Employment Reconciliation and Nowcasting*

Eiji Goto
University of Missouri-St. Louis
egoto@umsl.edu

Jan P.A.M. Jacobs
University of Groningen, CAMA and CIRANO
j.p.a.m.jacobs@rug.nl

Tara M. Sinclair
The George Washington University and CAMA
tsinc@gwu.edu

Simon van Norden
HEC Montréal, CIRANO and CIREQ
simon.van-norden@hec.ca

December 2021

Abstract

The monthly release of employment data for the U.S. includes two different estimates from two different surveys. One is based on a survey of establishments (payroll) and the other is based on a survey of households. The presence of two different sources of information on broadly the same theoretical concept leads to an obvious question: can we combine the information to obtain an improved estimate of employment?

In this paper we build on the research on combining different measures of output to instead combine different measures of employment. We construct a latent employment estimate which reconciles the information from the two separate surveys as well as incorporating the preliminary data revision process of the payroll data. We find that our reconciled latent employment series looks different than the initial release of payroll employment and is closer to the fully-revised data (benchmarked to a near census of employment), particularly during the Great Recession. Once we move to a real-time exercise, however, our findings suggest that the reconciled employment estimate is remarkably similar to the initial release of payroll employment with near zero weight on the household survey information.

JEL classification: C22, C53, C82, E24
Keywords: employment, United States, real-time data, news, noise

*The authors thank Matthew Zahn (whose undergraduate thesis inspired this project), Elmar Mertens, Todd Clark, James Mitchell and other participants at the FriendlyFaces seminar. In additional we thank participants at the Real-time Economics Conference 2021.

1 Introduction

Employment is a key economic indicator that is closely watched by governments, businesses, journalists, financial analysts, and many others. In the U.S. in particular, the releases by the Bureau of Labor Statistics (BLS) have a following of data watchers and analysts on Twitter who celebrate "jobs day" each month with a race to make the most interesting charts and analyses based on the latest data. The monthly BLS Employment Situation news release includes two different estimates of employment from two different surveys. The payroll employment estimates are based on the monthly survey of businesses and government agencies, whereas the household employment estimates are based on a monthly survey of households (see `https://www.bls.gov/web/empsit/ces_cps_trends.htm` for more details on the two surveys). The data are typically released at 8:30 am ET on the first Friday following the reference month. In the short term, the household estimates are not revised but the payroll estimates are revised to incorporate additional data over the next two months. In the longer term, the payroll estimates are revised each year in a benchmarking process that re-anchors the employment estimates to the full population counts based on unemployment insurance records which cover approximately 97 percent of employment (see `www.bls.gov/opub/mlr/2017/article/benchmarking-the-current-employment-statistics-national-estimates.htm`). The household numbers are revised for seasonal factors for the seasonally adjusted data and for annual population adjustments. Wu (2004) discusses the pros and cons of the two different surveys for tracking employment in real time.

An obvious question arises with two different surveys focused on the same underlying economic variable—is there a way to combine the information from the surveys to obtain an improved estimate of employment in real time?[1] This question is particularly important in the US employment case because the payroll survey is generally the one that is given more attention due to its larger sample size, but the revision patterns in the payroll data may make it not as useful in real time.[2] For example, Neumark and Washer (1991) find

[1] Justin Wolfers has suggested a rule of thumb on Twitter with 80 percent on the payroll survey, 20 percent on the household survey: `https://twitter.com/JustinWolfers/status/431783260524265472?s=20`.

[2] A third series, the ADP employment report (`https://adpemploymentreport.com/`), is an estimate of US employment from a private payroll processing company and is released a couple of days before the BLS

that preliminary payroll employment estimates are not "efficient forecasts" of the revised estimates, Phillips and Nordlund (2012) find evidence of a cyclical bias in the early payroll employment estimates, and Owyang and Vermann (2014) document systematic bias in the revisions of the payroll employment data. Furthermore, although we are focused on shorter term revisions, Haltom et al. (2005) find that previous (longer term) benchmark data revisions help predict future employment benchmark data.

Several studies have tried to determine which measure should be preferred in various contexts, but there is a growing literature instead focusing on reconciliation of multiple measures in order to incorporate more information. Much of the argument for reconciliation has focused on different measures of output, for example, Stone et al. (1942), Weale (1992) and Diebold (2010). Early models of reconciliation relied on the assumption that measurement errors are "noise", which in turn forces the reconciled estimate of the latent variable to be less variable than any of the individual series being reconciled. Aruoba et al. (2013) consider the problem from a forecast combination perspective, assuming "news" errors and imposing priors to address identification, while Aruoba et al. (2016) consider alternative identifying assumptions and propose the addition of an instrumental variable. Almuzara et al. (2021) investigate a dynamic factor model (DFM) with cointegration restrictions while Anesti et al. (2021) propose a mixed-frequency Release-Augmented DFM. See Jacobs et al. (2020) for details. Our contribution to this literature is to extend the analysis to focus on employment instead of output and to incorporate the unique features of the payroll and household surveys used to measure U.S. employment.

In this paper we construct a latent employment estimate which reconciles the information from the two separate surveys as well as incorporating the preliminary data revision process of the payroll data. Our model builds upon a version of Jacobs et al. (2020). We find, similar to the Council of Economic Advisors (2017), that the household survey is in general not very informative, but we also find that our reconciled latent employment series looks different than the initial release of payroll employment, particularly during the Great Recession when it is closer to the benchmarked data. This is only true, however, for the

estimates. We do not include this series in our analysis because the publicly available series is a forecast of employment rather than the microdata from the firm. Researchers from the Federal Reserve have access to the microdata and have combined it with the CES data in Cajner et al. (2019).

full sample. Once we move to a real-time exercise, our findings suggest that the reconciled employment estimate is remarkably similar to the initial release of payroll employment. Furthermore, from our estimates we find that the payroll series is predominantly news, with much of the news in the first estimate, whereas the household series is almost all noise. These results confirm the approach of focusing attention on the preliminary release of payroll data when analyzing the US labor market.

The paper is structured as follows. In Section 2 we present our econometric framework. In Section 3, we describe our data and estimation method. Results are shown in Section 4 and Section 5 concludes. An Appendix provides a detailed explanation of the Bayesian estimation method.

2 Econometric Framework

2.1 Some assumptions and notation

We have two data series: Payroll (PR) employment x_t and Household (HH) employment z_t. We model them in first differences, since that is how employment is typically discussed. Working with first differences also ensures that both series are stationary and mitigates problems associated with benchmark or historical revisions. See e.g. Croushore (2011). We also estimated models in growth rates which produced similar results.

Both series are reported at the same frequency, and are released at the same time. We have data from the Household survey starting with vintages starting in February 1961 (target date January 1961, but taking first difference means the sample starts February 1961). Our latest available data currently is from the February 2021 vintage.

We use seasonally adjusted data[3] and ignore seasonal factor changes to treat z_t as if it is not revised (although we do use real-time data in the estimates). We explicitly model

[3]Employment data are typically seasonally adjusted with the Census X13-ARIMA-SEATS method, the combination of Census X12-ARIMA and TRAMO-Seats which has become the industry standard (Department of Commerce Census Bureau `https://www.census.gov/data/software/x13as.html`). Recently, the method has become under fire. It produces the best trend/cycle and season estimates in 'normal' circumstances. However its performance during crises is strange. For details see Abeln and Jacobs (2021).

the more substantive revisions of x_t for which we use three consecutive monthly releases.[4] Let x_t^j be the jth published estimate of x_t and $\boldsymbol{X}_t \equiv [x_t^1, \ldots, x_t^l]$.

Both the HH and PR series are measures of the same theoretical concept but do so with some degree of error as well as differences in definition. We call the underlying concept "employment" m_t and treat it as a latent variable. The Bureau of Labor Statistics (BLS) also releases (on the same day) a version of the HH survey that is modified to match the payroll definition, but that data is only available back to 1994, so for our analysis we use the unadjusted HH data.

Since there are differences between the HH and PR employment concepts, we also estimate models including a constant (μ_Z) and a slope (β_Z) in the relationship between HH and the latent variable, since we think of the payroll definition as the one to apply to the latent variable.

The measurement error in PR for each release j is $e_t^j \equiv x_t^j - m_t$. Each e_t^j has mean zero and could be the sum of a pure news (ν_t^j) and a pure noise (ϵ_t^j) component. The measurement error in HH is defined similarly, has a mean of zero, and could be the sum of a pure news (ν_t^H) and a pure noise (ϵ_t^H) component. Those components could be correlated (to varying degrees) with ν_t^j or ϵ_t^j.

We evaluate the quality of the latent employment estimate by comparing it to the benchmarked Establishment Survey Estimates: `https://www.bls.gov/web/empsit/cesbmart.htm`. The latest benchmarked data was released on February 5, 2021 (`https://www.bls.gov/news.release/archives/empsit_02052021.htm`). For our evaluation we will treat these latest numbers as the final numbers, although small revisions for seasonal adjustment happen for 5 years and occasionally more for benchmark or historical revisions which change the whole vintage. The latest benchmark level provides data that is considered fully revised through March of 2019. Therefore when we evaluate the latent series we will compare it with the data available from the February 5, 2021 vintage that goes through March of 2019.

[4] Our modeling framework allows for l different releases of x_t. There's nothing that requires these l releases to be consecutive, but they need to be in order. For example, when $l = 3$ we could use releases $(1, 2, 3)$ or $(1, 3, 12)$ but not $(1, 12, 3)$.

2.2 State Space Form

In the general case with l vintages of x_t, we can write the state-space form of our model as

State Vector:

$$\boldsymbol{\alpha}_t = \left[m_t, e_t^H, e_t^1, \ldots, e_t^l\right]' \qquad \text{which is } (2+l) \times 1. \tag{1}$$

Measurement Equation:

$$\begin{bmatrix} z_t \\ x_t^1 \\ \vdots \\ x_t^l \end{bmatrix} = \begin{bmatrix} \mu_z \\ 0 \\ \vdots \\ 0 \end{bmatrix} + \begin{bmatrix} \beta_z & 1 & 0 & \cdots & 0 \\ 1 & 0 & \ddots & \ddots & \vdots \\ \vdots & \vdots & \ddots & 1 & 0 \\ 1 & 0 & \cdots & 0 & 1 \end{bmatrix} \cdot \boldsymbol{\alpha}_t. \tag{2}$$

In the simplest model we impose $\mu_Z = 0$ and $\beta_Z = 1$.

State Equation:

$$\boldsymbol{\alpha}_t = \boldsymbol{T} \cdot \boldsymbol{\alpha}_{t-1} + \boldsymbol{R} \cdot \boldsymbol{\eta}_t, \tag{3}$$

where

$$\boldsymbol{T} = \begin{bmatrix} \rho & 0 & \cdots & 0 \\ 0 & 0 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & 0 \end{bmatrix}, \qquad \text{which is } (2+l) \times (2+l);$$

$$\boldsymbol{R} = \begin{bmatrix} \sigma_\nu^H & \sigma_\nu^1 + \sigma_\nu^{H1} & \ldots & \ldots & \sigma_\nu^l + \sigma_\nu^{Hl} & 0 & 0 & \ldots & \ldots & 0 \\ \sigma_\nu^H & \sigma_\nu^{H1} & \ldots & \ldots & \sigma_\nu^{Hl} & \sigma_\epsilon^H & \sigma_\epsilon^{H1} & \ldots & \ldots & \sigma_\epsilon^{Hl} \\ 0 & -\sigma_\nu^1 & -\sigma_\nu^2 & \ldots & -\sigma_\nu^l & 0 & \sigma_\epsilon^1 & 0 & \ldots & 0 \\ 0 & 0 & -\sigma_\nu^2 & \ldots & -\sigma_\nu^l & 0 & 0 & \sigma_\epsilon^2 & \ddots & \vdots \\ \vdots & \vdots & \ddots & \ddots & \vdots & \vdots & \vdots & \ddots & \ddots & 0 \\ 0 & 0 & \ldots & 0 & -\sigma_\nu^l & 0 & 0 & \cdots & 0 & \sigma_\epsilon^l \end{bmatrix},$$

which is $(2+l) \times 2(1+l)$;

$$\boldsymbol{\eta}_t = \left[\nu_t^H, \nu_t^1, \ldots, \nu_t^l, \epsilon_t^H, \epsilon_t^1, \ldots, \epsilon_t^l\right]' \qquad \boldsymbol{\eta}_t \sim \text{i.i.d. } N(0, \boldsymbol{I}_{2(1+l)}).$$

We can also rewrite matrix $\boldsymbol{R}$ as

$$\boldsymbol{R} = \begin{bmatrix} \boldsymbol{\sigma}_\nu^H & \boldsymbol{R}_1 + \boldsymbol{R}_2 & 0 & \mathbf{0}_{1\times l} \\ \boldsymbol{\sigma}_\nu^H & \boldsymbol{R}_2 & \sigma_\epsilon^H & \boldsymbol{R}_4 \\ \mathbf{0}_{l\times 1} & -\boldsymbol{V}_l \cdot \text{diag}(\boldsymbol{R}_1) & \mathbf{0}_{l\times 1} & \text{diag}(\boldsymbol{R}_3) \end{bmatrix}, \tag{4}$$

where

$\boldsymbol{V}_l$ is a $l \times l$ matrix with zeroes below the main diagonal and ones everywhere else;

$\boldsymbol{R}_1 = \left[\sigma_\nu^1, \ldots, \sigma_\nu^l\right]$;

$\boldsymbol{R}_2 = \left[\sigma_\nu^{H1}, \ldots, \sigma_\nu^{Hl}\right]$;

$\boldsymbol{R}_3 = \left[\sigma_\epsilon^1, \ldots, \sigma_\epsilon^l\right]$;

$\boldsymbol{R}_4 = \left[\sigma_\epsilon^{H1}, \ldots, \sigma_\epsilon^{Hl}\right]$.

All the $\boldsymbol{R}_j$ are $1 \times l$ row vectors; diag($\boldsymbol{v}$) is a diagonal matrix with the vector $\boldsymbol{v}$ on the main diagonal.

Parameters:

The $5 + 4l$ parameters to estimate are

$$\boldsymbol{\theta} = [\mu_z, \beta_Z, \rho, \sigma_\nu^H, \sigma_\nu^1, \ldots, \sigma_\nu^l, \sigma_\nu^{H1}, \ldots, \sigma_\nu^{Hl}, \sigma_\epsilon^H, \sigma_\epsilon^1, \ldots, \sigma_\epsilon^l, \sigma_\epsilon^{H1}, \ldots, \sigma_\epsilon^{Hl}]', \tag{5}$$

where ρ controls the AR(1) serial correlation in m_t; σ's capture the standard deviations of the different types of measurement errors, σ^j captures the measurement error in x_t^j, σ^H captures the measurement error in z_t that is orthogonal to the measurement error in all releases x_t^j, σ^{Hj} captures the measurement error in z_t that is perfectly correlated with the measurement error in release x_t^j, σ_ν captures news measurement errors, and σ_ϵ

captures noise measurement errors.[5] μ_z and β_Z capture the (assumed linear) relationship between the household employment series and the latent employment series since there are definitional differences between household and payroll employment and the latent series is assumed to follow the definitions of the payroll series.

3 Data and Estimation

3.1 Data

Since our objective is to reconcile employment estimates in real time we take the revisions process of our data seriously. All our data are from the Employment Situation releases from the Bureau of Labor Statistics and we source the real-time data from ALFRED, ArchivaL Federal Reserve Economic Data, provided by the Federal Reserve Bank of St. Louis. We focus on two seasonally adjusted data series: (1) Payroll employment (PR): total nonfarm payroll employment from the Current Employment Statistics (Establishment Survey) and (2) Household employment (HH): civilian employment from the Current Population Survey (Household Survey). In order to capture the appropriate real-time relationships, we focus on the initial release of the (change in) employment from both the household and the payroll surveys. Since the payroll survey follows a regular pattern of revisions in the following two months, we also include the second and third releases of payroll survey data. Since we are modelling the relationship in first differences, the level differences in definition may not matter, but we allow for differences in both intercept and slope in our main models below.

Our baseline analysis focuses on data from 1961 through 2019. In Section 4.4, we extend the analysis to include data for 2020 to explore the robustness of our results to unprecedented effects of the recent pandemic on employment growth.

Figure 1 presents the initial releases of the change in employment from the two surveys. From this figure we can see the greater volatility of the household numbers (due to the smaller sample size). This greater volatility can also be seen in the descriptive statistics reported in Table 1. The standard deviation for the household data is substantially larger

[5] We allow for positive covariances between measurement errors in z_t and x_t^j by making the measurement errors in z_t equal the sum of its own component and a multiple of the component in x_t^j.

than that for the payroll data, however, the means and medians of the first differences are broadly similar across the initial releases of the two series as well as the revised payroll data.

[Figure 1 about here.]

[Table 1 about here.]

To evaluate the properties of our reconciled latent employment estimates we compare them to the benchmarked data that are available through March of 2019. Table 2 shows that there is little impact of excluding the remainder of 2019 on the data series that enter the model. For the benchmarked data, the mean and median are both higher than initial release data. Notably, the benchmarked PR data is statistically significantly higher on average than the initial PR data by about 20k. This difference becomes less for the second revision (but still statistically significant) and by the third revision the difference is less than 3k and statistically insignificant.

[Table 2 about here.]

One of the key time periods in our sample is the Great Recession. According to the NBER, the U.S. economy peaked in December of 2007 and reached a trough in June of 2009. The economy was in recession for 18 months. Table 3 provides the descriptive statistics for this period. Here the mean and median of the benchmarked data are lower than the other series, and the difference is statistically significant for the three earlier payroll releases. Notably, the difference is not statistically different for the household series, even though the difference is larger than for the second revision of the payroll series. It is important to remember that the household series has a larger variance which affects the statistical significance, particularly in this small sample. In real time, however, it may be the case that the initial household data is more informative in recessions than the inital payroll data, as discussed by Chauvet and Piger (2013).

[Table 3 about here.]

3.2 Estimation

We estimate the model with $l = 3$ vintages for PR x_t^j and data and one release of HH z_t.

State Vector:

$$\boldsymbol{\alpha}_t = \left[m_t, e_t^H, e_t^1, e_t^2, e_t^3\right]', \tag{6}$$

where the measurement errors e_t^H (for the household survey) and e_t^j, $j = 1, 2, 3$ (for the three payroll survey releases) are the sum of news and noise errors; $e_t^i = \nu_t^i + \varepsilon_t^i$ where $i = H, 1, 2, 3$.

Measurement Equation:

$$\begin{bmatrix} z_t \\ x_t^1 \\ x_t^2 \\ x_t^3 \end{bmatrix} = \begin{bmatrix} \mu_z \\ 0 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} \beta_z & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 1 \end{bmatrix} \cdot \boldsymbol{\alpha}_t. \tag{7}$$

State Equation:

$$\boldsymbol{\alpha}_t = \boldsymbol{T} \cdot \boldsymbol{\alpha}_{t-1} + \boldsymbol{R} \cdot \boldsymbol{\eta}_t, \tag{8}$$

where

$$\boldsymbol{T} = \begin{bmatrix} \rho & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix},$$

$$\boldsymbol{R} = \begin{bmatrix} \sigma_\nu^H & \sigma_\nu^1 + \sigma_\nu^{H1} & \sigma_\nu^2 + \sigma_\nu^{H2} & \sigma_\nu^3 + \sigma_\nu^{H3} & 0 & 0 & 0 & 0 \\ \sigma_\nu^H & \sigma_\nu^{H1} & \sigma_\nu^{H2} & \sigma_\nu^{H3} & \sigma_\epsilon^H & \sigma_\epsilon^{H1} & \sigma_\epsilon^{H2} & \sigma_\epsilon^{H3} \\ 0 & -\sigma_\nu^1 & -\sigma_\nu^2 & -\sigma_\nu^3 & 0 & \sigma_\epsilon^1 & 0 & 0 \\ 0 & 0 & -\sigma_\nu^2 & -\sigma_\nu^3 & 0 & 0 & \sigma_\epsilon^2 & 0 \\ 0 & 0 & 0 & -\sigma_{\boldsymbol{\nu}}^3 & 0 & 0 & 0 & \sigma_\epsilon^3 \end{bmatrix},$$

$$\eta_t = \left[\nu_t^H, \nu_t^1, \nu_t^2, \nu_t^3, \epsilon_t^H, \epsilon_t^1, \epsilon_t^2, \epsilon_t^3\right]'.$$

3.3 Bayesian estimation

For our preferred specification, we estimate the parameters using Bayesian methods. We generate draws from the posterior distributions using a random walk Metropolis-Hastings algorithm.[6] Table 4 summarizes the specification of the prior distribution. We make 50,000 draws in total but throw away the first 5,000 draws as a burn-in period. After the burn-in period, we draw the median, 25th percentile, and 75th percentile of the draws from the posterior distribution. A detailed explanation of the Bayesian estimation is in Appendix A. Given the median parameter draw, we estimate the state variables using the Kalman filter.

[Table 4 about here.]

4 Results

4.1 Results for the 1961–2019 sample

Table 5 provides the parameter estimates for our different models all for first-differenced data for the sample of February 1961 through December of 2019. Our preferred model that we will use for reporting further estimates is the last column of the table, the Bayesian MCMC model including both a constant and a slope for the connection between the household data and the latent series, our most general and flexible model. The results, however, are remarkably similar across all the models. From these estimates we observe that the payroll series is predominantly news, whereas the household series is almost all noise. Furthermore for the payroll series, most of the news is in the first estimate, but the later estimates provide some further news. This conclusion is based on the σ_ν parameters capturing the news and the σ_ϵ parameters capturing the noise. The news parameters show large estimates for the payroll releases, particularly the first release, but small estimates

[6] Chib (1995) contains a comprehensive introduction of the Metropolis-Hastings algorithm. This method has been intensively used to estimate state space models (see e.g. Smets and Wouters (2007) and Aruoba et al. (2016)).

for the household series. The noise parameter estimates show the opposite pattern, with a large value for the household series but small values for the payroll series. The model includes covariance parameters, but their estimates are small so the interpretation is not much affected.[7]

[Table 5 about here.]

It is important to note that the latent series reported here is estimated over the full sample. We can see from the Kalman gains reported in Table 6 that the model puts almost all weight on the second revision of the payroll data. This may seem unsurprising since the objective of the revisions is to improve the estimates, but it is notable that our reconciliation model picks this up so clearly without being provided any final target.

[Table 6 about here.]

Figure 2 presents our smoothed estimates of the latent employment change series based on our preferred model along with the initial payroll estimate.[8] These two series appear similar, but there are some notable differences as shown in Figure 3. Two time periods are important to examine to understand the results of our estimates and both are visible clearly in Figure 3: the Great Recession of 2008–2009 and the period in the early 1980s where we see the most volatility in the differences between the latent series and the initial estimate.

[Figure 2 about here.]

[Figure 3 about here.]

[7] Aruoba et al. (2016) argue that measurement errors from household surveys should be orthogonal to those from business surveys and use this to justify their choice of instrumental variable. The parameter estimates shown above in Table 5 imply correlations between each of the noise components in PR and that in HH are all much less than 1%. We also re-estimated our models assuming zero correlations between the noise shocks from the two different surveys. These restricted models produced Kalman gains and figures that were strikingly similar to those of the models in Table 5.

[8] A smoothed estimate for period t is the expected value of the latent series m_t conditional on the parameter estimates $\hat{\theta}$ and all our data series from $t = 1$ to T.

4.2 Nowcasting results

The previous results showed smoothed estimates for the full sample. A key question is whether or not we can improve upon initial releases of payroll employment in real time. In order to explore this question we produce filtered estimates of the series for each month, and take into account the missing observations for the revised estimates. Thus for the nowcast (technically a back-cast since it's the estimate for the month prior, but we'll consider it a nowcast here) we are missing the two revisions for the current (change in) payroll as well as the second revision for the previous month's payroll number. Note that we are still using the estimates from the full sample for the parameters, though. We find two key results from these estimates. First, the nowcast series looks very similar to the initial payroll release, as can be seen from Figure 4. Second, there remains little weight on the household series, even in real time when we don't have the later revisions of the payroll series, as can be seen from the Kalman gains reported in Table 7.

[Figure 4 about here.]

[Table 7 about here.]

4.3 Comparing the latent employment series with benchmarked employment

An interesting feature of US payroll employment is that the data are eventually benchmarked to a near census of employment. As discussed in the Data section above, the latest benchmarked employment data as of this analysis is through March of 2019. In order to evaluate the performance of our latent series, we compare our series with the benchmarked data.

Table 8 presents the mean absolute errors (MAE) of using the latent series for the benchmarked series based on our full sample and nowcast models. It also shows the results of using initial payroll release data as our latent series. This analysis is consistent with our previous findings that with the full sample we can find some improvements from our model, but for nowcasts we might just as well use just the initial payroll release.

[Table 8 about here.]

We also considered the performance of the different models in different states of the economy.[9] The NBER only announces the peak and trough dates of the business cycle with a lag so this information is not available in real time, but an ex post analysis of how the model performs in the different periods is insightful. In general the models perform roughly 25% worse (in terms of MAE for the fully revised data) in recessions than in expansions.

4.4 Extending the analysis through 2020

Based on the results during the Great Recession, it may be the case that, at least in sample, our latent series performs differently than the initial release of payroll employment in recessions. Since we experienced a dramatic and unusual recession in 2020, we re-estimated the model using data through 2020 to compare the parameter estimates and the estimates of the latent series.

[Figure 5 about here.]

To put these results into context, note that they are based on adding 13 new observations to arrive at a sample size of 720 months. However, owing to the start of the Covid-19 pandemic in early 2020, some of those observations are extreme, as Figure 5 illustrates. (The upper and lower panels of the Figure show the same series, however the vertical scaling of the lower panel has been adjusted so that all the values fit on the chart.) The values for April 2020 are notable, as those shown for the Payroll series are roughly 100 standard deviations below the series mean.[10]

[9] Results available from the authors upon request.

[10] The sample mean of the 3rd release of the PR series up to the end of 2019 is 136,200 while its sample standard deviation is 201,800. The value for April 2020 is -20,787,000, which is 103.7 standard deviations below the mean. To put this into perspective, Dowd et al. (2008) calculate that (if we assume normality) the probability of observing an event at least 8 sigma below the mean is roughly equal to 1 divided by the number of days since the Big Bang, while observing an event 20 sigma below the mean is roughly equal to 1 divided by the number of particles in the universe. The probability of a 38-sigma event is roughly 10^{-316} and beyond this point the limitations of IEEE double-precision values (which cannot represent anything closer to zero than 4.94066×10^{-324}) complicate calculation of the probability. Using arbitrary-precision computations provides a estimated probability for events at least 103.7 standard deviations below the mean of roughly 6×10^{-2337}. Note that recently Serena Ng (2021) independently makes a similar observation.

[Table 9 about here.]

Adding the 2020 data changes the estimated parameters as shown in Table 9. Despite only adding 13 more observations, we find that the serial correlation in m_t drops dramatically compared to the earlier estimates. We also see dramatic changes in terms of the role of news and noise. The news content of the initial PR release (σ_ν^1) falls by more than an order of magnitude while the news content of its other two releases show little change. At the same time, the independent news content (σ_ν^H) of the HH release increases by a bit under three orders of magnitude, similar to the increase in its correlated news content (σ_ν^{H1} through σ_ν^{H3}). While its independent noise content (σ_ϵ^H) falls by almost two-thirds, its correlated noise components (σ_ϵ^{H1} through σ_ϵ^{H3}) become larger still. More tellingly, its loading on m_t (β_z) falls from 0.580 to 0.092, somewhat offsetting the improvement in the noise-to-signal ratio of the HH series.

[Figure 6 about here.]

[Figure 7 about here.]

Despite the dramatic changes in parameter estimates, the estimates of reconciled employment are largely robust to these changes. Figure 6 compares nowcasts of reconciled employment based on the two sets of parameters, while Figure 7 compares the smoothed estimates using the full sample. Both show the same series in their upper and lower panels, rescaling them in the lower panel so that all values fit on the chart. In both cases, the parameter changes have little effect on the estimated series and the correlation between those from the pre-2020 and the updated models are over 99 per cent. This reflects the fact that, despite changes in the noise-to-signal ratio of the HH series, the Kalman Gain continues to give it near zero weight as the overwhelming majority of its news content simply duplicates the news content of the PR series.[11]

Figures 8 and 9 show the comparison between the initial payroll estimates and the latent series produced with the full sample. From Figure 8 we can see that the results up

[11] For the nowcast, the new parameters reduce the weight on innovations in HH from 2.8 to 0.8 per cent while that on PR rises from 98.4 to 100. For the smoothed full-sample estimates, however, the weight on the 3rd release of PR falls from 99.9 per cent to 84.6, with the weight on its 2nd release rising to 16.8. However, given the small differences between these releases, the changes in the estimates are trivial.

through 2019 are not much different when including the 2020 data, but the 2020 estimates show the largest difference between the preliminary data and the latent series of anywhere in the sample. Zooming in on the 2020 results in Figure 9 we see that the big difference is in March of 2020.

[Figure 8 about here.]

[Figure 9 about here.]

5 Conclusion

Unlike earlier work using payroll and household survey data to forecast employment growth, we've introduced a structural model of news and noise measurement errors to explore how these two data sources may best be reconciled to produce combined estimates of employment growth. The state-space framework used shows how those estimates should be formulated when initial estimates are released, and how they should be updated as revised payroll estimates become available.

Our results are strongly consistent with conclusions from work on employment forecasting and strikingly different from those on reconciling expenditure and income estimate of real GDP. We find that data from the household survey receive only vanishingly small weight in reconciliation, regardless of whether or not a full set of revised payroll estimates are available. Our reconciled estimate closely tracks the latest available estimate of the payroll series, which has notable deviations from the initial estimates at times, particularly in the recessions of the early 1980s and 2008-09. However, it closely tracks the official reconciliation of the household and establishment survey results. The near-exclusive weight on the best available payroll estimates is also robust to extending our analysis to include the extreme shocks of the early Covid-19 pandemic.

References

Abeln, Barend and Jan P.A.M. Jacobs (2012), “Covid19 and seasonal adjustment”, CAMA Working Paper Series 23/2021, Centre for Applied Macroeconomic Analysis, Australian National University.

Almuzara, Tincho, Gabriele Fiorentini, and Enrique Sentana (2021), “U.S. aggregate output measurement: A common trend approach”, Staff Report 962, FRB of New York.

Anesti, Nikoleta, Ana Galvão, and Silvia Miranda-Agrippino (2021), “Uncertain Kingdom: Nowcasting GDP and its revisions”, *Journal of Applied Econometrics* [forthcoming].

Aruoba, S. Borağan, Francis X. Diebold, Jeremy Nalewaik, Frank Schorfheide, and Dongho Song (2013), “Improving US GDP measurement: A forecast combination perspective”, in X. Chen and N. Swanson, editors, *Recent Advances and Future Directions in Causality, Prediction, and Specification Analysis*, Springer, New York, 1–26.

Aruoba, S. Borağan, Francis X. Diebold, Jeremy Nalewaik, Frank Schorfheide, and Dongho Song (2016), “Improving GDP measurement: A measurement-error perspective”, *Journal of Econometrics*, **191**(2), 384–397.

Cajner, Tomaz, Leland D Crane, Ryan A Decker, Adrian Hamins-Puertolas, and Christopher Kurz (2019), “Improving the accuracy of economic measurement with multiple data sources: The case of payroll employment data”, Working Paper 26033, National Bureau of Economic Research.

Chauvet, Marcelle and Jeremy Piger (2013), “Employment and the Business Cycle”, *The Manchester School*, **81**(S2), 16–42.

Chib, Siddhartha and Edward Greenberg (1995), “Understanding the Metropolis-Hastings algorithm”, *The American Statistician*, **49**(4), 327–335.

Council of Economic Advisers (CEA) (2017), “Assessing the state of the economy in real time using headline economic indicators”, Issue Brief.

Croushore, Dean (2011), “Frontiers of real-time analysis”, *Journal of Economic Literature*, **49**, 72–100.

Diebold, Francis X. (2010), "Comment", *Brookings Papers on Economic Activity*, **Spring**, 107–112.

Dowd, Kevin, John Cotter, Christopher Humphrey, and Margaret Woods (2008), "How unlucky is 25-sigma?", *Journal of Portfolio Management*, **34**, 76–80.

Haltom, Nicholas L., Vanessa D. Mitchell, Ellis W. Tallman, et al. (2005), "Payroll employment data: Measuring the effects of annual benchmark revisions", *Economic Review-Federal Reserve Bank of Atlanta*, **90**(2), 1.

Jacobs, Jan P.A.M., Samad Sarferaz, Jan-Egbert Sturm, and Simon van Norden (2020), "Can GDP measurement be further improved? Data revision and reconciliation", *Journal of Business Economics & Statistics* [forthcoming].

Neumark, David and William L Wascher (1991), "Can we improve upon preliminary estimates of payroll employment growth?", *Journal of Business & Economic Statistics*, **9**(2), 197–205.

Ng, Serena (2021), "Modeling macroeconomic variations after COVID-19", arXiv:2130:02732.

Owyang, Michael and E. Katarina Vermann (2014), "Employment revision asymmetries", *Economic Synopses*, **11**(4), 1–2.

Phillips, Keith R. and James Nordlund (2012), "The efficiency of the benchmark revisions to the current employment statistics (CES) data", *Economics Letters*, **115**(3), 431–434.

Smets, Frank and Rafael Wouters (2007), "Shocks and frictions in US business cycles: A Bayesian DSGE approach", *American Economic Review*, **97**(3), 586–606.

Stone, Richard, D. G. Champernowne, and J. E. Meade (1942), "The precision of national income estimates", *Review of Economic Studies*, **9**, 111–125.

Weale, Martin (1992), "Estimation of data measured with error and subject to linear restrictions", *Journal of Applied Econometrics*, **7**, 167–174.

Wu, Tao (2004), "Two measures of employment: How different are they?", Technical report, Federal Reserve Bank of San Francisco.

Appendix

A Detailed Explanation of the Bayesian Estimation

Given the assumption that the error terms of the state space model follow the Gaussian distribution, the density of the data $f(\boldsymbol{Y}_t|\boldsymbol{\alpha}_t)$, where $\boldsymbol{Y}_t = \left[z_t x_t^1 \dots x_t^l\right]'$, is:

$$f(\boldsymbol{Y}_t|\boldsymbol{\alpha}_t) = (2\pi)^{-\frac{1}{2}}|\boldsymbol{f}_{t,t-1}|^{-\frac{1}{2}} exp(-0.5\boldsymbol{u}_{t,t-1}\boldsymbol{f}_{t,t-1}^{-1}\boldsymbol{u}_{t,t-1}),$$

where $\boldsymbol{u}_{t,t-1}$ is the predictive error and $\boldsymbol{f}_{t,t-1}$ is the variance of the predictive error of the Kalman filter. The likelihood function of the model is:

$$f(\boldsymbol{Y}|\boldsymbol{\theta}) = \prod_{t=1}^{T} f(\boldsymbol{Y}_t|\boldsymbol{\alpha}_t),$$

where $\boldsymbol{\theta} = \{\rho, \sigma_\nu^H, \sigma_\nu^1, \sigma_\nu^2, \sigma_\nu^3, \sigma_\nu^{H1}, \sigma_\nu^{H2}, \sigma_\nu^{H3}, \sigma_\epsilon^H, \sigma_\epsilon^1, \sigma_\epsilon^2, \sigma_\epsilon^3, \sigma_\epsilon^{H1}, \sigma_\epsilon^{H2}, \sigma_\epsilon^{H3}, \mu_z, \beta_z\}$.

We conduct the random walk Metropolis-Hastings approach in the following steps:

Step 1: Specify a starting value $\boldsymbol{\theta}_0$ and variance of the shock $\boldsymbol{\Sigma}$.

Step 2: Draw a new parameter vector from the random walk equation:

$$\boldsymbol{\theta}_{NEW} = \boldsymbol{\theta}_{OLD} + \boldsymbol{e} \qquad \boldsymbol{e} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{\Sigma}).$$

Step 3: Compute the acceptance probability:

$$\alpha = \min\left(\frac{f(\boldsymbol{Y}|\boldsymbol{\theta}_{NEW})p(\boldsymbol{\theta}_{NEW})}{f(\boldsymbol{Y}|\boldsymbol{\theta}_{OLD})p(\boldsymbol{\theta}_{OLD})}, 1\right),$$

where $p(\boldsymbol{\theta}_i)$ is the prior density.

Step 4: If $\alpha > a \sim \mathcal{U}(0, 1)$, obtain $\boldsymbol{\theta}_{NEW}$. Otherwise $\boldsymbol{\theta}_{NEW} = \boldsymbol{\theta}_{OLD}$.

Repeat steps 2, 3, and 4 50,000 times. The first 5,000 draws are a burn-in period. After the burn-in period, we draw the median, 25th percentile, and 75th percentile of the draws from the posterior distribution.

List of Figures

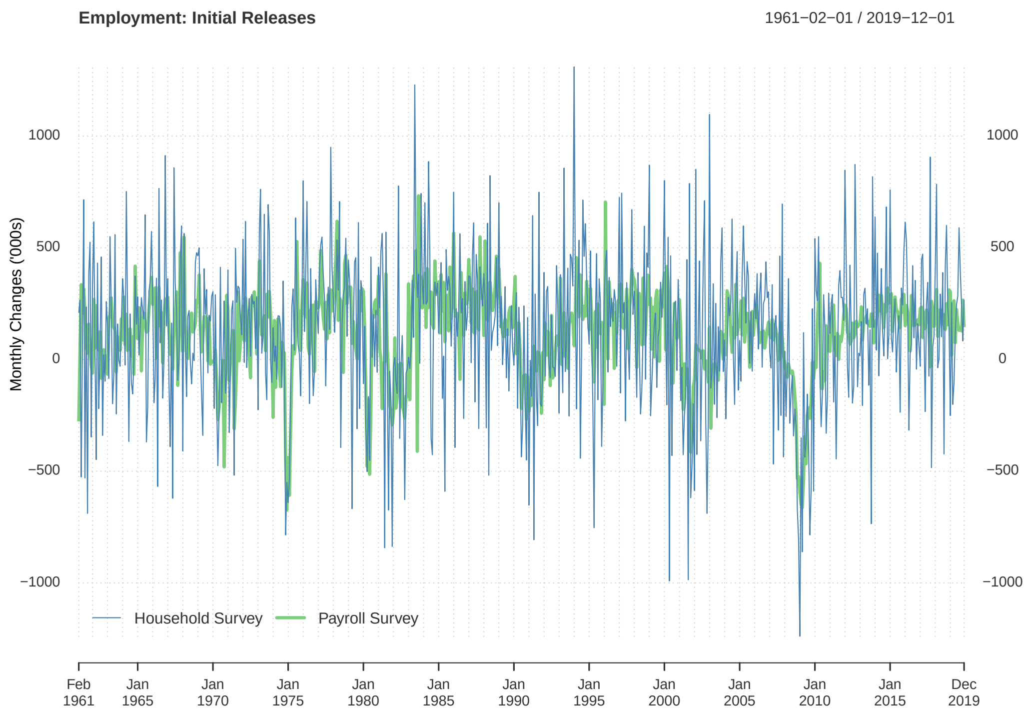

Figure 1: U.S. Employment: Initial Release of Monthly Changes, Thousands of Persons, Feb 1961–Dec 2019

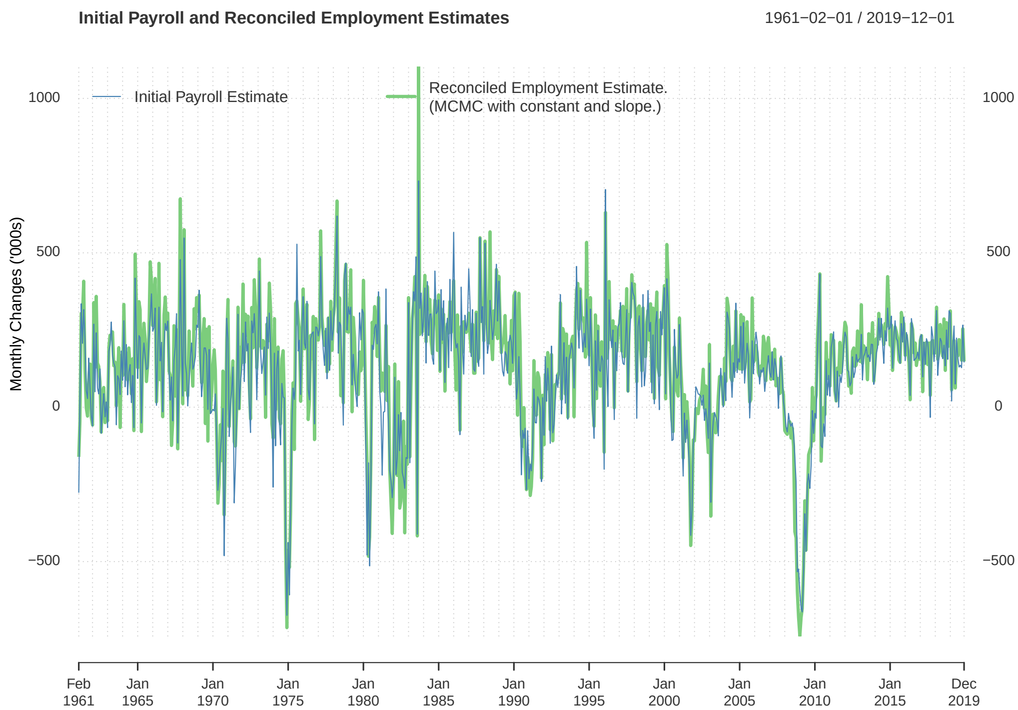

Figure 2: Comparing Initial Release of Payroll Employment Change with Latent Estimate based on MCMC estimate including constant and slope, Thousands of Persons, Feb 1961–Dec 2019

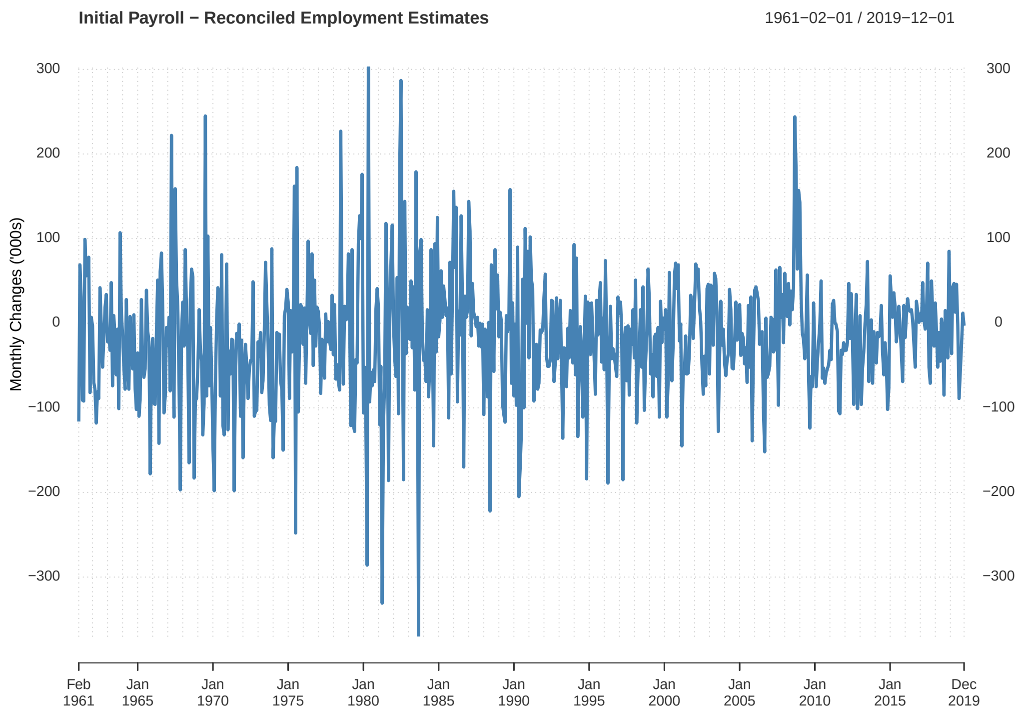

Figure 3: Difference Between Initial Release and Latent Estimate of Employment Change, Thousands of Persons, Feb 1961–Dec 2019

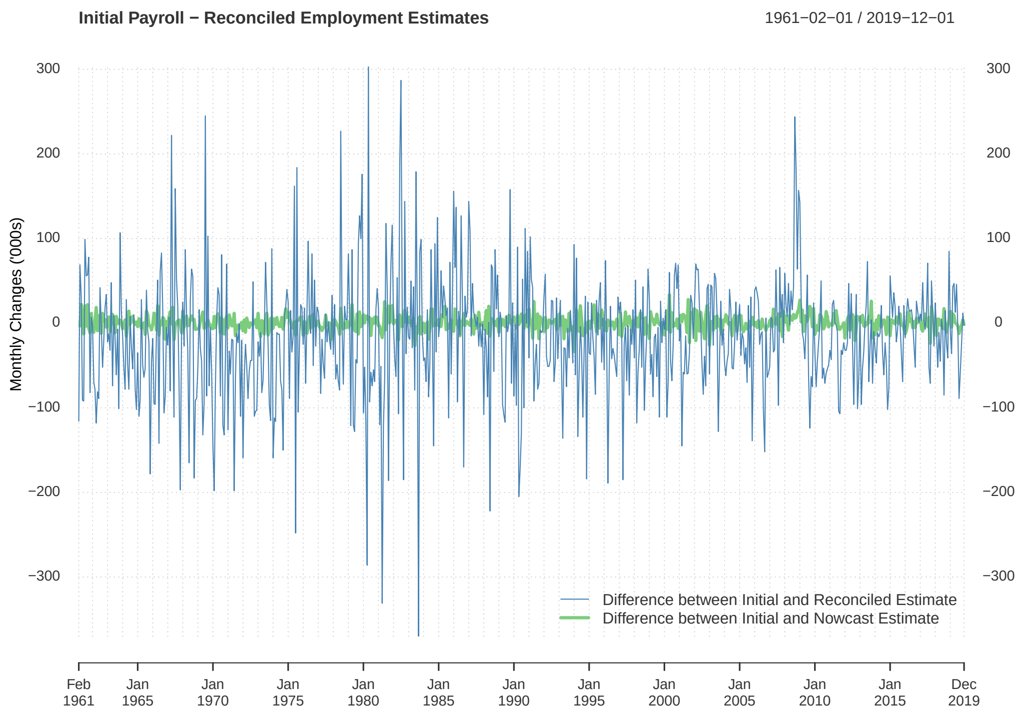

Figure 4: Comparing Difference Between Latent and Initial Release and Nowcast and Initial Release of Employment Change, Thousands of Persons, Feb 1961–Dec 2019

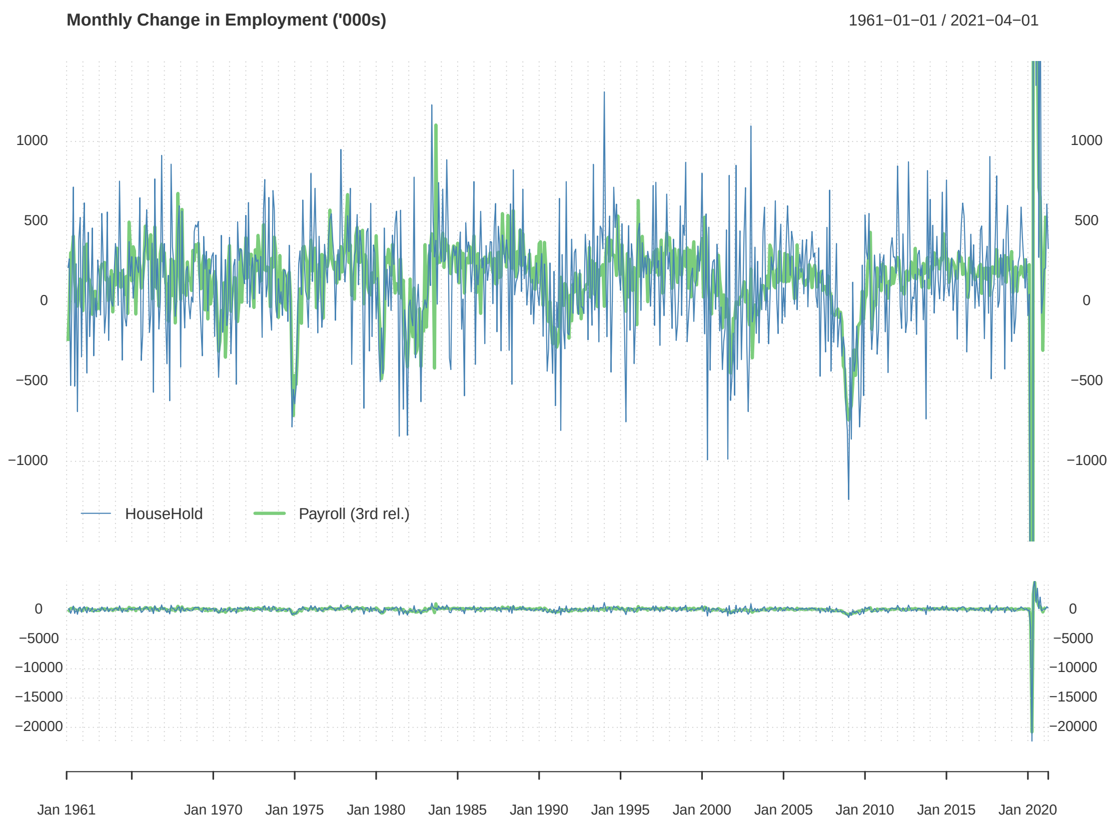

Figure 5: The Covid-19 Pandemic in Perspective

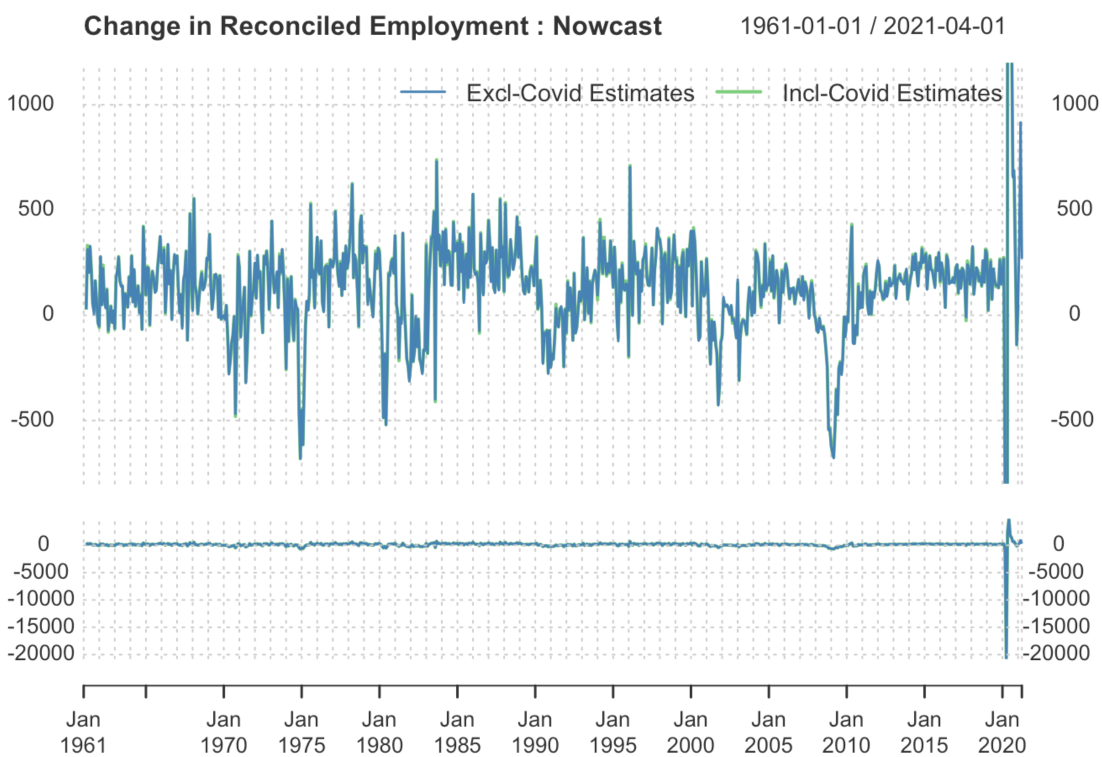

Figure 6: Two Nowcasts of Reconciled Employment Growth

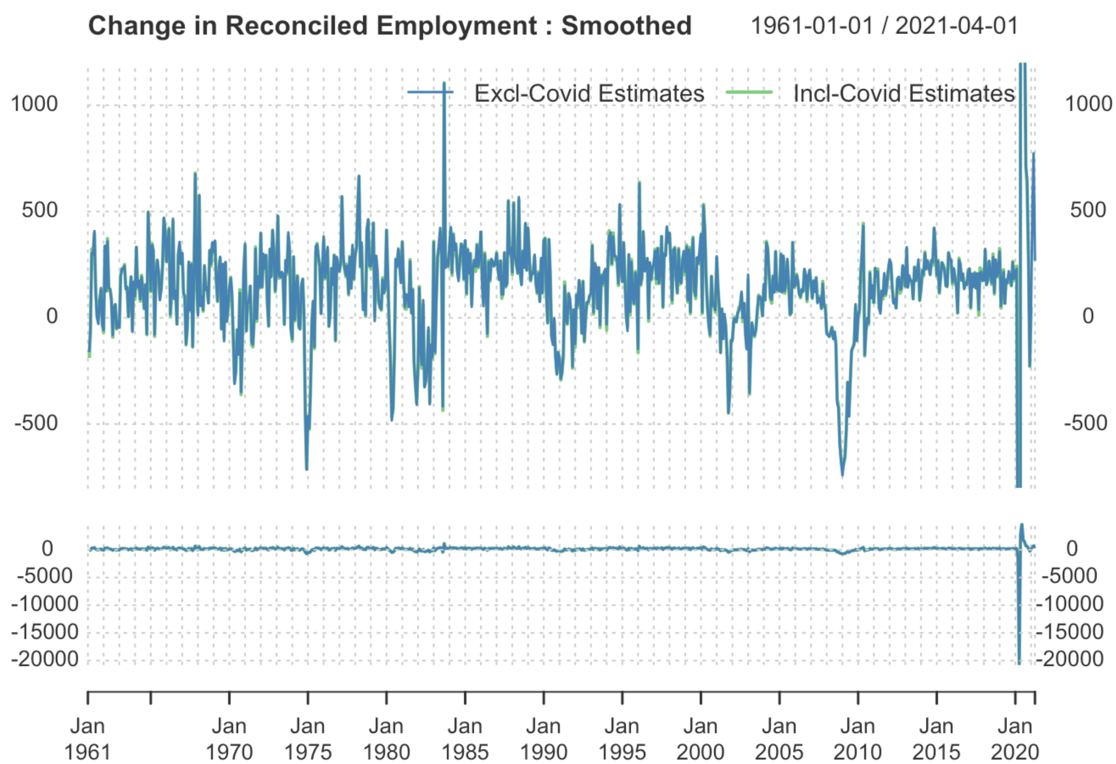

Figure 7: Two Smoothed Estimates of Reconciled Employment Growth

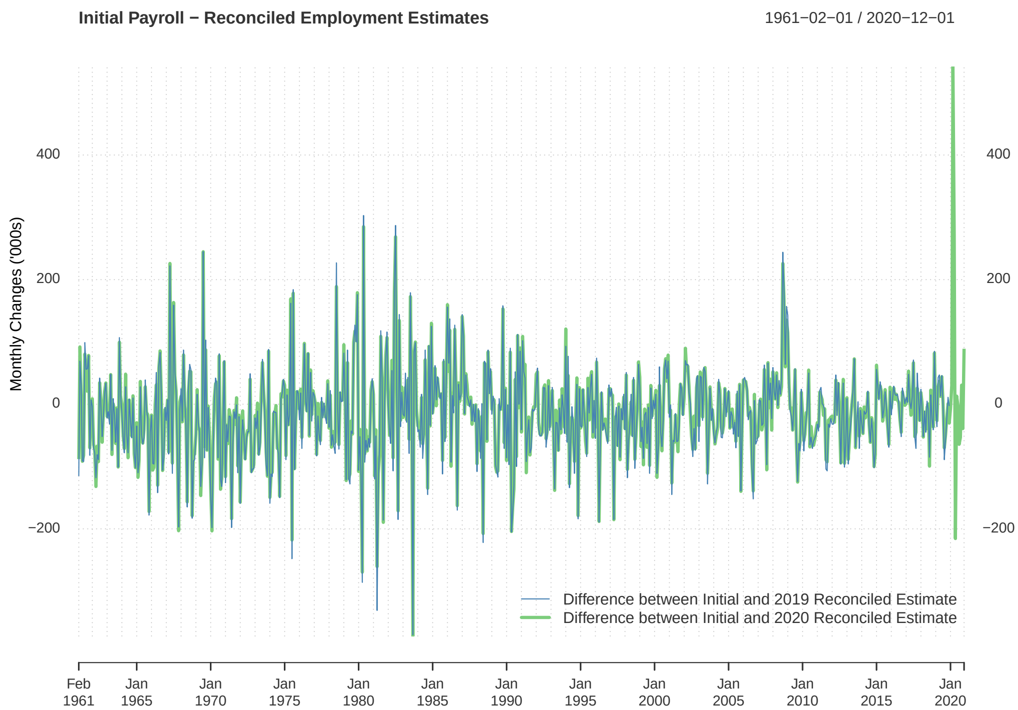

Figure 8: Comparing Initial Release of Payroll Employment Change with Latent Estimate based on MCMC estimate including constant and slope, Thousands of Persons, 2019 sample: Feb 1961–Dec 2019, 2020 sample: Feb 1961–Dec 2020

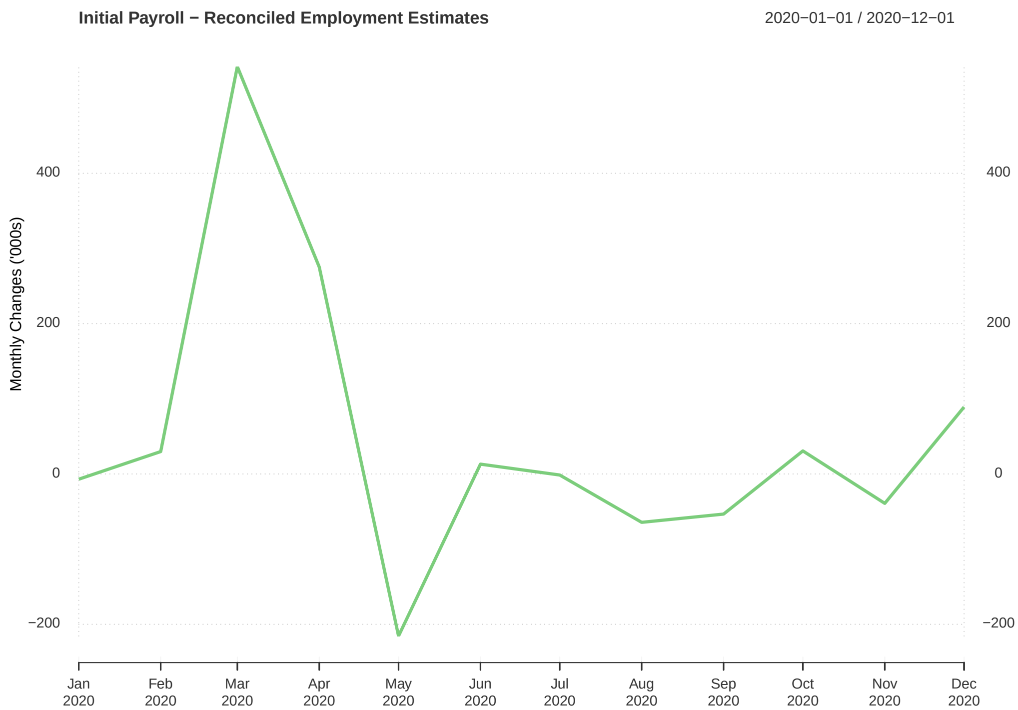

Figure 9: Comparing Initial Release of Payroll Employment Change with Latent Estimate based on MCMC estimate including constant and slope, Thousands of Persons, zooming in on 2020

List of Tables

Table 1: Descriptive statistics for first difference employment data Feb 1961–Dec 2019

	Initial Household dz	Initial Payroll dx1	First Revision Payroll dx2	Second Revision Payroll dx3
Mean	125	119	124	136
Median	157	145	157	165
Standard Deviation	350	190	197	202
Max	1310	733	1018	1103
Min	−1239	−674	−699	−741
Observations	707	707	707	707

Change in employment in ’000s.

Table 2: Descriptive statistics for first difference employment data covering Feb. 1961–Mar. 2019 for evaluation

	Initial Household dz	Initial Payroll dx1	First Revision Payroll dx2	Second Revision Payroll dx3	Benchmarked Payroll dfinal
Mean	124	119	124	136	138
Median	156	145	156	164	173
Standard Deviation	352	191	198	203	197
Max	1310	733	1018	1103	1118
Min	−1239	−674	−699	−741	−800
Observations	698	698	698	698	698

Change in employment in '000s.

Table 3: Descriptive statistics for first difference employment data during the Great Recession Jan 2008–June 2009

	Initial Household dz	Initial Payroll dx1	First Revision Payroll dx2	Second Revision Payroll dx3	Benchmarked Payroll dfinal
Mean	-326	-286	-305	-340	-410
Median	-291	-199.5	-302	-353	-401
Standard Deviation	372	241	248	253	270
Max	362	-17	-22	-47	11
Min	-1239	-663	-699	-741	-800
Observations	18	18	18	18	18

Change in employment in '000s.

Table 4: Prior Specification

Parameter	Domain	Density	Parameter 1	Parameter 2
ρ	$(0,1]$	Normal	MLE estimates	0.2
σ_ν^H	$(0,\infty)$	Inverse-Gamma	MLE estimates	10
σ_ν^1	$(0,\infty)$	Inverse-Gamma	MLE estimates	10
σ_ν^2	$(0,\infty)$	Inverse-Gamma	MLE estimates	10
σ_ν^3	$(0,\infty)$	Inverse-Gamma	MLE estimates	10
σ_ν^{H1}	$(0,\infty)$	Inverse-Gamma	MLE estimates	10
σ_ν^{H2}	$(0,\infty)$	Inverse-Gamma	MLE estimates	10
σ_ν^{H3}	$(0,\infty)$	Inverse-Gamma	MLE estimates	10
σ_ϵ^H	$(0,\infty)$	Inverse-Gamma	MLE estimates	10
σ_ϵ^1	$(0,\infty)$	Inverse-Gamma	MLE estimates	10
σ_ϵ^2	$(0,\infty)$	Inverse-Gamma	MLE estimates	10
σ_ϵ^3	$(0,\infty)$	Inverse-Gamma	MLE estimates	10
σ_ϵ^{H1}	$(0,\infty)$	Inverse-Gamma	MLE estimates	10
σ_ϵ^{H2}	$(0,\infty)$	Inverse-Gamma	MLE estimates	10
σ_ϵ^{H3}	$(0,\infty)$	Inverse-Gamma	MLE estimates	10
μ_z	$\mathbb{R}$	Normal	MLE estimates	10
β_z	$(0,\infty)$	Normal	MLE estimates	0.2

Note: Parameter 1 is the mean of the normal distribution and the mode value of the Inverse-Gamma distribution. Parameter 2 is the standard deviation of the normal distribution and the Inverse-Gamma distribution.

Table 5: Parameter estimates based on the Feb 1961–Dec 2019 sample

Estimation	MLE	MLE	MLE	MCMC	MCMC	MCMC
Constant	No	Yes	Yes	No	Yes	Yes
Slope	No	No	Yes	No	No	Yes
ρ	0.651	0.640	0.640	0.673	0.646	0.644
	(0.026)	(0.025)	(0.025)	(0.655,0.681)	(0.632,0.655)	(0.635,0.654)
σ_ν^H	0.116	0.120	0.138	0.219	0.216	0.229
	(9.993)	(11.153)	(2.448)	(0.136,0.435)	(0.134,0.360)	(0.151,0.387)
σ_ν^1	161.386	161.641	161.428	158.991	163.576	162.398
	(4.325)	(4.302)	(3.999)	(157.554,160.427)	(162.521,164.272)	(160.479,163.405)
σ_ν^2	59.757	59.782	59.781	59.646	59.900	59.642
	(1.590)	(1.591)	(1.588)	(58.869,60.720)	(59.646,60.208)	(58.636,60.650)
σ_ν^3	31.570	39.116	39.114	27.447	39.186	38.905
	(4.435)	(1.047)	(1.033)	(25.374,29.905)	(39.099,39.461)	(38.520,39.625)
σ_ν^{H1}	0.009	0.002	0.236	0.015	0.003	0.331
	(0.196)	(0.033)	(1.563)	(0.010,0.025)	(0.002,0.005)	(0.225,0.528)
σ_ν^{H2}	0.010	0.007	0.048	0.019	0.014	0.087
	(0.275)	(0.330)	(2.436)	(0.011,0.031)	(0.009,0.027)	(0.053,0.160)
σ_ν^{H3}	0.012	0.003	0.039	0.021	0.005	0.061
	(0.541)	(0.125)	(0.308)	(0.014,0.036)	(0.003,0.009)	(0.039,0.089)
σ_ϵ^H	312.151	317.102	316.681	310.432	315.400	321.766
	(10.304)	(8.826)	(8.570)	(304.651,314.450)	(314.447,316.437)	(317.941,324.715)
σ_ϵ^1	0.000	0.118	0.330	0.000	1.384	0.579
	(0.011)	(4.796)	(17.777)	(0.000,0.001)	(0.217,2.157)	(0.365,0.997)
σ_ϵ^2	0.465	0.118	0.153	0.998	0.220	0.282
	(2.949)	(4.109)	(1.940)	(0.514,1.540)	(0.138,0.353)	(0.178,0.462)
σ_ϵ^3	23.130	0.241	0.369	28.364	0.437	0.664
	(5.981)	(19.298)	(14.166)	(24.590,29.989)	(0.269,0.745)	(0.406,1.197)
σ_ϵ^{H1}	10.847	0.150	0.208	14.085	0.278	0.365
	(124.902)	(33.655)	(7.859)	(10.493,19.525)	(0.175,0.509)	(0.229,0.647)
σ_ϵ^{H2}	4.864	0.157	0.227	7.170	0.300	0.404
	(33.455)	(6.442)	(17.168)	(4.980,11.378)	(0.178,0.547)	(0.254,0.704)
σ_ϵ^{H3}	0.000	0.135	0.147	0.000	0.228	0.250
	(0.000)	(17.591)	(0.891)	(0.000,0.000)	(0.145,0.389)	(0.160,0.370)
μ_z		30.457	79.043		30.797	78.866
		(19.441)	(10.176)		(29.574,32.299)	(77.734,80.206)
β_z			0.564			0.580
			(0.052)			(0.566,0.595)
log likelihood	−17229.2	−17227.3	−17197.4	-17230.2	-17227.5	-17197.8

Notes. Model applied to data in first differences. Values between parentheses represent standard errors for MLE estimates and 25th% and 75th% draws for MCMC estimates. Prior: ρ: N(MLE estimate, 0.2), standard deviation: IG (mode: MLE estimates, SD: 10), Constant N (MLE estimate, 10), Slope: N (MLE estimate, 0.2).

Table 6: Kalman Gain full sample estimates Feb 1961–Dec 2019

Estimation	MLE	MLE	MLE	MCMC	MCMC	MCMC
Constant	No	Yes	Yes	No	Yes	Yes
Slope	No	No	Yes	No	No	Yes
Initial Household (dz)	0.0035	0.0000	0.0000	0.0040	0.0000	0.0000
Initial Payroll (dx1)	0.0000	0.0000	0.0000	0.0002	0.0000	0.0000
First Revision Payroll (dx2)	0.3480	0.0000	0.0001	0.5139	0.0001	0.0003
Second Revision Payroll (dx3)	0.6484	1.0000	0.9999	0.4820	0.9999	0.9997

Table 7: Kalman Gain nowcasting estimates Feb 1961–Dec 2019

Estimation	MLE	MLE	MLE	MCMC	MCMC	MCMC
Constant	No	Yes	Yes	No	Yes	Yes
Slope	No	No	Yes	No	No	Yes
Initial Household (dz)	0.0447	0.0483	0.0283	0.0427	0.0490	0.0280
Initial Payroll (dx1)	0.9553	0.9517	0.9841	0.9573	0.9510	0.9840
Initial Household (dz)	0.0101	0.0150	0.0086	0.0077	0.0152	0.0085
Initial Payroll (dx1)	0.0001	0.0000	0.0001	0.0003	0.0000	0.0001
First Revision Payroll (dx2)	0.9898	0.9850	0.9951	0.9920	0.9848	0.9951
Initial Household (dz)	0.0035	0.0000	0.0000	0.0040	0.0000	0.0000
Initial Payroll (dx1)	0.0000	0.0000	0.0000	0.0002	0.0000	0.0000
First Revision Payroll (dx2)	0.3480	0.0000	0.0001	0.5139	0.0001	0.0003
Second Revision Payroll (dx3)	0.6484	1.0000	0.9999	0.4820	0.9999	0.9997

Table 8: Mean Absolute Errors Comparing Latent with Benchmarked Employment

Estimation	MLE	MLE	MLE	MCMC	MCMC	MCMC
Constant	No	Yes	Yes	No	Yes	Yes
Slope	No	No	Yes	No	No	Yes
Full sample Feb 1961–Dec 2019	73.880	74.183	74.226	74.328	74.183	74.223
Nowcast	82.124	82.443	82.582	82.100	82.443	82.576
PR First release	82.504	82.504	82.504	82.504	82.504	82.504

Table 9: Parameter estimates comparing samples

Sample	1961–2020	1961-2019
Estimation	MCMC	MCMC
Constant	Yes	Yes
Slope	Yes	Yes
ρ	0.054	0.644
	(0.044,0.064)	(0.635,0.654)
σ_ν^H	182.168	0.229
	(175.874,189.106)	(0.151,0.387)
σ_ν^1	14.417	162.398
	(11.050,19.013)	(160.479,163.405)
σ_ν^2	55.857	59.642
	(54.025,57.549)	(58.636,60.650)
σ_ν^3	38.987	38.905
	(37.832,40.124)	(38.520,39.625)
σ_ν^{H1}	778.891	0.331
	(772.627,785.072)	(0.225,0.528)
σ_ν^{H2}	115.911	0.087
	(109.256,121.973)	(0.053,0.160)
σ_ν^{H3}	73.090	0.061
	(66.473,80.617)	(0.039,0.089)
σ_ϵ^H	115.438	321.766
	(109.164,121.950)	(317.941,324.715)
σ_ϵ^1	21.891	0.579
	(16.862,26.992)	(0.365,0.997)
σ_ϵ^2	5.555	0.282
	(4.628,6.461)	(0.178,0.462)
σ_ϵ^3	18.350	0.664
	(16.256,20.335)	(0.406,1.197)
σ_ϵ^{H1}	4.758	0.365
	(2.882,8.067)	(0.229,0.647)
σ_ϵ^{H2}	270.721	0.404
	(265.245,275.988)	(0.254,0.704)
σ_ϵ^{H3}	173.949	0.250
	(167.823,180.454)	(0.160,0.370)
μ_z	1.497	78.866
	(-0.018,2.951)	(77.734,80.206)
β_z	0.092	0.580
	(0.081,0.107)	(0.566,0.595)
log likelihood	-18806.5	-17197.8

Notes. Model applied to data in first differences. Parentheses represent 25th% and 75th% draws for MCMC estimates. Prior: ρ: N(MLE estimate, 0.2), standard deviation: IG (mode: MLE estimates, SD: 10), Constant N (MLE estimate, 10), Slope: N (MLE estimate, 0.2).