these variables. These densities follow a t-distribution with mean $\mu_{i,t+h}$ and scale $\sigma_{i,t+h}$. Here $\mu_{i,t+h} = X_{i,t}\widehat{B}'_i$ where for our AR(p) models $X_{i,t} = \begin{bmatrix} 1 & x_t & .... & x_{t-p+1} \end{bmatrix}$ and $B = \begin{bmatrix} \alpha & \beta_1 & ... & \beta_p \end{bmatrix}$ and

$$\sigma_{t+h} = \frac{1}{\nu}\left(1 + X_t M^{-1} X'_t\right)\widehat{u}'_t\widehat{u}_t \tag{2}$$

where $\nu = T + 1 - k - 2$ with $k = p + 1$ being the number of parameters in the AR(p) model and T being the sample size used estimate equation (1). Finally $M = X'X$, where X is the data used to estimate (1) and $\widehat{u}_t$ is vector of residuals.

We begin our psuedo real-time forecasting exercise in January 1927. We estimate equation (1) using data from 1919M2 to 1927M1 for $h = 1, 3, 6$ and then forecast inflation or industrial production growth in February 1927, April 1927 and July 1927.. In other words we generate 1-month, 3-month and 6-month ahead forecasts. In each case we generate forecasts from AR(p) models with $p = 1, 2, 3, 4$ and so for each horizon and variable we have four forecasts. We then move forward one period and estimate (1) using data from 1919M3 to 1927M2 and generate one, three and six month ahead forecasts for February 1927. Note that when generating the forecasts for February 1927 we move the beginning of our estimation sample forward one period to 1919M3, and so we use a rolling window (of 96 observations) for our recursive estimation.[3] We continue this recursive forecasting exercise to obtain one, three and six month ahead forecasts for the months January 1927 to December 1934 inclusive. Note that in each case we only use data up to and including the current month in order to estimate model parameters and to project forward to obtain our forecasts.

As Croushore and Stark (2001) note, many macroeconomic variables are subject to data revisions. Failing to account for this by using heavily revised data often masks

---

[3]An alternative would be to keep the beginning of the sample fixed at 1919M2 and use an expanding window for recursive estimation. We do this and obtain quantitatively similar results to those reported in the next section.

real-time predictive content. While real-time data are readily available for a number of macroeconomic variables for the late twentieth and early twenty-first centuries, this is not generally the case for the inter-war period. Having said that, real time data for our sample period does exist for one of the variables that we study, namely industrial production. In this draft of the paper we focus on a single vintage from January 22nd, 1942. This date is sufficiently past the end of our period of interest that the effect of data revisions will be minimal. In subsequent drafts we plan to fully exploit the real-time data on industrial production in order to get as close as possible to the information set available to agents in real time.

2.2   Combining Forecasts Using the LOP

We utilize a linear opinion pool (LOP) to combine the forecast densities from our different models; see, among others, Timmermand (2006) or Coe and Vahey (2014). The opinion pool approach has a long tradition in management science for effective combination of expert opinions. Wallis (2005) discusses the foundations for and development of prediction with opinion pools in various applied statistical fields. One appealing feature of the LOP is that the combination of Gaussian forecast densities is not necessarily (or generally) Gaussian. That is, the LOP affords additional flexibility to the combination, with scope for asymmetries in the forecast densities.

We approximate the unknown "true" forecast densities for x_{t+h} by using (time-varying) combinations of the individual model forecasts. Given $i = 1, \ldots, N$ models, the forecast combination density is defined by the LOP:

$$p(x_{t+h} \mid I_t) = \sum_{i=1}^{N} w_{i,h,t}\, g(x_{t+h} \mid I_{i,t}), \qquad t = \underline{t}, \ldots, \bar{t}, \tag{3}$$

where $g(x_{t+h} \mid I_{i,t})$ is the $h-$step ahead forecast density from model i for in period t, conditional on the information set $I_{i,t}$. The non-negative weights, $w_{i,h,t}$, in this finite

mixture sum to unity. Furthermore, the weights may change with each recursion in the evaluation period $t = \underline{t}, \ldots, \bar{t}$. We adopt weights based on the out of sample fit of the individual model's forecast densities. The logarithmic scoring rule provides an intuitively appealing method to gauge fit, giving a high score to a density forecast that assigns a high probability to the realized value of the variable of interest. The logarithmic score of the i^{th} density forecast, $\ln g(x^o_{t+h} \mid I_{i,t})$, is the natural logarithm of the probability density function $g(. \mid I_{i,t})$, evaluated at the realization (out-turn), x^o_{t+h}.

The rolling weights for the $h-$step ahead densities are given by:

$$w_{i,h,t} = \frac{\exp\left[\sum_{s=t-\kappa}^{s=t-1} \ln g(x^o_{t+h} \mid I_{i,t})\right]}{\sum_{i=1}^{N} \exp\left[\sum_{s=t-\kappa}^{s=t-1} \ln g(x^o_{t+h} \mid I_{i,t})\right]}, \qquad t = \underline{t} + h, \ldots, \bar{t} \qquad (4)$$

We set $\underline{t} = 1928\text{M}1$ and $\bar{t} = 1934\text{M}12$ and so calculate weights and generate 1, 3, 6 and 12-month ahead expert combination forecasts for the period January 1928 December 1934. Here κ is the length of the rolling window used to determine the weights. We report results below with a window length of $\kappa = 12$; that is a rolling twelve month window, and so 1927M1 to 1927M12 comprises the training period used to initialize these weights. The intuition being that models that produce superior forecasts over the previous year are given more weight in the combination forecast that those that do not. It is also important to stress that the weights on the various specifications vary through time. As such, the combination forecast exhibits greater flexibility than any single linear model specification.

3   Preliminary Results

As mentioned in the previous section we generate 1,3 and 6 month ahead forecast densities from AR(p) models with $p = 1, \ldots 4$ for inflation and industrial production growth

for the period 1927M1 to 1934M12 and also the combination forecast for 1928M1-1934M12. In this section we first discuss the means of the forecast densities as is the case in the existing literature. After this we move on and present the probabilities of deflation and output declines derived from our forecast densities.

3.1 Forecast Density Means

The means of the combination forecast densities are reported in figure 3 for inflation and figure 4 for industrial production growth. The forecast density means are similar for each our $\text{AR}(p)$ models with $p = 1, 2, 3, 4$ and so we focus our attention on the combination forecast. Also reported in these figures are the out-turns for each variable. The existing literature which focuses on point forecasts judges the extent to which deflation was anticipated by whether or not the forecast mean is close to the out-turn. Figure 3 suggests the that the deflation of the early 1930s was largely unpredicted at the 1, 3 and 6 month horizons. The mean of the forecast density does not become negative at any of the forecast horizons until mid-1930 and for the most part remains above the out-turn until early 1933. It is only in 1933 when inflation becomes positive again that the forecast density mean moves below the out-turn. On the whole the means of the forecast densities look similar at all three horizons, although at the three and six month horizons they are typically slightly below where they are at the one month horizon and therefore slightly closer to the out-turn on average.

Figure 4 paints a similar picture for industrial production growth. At the one month horizon the mean of the forecast density does turn negative in late 1929 and remains negative for most months until late 1932. However, the mean of the forecast density typically fails to match the depth of the decline in industrial production. At the three and six month horizons the forecast density mean does not turn negative until early to mid 1930 and does an even poorer job of matching the depth of the decline in industrial production than it does at the one month horizon. These plots are consistent with the

10

results of Dominguez et al (1988) where VAR models fail to generate point forecasts that match the depth of the Great Depression.

The means of the forecast densities for inflation in figure 3 contrast with those reported by Cecchetti (1992) using similar univariate time series models. This difference is illustrated in the top panel of figure 5. This figure plots two one month ahead forecasts and the out-turn for the period 1928M1 - 1934M12. The first forecast is from an AR(1) model estimated recursively using an eight year window. This is the one of the four models that contributes to our combination forecast in figure 3. As with the combination forecast it does not suggest a movement to deflation until late 1930 and does not match the depth of realized deflation during the Great Depression. In contrast the second forecast comes from an AR(1) model in which the parameters are estimated over the sample 1919M1 - 1928M12, but not updated as observations after 1928M12 become available. This is one of the two samples used to Cecchetti to estimate the model used to forecast inflation.[4] The forecast density means from this model imply an move to deflation in late 1929, rather than late 1930 as implied by our combination forecast. Also throughout 1930 and the first half of 1931 these forecast density means are much closer to the out-turns of inflation than those of the recursively estimated model.

The difference in predictive ability implied by these two forecast density means is explained by the second panel of figure 5. The mean of a one step ahead forecast density for inflation from an AR(1) model is given by:

$$\widehat{\alpha} + \widehat{\beta}_1 x_t \tag{5}$$

where x_t is the current inflation rate. The relative success of the AR(1) model with the parameters fixed at those estimated on data from 1919M1 - 1928M12 comes from the implied persistence in the inflation process. With $\widehat{\beta} = 0.48$ and $\widehat{\alpha} = 0.027$, this model

---

[4]The other is 1919M1-1941M12 which yields similar parameter estimates.

will return a point forecast that implies deflation at the one-step ahead horizon for any $x_t < -0.056$ (an annual inflation rate of about -0.7%). Therefore, when deflation starts in mid-1929 the forecast density mean for this model immediately switches from positive to negative and predicts that deflation will continue, until it actually ends in early 1933. As Evans and Wachtel point out this forecast is only a good proxy for agents expectations if agents believed that prices would continue to follow this process. The second panel of figure 5 shows recursive estimates of β. As discussed in the previous section these estimates are constructed using a moving window of 96 months. This figure shows that these estimates imply a decline in the persistence of the inflation process in the period preceding the Great Depression. As we move the end of the sample forward from 1927M12 to until early 1929 the estimate of β falls from around 0.4 to zero. It then remains close to zero until mid-1930, when it starts to climb again. This explains the difference in the forecast density means in the top panel of figure 5, as β moves close to zero the mean of the forecast density becomes flat. The bottom panel of figure 5 is also consistent with the argument of Evans and Wachtel, that not allowing for uncertainty in the model parameters overstates the predictability of the deflation during the Great Depression.

Note, we do not see the same pattern in the AR(1) parameter for the industrial production growth processes. At the one month horizon the estimate of β is around 0.35 at the beginning of our sample, but reaches 0.5 by early-1933. This helps to explain how the mean of the forecast density is closer to the out-turn in panel (a) of figure 4. At the three and six month horizons the estimate of β is much closer to zero for most of the sample, consistent with the pattern of the forecast density means reported in panels (b) and (c) of figure 4.

3.2 Density Forecasting

In this section we shift our focus away from the mean of the forecast density and look at at the probabilities of certain events that are implied by these densities. Therefore, we are no longer judging the extent to which an event could have been predicted by whether or not the mean of the forecast density is to the out-turn. Rather we judge how predictable events could have been by the probabilities that our forecast densities assign to them. In particular we are interested in whether the probability of deflation or output declines increase in the run up to the Great Depression.

The three panels on the left hand side in figure 6 show the probability of deflation at the one, three and six month horizons using the LOP combination forecast, that is, they plot $p(\Delta p_{t+h} < 0 \mid I_t)$. Not surprisingly given the mean of the inflation series in figure one, and the estimates of β reported in panel (b) of figure 5 this probability is around 0.5 until late 1930 when it begins to rise. It continues to rise until 1933 when it drops back to about 0.5 where it remains until the end of 1934. While these figures suggest that it was reasonable to believe that agents might have believed deflation was a very real possibility the do not show a sharp increase in the probability of deflation until late 1930, by which time prices had already been falling for a year.

The three panels on the right hand side in figure 6 shows the probability of a one-month decline in industrial production at the one, three and six month horizons using the LOP combination forecast, that is, $p(\Delta y_{t+h} < 0 \mid I_t)$. They do show an increase in the probability assigned to the negative growth event beginning around mid-1929. This is particularly apparent at the one month horizon. The increase is less apparent at the three and six month horizons, but in a each case a probability of at least 0.5 is assigned to negative growth before the end of 1929.

While figure 6 suggests that declines in the price level and output were not events that should have been considered particularly unlikely in the run up to the Great Depression, a more interesting question is whether or not declines of the magnitude

observed between 1929 and 1933 would have been considered unlikely. In the 44 months between July 1929 and March 1933 the consumer price index fell by an average of 0.71% per month and the industrial production index fell by an average of 1.68% per month.[5] Figure 7 explores the question of how likely declines of these magnitudes should have been considered by plotting the objects $p(\Delta p_{t+h} < -0.71\% \mid I_t)$ and $p(\Delta y_{t+h} < -1.68\% \mid I_t)$ implied by our forecast densities. As with figure 6, the three panels on the left depict the probabilities for the deflation event at the three horizons and the three panels on the right depict the probabilities for the output event.

The probability assigned to a one-month fall in prices of 0.71% at each horizon is slightly above 0.2 at the beginning of 1928, but it then falls steadily until the fall of 1929. It then remains constant at around 0.1 until early-1930 before it starts to rise, reaching a peak of around 0.5 by early 1932. While these probabilities do not suggest agents using a rule of thumb to form their price expectations would have been particularly fearful of a large deflation in the early days of the Great Depression, they do imply a greater risk of deflation starting in early to mid-1930. Therefore while a deflation in the order of 0.71% per month may not have been considered the most likely outcome according to rule of thumb expectations, by the summer of 1930 it did not look like a remote possibility.

The figures on the right hand side of figure 7 depict a similar to those on the right hand side of figure 6. That is the probability assigned to the event of interest, in this a case one month industrial production decline of 1.68%, begins to rise in mid 1929. By 1930 it reaches 0.4 at the one month horizon and 0.3 at the longer horizon. Again this suggests that while output declines of this magnitude might not have been most likely outcome, they were not zero-probability outcome either.

---

[5]In July 1929 the (seasonally adjusted) Bureau of Labor Statsitics' Consumer Price Index was 17.33 and in March 1933 it was 12.68. The corresponding figures of industrial production were 114 and 54 respectively.

4 Discussion and Next Steps

Overall our results suggest that while deflation and output growth of the magnitudes seen during the Great Depression may not have been the most likely outcomes in 1929 and early 1930, it would not have made sense to rule them out completely as possibilities. However, one obvious shortcoming of the simple AR(p) models that we have used to date to generate forecasts of future inflation and output growth is that they are based purely on past realizations on the variable of interest. As such any predictability and forecast uncertainty comes from the series own degree of persistence. In particular such models will ignore the information and forecast uncertainty that comes from other variables.

Therefore, our next step is to expand the model space to ask whether in a density forecasting exercise there is additional information that can be gained from expanding the information set. By moving away from univariate models we can exploit any correlation in the data between output and prices. We can also ask whether other variables such as monetary aggregates, interest rates, gold flows and leading indicators such as those used by Dominguez et al (1988) contain predictive content for inflation and output growth in a density forecasting context.

Bibliography

Bernanke, Ben S. (1983) "Non-monetary Effects of the Financial Crisis in the Propagation of the Great Depression" *American Economic Review* 73, pp257-276.

Cecchetti, Stephen G. (1992) "Prices During the Great Depression: Was the Deflation of 1930-32 Really Unanticipated?" *American Economic Review* 82, pp141-156.

Coe, Patrick J. and Shaun P. Vahey (2014) "Probabilistic Prediction of the US Great Recession Using Historical Experts" *Warwick Business School Economic Modelling and Forecasting Group, Working Paper 5.*

Croushore, D. and T. Stark (2001) "A Real-time Data Set for Macroeconomists", *Journal of Econometrics*, 105, pp111-130.

Dominguez, Kathryn M., Ray C. Fair and Matthew D. Shapiro (1988) "Forecasting the Depression: Harvard versus Yale" *American Economic Review* 78, pp595-612.

Evans, Martin and Paul Wachtel (1993) "Were Price Changes During the Great Depression Anticipated?" *Journal of Monetary Economics* 32, pp3-34.

Fisher, Irving (1933) "The Debt-Deflation Theory of Great Depressions" *Econometrica* 1, pp337-357.

Friedman, Milton and Anna J. Schwartz (1963) *A Monetary History of the United States, 1867-1960*, Princeton University Press, Princeton.

Hamilton, James D. (1987) "Monetary Factors in the Great Depression" *Journal of Monetary Economics* 19, pp145-169.

Hamilton, James D. (1992) "Was the Deflation During the Great Depression Anticipated" *American Economic Review* 82, pp157-178.

Romer, Christina D. and David H. Romer (2103) "The Missing Transmission Mechanism in the Monetary Explanation of the Great Depression" *American Economic Review Papers and Proceedings* 103, pp66-72.

Timmermann, A. (2006) "Forecast Combination" in G. Elliot, C.W.J. Granger and A. Timmermann, eds, *Handbook of Economic Forecasting*, North-Holland, pp197-284.

Wallis, K.F. (2005) "Combining Density and Interval Forecasts: a Modest Proposal" *Oxford Bulletin of Economics and Statistics*, 67, pp983-994.
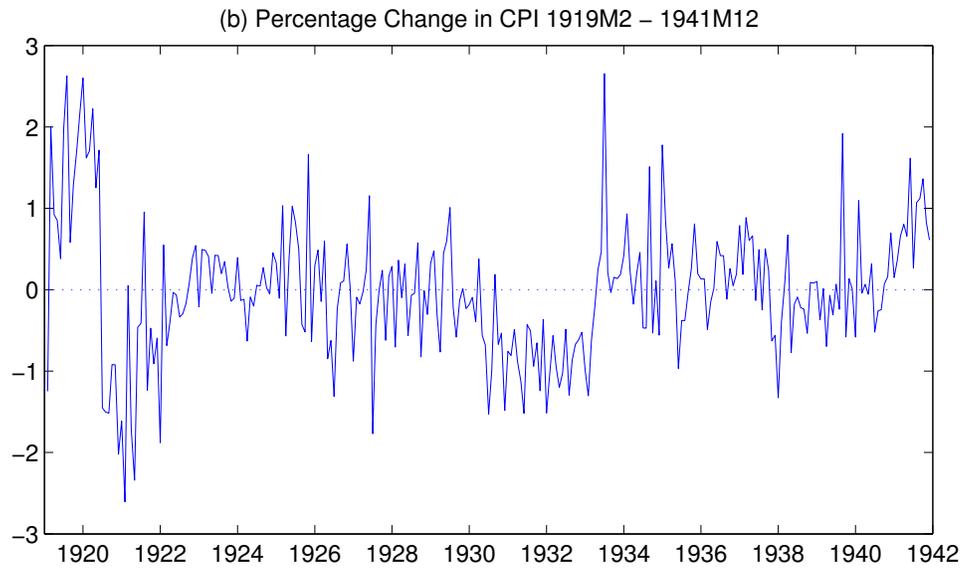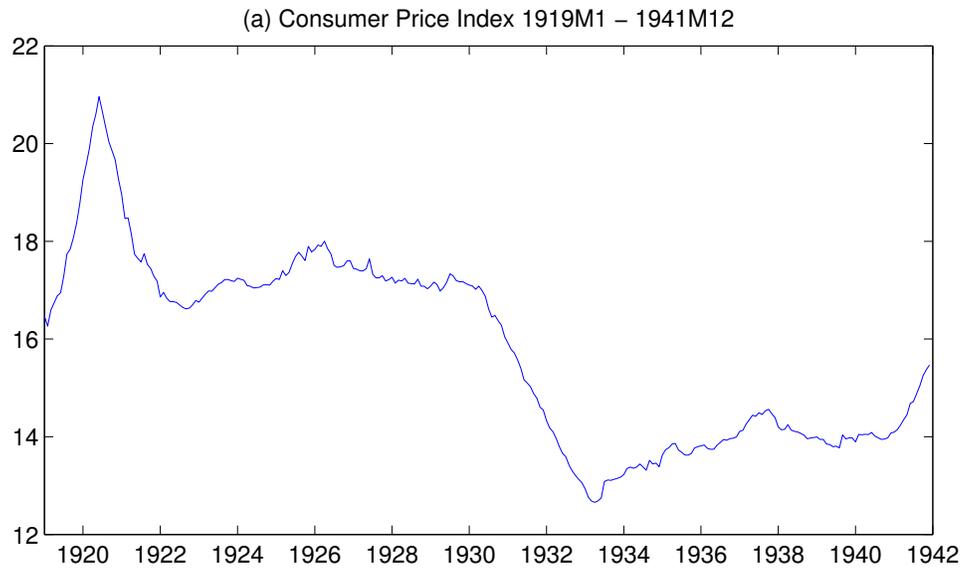
(a) Consumer Price Index 1919M1 – 1941M12

(b) Percentage Change in CPI 1919M2 – 1941M12

FIGURE 2: INDUSTRIAL PRODUCTION INDEX 1919 - 1941

(a) Industrial Production Index 1919M1 – 1941M12

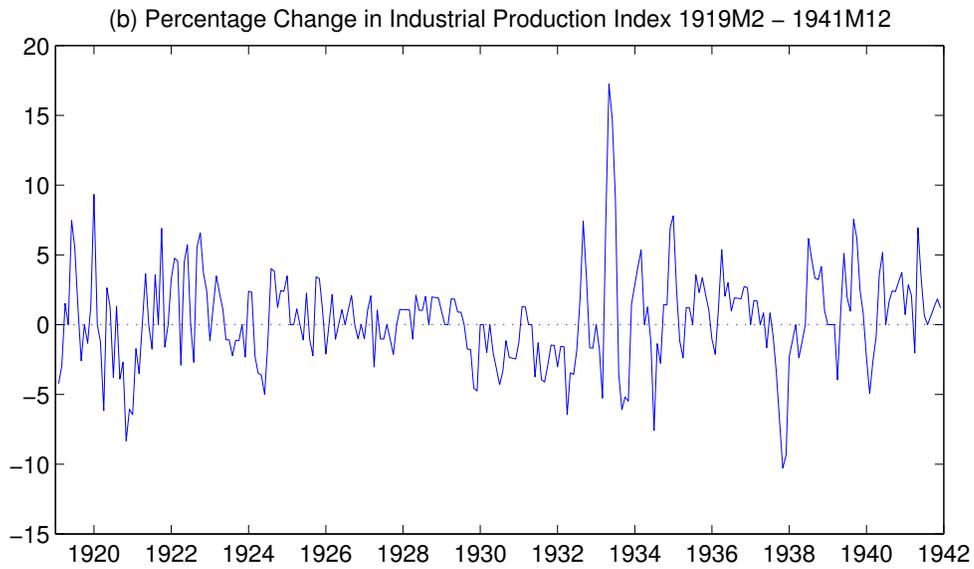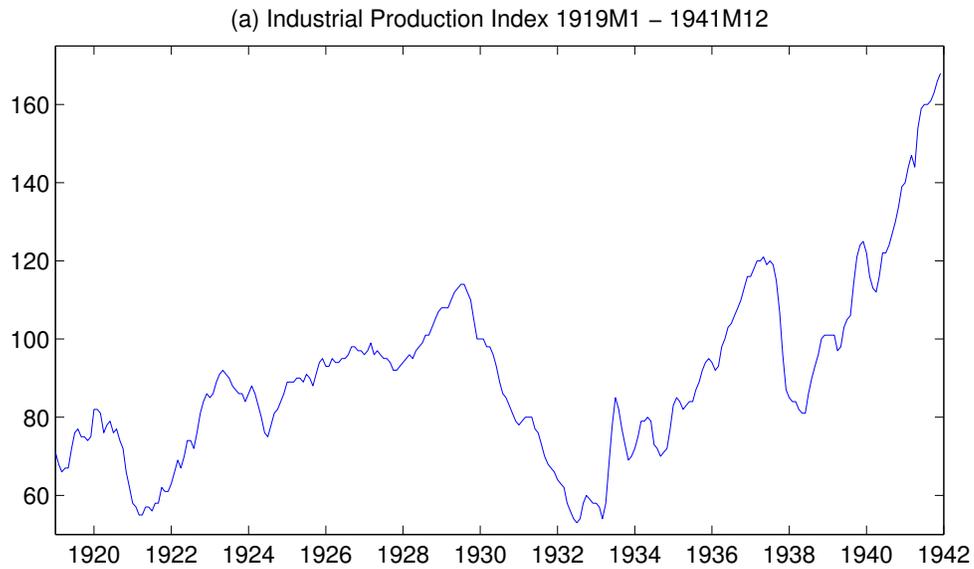(b) Percentage Change in Industrial Production Index 1919M2 – 1941M12
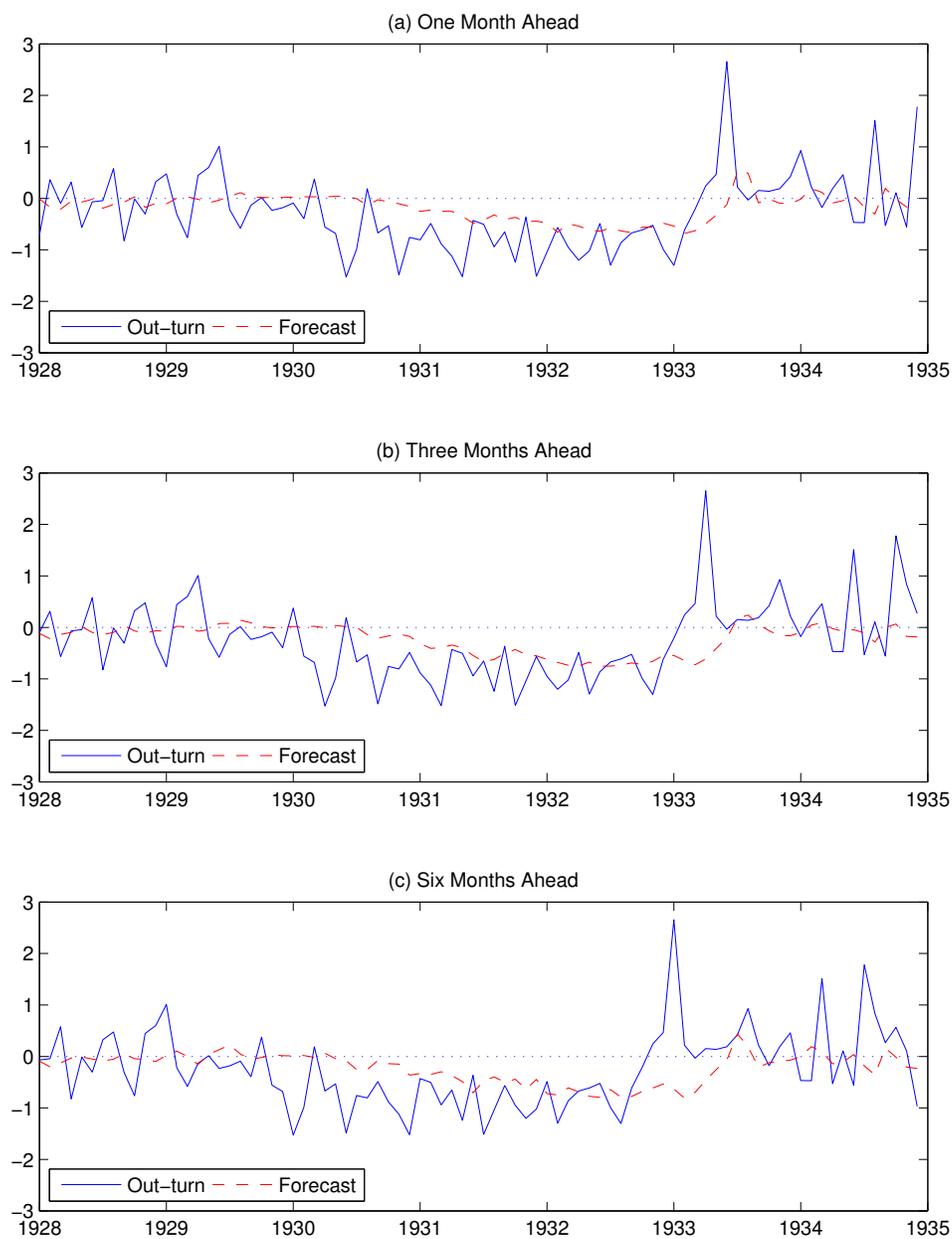
FIGURE 3: INFLATION FORECASTS
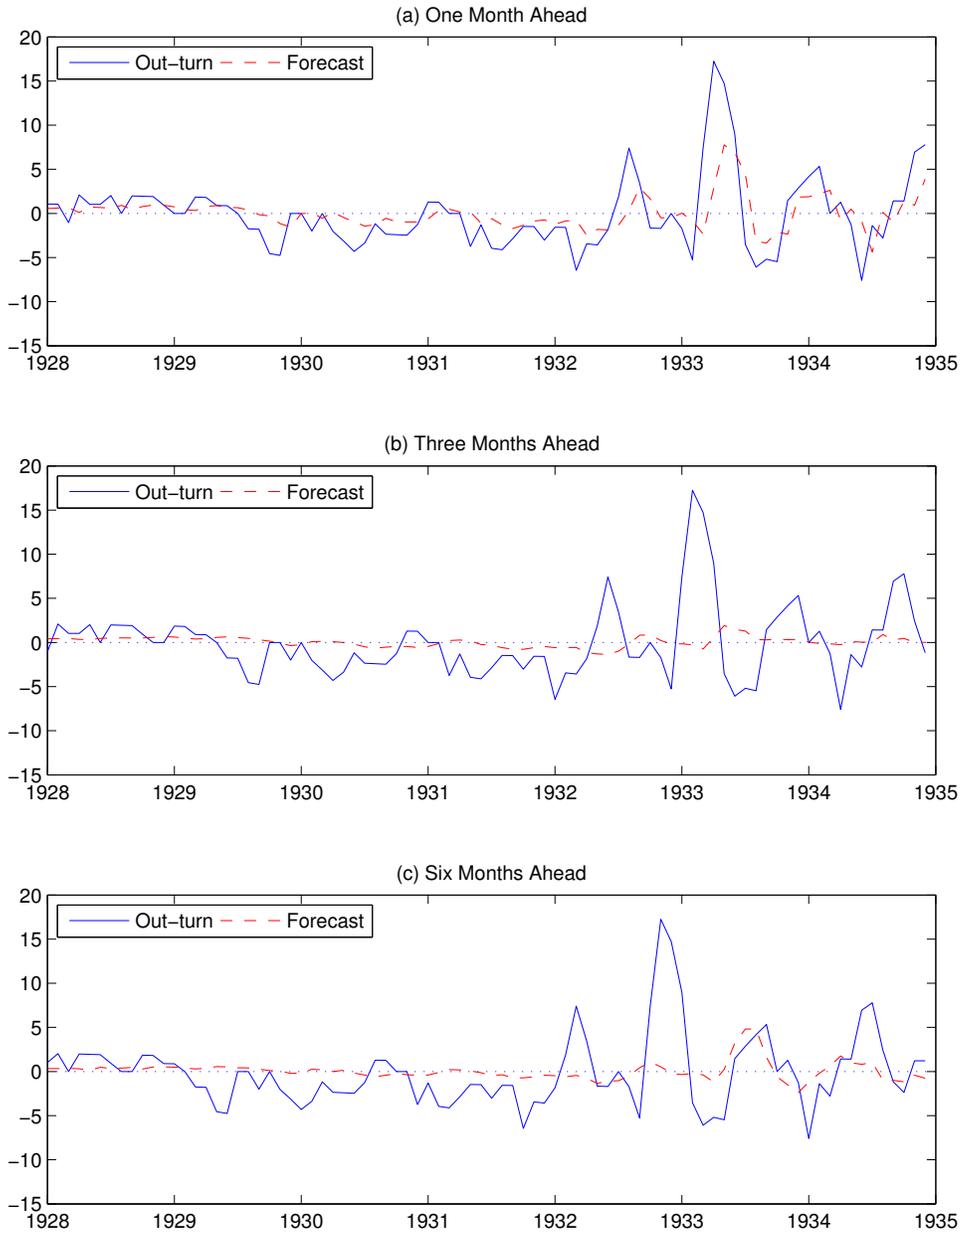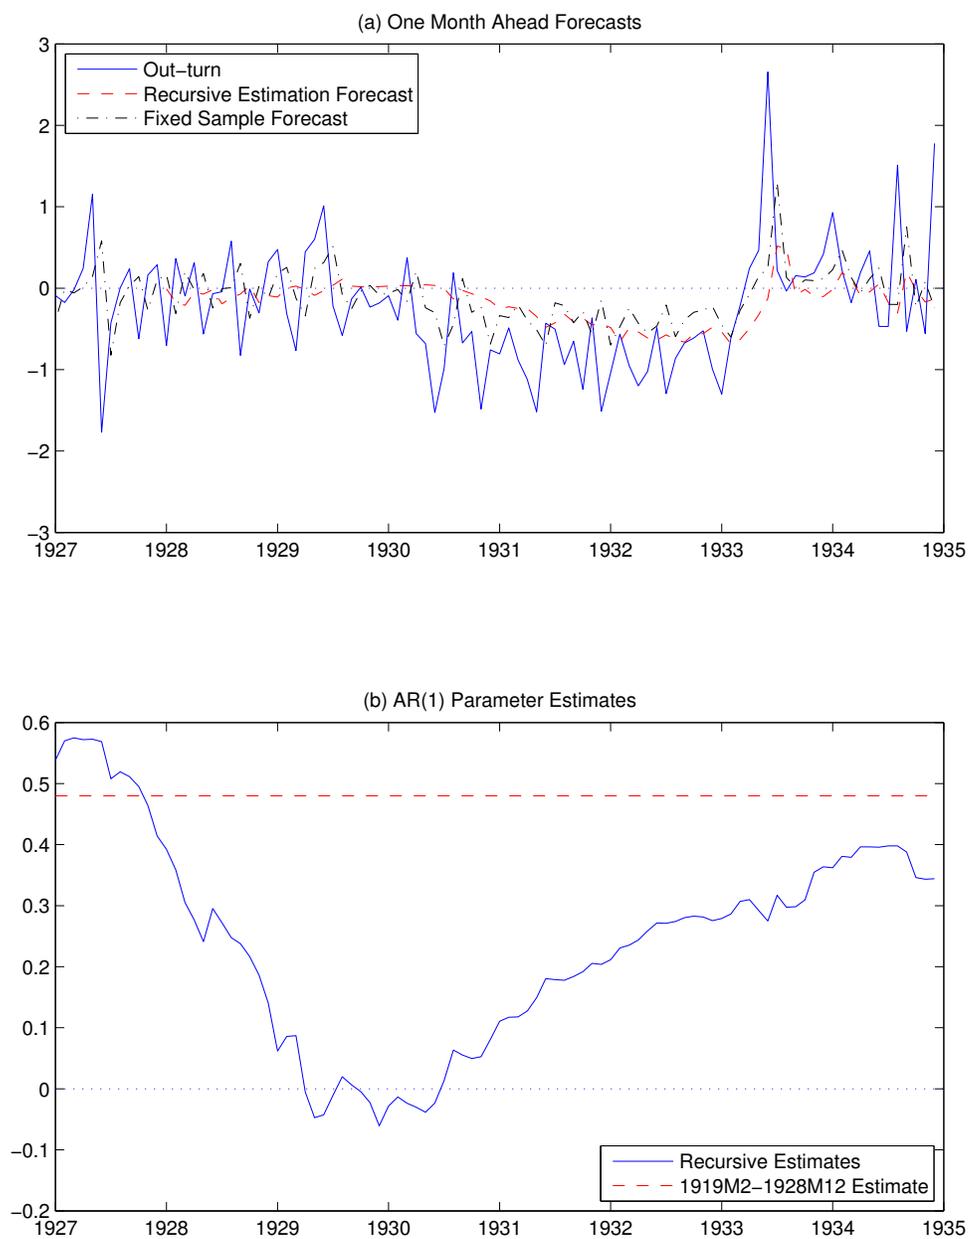
FIGURE 4: INDUSTRIAL PRODUCTION GROWTH FORECASTS

(a) One Month Ahead

(b) Three Months Ahead

(c) Six Months Ahead

FIGURE 5: COMPARISON OF FIXED SAMPLE AND RECURSIVE SAMPLE FORECASTS

(a) One Month Ahead Forecasts

(b) AR(1) Parameter Estimates

Notes: The recursive estimation forecast is from an AR(1) estimated recursively as described in section 2. The fixed sample forecast is generated using with parameters from an AR(1) fixed at those estimated on data from 1919M1 - 1928M12.

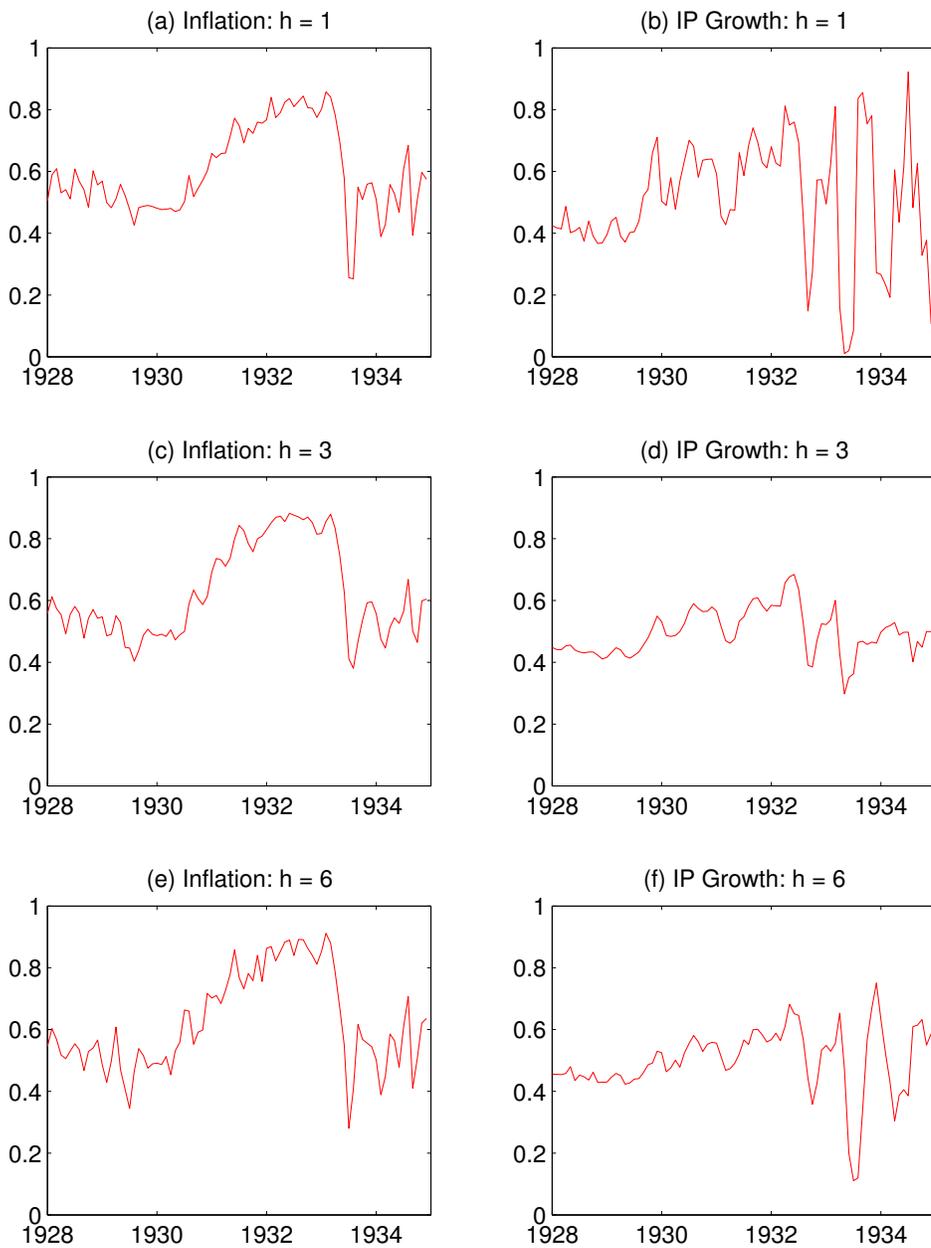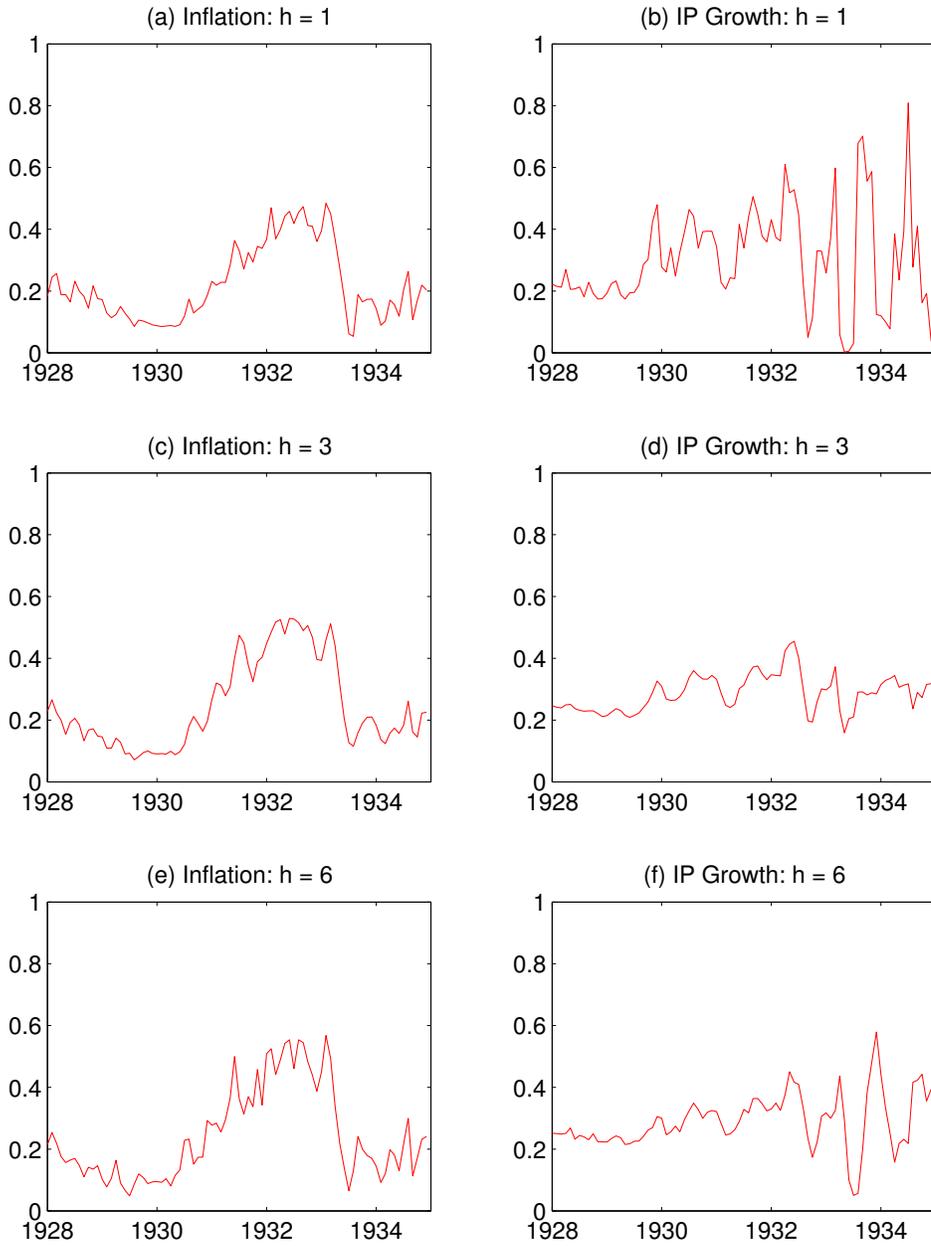FIGURE 6: PROBABILITIES OF ONE MONTH DEFLATION AND OUTPUT DECLINES

FIGURE 7: PROBABILITIES OF "LARGE" ONE MONTH DEFLATION AND OUTPUT DECLINES

Notes: A "large" one month deflation is interpreted as a deflation of 0.71% and a "large" one month decline in output is interpreted as a decline in industrial production of 1.68%.