Macroeconomic Forecasting in the Time of COVID-19*

Giorgio E. Primiceri† Andrea Tambalotti‡

June 2020

Abstract

Forecasting the evolution of the U.S. economy following the outbreak of COVID-19 requires unusually strong assumptions. We propose a possible set of such assumptions. The main features of the resulting projections are that employment continues to fall for a few months even after the epidemic is in retreat, and that consumption does not recover to its pre-pandemic trend.

1 Introduction

The COVID-19 pandemic has been ravaging the world and its economies since March 2020. In April, the U.S. unemployment rate shot up to 14.7 percent, while personal consumption expenditures were almost 20 percent lower than at their peak in February, with entire sectors—such as entertainment and air transportation—grinding to a halt in large swaths of the country. Forecasting the macroeconomic effects of the pandemic over the next few years is crucial to direct the appropriate policy response, as well as to guide business and household decisions, but it is extremely challenging since this kind of economic shock is unprecedented in its type and scale. This is the task that we tackle in this paper, focusing more on the strong assumptions that are necessary to make some progress on the problem, than on the forecasts that we produce. Even if those forecasts turn out to be wildly off the mark, being explicit about the set of assumptions that

*We thank Domenico Giannone for numerous conversations on the topic. The views expressed in this paper are those of the authors and do not necessarily represent those of Amazon, the Federal Reserve Bank of New York or the Federal Reserve System.

†Northwestern University, CEPR and NBER

‡Federal Reserve Bank of New York

produce them will result in helpful insights on the underlying sources of economic fluctuations and their propagation.

There are at least two approaches to the problem of forecasting the macroeconomic effects of the Covid shock. One is to model the dynamics of the epidemic and its impact using explicit economic assumptions, as an exploding literature has done (e.g. Eichenbaum et al., 2020a, Eichenbaum et al., 2020b, Acemoglu et al., 2020, Baqaee and Farhi, 2020, Baqaee et al., 2020, Favero et al., 2020). The other approach is to use time-series data and techniques to learn from history about the dynamic effects of the Covid shock on the economy. Taken at face value, this approach might seem hopeless, since there is clearly no event in recent history that is directly comparable to the COVID-19 pandemic, with the possible exception of the Spanish influenza of 1918 and 1919 (e.g. Barro et al., 2020, Barro, 2020, Velde, 2020). This consideration might explain why the work that follows this approach is much sparser than the one based on more structural models. In this respect, the closest paper to ours is Ludvigson et al. (2020), who project the economic impact of COVID-19 based on that of deadly disasters in recent U.S. history.

In this paper, we show that it is possible to make some progress on the problem of predicting the economic effects of the Covid shock based on the historical dynamics of macroeconomic variables, without any behavioral assumptions. However, this progress requires unusually strong assumptions for the standards of reduced-form, time-series modeling. But since the assumptions of more structural models are even stronger, our approach is a viable complement to the augmented SIR models that many economists have recently developed.

Intuitively, our methodology is based on the idea of "synthesizing" a Coronavirus shock from the more typical disturbances that have historically driven macroeconomic fluctuations. This synthesis requires three key ingredients. First, we assume that the Covid "shock" was the dominant source of variation in macroeconomic variables in March and April 2020. This is the same assumption underlying event studies and narrative identification strategies in SVARS (e.g. Gurkaynak and Wright, 2013, Ben Zeev, 2018 and Antolin-Diaz and Rubio-Ramirez, 2018). This assumption never holds exactly in any event study, but it is an especially compelling approximation for the recent months, given the unprecedented effects of the COVID-19 pandemic on the economy since March 2020.

The second ingredient in the procedure is the assumption that the Covid shock will propagate over the next few months like a combination of more typical macroeconomic disturbances, whose transmission has been observed in the past. We choose this combination of historical shocks as the one that produces the same forecast error in the variables of interest as the one generated by the Covid shock in the early months of the epidemic. For example, suppose that a typical (negative) demand shock drives unemployment up by 5 and inflation down by 1. In contrast, a (negative) supply shock raises both unemployment and inflation by 1. If unemployment in March 2020 went up by 11, while inflation went down by 1 (to use very round numbers), we can reproduce the effect of the Covid shock with a combination of a demand shock of size 2 and a supply shock of size 1.

The two steps above produce a candidate synthetic Covid shock. However, if we simply projected its effects forward, the result would be a Great Depression, or worse. The reason is that these historical shocks are associated with persistent, often hump-shaped dynamics, which would imply a further deterioration in the economy over the coming months, rather than the fast improvement that most forecasters expect, and that some early data also seem to suggest. To address this issue, the third step in our procedure consists of "tilting" the propagation of the synthetic Covid shock to take into account that the diffusion path of the disease will likely differ from that of the macroeconomic disturbances responsible for conventional recessions. We implement this dynamic tilting by postulating some possible scenarios about the future evolution of the epidemic, and comparing them with the approximate persistence of the exogenous processes that dominate the fluctuations in estimated DSGE models.

In particular, we formulate three alternative scenarios on the development of the pandemic through the end of 2020 and beyond. Under the baseline scenario, the epidemic is gone by the end of 2020. In a more pessimistic scenario, a second wave of infections develops in the fall, while in the more optimistic one contagion ends in the fall. Under each of these scenarios, we project the implied behavior of the U.S. economy over the next three years. The main features of these projections are that employment continues to fall for a few months even as the epidemic retreats, and that consumption does not recover to its pre-pandemic trend. These recession patterns are more severe and prolonged under the second-wave scenario, but they remain present even under the optimistic scenario.

The rest of the paper proceeds as follows. Section 2 lays out the econometric framework for our forecasting exercise, and the necessary assumptions to synthetize the Covid shock and project its economic impact. Section 3 illustrates the VAR used to model historical macroeconomic dynamics, and presents the forecasting results based on various epidemic scenarios. Section 4 concludes.

2 The Problem of Prediction in the Time of COVID-19

Predicting the macroeconomic impact of the COVID-19 epidemic with reduced-form time-series models is challenging, because a pandemic of this scale was never directly observed in the available sample. To make some progress on this difficult prediction problem, we propose to "synthesize" the Covid shock through a combination of the macroeconomic disturbances that have been observed in the past, augmented with assumptions on the future dynamics of the pandemic. This section describes this approach formally, focusing on the strong assumptions needed to translate statements about the evolution of the epidemic over the next few months into a plausible account of its impact on the evolution of the US economy. We frame the discussion in the context of the problem of shock identification in state-space models, but subsequently focus on vector autoregressions to obtain our empirical results.

Consider the following two equations, describing the dynamics of an $n \times 1$ vector of macroeconomic variables y_t:

$$y_t = G(L)\,\xi_t \tag{1}$$

$$\xi_t = F(L)\,\varepsilon_t. \tag{2}$$

The first expression relates the evolution of y_t to a vector of exogenous variables ξ_t and their lags. The second equation states that ξ_t is a moving average of an $n \times 1$ vector of shocks ε_t, whose covariance matrix is normalized to I_n. The vector ε_t contains a set of orthogonal structural disturbances with an economic interpretation: they might represent taste, technology and policy shocks. Both $G(L) \equiv G_0 + \sum_{i=1}^{\infty} G_i L^i$ and $F(L) \equiv I_n + \sum_{i=1}^{\infty} F_i L^i$ are lag polynomial matrices of suitable dimensions and of potentially infinite order. We refer to the dynamics induced by $G(L)$ and $F(L)$ as the internal and external—or endogenous and exogenous—propagation of ε_t, being aware that their neat separation is typically only possible within a fully specified structural model. For instance, DSGE models usually include a number of exogenous disturbances that are often assumed to follow independent AR or ARMA processes (e.g. Smets and Wouters, 2007). In that case, $F(L)$ represents those AR or ARMA structures, while $G(L)$ captures the endogenous propagation of the exogenous processes due to the presence of persistent state variables, such as capital, and of forward looking expectations. We assume that y_t and ε_t have the same dimension because our empirical results are based on a VAR with the same number of shocks and observables. For the same reason, we assume that $G(L)$ and $F(L)$ imply a fundamental representation of y_t as a moving average of ε_t.

Under these assumptions, the combination of (1) and (2) yields the Wold representation

$$y_t = \Theta(L) G_0 \varepsilon_t, \tag{3}$$

where $u_t \equiv G_0\varepsilon_t$ is the vector of forecast errors and $\Theta(L) \equiv G(L) F(L) G_0^{-1}$, with $\Theta_0 = I_n$. According to (3), the effect of ε_t on y_{t+h} is given by $\Theta_h G_0$. Therefore, we must know $\{\Theta_i\}_{i=1}^{\infty}$ and G_0 to infer the dynamic impact of a structural shock on the endogenous variables. The coefficients $\{\Theta_i\}_{i=1}^{\infty}$ can typically be estimated from the data. In contrast, the identification of G_0 requires additional assumptions, since the data are only informative about the covariance matrix of the forecast errors, $G_0 G_0'$, but not about the impact matrix G_0. Therefore, inference about the dynamic effects of any one structural shock requires disentangling its impact on the endogenous variables from that of other disturbances.[1] Producing credible restrictions that yield the desired interpretation of (some of) the elements ε_t is among the most controversial steps in time-series analysis. This is the well-known identification problem in structural VARs.

In the case of the Covid shock, the problem is the opposite. On the one hand, the pandemic's short-term impact is straightforward to identify because the Covid shock is the overwhelming source of variation in the data around the time of the outbreak. Most standard macroeconomic variables have undergone record changes in and right after March 2020, which are clearly attributable to the effects of the virus and related shut-downs. These effects are easy to disentangle from those of other disturbances, since the latter are at least an order of magnitude smaller. On the other hand, the historical novelty of the Covid pandemic makes it difficult to estimate its future propagation from historical data, even conditional on a path for the pandemic itself. As a consequence, producing reasonable macroeconomic forecasts with time-series tools requires making some unusually strong assumptions. The need for these assumptions is the first message of this paper. In what follows, we describe a possible set of such assumptions and the forecasts that they produce.

We modify equation (3) to account for the effects of the pandemic by including a virus shock v_t as follows:

$$y_t = \Theta(L) G_0 \varepsilon_t + \theta(L) r(L) \circ r_0 v_t, \tag{4}$$

where the symbol "$\circ$" denotes the element-wise product of two vectors, $\theta(L) \equiv I_{n\times n} + \sum_{i=1}^{\infty} \theta_i L^i$

[1] For instance, a popular assumption for the identification of monetary policy shocks is that they move the policy rate, but not the other macroeconomic variables on impact (e.g. Christiano et al., 1999). This triangular assumption is imposed directly on the object of interest, G. However, G can also be identified through assumptions on the delayed response of the endogenous variables to some shocks (e.g. Blanchard and Quah, 1989, Gali, 1999, Mountford and Uhlig, 2009).

is an $n \times n$ lag-polynomial matrix, $r(L) \equiv 1_{n\times 1} + \sum_{i=1}^{\infty} r_i L^i$ is an $n \times 1$ lag-polynomial vector, r_0 is an $n \times 1$ vector, and v_t represents the Covid shock. This shock is equal to zero in all time periods except for the month in which the Coronavirus started to affect the U.S. economy, namely March 2020. We denote it by t^*. Therefore, r_0 represents the initial impact of the virus on y_t. Following t^*, we model the economic effects of the epidemic as unfolding deterministically, according to the lag polynomial $\theta(L) r(L) \circ r_0$, with no more "surprises". Of course, we are currently far from certain about those effects. In fact, they are what we want to estimate! But conceptually, the fact that the Coronavirus is spreading through the population and affecting the economy is now known. We capture the significant uncertainty on how that spread will occur —for instance regarding a potential second wave in the fall—through alternative scenarios, which we translate into coefficients of $r(L)$ according to a procedure described below.[2]

This procedure is based on the following three assumptions.

Assumption 1: v_t accounts for all the unexpected variation of $y_{t^*},\ y_{t^*+1}, ...,\ y_{t^*+j^*}$, where $t^* + j^*$ is the last period for which y_t is observable.
All event studies make a similar assumption, which is often valid only in a restricted window of time after the event (Gurkaynak and Wright, 2013). In the case of the Covid shock, its macroeconomic effects have been so large that the assumption is likely to remain valid for several months after March 2020.

Assumption 2: $\theta(L) = \Theta(L)$.
This assumption states that, if $r(L)$ was a lag polynomial of order zero, the pandemic shock v_t would propagate like the specific combination of structural disturbances ε_t that produces the same forecast error for the observables. Absent $r(L)$, this assumption would thus allow to predict the effects of the Covid shock as in a standard forecasting exercise, starting from the current value of the observable variables, and projecting them into the future using their historical dynamics. The lag polynomial $r(L)$ then tilts this propagation mechanism.

But where does $r(L)$ come from? $\{r_j\}_{j=0}^{j^*}$ can actually be easily estimated. To see how, rewrite (4) as

$$\Theta(L)^{-1} y_{t^*} = G_0 \varepsilon_{t^*} + r(L) \circ r_0 \approx r(L) \circ r_0 \tag{5}$$

where $\Theta_0 \equiv I$. This expression uses $\theta(L) = \Theta(L)$ from assumption 2, $v_{t^*} = 1$ as a normalization,

[2]There are two alternative approaches that could be adopted to model the effects of Covid. One strategy would be to capture the unfolding of the pandemic in future months as a sequence of shocks, as opposed to a single shock that hits in March 2020 and then propagates. Another method would be to explicitly model the uncertainty faced by economic agents about the possible dynamics of the virus following the initial shock in March 2020. Our approach based on the analysis of several pandemic scenarios should be interpreted as a simple and viable version of the latter.

and the approximate equality follows from assumption 1. Given an estimate of $\{\Theta\}_{j=1}^{\infty}$, which is easy to obtain, (5) implies that $\{r_j\}_{j=0}^{j^*}$ can be inferred from the forecast errors at time $\{t^*+j\}_{j=0}^{j^*}$, without additional restrictions. In contrast, the data are not informative about $\{r_j\}_{j>j^*}$, which motivates a third assumption.

Assumption 3: $\{r_j\}_{j>j^*} = \{\hat{r}_j\}_{j>j^*}$, where $\{\hat{r}_j\}_{j>j^*}$ must be specified a priori.
To choose $\{\hat{r}_j\}_{j>j^*}$, it is useful to clarify the interpretation of $r(L)$. Use assumption 2 and the definition of $\Theta(L)$ to write (4) as

$$y_t = G(L)\left[F(L)\,G_0^{-1}G_0\varepsilon_t + F(L)\,G_0^{-1}\,r(L)\circ r_0 v_t\right].$$

In this expression, $F(L)\,G_0^{-1}$ captures the exogenous propagation of the forecast error $G_0\varepsilon_t$. Along similar lines, $f(L) \equiv F(L)\,G_0^{-1}\,r(L)$ characterizes the external propagation of $r_0 v_t$. It follows that $r(L)$ should be interpreted as capturing the exogenous dynamics of the forecast errors generated by the Covid shock, *relative* to those associated to the forecast errors generated by the standard macroeconomic shocks. The next section shows how to use this insight to calibrate$\{r_j\}_{j>j^*}$ based on alternative pandemic scenarios.

The ability to consider different scenarios about the unfolding of the pandemic represents an advantage of our approach. One important limitation is that we are implicitly assuming that agents form expectations about the future as if the external propagation of the Covid shock were similar to that of standard forecast errors. Sizable deviations from this benchmark, as captured by the complexity of the lag polynomial $r(L)$,would violate the Lucas' critique. In this respect, our approach is similar to the analysis of "modest policy interventions," which are defined as having a modest impact on agents' beliefs about the prevalence of a certain policy regime (Leeper and Zha, 2003). Another recent example of an approach in this spirit is the approximation of macroeconomic dynamics at the zero lower bound based on the assumption that interest rates will predictably remain at zero for a pre-specified period of time. This strategy, which is usually implemented through the addition of anticipated shocks in the interest rate rule, has been widely adopted in the DSGE literature (e.g. Laseen and Svensson, 2011, Campbell et al., 2012, Cocci et al., 2013).

A final note of caution concerns the linearity of the model. This assumption is arguably more problematic now then in normal times, given the exceptionally large macroeconomic volatility induced by the Covid shock, relative to historical standards.

3 A Monthly Time-Series Model of the U.S. Economy

This section levers the two assumptions discussed above to forecast the evolution of the U.S. economy following the Covid outbreak. To obtain the parameters of the vector moving-average process in (3), we invert the vector autoregression (VAR)

$$y_t = c + B_1 y_{t-1} + ... + B_p y_{t-p} + u_t,$$

where c is a vector of constant terms. This VAR is estimated with six key U.S. macroeconomic indicators observed at monthly frequency: (i) *employment*, measured by the logarithm of the total number of nonfarm employees (ii) *unemployment*, measured by the civilian unemployment rate; (iii) *consumption*, measured by the logarithm of real personal consumption expenditures; (iv) *industrial production*, measured by the logarithm of the industrial production index; (v) *CPI*, measured by the logarithm of the consumer price index; and (vi) *core PCE*, measured by the logarithm of the price index of personal consumption expenditures excluding food and energy.

The VAR includes 13 lags and it is estimated on the sample 1988:12 to 2020:2. The analysis starts in 1988:12 and not earlier to capture the muted reaction of inflation to economic fluctuations that characterizes the last three decades, as documented in Del Negro et al. (2020b). The estimation stops in 2020:2 to prevent the massive movements in the variables observed starting in March 2020 from affecting the parameter estimates. The model is estimated with Bayesian methods and a standard Minnesota prior, whose tightness is chosen based on the data-driven approach of Giannone et al. (2015).

3.1 Scenarios

The crucial element of our procedure to forecast the economic impact of Covid-19 is the choice of $r(L)$. This choice consists of two parts. First, we infer $\{r_j\}_{j=0}^{j^*}$ using data through April 2020. Second, we calibrate $\{r_j\}_{j>1}$ by designing scenarios for $f(L)$, which captures the exogenous dynamics of the epidemic, and translating $f(L)$ into $r(L)$ as described below.

At the time of writing, we have data for March 2020—the month in which the Covid shock first hit the U.S. economy—and April 2020. As a result, $j^* = 1$ and we can easily estimate r_0 and r_1 using the VAR forecast errors in March and April, as described in section 2. As for the parameters $\{r_j\}_{j>1}$, we select them by considering three alternative choices of $f(L)$, corresponding to three possible scenarios about the evolution of the pandemic.

1. A baseline scenario, in which the pandemic gradually disappears by the end of 2020, following the path
$$f_j = \frac{1}{1+e^{j-5.5}} f_1, \quad \forall j \geq 2.$$

2. An optimistic scenario, in which the pandemic gradually disappears by September 2020, following the path
$$[f_2, f_3, f_4, f_5] = [.9, .6, .3, .1]\, f_1 \text{ and } f_j = 0_{n\times 1}, \quad \forall j \geq 6.$$

3. A pessimistic scenario that starts like the baseline, but exhibits a second wave of infections starting in November 2020 and peaking in February 2021, according to the path
$$f_j = \left[\frac{1}{1+e^{j-5.5}} f_1 + e^{-\frac{1}{4}(j-11)^2}\right] f_1, \quad \forall j \geq 2.$$

Figure 3.1 depicts the three pandemic scenarios, by plotting the coefficients that relate f_j to f_1. These scenarios should be interpreted as statements about the "intensity" of the epidemic, relative to its intensity in April 2020. For now, we choose these scenarios a priori, but one could generate them by running an epidemiological model and looking at a summary statistic of its outcomes, such as the number of infections. In this sense, the fact that the second wave under the pessimistic scenario eventually reaches a peak of one can be loosely interpreted as implying that the pandemic will be as "bad" in February 2021 as it was in April 2020, without being very precise on how "bad" should be measured.

The final step in the procedure consists in transforming $\{f_j\}_{j>1}$ into the corresponding $\{r_j\}_{j>1}$, using the relation $r(L) = G_0 F(L)^{-1} f(L)$. This translation is challenging because computing $G_0 F(L)^{-1}$ would require identifying the structural macroeconomic shocks (i.e. knowledge of G_0), as well as taking a stance on their exogenous transmission (i.e. knowledge of $F(L)$). This is only possible within a fully specified structural model. As a less restricted alternative, we approximate the exogenous dynamics of the forecast errors in normal times with those of the dominant exogenous process in a benchmark DSGE model. This approximation is motivated by the fact that most estimated DSGE models typically attribute the bulk of variation in the data to a single exogenous process (e.g. Smets and Wouters, 2007, Justiniano et al., 2010, Justiniano et al., 2011). If these models are correct on the sources of business cycle fluctuations, the dynamics implied by $G_0 F(L)^{-1}$ should be similar to those of this dominant exogenous process. For example, Justiniano et al. (2010) find that the dominant business cycle shock is an investment shock, which

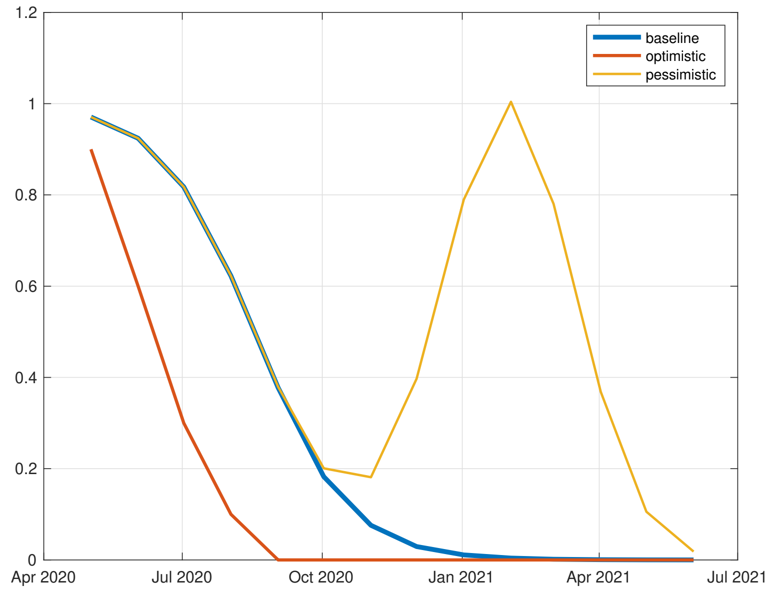

Figure 1: Pandemic scenarios, relative to April 2020.

follows an AR(1) process with coefficient 0.72 when estimated on quarterly data. In monthly data, this estimate implies a persistence coefficient of 0.9, which we use as reference for the calculations below. Therefore, we set $G_0F\left(L\right)^{-1} = G_0\left(1 - 0.9L\right)$, so that $r_j = G_0\left(f_j - 0.9f_{j-1}\right)$. The last expression can be rewritten as $r_j = G_0f_1\left(s_j - 0.9s_{j-1}\right)$, where s_j are the coefficients used to link f_j to f_1 in the scenarios of section 3.1. Finally, since $f_1 = G_0^{-1}\left(r_1 + 0.9 \cdot 1_{n\times 1}\right)$, it follows that $r_j = \left(r_1 + 0.9 \cdot 1_{n\times 1}\right)\left(s_j - 0.9s_{j-1}\right)$.

3.2 Forecasting results

Figure 2 plots the projected impact of the Covid shock on the six variables included in our VAR under the baseline scenario. The most notable feature of these responses is that the performance of the economy continues to deteriorate for a few months after April 2020, when the epidemic starts to unwind, according to our assumptions. Employment falls by more than 20 percent, reaching a through in August 2020 and the unemployment rate shoots above 20 percent before declining steadily, reaching 5 percent at the beginning of 2023. Consumption also declines by almost 20 percent and industrial production by 40 percent. Inflation reacts less than the real variables, consistent with the evidence presented in Del Negro et al. (2020b) and Del Negro et al. (2020a). Core PCE inflation does become negative, but it returns to 2 percent in mid 2021. Another feature of these forecasts is that consumption starts growing after July 2020, but it

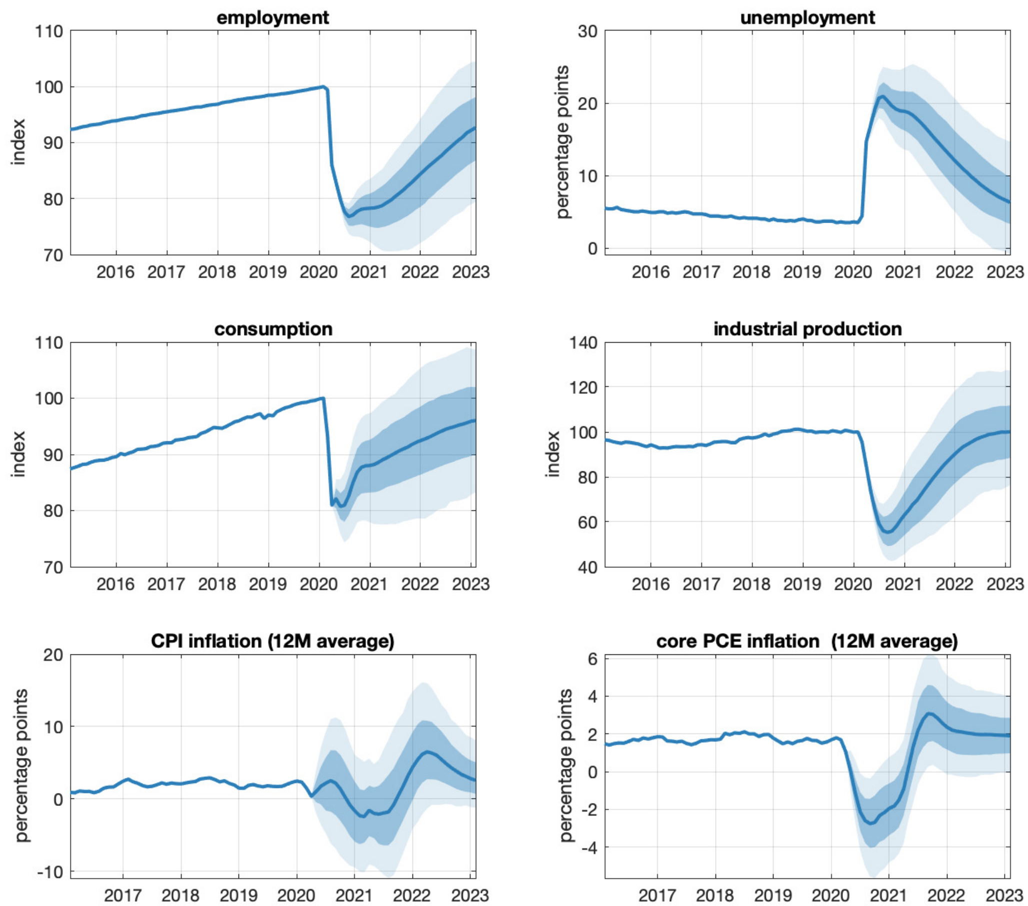

Figure 2: Baseline scenario: Forecasts of the impact of the Covid shock. The solid lines represent actual data until April 2020, and posterior medians of the predictive density afterwards. The shaded areas correspond to 68- and 95-percent posterior credible regions.

never fully return to its pre-Covid trend. This pattern is similar to that followed by consumption during the recovery from the Great Recession. As a complement to the figure, table 3.2 translates the projected path of consumption into (annualized) quarterly and annual growth rates for the period 2020-2021, in a format similar to the Survey of Professional Forecasters (SPF).

Figures 3 and 4 present the outcome of our forecasting exercise under the more optimistic and pessimistic scenarios of figure 3.1. As expected, the economy recovers at a faster pace when the pandemic dissipates more quickly. Relative to the baseline scenario, all the variables reach their trough two months earlier, in May 2020. In addition, the fall in employment and industrial production are less pronounced and unemployment peaks near 15 percent. In comparison, the forecasts under the scenario that includes a second wave of infections in the fall are considerably

	2020			2021	2022
scenario	Q2	Q3	Q4		
baseline	-67.2 $[-72.4 \quad -61.7]$	7.3 $[-5.3 \quad 19.4]$	22.2 $[10.8 \quad 33.0]$	3.1 $[-1.8 \quad 7.5]$	4.7 $[1.7 \quad 7.9]$
optimistic	-57.2 $[-62.5 \quad -51.6]$	40.8 $[31.4 \quad 51.8]$	2.9 $[-5.3 \quad 11.5]$	3.9 $[1.1 \quad 6.4]$	3.9 $[2.1 \quad 5.8]$
pessimistic	-67.2 $[-72.4 \quad -61.7]$	7.3 $[-5.3 \quad 19.4]$	7.4 $[-3.4 \quad 17.8]$	-8.1 $[-14.4 \quad -1.7]$	8.6 $[3.1 \quad 13.3]$

Table 1: Projected annualized quarterly growth rate (in percentage points) of real personal consumption expenditures. Medians and 68-percent credible intervals of the predictive density.

gloomier, both in terms of duration and severity of the recession.

The translation of the various epidemic scenarios into the dynamic impact of the virus on the economy is conditional on the assumed approximation of $F(L)^{-1}$, which we based on a simple back-of-the-envelope calculation. However, the relative responses across scenarios and across variables are less dependent on the choice of this matrix, which is mostly relevant to determine the overall dynamic profile of the responses. We exploit this fact in one of the robustness exercises below, in which we approximate $F(L)^{-1}$ so that the VAR forecast for unemployment resembles the consensus in the SPF. In that exercise, the overall dynamics of the economy in response to the pandemic are pinned down by the views embedded in the SPF consensus for unemployment, but the VAR is left to determine the implications of those dynamics for the other variables.

3.3 Robustness

This subsection studies the sensitivity of our forecasts to some modeling choices. Figure 5 displays some robustness to the approximation of $F(L)^{-1}$. The blue lines correspond to the approximation $F(L)^{-1} \approx I_n(1-0.9L)$, which underlies the forecasts in the previous subsection. The grey and red lines correspond to forecasts obtained with coefficients of 0.85 and 0.95 respectively. The latter implies a path of unemployment that is closer to the consensus SPF forecast, as mentioned above. The duration and severity of the recession varies somewhat across calibrations, but the relative response of the different variables is more similar.

Figure 6 presents forecasts under the baseline scenario, based on a VAR estimated with data

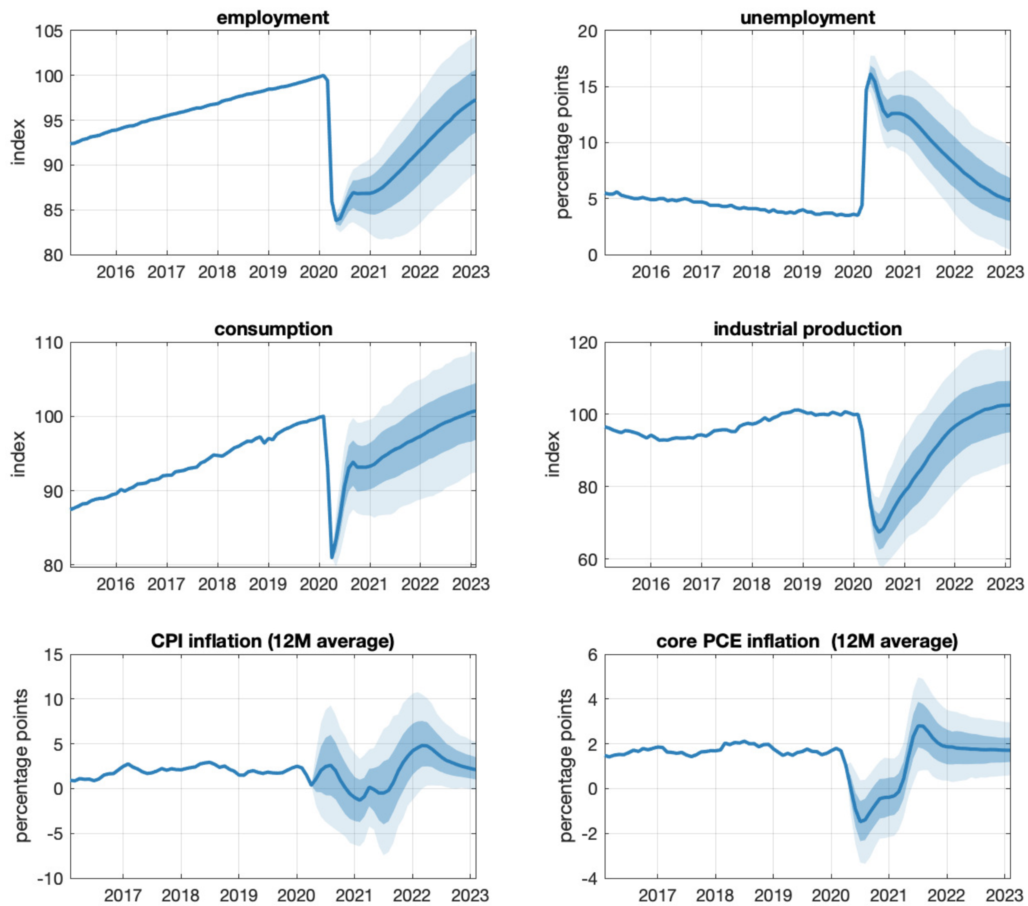

Figure 3: Optimistic scenario: Forecasts of the impact of the Covid shock. The solid lines represent actual data until April 2020, and posterior medians of the predictive density afterwards. The shaded areas correspond to 68- and 95-percent posterior credible regions.

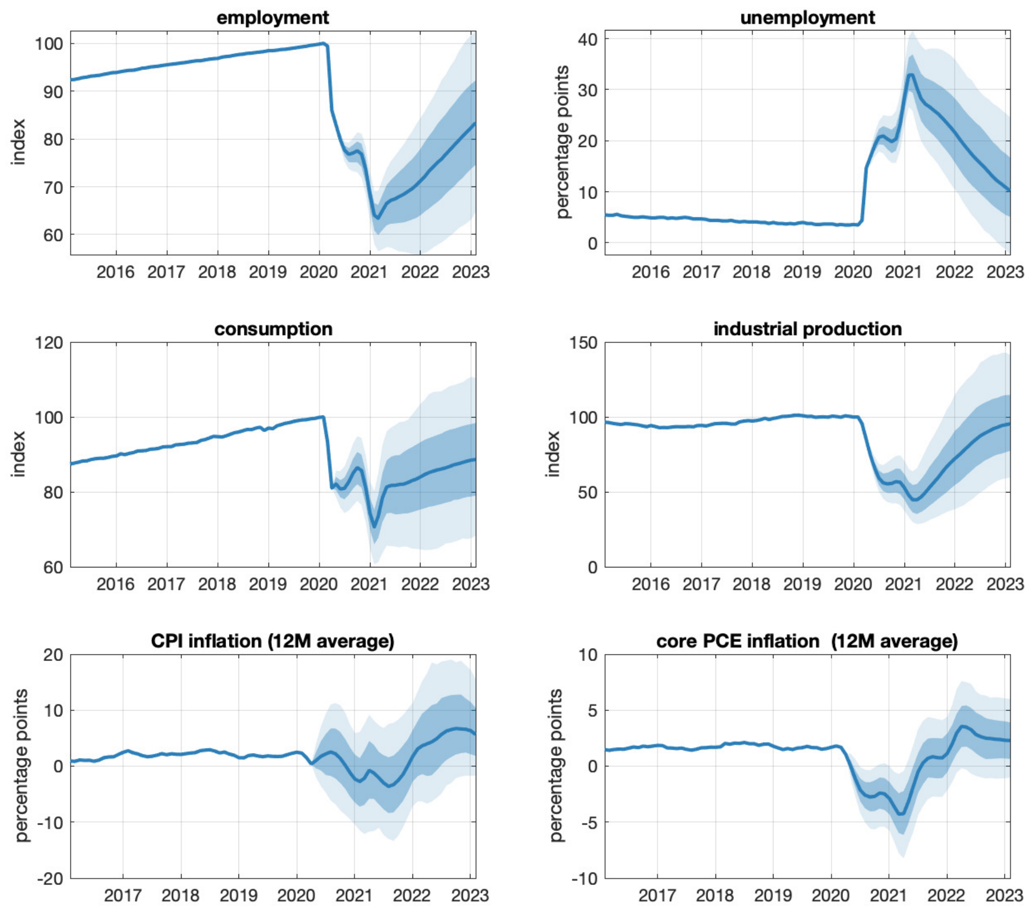

Figure 4: Pessimistic scenario: Forecasts of the impact of the Covid shock. The solid lines represent actual data until April 2020, and posterior medians of the predictive density afterwards. The shaded areas correspond to 68- and 95-percent posterior credible regions.

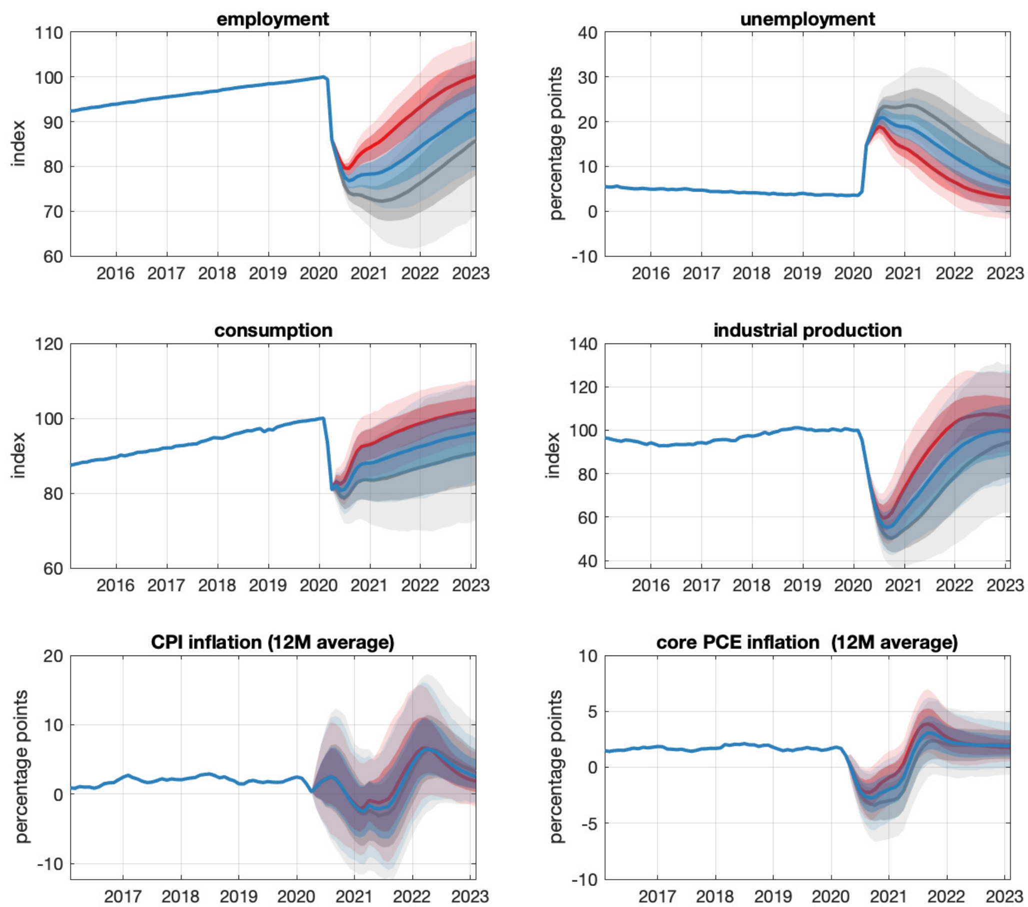

Figure 5: Baseline scenario: Forecasts of the impact of the Covid shock under three different assumptions about the external dynamics of shocks in normal times. The solid lines represent actual data until April 2020, and posterior medians of the predictive density afterwards. The shaded areas correspond to 68- and 95-percent posterior credible regions.

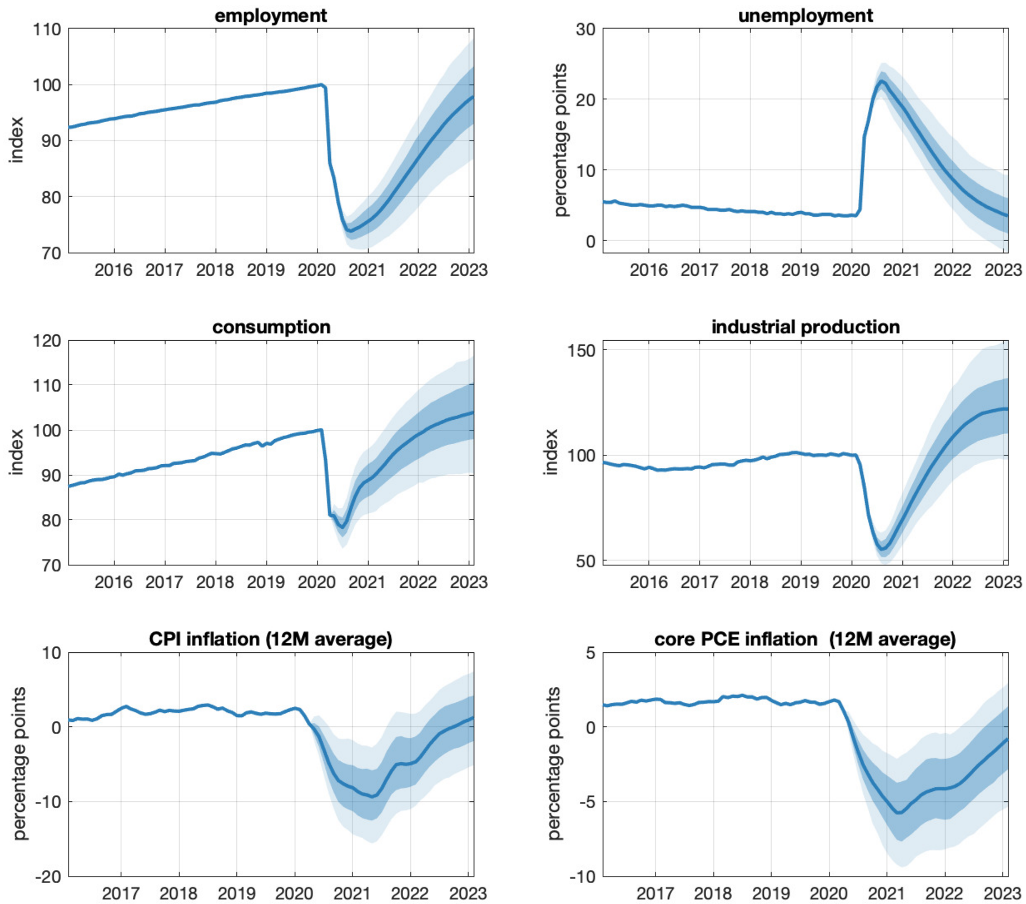

Figure 6: Baseline scenario: Forecasts of the impact of the Covid shock when the VAR is estimated using data from 1959:1. The solid lines represent actual data until April 2020, and posterior medians of the predictive density afterwards. The shaded areas correspond to 68- and 95-percent posterior credible regions.

starting in 1959. These projections have two notable features. First, deflation is deeper and more entrenched than under the post-1990 dynamics, persisting through 2023. Second, consumption recovers faster, eventually returning to its pre-Covid trend. Nonetheless, the effects on employment and unemployment are similar to those in the baseline forecasts.

Figure 7 is based on a VAR that includes Real Disposable Personal Income and the number of Job Losers on Layoff from the Employment Report. This information should help to capture the effects of the massive intervention to sustain incomes by the fiscal authority, as well as the possible change in business cycle dynamics associated with the unusually large number of temporary layoffs during the pandemic, relative to more conventional recessions. However, including these variables does not change the projections much. This is also the case when the baseline VAR is

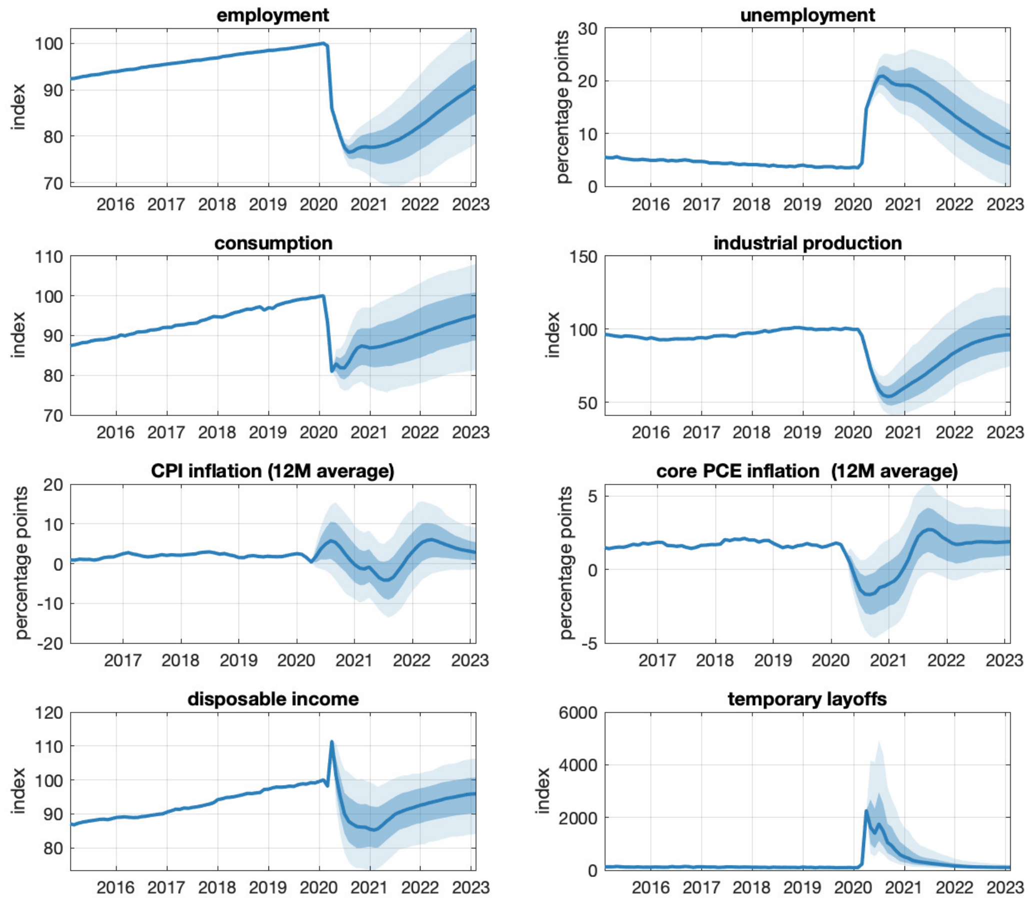

Figure 7: Baseline scenario: Forecasts of the impact of the Covid shock when the VAR is augmented with data on the logarithm of Real Disposable Personal Income and the number of Job Losers on Layoff. The solid lines represent actual data until April 2020, and posterior medians of the predictive density afterwards. The shaded areas correspond to 68- and 95-percent posterior credible regions.

augmented to include the corporate bond spread of Gilchrist and Zakrajsek (2012), a popular measure of financial stress, as shown in figure 8.

The final set of additional results is motivated by the fact that the pandemic has had a markedly different impact across sectors, with those producing "social" goods hit much harder than the rest. . To capture the effects of the asymmetric sectoral distribution of the shock , we estimated a large-scale VAR that includes employment data for eleven major industry sectors. Given its relatively large number of variables, we estimate this model on data from 1959. The forecasts it produces are in figures 9 and 10. Perhaps surprisingly, the projections for the aggregate variables are similar to those in the baseline, although they are noisier due to the large number of

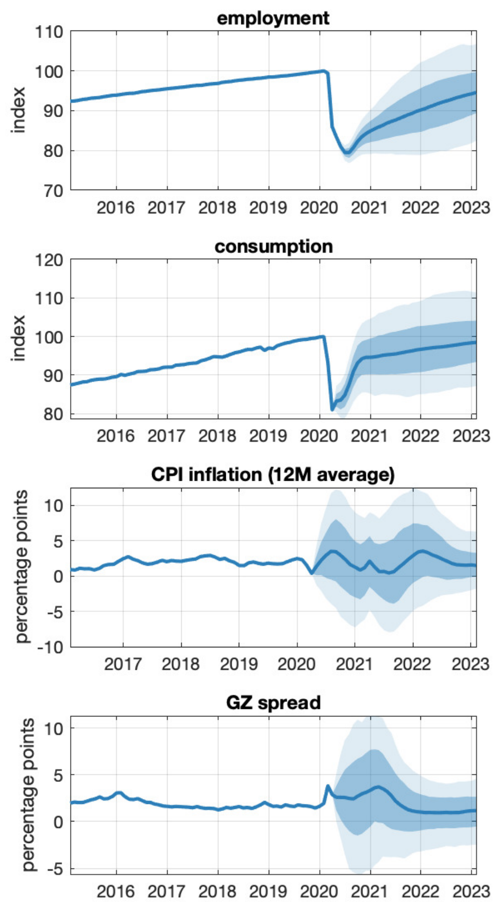

Figure 8: Baseline scenario: Forecasts of the impact of the Covid shock when the VAR is augmented with data on the Gilchrist and Zakrajšek (2012) spread. The solid lines represent actual data until April 2020, and posterior medians of the predictive density afterwards. The shaded areas correspond to 68- and 95-percent posterior credible regions.

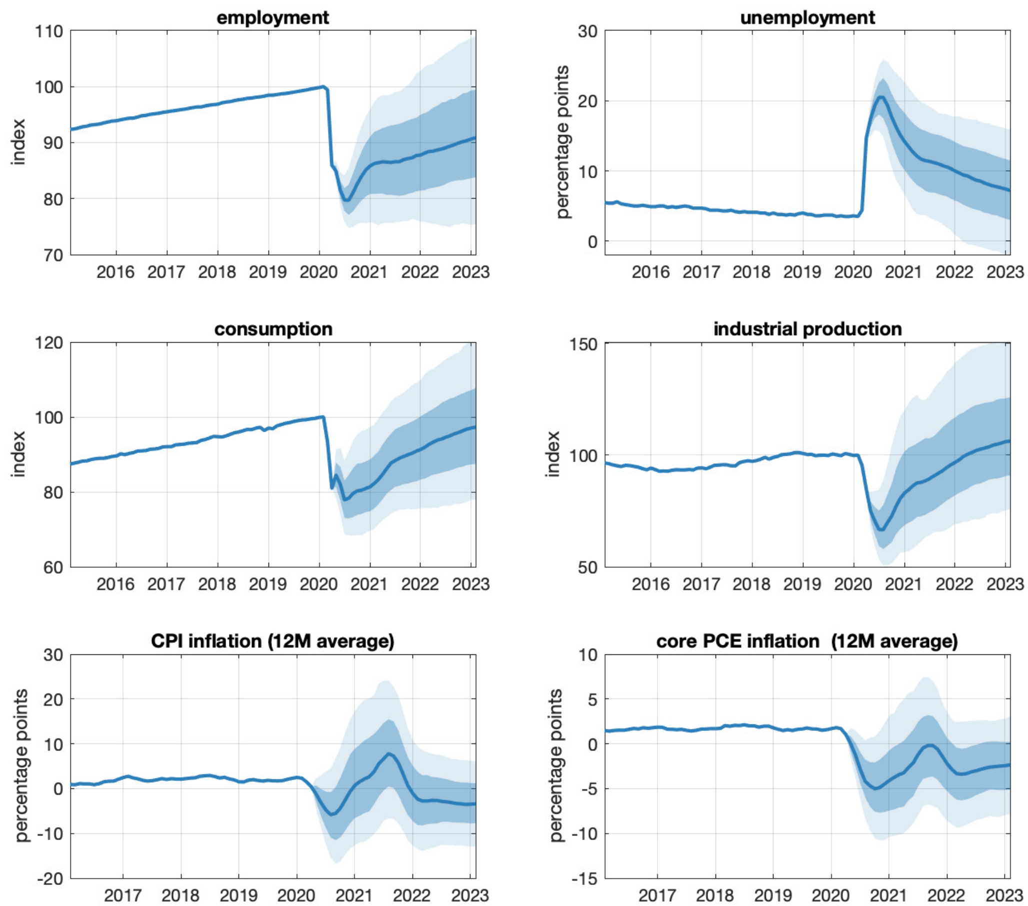

Figure 9: Baseline scenario: Forecasts of the impact of the Covid shock when the VAR is augmented with employment sectoral data. The solid lines represent actual data until April 2020, and posterior medians of the predictive density afterwards. The shaded areas correspond to 68- and 95-percent posterior credible regions.

estimated parameters.

4 Concluding remarks

Forecasting the evolution of the U.S. economy in response to the COVID-19 pandemic is crucial, but challenging, since an economic shock of this kind was never observed in the recorded past. This paper has articulated a set of assumptions that allow to build a forecast distribution for a vector of monthly macroeconomic variables based on their historical evolution. Given the uncertainty on the future of the epidemic, as well as the difficulties in translating a path for the disease into its implications for the economy, the forecasts presented above are not more than a

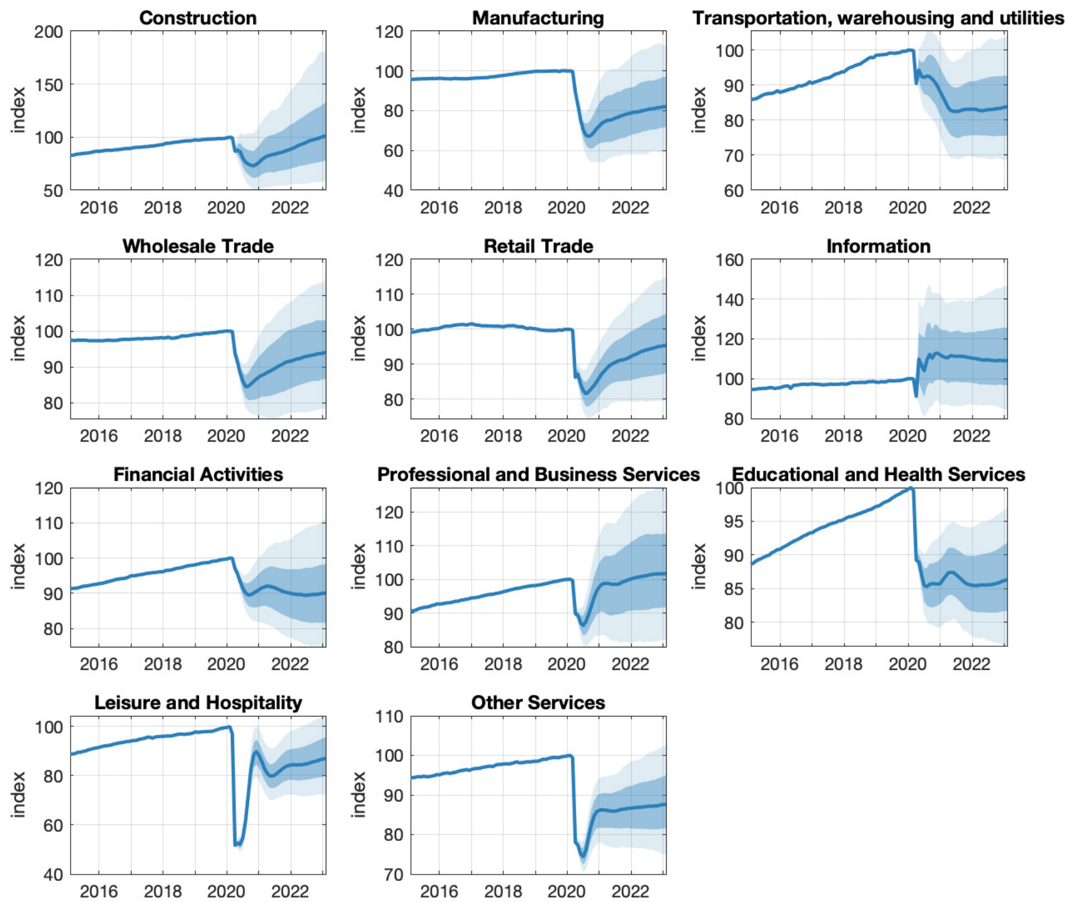

Figure 10: Baseline scenario: Forecasts of the impact of the Covid shock when the VAR is augmented with employment sectoral data. The solid lines represent actual data until April 2020, and posterior medians of the predictive density afterwards. The shaded areas correspond to 68- and 95-percent posterior credible regions.

rough guide to what might happen to the U.S. economy over the next year or two, conditional on the assumptions on which they are built. Even if those assumptions and the resulting forecasts are not borne out by the data, this exercise should produce valuable insights into the origins of economic fluctuations and their propagation across sectors of the economy and over time.

References

ACEMOGLU, D., V. CHERNOZHUKOV, I. WERNING, AND M. D. WHINSTON (2020): "Optimal Targeted Lockdowns in a Multi-Group SIR Model," NBER Working Papers 27102, National Bureau of Economic Research, Inc.

ANTOLIN-DIAZ, J. AND J. F. RUBIO-RAMIREZ (2018): "Narrative Sign Restrictions for SVARs," *American Economic Review*, 108, 2802–29.

BAQAEE, D., E. FARHI, M. J. MINA, AND J. H. STOCK (2020): "Reopening Scenarios," NBER Working Papers 27244, National Bureau of Economic Research, Inc.

BAQAEE, D. R. AND E. FARHI (2020): "Supply and Demand in Disaggregated Keynesian Economies with an Application to the Covid-19 Crisis," CEPR Discussion Papers 14743, C.E.P.R. Discussion Papers.

BARRO, R. J. (2020): "Non-Pharmaceutical Interventions and Mortality in U.S. Cities during the Great Influenza Pandemic, 1918-1919," NBER Working Papers 27049, National Bureau of Economic Research, Inc.

BARRO, R. J., J. F. URSUA, AND J. WENG (2020): "The Coronavirus and the Great Influenza Pandemic: Lessons from the "Spanish Flu" for the Coronavirus's Potential Effects on Mortality and Economic Activity," NBER Working Papers 26866, National Bureau of Economic Research, Inc.

BEN ZEEV, N. (2018): "What can we learn about news shocks from the late 1990s and early 2000s boom-bust period?" *Journal of Economic Dynamics and Control*, 87, 94–105.

BLANCHARD, O. J. AND D. QUAH (1989): "The Dynamic Effects of Aggregate Demand and Supply Disturbances," *American Economic Review*, 79, 655–673.

CAMPBELL, J. R., C. L. EVANS, J. D. FISHER, AND A. JUSTINIANO (2012): "Macroeconomic Effects of Federal Reserve Forward Guidance," *Brookings Papers on Economic Activity*, 43, 1–80.

CHRISTIANO, L. J., M. EICHENBAUM, AND C. L. EVANS (1999): "Monetary policy shocks: What have we learned and to what end?" in *Handbook of Macroeconomics*, ed. by J. B. Taylor and M. Woodford, Elsevier, vol. 1 of *Handbook of Macroeconomics*, chap. 2, 65–148.

COCCI, M., M. D. NEGRO, S. EUSEPI, M. GIANNONI, R. B. HASEGAWA, M. H. LINDER, A. M. SBORDONE, AND A. TAMBALOTTI (2013): "The FRBNY DSGE model," Staff Reports 647, Federal Reserve Bank of New York.

DEL NEGRO, M., M. LENZA, G. E. PRIMICERI, AND A. TAMBALOTTI (2020a): "What's up with Inflation and the Business Cycle after the COVID-19 Shock?" Mimeo, Northwestern University.

——— (2020b): "What's up with the Phillips Curve?" *Brookings Papers on Economic Activity*, forthcoming.

EICHENBAUM, M. S., S. REBELO, AND M. TRABANDT (2020a): "The Macroeconomics of Epidemics," NBER Working Papers 26882, National Bureau of Economic Research, Inc.

——— (2020b): "The Macroeconomics of Testing and Quarantining," NBER Working Papers 27104, National Bureau of Economic Research, Inc.

FAVERO, C. A., A. ICHINO, AND A. RUSTICHINI (2020): "Restarting the economy while saving lives under Covid-19," CEPR Discussion Papers 14664, C.E.P.R. Discussion Papers.

GALI, J. (1999): "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review*, 89, 249–271.

GIANNONE, D., M. LENZA, AND G. E. PRIMICERI (2015): "Prior Selection for Vector Autoregressions," *The Review of Economics and Statistics*, 97, 436–451.

GILCHRIST, S. AND E. ZAKRAJSEK (2012): "Credit Spreads and Business Cycle Fluctuations," *American Economic Review*, 102, 1692–1720.

GURKAYNAK, R. S. AND J. H. WRIGHT (2013): "Identification and Inference Using Event Studies," *Manchester School*, 81, 48–65.

JUSTINIANO, A., G. PRIMICERI, AND A. TAMBALOTTI (2011): "Investment Shocks and the Relative Price of Investment," *Review of Economic Dynamics*, 14, 101–121.

JUSTINIANO, A., G. E. PRIMICERI, AND A. TAMBALOTTI (2010): "Investment shocks and business cycles," *Journal of Monetary Economics*, 57, 132–145.

LASEEN, S. AND L. E. SVENSSON (2011): "Anticipated Alternative Policy-Rate Paths in Policy Simulations," *International Journal of Central Banking*, 7, 1–35.

LEEPER, E. M. AND T. ZHA (2003): "Modest policy interventions," *Journal of Monetary Economics*, 50, 1673–1700.

LUDVIGSON, S. C., S. MA, AND S. NG (2020): "Covid19 and the Macroeconomic Effects of Costly Disasters," NBER Working Papers 26987, National Bureau of Economic Research, Inc.

MOUNTFORD, A. AND H. UHLIG (2009): "What are the effects of fiscal policy shocks?" *Journal of Applied Econometrics*, 24, 960–992.

SMETS, F. AND R. WOUTERS (2007): "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, 97, 586 – 606.

VELDE, F. (2020): "What Happened to the US Economy During the 1918 Influenza Pandemic? A View Through High-Frequency Data," Federal Reserve Bank of Chicago Working Paper No. 2020-11.