Nowcasting in Real Time Using Popularity Priors

Monokroussos, George

1 November 2015

Nowcasting in Real Time Using Popularity Priors*

George Monokroussos

Amazon - Seattle, WA, USA†

November 2015

Abstract

This paper proposes a Bayesian nowcasting approach that utilizes information coming both from large real-time data sets and from priors constructed using internet search popularity measures. Exploiting rich information sets has been shown to deliver significant gains in nowcasting contexts, whereas popularity priors can lead to better nowcasts in the face of model and data uncertainty in real time, challenges which can be particularly relevant during turning points. It is shown, for a period centered on the latest recession in the United States, that this approach has the potential to deliver particularly good real-time nowcasts of GDP growth.

JEL Classification: C11, C22, C53, E37, E52.

KEYWORDS: Nowcasting, Gibbs Sampling, Factor Models, Kalman Filter, Real-Time Data, Google Trends, Monetary Policy, Great Recession.

*I thank seminar participants at the 5th European Seminar on Bayesian Econometrics (ESOBE 2014), at the 34th International Symposium on Forecasting, and at the Joint Research Centre of the European Commission for useful comments and discussions. Furthermore, I am grateful to David Small for making available his real-time data. This research was conducted while the author was with the Joint Research Centre of the European Commission. The content of this study does not reflect the official opinion of the European Union. Responsibility for the information and views expressed in the study lies entirely with the author.

†Email address: monokroussos@gmail.com

1 Introduction: Nowcasting Around Turning Points

Obtaining accurate and timely forecasts of turning points in GDP growth is a central preoccupation of macroecononometrics. This is quite justifiable, given the importance of the task for the private sector and for policy makers alike. However, good such forecasts can be difficult to obtain. For example, there is a substantial, decades-old literature (see, inter alia, Stekler (1972), Zarnowitz (1986), Loungani (2001), Lahiri and Wang (2013)) that presents ample evidence of such predictive failures spanning several countries, historical episodes of recessions, forecast horizons, and types of forecasters.

On the other hand, recent literature has demonstrated that it is possible to produce good *nowcasts* of GDP growth. For instance, the influential paper of Giannone, Reichlin and Small (2008) has convincingly demonstrated the gains to be made when nowcasting US GDP growth using dynamic factors to exploit information coming from a large data set.

In view of both of the above, one question that arises naturally is how well we can detect turning points in GDP growth in real time during the current quarter. Of course, this is a key question: Being able to accurately assess the present state of the economy, especially at the onset of recessions, as this is summarized by current-quarter GDP growth, is of central importance for the timely conduct of monetary policy, among other purposes. The most recent historical episode of the Great Recession is particularly telling.

The NBER dates for the latest recession are December of 2007 for its peak and June of 2009 for its trough. One particularly consequential quarter in this recession was 2008 Q3, as it included events such as the collapse of Lehman Brothers, and it was also the first in a series of consecutive quarters with negative GDP growth. Of course, official estimates on 2008 Q3's negative growth only became available in the following (fourth) quarter. Similarly, the NBER called the recession on December 1, 2008. It is expected though that NBER recession announcements will not be the most timely possible.

The same cannot be said however regarding the deliberations and decision making process of the monetary policy maker. Given likely lags in the monetary transmission mechanism, central banks have to rely heavily on macroeconomic forecasts. In particular, the Federal Reserve forecasts are generally perceived as being quite good. However, the recently released minutes of the FOMC meetings covering the crucial period of the summer and fall of 2008 paint a picture of insufficient appreciation of (the extent of) the slowdown in real time and thus of the consequent policy risks and trade offs faced[1].

Given this, how would an approach that relies explicitly and exclusively on an econometric model fare? As was discussed above, the large-data, factor-based model of Giannone, Reichlin and Small (2008, henceforth GRS) is arguably at the peak of what we can achieve in GDP nowcasting contexts. Figure 1 shows the nowcasts of GDP growth obtained in real time using the model of GRS with historical vintages of close to 200 variables. As can be seen there, the model-based nowcasts do not turn negative until December of 2008 again. Furthermore, Giannone et al (2010) provide real-time nowcasts that turn negative late in the fall of 2008 and which are, however, more timely than either the respective Greenbook forecasts or the respective figures released by the Survey of Professional Forecasters.

In view of all of the above, one may tend to conclude that in the face of challenges such as model and data uncertainty, problems which can be particularly acute in real time during turning points, there may be little more that we can do. We may have to settle for nowcasts which are inferior around such turning points (when they are arguably needed the most) than at other times.

A central contribution of this paper is to show that such pessimistic assessments do not have to be true any longer. Given the contemporary prevalence of the internet and internet search engines, we now have forecasting tools at our disposal that were not available until recently. GDP growth turning from positive to negative typically entails a widespread slowdown in economic activity. Workers, investors, employers, etc. who experience a change

[1]Matthew O'Brien scrutinized the recently released minutes of the June 24-25, August 5, and September 16, 2008 FOMC meetings in terms of keyword counts (e.g. frequency of the word "inflation" vs. "unemployment" or "systemic risks/crises") as well as in terms of specific statements by participants in the FOMC meetings and provides a substantial series of evidence along these lines (O'Brien (2014)). The Greenbook forecasts for 2008 Q3 GDP growth associated with these three FOMC meetings were all positive and indeed close to 1%.

in their conditions that is associated with the slowdown are more likely to conduct internet searches using keywords related to the slowdown than at other times. Internet-based services such as *Google Trends* construct normalized time series indices reflecting the relative volume of such keyword searches. Such measures may provide a valuable gauge of the economy in real time as they have the potential to capture widespread changes in conditions in a timely manner and are also not subject to revisions and real-time inaccuracies like many of the more traditional variables are.

This paper proposes and estimates a new Bayesian model for a policy maker or any forecaster in general who, rather than operate within the confines of a traditional data set, instead "listens to hoi polloi" too, that is, lets her prior beliefs be influenced by internet search popularity measures. Forecasts emerge from posterior estimates that reflect both such prior beliefs and information coming from large traditional data sets. Dynamic factors are used to capture the collinearities and summarize these large data sets in a parsimonious way without throwing away information.

The main empirical result is that for a time period centered around the Great Recession, this Bayesian factor-based nowcasting approach with popularity priors delivers a more timely detection of the 2008 Q3 turning point than all of the other alternatives discussed above (as is illustrated in Figure 8). Furthermore, it achieves a substantially better outcome regarding this consequential turning point while its nowcasts for the rest of the time are (at least) as good as the ones obtained from alternative approaches in real time.

The Bayesian approach advocated in this paper points to a similar direction as Wright's (2013) innovative study. Wright proposes a Bayesian VAR forecasting approach and demonstrates the significant macro forecasting gains to be made when the prior is constructed using survey responses (hence dubbed "democratic priors").

From a methodological perspective, there are other contributions employing likelihood-based and Bayesian approaches to estimate macroeconomic factor models, including Kose, Otrok and Whiteman (2003), Boivin and Giannoni (2005), Bernanke, Boivin and Eliasz

(2005), Doz, Giannone and Reichlin (2012). Kim and Nelson (1998) is an important early contribution employing a Gibbs sampler to extract one factor from a full data set of four variables, which, when augmented with regime-switching features, serves as their coincident index. Thus, their framework is not designed for purposes of producing high-frequency nowcasts in real time using large data sets. D' Agostino et al (2015) employ a Gibbs sampling approach in a business cycle nowcasting context, using a small set of variables as well. Our framework here is one with tens of variables (close to two hundred variables in our application), with missing values, from which more than one factors, and with a general covariance structure, are extracted. It is the "Bayesian counterpart" to factor-based nowcasting models such as that of GRS, and can be used to nowcast using incomplete large data in real time, *irrespective of prior choice.*

The plan for the rest of the paper is as follows: The following section describes the theory and the details of the model, whereas Section 3 provides the specifics of the estimation and the Gibbs sampler. Section 4 discusses the application of the model to nowcasting US GDP growth in a period centered around the Great Recession. Finally, Section 5 provides some concluding remarks.

2 A Bayesian Dynamic Factor Approach for Large, Evolving, Jagged-Edge Data Sets

We seek to generate h-period ahead forecasts (or nowcasts, when $h = 0$) of the (stationary) macroeconomic aggregate y we are interested in by exploiting a rich information set coming from (potentially) hundreds of relevant variables[2]. Furthermore, our goal is to create a framework that allows us to generate such forecasts at *any* point in time using the most up-to-date information that is available at that point; that is, we seek to generate real-time forecasts or nowcasts.

[2]In a standard fashion we assume that our variables are stationary or transformed appropriately to induce stationarity.

Such a real-time, rich-information approach can certainly deliver efficiency gains. Furthermore, it arguably offers a realistic depiction of the tasks undertaken by policy makers and forecasters who need to keep updating their forecasts of the macro aggregates they track in the face of a high-frequency stream of news releases on the evolving conditions in the economy. However, there are two central challenges associated with such attempts to generate forecasts by employing all the latest information coming from many variables. First, any model that requires estimating separate parameters for all or most of these variables will run into severe degrees-of-freedom problems and hence, as is well known, will deliver poor forecasts out of sample. Second, as these variables get updated at different points in time, with different frequencies, and with various coverage lags, there will be missing values in the data set that is available for estimation and forecasting at any point in time, especially for observations corresponding to the most recent time periods. That is, we will always have to deal with a "jagged-edge" data set.

In their seminal contribution, Giannone, Reichlin, and Small (2008) propose a dynamic factor model that can deal with the challenges outlined above and use it to nowcast current-quarter GDP growth from a large jagged-edge macroeconomic data set. Here we build on their approach.

More specifically, since many macro variables are monthly and get updated once a month, at (typically) the same business day of the month, we assume that time is measured in months t, and days τ_t, that is business days of month t when new data releases on one or more variables in our data set become available[3]. We have data on N variables x_i with $i = 1, ..., N$ that contain information that is potentially useful for the macroeconomic aggregate we wish to forecast. These variables get updated at different days of the month, and the new data releases

[3]So, we assume that all our variables are either monthly or converted to a monthly frequency. Furthermore, we ignore heterogeneity issues in data releases within business days and across months and assume that the same variables get updated at the same time, on the same business day every month. Our assumptions on the time units are the same as those of GRS and are appropriate for numerous macro contexts, including the application of this paper, where we nowcast GDP using a set of macro variables. Our discussion in this section and our estimation algorithm can be applied, however, to contexts where different frequencies are assumed.

through which they get updated reflect various coverage lags. So, at day τ of month t the latest available observation for variable x_i covers the present month, or last month, or some earlier month, depending on the variable's update schedule and publication lag. That is, $x_{it|\tau_t}$ is either available or missing, and thus, inevitably the $N \times 1$ vector $x_{t|\tau_t} = (x_{1t|\tau_t}, ..., x_{Nt|\tau_t})'$ has missing values and the $t \times N$ panel $x_{|\tau_t}$ is an unbalanced, jagged-edge panel. Note that $x_{|\tau_t}$, which includes all the information that is available at day τ_t, can differ from the information set available at $(\tau - 1)_t$ or any earlier day for one or both of the following reasons: New observations for one or more variables have become available (through news release(s) at day τ_t) or existing observation(s) for one or more variables have been revised (again through news release(s) at day τ_t).

We further assume that collinearities of the variables in our panel are captured by a few (common, latent) factors $F_{1t|\tau_t}, ..., F_{rt|\tau_t}$, which summarize all information that is available at time τ_t (so, in any given month t, we expect to estimate the month's factors, $F_{1t}, ..., F_{rt}$ as many times as there are days τ_t with new data releases). These factors thus produce a parsimonious specification that enables us to deal with the degrees-of-freedom problem and deliver (potentially) efficient forecasts.

The above discussion is summarized by the following equation:

$$x_{t|\tau_t} = \mu + \Lambda F_t + \xi_{t|\tau_t} \tag{1}$$

where μ is an $N \times 1$ vector of constants, and $\xi_{t|\tau_t}$ is an $N \times 1$ vector of idiosyncratic error terms which are Gaussian white noises, and are also cross-sectionally orthogonal. That is, $E(\xi_{t|\tau_t}\xi'_{t-s|\tau_t}) = 0$, for all $s > 0$, and all τ, t, and $E(\xi_{t|\tau_t}\xi'_{t|\tau_t}) = \Sigma^{\xi} = diag(\sigma^2_{\xi_1}, ..., \sigma^2_{\xi_N})$. Λ is an $N \times r$ vector of factor loadings, and the $r \times 1$ vector $F_t = (F_{1t}, ..., F_{rt})'$, is the set of factors that are orthogonal to the error terms.

Regarding these factors, the literature commonly employs dynamic specifications[4] that allow for inertia and that can capture intertemporal relationships among variables during the

[4]Indeed most of the studies mentioned in the Introduction employ dynamic factor specifications.

business cycle, and here we follow this standard paradigm by specifying a first order vector autoregression for the factors:

$$F_t = AF_{t-1} + \zeta_t \tag{2}$$

where A is an $r \times r$ coefficient matrix with all the roots of $\det(I_r - Az)$ lying outside the unit circle, ζ_t is an $r \times 1$ vector of "common shocks"[5], Gaussian white noises that are independent from the idiosyncratic error terms and with $E(\zeta_t \zeta_t') = \Sigma^{\zeta}$ and where F_t, F_{t-1} are the current and last month's factors.

Finally the model is complete with a so-called "bridge equation" that delivers forecasts or nowcasts of the macro aggregate in question, y, as a function of the factors:

$$y_{t+h|\tau_t} = \alpha + \beta' F_{t+h} + \varepsilon_{t+h|\tau_t} \tag{3}$$

where α is scalar and β is an $r \times 1$ vector of coefficients, $h = 0$ (for nowcasts) or higher, and $\varepsilon_{t+h|\tau_t}$ is also a Gaussian white noise with variance σ_ε^2. Note that we assume here without loss of generality that y is observed at the same frequency as x. In cases where this is not the case, such as that of the application that follows in which y is GDP and x consists of monthly macro quantities, several options are available. For instance, GRS filter the data using the filter suggested by Mariano and Murasawa (2003, approximating quarterly arithmetic means by geometric means, in order to preserve the linear state-space structure of the model and thus simplify the estimation) so as to convert the variables into quarterly quantities when observed at the end of a quarter. Another useful approach is based on Mixed Data Sampling (see, inter alia, Ghysels et al (2004), Andreou et al (2010)). Standard MIDAS techniques can be applied to bridge equations such as equation (3) above to deliver a way of dealing with the mixed frequency issue that is both parsimonious and data-driven (Marcellino and Schumacher (2010) and Kuzin, Marcellino and Schumacher (2011) discuss

[5]As discussed in GRS the assumption that there are as many common shocks as there are factors can be relaxed to allow for fewer common shocks than factors.

interesting applications in related nowcasting contexts).

As is clear from equations (1) and (2), we have a state-space framework, in which, however, standard Kalman filtering techniques are not readily applicable, because of the missing data issue discussed above. A modification, possibly based on the EM algorithm (see, for example, Stock and Watson (2002a)) is necessary. GRS deal with this problem in a similar way by adopting a multi-step estimation approach; they obtain preliminary estimates of the factors by extracting principal components from a balanced subset of the full jagged-edge data set, which they use to obtain parameter estimates, which they then plug into their Kalman smoother to produce final estimates of the factors. These final estimates are then plugged into the bridge equation, which delivers forecasts of the macro aggregate in question.

One issue with such estimation strategies is that they take point estimates obtained from the previous step and plug them directly into the next step, thus ignoring parameter estimation uncertainty. This, in a sense, is a consequence of a classical estimation approach that treats the model's hyperparameters differently from the factors; however, as the factors are also unobservables, this could be viewed as an arbitrary distinction. A Bayesian approach removes such asymmetries, and utilizes the joint posterior distribution - conditioning on day τ_t's information set - of the model's variables to make inferences on the factors (and on the other model variables). Indeed we believe that it can be especially suitable in the present context where the focus is less about hoping for a large enough sample size in order to uncover true parameter values and more about updating estimates and forecasts as soon as new information becomes available on day τ_t.

Furthermore, a Bayesian estimation approach readily delivers the entire posterior distribution of $y_{t+h|\tau_t}$, reflecting prior beliefs too, so a substantial amount of information that can be particularly useful during turning points and other challenging times for forecasting and nowcasting. Different clients can utilize this information to form their own forecasts depending on their own loss functions[6].

[6] Of course this assumes that the clients have the same priors and model, which is not necessarily the case. However, Geweke and Whiteman (2006) discuss a reweighing methodology that remote clients can use to

Moreover, in our context Bayesian estimation is implemented with a Gibbs sampling simulation algorithm and has several additional desirable features, that are associated with the computational simplifications resulting from the conditional block structure of the Gibbs sampler. Time varying parameters or nonlinear features can be introduced by appropriately augmenting the Gibbs sampler with additional blocks. Such additions could be computationally cumbersome or infeasible with classical estimation techniques. For example, Kim and Nelson (1998) discuss how Hamilton's (1989) regime switching cannot be embedded in practice within an exact Kalman filtering framework when the time series dimension is moderate or high. Furthermore, the Gibbs sampler can accommodate data augmentation steps, whereby missing values are filled in with simulated data that are generated from their-model implied conditional distribution, resulting in a balanced panel. The following section provides the details of this and of the entire algorithm that we implement.

3 Estimation Using a Gibbs Sampling Algorithm

For reasons such as the ones discussed above, Gibbs Sampling techniques are becoming increasing popular in macroeconomics and other fields. They are useful in cases where the set of a model's variables can be broken into appropriately chosen subsets, that is subsets from the conditional distributions of which we know how to sample. They basically entail generating, through Monte Carlo iterations, simulated samples from the joint distribution (conditional on the data) of model variables by drawing samples from the conditional distributions of subsets of variables (conditioning on the other subsets and the data). After a sufficient number of iterations the algorithm converges[7] and then the joint and marginal distributions of the model variables can be approximated arbitrarily well by their simulated counterparts (see, e.g. Gelfand and Smith (1990), and Robert and Casella (1999)).

adjust priors without having to repeat all the estimation work.

[7]Establishing convergence is a crucial part of the whole task and several approaches have been proposed in the literature to detect convergence. Here we follow McCulloch and Rossi (1994) and we compare the empirical posterior distributions that we get based on increasing number of iterations, with little changes in these distributions (after the additional iterations) serving as evidence of convergence.

In the specific context of the model described in the previous section, the Gibbs sampler we implement consists of blocks for the following nine groups of model variables: Σ^{ξ}, μ and Λ, A, Σ^{ζ}, $\{F_1, ..., F_{t+h}\}$, x^{miss}, $\theta = (\alpha, \beta')$, σ^2_{ε}, $\{y^f_t, ..., y^f_{t+h}\}$.

Note that following Carter and Kohn (1994) we generate all the factors in one multi-move Gibbs sampling block, which is computationally more efficient (when compared to single-move alternatives) and also converges faster. Furthermore, note that x^{miss} stands for the missing values, typically found at the end of jagged-edge data set $x_{|\tau_t}$ on day τ_t. We generate these in a data augmentation step following Tanner and Wong (1987). Finally, note that for each of the blocks, we condition on the simulated samples from the other subsets of model variables, as well as on $\Upsilon_{|\tau_t}$, which consists of $x_{|\tau_t}$ and of the dependent variable y as this is available on day τ_t. Quite clearly, the implication here is that this estimation can be performed every day τ_t when new information becomes available and all variable estimates and forecasts reflect the information that is available on day τ_t.

In what follows we provide all the relevant details in turn regarding how each group of model variables are generated, suppressing all conditioning for notational convenience:

3.1 Generating μ and Λ

Given $F = \begin{pmatrix} F_1' \\ \vdots \\ F_t' \end{pmatrix}$, we can focus on Equation (1) in order to generate μ, Λ, and Σ^{ξ} and then we have a series (for each of the N variables) of standard Bayesian linear regressions for which the standard Normal/Inverted Gamma setup for priors and posteriors is available.

More specifically, given $\Sigma^{\xi} = diag(\sigma^2_{\xi_1}, ..., \sigma^2_{\xi_N})$, we can rewrite Equation (1) in stacked form for periods $1, ..., t$, and for each of the N variables:

$x_i = \mu_i + F\Lambda_i + \xi_i = \Phi\Lambda_i^* + \xi_i$, $i = 1, ..., N$, where $\Phi = \binom{\iota'}{F'}'$ is an $t \times (r+1)$ matrix with ι being a $t \times 1$ vector of 1s, x_i is a $t \times 1$ vector with the observations for variable i (as they are available on day τ_t), μ_i and ξ_i are also $t \times 1$, and $\Lambda_i^* = \binom{\mu}{\Lambda_i}$ is an $(r+1) \times 1$ vector.

Then, if the prior for Λ_i^* is $N(\Lambda_{i0}^*, \Sigma_0^{\Lambda_i^*})$, where Λ_{i0}^* and $\Sigma_0^{\Lambda_i^*}$ are known, its posterior will also be normal, $N(\Lambda_{i1}^*, \Sigma_1^{\Lambda_i^*})$, where $\Lambda_{i1}^* = ((\Sigma_0^{\Lambda_i^*})^{-1} + (\sigma_{\xi i}^2)^{-1}\Phi'\Phi)^{-1}((\Sigma_0^{\Lambda_i^*})^{-1}\Lambda_{i0}^* + (\sigma_{\xi i}^2)^{-1}\Phi' x_i)$, and $\Sigma_1^{\Lambda_i^*} = ((\Sigma_0^{\Lambda_i^*})^{-1} + (\sigma_{\xi i}^2)^{-1}\Phi'\Phi)^{-1}$.

3.2 Generating Σ^ξ

Given μ and Λ, and in order to generate the variances of the $N \times N$ diagonal matrix Σ^ξ, we proceed as follows:

If the prior for $\sigma_{\xi_i}^2, i = 1, ..., N$, is an inverted gamma distribution $IG(\frac{\nu_0^{\xi_i}}{2}, \frac{\delta_0^{\xi_i}}{2})$, where $\nu_0^{\xi_i}$ and $\delta_0^{\xi_i}$ are known, the posterior for $\sigma_{\xi_i}^2, i = 1, ..., N$ will also be inverted gamma, $IG(\frac{\nu_1^{\xi_i}}{2}, \frac{\delta_1^{\xi_i}}{2})$, where $\nu_1^{\xi_i} = \nu_0^{\xi_i} + t$, and $\delta_1^{\xi_i} = \delta_0^{\xi_i} + (x_i - \Phi\Lambda_i^*)'(x_i - \Phi\Lambda_i^*)$, with Φ as defined above. Quite clearly then Σ^ξ is the diagonal matrix whose i^{th} entry is $\sigma_{\xi_i}^2, i = 1, ..., N$.

3.3 Generating A

Given F, and since ζ_t is independent from the idiosyncratic error terms, we can focus on the Bayesian VAR context of Equation (2) in isolation from the measurement equation, and generate A and Σ^ζ by employing a pair of Normal-Wishart priors:

Specifically, and given Σ^ζ, we can rewrite Equation (2) in stacked form as follows:

$F^\dagger = (I_r \otimes \Gamma)\alpha + \zeta$, where $F^\dagger = \begin{pmatrix} F^1 \\ \vdots \\ F^r \end{pmatrix}$, i.e. $F^\dagger$ is an $r(t-1) \times 1$ vector which stacks all $2, ..., t$ observations of the first factor, then all $t-1$ observations of the second factor, and so on until the r^{th} factor, $\Gamma = \begin{pmatrix} F_1' \\ \vdots \\ F_{t-1}' \end{pmatrix}$, i.e. Γ is a $(t-1) \times r$ matrix, $\alpha_A = vec(A)$, which is an $r^2 \times 1$ vector that stacks all the coefficients of A into a vector, and ζ is the $r(t-1) \times 1$ vector stacking the corresponding common shocks.

Then, if the prior for α_A is $N(\alpha_0, \Sigma_0^\alpha)$[8], where α_0 and Σ_0^α are chosen by the modeler, the

[8]Note that the stationarity condition that all the roots of $\det(I_r - Az)$ lie outside the unit circle can be

resulting posterior will be $N(\alpha_1, \Sigma_1^\alpha)$, where $\alpha_1 = [(\Sigma_0^\alpha)^{-1} + ((\Sigma^\zeta)^{-1} \otimes (\Gamma'\Gamma))]^{-1}[(\Sigma_0^\alpha)^{-1}\alpha_0 + ((\Sigma^\zeta)^{-1} \otimes \Gamma)'F^\dagger]$,

and $\Sigma_1^\alpha = [(\Sigma_0^\alpha)^{-1} + ((\Sigma^\zeta)^{-1} \otimes (\Gamma'\Gamma))]^{-1}$. Finally, A can be constructed by unstacking α_A.

3.4 Generating Σ^ζ

Given A, we proceed as follows:

If the prior for $(\Sigma^\zeta)^{-1}$ is Wishart $W(S_0^{-1}, \nu_0^\zeta)$, where S_0^{-1}, and ν_0^ζ are known, the posterior for $(\Sigma^\zeta)^{-1}$ will also be Wishart, $W(S_1^{-1}, \nu_1^\zeta)$, where $\nu_1^\zeta = \nu_0^\zeta + (t-1)$,

and $S_1 = S_0 + \sum_{j=2}^{t}(F_j - AF_{j-1})(F_j - AF_{j-1})'$.

3.5 Generating $\{F_1, ..., F_{t+h}\}$

Equations (1) and (2) constitute a state space model, so given the parameters, $\{F_1, ..., F_{t+h}\}$ are generated in one multi-move block as in Carter and Kohn (1994) by applying a Kalman smoothing algorithm as Kim and Nelson (1999) illustrate.

Specifically, for $j = 1, ..., t$ we have the following prediction equations going forward:

$F_{j|j-1} = AF_{j-1|j-1}$

$P_{j|j-1} = AP_{j-1|j-1}A' + \Sigma^\zeta$

$\eta_{j|j-1} = x_j - x_{j|j-1} = x_j - \mu - \Lambda F_j$

$f_{j|j-1} = \Lambda P_{j|j-1}\Lambda' + \Sigma^\xi$

Given that we are in a stationary environment we initiate the above with the unconditional mean $F_{0|0}$ and variance $P_{0|0}$. Now, given the Kalman gain $K_j = P_{j|j-1}\Lambda' f_{j|j-1}^{-1}$ we have the following updating equations:

$F_{j|j} = F_{j-1|j-1} + K_j\eta_{j|j-1}$

$P_{j|j} = P_{j|j-1} - K_j\Lambda P_{j|j-1}$

satisfied by choosing a prior that assigns 0 probability to the non-stationary part of the support space for A. The posterior will then also satisfy the stationarity condition.

We then draw F_t from $N(F_{t|t}, P_{t|t})$ and going backwards, for $j = T-1, ..., 1$, we draw $F_j|F_{j+1}$ from $N(F_{j|j,F_{j+1}}, P_{j|j,F_{j+1}})$ where

$$F_{j|j,F_{j+1}} = F_{j|j} + P_{j|j}A'(AP_{j|j}A' + \Sigma^{\zeta})^{-1}(F_{j+1} - AF_{j|j})$$

$$P_{j|j,F_{j+1}} = P_{j|j} - P_{j|j}A'(AP_{j|j}A' + \Sigma^{\zeta})^{-1}AP_{j|j}$$

3.6 Generating x^{miss}

This is the data augmentation step which entails generating simulated values for missing $x_{ij|\tau_t}$, $i = 1, ..., N$, $j = 1, ..., t+h$ given the information set on day τ_t. Given $\{F_1, ..., F_{t+h}\}$, μ, Λ, and Σ^{ξ}, simulated values for x^{miss} can be obtained from the left hand side variable of Equation (1) for periods $1, ..., t+h$.

3.7 Generating $\theta = (\alpha, \beta')$

Note that given $F = \begin{pmatrix} F_1' \\ \vdots \\ F_t' \end{pmatrix}$, we can focus on Equation (3) in order to generate $\theta = (\alpha, \beta')$ and σ_{ε}^2 and this collapses again to a standard Bayes regression for which we employ again the standard Normal/Inverted Gamma pair of priors.

First, and given σ_{ε}^2, let's rewrite Equation (3) in stacked form for periods $1, ..., t$: $y = \alpha + F\beta + \varepsilon = \Phi\theta' + \varepsilon$, where $\Phi = \binom{\iota'}{F'}'$ is an $t \times (r+1)$ matrix with ι being a $t \times 1$ vector of 1s.

Then, if the prior for $\theta = (\alpha, \beta')$ is $N(\theta_0, \Sigma_0^{\theta})$, where θ_0 and Σ_0^{θ} are known, its posterior will also be normal, $N(\theta_1, \Sigma_1^{\theta})$, where $\theta_1 = ((\Sigma_0^{\theta})^{-1} + (\sigma_{\varepsilon}^2)^{-1}\Phi'\Phi)^{-1}((\Sigma_0^{\theta})^{-1}\theta_0 + (\sigma_{\varepsilon}^2)^{-1}\Phi' y)$, and $\Sigma_1^{\theta} = ((\Sigma_0^{\theta})^{-1} + (\sigma_{\varepsilon}^2)^{-1}\Phi'\Phi)^{-1}$.

3.8 Generating σ_{ε}^2

Second, and given θ, if the prior for σ_{ε}^2 is an inverted gamma distribution $IG(\frac{\nu_0^{\varepsilon}}{2}, \frac{\delta_0^{\varepsilon}}{2})$, where ν_0^{ε} and δ_0^{ε} are known, its posterior will also be inverted gamma, $IG(\frac{\nu_1^{\varepsilon}}{2}, \frac{\delta_1^{\varepsilon}}{2})$, where $\nu_1^{\varepsilon} = \nu_0^{\varepsilon} + t$, and $\delta_1^{\varepsilon} = \delta_0^{\varepsilon} + (y - \Phi\theta')'(y - \Phi\theta')$, with Φ as defined above.

3.9 Generating $\{y_t^f, ..., y_{t+h}^f\}$

Given $\{F_1, ..., F_{t+h}\}$, θ, and σ_{ε}^2, estimated from the blocks above conditional on the information set of day τ_t, $y_t^f, ..., y_{t+h}^f$ emerge simply as the left hand side variable from Equation (3) for periods $t, ..., t+h$.

4 An Application: Nowcasting US GDP During the "Great Recession"

We employ the model and estimation algorithm described in the previous sections to nowcast US GDP growth, with a focus on the latest recession. As is well known, the Bureau of Economic Analysis does not release its estimates of current quarter GDP growth until next quarter; it releases a preliminary ("Advance") estimate towards the end of the first month of the following quarter, and then it updates this figure one and two months later ("Second" and "Third" estimates, respectively). This constitutes a significant lag, especially for monetary policy purposes. Thus the task of nowcasting GDP growth (as well as other aggregates subject to similar lags in releases), of obtaining that is current-quarter GDP growth estimates during the current quarter, becomes particularly relevant, and the attention that the burgeoning nowcasting literature has been receiving is well justified.

We adopt here the perspective of a monetary policy maker, or of a professional forecaster, who needs to assess the current state of the economy in real time as accurately as possible and given all the available information. Assuming a formal modeling approach is adopted, it is arguably desirable to consider frameworks that, as discussed earlier, can handle large data sets, with jagged edges, and possibly mixed frequencies. Furthermore, if we hope to provide a realistic depiction of the nowcasting environment, with the regular influx of possibly inaccurate real-time information, we should be employing real-time data, rather than revised series that only became available ex-post (see, inter alia, Croushore and Stark (2001) and Orphanides (2001)). So, in this section we use a real-time data set - indeed a series of weekly, Friday jagged-edge data sets reflecting all the updates (new observations or revisions of existing observations) that took place during the week. There are close to 200 (mostly monthly) macroeconomic variables, including monetary aggregates, prices, employment statistics, survey data, housing, banking balance sheet figures, etc. They include observations starting in January of 1982, with the weekly real time vintages starting in March of 2005[9].

We perform stationarity-inducing transformations and standardize the data. Furthermore, and since GDP growth is a quarterly quantity, whereas most of the macro series that are used to obtain the factors are monthly, we follow GRS and filter the data using the Mariano-Murasawa (2003) approach. We also choose the same number of factors (two) and number of common shocks (also two) as they do. One could consider alternatives for all these modeling choices[10], but one of our goals is to make our results directly comparable to theirs. Furthermore, such specification choices are quite standard in the literature. Examples where one factor is used include Stock and Watson (1999), and the Chicago Fed's National Activity Index - CFNAI (Federal Reserve Bank of Chicago (2012)) where the factor is extracted from a smaller set of variables capturing economic activity[11]. Here we consider a wider information set including various price variables.

4.1 Uninformed Priors

We begin our investigation by considering *uninformed* priors, whose intent is to convey a largely agnostic approach towards prior beliefs. These priors (distributions and values for hyperparameters for the various blocks discussed in the previous section) are summarized in Table 1.

[9]The data, which were kindly provided by David Small, are used by Giannone et al (2010), so the reader is referred to that study, as well as to GRS for more details.

[10]We estimated alternative specifications with one factor as well ($r = q = 1$), but chose to focus on the $r = q = 2$, case for the reasons discussed here.

[11]The CFNAI is based on 85 variables on production and income, employment, personal consumption and housing, and sales, orders and inventories.

Our approach towards establishing convergence of the Gibbs sampler is that of McCulloch and Rossi (1994): We compare the simulated posterior distributions we obtain with different numbers of iterations, with only small changes serving as evidence of convergence. Figures 2 and 3 summarize some of these exercises: they plot means of the posterior distributions for the fitted values of GDP growth and for the variances of the idiosyncratic error terms (Equation (1)) using 1000 burn-in iterations (to be discarded) and 2000 iterations to be retained (on the basis of which the empirical posterior distributions are obtained), and respective means using 2000-10000 iterations. In both cases the lines (corresponding to low and high numbers of iterations) are almost indistinguishable from each other. Figure 4 has two panels showing the entire posterior distributions (based on the low and high numbers of iterations) for the variance of the error term from the bridge regression (Equation (3)). Again the two distributions are almost identical, and the two overall means are 2.598 and 2.593 respectively.

When it comes to overall fit, Figure 5 plots the means of fitted values we obtain with our Bayesian model, the respective ones using the GRS approach (using the same data, number of factors and common shocks), and actual GDP growth. Two observations are in order here: First, the two models deliver roughly similar results, with comparable fits and no clear evidence of either approach dominating in terms of in-sample fit. So, if the objective is to employ means as the nowcasts, there may be little reason in this particular case to choose one model over the other. Second, the overall fit is generally good[12], with the exception of sharp drops or rises of GDP and business cycle turning points such as the 1990-91 recession, for instance, during which times the models may undershoot.

We investigate this issue a little more closely by focusing on the latest episode of high volatility, the Great Recession of 2007-2009. We produce several real-time nowcasts of GDP growth during a period that includes this recession (specifically, 2007Q1-2009Q4) by adjusting our series of data vintages described above so as to re-create the information sets faced by

[12] GRS provide convincing evidence that the model fares well when compared to Federal Reserve Greenbook data and forecasts obtained from the Survey of Professional Forecasters.

a nowcaster in real time[13]. We can produce such estimates for any day of the week, and Figure 6 plots (advance releases of) GDP growth together with its nowcasts obtained on the first business day of the month (for every month of the quarter), that is the estimates that a nowcaster would have been able to compute with the GRS model or with our Bayesian model, and using the information that was available to her on that day[14]. While the model generally does a decent job tracking the movements of GDP growth, it does not do well for the quarter during which GDP growth turns negative for the first time in a series of consecutive quarters, namely 2008 Q3. Indeed, if our criterion is both how far the nowcasts are from the actual figures, and how much of an improvement one can see in successive nowcasts as the quarter progresses and more information gathers on the state of the economy, then 2008 Q3 is by far the worst quarter: The nowcasts are positive and far from the negative GDP growth for the quarter, and stay like that during the entire quarter and even in October of 2008.

This perhaps may come as little surprise. As was discussed in the introduction, we expect that many models will not perform as well during turning points as during other times. Indeed, the pessimistic conclusion that one may be tempted to reach is that model and data uncertainties are insurmountable challenges around turning points, and that this is true even when nowcasting in real time. However, in what follows we present a more optimistic picture in terms of what can be achieved by exploiting internet-based tools that were not available until recently, such as Google Trends.

4.2 Popularity Priors

At the core of our approach is the realization that when aggregate growth turns from positive to negative, this typically reflects a widespread slowdown in economic activity, which thus

[13] This requires some tedious manual manipulations whereby we add or delete observations from the Friday weekly data sets (for one or more variables, and on the basis of the stylized schedule of data releases), so as to create the real-time information set of any day of the week for which we wish to estimate the real-time nowcast of GDP growth.

[14] Note that the figure includes four such beginning-of-month estimates per quarter, as the nowcast produced on the first business day of the first month of the following quarter pertains to information covering the quarter that just ended. Furthermore, the official "Advance" estimate of GDP growth for the quarter that just ended will not be available until the last week of the month that just started.

affects a wide spectrum of people, either directly or indirectly. Given the prevalence of the internet in our time, many people may turn to internet search engines in an attempt to better understand the economy and their changing conditions. They may search related keywords in higher volumes than at other times. Tools such as Google Trends measure the volume of keyword-based searches and have begun compiling and making publicly available normalized time series indices based on the volume of searches. There are already interesting forecasting/nowcasting applications using such indices including Askitas and Zimmermann (2009), D'Amuri and Marcucci (2010) and Scott and Varian (2014a,b). The growing interest is certainly justified given that such indices have several key features, including that they reflect relative, and not absolute, volumes of searches, that they are based on real-time information, and that they, for certain keywords, may be able to capture widespread changes in financial and economic conditions in a timely manner, without being plagued by real-time inaccuracies and revision issues.

Helpful keywords in the present context would lead to Google indices that stay relatively flat at times when models such as the ones discussed above nowcast well, and that spike up when we do not nowcast well, primarily in the fall of 2008. An example of such a helpful keyword is "recession", whose Google index is shown in Figure 7. We can notice there that it spikes up not only in the fall of 2008, but also in the period around the end of 2007 and the beginning of 2008 (which includes the NBER peak).

A policy maker/nowcaster, well aware of the above, may conclude that "listening to the people too" in addition to consulting information sets based on "traditional" large data sets, is a promising nowcasting approach, with the potential to deliver superior results around turning points. Our approach allows such a nowcaster to take both into account. She lets internet search popularity measures such as Google Trends inform her prior beliefs. Thus her nowcasts reflect both such "popularity priors" and information coming from the dynamic factors and the large data sets as discussed above.

This approach can be implemented in various ways. In the example that follows, we propose one such possibility, using the "recession" keyword, and a simple, conservative approach towards constructing priors. The model and all the priors remain the same as above (see Table 1), with the only one changing being that of α, the intercept of equation 3. Specifically, the mean of the prior of α is set to the post-WWII recession (duration-weighted) average. Its prior variance follows the schedule below:

$$Variance(\alpha) : \begin{cases} 1 & when\, GRI\, is\, between\, 0\, and\, 9 \\ 0.9 & when\, GRI\, is\, between\, 10\, and\, 19 \\ \vdots & \vdots \\ 0.1 & when\, GRI\, is\, between\, 90\, and\, 99 \\ 0.05 & when\, GRI\, is\, 100 \end{cases} \tag{4}$$

Quite clearly, this postulates an increasing level of certainty that we are indeed in a recession as *GRI, the Google Recession Index* (Figure 7) increases.

The resulting real-time nowcasts obtained with this model are provided in Figure 8, and are contrasted with those coming from the other models discussed above. The key finding is that the Bayesian model with the above popularity prior delivers a more timely recognition of the turning point in GDP growth in the fall of 2008, and substantially so, as the nowcast now turns negative by the end of Q3/beginning of Q4 of 2008[15]. It is interesting to note here that if the benchmark by which nowcasts are to be judged is advance GDP releases (depicted in Figure 8), then the nowcast in question (produced on October 1, 2008) undershoots by a significant amount. However, these estimates can be subject to significant revisions themselves, especially around turning points, while this is not a concern with the GRI. Figure 9 depicts revised GDP estimates; the revised estimate corresponding to 2008 Q3 is indeed much closer to our nowcast using popularity priors.

Furthermore, the improvement discussed in the previous paragraph, which can be quite

[15]This is the nowcast we could have obtained in the morning of October 1, 2008, and thus reflects information available up to Quarter 3. Recall that in this application we produce one nowcast per month, in the very beginning of each month. Of course, our methodology can generate updated nowcasts many times each month, any time new information becomes available.

consequential for monetary policy purposes in real time, does not come at the cost of deteriorated performance at other times, when compared to either the GRS nowcasts or the Bayesian nowcasts with the uninformed priors. Indeed, a policy maker with popularity priors would have actually been closer to the "truth" more often during the three years in question (56.3% of the time closer to advance GDP estimates and 63% of the time closer to revised GDP estimates) than she would have been had she relied exclusively on the traditional data. Table 2 provides the nowcast errors (with respect to both advance and revised GDP releases for 2007Q1-2009Q4) associated with both GRS and the model with popularity priors, as well as respective ratios of Mean Square Errors and corresponding Diebold Mariano (1995) t-statistics. Quite clearly, we cannot reject the null hypothesis that the MSEs coming from the two alternative approaches are equal.

Taking all of the above into account, we can conclude that the Bayesian nowcasting model with popularity priors achieves a substantially better outcome regarding the consequential turning point of the fall of 2008, while its nowcasts for the rest of the time are (at least) as good as the ones obtained from alternative approaches in real time.

The above is indicative of what can be achieved using this approach, which is quite flexible, as it can accommodate other specifications as well. Another example is that of a prior for α that is informed not just by a single Google Recession Index, but by a combination of GRIs coming from other useful keywords or phrases. For example, the phrase "end of recession," whose GRI is depicted in Figure 10, could be potentially useful in improving the nowcasts obtained with the Bayesian model around the end of the recession. An alternative prior for α that combines both the recession and the "end of recession" GRIs can be as follows:

- Mean of α: $\begin{cases} recession\,average & when\,GRI_{recession} > GRI_{end\,of\,recession} \\ non-recession\,average & when\,GRI_{end\,of\,recession} > GRI_{recession} \end{cases}$

with both averages referring to the post-WWII era, and

- GRI: $\begin{cases} GRI_{recession} & when\,GRI_{recession} > GRI_{end\,of\,recession} \\ GRI_{end\,of\,recession} & when\,GRI_{end\,of\,recession} > GRI_{recession} \end{cases}$

with the variance of α being the same as above. Figure 11 depicts the nowcasts obtained using this alternative popularity prior, and as we can see there we have a modest improvement in the months around the end of the recession, when the prior shifts to the end-of-recession GRI. More specifically, if the criterion is the distance of the nowcast from the revised GDP estimates, then we have (for the 13 nowcasts for which the end-of-recession GRI kicks in) an average improvement of 0.25% and a median improvement of 0.30% of GDP growth.

Of course, these are just two examples of popularity priors and alternative ones are possible, and regarding not just α, but other model parameters as well.

5 Concluding Remarks

An important and rapidly expanding literature has made a strong case for resorting to large data sets when forecasting or nowcasting macroeconomic aggregates such as GDP, and often recommends basing those nowcasts on factor models, that avoid the proliferation of parameters and efficiently summarize the comovements and dynamics in the data without discarding essential information. From a policy maker's or a macro forecaster's perspective, one implication of looking at many, possibly hundreds of variables in real time is that an update to the existing forecast or nowcast can be produced every time new information arrives (in the form of new observations or revisions to existing observations). This amounts to updating model estimates and resulting forecasts conditional on the latest data at any point in time, and we view a Bayesian approach such as the one we propose in this paper as a natural way to proceed in such contexts.

Furthermore, such a Bayesian approach allows the policy maker to move beyond the strict confines of traditional data sets, and to do so in a systematic and quantifiable way, by allowing her prior beliefs to be influenced by internet search popularity indices. Such measures can provide a useful early gauge of important and widespread changes in the economy and can thus mitigate the problems associated with model and data uncertainty in real time, which can be particularly pernicious around turning points. The nowcasting example discussed in this study of the US GDP turning point in the fall of 2008 is telling.

References

Andreou, E., E. Ghysels, and A. Kourtellos (2010): "Regression models with mixed sampling frequencies," *Journal of Econometrics*, 158, 246-61.

Askitas, N. , and K. Zimmermann (2009): "Google Econometrics and Unemployment Forecasting," *IZA Discussion Paper* No. 4201.

Banbura, M., D. Giannone, and L. Reichlin (2011): "Nowcasting," in Michael P. Clements and David F. Hendry, editors, *Oxford Handbook on Economic Forecasting*, forthcoming.

Bernanke, B., J. Boivin, and P. Eliasz (2005): "Measuring Monetary Policy: A Factor Augmented Autoregressive (FAVAR) Approach," *Quarterly Journal of Economics*, 120: 387–422.

Boivin, J. and M.P. Giannoni (2005): "DSGE Models in a Data-Rich Environment," *mimeo, Columbia University.*

Camacho, M. and G. Perez-Quiros (2010): "Introducing the EURO-STING: Short Term Indicator of Euro Area Growth," *Journal of Applied Econometrics* (forthcoming).

Carter, C.K., and P. Kohn (1994): "On Gibbs Sampling for State Space Models," *Biometrika* 81, 541-553.

Croushore, D. and T. Stark (2001): "A Real-Time Data Set for Macroeconomists," *Journal of Econometrics* 105(1), 111-130.

D' Agostino, A., D. Giannone, M. Lenza and M. Modugno (2015): "Nowcasting Business Cycles: a Bayesian Approach to Dynamic Heterogeneous Factor Models," *Finance and Economics Discussion Series* 2015-066. Washington: Board of Governors of the Federal Reserve System.

D'Amuri, F. and J. Marcucci (2010): "Google it! Forecasting the US Unemployment Rate with a Google Job Search Index," *Nota di Lavoro della Fondazione Eni Enrico Mattei* 31.2010.

Doz, C., D. Giannone, D. and L. Reichlin (2011): "A Two-step Estimator for Large Approximate Dynamic Factor Models Based on Kalman Filtering," *Journal of Econometrics* (164), 188-205.

Doz, C., D. Giannone, D. and L. Reichlin (2012): "A Quasi–Maximum Likelihood Approach for Large, Approximate Dynamic Factor Models," *Review of Economics and Statistics* 94(4), 1014-1024.

Federal Reserve Bank of Chicago (2012): "Background on the Chicago Fed National Activity Index," www.chicagofed.org

Gelfand, A. E. and A. F. M. Smith (1990): "Sampling-Based Approaches to Calculating Marginal Densities," *Journal of the American Statistical Association*, June 1990, 85(410), 398-409.

Geweke, J., and C. Whiteman (2006): "Bayesian Forecasting," in Elliott, G., C. W. J. Granger, and A. Timmermann (Eds.) *Handbook of Economic Forecasting, Vol. 1*, Elsevier, 3-80.

Ghysels, E., P. Santa-Clara, and R. Valkanov (2004): "The MIDAS Touch: Mixed Data Sampling Regression Models," *discussion paper, UN Chapel Hill and UCLA*.

Giannone, D., M. Modugno, L. Reichlin, and D. Small (2010): "Nowcasting in Real-Time," *mimeo*, presented at the *16th International Conference on Computing in Economics and Finance, Society for Computational Economics*, London, UK. July 13.

Giannone, D., L. Reichlin, and L. Sala (2004): "Monetary Policy in Real Time," in: Gertler, Mark, Rogoff, Kenneth (Eds.), *NBER Macroeconomics Annual.* MIT Press, 161–200.

Giannone, D., L. Reichlin, and D. Small (2008): "Nowcasting: The Real-Time Informational Content of Macroeconomic Data," *Journal of Monetary Economics* (55), 665-676.

Hamilton, J.D. (1989): "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle," *Econometrica* 57(2), 357-384.

Kim, Chang-Jin, and Charles R Nelson (1998): "Business Cycle Turning Points, a New Coincident Index, and Tests of Duration Dependence Based on a Dynamic Factor Model with Regime Switching," *The Review of Economics and Statistics* 80, 188-201.

Kim, C., and C. Nelson (1999): *State-Space Models With Regime Switching*, MIT Press, Cambridge.

Kose, M.A, C. Otrok, and C.H. Whiteman (2003): "International Business Cycles: World, Region, and Country-Specific Factors," *American Economic Review*, 93: 1216–1239.

Kuzin, V., M. Marcellino, and C. Schumacher (2011): "MIDAS vs Mixed-Frequency VAR for Nowcasting GDP in the Euro Area," *International Journal of Forecasting* 27, 529-542.

Lahiri, K., and J.G. Wang (2013): "Evaluating Probability Forecasts for GDP Declines Using Alternative Methodologies," *International Journal of Forecasting* 29(1), 175-190.

Loungani, P. (2001): "How Accurate are Private Sector Forecasts? Cross-country Evidence from Consensus Forecasts of Output Growth," *International Journal of Forecasting* 17(3), 419-432.

Marcellino, M., and C. Schumacher (2010): "Factor-MIDAS for Now- and Forecasting with Ragged-Edge Data: A Model Comparison for German GDP," *Oxford Bulletin of Economics and Statistics* (72), 518-550.

Mariano, R., and Y. Murasawa (2003): "A New Coincident Index of Business Cycles Based on Monthly and Quarterly Series," *Journal of Applied Econometrics*, 18, pp. 427-443.

McCulloch, R., and P.E. Rossi. (1994): "An Exact Likelihood Analysis of the Multinomial Probit Model." *Journal of Econometrics* 64: 207-240.

O'Brien, M. (2014): "How the Fed Let the World Blow up in 2008," *The Atlantic*, February 26, 2014.

Orphanides, A. (2001): "Monetary Policy Rules Based on Real-Time Data," *American Economic Review* 91(4), 964-985.

Robert, C. P., & G. Casella (1999): *Monte Carlo Statistical Methods*, New York: Springer.

Scott, S.L. and H.R. Varian (2014a): "Predicting the Present with Bayesian Structural Time Series," *Int. J. Mathematical Modeling and Numerical Optimisation,* Vol. 5, Nos. 1/2.

Scott, S.L. and H.R. Varian (2014b): "Bayesian Variable Selection for Nowcasting Economic Time Series," *mimeo*.

Stekler, H. (1972): "An Analysis of Turning Point Forecasts," *American Economic Review* 62 (4), pp. 724-729.

Stock, J.H. and M.W. Watson (1999): "Forecasting Inflation," *Journal of Monetary Economics*, 44, 293-335.

Stock, J.H. and M.W. Watson (2002a): "Macroeconomic Forecasting Using Diffusion Indexes," *Journal of Business and Economic Statistics* 20 (2), 147–162.

Tanner, M. A., and W. H. Wong (1987): "The Calculation of Posterior Distributions by Data Augmentation." *Journal of the American Statistical Association*, 82, 528-549.

Wright, J. (2013): "Evaluating Real-Time VAR Forecasts with an Informative Democratic Prior," *Journal of Applied Econometrics*, 28, pp.762-776.

Zarnowitz, V. (1986): "The Record and Improvability of Economic Forecasting," *NBER Working Paper 2099.*

Table 1: Uninformed Priors

Parameter	*Prior Distribution*	*Hyperparameter Values*
$\Lambda_i^* = \binom{\mu}{\Lambda_i}$	$N(\Lambda_{i0}^*, \Sigma_0^{\Lambda_i^*})$	$\Lambda_{i0}^* = \mathbf{0_3}$, $\Sigma_0^{\Lambda_i^*} = \mathbf{I_3}$, $i = 1, ..., N$
$\Sigma^\xi = \begin{bmatrix} \sigma_{\xi_1}^2 & & \\ & \ddots & \\ & & \sigma_{\xi_N}^2 \end{bmatrix}$	$IG(\frac{\nu_0^{\xi_i}}{2}, \frac{\delta_0^{\xi_i}}{2})$	$\nu_0^{\xi_i} = 0$, $\delta_0^{\xi_i} = 0$, $i = 1, ..., N$
$\alpha_A = vec(A)$	$N(\alpha_0, \Sigma_0^\alpha)$	$\alpha_0 = \mathbf{0_4}$, $\Sigma_0^\alpha = \mathbf{I_4}$
$(\Sigma^\zeta)^{-1}$	$W(S_0^{-1}, \nu_0^\zeta)$	$S_0 = 0$, $\nu_0^\zeta = 0$
$\theta = (\alpha, \beta')$	$N(\theta_0, \Sigma_0^\theta)$	$\theta_0 = \mathbf{0_3}$, $\Sigma_0^\theta = \mathbf{I_3}$
σ_ε^2	$IG(\frac{\nu_0^\varepsilon}{2}, \frac{\delta_0^\varepsilon}{2})$	$\nu_0^\varepsilon = 0$, $\delta_0^\varepsilon = 0$

Table 2: Real-Time Nowcast Errors

Quarter/Month	GRS - "Actual" GDP growth	Pop.Prior - "Actual" GDP growth	GRS - "Revised" GDP growth	Pop.Prior - "Revised" GDP growth
07Q1/1	1.019827935	0.707016776	1.065454	0.752642841
07Q1/2	0.949127935	0.69139779	0.994754	0.737023855
07Q1/3	1.265727935	0.803092403	1.311354	0.848718468
07Q1/4	1.792227935	1.513593497	1.837854	1.559219562
07Q2/1	-0.316299136	-0.855373827	-0.20599	-0.745064691
07Q2/2	-0.225199136	-0.950102102	-0.11489	-0.839792966
07Q2/3	0.177100864	-0.058270733	0.28741	0.052038403
07Q2/4	0.390000864	-0.056414795	0.50031	0.053894341
07Q3/1	0.12087276	-0.182532089	0.346918	0.043513151
07Q3/2	0.28407276	-0.075014371	0.510118	0.151030869
07Q3/3	-0.06862724	-0.542146098	0.157418	-0.316100858
07Q3/4	-0.67172724	-0.912242476	-0.445682	-0.686197236
07Q4/1	2.674194788	2.318856002	1.185576	0.830237215
07Q4/2	1.946494788	1.885033762	0.457876	0.396414974
07Q4/3	1.264294788	1.108352516	-0.224324	-0.380266271
07Q4/4	0.835194788	-0.080937934	-0.653424	-1.569556722
08Q1/1	1.122033874	-0.061289788	2.4440693	1.260745638
08Q1/2	0.844933874	0.957621647	2.1669693	2.279657073
08Q1/3	0.771733874	0.645340434	2.0937693	1.967375861
08Q1/4	0.264233874	0.118374067	1.5862693	1.440409493
08Q2/1	-0.56983965	-1.429657998	-0.155173	-1.014991348
08Q2/2	-0.84983965	-0.816082813	-0.435173	-0.401416163
08Q2/3	-0.93813965	-1.076642882	-0.523473	-0.661976232
08Q2/4	-0.22973965	-0.631370996	0.184927	-0.216704346
08Q3/1	2.490380027	2.140719343	4.913985	4.564324316
08Q3/2	3.083480027	2.293092087	5.507085	4.716697059
08Q3/3	3.244380027	2.385108896	5.667985	4.808713868
08Q3/4	2.378680027	-3.38130329	4.802285	-0.957698317
08Q4/1	6.382294802	0.489304074	7.876641	1.983650272
08Q4/2	4.429094802	3.129256761	5.923441	4.623602959
08Q4/3	1.825994802	-0.721237799	3.320341	0.7731084
08Q4/4	-1.385805198	-3.697672438	0.108541	-2.203326239
09Q1/1	4.251908037	-2.939522473	4.340235	-2.85119551
09Q1/2	3.656008037	-2.573499108	3.744335	-2.485172144
09Q1/3	2.550208037	-1.756664223	2.638535	-1.66833726
09Q1/4	2.245508037	0.739908037	2.333835	0.828235
09Q2/1	-0.85294911	-6.84684911	-1.1377359	-7.1316359
09Q2/2	-1.37204911	-4.361270106	-1.6568359	-4.646056896
09Q2/3	0.33285089	-2.282566995	0.0480641	-2.567353786
09Q2/4	1.50245089	-0.68084911	1.2176641	-0.9656359
09Q3/1	-0.709246958	-4.110146958	0.527015	-2.873885
09Q3/2	0.503853042	-2.351642112	1.740115	-1.115380154
09Q3/3	-0.236846958	-1.972942892	0.999415	-0.736680934
09Q3/4	-0.405946958	-1.317555104	0.830315	-0.081293146
09Q4/1	-0.197155961	-0.604612804	-0.178495	-0.585951843
09Q4/2	-0.503355961	-1.281040176	-0.484695	-1.262379215
09Q4/3	-0.939755961	-0.689476305	-0.921095	-0.670815344
09Q4/4	-1.343055961	-1.945042448	-1.324395	-1.926381487
Ratio of MSEs	1.134687308		0.773307705	
Diebold-Mariano	-0.331		0.794	

Notes: "GRS" and "Pop.Prior" refer to the nowcasts obtained with the GRS and Popularity Prior approaches, respectively. The "fourth month" of a quarter indicates the nowcast obtained at the very beginning of the first month of the following quarter. The "Ratio of MSEs" refers to the Pop.Prior/GRS ratio. "Diebold Mariano" refers to the Diebold-Mariano t-statistic testing the null hypothesis of equal mean square nowcast errors between the two approaches (GRS and Pop.Prior).

Figure 1: Real-Time Nowcasts and GDP Growth

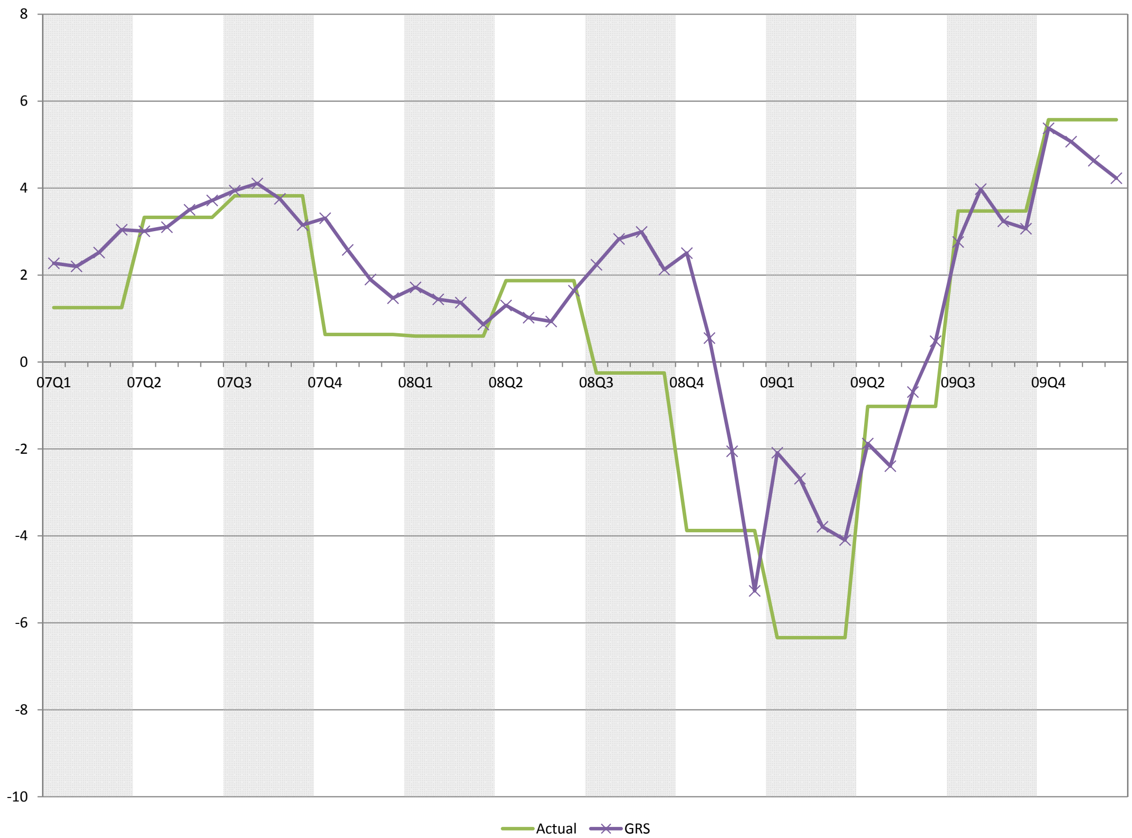

Figure 2: Posterior Means for Fitted Values of GDP Growth

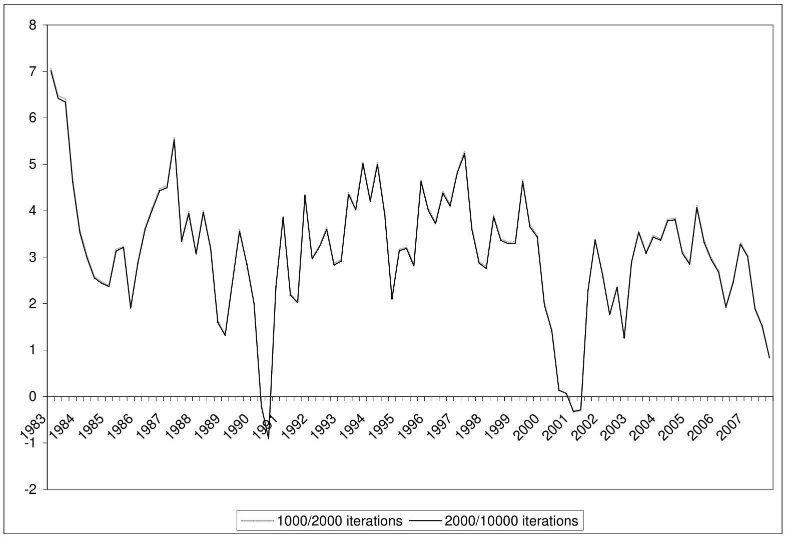

Figure 3: Posterior Means for Variances of Idiosyncratic Error Terms

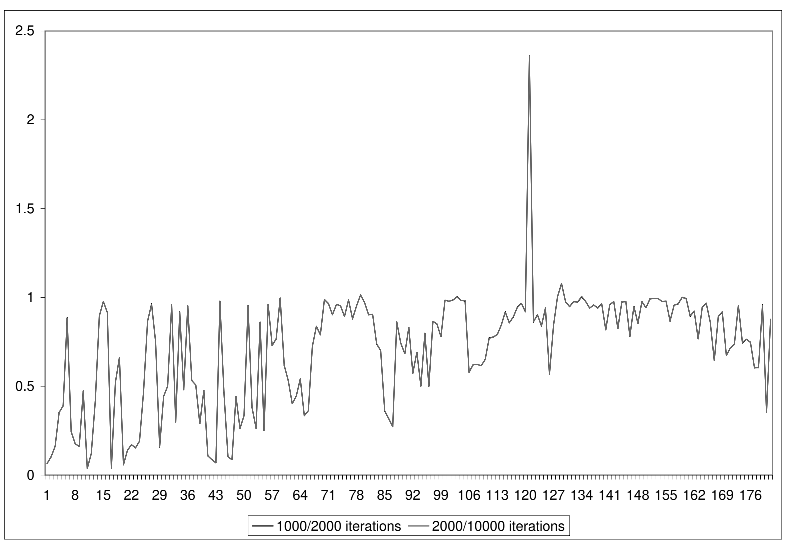

Figure 4: Posterior Distribution for σ_ε^2, Low Number of Iterations

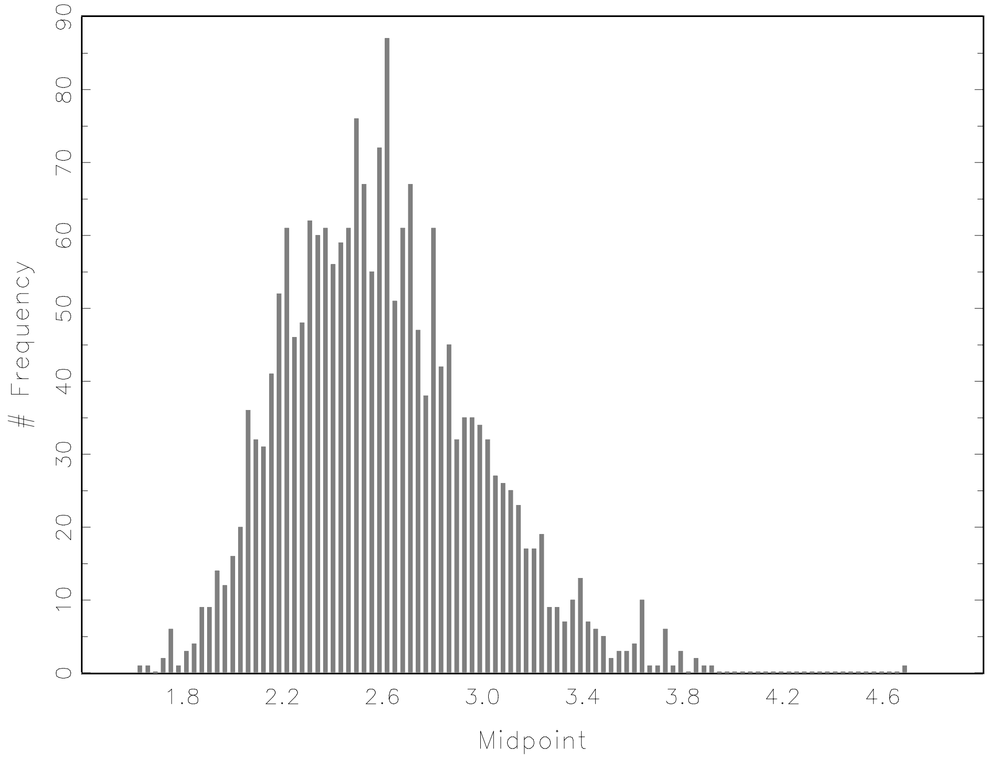

Figure 4, second panel: Posterior Distribution for σ_ε^2, High Number of Iterations

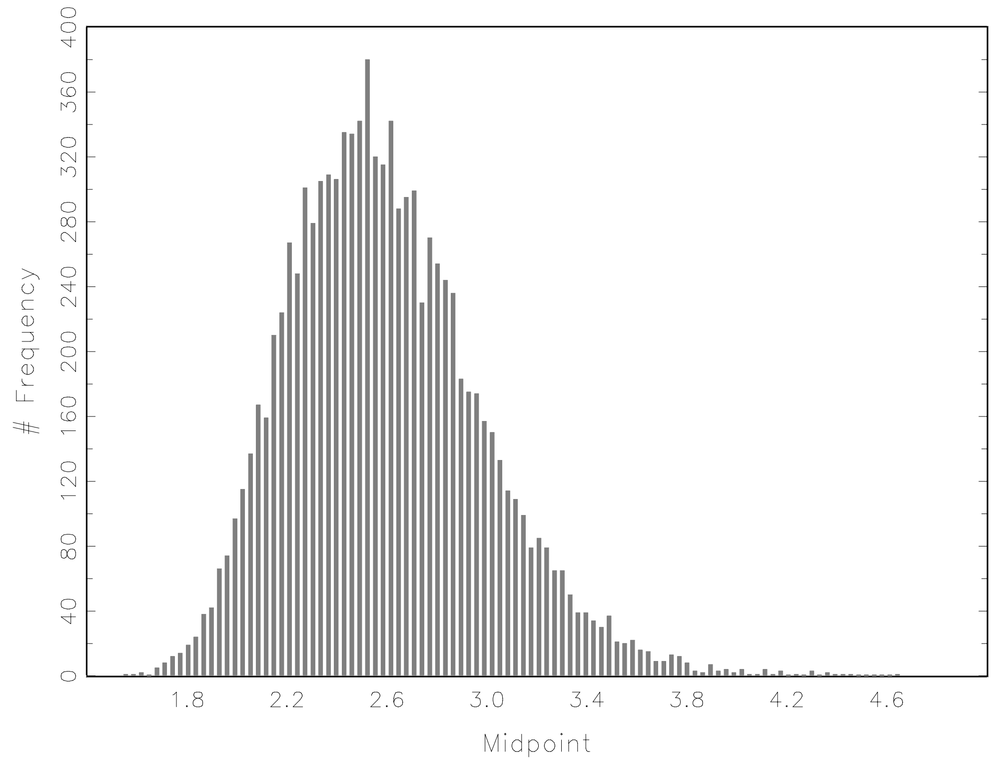

Figure 5: Fitted Values vs GDP Growth

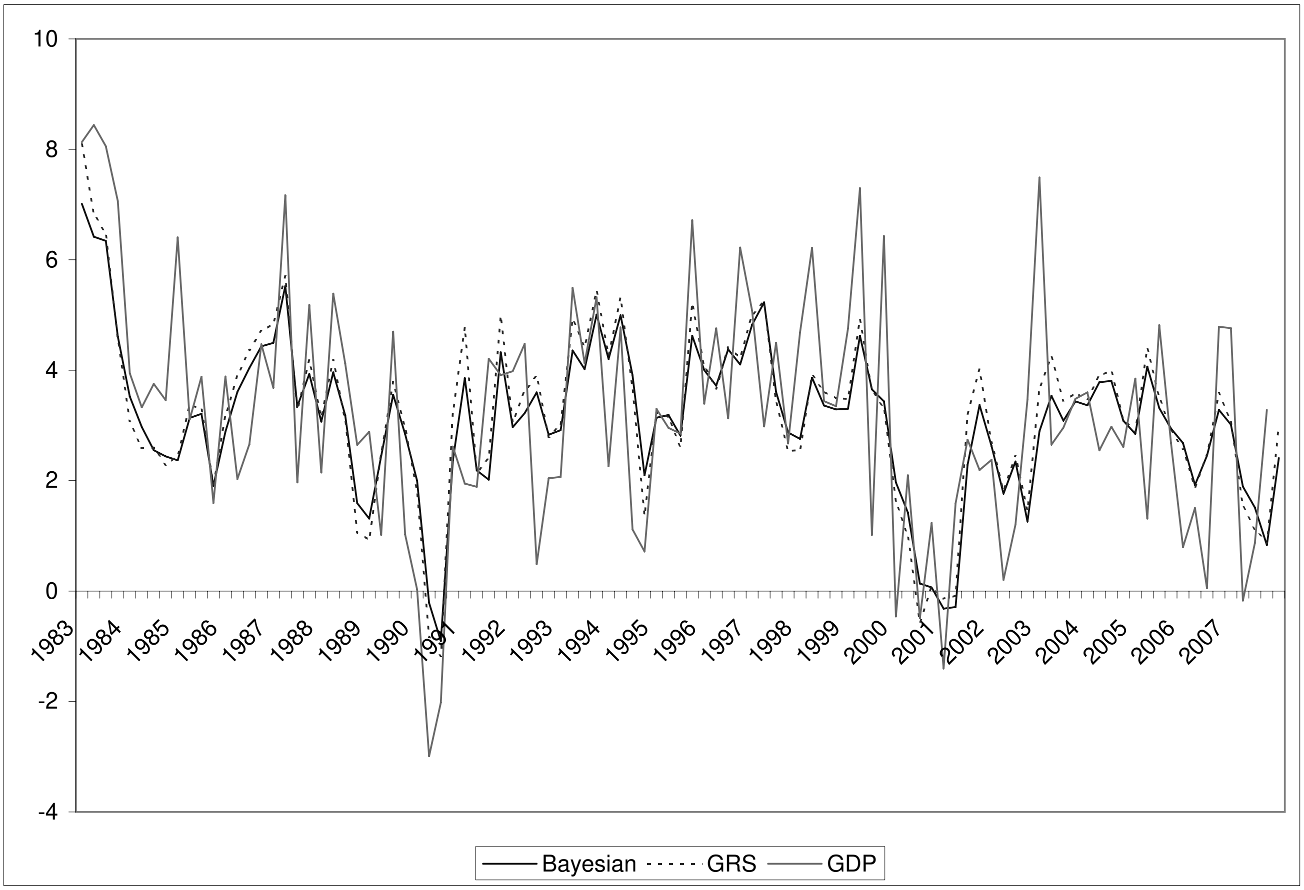

Figure 6: Real-Time Nowcasts (Bayesian and GRS) and GDP Growth

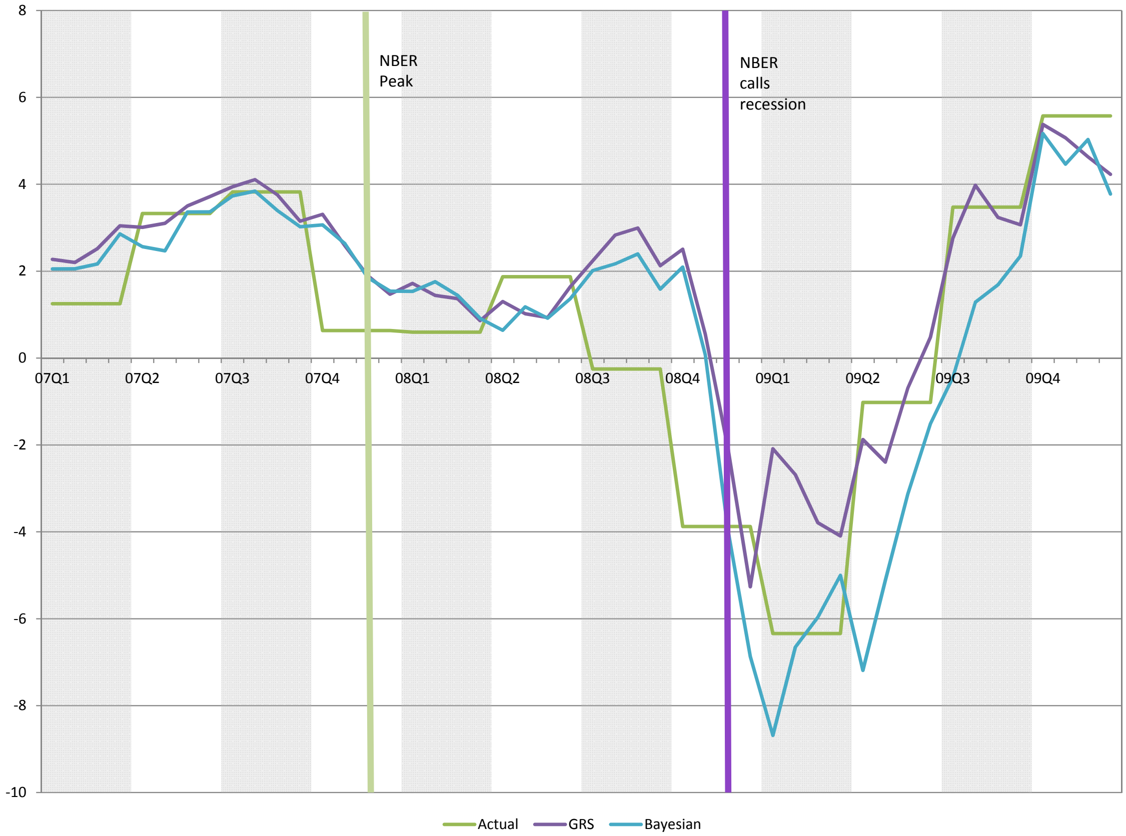

Figure 7: Google Searches for Keyword “Recession”

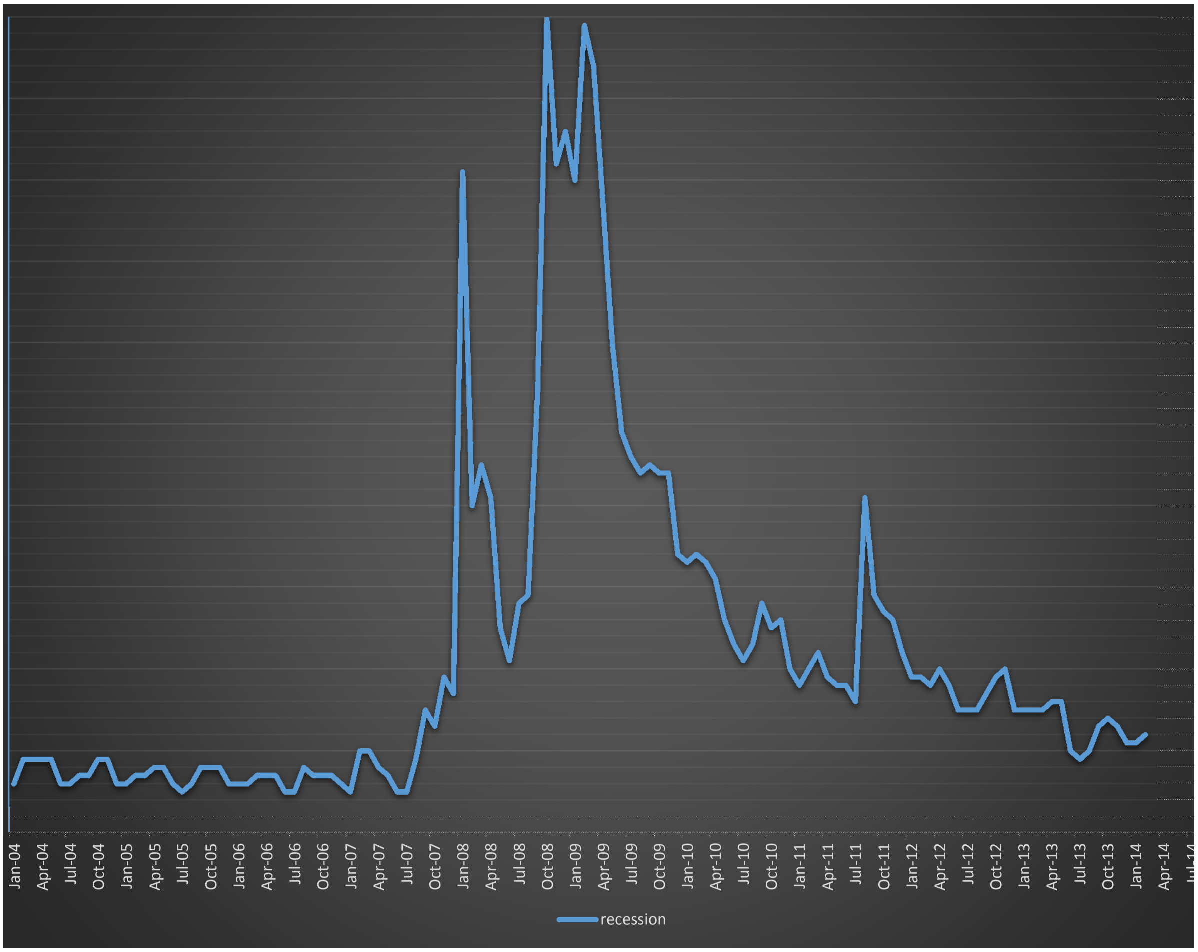

Figure 8: Real-Time Nowcasts with Popularity Priors

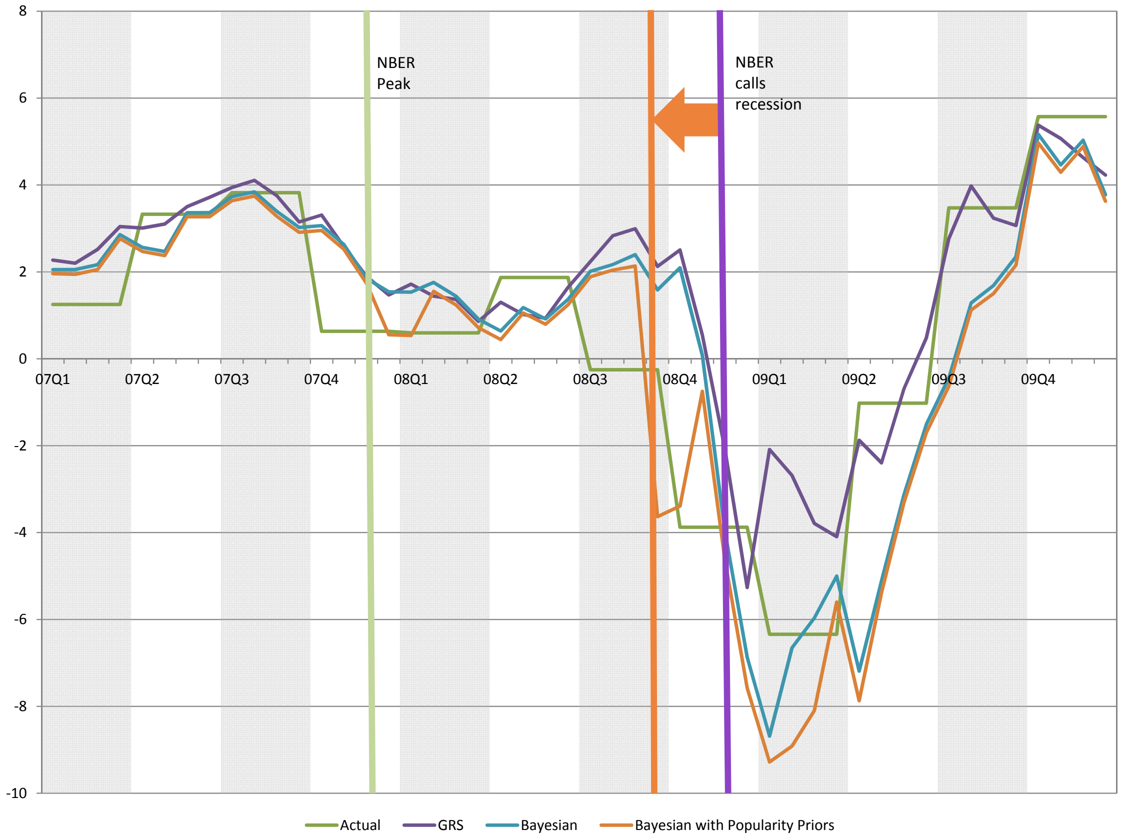

Figure 9: Real-Time Nowcasts with Popularity Priors and Revised Estimates of GDP Growth

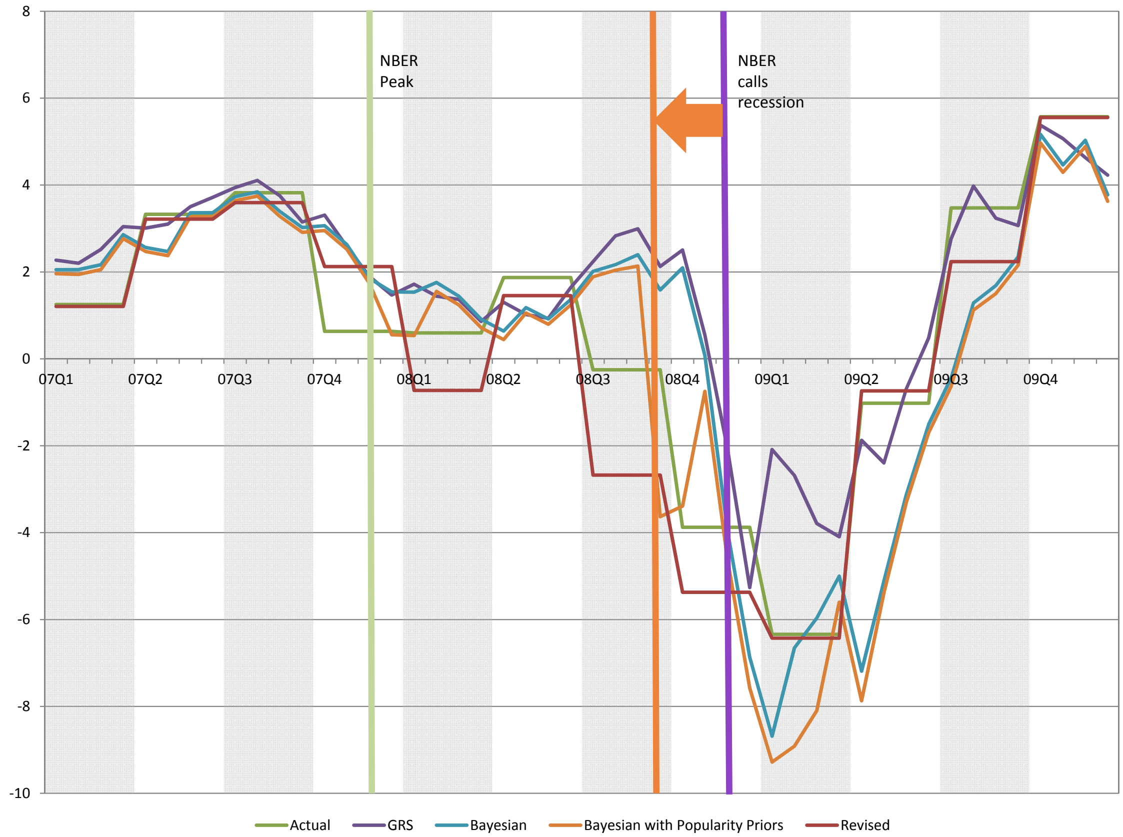

Figure 10: Google Searches for Keyword “End of Recession”

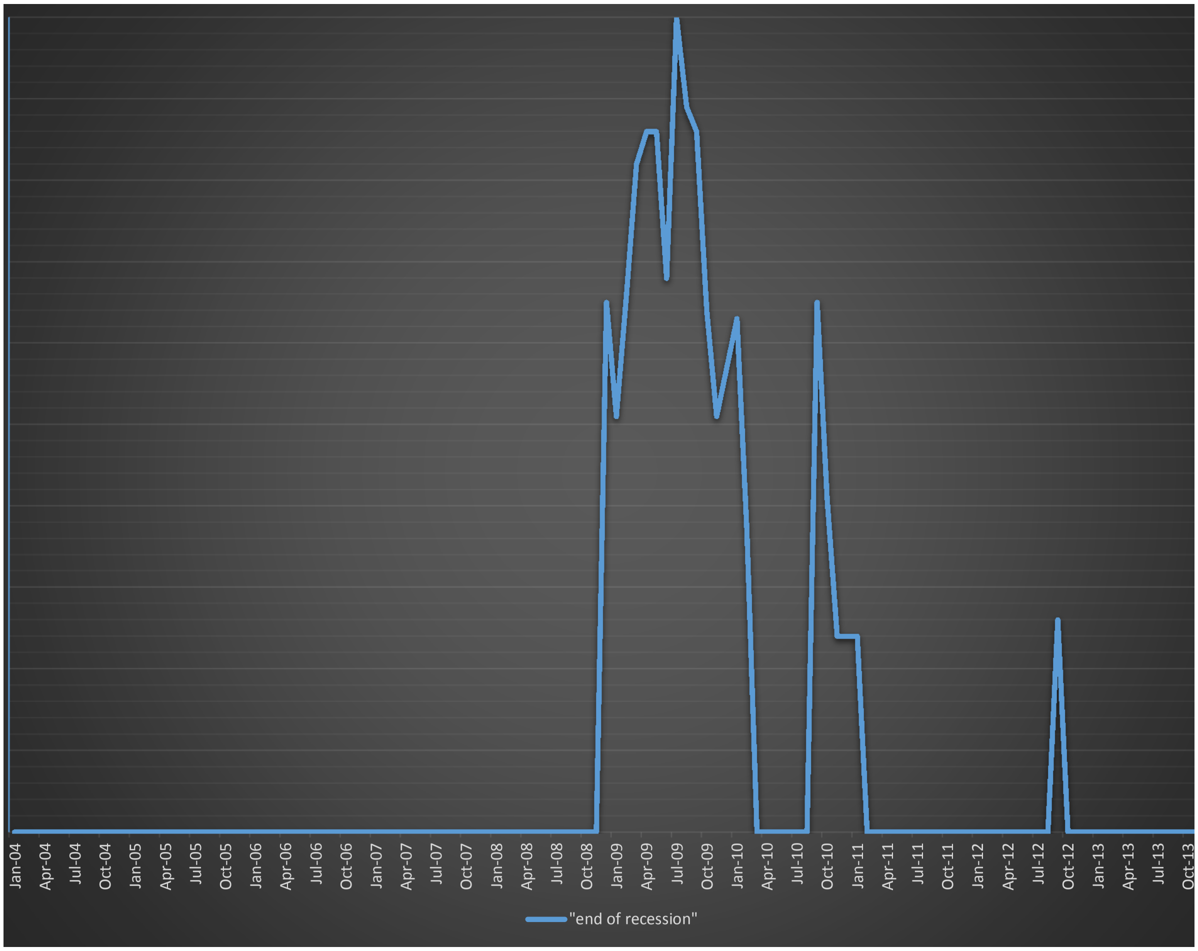

Figure 11: Real-Time Nowcasts with Alternative Popularity Priors