Federal Reserve Bank of Chicago

Unusual Shocks in Our Usual Models

Filippo Ferroni, Jonas D. M. Fisher, and Leonardo Melosi

REVISED
April 30, 2024

WP 2022-39

https://doi.org/10.21033/wp-2022-39

Unusual shocks in our usual models[*]

Filippo Ferroni   Jonas D. M. Fisher   Leonardo Melosi

April 30, 2024

Abstract

We propose a method to allow usual business cycle models to account for the unusual COVID episode. The pandemic and the public and private responses to it are represented by a new shock called the *Covid shock*, which loads onto wedges that underlie the usual shocks and comes with news about its evolution. We apply our method to a standard medium-scale model, estimating the loadings with 2020q2 data and the evolving news using professional forecasts. It accounts for most of the early macroeconomic dynamics, was inflationary and a persistent drag on activity, and the majority of its effects were unanticipated. We also show how the Covid shock can be used to estimate DSGE models with data before, during, and after the pandemic.

JEL Classification Numbers: C51, E10, E31, E32, E52
Keywords: COVID-19, pandemic, DSGE models, Survey of Professional Forecasters, business cycles

[*]All views expressed are the authors' and do not necessarily represent those of the Federal Reserve Bank of Chicago or the Federal Reserve System. Ferroni and Fisher are affiliated with Federal Reserve Bank of Chicago. Melosi is affiliated with the Federal Reserve Bank of Chicago, the University of Warwick, and the CEPR. Our e-mail addresses are: jonas.fisher@chi.frb.org, fferroni@frbchi.org, and lmelosi@gmail.com. The not-for-publication appendix that begins on page 36 is co-authored with Jeffrey Campbell, University of Notre Dame and Tilberg University, jcampbel24@nd.edu. We thank our referees, Lisa Barrow, Martin Eichenbaum, Domenico Giannone, Spencer Krane, Michele Lenza, Giorgio Primiceri, Sergio Rebelo, Barbara Rossi, Frank Schorfheide, Andrea Tambalotti, Mathias Trabandt, and seminar participants at the 2021 NBER Summer Institute Dynamic Equilibrium Models workshop, CEPR-Bundesbank conference on Challenges in empirical macro since 2020, and Joint Research Center of the European Commission for their helpful comments and suggestions. We also thank Han Choi for his editorial comments.

Dynamic stochastic general equilibrium (DSGE) models have proven to be a valuable empirical framework for understanding aggregate economic dynamics. Since these models are estimated using historical data, they are suitable to study recurrent dynamics, such as the business cycle. Can these models remain useful in the face of the unusual COVID pandemic shock? We propose a methodology to incorporate unusual shocks into our usual models and use it to study the COVID recession and recovery within the context of a medium-scale New Keynesian business cycle model.

The COVID pandemic and the ensuing public and private responses to it certainly were unusual. The combination of pandemic, social distancing, government restrictions on economic activity, and fiscal policy in the second quarter of 2020 were a shock to the economy without precedent. This unusual shock lead to highly unusual macroeconomic dynamics. Figure 1 compares the COVID recession and recovery to all previous recessions and recoveries since 1947. The COVID recession is far deeper and the subsequent recovery is much faster than those of a typical business cycle. Furthermore, there was unusual co-movement. Output, consumption, investment, and hours were highly synchronized early on and consumption initially fell by more than investment and hours rather than by less.[1]

The pre-pandemic dynamics are embedded in contemporaneous private sector forecasts. Figure 2 shows forecast revisions for real GDP growth in the first quarter of a recession taken from the *Survey of Professional Forecasters* (SPF) starting from the beginning of the survey in 1968.[2] Before the pandemic, forecasters would revise down their forecasts for two quarters out. In contrast, forecasters surveyed in May 2022 correctly predicted an immediate start to a rapid recovery, with forecast revisions large and positive for three quarters out. Clearly, the decision by public authorities to close down whole swaths of the economy and mandate people to stay home led to the huge contraction in 2020q2. These containment measures were expected to last only a short time and hence forecasters expected a large rebound in real activity for 2020q3 and beyond.

[1] Eichenbaum, Rebelo, and Trabandt (2021) study the unique behavior of consumption in the COVID recession.

[2] The revision in period 0 is the difference between the SPF nowcast and the forecast from the previous period when the economy was at its business cycle peak.

Figure 1: Recessions and recoveries, 1947–2021

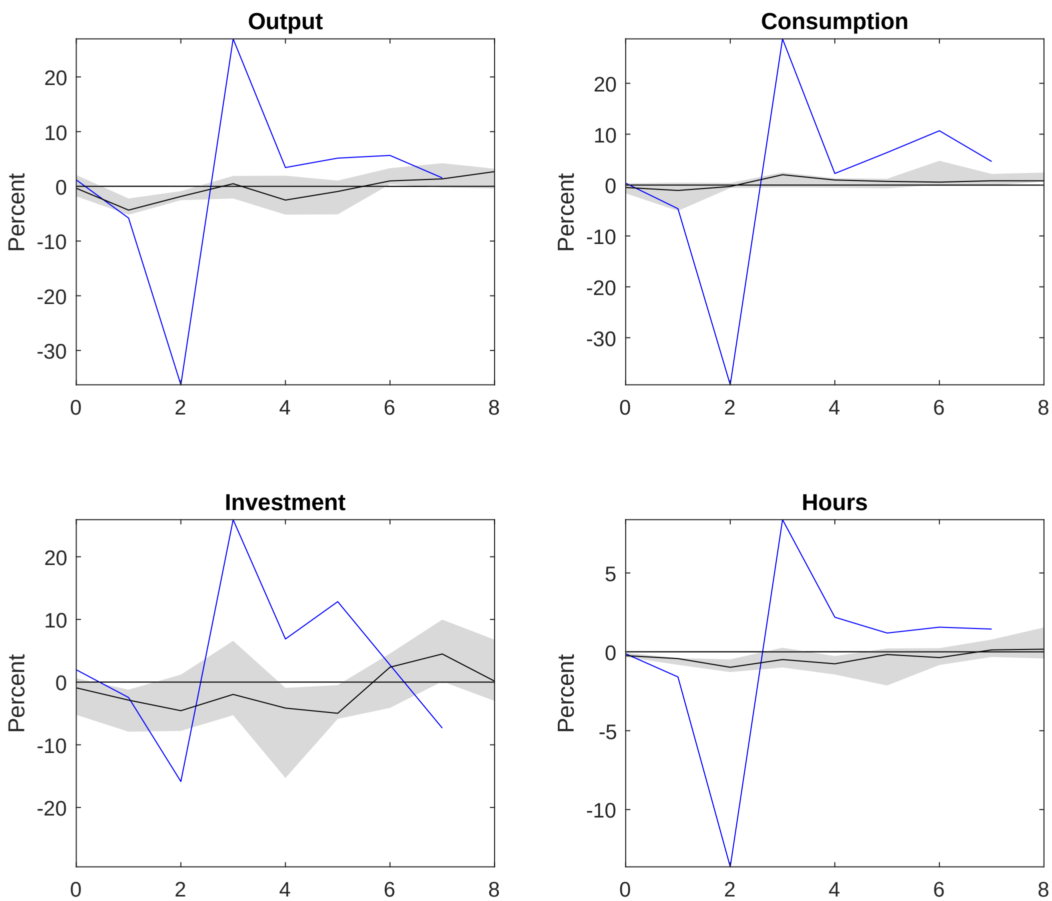

Note: The figure shows recessions and recoveries in the growth rates of per capita real GDP, consumption, investment and hours since 1947. Consumption includes non-durables and services consumption plus the service flow from the stock of consumer durables. Investment includes gross private domestic investment, expenditures on consumer durables, and government investment. GDP is the same as the BEA measure except it also includes the service flow from consumer durables. Details of our measurement are in E. Hours is for the non-farm business sector. The black line corresponds to medians and the shading is the 25 to 75 percentile range from before the Covid recession. The blue line corresponds growth rates of the indicated variable during the Covid recession and recovery. Source: Haver Analytics and author's calculations.

These COVID dynamics posed challenges to the viability of our usual models. We discuss this in the context of a standard DSGE model. Given the flexibility allowed by the shocks in this model, it can account for the magnitude of the fluctuations and the unique co-movement,

Figure 2: GDP growth forecast revisions at the onset of recessions, 1968–2021

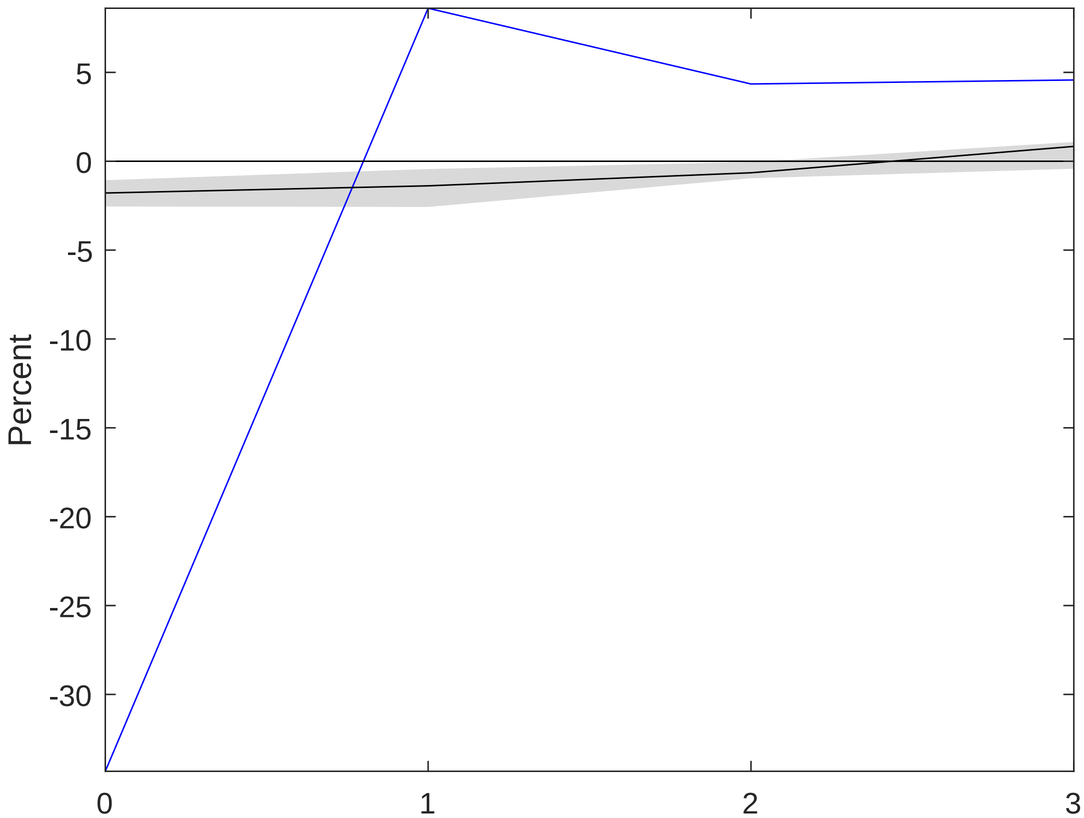

Note: Real GDP growth forecast revisions in the quarter of the onset of a recession compared to the quarter before taken from the *Survey of Professional Forecasters*, which begins in 1968. The horizontal axis indicates the horizon of the forecast. In period 0 the revision is the difference between the one-quarter-ahead growth forecast in the quarter before the recession starts and the nowcast in the quarter it starts. The black line is medians and shading is the 25 to 75 percentile range from before the Covid recession. The blue line corresponds to forecast revisions at the onset of the Covid recession. Source: Survey of Professional Forecasters and author's calculations.

but the shocks have to be very far outside their historical distributions. However, the model struggles to account for the forecast revisions regardless of the size of the shocks.

We contribute to addressing these challenges in two ways. First, we introduce a framework to estimate a new shock called the *Covid shock* to help capture the unusual macroeconomic dynamics of the COVID period. Including this single new shock goes a long way toward capturing the unusual dynamics while making the usual shocks more normal-sized. Our second contribution is to show how to incorporate the Covid shock into a more general framework that allows for the estimation of the structural parameters of business cycle models with data before, during, and after the COVID period.

We define the Covid shock using the wedges that underly the usual shocks that enter a DSGE model. We focus on the wedges underlying the shocks that drive typical business cycle co-movement as well as other shocks that play a major role in explaining consumption, investment, and inflation. The chosen wedges and corresponding loadings define the *nature* of the Covid shock. In each period, the model's agents update their beliefs about the path of the Covid shock as current surprises about the shock come with news about its evolution. The surprise and news structure of the Covid shock provides the flexibility not provided by the usual shocks to account for the dramatic fall in economic activity in 2020q2 and the forecasts of a sharp rebound and rapid recovery starting in 2020q3.

We adopt an event-study approach to identify the nature of the Covid shock. The macroeconomic data in 2020q2, including private sector forecasts, are clearly dominated by the actual and anticipated public and private sector responses to the pandemic and so should be particularly informative about the nature of the new shock. We acknowledge this by estimating the parameters defining the nature of the Covid shock with data from this quarter alone.[3] We use revisions to SPF forecasts of output growth and inflation to identify the anticipated component of the Covid shock over time.[4] Including news and holding fixed the nature of the Covid shock over time enables us to distinguish the macroeconomic effects of the Covid shock from those of the usual shocks.

Observing professional forecasts is particularly helpful when studying an unusual shock. Forecasters recognize that by its very nature the propagation of an unusual shock is not captured by past data. As such their forecasts will not rely as much on the historical

[3] The economic impact of the pandemic began to take hold in March 2020. This only shows up as a small contraction in activity in 2020q1.

[4] In standard DSGE models agents are rational and perfectly informed and so their expectations are inconsistent with the overreaction that has been found to characterize the SPF expectations, e.g. Bordalo, Gennaioli, Ma, and Shleifer (2020), Ansgar and Walther (2021), and Bianchi, Ilut, and Saijo (2023). We mitigate this issue by assuming that SPF expectations are observed with measurement error. Due to the relatively short sample, it is unclear whether SPF expectations were affected by over-reaction in the pandemic period we use to estimate the Covid shock.

dynamics and will incorporate any new information that they are absorbing in real time about how the shock will propagate through the economy. As forecasters update their beliefs about the propagation of the shock, this will get reflected in their forecast revisions.[5]

We use our model to isolate the effects of the Covid shock on aggregate activity and inflation over the period 2020q2 – 2021q3. The Covid shock explains about two-thirds of the massive decline in hours in 2020q2 and contributes considerably to the extraordinary economic rebound in the next quarter. Over the following four quarters, the shock is a significant drag on economic activity even as the usual business cycle dynamics take hold. The inflationary effects of the Covid shock are more muted, mostly because of the model's flat Phillips curve. Nonetheless, the Covid shock is inflationary throughout the sample period, significantly so in 2020q2.

An advantage of our methodology is the ability it provides to quantify the role of beliefs about the propagation of the Covid shock. This is due to the shock's surprise and news structure and the observation of revisions to professional forecasts. We find that beliefs about the future path of the Covid shock reduced the magnitude of the 2020q2 contraction in GDP by 10 percentage points, but were a drag on activity for the remainder of 2020. We also find that the economic effects of the Covid shock were mostly hard to anticipate. In other words, professional forecasters were continually surprised by the actual and anticipated effects of the shock.

The news structure of the novel shock is also helpful when unusual events repeat themselves. In the case of COVID, we observed recurrent waves of infections as well as the emergence of new variants. A prime example is the Delta wave that began in 2021q3. The Delta wave was not accompanied by the same fiscal interventions and new restrictions on economic activity that characterized the initial outbreak. However, to some extent the effects could be anticipated at the time to be similar to the initial outbreak. For example, it was reasonable to expect voluntary social-distancing to increase. As such agents could use their

[5]Professional forecasts are also valuable when the effects of an unusual shock are studied in real time. Given the large uncertainty and the scarcity of data that characterize the initial periods of an unusual episode, observing such data allows for the best possible real-time estimation of the effects of the unusual events.

knowledge about the propagation of the original Covid shock to anticipate how the onset of the Delta wave would propagate through the economy. We illustrate how our framework can take this into account.

1. Related literature

Primiceri and Tambalotti (2020) identify a surprise Covid shock in a monthly vector autoregressive model (VAR). Their shock is defined as a linear combination of the VAR's reduced form shocks with weights estimated using data from March and April 2020, when aggregate dynamics clearly were dominated by the Covid shock. We differ in two respects: First, our shock is a linear combination of wedges in a structural model, and second, we include forward-looking information to identify news shocks that come with the surprise. Including the news shocks turns out to be crucial to our estimation of the Covid shock.

Lenza and Primiceri (2022) model Covid in a VAR by scaling the variances of the usual independently and identically distributed (i.i.d.) residuals by a common scaling parameter that decays exponentially over time. The scaling parameter and its rate of decay are estimated using data from March, April, and May 2020. They use their framework to demonstrate that one obtains similar VAR parameter estimates by dropping those observations. This will be helpful going forward to estimate VARs with data that includes the pandemic period. We provide a way to estimate structural models with these data. Note that Lenza and Primiceri (2022) do not exploit the information in private sector forecasts. Our structure allows us to measure the sensitivity of private sector decisions to expectations about the future effects of the pandemic that are revealed through news.

In a short note on inflation based on the NY Fed's DSGE model Del Negro, Gleich, Goyal, Johnson, and Tambalotti (2022) describe how they account for the pandemic in that model. Their approach also involves introducing shocks to some of the model's existing wedges, but it differs in several respects: they do not exploit the parsimony of a factor structure, they do not use data on professional forecasts to identify the new shock, and they calibrate instead of estimate key parameters, including scaling down the volatility of the usual shocks when they estimate their i.i.d. wedge shocks' variances using data in 2020q2. Cardani, Croitorov, Giovannini, Pfeiffer, Ratto, and Vogel (2022) also introduce a novel Covid shock into an off-the-shelf DSGE model, in their case a model of the Euro area, but their shock is based on modifying the model's structure to include forced savings and labor hoarding rather than leveraging the model's pre-existing wedges.

Our methodology addresses model misspecification during the COVID period. There are a variety of ways to address misspecification that have been proposed in the literature. Ireland (2004) uses measurement error that does not appear anywhere in the structural model. We discuss how adding measurement error to address the COVID period can lead to implausible findings. Similar to our methodology, den Haan and Drechsel (2021) and Inoue, Kuo, and Rossi (2020) use shocks to wedges to address misspecification. The former includes shocks in every structural equation while the latter is more similar to our approach as it adds shocks to a subset of wedges. These papers focus on detecting and correcting potential sources of misspecification of models estimated over long sample periods while we are focused on a short time period. A critical challenge we face is precisely the fact that information about the misspecification to correct is concentrated in a short period of time, requiring a more parsimonious approach. We do not have enough data to apply the den Haan and Drechsel (2021) and Inoue et al. (2020) methods for such a short period. Another difference with these papers is that we do not estimate the usual structural parameters simultaneously with the misspecification parameters. However, one could straightforwardly extend our methods to incorporate the approaches developed in these papers to estimate the structural parameters and we discuss this below.

The Covid shock loads significantly on wedges that generate supply and demand effects. Guerrieri, Lorenzoni, Straub, and Werning (2022) show how a supply shock can cause demand shortages in a two-sector New Keynesian model. In our setting the endogenous effects of the supply shock on demand they describe would be captured by the loadings of the common factor on the model's wedges. The wedges also can be viewed as a reduced-form characterization of the interaction of demand and supply shocks in the New Keynesian model with input-output linkages studied by Baqaee and Farhi (2022).

Our analysis is complementary to the large literature that embeds epidemiological models within otherwise standard business cycle models to study the COVID pandemic, for example Eichenbaum et al. (2021) and Acemoglu, Chernozhukov, Werning, and Whinston (2021). These "epi-mac" models yield important new insights but add considerable complexity. Our framework does not involve changing our usual models, but leverages their existing structure to synthesize a new shock that can in principle capture the dynamics resulting from the pandemic and related developments. By basing our analysis on a standard DSGE model we can assess the empirical relevance of the new shock relative to the usual shocks that have proved to be useful in accounting for U.S. business cycles.

Our approach addresses the absence of a major pandemic in the postwar data before COVID that could be used to identify some of the effects of COVID. Alternatively one could use additional time-series data to learn from history about the possible effects of the Covid shock on the economy. The only comparable major pandemic was the Spanish influenza of 1918 and 1919. Barro, Ursua, and Weng (2020), Barro (2020), and Velde (2020) use this episode to shed light on the economic effects of a pandemic. Ludvigson, Ma, and Ng (2020) project the economic impact of COVID based on estimates of the impact of deadly disasters in recent U.S. history.

2. The Covid shock and its estimation

This section describes the new Covid shock and how we estimate it. First, we describe the shock and introduce enough structure to separately identify the shock and news about its evolution. The incorporation of news allows agents' beliefs about the evolution of the Covid shock to vary over time. We then show how to incorporate the shock into a general linearized DSGE model. Next, we describe how we estimate the shock taken the usual parameters as given, which is what we do in this paper. We finish by describing how estimation of the Covid shock can be integrated into a general procedure where the usual parameters are estimated with data that sandwiches the period the Covid shock is active.

2.1. Definition of the shock

We introduce an exogenous scalar i.i.d. random variable

$$f_t \sim \mathcal{N}(0, \sigma^2(t)), \tag{1}$$

with time-varying volatility (known to agents in the model), whose realization captures the time-varying quantity of *news* about the Covid shock that arrives in the economy each period. We assume that $\sigma(t) = 0$ for $t < t^*$ and $t > T$, where t^* is the period when the Covid shock first impacts the economy and T is the last period agents receive news about the Covid shock.

For a given quantity of news, f_t, we need to specify how that news is distributed over time, i.e. whether that news is about current or future values of the Covid shock, and which wedges in the economy the Covid shock will influence. The news about the Covid shock consists of surprise and anticipated components. The time t realization of the Covid shock — the impact of the Covid shock on the values of contemporaneous wedges — is given by the scalar

$$\Psi_t = \sum_{j=0}^{N} \psi_{t-j}^{j}, \;\; N \geq 0, \tag{2}$$

where the random variables ψ_{t-j}^{j} measures the current time t impact of the Covid shock that is anticipated j periods earlier and N denotes the anticipation horizon. Notice that the random variables ψ_t^j equal date t revisions to expectations about the evolution of the Covid shock Ψ_{t+j}. Specifically, from (2) we have

$$\psi_t^j = E_t \Psi_{t+j} - E_{t-1} \Psi_{t+j} \;\, j \in \{0, 1, \ldots, N\}. \tag{3}$$

We allow for time-variation in the degree to which agents anticipate Covid effects, to capture the idea that agents may be learning each period about how the effects of the Covid shock will play out. Hence, we assume that

$$\psi_t^j \;=\; \lambda_j(t) f_t,\; j \in \{0, 1, \ldots, N\}, \tag{4}$$

where the parameters $\lambda_j(t)$ are fully known to agents and allowed to vary both by time the information arrives, t, and by the horizon of anticipation, j. We normalize the loading onto the surprise component of the Covid shock to one, i.e. $\lambda_0 = 1$. Given the structure of the Covid shock summarized by equations (2) and (4) we can write it as

$$\Psi_t = \sum_{j=0}^{N} \lambda_j(t-j) f_{t-j}.$$

It follows that the Covid shock Ψ_t is time-varying serially correlated as it depends on current and past realizations of f_t and the lag structure is time-varying.

We assume a time-invariant mapping between the scalar Covid shock, Ψ_t, and M wedges in the model economy, $\Upsilon_t(i),\; i \in \{1, 2, \ldots, M\}$. These wedges enter into the usual DSGE model identically to shocks already present. To be concrete, suppose the model is a production economy and i refers to total factor productivity, $\exp(A_t)$, in a production function that is linear in hours worked, h_t. The usual shock would be A_t. The new wedge $\Upsilon_t(i)$ would enter the model as $\exp(A_t)\exp(\Upsilon_t(i))h_t$. The additional component of each wedge $\Upsilon_t(i)$ is determined according to

$$\Upsilon_t(i) = \phi_i \Psi_t,\; \in \{1, 2, \ldots, M\}. \tag{5}$$

The scalar parameters ϕ_i known to agents are the loadings of the Covid shock Ψ_t onto the wedges. We refer to the choice of wedges and the loadings as the *nature* of the Covid shock. Note that the loadings $\phi = \{\phi_i\}_{i=1}^M$ do not depend on the anticipation horizon of the wedges so that the combination of the DSGE wedges does not vary across anticipation horizons. We think this assumption is natural but it also provides parsimony.[6]

[6]This assumption also makes it possible to identify the Covid news separately from news of the usual shocks if it is already present in the DSGE model, for example as in Schmitt-Grohé and Uribe (2012), provided that news of each usual shock is not perfectly correlated as the Covid wedges are in our framework.

To sum up, we capture the dynamics of the Covid shock with the loadings ϕ and $\lambda(t)$ and the common factor f_t. The vector ϕ describes the nature of the Covid shock, defined as a particular combination of wedges that enter into the DSGE model in the same way as a subset of the usual shocks. The loadings $\lambda(t)$ capture evolving news about the Covid shock. The variance $\sigma(t)$ measures the quantity of news in any given period. Individual realizations of the exogenous variable f_t account for revisions to agents' expectations of the future path of the Covid shock.

2.2. Introducing the Covid shock into a usual DSGE model

Consider a general linearized DSGE model of the form

$$Ay_{t-1} + By_t + CE_t y_{t+1} + Dx_t = 0, \tag{6}$$

where y_t is a $K \times 1$ state vector of endogenous variables (e.g. consumption) and exogenous shocks (e.g. total factor productivity), and x_t is a $P \times 1$ vector of i.i.d. innovations to the exogenous shocks with $P \geq M$. The matrices A, B, C and D are conformable with y_t and x_t and are composed of the coefficients of the linearized equations that characterize the equilibrium of the model.

The onset of the new shock process is an unanticipated event. Prior to t^*, the model's dynamics are characterized by equation (6). At date $t = t^*$, the model changes by appending y_t with the Covid surprise and news components and x_t with the new source of randomness f_t, and adjusting A, B, C and D accordingly. For simplicity, assume that $N = M = 1$ so that a single wedge defines the Covid shock and there is one period of anticipation. Define X_m as column m of matrix X. Now we can write the DSGE model augmented with the Covid shock as

$$\begin{pmatrix} A & \mathbf{0} & \phi A_m \\ \mathbf{0}' & 0 & 0 \\ \mathbf{0}' & 0 & 0 \end{pmatrix} \begin{pmatrix} y_{t-1} \\ \psi^0_{t-1} \\ \psi^1_{t-1} \end{pmatrix} + \begin{pmatrix} B & \phi B_m & \mathbf{0} \\ \mathbf{0}' & 1 & 0 \\ \mathbf{0}' & 0 & 1 \end{pmatrix} \begin{pmatrix} y_t \\ \psi^0_t \\ \psi^1_t \end{pmatrix}$$

$$+\begin{pmatrix} C & \mathbf{0} & \mathbf{0} \\ \mathbf{0}' & 0 & 0 \\ \mathbf{0}' & 0 & 0 \end{pmatrix} E_t \begin{pmatrix} y_{t+1} \\ \psi^0_{t+1} \\ \psi^1_{t+1} \end{pmatrix} + \begin{pmatrix} D & \mathbf{0} & \mathbf{0} \\ \mathbf{0}' & -1 & 0 \\ \mathbf{0}' & 0 & -1 \end{pmatrix} \begin{pmatrix} x_t \\ f_t \\ \lambda_1(t) f_t \end{pmatrix} = 0,$$

where $\mathbf{0}$ is a $K \times 1$ vector of zeros. In this example, columns m of A and B contain the coefficients of the model's linearized equations that multiply the usual shock associated with the wedge the Covid shock loads onto and that appears in row m of the state vector y_t.

This representation of the Covid-augmented model highlights two key points. First, the matrices that pre-multiply the appended state and innovation vectors include the A, B, C, and D matrices from the original model. This means that the parameters of the original model do not change with the introduction of the new shock process. Second, the fact that A and B are appended to include elements from these matrices that pre-multiply the usual shock reflects our assumption that the wedges that define the Covid shock enter the model in the same way as the associated usual shock.

2.3. Estimation taking the usual parameters as given

To measure the Covid shock Ψ_t we need to estimate ϕ and $\Xi(t) = [\lambda(t), \sigma(t)]$ for the number periods we determine that agents update their beliefs about the Covid shock. We apply an event-study approach to identify ϕ. In 2020q2 there was an unusually large drop in economic activity — far beyond the bounds of a typical business cycle peak to trough — and an unusual expected rebound. We assume the dramatic variation in 2020q2 is due chiefly to the Covid shock. Therefore 2020q2 data on current and expected future activity and inflation will be particularly informative about ϕ and this guides our estimation strategy. We assume 2020q2 is the first date of the new Covid shock process, i.e. $t^* = 2020q2$.

Let Θ denote the usual parameters in our DSGE model, which are taken as given. Note that this includes the volatilities of the usual shocks. This is important because it means our estimation in effect lets the data speak about the relative volatility of the Covid shock. We use Bayes' theorem to obtain a distribution of $\Xi(t)$ and ϕ conditional on the usual data up to date t, denoted X^t. At date $t = t^*$ we have,

$$p\left(\Xi(t), \phi | X^t, \Theta, s_{t-1}\right) \propto \mathcal{L}\left(X^t | \Xi(t), \phi, \Theta, s_{t-1}\right) p\left(\Xi(t), \phi\right), \tag{7}$$

where s_{t-1} is the model's state vector estimated one quarter earlier. The density $p(\cdot)$ is our prior on the new parameters capturing the nature of and beliefs about the Covid shock. The density $\mathcal{L}(\cdot)$ is the likelihood function associated with the data X^t. With ϕ estimated with the data at date $t = t^*$, for $t > t^*$ we have

$$p\left(\Xi(t^* + j) | X^{t^*+j}, \phi, \Theta, s_{t^*+j-1}\right) \propto \mathcal{L}\left(X^{t^*+j} | \Xi(t^* + j), \phi, \Theta, s_{t^*+j-1}\right) p\left(\Xi(t^* + j)\right), \tag{8}$$

for $j = 1, 2, \ldots, T - t^*$.

We estimate ϕ and $\Xi(t)$ sequentially by maximizing the posterior modes in (7) and (8) for a given T. For $t = t^*$, the intuition is to find the combination of the wedges $\Upsilon_t(i)$ that, along with the usual shocks, best explain the one-step-ahead forecast error of the usual data that includes current activity and professional forecasts. For $t > t^*$ (and $t \leq T$) the $\Xi(t)$ are identified by the revisions to the professional forecasts of output and inflation. The last date T is the period before the first time the marginal liklihoods of the Covid parameters are too flat to identify $\Xi(t)$ reliably.

We use the Kalman smoother to estimate f_t. With f_t and our estimates of ϕ and $\lambda(t)$ we obtain estimates of the Covid shock and its anticipated and unanticipated components from (2) and (4).

2.4. Estimation with data that extends beyond the COVID period

We now describe how to integrate our framework into a more general setup to estimate DSGE models with datasets that include data from before, during, and after the COVID pandemic or any other unusual event. The method begins with dividing the data into three sub-samples: the pre-COVID sample when $f_t = 0$ (Y^1), the COVID sample during which f_t is active (Y^2), and the post-COVID sample when f_t is set to zero again (Y^3). The general estimation proceeds in three steps.

In the first step, the model's usual parameters, Θ, are estimated with data from the pre-COVID sample using standard Bayesian methods. The posterior kernel can be expressed as:

$$\mathcal{P}(\Theta|Y^1) \propto \mathcal{L}(Y^1|\Theta)p(\Theta),$$

where the first density on the right hand side is the likelihood followed by the prior distribution. We obtain an estimate of Θ by maximizing the mode of the posterior kernel.

The second step involves following the method described in Section 2.3 to estimate the parameters characterizing the Covid shock, summarized here by Φ, holding Θ fixed from the first step and using data from the first two sub-samples, Y^1 and Y^2. This estimation yields the state vector s_T from the last period of news about the Covid shock arrives.

In the third step, we set $f_t = 0$ and update our estimate of Θ to account for the information contained in the post-COVID sample. To do this, we use the posterior for Θ from the first step as the prior and use the state vector from the last period of the second step as the initial condition. Formally,

$$\mathcal{P}(\Theta|\Phi, Y^1, Y^3, s_T) \propto \mathcal{L}(Y^3|\Theta, \Phi, s_T)\mathcal{P}(\Theta|Y^1)$$

where the last term on the left hand side is the posterior kernal for the parameters Θ conditional on the data observed in the pre-COVID and post-COVID samples. The likelihood and posterior on the right hand side depend on the Covid parameters Φ as these contain news about future values of the Covid shock's wedges that are realized in the first several periods of the post-COVID sample. The data Y^2 are not used to estimate Θ except through their influence on the initial state, s_T, and the residual impact of the news received during the COVID period. They are excluded from estimating Θ precisely because the COVID period is unusual. We obtain an updated estimate of Θ by maximizing the posterior mode.

The first and third steps do not rely on a particular approach to structural estimation. One could apply the traditional methods that we use to estimate our model with pre-pandemic data. This approach takes the model to be correctly specified outside of the COVID period. However, it would be straightforward to extend our methodology to use the strategies developed by den Haan and Drechsel (2021) and Inoue et al. (2020) to address model misspecification in normal times.

This generalization of our procedure to estimate the Covid shock not only helps to estimate the parameters using pre- and post-COVID observations, but it also allows us to obtain a more compelling estimate of the model's states (e.g. the output gap, which is critical for appropriate monetary policy.) Later on, we will show that estimation of the model's states can go awry if a set of measurement errors instead are used to estimate the model during the COVID period.

3. The DSGE Model

We study the Covid shock within Campbell, Fisher, Justiniano, and Melosi (2016)'s medium-scale model. The model is familiar as it is a variant of Christiano, Eichenbaum, and Evans (2005) and Smets and Wouters (2007). It is closest Justiniano, Primiceri, and Tambalotti (2013)'s model. The main differences are the inclusion of a preference for real government bonds, investment-specific technological change, and anticipated deviations from the monetary policy rule. The details of the model are provided in A – D.

The model includes familiar shocks to various wedges in the model. Since the wedges are an important input into the Covid shock, we summarize the shocks here. They include shocks to the discount factor, government spending, the marginal efficiency of investment (MEI), liquidity preference, wage and price markups, neutral technology, and investment-specific technology.[7] The model also includes the usual surprise deviation from the monetary policy rule as well as news future deviations up to 10 quarters ahead for a total of 19 structural shocks. Our method adds just one additional shock to this list.

[7]The liquidity preference shock is a shock to the preference for government bonds. It enters the representative household's consumption Euler equation similarly to the "risk premium" shock in Smets and Wouters (2007) as shown by Fisher (2015). Such preferences are now common in the literature, for example Michaillat and Saez (2021), Eichenbaum, Johannsen, and Rebelo (2021), and Anzoategui, Comin, Gertler, and Martinez (2019).

4. Estimation of the DSGE model with the Covid shock

We estimate the model's usual structural parameters using Bayesian methods with data from before the pandemic.[8] We use data starting from the onset of the pandemic to estimate the nature and propagation of the Covid shock as described in Section 2.3.

4.1. Pre-pandemic period

Our pre-pandemic estimation follows Campbell, Ferroni, Fisher, and Melosi (2019) and is described in F. The model is estimated with 32 time series including up to 10 quarters of interest rate futures data to identify the anticipated deviations from the monetary policy rule.[9] The nonfinancial data include per capita GDP, consumption and investment growth, per capita hours, multiple wage and price inflation series, and various forecasts from the SPF. The forecasts are for one- to four-quarter-ahead inflation and GDP growth and ten-year-ahead-average inflation. We allow for measurement error in the SPF forecasts to address that agents in our model are rational, but evidence suggests that SPF forecasts deviate from rationality.[10] GDP, investment, and consumption are measured in a model-consistent way. Consumption includes nondurables and services consumption and the service flow from the stock of consumer durables; investment includes gross private domestic investment and durables consumption; and GDP is the BEA measure except that it includes the services flow from consumer durables. More detail about the construction of these measures is provided in E. The sample period for the pre-pandemic estimation is 1993q1–2016q4. The starting date is determined by the availability of the interest rate futures data.

[8] A general overview of Bayesian estimation is provided in Herbst and Schorfheide (2015) and Fernandez-Villaverde, Rubio-Ramirez, and Schorfheide (2016)

[9] The interest rate futures data is from the Chicago Fed. Unless otherwise noted all other source data are from Haver Analytics.

[10] See footnote 4 for references. We also include two auxiliary inflation measures (which do not enter the DSGE model) to map the model concepts of output and inflation to the their empirical counterparts. See E for the details.

4.2. The first wave of the pandemic

We assume the Covid shock is formed from the liquidity preference, neutral technology, MEI, discount factor, and price markup wedges ($M = 5$). Shocks to the first two wedges are major sources of co-movement in the model. Shocks to the MEI and discount factor wedges are important determinants of consumption and investment, but drive them in opposite directions. We include the price markup wedge to address the supply disruptions that followed the onset of the pandemic. We assume that agents try to anticipate the effects of the Covid shock up to four quarters ahead ($N = 4$). Note that the time horizon of the SPF forecasts used in our estimation exactly matches the horizon of the anticipated Covid shocks in the model. As we shall explain, observing these forecasts is key to identifying the evolution of the Covid shock.

Recall that the first date of the Covid shock process is assumed to be $t^* = 2020\text{q}2$. Using the Kalman filter, data prior to t^*, and our pre-pandemic parameter estimates, we obtain the state vector in $t^* - 1$. We then follow the strategy outlined in Section 2.3 starting in 2020q2 to obtain estimates of the Covid parameters. Since we find that $\lambda(t)$ is poorly identified in 2021q1, we estimate $T = 2020\text{q}4$.[11]

The fact that $f_t = 0$ after date T *does not* mean that the Covid shock has run its course. In particular, news about the Covid shock received by agents in 2020 appears in Ψ_t after date T because agents receive news about the Covid shock up to four quarters ahead. Therefore setting $f_t = 0$ after date T does not nullify the Covid shock over that period. Indeed, we show below that the Covid shock was a persistent drag on activity in 2021, although some of these effects are due to the endogenous propagation of past values of the Covid shock.

Our identification of $\lambda(t)$ relies on forecast revisions. DSGE (and in general autoregressive) model forecasts are strongly influenced by historical data. Since the propagation of

[11] By poorly identified we mean the marginal likelihoods of the parameters become very flat. In 2021q1 there was a new and large fiscal package. We tried estimating a new unusual shock in this quarter, but were unable to achieve identification because the macro data and expectations behave quite normally in that quarter. A model with a more substantial fiscal sector might make it possible to continue to identify the original Covid shock after 2020q4 by exploiting the record increase in govt transfers induced by the 2021 fiscal stimulus.

the Covid shock does not resemble the typical business cycle dynamics, we want to avoid the model relying too heavily on past dynamics in predicting the likely course of economy during the pandemic. One way to accomplish that is to give more importance to forecasts that are external to the model. These forecasts factor in the unusual propagation of the Covid shock that is not present in the pre-pandemic data. We give more importance to the external forecasts by reducing the standard deviations of the measurement error shocks on the SFP expectations by a factor of ten in 2020q2, 2020q3, and 2020q4.

4.3. The Delta wave

The Delta wave that began in 2021q3 was not accompanied by the fiscal interventions and new restrictions on economic activity that characterized the initial COVID outbreak. We attempted to estimate a new unusual shock in this quarter, but it was poorly identified. However, to some extent the effects could be anticipated at the time to be similar to the initial outbreak. For example, it was reasonable to expect voluntary social-distancing to increase. As such agents could use their knowledge about the propagation of the original Covid shock to anticipate how the onset of the Delta wave would propagate through the economy. Our framework allows us to take this into account.

Specifically, once the nature of the Covid shock is estimated, its surprise and news structure can be used to make different assumptions regarding how much agents have learned from previous experience. We illustrate how this feature can be applied by considering the hypothetical scenario in which agents assess the Delta wave starting in 2021q3 to be similar to the original Covid shock and that they have perfect ex-ante knowledge about the propagation of the shock as given by the estimates from the initial wave.[12]

To implement this scenario, we assume that when the second wave hits in period t^{**}, the following Covid news is realized: $\psi^{j}_{t^{**}} = \delta \cdot \Psi_{t^{*}+j}$ for $j \in \{0, 1, ..., T - t^{*}\}$, where $\Psi_{t^{*}+j}$ denote the value of the Covid shock j periods after the start of the first wave in period t^{*}. The parameter δ is a scaling factor that determines whether the second wave is more severe

[12] Our methodology is flexible enough to address alternative scenarios about the flow of information that relax the perfect foresight assumption.

($\delta > 1$) or less severe ($\delta < 1$) than the first wave. In the subsequent periods ($t^{**} + 1, t^{**} + 2, \ldots, t^{**} + T - t^* + 1$), there will be neither surprise nor news since all the effects of the second wave are assumed to be correctly anticipated from the start and hence there are no revisions to agents' expectations after t^{**}. In symbols, $\psi^j_{t^{**}+i} = 0$ for any $i \in \{1, 2, \ldots, T - t^*\}$ and $j \in \{1, 2, \ldots, N\}$.

5. The estimated effects of the Covid shock

Our estimates of the Covid parameters are displayed and discussed in G. In this section, we study the estimated Covid shock which depends on those parameters. First, we study the contributions of the unusual and usual shocks to the one-quarter-ahead forecast errors and forecast revisions of output and inflation in 2020q2 and the importance of including the Covid shock to explain the dynamics in Figure 2. Next, we examine the effects of the unusual and usual shocks on aggregate activity, inflation, and co-movement over the pandemic period 2020q2 to 2021q3. After this we discuss our finding that, perhaps surprisingly, the usual shocks play a relatively large role in explaining the aggregate dynamics. Lastly, we study the role of beliefs in the propagation of the Covid shock.

5.1. Forecast errors and revisions in 2020q2 with the Covid shock

The colored bars in Figure 3 show the decomposition of the forecast errors and revisions in 2020q2 into contributions from the Covid shock, including both surprise and news (red), the usual business cycle shocks (blue), surprise and news shocks to monetary policy (green), and other (gray).[13] The forecasts displayed are very close to the SPF forecasts and so Figure 3 shows that the Covid shock explains almost all of the sharp reversal in GDP growth forecasted in 2020q2 that we highlighted in Figure 2. Note that we did not scale down the volatility of the usual shocks from their pre-pandemic estimated values and so we let the data speak about the relative contribution of the Covid shock.

[13] Other includes shocks to variables from outside the model used to measure model-consistent GDP and core PCE inflation and classical measurement error on the inflation series in our dataset.

Figure 3: Decomposition of one-step-ahead forecast error and revisions in 2020q2

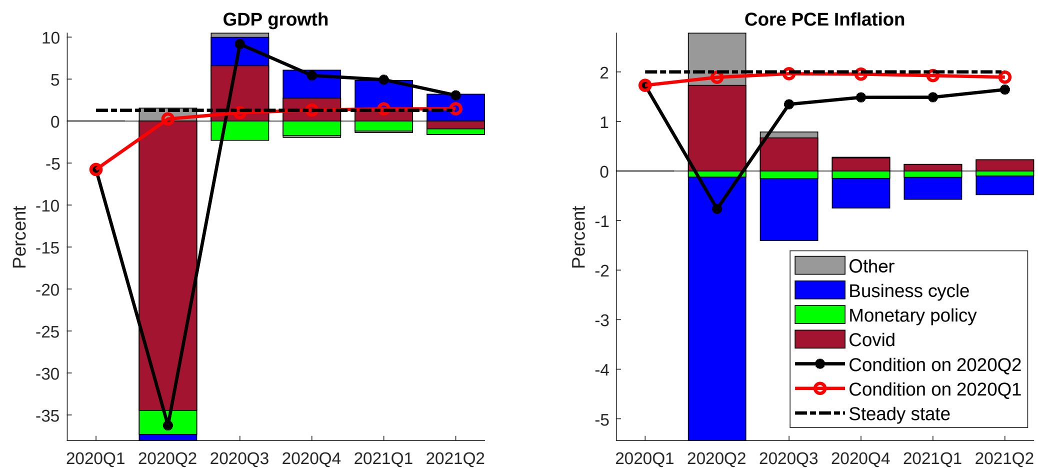

Note: Decomposition of one-step-ahead forecast error and revisions of output and inflation into the parts attributed to the Covid shock, usual business cycle shocks, surprise and news monetary policy shocks, and other shocks which include variables from outside the model used to measure model-consistent GDP and core PCE inflation. The black lines indicate 2020q2 output and inflation and SPF forecasts of these variables one to four quarters ahead. The red line is the forecast conditioned on 2020q1 data. The colored bars show the contribution to the forecast error and revisions of the indicated shocks. Source: Author's calculations, SPF, and Haver Analytics.

The right plot of Figure 3 shows Covid pushed prices higher in 2020q2 and was expected to put upward pressure on prices through 2021q2. The model attributes the decline in current and expected inflation to the usual business cycle shocks. The fall in output and rise in prices attributed to the Covid shock in 2020q2 indicate that the model interprets the initial shock, on net, as a supply shock. The accumulated effect of the Covid shock on the levels of output and prices indicate that in 2020q2 agents expected its relatively strong supply effects to persist.

Why does the likelihood function attribute such a large role to the Covid shock? Figure 4 provides some insight into this question. The red stars denote SPF forecasts in 2020q2 for GDP growth over the next four quarters. The black line shows the forecast based on inferring the usual shocks in 2020q2 from our model without the Covid shock. The difference between them is measurement error. This error is very large. For example, the measurement error attached to the forecast of growth in 2020q3 is 100 times larger than its pre-pandemic estimated standard deviation. Evidently, our medium-scale DSGE model with only its usual set of shocks struggles to account for the unusual nature of the anticipated recovery.

Figure 4: Forecast in 2020q2 of GDP growth from the SPF and the DSGE model that excludes the Covid shock

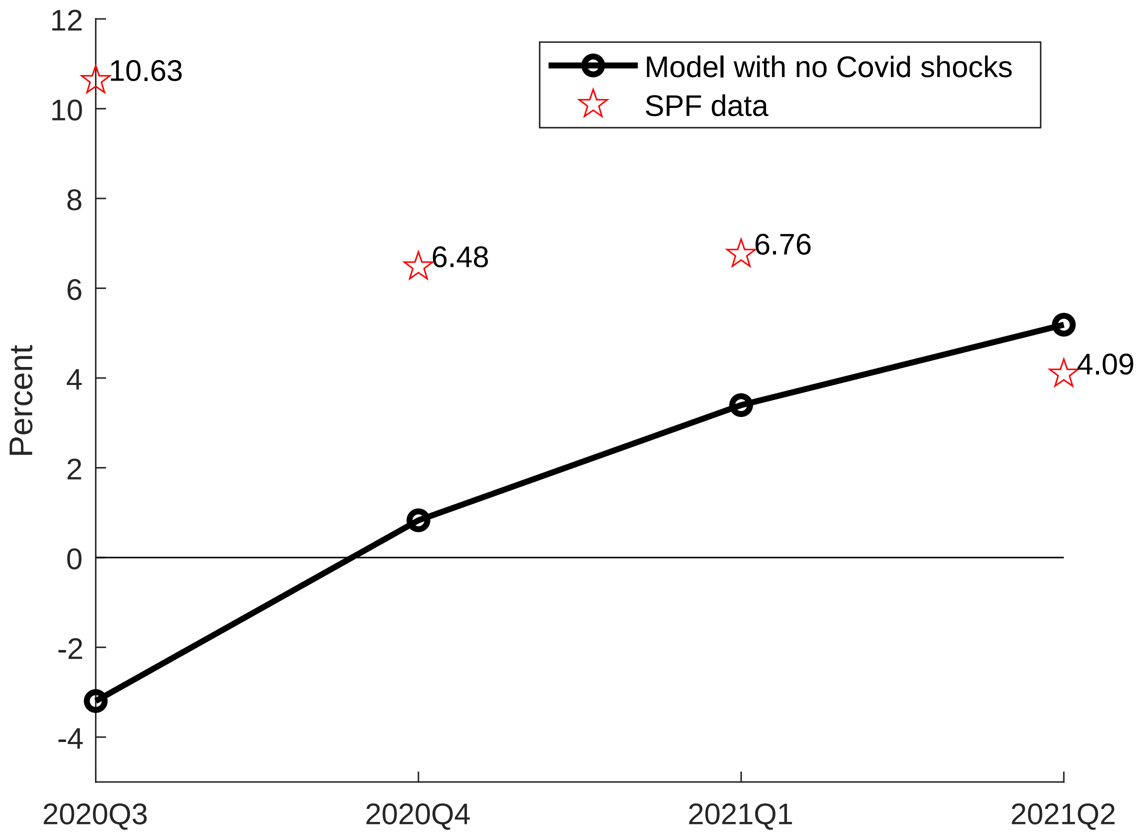

Note: Red stars indicate the median SFP forecast in 2020q2 of GDP growth in 2020q3 to 2021q2. The black line shows the DSGE model's forecast without the Covid shock but with otherwise identical parameters. The difference between the black line and the red stars is due to measurement error identified by the Kalman filter. Source: Survey of Professional Forecasters and authors' calculations.

5.2. The effects of the shocks from 2020q2 through 2021q3

We now study the Covid shock's contributions to aggregate outcomes alongside the usual shocks over the period 2020q2 to 2021q3. We measure the shocks with the Kalman smoother (the results so far are based on the Kalman filter). For ease of interpretation we group the model's usual shocks into five categories: demand, transitory supply, persistent supply, monetary policy, and other.[14] The composition of each category is summarized in Table 1.

Table 1: Categories of usual shocks

Category	Usual shock
Demand	Liquidity preference + Discount factor
Transitory supply	Wage and price cost push
Persistent supply	Neutral technology + IS technology + MEI
Monetary policy	Unanticipated and anticipated
Other	Residual (government) spending + measurement

Note: IS denotes investment-specific. Residual spending includes net exports, inventory investment, and government spending. Measurement includes measurement error in core PCE, as well as shocks to consumer durable inflation and inflation in the consumption price of residual output. The latter two variables are used to measure model-consistent GDP and core PCE inflation and do not enter the model as structural shocks.

We focus on the contributions of the shocks to log per capita hours worked and core PCE inflation.[15] Our measure of hours is a good indicator of the cyclical position of the US economy. The contributions are displayed in Figure 5 which is constructed similarly to Figure 3.

The left-hand plot in Figure 5 shows the sharp contraction and fast recovery of the labor market. The Covid shock is the largest factor contributing to the sharp downturn in 2020q2. This shock is also largely responsible for the initial recovery in 2020q3. If not for demand shocks, hours would have been 10 percentage points higher in that quarter. The Covid shock is a persistent drag on hours, significantly so in 2021. This seems consistent with the impact on labor supply often attributed to the pandemic (for example, the lower supply due to fear of the virus and the need to stay home to care for children who would otherwise be in school

[14] A demand shock is a shocks that drives output and prices in opposite directions while a supply shock is a shock that drives output and prices in the same direction.

[15] Our empirical measure of hours is de-trended from outside the model using estimates of secular trends in labor force participation and average hours per worker, as well as estimates of the natural rate of unemployment. Our measurement follows Campbell et al. (2016) and is described in E.

Figure 5: The estimated effects of the Covid shock and the usual shock, 2020q2–2021q3

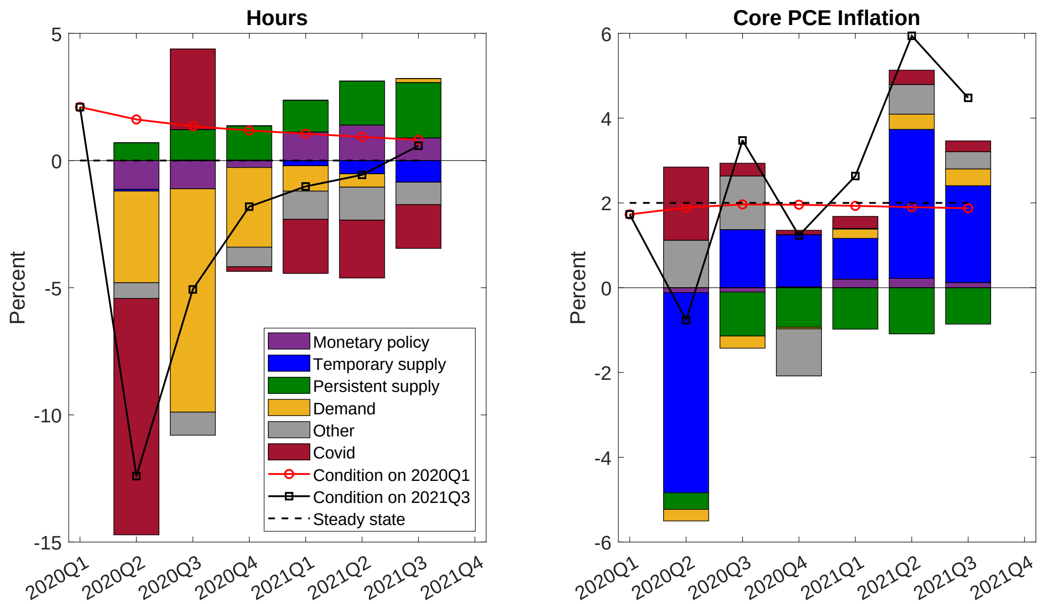

Note: Kalman smoothed decomposition of the contributions of the Covid shock and the model's usual shocks over the period 2020q2 to 2021q3. The black lines are data, and the red line is the forecast as of 2020q1. The colored bars are the contributions of the model's shocks that sum to the difference between the red and black lines. Log per capita hours is the deviation from steady state. Steady-state inflation equals two. Source: Haver Analytics, Board of Governors, Chicago Fed, and authors' calculations.

or daycare). The impact of the Covid shock used to capture the effects of the Delta wave is very small (we estimate $\delta = .03$). While overall the contributions of the Covid shock are substantial, the usual shocks play an important role as well. Demand shocks are a large drag on the labor market early on. Persistent supply shocks provide a notable boost to activity later on. Monetary policy is initially contractionary, presumably due to a binding ZLB, but its effects turn positive at the beginning of 2021.

The right-hand plot in Figure 5 shows the volatility of core PCE inflation and its sharp rise in the middle of 2021. Our estimates show Covid having a large positive impact on inflation in 2020q2 and remaining inflationary throughout the period, consistent with our earlier interpretation of it as being, on net, a supply shock. While the Covid shock pushes up inflation, it does so by relatively little after 2020q2. The model attributes most of the gyrations in inflation to transitory supply shocks — in this case, price markup shocks. This latter finding is consistent with Del Negro et al. (2022).

Adding the Covid shock also captures the key aspects of co-movement we highlighted in Figure 1 of the introduction. This shown in Figure 6 which shows the cumulative effects of the estimated Covid shocks from 2020q2 through 2020q4 (dashed lines) along with the corresponding data (stars). The figure shows that the Covid shock leads to consumption falling a lot more relative to investment compared to a typical recession and by more than hours. Furthermore, consumption, output, investment, and hours are all highly synchronized over the first three periods of the COVID episode with all four variables showing sharp reversals in growth in 2020q3. While reproducing fairly well the joint dynamics of output, consumption, investment, and hours, the single new shock cannot explain *all* the co-movement during the COVID period. As there is a lot that is unusual about the co-movement the model relies on the usual structural shocks quite a bit. It is rarely if ever the case with empirical DSGE models that a single shock can explain all the co-movement.

5.3. The Covid shock versus the usual shocks

Why do we find that the usual shocks play a significant role during our sample period? It is not because of our priors on the Covid shock's parameters.[16]. The 19 structural shocks already in the model provide it with a very rich stochastic structure that turns out to be well-suited to explain the data. The data, including the SPF expectations, observed in 2020 are highly anomalous; that is, largely inconsistent with the propagation of the usual business cycle shocks. Consequently, the model benefits from the additional shock in 2020.

The necessity of adding the Covid shock peters out as macroeconomic dynamics and the SPF expectations normalize, which from the perspective of our estimated model happens already in 2021. Despite the fact that COVID was a very articulated and powerful event that set off a number of concatenated reactions (three fiscal stimuli, supply chain disruptions, various rounds of vaccination, recrudescence due to the emergence of new variants of the virus etc.) spreading over many quarters, after the initial few quarters the macroeconomic effects of these developments can be accounted for by the usual shocks. In a sense, these subsequent

[16] We discuss the role of the priors in detail in H

Figure 6: The Covid shock and unusual co-movement during the COVID episode

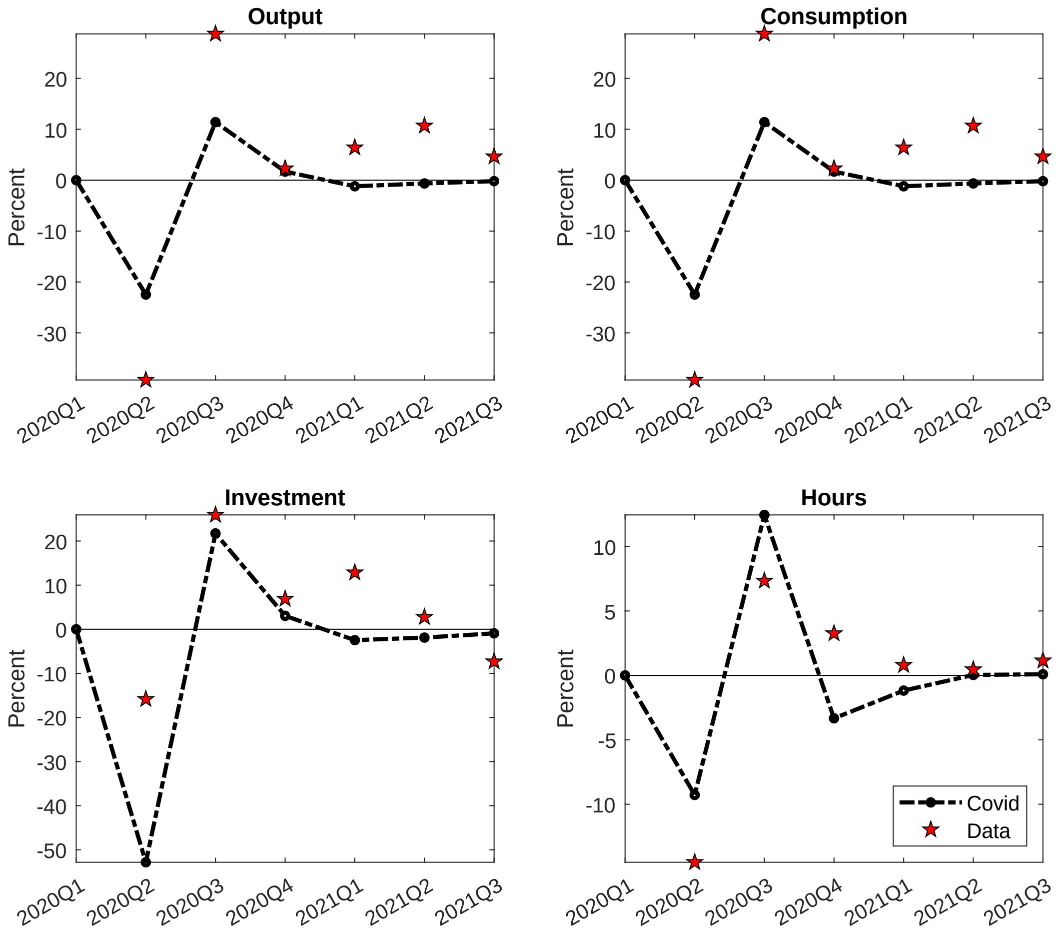

Note: The dashed lines shows the cumulative effects of the Covid shocks realized in 2020q2, 2020q3 and 2020q4 and the red stars indicate data, for the indicated variable. Source: Haver Analytics and author's calculations.

events are not so unusual from the perspective of our model estimated with pre-COVID data to demand the addition of a new shock.

A related issue is that the Covid shocks we estimate in 2020 do not explain most of the economic fluctuations.[17] As is shown in Figure 5, while the Covid shock explains a significant fraction of the hours fluctuations, demand shocks play a large role in 2020q3 and 2020q4.

[17] We explored reducing the standard deviations of the usual shocks by factors of 10 and 100. Doing so did very little to increase the role of the Covid shocks.

This result is not too surprising if one interprets it in light of the fact that recessions and recoveries are almost never explained by a single shock in empirical DSGE models. Indeed, it is typical of these models to explain the data using a combination of shocks.

Because of their flexibility the usual shocks can account for the magnitude of the fluctuations and the co-movement during the COVID episode illustrated in Figure 1 on their own, but the shocks have to be far outside the historical distribution. However these shocks cannot account for the forecast revisions shown in Figure 2 as demonstrated by Figure 4. By adding just one shock the usual shocks are more normal-sized.[18].

Finally, one may be concerned that the Covid shock is merely adding degrees of freedom that would help match the data at any time. In other words, is the Covid shock capturing something genuinely unusual or is it merely a way to fit the data better more generally. To assess this possibility we conducted the following placebo experiment. We assume new shocks arrive in 2019q1, 2019q2 and 2019q3 and reduce the measurement error on the SPF forecasts to estimate their parameters just as we do for the COVID period. In J we show the smoothed shock decomposition of hours and inflation in terms of the structural shocks and the new shocks over the period 2017q4-2019q4. We find the additional shock has very small effects, suggesting that the Covid shock is indeed an unusual shock.

5.4. The role of beliefs in the propagation of the Covid shock

The macroeconomic effects of the Covid shock in any given period can be fully decomposed into three parts: those due the current surprise (ψ_t^0), news about the future path of the Covid shock (ψ_t^1, ...,ψ_t^N), and the propagation through the economy of the past surprises and news. The latter is simply the sum of the impulse response functions of past surprise and news shocks. Figure 7 shows the contribution of the smoothed Covid shock to output growth and inflation decomposed into surprise (blue), news (yellow), and propagation (gray). The dashed line is the overall effect of Covid (the sum of the bars) and the red stars indicate data.

[18]See Figure 14 compared to Figure 13 in I

Figure 7: The role of beliefs about the Covid shock, 2020q2–2021q3

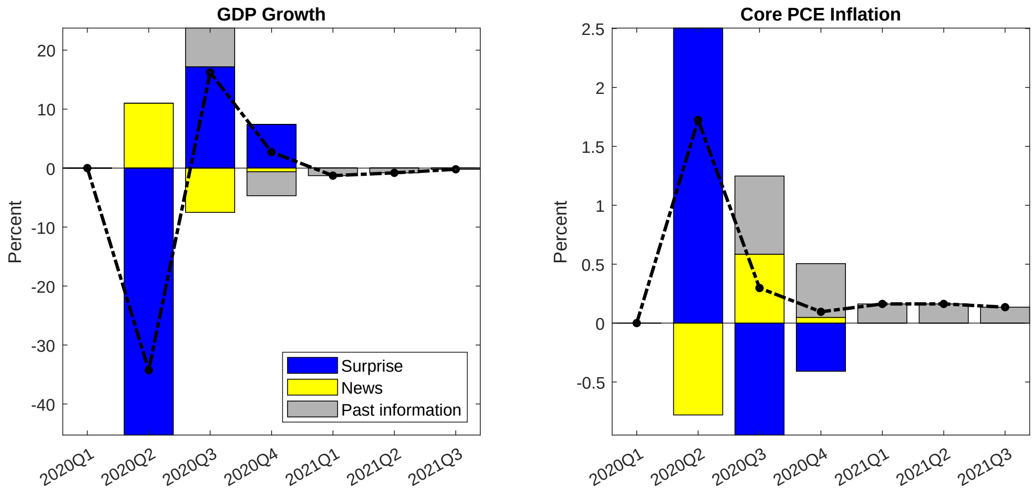

Note: Kalman smoothed decomposition of the contributions of the surprise, news, and past surprise and news shocks to output as described in the main text. The dashed lines shows the cumulative effects of the Covid shocks realized in 2020q2, 2020q3, and 2020q4 on output and inflation. The bars decompose the total effects into three components. The blue bars show the contribution of the surprise shock in the period it is realized. The yellow bars show the contribution of news in the period it is received. The gray bars shows the contribution of past surprises and news through their propagation. The red stars indicate data. Source: Haver Analytics and author's calculations.

The left plot shows that news about the future path of the Covid shock reduced the magnitude of the 2020q2 contraction in GDP by roughly 10 percentage points. In 2020q3 and 2020q4 the news about the pandemic dragged down activity, by 8 and 0.7 percentage points, respectively. This finding is consistent with concerns about the future path of the pandemic summarized in releases of Wolters Kluwer's *Blue Chip Economic Indicators* at the time. According to our model, beliefs of a recrudescence of the pandemic later on lowers current activity.

The left plot in Figure 7 also reveals that the macroeconomic consequences of the Covid shock were hard to anticipate. This is captured visually by the predominance of blue and yellow (compared to the gray). Still, the size of the gray bars suggests that at least some of the effects of Covid in 2020q3 and 2020q4 might be anticipated.[19] The contribution of news and surprise on inflation, as shown in the right plot, is the mirror image of that on GDP

[19] In the case of Delta, when we assume perfect foresight the yellow and blue would only appear in the period of the shock and the remainder of the path would be gray bars only.

growth. This suggests that on net both surprise and news about the path of Covid were dominated by supply effects.

6. Alternative approaches to addressing the COVID episode

Our methodology accounts for the unusual nature of the anticipated recovery, but there are other approaches one could take to try to achieve the same end, such as trimming the data, adding dummy variables, or introducing measurement error on the observables. These approaches share the property that they downplay the dynamics in the COVID episode. We now use a simple exercise to show why our approach might be preferred over these alternatives.

The exercise involves taking our DSGE model without the Covid shock and adding error terms to the measurement equations for output, consumption, investment and hours in 2020q2, 2020q3, and 2020q4 and then estimating the measurement error with loose priors. Figure 8 shows the actual data for GDP and the SPF forecasts (black) alongside the data the model infers (blue). The figure shows the model chooses to completely sidestep the GDP data. Rather than explaining the severe recession, the model interprets the data as masking a boom. The model's forecast is close to the SPF forecast even though we have left the standard deviations of the measurement errors on the forecasts at their pre-pandemic estimated values.

Figure 8: GDP growth and forecasts in 2020q2 from the model with measurement error

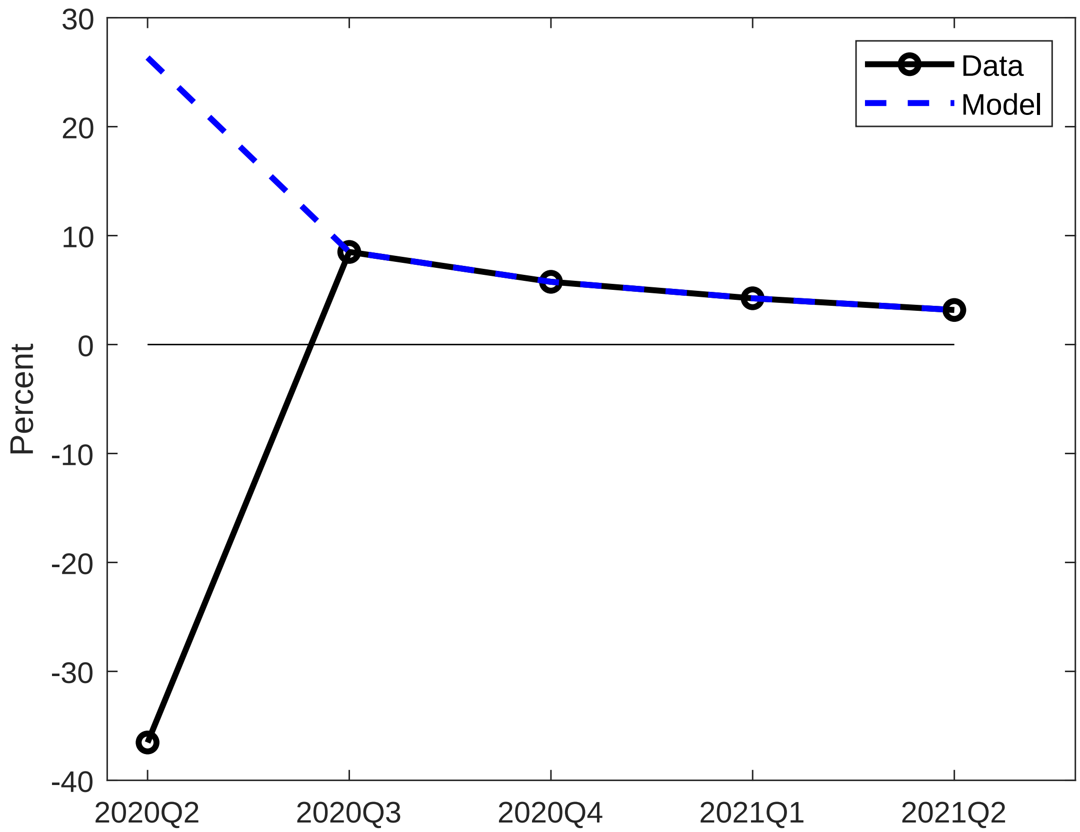

Note: Actual GDP growth and SPF forecasts of GDP growth are indicated by the black line. GDP growth implied by the model with measurement error and the forecasts of GDP with that model are shown by the blue dashed line.

Why does the model ignore the output data in 2020q2? The model has a choice if it wants to conform with the historical experience which does not include forecasts of sharp reversals in output growth. It could assign a lot of error to the SPF forecasts so the model does not see an anticipated sharp reversal in GDP growth and try to explain the massive drop in GDP in 2020q2. Or, it could stick with the SPF forecasts and ignore the output data. It chooses the latter because of the extreme GDP growth in 2020q2 compared to the more modest forecasts. Bringing the extreme GDP observation in would require the model to engineer a mix of very large structural shocks, which are highly penalized by the likelihood.

The SPF expectations for 2020q3 are large by historical standards, but much less extreme than the drop in GDP observed in 2020q2. So, the model chooses the path of least resistance. Note that the model could have chosen large measurement errors for the forecasts – this is evident in Figure 4.

This finding that the model chooses to sidestep the GDP data in 2020q2 when given the opportunity to do so has important consequences for addressing questions other than the one considered in this paper, i.e. the role of the Covid in explaining the economic dynamics in the COVID episode. For example, suppose we are interested in the dynamics of the output gap during the COVID recovery. Sidestepping the data is highly problematic for answering this question. Specifically, since the model with measurement errors does not see the gigantic output contraction in 2020q2 and instead sees a massive boom, it infers a large positive output gap in 2020q2 that agents expect to be positive and rising going forward. This dramatic failure of the model with measurement errors to provide a reasonable estimate of the output gap during the post-pandemic recovery is shown by the black line in Figure 9. The model estimated with the Covid shock sees the recession and is able to understand that even though the US economy improved dramatically in 2020q3, on net, the output gap is negative due to the large contraction observed in the previous quarter. This is shown with the red line in Figure 9. This configuration of the output gap seems much more plausible.

Figure 9: Output gap and its forecast as of 2020q2 with the Covid shock and with measurement errors

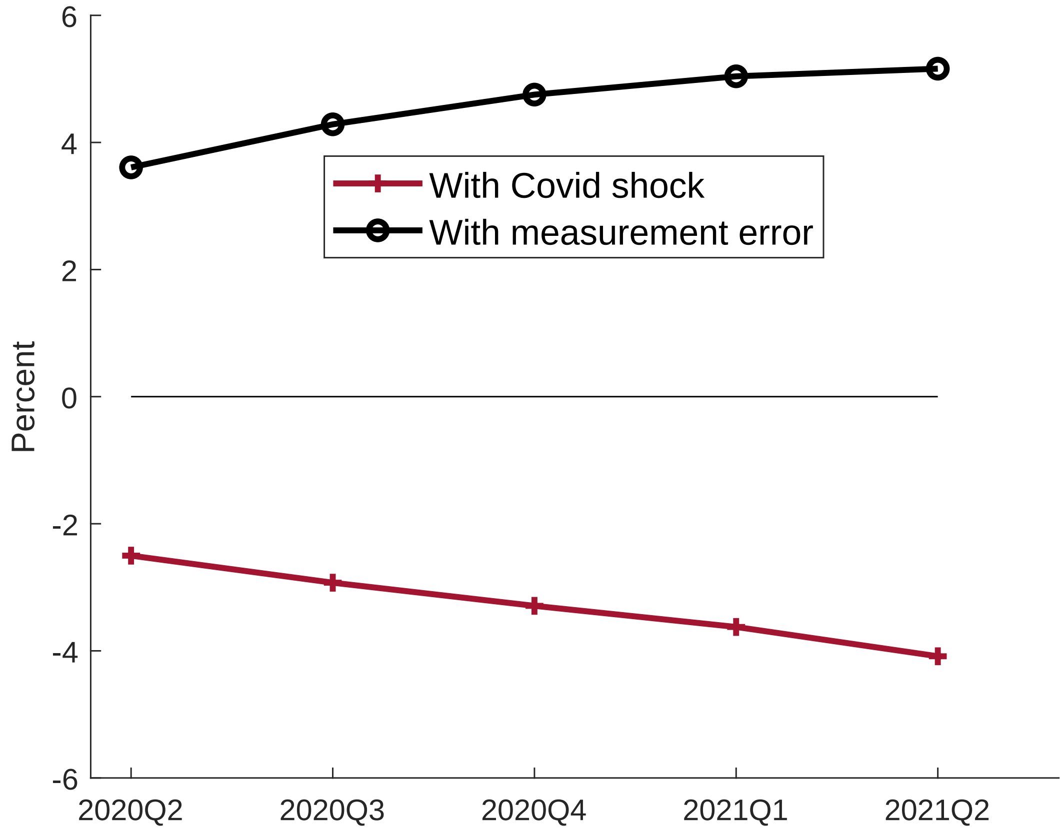

Note: We use standard measures of the output gap. They are the difference between output in the same model with and without nominal rigidities using the same shocks from the model with nominal rigidities except the real model does not have markup or monetary shocks.

The findings from this exercise also have important implications for step 3 of the estimation described in Section 2.4. The initial state is a key input to step 3. By going the measurement error route the initial state will be implausible as it reflects a period of above potential growth rather than growth below potential. This suggests the Covid shock is essential if one wishes to estimate a model with data from before, during, and after the COVID period.

7. Conclusion

We proposed the Covid shock to add to our usual business cycle models, applied it to study the macroeconomic consequences of the COVID pandemic within the context of a medium-scale DSGE model, and show how it can be used to estimate DSGE models with data from before, during, and after the COVID period. The Covid shock is most useful for capturing the unusual dynamics in 2020. The economic dynamics after the initial phase are not dissimilar to data used to identify the usual shocks and so these shocks eventually take over. For instance, the severe supply chain issues that were not apparent early on are well-captured by the usual markup shocks. Another interesting development of the COVID period was the shift in demand from services to durable consumption. If the shift in demand is important for predicting output and inflation then our use of professional forecasts and news about the Covid shock will, in principle, capture the effects.

Our framework can be applied to estimate any DSGE model with data that includes the pandemic period.[20] For example, it can be applied to models that are more suitable to study fiscal shocks than the one studied in this paper. Such models could be used to isolate the economic dynamics due to the unusual COVID fiscal policy from those of other factors driving the effects of the Covid shock. There will be plenty to learn about the propagation of fiscal shocks from the COVID episode, but it will be necessary to account for the pandemic, social distancing, and government restrictions on economic activity to separately identify their effects. Our framework provides a way to do this without modeling the underlying structure.

[20] This includes non-linear DSGE models such as those studied by Aruoba, Cuba-Borda, and Schorfheide (2018) and Gust, Herbst, López-Salido, and Smith (2017). Non-linearities could be important in the case of the COVID pandemic as the variation in macroeconomic data is so large.

References

Acemoglu, D., V. Chernozhukov, I. Werning, and M. Whinston (2021). Optimal targetd lockdowns in a multigroup SIR model. *American Economic Review: Insights 3*(4), 487–502.

Ansgar, A. N. K. and Walther (2021, September). Asymmetric attention. *American Economic Review 111*(9), 2879–2925.

Anzoategui, D., D. Comin, M. Gertler, and J. Martinez (2019). R&D as sources of business cycle persistence. *American Economic Journal: Macroeconomics 11*(3), 67–110.

Aruoba, B., P. Cuba-Borda, and F. Schorfheide (2018). Macroeconomic dynamics near the ZLB: A tale of two countries. *Review of Economic Studies 85*(1), 87–118.

Baqaee, D. and E. Farhi (2022). Supply and demand in disaggregated Keynesian economies with an application to the COVID-19 crisis. *American Economic Review 112*(5), 1397–1436.

Barro, R. (2020). Non-parmaceutical interventions and mortality in us cities during the Great Influenza Pandemic, 1918–1919. NBER Working Paper 27049.

Barro, R., J. Ursua, and J. Weng (2020). The coronavirus and the great influenza pandemic: Lessons from the "Spanish flu" for the coronavirus' potential effects on mortality and economic activity. NBER Working Paper 26966.

Bianchi, F., C. Ilut, and H. Saijo (2023, February). Diagnostic business cycles. *The Review of Economic Studies*.

Bordalo, P., N. Gennaioli, Y. Ma, and A. Shleifer (2020, September). Overreaction in macroeconomic expectations. *American Economic Review 110*(9), 2748–82.

Calvo, G. A. (1983). Staggered Prices in a Utility-Maximizing Framework. *Journal of Monetary Economics 12*, 383–398.

Campbell, J., F. Ferroni, J. Fisher, and L. Melosi (2019). The limits of forward guidance. *Journal of Monetary Economics 108*, 118–134.

Campbell, J., J. Fisher, A. Justiniano, and L. Melosi (2016). Forward guidance and macroeconomic outcomes since the financial crisis. In M. Eichenbaum and J. Parker (Eds.), *NBER Macroeconomics Annual 2016*, Volume 31. University of Chicago Press.

Cardani, R., O. Croitorov, M. Giovannini, P. Pfeiffer, M. Ratto, and L. Vogel (2022). The Euro area's pandemic recession: A DSGE-based interpretation. *Journal of Economic Dynamics and Control 143*.

Christiano, L., M. Eichenbaum, and C. Evans (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy 113*(1), 1–45.

Del Negro, M., A. Gleich, S. Goyal, A. Johnson, and A. Tambalotti (2022). Drivers of inflation: The New York Fed DSGE Model's Perspective. *Liberty Street Economics*.

den Haan, W. and T. Drechsel (2021). Agnostic structural disturbances (ASDs): Detecting and reducing misspecification in empirical macroeconomic models. *Journal of Monetary Economics 117*, 258–277.

Eichenbaum, M., B. Johannsen, and S. Rebelo (2021). Monetary policy and the predictability of nominal exchange rates. *The Review of Economic Studies 88*(1), 192–228.

Eichenbaum, M., S. Rebelo, and M. Trabandt (2021). The macroeconomics of epidemics. *Review of Financial Studies 34*(11), 5149–5187.

Fernandez-Villaverde, J., J. Rubio-Ramirez, and F. Schorfheide (2016). Solution and estimation methods for DSGE models. In J. Taylor and H. Uhlig (Eds.), *Handbook of Macroeconomics*, Volume 1 of *Handbooks in Economics*, pp. 527–724. North-Holland: Elsevier Science, North-Holland.

Fisher, J. (2015, March-April). On the structural interpretation of the Smets-Wouters "Risk Premium" shock. *Journal of Money, Credit and Banking 47*(2-3), 511–516.

Guerrieri, V., G. Lorenzoni, L. Straub, and I. Werning (2022). Macroeconomic implications of COVID-19: Can negative supply shocks cause demand shortages? *American Economic Review 112*(5), 1437–1474.

Gürkaynak, R. S., B. Sack, and E. T. Swanson (2005, May). Do actions speak louder than words? The response of asset prices to monetary policy actions and statements. *International Journal of Central Banking 1*(1), 55–93.

Gust, C., E. Herbst, D. López-Salido, and M. Smith (2017). The empirical implications of the interest-rate lower bound. *American Economic Review 107*(7), 1971–2006.

Herbst, E. and F. Schorfheide (2015). *Bayesian estimation of DSGE models*. Princeton, NJ: Princeton University Press.

Inoue, A., C. Kuo, and B. Rossi (2020). Identifiying the sources of model misspecification. *Journal of Monetary Economics 110*, 1–18.

Ireland, P. (2004). A method for taking models to the data. *Journal of Economic Dynamics and Control 28*(6), 1205–1226.

Justiniano, A., G. E. Primiceri, and A. Tambalotti (2013, April). Is there a trade-off between inflation and output stabilization? *American Economic Journal: Macroeconomics 5*(2), 1–31.

Lenza, M. and G. Primiceri (2022). How to estimate a VAR after March 2020. *Journal of Applied Econometrics 37*, 688–699.

Ludvigson, S., S. Ma, and S. Ng (2020). Covid-19 and the macroeconomic effects of costly disasters. NBER Working Paper 26987.

Michaillat, P. and E. Saez (2021). Resolving New Keynesian anomolies with wealth in the utility function. *The Review of Economics and Statistics 103*(2), 197–215.

Primiceri, G. and A. Tambalotti (2020). Macroeconomic forecasting in the time of COVID-19. Manuscript.

Schmitt-Grohé, S. and M. Uribe (2012). What's news in business cycles. *Econometrica 80*(6), 2733–2764.

Smets, F. and R. Wouters (2007). Shocks and frictions in US business cycles: A Bayesian DSGE approach. *The American Economic Review 97*(3), 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.

Yi, K.-M. and J. Zhang (2017). Understanding global trends in long-run real interest rates. *Economic Perspectives 41*(2).

Appendix to "Unusual shocks in our usual models" by Filippo Ferroni, Jonas Fisher, and Leonardo Melosi[21]

This appendix describes the DSGE model, estimation, and measurement in detail: A – F.[22] We also discuss the Covid shock parameter estimates (G), the role of the prior on the Covid shock's parameters (H), the effect of adding the Covid shock on the realizations of the usual shocks (I), and the placebo experiment mentioned in the main text (J).

A. The Model's Primitives

Eight kinds of agents populate the model economy: Households, investment producers, competitive final goods producers, monopolistically-competitive differentiated goods producers, labor packers, monopolistically-competitive guilds, a fiscal authority, and a monetary authority. These agents interact with each other in markets for: final goods used for consumption and investment, investment goods used to augment the stock of productive capital, differentiated intermediate goods, capital services, raw labor, differentiated labor, composite labor, government bonds, privately-issued bonds, and state-contingent claims.

A.1. Households

Our model's households are the ultimate owners of all assets in positive net supply (the capital stock, differentiated goods producers, and guilds). They provide labor and divide their current after-tax income (from wages and assets) between current consumption, investment in productive capital, and purchases of financial assets, both those issued by the government and those issued by other households. The individual household divides its current resources between consumption and the available vehicles for intertemporal substitution (capital and financial assets) to maximize a discounted sum of current and expected future felicity.

$$\mathbb{E}_t \left[\sum_{\tau=0}^{\infty} \beta^{\tau} \varepsilon^b_{t+\tau} \left(U_{t+\tau} + \varepsilon^s_{t+\tau} L \left(\frac{B_{t+\tau}}{P_{t+\tau} R_{t+\tau}} \right) \right) \right]$$

with

$$U_t = \frac{1}{1-\gamma_c} \left((C_t - \varrho \bar{C}_{t-1})(1 - H_t^{1+\gamma_h}) \right)^{(1-\gamma_c)} \tag{9}$$

The function $L(\cdot)$ is strictly increasing, concave, and differentiable everywhere on $[0, \infty)$. In particular, $L'(0)$ exists and is finite. Without loss of generality, we set $L'(0)$ to one. The argument of $L(\cdot)$ equals the real value of government bonds in the household's portfolio: their period $t+1$ redemption value B_t divided by their nominal yield R_t expressed in units of the consumption good with the nominal price index P_t. The time-varying coefficient

[21]The first parts of this appendix are co-authored with Jeffrey Campbell, University of Notre Dame and Tilburg University, jcampbel24@nd.edu.

[22]Some of the notation in this Appendix overlaps with the main text. To avoid confusion note that none of the notation here refers to variables discussed in the main tex.

multiplying this felicity from bond holdings, ε_t^s, is the liquidity preference shock introduced by Fisher (2015). A separate shock influences the household's discounting of future utility to the present, ε_t^b. Specifically, the household discounts a certain utility in $t+\tau$ back to t with $\beta^\tau \mathbb{E}_t\left[\varepsilon_{t+\tau}^b/\varepsilon_t^b\right]$. In logarithms, these two preference shocks follow independent autoregressive processes.

$$\ln \varepsilon_t^b = (1-\rho_b)\ln \varepsilon_*^b + \rho_b \ln \varepsilon_{t-1}^b + \eta_t^b, \eta_t^b \sim \mathbb{N}(0,\sigma_b^2) \tag{10}$$

$$\ln \varepsilon_t^s = (1-\rho_s)\ln \varepsilon_*^s + \rho_s \ln \varepsilon_{t-1}^s + \eta_t^s, \eta_t^s \sim \mathbb{N}(0,\sigma_s^2). \tag{11}$$

A household's wealth at the beginning of period t consists of its nominal government bond holdings, B_t, its net holdings of privately-issued financial assets, and its capital stock K_{t-1}. The household chooses a rate of capital utilization u_t, and the capital services resulting from this choice equal $u_t K_{t-1}$. The cost of increasing utilization is higher depreciation. An increasing, convex and differentiable function $\delta(U)$ gives the capital depreciation rate. We specify this as

$$\delta(u) = \delta_0 + \delta_1(u-u_\star) + \frac{\delta_2}{2}(u-u_\star)^2 .$$

A household can augment its capital stock with investment, I_t. Investment requires paying adjustment costs of the "i-dot" form introduced by Christiano, Eichenbaum, and Evans (2005). Also, an *investment demand shock* alters the efficiency of investment in augmenting the capital stock. Altogether, if the household's investment in the previous period was I_{t-1}, and it purchases I_t units of the investment good today, then the stock of capital available in the *next* period is

$$K_t = (1-\delta(u_t))\,K_{t-1} + \varepsilon_t^i\left(1 - S\left(\frac{A_{t-1}^K I_t}{A_t^K I_{t-1}}\right)\right) I_t. \tag{12}$$

In (12), A_t^K equals the productivity level of capital goods production, described in more detail below, and ε_t^i is the investment demand shock. In logarithms, this follows a first-order autoregression with a normally-distributed innovation.

$$\ln \varepsilon_t^i = (1-\rho_i)\ln \varepsilon_*^i + \rho_i \ln \varepsilon_{t-1}^i + \eta_t^i, \eta_t^i \sim \mathbb{N}(0,\sigma_i^2) \tag{13}$$

A.2. Production

The producers of investment goods use a linear technology to transform the final good into investment goods. The technological rate of exchange from the final good to the investment good in period t is A_t^I. We denote $\Delta \ln A_t^I$ with ω_t, which we call the *investment-specific technology shock* and which follows first-order autogregression with normally distributed innovations.

$$\omega_t = (1-\rho_\omega)\omega_\star + \rho_\omega \omega_{t-1} + \eta_t^\omega, \eta_t^\omega \sim \mathbb{N}(0,\sigma_\omega^2) \tag{14}$$

Investment goods producers are perfectly competitive.

Final good producers also operate a constant-returns-to-scale technology; which takes as inputs the products of the differentiated goods producers. To specify this, let Y_{it} denote the quantity of good i purchased by the representative final good producer in period t, for $i \in [0, 1]$. The representative final good producer's output then equals

$$Y_t \equiv \left(\int_0^1 Y_{it}^{\frac{1}{1+\lambda_t^p}} di \right)^{1+\lambda_t^p} .$$

With this technology, the elasticity of substitution between any two differentiated products equals $1 + 1/\lambda_t^p$ in period t. Although this is constant across products within a time period, it varies stochastically over time according to an ARMA$(1, 1)$ in logarithms.

$$\ln \lambda_t^p = (1 - \rho_p) \ln \lambda_\star^p + \rho_p \ln \lambda_{t-1}^p - \theta_p \eta_{t-1}^p + \eta_t^p, \eta_t^p \sim \mathbb{N}(0, \sigma_p^2) \tag{15}$$

Given nominal prices for the intermediate goods P_{it}, it is a standard exercise to show that the final goods producers' marginal cost equals

$$P_t = \left(\int_0^1 P_{it}^{-\frac{1}{\lambda_t^p}} di \right)^{-\lambda_t^p} \tag{16}$$

Just like investment goods firms, the final goods' producers are perfectly competitive. Therefore, profit maximization and positive final goods output together require the competitive output price to equal P_t. Therefore, we can define inflation of the nominal final good price as $\pi_t \equiv \ln(P_t/P_{t-1})$.

The intermediate goods producers each use the technology

$$Y_{it} = (K_{it}^e)^\alpha \left(A_t^Y H_{it}^d \right)^{1-\alpha} - A_t \Phi \tag{17}$$

Here, K_{it}^e and H_{it}^d are the capital services and labor services used by firm i, and A_t^Y is the level of neutral technology. Its growth rate, $\nu_t \equiv \ln(A_t^Y / A_{t-1}^Y)$, follows a first-order autogregression.

$$\nu_t = (1 - \rho_\nu) \nu_* + \rho_v \nu_{t-1} + \eta_t^\nu, \eta_t^\nu \sim \mathbb{N}(0, \sigma_\nu^2), \tag{18}$$

The final term in (17) represents the fixed costs of production. These grow with

$$A_t \equiv A_t^Y \left(A_t^I \right)^{\frac{\alpha}{1-\alpha}} . \tag{19}$$

We demonstrate below that A_t is the stochastic trend in equilibrium output and consumption, measured in units of the final good. We denote its growth rate with

$$z_t = \nu_t + \frac{\alpha}{1-\alpha} \omega_t \tag{20}$$

Similarly, define

$$A_t^K \equiv A_t A_t^I \tag{21}$$

In the specification of the capital accumulation technology, we labelled A_t^K the "productivity level of capital goods production." We demonstrate below that this is indeed the case with the definition in (21).

Each intermediate goods producer chooses prices subject to a Calvo (1983) pricing scheme. With probability $\zeta_p \in [0,1]$, producer i has the opportunity to set P_{it} without constraints. With the complementary probability, P_{it} is set with the indexing rule

$$P_{it} = P_{it-1}\pi_{t-1}^{\iota_p}\pi_\star^{1-\iota_p}. \tag{22}$$

In (22), $\pi_\star$ is the gross rate of price growth along the steady-state growth path, and $\iota_p \in [0,1]$.[23]

A.3. Labor Markets

Households' hours worked pass through two intermediaries, guilds and labor packers, in their transformation into labor services used by the intermediate goods producers. The guilds take the households' homogeneous hours as their only input and produce differentiated labor services. These are then sold to the labor packers, who assemble the guilds' services into composite labor services.

The labor packers operate a constant-returns-to-scale technology with a constant elasticity of substitution between the guilds' differentiated labor services. For its specification, let H_{it} denote the hours of differenziated labor purchased from guild i at time t by the representative labor packer. Then that packer's production of composite labor services, H_t^s are given by

$$H_t^s = \left(\int_0^1 (H_{it})^{\frac{1}{1+\lambda_t^w}}\, di\right)^{1+\lambda_t^w}.$$

As with the final good producer's technology, an ARMA$(1,1)$ in logarithms governs the constant elasticity of substitution between any two guilds' labor services.

$$\ln\lambda_t^w = (1-\rho_w)\ln\lambda_\star^w + \rho_w \ln\lambda_{t-1}^w - \theta_w\eta_{t-1}^w + \eta_t^w, \eta_t^w \sim \mathbb{N}(0,\sigma_w^2) \tag{23}$$

Just as with the final goods producers, we can easily show that the labor packers' marginal cost equals

$$W_t = \left(\int_0^1 (W_{it})^{-\frac{1}{\lambda_t^w}}\, di\right)^{-\lambda_t^w}. \tag{24}$$

Here, W_{it} is the nominal price charged by guild i per hour of differentiated labor. Since labor packers are perfectly competitive, their profit maximization and positive output together require that the price of composite labor services equals their marginal cost.

Each guild produces it's differentiated labor service using a linear technology with the household's hours worked as its only input. A Calvo (1983) pricing scheme similar to that

[23] To model firms' price-setting opportunities as functions of s_t, define a random variable u_t^p which is independent over time and uniformly distributed on $[0,1]$. Then, firm i gets a price-setting opportunity if either $u_t^p \geq \zeta_p$ and $i \in [u_t^p - \zeta_p, u_t^p]$ or if $u_t^p < \zeta_p$ and $i \in [0, u_t^p] \cup [1 + u_t^p - \zeta_p, 1]$.

of the differentiated goods producers constrains their nominal prices. Guild i has an unconstrained opportunity to choose its nominal price with probability $\zeta_w \in [0, 1]$. With the complementary probability, W_{it} is set with an indexing rule based on π_{t-1} and last period's trend growth rate, z_{t-1}.

$$W_{it} = W_{it-1} \left(\pi_{t-1} e^{z_{t-1}}\right)^{\iota_w} \left(\pi_\star e^{z_\star}\right)^{1-\iota_w}. \tag{25}$$

In (25), $z_\star \equiv \nu_\star + \frac{\alpha}{1-\alpha}\omega_\star$ is the unconditional mean of z_t and $\iota_w \in [0, 1]$.

A.4. Fiscal and Monetary Policy

The model economy hosts two policy authorities, each of which follows exogenously-specified rules that receive stochastic disturbances. The fiscal authority issues bonds, B_t, collects lump-sum taxes T_t, and buys "wasteful" public goods G_t. Its period-by-period budget constraint is

$$G_t + B_{t-1} = T_t + \frac{B_t}{R_t}. \tag{26}$$

The left-hand side gives the government's uses of funds, public goods spending and the retirement of existing debt. The left-hand side gives the sources of funds, taxes and the proceeds of new debt issuance at the interest rate R_t. We assume that the fiscal authority keeps its budget balanced period-by-period, so $B_t = 0$. Furthermore, the fiscal authority sets public goods expenditure equal to a stochastic share of output, expressed in consumption units.

$$G_t = (1 - 1/g_t)Y_t, \tag{27}$$

with

$$\ln g_t = (1 - \rho_g) \ln s_\star^g + \rho_g \ln g_{t-1} + \eta_t^g, \eta_t^g \sim \mathbb{N}(0, \sigma_g^2). \tag{28}$$

The monetary authority sets the nominal interest rate on government bonds, R_t. For this, it employs a Taylor rule with interest-rate smoothing and forward guidance shocks.

$$\ln R_t = \rho_R \ln R_{t-1} + (1 - \rho_R) \ln R_t^n + \sum_{j=0}^{M} \xi_{t-j}^j. \tag{29}$$

The monetary policy disturbances in (29) are $\xi_t^0, \xi_{t-1}^1, \ldots, \xi_{t-M}^M$. The public learns the value of ξ_{t-j}^j in period $t - j$. The conventional unforecastable shock to current monetary policy is ξ_t^0, while for $j \geq 1$, these disturbances are *forward guidance shocks*. We gather all monetary shocks revealed at time t into the vector ε_t^R. This is normally distributed and *i.i.d.* across time. However, *its elements may be correlated with each other.* That is,

$$\varepsilon_t^R \equiv \left(\xi_t^0, \xi_t^1, \ldots, \xi_t^M\right) \sim \mathbb{N}(0, \Sigma_\varepsilon). \tag{30}$$

The off-diagonal elements of Σ^1 are not necessarily zero, so forward-guidance shocks need not randomly impact expected future monetary policy at two adjacent dates independently.

Current economic circumstances influence R_t through the notional interest rate, R_t^n.

$$\ln R_t^n = \ln r_\star + \ln \pi_t^\star + \frac{\phi_1}{4}\mathbb{E}_t \sum_{j=-2}^{1} (\ln \pi_{t+j} - \ln \pi_t^\star) + \frac{\phi_2}{4}\mathbb{E}_t \sum_{j=-2}^{1} (\ln Y_{t+j} - \ln y^\star - \ln A_{t+j}) . \quad (31)$$

The constant $r_\star$ equals the real interest rate along a steady-state growth path, and $\pi_t^\star$ is the central bank's intermediate target for inflation. We call this the *inflation-drift shock*. it follows a first-order autoregression with a normally-distributed innovation. Its unconditional mean equals $\pi_\star$, the inflation rate on a steady-state growth path.

$$\ln \pi_t^\star = (1 - \rho_\pi)\pi_\star + \rho_\pi \ln \pi_{t-1}^\star + \eta_t^\pi, \eta_t^\pi \sim \mathbb{N}(0, \sigma_\pi^2) \quad (32)$$

Allowing $\pi_t^\star$ to change over time enables our model to capture the persistent decline in inflation from the early 1990s through the early 2000s engineered by the Greenspan FOMC.

A.5. Other Financial Markets and Equilibrium Definition

All households participate in the market for nominal risk-free government debt. Additionally, they can buy and sell two classes of privately issued assets without restriction. The first is one-period nominal risk-free *private* debt. We denote the value of household's net holdings of such debt at the beginning of period t with B_{t-1}^P and the interest rate on such debt issued in period t maturing in $t+1$ with R_{t+1}^P. The second asset class consists of a complete set of *real* state-contingent claims. As of the end of period t, the household's ownership of securities that pay off one unit of the aggregate consumption good in period τ if history s^τ occurs is $Q_t(s^\tau)$, and the nominal price of such a security in the same period is $J_t(s^\tau)$.

We define an equilibrium for our economy in the usual way: Households maximize their utility given all prices, taxes, and dividends from both producers and guilds; final goods producers and labor packers maximize profits taking their input and output prices as given; differentiated goods producers and guilds maximize the market value of their dividend streams taking as given all input and financial-market prices; differentiated goods producers and guilds produce to satisfy demand at their posted prices; and otherwise all product, labor, and financial markets clear.

B. Detrending

To remove nominal and real trends, we deflate nominal variables by their matching price deflators, and we detrend any resulting real variables influenced permanently by technological change. All scaled versions of variables are the lower-case counterparts.

$$c_t = \frac{C_t}{A_t} \qquad i_t = \frac{I_t}{A_t A_t^I}$$

$$k_t = \frac{K_t}{A_t A_t^I} \qquad k_t^e = \frac{K_t^e}{A_t A_t^I}$$

$$w_t = \frac{W_t}{A_t P_t} \qquad \tilde{w}_t = \frac{\tilde{W}_t}{A_t P_t}$$
$$\tilde{p}_t = \frac{\tilde{P}_t}{P_t} \qquad \pi_t = \frac{P_t}{P_{t-1}}$$
$$y_t = \frac{Y_t}{A_t} \qquad mc_t = \frac{MC_t}{P_t}$$
$$r_t^k = \frac{R_t^k A_t^I}{P_t} \qquad w_t^h = \frac{W_t^h}{A_t P_t}$$
$$\lambda_t^1 = \Lambda_t^1 A_t^{\gamma_C} \qquad \lambda_t^2 = \Lambda_t^2 A_t^{\gamma_C} A_t^I$$
$$\varepsilon_t^s = A_t^{\gamma_C} \varepsilon_t^s$$

B.1. Detrended Equations

The detrended equations describing our model are listed in the following sections.

Households' FOC

$$\lambda_t^1 = \varepsilon_t^b \left[\left(c_t - \varrho \frac{c_{t-1}}{e^{z_t}} \right) \left(1 - \varepsilon_t^h h_t^{1+\gamma_h} \right) \right]^{-\gamma_c} \left(1 - \varepsilon_t^h h_t^{1+\gamma_h} \right)$$
$$\lambda_t^1 w_t^h = (1+\gamma_h) \varepsilon_t^b \left[\left(c_t - \varrho \frac{c_{t-1}}{e^{z_t}} \right) \left(1 - \varepsilon_t^h h_t^{(1+\sigma_h)} \right) \right]^{-\gamma_c} \left(c_t - \varrho \frac{c_{t-1}}{e^{z_t}} \right) \varepsilon_t^h h_t^{\gamma_h}$$
$$\frac{\lambda_t^1}{R_t^P} = \beta E_t \left[\frac{\lambda_{t+1}^1 e^{-\gamma_C z_{t+1}}}{\pi_{t+1}} \right]$$
$$\frac{\lambda_t^1}{R_t} - L'(0) \frac{\varepsilon_t^b \varepsilon_t^s}{R_t} = \beta E_t \frac{\lambda_{t+1}^1}{\pi_{t+1}} e^{-z_{t+1}\gamma_C}$$
$$\lambda_t^1 = \varepsilon_t^i \lambda_t^2 \left((1 - S_t(\cdot)) - S_t'(\cdot) \frac{i_t}{i_{t-1}} \right) + \beta E_t \left[\varepsilon_{t+1}^i e^{(1-\gamma_C) z_{t+1}} \lambda_{t+1}^2 S_{t+1}'(\cdot) \frac{i_{t+1}^2}{i_t^2} \right]$$
$$\lambda_t^2 = \beta E_t \left[e^{-\gamma_C z_{t+1} - \omega_{t+1}} \left(\lambda_{t+1}^1 r_{t+1}^k u_{t+1} + \lambda_{t+1}^2 (1 - \delta(u_{t+1})) \right) \right]$$
$$\lambda_t^1 r_t^k = \lambda_t^2 \delta'(u_t)$$
$$k_t = (1 - \delta(u_t)) k_{t-1} e^{-z_t - \omega_t} + \varepsilon_t^i (1 - S(\cdot)) i_t$$
$$k_t^e = u_t k_{t-1} e^{-z_t - \omega_t}$$

Final Goods Price Index

$$1 = \left[(1-\zeta_p) \tilde{p}_t^{\frac{1}{1-\lambda_{p,t}}} + \zeta_p (\pi_{t-1}^{\iota_p} \pi^{*(1-\iota_p)} \pi_t^{-1})^{\frac{1}{1-\lambda_{p,t}}} \right]^{1-\lambda_{p,t}}$$

Intermediate Goods Firms: Capital-Labor Ratio

$$\frac{k_t^e}{h_t^d} = \frac{\alpha}{1-\alpha}\frac{w_t}{r_t^k}$$

Intermediate Goods Firms: Real Marginal Costs

$$mc_t = \frac{w_t^{1-\alpha}\left(r_t^k\right)^{\alpha}}{\varepsilon_t^a \alpha^{\alpha}(1-\alpha)^{1-\alpha}}$$

Intermediate Goods Firms: Price-Setting Equation

$$0 = E_t \sum_{s=0}^{\infty} \zeta_p^s \beta^s \lambda_{t+s}^1 \frac{\tilde{y}_{t,t+s}}{\lambda_{p,t+s}-1} \left(\frac{A_{t+s}}{A_t}\right)^{1-\gamma_C} \left[\lambda_{p,t+s} mc_{t+s} - \tilde{X}_{t,s}^p \tilde{p}_t\right]$$

where

$$\tilde{X}_{t,s}^p = \left\{ \begin{array}{ll} 1 & : s = 0 \\ \frac{\prod_{j=1}^{s} \pi_{t+j-1}^{1-\iota_p} \pi_*^{\iota_p}}{\prod_{j=1}^{s} \pi_{t+j}} & : s = 1, \ldots, \infty \end{array} \right\}$$

$\tilde{y}_{t,t+s}$ denotes the time $t+j$ output sold by the producers that have optimized at time t the last time they have reoptimized. Since it can be shown that optimizing producers all choose the same price, then we do not have to carry the i-subscript.

Labor Packers: Aggregate Wage Index

$$w_t = \left[(1-\zeta_w)\tilde{w}_t^{-\frac{1}{\lambda_{w,t}}} + \zeta_w \left(e^{\iota_w z_{t-1} - z_t} e^{(1-\iota_w) z_*} \pi_{t-1}^{\iota} \pi_t^{-1} \pi_*^{1-\iota_w} w_{t-1}\right)^{-\frac{1}{\lambda_{w,t}}}\right]^{-\lambda_{w,t}}$$

Guilds: Wage-Setting Equation

$$0 = E_t \sum_{s=0}^{\infty} \zeta_w^s \beta \lambda_{t+s}^1 \left(\frac{A_{t+s}}{A_t}\right)^{1-\gamma_C} \frac{\tilde{h}_{t,t+s}}{\lambda_{w,t+s}} \left((1+\lambda_{w,t+s}) w_{t+s}^h - \tilde{X}_{t,s}^l \tilde{w}_t\right)$$

where

$$\tilde{X}^l_{t,s} = \left\{ \begin{array}{ll} 1 & : s = 0 \\ \frac{\prod_{j=1}^s \left(\pi_{t+j-1} e^{z_{t+j-1}}\right)^{1-\iota_w} (\pi\gamma)^{\iota_w}}{\prod_{j=1}^s \pi_{t+j} e^{z_{t+j}}} & : s = 1, \ldots, \infty \end{array} \right\}$$

$\tilde{h}_{t,t+s}$ denotes the time $t+j$ labor supplied by the guild that have optimized at time t the last time they have reoptimized. Since it can be shown that optimizing guilds all choose the same wage, then we do not have to carry the i-subscript.

Monetary Authority

$$R_t = R_{t-1}^{\rho_R} \left[r_* \pi_t^* \left(\prod_{j=-2}^{1} \frac{\pi_{t+j}}{\pi_t^*} \right)^{\frac{\psi_1}{4}} \left(\prod_{j=-2}^{1} \frac{y_{t+j}}{y^*} \right)^{\frac{\psi_2}{4}} \right]^{1-\rho_R} \prod_{j=0}^{M} \xi_{t-j,j}$$

The Aggregate Resource Constraint

$$\frac{y_t}{g_t} = c_t + i_t$$

Production Function

$$y_t = \varepsilon_t^a \left(k_t^e\right)^\alpha \left(h_t^d\right)^{1-\alpha} - \Phi$$

Labor Market Clearing Condition

$$h_t = h_t^d$$

C. Steady State

We normalize most shocks and the utilization rate:

$$u_\star = 1 \qquad \varepsilon^i = 1$$
$$\varepsilon^a = 1 \qquad \varepsilon^b = 1$$

Next, we set the following restriction on adjustment costs:

$$S(\cdot_*) \equiv 0$$
$$S'(\cdot_*) \equiv 0$$

C.1. Prices and Interest Rates

Given β, z_*, γ_C, and π_*, we can solve for the steady-state nominal interest rate on private bonds R^P_* by using the FOC on private bonds:

$$R^P_* = \frac{\pi_*}{(\beta e^{-\gamma_C z_*})} \tag{33}$$

From the definition of $\delta(u)$, we have

$$\begin{aligned}\delta(1) =& \delta_0 \\ \delta'(1) =& \delta_1.\end{aligned}$$

Next, given ω_*, δ_0, and the above, we can solve for the real return on capital r^k_* using the FOC on capital:

$$r^k_* = \frac{e^{\gamma_C z_* + \omega_*}}{\beta} - (1 - \delta_0) \tag{34}$$

C.2. Ratios

Moving to the production side, we can use the aggregate price equation to solve for $\tilde{p}_*$:

$$\tilde{p}_* = 1$$

Using this result and given $\lambda_{p,*}$, we can use the price Phillips curve to solve for mc_*:

$$mc_* = \frac{1}{1 + \lambda_{p,*}} \tag{35}$$

Given values for α and ε^a_*, we can use the marginal cost equation to solve for w_*:

$$w_* = \left(mc_* \alpha^\alpha (1 - \alpha)^{1-\alpha} (r^k_*)^{-\alpha}\right)^{\frac{1}{1-\alpha}} \tag{36}$$

The definition of effective capital gives us a value for k^e_* in terms of k_*:

$$k^e_* = k_* e^{-z_* - \omega_*}$$

Calculating y_* using the labor share of output $1 - \alpha$:

$$y_* = \frac{w_* h_*}{1 - \alpha}$$

Using capital shares based off our value of α, we can calculate the output to capital ratio as follows:

$$\frac{y_*}{k^e_*} = \frac{r^k_*}{\alpha}$$

$$\frac{y_*}{k_*} = e^{-z_*-\omega_*}\frac{r^k_*}{\alpha}$$

Using the capital accumulation equation, we can get a value for $\frac{i_*}{k_*}$

$$\frac{i_*}{k_*} = 1 - (1-\delta_0)e^{-z_*-\omega_*}$$

Using the resource constraint, we can get $\frac{c_*}{k_*}$:

$$\frac{c_*}{k_*} = \frac{y_*}{k_* s^g_\star} - \frac{i_*}{k_*}$$

These ratios will give us the remaining steady-state levels and ratios:

$$k_* = y_*\left(\frac{y_*}{k_*}\right)^{-1} \qquad i_* = \frac{i_*}{k_*}k_*$$
$$c_* = \frac{c_*}{k_*}k_* \qquad g_* = g_y y_*$$

C.3. Liquidity Premium

Using the aggregate wage equation, we can get the following for $\tilde{w}_*$:

$$\tilde{w}_* = w_*$$

Combining this result with the wage Phillips curve, we get the following:

$$w^h_* = \frac{w_*}{1+\lambda_{w,*}}$$

We can use the FOC for consumption and the labor supply to pin down ε^h and λ^1_*

$$\varepsilon^b\left[c_*\left(1-\frac{\varrho}{e^z}\right)\right]^{-\gamma_c}\left(1-\varepsilon^h h_*^{(1+\gamma_h)}\right) - \lambda^1_* = 0$$
$$-(1+\gamma_h)\varepsilon^b c_*^{(1-\gamma_c)}\left(1-\frac{\varrho}{\varepsilon^z}\right)^{(1-\gamma_c)}\left(1-\varepsilon^h h_*^{(1+\gamma_h)}\right)^{-\gamma_c}\varepsilon^h h_*^{\gamma_h} + \lambda^1_* w^h_* = 0$$

Finally, the government bond rate is calculated from

$$\lambda^1_* - \varepsilon^b_*\varepsilon^s_* = \beta R_*\frac{\lambda^1_*}{\pi_*}e^{-\gamma_C z}$$
$$\underbrace{\frac{\pi_*}{\beta e^{-\gamma_C z}}}_{R^P_*} - \varepsilon^b_*\varepsilon^s_*\frac{\pi_*}{\beta e^{-\gamma_C z}\lambda^1_*} = R_*$$

Noting that $R^P_* = \frac{\pi_*}{\beta e^{-\gamma_C z}}$ we can write

$$\frac{R^P_* - R_*}{R^P_*} = \frac{\varepsilon^b_* \varepsilon^s_*}{\lambda^1_*}.$$

This is the liquidity premium in steady state.

D. Log Linearization

Hatted variables refer to log deviations from steady-state ($\hat{x} = \ln\left(\frac{x_t}{x_*}\right)$). In the cases of z_t, ω_t, and ν_t, we have that $\hat{x} = x_t - x_*$ as these variables are already in logs.

Households' First Order Conditions

$$\hat{\varepsilon}^b_t - \hat{\lambda}^1_t - \gamma_c \frac{1}{1-\frac{\varrho}{e^z}}\hat{c}_t + \gamma_c \frac{\frac{\varrho}{e^z}}{1-\frac{\varrho}{e^z}}(\hat{c}_{t-1} - \hat{z}_t) \tag{37}$$

$$\hat{\lambda}^1_t + \hat{w}^h_t - \hat{\varepsilon}^b_t - \hat{\varepsilon}^h_t - \frac{1-\gamma_c}{1-\frac{\varrho}{e^z}}\hat{c}_t + (1-\gamma_c)\frac{\frac{\varrho}{e^z}}{1-\frac{\varrho}{e^z}}(\hat{c}_{t-1} - \hat{z}_t) \tag{38}$$

$$-\left(\gamma_h + \gamma_c\left(1+\gamma_h\right)\frac{\varepsilon^h h_*^{1+\gamma_h}}{(1-\varepsilon^h h_*^{1+\gamma_h})^2}\right)\hat{h}_t = 0$$

$$\hat{\lambda}^1_t = \frac{R^P_* - R_*}{R^P_*}(\hat{\varepsilon}^s_t + \hat{\varepsilon}^b_t) + \frac{R_*}{R^P_*}(\hat{R}_t + E_t[(\hat{\lambda}^1_{t+1} - \hat{\pi}_{t+1} - \gamma_C \hat{z}_{t+1}]) \tag{39}$$

$$\hat{\lambda}^1_t = E_t\left[\hat{\lambda}^1_{t+1} - \gamma_C \hat{z}_{t+1} + \hat{R}_t - \hat{\pi}_{t+1}\right] \tag{40}$$

$$\hat{\lambda}^1_t = \left(\ln \varepsilon^i_t + \hat{\lambda}^2_t\right) - S''\left(\hat{\imath}_t - \hat{\imath}_{t-1}\right) + \beta e^{(1-\gamma_C)\gamma} S'' E_t\left(\hat{\imath}_{t+1} - \hat{\imath}_t\right) \tag{41}$$

$$\lambda^2_* \hat{\lambda}^2_t = \beta e^{-\gamma_C z_* - \omega_*}\left[\lambda^1_* u_* r^k_* E_t\left(-\gamma_C \hat{z}_{t+1} - \hat{\omega}_{t+1} + \hat{\lambda}^1_{t+1} + \hat{r}^k_{t+1} + \hat{u}_{t+1}\right)\right] + \tag{42}$$

$$+\beta e^{-\gamma_C z_* - \omega_*}\left[(1-\delta_0)\lambda^2_* E_t\left(-\gamma_C \hat{z}_{t+1} - \hat{\omega}_{t+1} + \hat{\lambda}^2_{t+1}\right) - \lambda^2_* \delta_1 u_* E_t \hat{u}_{t+1}\right]$$

$$\hat{\lambda}^1_t = \hat{\lambda}^2_t + \frac{\delta_2}{\delta_1} u_* \hat{u}_t - \hat{r}^k_t \tag{43}$$

$$\hat{k}_t = \left(1 - \frac{\varepsilon^i_* i_*}{k_*}\right)\left(\hat{k}_{t-1} - \hat{z}_t - \hat{\omega}_t\right) + \frac{\varepsilon^i_* i_*}{k_*}\left(\hat{\varepsilon}^i_t + \hat{\imath}_t\right) - \delta_1 u_* e^{-z_* - \omega_*}\hat{u}_t \tag{44}$$

$$\hat{k}^e_t = \hat{u}_t + \hat{k}_{t-1} - \hat{z}_t - \hat{\omega}_t \tag{45}$$

Capital-Labor Ratio

$$\hat{k}^e_t = \hat{w}_t - \hat{r}^k_t + \hat{h}^d_t \tag{46}$$

Real Marginal Costs

$$\widehat{mc}_t = (1-\alpha)\,\hat{w}_t + \alpha \hat{r}_t^k - \hat{\varepsilon}_t^a \tag{47}$$

The New Keynesian Phillips Curve for Inflation

$$\begin{aligned}\hat{\pi}_t = & \frac{(1-\beta\zeta_p e^{(1-\gamma_C)z_*})(1-\zeta_p)}{(1+\beta\iota_p e^{(1-\gamma_C)z_*})\zeta_p}\left[\frac{\lambda_{p,*}}{1+\lambda_{p,*}}\hat{\lambda}_{p,t} + \widehat{mc}_t\right] + \\ & + \frac{\iota_p}{1+\beta\iota_p e^{(1-\gamma_C)z_*}}\hat{\pi}_{t-1} + \frac{\beta e^{(1-\gamma_C)z_*}}{1+\beta\iota_p e^{(1-\gamma_C)z_*}}E_t\hat{\pi}_{t+1}\end{aligned} \tag{48}$$

Wage Mark-Up

$$\hat{\mu}_t^w = \hat{w}_t - \hat{w}_t^h \tag{49}$$

The New Keynesian Phillips Curve for Wages

$$\begin{aligned}\hat{w}_t = & \frac{1}{1+\beta e^{(1-\gamma_C)z_*}}\hat{w}_{t-1} + \frac{\beta e^{(1-\gamma_C)z_*}}{1+\beta e^{(1-\gamma_C)z_*}}\hat{w}_{t+1} + \frac{\beta e^{(1-\gamma_C)z_*}}{1+\beta e^{(1-\gamma_C)z_*}}(E_t\hat{\pi}_{t+1} + E_t\hat{z}_{t+1}) + \\ & \frac{\iota_w}{1+\beta e^{(1-\gamma_C)z_*}}(\hat{\pi}_{t-1} + \hat{z}_{t-1}) - \frac{1+\iota_w\beta e^{(1-\gamma_C)z_*}}{1+\beta e^{(1-\gamma_C)z_*}}(\hat{\pi}_t + \hat{z}_t) + \\ & \frac{1-\beta\zeta_w e^{(1-\gamma_C)z_*}}{1+\beta e^{(1-\gamma_C)z_*}}\frac{1-\zeta_w}{\zeta_w}\left[\frac{\lambda_{w,*}}{1+\lambda_{w,*}}\hat{\lambda}_{w,t} - \hat{\mu}_t^w\right]\end{aligned} \tag{50}$$

The Aggregate Resource Constraint

$$\frac{y_*}{g_*}(\hat{y}_t - \hat{g}_t) = \frac{c_*}{c_* + i_*}\hat{c}_t + \frac{i_*}{c_* + i_*}\hat{\imath}_t \tag{51}$$

The Production Function

$$\hat{y}_t = \frac{1}{mc_*}\left(\ln\varepsilon_t^a + \alpha\hat{k}_t^e + (1-\alpha)\,\hat{h}_t^d\right) \tag{52}$$

Labor Market Clearing Condition

$$\hat{h}_t = \hat{h}_t^d \tag{53}$$

Monetary Authority's Reaction Function

$$\hat{R}_t = \rho_R \hat{R}_{t-1} + (1-\rho_R)\left[(1-\psi_1)\,\hat{\pi}^*_t + \frac{\psi_1}{4}\left(\sum_{j=-2}^{1}\hat{\pi}_{t+j}\right) + \frac{\psi_2}{4}\left(\sum_{j=-2}^{1}\hat{y}_{t+j}\right)\right] + \sum_{j=0}^{M}\hat{\xi}_{t-j,j} \tag{54}$$

E. Measurement

E.1. National Income Accounts

The model economy's basic structure, with the representative household consuming a single good and accumulating capital using a different good, differs in some important ways from the accounting conventions of the Bureau of Economic Analysis (BEA) underlying the National Income and Product Accounts (NIPA). In particular

1. The BEA treats household purchases of long-lived goods inconsistently. It classifies purchases of residential structures as investment and treats the service flow from their stock as part of Personal Consumption Expenditures (PCE) on services. The BEA classifies households purchases of all other durable goods as consumption expenditures. No service flow from the stock of household durables enters measures of current consumption. In the model, all long-lived investments add to the productive capital stock.

2. In our model all government purchases are consumption. In fact government spending includes investment goods purchased on behalf of the populace. In the model, these should be treated as additions to the single stock of productive capital.

3. The BEA sums PCE and private expenditures on productive capital (Business Fixed Investment and Residential Investment), with government spending, inventory investment, and net exports to create Gross Domestic Product. The model features only the first three of these.

To bridge these differences, we create four *model consistent* NIPA measures from the BEA NIPA data.

1. Model-consistent GDP. Since the model's capital stock includes both the stock of household durable goods and the stock of government-purchased capital, a model-consistent GDP series should include the value of both stocks' service flows. To construct these, we followed a five-step procedure.

   (a) We begin by estimating a constant (by assumption) service-flow rate by dividing the nominal value of housing services from NIPA Table 2.4.5 by the beginning-of-year value of the residential housing stock from the BEA's Fixed Asset Table 1.1. We use annual data and average from 1947 through 2014. The resulting estimate is 0.096. That is, the annual value of housing services equals approximately 10 percent of the housing stock's value each year.

(b) In the second step, we estimate estimate constant (by assumption) depreciation rates for residential structures, durable goods, and government capital. We constructed these by first dividing observations of value lost to depreciation over a calendar year by the end-of-year stocks. Both variables were taken from the BEA's Fixed Asset Tables. (Table 1.1 for the stocks and Table 1.3 for the deprecation values.) We then averaged these ratios from 1947 through 2014. The resulting estimates are 0.021, 0.194, and 0.044 for the three durable stocks.

(c) In the third step, we calculated the average rates of real price depreciation for the three stocks. For this, we began with the nominal values and implicit deflators for PCE Nondurable Goods and PCE Services from NIPA Table 1.2. We used these series and the Fisher-ideal formula to produce a chain-weighted implicit deflator for PCE Nondurable Goods and Services. Then, we calculated the price for each of the three durable good's stocks in consumption units as the ratio of the implicit deflator taken from Fixed Asset Table 1.2 to this deflator. Finally, we calculated average growth rates for these series from 1947 through 2014. The resulting estimates equal 0.0029, −0.0223, and 0.0146 for residential housing, household durable goods, and government-purchased capital.

(d) The fourth combines the previous steps' calculations to estimate constant (by assumption) service-flow rates for household durable goods and government-purchased capital. To implement this, we assumed that all *three* stocks yield the same financial return along a steady-state growth path. These returns sum the per-unit service flow with the appropriately depreciated value of the initial investment. This delivers two equations in two unknowns, the two unknown service-flow rates. The resulting estimates are 0.29 and 0.12 for household durable goods and government-purchased capital.

(e) The fifth and final step uses the annual service-flow rates to calculate real and nominal service flows from the real and nominal stocks of durable goods and government-purchased capital reported in Fixed Asset Table 1.1. This delivers an annual series. Since the stocks are measured as of the end of the calendar year, we interpret these as the service flow values in the *next* year's first quarter. We create quarterly data by linearly interpolating between these values.

With these real and nominal service flow series in hand, we create nominal model-consistent GDP by summing the BEA's definition of nominal GDP with the nominal values of the two service flows. We create the analogous series for model-consistent real GDP by applying the Fisher ideal formula to the nominal values and price indices for these three components.

2. Model-consistent Investment. The nominal version of this series sums nominal Business Fixed Investment, Residential Investment, PCE Durable Goods, and government investment expenditures. The first three of these come from NIPA Table 1.1.5, while government investment expenditures sums Federal Defense, Federal Nondefense, and

State and Local expenditures from NIPA Table 1.5.5. We construct the analogous series for real Model-consistent Investment by combining these series with their real chain-weighted counterparts found in NIPA Tables 1.1.3 and 1.5.3 using the Fisher ideal formula. By construction, this produces an implicit deflator for Model-consistent investment as well.

3. Model-consistent Consumption. The nominal version of this series sums nominal PCE Nondurable Goods, PCE Services, and the series for nominal services from the durable goods stock. The first two of these come from NIPA Table 1.1.5. We construct the analogous series for real Model-consistent consumption by combining these series with their real chain-weighted counterparts using the Fisher ideal formula. The two real PCE series come from NIPA Table 1.1.3. Again, this produces an implicit deflator for Model-consistent consumption as a by-product.

4. Model-consistent Government Purchases. Conceptually, the model's measure of Government Purchases includes all expenditures not otherwise classified as Investment or Consumption: Inventory Investment, Net Exports, and actual Government Purchases. We construct the nominal version of this series simply by subtracting nominal Model-consistent Investment and Consumption from nominal Model-consistent GDP. We calculate the analogous real series using "chain subtraction." This applies the Fisher ideal formula to Model-consistent GDP and the *negatives* of Model-consistent Consumption and Investment.

Our empirical analysis requires us to compare model-consistent series measured from the NIPA data with their counterparts from the model's solution. To do this, we begin by solving the log-linearized system above, and then we feed the model specific paths for all exogenous shocks starting from a particular initial condition. for a given such simulation, the growth rates of Model-consistent Consumption and Investment equal

$$\begin{aligned}\Delta \ln C_t^{obs} &= z_* + \Delta \hat{c}_t + z_t \text{ and} \\ \Delta \ln I_t^{obs} &= z_* + \omega_* + \Delta \hat{i}_t + z_t + \omega_t\end{aligned}$$

The measurement of GDP growth in the model is substantially more complicated, because the variables Y_t and y_t denote model output *in consumption units.* In contrast, we mimic the BEA by using a chain-weighted Fisher ideal index to measure model-consistent GDP. Therefore, we construct an analogus chain-weighted GDP index from model data. Since such an ideal index is invariant to the units with which nominal prices are measured, we can normalize the price of consumption to equal one and employ the prices of investment goods and government purchases relative to current consumption. Our model identifies the first of these relative prices as with investment-specific technology. However, the model characterizes only government purchases *in consumption units*, because private agents do not care about their division into "real" purchases and their relative price. For this reason, we use a simple autoregression to characterize the evolution of the price of government services in consumption units. Denote this price in quarter t with P_t^g. We construct this for the US economy by dividing the Fisher-ideal price index for model-consistent government purchases by that for model-consistent consumption. Then, our model for its evolution is

$$\pi_t^{g,obs} = \ln(P_t^g/P_{t-1}^t) = (1-\beta_{2,1}-\beta_{2,2})\pi_g^* + \beta_{2,1}\ln(P_{t-1}^g/P_{t-2}^g) + \beta_{2,2}\ln(P_{t-2}^g/P_{t-3}^g) + u_t^g. \quad (55)$$

Here, $u_t^g \sim \mathbb{N}(0, \sigma_g^2)$. Given an arbitrary normalization of P_t^g to one for some time period, simulations from (55) can be used to construct simulated values of P_t^g for all other time periods. With these and a simulation from the model of all other variables in hand, we can calculate the simulation's values for Fisher ideal GDP growth using

$$\frac{Q_t}{Q_{t-1}} \equiv \sqrt{\dot{Q}_t^P \dot{Q}_t^L}, \quad (56)$$

where the Paasche and Laspeyres indices of quantity growth are

$$\dot{Q}_t^P \equiv \frac{C_t + P_t^I I_t + P_t^G(G_t/P_t^G)}{C_{t-1} + P_t^I I_{t-1} + P_t^G(G_{t-1}/P_{t-1}^G)} \text{ and} \quad (57)$$

$$\dot{Q}_t^L \equiv \frac{C_t + P_{t-1}^I I_t + P_{t-1}^G(G_t/P_t^G)}{C_{t-1} + P_{t-1}^I I_{t-1} + P_{t-1}^G(G_{t-1}/P_{t-1}^G)}. \quad (58)$$

In both (57) and (58), P_t^I is the relative price of investment to consumption. In equilibrium, this always equals A_t^I.

The above gives a complete recipe for *simulating* the growth of model-consistent real GDP growth. However, we also embody its insights into our estimation with a log-linear approximation. For this, we start by removing stochastic trends from all variables in (57) and (58), and we proceed by taking a log-linear approximation of the resulting expression. Details are available from the authors upon request.

E.1.1. Output Growth Expectations

We also discipline our model's inferences about the state of the economy by comparing expectations of one- to four-quarter ahead real GDP growth from the Survey of Professional Forecasters with the analogous expectations from our model. The Survey of Professional Forecasters did not report these expectations prior to 2007, so we use them only in the second sample. As discussed in previous section, the quarterly per-capita model-consistent real GDP growth ($\Delta \ln Q_t$) does not map one-to-one with the SPF forecast of the BEA annual real GDP growth ($\Delta \ln Y_t^{BEA}$). So we transform the former into the latter by adding back population growth to the per-capita model-consistent real GDP growth and by fitting a linear regression model of BEA real GDP growth on model-consistent real GDP growth over the sample 1993:Q1-2016Q4. In particular, we estimate the following model

$$\Delta \ln Y_t^{BEA} = \underbrace{a}_{-0.14} + \underbrace{b}_{1.06}[4 \times (\Delta \ln Q_t^{obs} + pop_t)] \quad R^2 = 0.996$$

When we bridge model and SPF forecasts, we allow these two sets of expectations to differ from each other by including serially correlated measurement errors. The observation equations are

$$\Delta \ln Y_t^{l,obs} = a + 4b(\Delta \ln Q_t^{l,obs} + pop_t^l), \quad l = 1, 2, 3, 4;$$

and we assume that population forecast is at 1 percent at annual rate throughout. The measurement errors follow mutually-independent first-order autoregressive processes.

E.2. Hours Worked Measurement

Empirical work using DSGE models like our own typically measure labor input with hours worked per capita, constructed directly from BLS measures of hours worked and the civilian non-institutional population over age 16. However, this measure corresponds poorly with business cycle models because it contains underlying low frequency variation. This fact led us to construct a new measure of hours for the model using labor market trends produced for the FRB/US model and for the Chicago Fed's in-house labor market analysis.

We begin with a multiplicative decomposition of hours worked per capita into hours per worker, the employment rate of those in the labor force, and the labor-force participation rate. The BLS provides CPS-based measures of the last two rates for the US as a whole. However, its measure of hours per worker comes from the Establishment Survey and covers only the private business sector. If we use hours per worker in the business sector to approximate hours per worker in the economy as a whole, then we can measure hours per capita as

$$\frac{H_t}{P_t} = \frac{H_t^E}{E_t^E}\frac{E_t^C}{L_t^C}\frac{L_t^C}{P_t^C}.$$

Here, H_t and P_t equal total hours worked and the total population, H_t^E/E_t^E equals hours per worker measured with the Establishment survey, E_t^C/L_t^C equals one minus the CPS based unemployment rate, and L_t^C/P_t^C equals the CPS based labor-force participation rate. Our measure of *model-relevant* hours worked deflates each component on the right-hand side by an exogenously measured trend. The trend for the unemployment rate comes from the Chicago Fed's Microeconomics team, while those for hours per worker and labor-force participation come from the FRB/US model files.

E.3. Inflation

Our empirical analysis compares model predictions of price inflation, wage inflation, inflation in the price of investment goods relative to consumption goods, and inflation expectations with their observed values from the U.S. economy. We describe our implementations of these comparisons sequentially below.

E.3.1. Price Inflation

Our model directly characterizes the inflation rate for Model-consistent Consumption. In principle, this is close to the FOMC's preferred inflation rate, that for the implicit deflator of PCE. However, in practice the match between the two inflation rates is poor. In the data, short-run movements in food and energy prices substantially influences the short-run evolution of PCE inflation. Our model lacks such a volatile sector, so if we ask it to match observed short-run inflation dynamics, it will attribute those to transitory shocks to intermediate goods' producers' desired markups driven by λ_t^p.

To avoid this outcome, we adopt a different strategy for matching model and data inflation rates, which follows that of Justiniano, Primiceri, and Tambalotti (2013). This relates three observable inflation rates – core CPI inflation, core PCE inflation, and market-based PCE inflation – to Model-consistent consumption inflation using auxiliary observation equations. For core PCE inflation, this equation is

$$\pi_t^{1,obs} = \pi_* + \pi_*^1 + \beta^{\pi,1}\hat{\pi}_t + \gamma^{\pi,j}\pi_t^{d,obs} + u_t^{\pi,1}, \tag{59}$$

In (59) as elsewhere, π_* equals the long-run inflation rate. The constant π_*^1 is an adjustment to this long-run inflation rate which accounts for possible long-run differences between realized inflation and the FOMC's goal of $\pi_\star$ (for PCE inflation π_*^1 is set to zero). The right-hand side's inflation rates, $\hat{\pi}_t$ and $\pi_t^{d,obs}$ equal Model-consistent consumption inflation and PCE Durables inflation. We refer to the coefficients multiplying them, $\beta^{\pi,1}$ and $\gamma^{\pi,1}$, as the *inflation loadings*. We include PCE Durables inflation on the right-hand side of (59) because the principle adjustment required to transform Model-consistent inflation into core PCE inflation is the replacement of the price index for durable goods services with that for durable goods purchases. The disturbance term $u_t^{\pi,1}$ follows a zero-mean first-order autoregressive process.

The other two observed inflation measures, market-based PCE inflation and core CPI inflation, have identically specified observation equations. We use 2 and 3 in superscripts to denote these equations parameters and error terms, and we use the same expressions as subscripts to denote the parameters governing the evolution of their error terms. We assume that the error terms $u_t^{\pi,1}$, $u_t^{\pi,2}$, and $u_t^{\pi,3}$ are independent of each other at all leads and lags.

To produce forecasts of inflation with these these three observation equations, we must forecast their right-hand side variables. The model itself gives forecasts of $\hat{\pi}_t$. The forecasts of durable goods inflation come from a second-order autoregression.

$$\pi_t^{d,obs} = (1 - \beta_{1,1} - \beta_{1,2})\pi_*^d + \beta_{1,1}\pi_{t-1}^{d,obs} + \beta_{1,2}\pi_{t-2}^{d,obs} + u_t^d \tag{60}$$

Its innovation is normally distributed and serially uncorrelated.

E.3.2. Wage Inflation

Although observed wage inflation does not feature the same short-run variability as does price inflation, it does include the influences of persistent demographic labor-market trends which we removed ex ante from our measure of hours worked. Therefore, we follow the same general strategy of relating observed measures of wage inflation to the model's predicted wage inflation with a error-augmented observation equation. For this, we employ two measures of compensation per hour, Earnings per Hour and Total Compensation per Hour. In parallel with our notation for inflation measures, we use 1 and 2 to denote these two wage measures of wage inflation. The observation equation for Earnings per Hour is

$$\Delta \ln w_t^{1,obs} = z_* + w_*^j + \beta^{w,1}\left(\hat{w}_t - \hat{w}_{t-1} + \hat{z}_t\right) + u_t^{w,1}, \tag{61}$$

where "Δ" is the first difference operator. Just as with the price inflation measurement errors, $u_t^{w,1}$ follows a zero-mean first-order autoregressive process. The observation equation for Total Compensation per Hour is analogous to (61).

E.3.3. Relative Price Inflation

To empirically ground investment-specific technological change in the model, we use an error-augmented observation equation to relate the relative price of investment to consumption, both model-consistent measures constructed from NIPA and Fixed Asset tables as described above, with the model's growth rate of the rate of technological transformation between these two goods, ω_t.

$$\pi_t^{i,obs} = \omega_* + \hat{\omega}_t + u_t^{c/i};$$

Here, $\pi_t^{i,obs}$ denotes the price of *c*onsumption relative to *i*nvestment. The measurement error $u_t^{c/i}$ follows a i.i.d. zero-mean normally-distributed innovation.

We also discipline our model's inferences about the state of the economy by comparing expectations of one- to four-quarter ahead and 10-year inflation from the Survey of Professional Forecasters with the analogous expectations from our model. Just as with all of the other inflation measures, we allow these two sets of expectations to differ from each other by including serially correlated measurement errors. The observation equations are

$$\begin{aligned}
\pi_t^{l,j,obs} &= \pi_* + \pi_*^{l,j} + \beta^{l,j} E_t \hat{\pi}_{t+l} + u_t^{l,j,\pi}, \; j = 1,2, \; l = 1, ...4; \\
\pi_t^{l,j,obs} &= \pi_* + \pi_*^{l,j} + \frac{\beta^{l,j}}{l} \sum_{i=1}^{l} E_t \hat{\pi}_{t+i} + u_t^{l,j,\pi}, \; j = 1,2, \; l = 40;
\end{aligned}$$

The measurement errors follow mutually-independent first-order autoregressive processes.

E.4. Interest Rates and Monetary Policy Shocks

Since our model features forward guidance shocks, it has non-trivial implications for the current policy rate as well as for expected future policy rates. To discipline the parameters governing their realizations, the elements of Σ_ε, using data, we compare the model's monetary policy shocks to high-frequency interest-rate innovations informed by event studies, such as that of Gürkaynak, Sack, and Swanson (2005). Those authors applied a factor structure to innovations in implied expected interest rates from futures prices around FOMC policy announcement dates. Specifically, they show that the vector of M implied interest rate changes following an FOMC policy announcement, Δr_t, can be written as

$$\Delta r_t = \Lambda f_t + \eta_t$$

Where f is a 2×1 vector of factors, Λ is a $H \times 2$ matrix of factor loadings, and η is an $H \times 1$ vector of mutually independent shocks. Denoting the 2×2 diagonal variance covariance matrix of f with Σ_f and the $H \times H$ diagonal variance-covariance matrix of η with Ψ, we can express the observed variance-covariance matrix of Δr as $\Lambda \Sigma_f \Lambda' + \Psi$.

Our model has implications for this same variance covariance matrix. For this, use the model's solution to express the changes in current and future expected interest rates following monetary policy shocks as $\Delta r = \Gamma_1 \varepsilon^R$. Here, ε_t^R is the vector which collects the current monetary policy shock with $M-1$ forward guidance shocks, and Γ_1 is an $H \times H$ matrix. In general, Γ_1 does *not* simply equal the identity matrix, because current and future inflation and output gaps respond to the monetary policy shocks and thereby influence future monetary policy "indirectly" through the interest rate rule.

We assume that a factor structure determines the cross-correlations among monetary policy shocks. Specifically, we assume

$$\varepsilon_{R,t}^j = \alpha_j f_t^\alpha + \beta_j f_t^\beta + \eta_t^j,$$

where the factors f_t^α and f_t^β and factor loadings α_i and β_i are scalars, η_t^j is a measurement error. The factors and shocks have zero means and are independent and normally distributed. In matrix notation, we have

$$\varepsilon_t^R = \boldsymbol{\alpha} f_t^\alpha + \boldsymbol{\beta} f_t^\beta + \eta_t,$$

where $\boldsymbol{\alpha} = [\alpha_0, \ldots, \alpha_H]'$, $\boldsymbol{\beta} = [\beta_0, \ldots, \beta_H]'$. Let $\Sigma_\eta = E\left(\eta_t \eta_t'\right)$ denote the variance-covariance matrix of the idiosyncratic shocks, and σ_α^2 (σ_β^2) denote the variance of f_t^α (f_t^β). Therefore we have that

$$\Lambda \Sigma_f \Lambda' + \Psi = \Gamma_1 (\boldsymbol{\alpha}\boldsymbol{\alpha}' \sigma_\alpha^2 + \boldsymbol{\beta}\boldsymbol{\beta}' \sigma_\beta^2) \Gamma_1' + \Gamma_1 \Sigma_\eta \Gamma_1'$$

E.5. Measurement Equations Synthesis

To summarize the measurement equations are as follows:

$$\begin{aligned}
\Delta \ln Q_t^{obs} &= f\left(\hat{c}_t, \hat{c}_{t-1}, \hat{i}_t, \hat{i}_{t-1}, \hat{g}_t, \hat{\omega}_t, \hat{\pi}_t^{g,obs}\right) \equiv \Delta \ln Q_t^j; \\
\Delta \ln Y_t^{l,obs} &= a + 4b(\Delta \ln Q_t^l + pop_t^l), \quad l = 1,2,3,4; \\
\Delta \ln C_t^{obs} &= z_* + \Delta \hat{c}_t + \hat{z}_t; \\
\Delta \ln I_t^{obs} &= z_* + \omega_* + \Delta \hat{\imath}_t + \hat{z}_t + \hat{\omega}_t; \\
\log H_t^{obs} &= \hat{H}_t; \\
\pi_t^{i,obs} &= \omega_* + \hat{\omega}_t + u_t^i; \\
R_t^{obs} &= R_* + \hat{R}_t;
\end{aligned}$$

$$
\begin{aligned}
R_t^{j,obs} &= R_* + E_t\hat{R}_{t+j}, \;\; j = 1, 2, \ldots, H; \\
\pi_t^{l,j,obs} &= \pi_* + \pi_*^{l,j} + \beta^{l,j} E_t\hat{\pi}_{t+l} + u_t^{l,j,\pi}, \;\; j = 1, 2, \;\; l = 1, ...4; \\
\pi_t^{l,j,obs} &= \pi_* + \pi_*^{l,j} + \frac{\beta^{l,j}}{l}\sum_{i=1}^{l} E_t\hat{\pi}_{t+i} + u_t^{l,j,\pi}, \;\; j = 1, 2, \;\; l = 40; \\
\pi_t^{j,obs} &= \pi_* + \pi_*^{j} + \beta^{\pi,j}\hat{\pi}_t + \gamma^{\pi,j}\pi_t^{d,obs} + u_t^{j,p}, \;\; \text{with } \beta^{\pi,1} = 1, j = 1, 2, 3; \\
\Delta \ln w_t^{j,obs} &= z_* + w_*^{j} + \beta^{w,j}\left(\hat{w}_t - \hat{w}_{t-1} + \hat{z}_t\right) + u_t^{j,w}, \;\; \text{with } \beta^{w,1} = 1, j = 1, 2; \\
\pi_t^{d,obs} &= (1 - \beta_{1,1} - \beta_{1,2})\pi_*^{d} + \beta_{1,1}\pi_{t-1}^{d,obs} + \beta_{1,2}\pi_{t-2}^{d,obs} + u_t^{d}; \\
\pi_t^{g,obs} &= (1 - \beta_{2,1} - \beta_{2,2})\pi_*^{g} + \beta_{2,1}\pi_{t-1}^{g,obs} + \beta_{2,2}\pi_{t-2}^{g,obs} + u_t^{g}.
\end{aligned}
$$

The left hand side variables represent data (Q denotes chain-weighted GDP). The function f in the first equation represents the linear approximation to the chain-weighted GDP formula. As previously discussed, two variables are included to complete the mapping from model to data but are not endogenous to the model. Specifically, the consumption price of government consumption plus net exports, $\pi_t^{g,obs}$, helps map model GDP to our model-consistent measure of chain-weighted GDP, and inflation in the consumption price of consumer durable goods, $\pi_t^{d,obs}$, is used to complete the mapping from model inflation to measured inflation.

The measurement equations indicate we use 21 time series to estimate the model in the first sample. In addition to the real quantities and federal funds rate that are standard in the literature our estimation includes multiple measures of wage and consumer price inflation, two measures each of average inflation expected over the next ten years and over one quarter, and $H = 4$ quarters of interest rate futures. Our second sample estimation is restricted to estimating the parameters of the stochastic process for forward guidance news with $H = 10$ plus the processes driving $\pi_t^{g,obs}$ and $\pi_t^{d,obs}$ (only the constant and the standard deviation). This estimation uses the measurement equations involving the current federal funds rate and 10 quarters of expected future policy rates plus the last two equations. We take into account the change in steady state but keep the remaining structural parameters at their first sample values. Because our estimation forces data on real activity, wages and prices to coexist with the interest rate futures data, we expect the estimation to mitigate the forward guidance puzzle. Finally, it is worth stressing that our estimation respects the ELB in the second sample. This is because we measure expected future rates in the model, the $E_t\hat{R}_{t+j}$, using the corresponding empirical futures rates, $R_t^{j,obs}$, and we use futures rates extending out 10 quarters. Finally, in the second sample we extend the use the Survey of Professional Forecasts about near term inflation expectations using the 1Q-4Q ahead CPI and PCE inflation expectations, and introduce the SPF expectations about near term real GDP growth expectations, i.e. 1Q ahead until 4Q ahead.

E.6. Data Synopsis

[24]

[24] Unless otherwise indicated all data are from Haver Analytics.

Model-Consistent Output

- The DSGE model output is the chained sum of conventional GDP with government capital services and durable goods services. This series is de-trended by population growth.

Model-Consistent Consumption

- DSGE consumption is defined as the chained sum of conventional PCE nondurable goods with PCE services and durable goods services. This series is de-trended by population growth.

Model-Consistent Investment

- Model-consistent Investment is the chained sum of durable goods purchases, fixed investment, and government investment. This series is de-trended by population growth.

Model-Consistent Residual Output Inflation

- The residual output is the chained difference of model consumption and investment from model GDP. Residual output reflects government spending and net exports.

Relative Price of Consumption to Investment

- The relative price is constructed by dividing the consumption price series and investment price series.

Deflators for Consumer Durables

- We take the log difference[25] of the PCE Durable Goods Chain Price Index for the deflators for consumer durables.

Inflation Expectations

- Our inflation expectations series are quarterly inflation expectations data from the Survey of Professional Forecasters at the Philadelphia Fed. They report inflation expectations at various horizons for both PCE and CPI measures. We use their aggregate measures of 1Q to 4Q ahead core CPI and core PCE inflation expectations, 40Q ahead average inflation expectations for CPI and PCE. The SPF did not report expectations for PCE prior to 2007, so we do not have many observations for the first sample of our data. However, we continue to include these few observations in order to initialize the kalman filter for second sample estimation. We have the full data for CPI expectations.

Real GDP Growth Expectations

[25] All log differenced series are multipled by 100.

- Our real GDP growth expectations series are annualized expectations data from the Survey of Professional Forecasters at the Philadelphia Fed. We use their BEA real GDP growth expectations from 1Q to 4Q ahead. The SPF reports these expectations throughout our sample period. We use them only in the second sample because the inflation data are only available for the second sample.

Real Wages

- We have two different measures of wages in the model - average hourly earnings and employment compensation. We take the average hourly earnings and divide by the chain price index of core PCE, then take the log difference.
- We repeat the same steps to calculate employment compensation but use the employment cost index for the compensation of civilian workers.

Price Inflation

- We use three different measures of price inflation: Core PCE, Market-Based Core PCE, and Core CPI.

Hours

- We construct our hours series with the methodology as described in *Forward Guidance and Macroeconomic Outcomes Since the Financial Crisis* (Campbell et al., 2016).

Effective Federal Funds Rate

- For the first sample (1993q1-2008q3), we use the federal funds target rate observed as the average over the last month of the quarter.
- For the second sample (2008q4-2018q4), we use the federal funds target rate observed at the end of the quarter.
- We divide the series by 4 to convert to quarterly rates.

Expected Federal Funds Rate (FFR)

- From 1993Q1 to 2005Q4, our 4-quarter ahead path comes from Eurodollar futures. Eurodollar futures have expiration dates that lie about two weeks before the end of each quarter. Eurodollar rate is closely tied to expectations for the Federal Funds rates over the same period, so the Eurodollar futures rate corresponds with the Fed Funds rate at the middle of the last month of each quarter.
- Beginning with 2006Q1, our 4-quarter ahead, and later, 10-quarter ahead path comes from the Overnight Index Swaps (OIS). The OIS data are converted into a point estimate of the Fed Funds for a particular date using a Svensson term structure model. The dates of the OIS data reflect the middle of the quarter values, and we interpolate to obtain the end of quarter values.

- From 2014Q1, we began to use the expected Fed Funds from the Survey of Market Participant (SMP). The SMP correspond to the survey participants' expected Fed Funds at the end of the quarter.
- All expected FFR series are in quarterly rates.

F. Calibration and Bayesian Estimation

As we discussed, we follow a two-stage approach to the estimation of our model's parameters. In a calibration stage, we set the values of selected parameters so that the model has empirically-sensible implications for long-run averages from the U.S. economy. In this stage, we also enforce several normalizations and a judgemental restriction on one of the measurement error variances. In the second stage, we estimate the model's remaining parameters using standard Bayesian methods.

F.1. Calibration

Our calibration strategy is the same as in Campbell, Fisher, Justiniano, and Melosi (2016) except that we address the well-known evidence of secular declines in economic growth and rates of return on nominally risk free assets. We address these developments by imposing a change in steady state in 2008q4 (the choice of this date is motivated in the next subsection). Steady state GDP growth is governed by the mean growth rates of the neutral and investment-specific technologies, ν_* and ω_*. We adjust ω_* down to account for the slower decline in the relative price of investment since 2008q4. Given this change we then lower ν_* so that steady state GDP growth is reduced to 2%. To match a lower real risk-free rate of 1% we increase the steady state marginal utility of government bonds using ε^s_*.[26] These adjustments leave the other calibrated parameters unchanged but do change the steady state values of the endogenous variables and therefore the point at which the economy is log-linearized.[27]

We observe the long-run average of the following aggregates: nominal federal funds rate, labor share, government spending share, investment spending share, the capital-output ratio, real per-capita GDP growth (g_y), inflation in price of government, net exports and inventory investment relative to non-durables and services consumption, and the growth rate of the consumption-investment relative price.

- The labor share can be used to calibrate the parameter α.
- The government spending share determines s^g_*.
- The government price growth rate pins down π^g_*.

[26]The targets for steady state GDP growth and risk-free rate reflect a variety of evidence including the Fed's Summary of Economic Projections.

[27]Our re-calibration changes the return on private assets by a little. This small change is consistent with Yi and Zhang (2017) who show that rates of return on private capital have stayed roughly constant in the face of declines in risk free rates.

- The growth rate of the consumption-investment relative price pins down ω_*.
- The investment share pins down i_*/y_*.
- The capital output ratio pins down k_*/y_*.
- Calculate the consumption-output share

$$\frac{c_*}{y_*} = \left(1 - \frac{i_*}{y_*} - \frac{g_*}{y_*}\right). \tag{62}$$

- The growth rate of real chain-weighted GDP is used to pin down the growth rate of the common trend z_*. First

$$g_y = e^{z_*}\sqrt{\frac{\frac{c_*}{y_*} + e^{\omega}\frac{i_*}{y_*} + \left(\pi_*^g\right)^{-1}\frac{g_*}{y_*}}{\frac{c_*}{y_*} + e^{-\omega}\frac{i_*}{y_*} + \pi_*^g\frac{g_*}{y_*}}}$$

All the variables in this equation are known except for z_*. So we can solve for z_*:

$$z_* \quad = \quad g_y - \frac{1}{2}\ln\left(\frac{\frac{c_*}{y_*} + e^{\omega}\frac{i_*}{y_*} + \left(\pi_*^g\right)^{-1}\frac{g_*}{y_*}}{\frac{c_*}{y_*} + e^{-\omega}\frac{i_*}{y_*} + \pi_*^g\frac{g_*}{y_*}}\right) \tag{63}$$

- The growth rate of the labor-augmenting technology ν_* can be easily obtained by exploiting the following equation:

$$z_* = v_* + \frac{\alpha}{1-\alpha}\omega_*. \tag{64}$$

- We are now in a position to identify the depreciation rate δ_0 using the steady-state equation pinning down the investment capital ratio:

$$\begin{aligned}\frac{i_*}{k_*} \quad &= \quad 1 - (1-\delta_0)e^{-z_*-\omega_*} \\ &\Rightarrow \quad \delta_0 = 1 + \left(\frac{i_*}{k_*} - 1\right)e^{z_*+\omega_*}\end{aligned}$$

where the investment capital ratio is obtained combining the investment share and the capital output ratio:

$$\frac{i_*}{k_*} = \frac{i_*/y_*}{k_*/y_*}. \tag{65}$$

- From the steady-state equilibrium we have that

$$\frac{y_*}{k_*} = e^{-z_*-\omega_*}\frac{\delta_1}{\alpha}. \tag{66}$$

Therefore

$$\delta_1 = \alpha \left(\frac{k_*}{y_*}\right)^{-1} e^{z_* + \omega_*} \tag{67}$$

where the capital output ratio is given above.

- In steady state, the real rate of return on private bonds is derived from the first order condition for private bonds:

$$r_*^p \equiv \frac{R_*^P}{\pi_*} = \frac{e^{\gamma_c z_*}}{\beta}. \tag{68}$$

In steady state the real rental rate of capital is derived from the first order condition for capital:

$$r_*^k = \left[\frac{e^{\gamma_c z_*}}{\beta}\right] e^{\omega_*} - (1 - \delta_0) \tag{69}$$

Combining these last two equations yields

$$r_*^k = r_*^p e^{\omega_*} - (1 - \delta_0)$$

and hence

$$r_*^p = \left[r_*^k + 1 - \delta_0\right] e^{-\omega_*}.$$

Note that $r_*^k = \delta_1$ from the first order condition for capacity utilization. It follows that

$$r_*^p = (1 - \delta_0 + \delta_1)\, e^{-\omega_*}$$

- The liquidity premium in steady state (i.e., $\frac{R_*/\pi_*}{r_*^p}$) can be computed now by assuming a *nominal* average federal funds rate, R_*, and an annualized average inflation rate.

- Using equation (69) and the fact that $r_*^k = \delta_1$, we can calibrate the discount factor β :

$$\beta = (1 - \delta_0 + \delta_1)^{-1} e^{\omega_*} e^{\gamma_c z_*}$$

where γ_c is a parameter of the utility function to be estimated.

F.2. Bayesian Estimation

Our Bayesian estimation uses the same split-sample strategy as in Campbell, Fisher, Justiniano, and Melosi (2016) except that we incorporate the change in steady state described above and one other change noted below. As in Campbell, Fisher, Justiniano, and Melosi (2016) our sample begins in 1993q1. This date is based on the availability and reliability of the overnight interest rate futures data. The sample period ends in 2016q4 but we impose a sample break in 2008q4. Our choice of this latter date is motivated by three main considerations. First, there is the evidence that points to lower interest rates and economic growth later in the sample. Second, it seems clear that the horizon over which forward guidance was communicated by the Fed lengthened substantially during the ELB period. Finally, the downward trends in inflation and inflation expectations from the early 1990s appear to come to an end in the mid-2000s. Splitting the sample in 2008q4 and assuming some parameters change at that date is our way of striking a balance between parsimony and addressing the multiple structural changes that seem to occur around the same time.

We estimate the full suite of non-calibrated structural parameters in the first sample under the assumption that forward guidance extends for $H = 4$ quarters. Starting in 2008q4 we assume the model environment changes in three ways. First we assume the change in the steady state described above. Second, forward guidance lengthens to $H = 10$ quarters Third, the time-varying inflation target from the first sample becomes a constant equal to the steady state rate of inflation, 2% at an annual rate. All three changes are assumed to be unanticipated and permanent.

We employ standard prior distributions, but those governing monetary policy shocks deserve further elaboration. Our estimation requires the variance-covariance matrix of monetary policy shocks to be consistent with the factor-structure of interest rate innovations used by Gürkaynak, Sack, and Swanson (2005), as described above. Therefore, we parameterize Σ_ε in terms of factors STD (σ_α and σ_β), factor loadings (α and β) and STD of the idiosyncratic errors ($\sigma_{\eta,j}$). We then center our priors for these parameters at their estimates from event-studies. However, we do not require our estimates to equal their prior values. Our Bayesian estimation procedure employs quarterly data on expected future interest rates, the posterior likelihood function includes them as free parameters. It is well known that factors STD and loadings are not separately identified, so we impose two scale normalizations and one rotation normalization on α and β. The rotation normalization requires that the first factor, which we label "Factor A", is the only factor influence the current policy rate. That is, the second factor, "Factor B" influences only future policy rates. Gürkaynak, Sack, and Swanson (2005) call Factors A and B the "target" and "path" factors.

F.3. Posterior Estimates

We report the results of our two-stage two-sample estimation in a series of tables. Table 2 reports our most notable calibration targets. The long-run policy rate equals 1.1 percent on a quarterly basis. We target a two percent growth rate of per capita GDP. Given an average population growth rate of one percent per year, this implies that our potential GDP growth rate equals three percent. The other empirical moments we target are a nominal investment to output ratio of 26 percent and nominal government purchases to output ratio of 15 percent. Finally, we target a capital to output ratio of approximately 10 on a quarterly basis.

Table 3 lists the parameters which we calibrate along with their given values. The table includes many more parameters than there are targets in Table 2. This is because Table 2 omitted calibration targets which map one-to-one with particular parameter values. For example, we calibrate the steady-state capital depreciation rate (δ_0) using standard methods applied to data from the Fixed Asset tables. It is also because Table 3 lists several parameters which are normalized prior to estimation. Most notable among these are the three factor loadings listed at the table's bottom. Tables 4 and 7 report prior distributions and posterior modes for the model's remaining parameters, for the first and second samples respectively.

Table 2: First Sample Calibration Targets

Description	Expression	Value
Fixed Interest Rate (quarterly, gross)	R^*	1.011
Per-Capita Steady-State Output Growth Rate (quarterly)	Y_{t+1}/Y_t	1.005
Investment to Output Ratio	I_t/Y_t	0.2597
Capital to Output Ratio	K_t/Y_t	10.7629
Fraction of Final Good Output Spent on Public Goods	G_t/Y_t	0.1532
Growth Rate of Relative Price of Consumption to Investment	P_C/P_I	0.371

Table 3: First Sample Calibrated Parameters

Parameter	Symbol	Value
Discount Factor	β	0.9857
Steady-State Measured TFP Growth (quarterly)	z_*	0.489
Investment-Specific Technology Growth Rate	ω_*	0.371
Elasticity of Output w.r.t Capital Services	α	0.401
Steady-State Wage Markup	λ_*^w	1.500
Steady-State Price Markup	λ_*^p	1.500
Steady-State Scale of the Economy	H_*	1.000
Steady-State Inflation Rate (quarterly)	π_*	0.500
Steady-State Depreciation Rate	δ_0	0.0162
Steady-State Marginal Depreciation Cost	δ_1	0.0385
Core PCE, 1Q Ahead and 10Y Ahead Expected PCE		
Constant	$\pi_*^1, \pi_*^{l,1}$	0.000
Loading 1	$\beta^{\pi,1}, \beta^{l,1}$	1.000
Core CPI, 1Q Ahead and 10Y Ahead Expected CPI		
Constant	$\pi_*^2, \pi_*^{l,2}$	0.122
10Y Ahead Expected CPI and PCE		
Standard Deviation of $u_t^{40,j,\pi}$		0.010
PCE Durable Goods Inflation		
1st Lag Coefficient	$\beta_{1,1}$	0.418
2nd Lag Coefficient	$\beta_{1,2}$	0.379
Inflation in Relative Price of Government, Inventories and Net Exports to Consumption		
1st Lag Coefficient	$\beta_{2,1}$	0.311
2nd Lag Coefficient	$\beta_{2,2}$	0.0057
Compensation		
Constant	w_*^1	-0.202
Loading	$\beta^{w,1}$	1.000
Earnings Constant	w_*^2	-0.237
Loading 0 Factor A	α_0	0.981
Loading 0 Factor B	β_0	0.000
Loading 4 Factor B	β_4	0.951

Table 4: First Sample Estimated Parameters

Parameter	Symbol	Density	Prior Mean	Std.Dev	Posterior Mode
Depreciation Curve	$\frac{\delta_2}{\delta_1}$	G	1.0000	0.150	0.474
Active Price Indexation Rate	ι_p	B	0.5000	0.150	0.409
Active Wage Indexation Rate	ι_w	B	0.5000	0.150	0.077
External Habit Weight	λ	B	0.7500	0.025	0.780
Labor Supply Elasticity	γ_H	N	0.6000	0.050	0.589
Price Stickiness Probability	ζ_p	B	0.8000	0.050	0.831
Wage Stickiness Probability	ζ_w	B	0.7500	0.050	0.914
Adjustment Cost of Investment	φ	G	3.0000	0.750	5.354
Elasticity of Intertemporal Substitution	γ_c	N	1.5000	0.375	1.319
Interest Rate Response to Inflation	ψ_1	G	1.7000	0.150	1.791
Interest Rate Response to Output	ψ_2	G	0.2500	0.100	0.398
Interest Rate Smoothing Coefficient	ρ_R	B	0.8000	0.100	0.801
Autoregressive Coefficients of Shocks					
Discount Factor	ρ_b	B	0.5000	0.250	0.813
Inflation Drift	ρ_π	B	0.9900	0.010	0.998
Exogenous Spending	ρ_g	B	0.6000	0.100	0.887
Investment-Demand	ρ_i	B	0.5000	0.100	0.791
Liquidity Preference	ρ_s	B	0.6000	0.200	0.887
Price Markup	ρ_{λ_p}	B	0.6000	0.200	0.136
Wage Markup	ρ_{λ_w}	B	0.5000	0.150	0.469
Neutral Technology	ρ_ν	B	0.3000	0.150	0.492
Investment Specific Technology	ρ_ω	B	0.3500	0.100	0.303
Moving Average Coefficients of Shocks					
Price Markup	θ_{λ_p}	B	0.4000	0.200	0.307
Wage Markup	θ_{λ_w}	B	0.4000	0.200	0.391
Standard Deviations of Innovations					
Discount Factor	σ_b	U	0.5000	2.000	1.768
Inflation Drift	σ_π	I	0.0150	0.0075	0.077
Exogenous Spending	σ_g	U	1.0000	2.000	4.139
Investment-Demand	σ_i	I	0.2000	0.200	0.549
Liquidity Preference	σ_s	U	0.5000	2.000	0.341
Price Markup	σ_{λ_p}	I	0.1000	1.000	0.101
Wage Markup	σ_{λ_w}	I	0.1000	1.000	0.035
Neutral Technology	σ_ν	U	0.5000	0.250	0.530
Investment Specific Technology	σ_ω	I	0.2000	0.100	0.259
Relative Price of Cons to Inv	$\sigma_{\frac{c}{i}}$	I	0.0500	2.000	0.675
Monetary Policy					
Unanticipated	σ_{η_0}	N	0.0050	0.0025	0.012

Notes: Distributions (**N**) Normal, (**G**) Gamma, (**B**) Beta, (**I**) Inverse-gamma-1, (**U**) Uniform

First Sample Estimated Parameters (Continued)

Parameter	Symbol	Density	Prior Mean	Prior Std.Dev	Posterior Mode
1Q Ahead	σ_{η_1}	N	0.0050	0.0025	0.012
2Q Ahead	σ_{η_2}	N	0.0050	0.0025	0.008
3Q Ahead	σ_{η_3}	N	0.0050	0.0025	0.009
4Q Ahead	σ_{η_4}	N	0.0050	0.0025	0.012
Compensation					
Standard Deviation of $u_t^{1,w}$		I	0.0500	0.100	0.194
AR(1) Coefficient of $u_t^{1,w}$		B	0.4000	0.100	0.458
Earnings					
Loading 1	$\beta^{w,2}$	N	0.8000	0.100	0.904
Standard Deviation of $u_t^{2,w}$		I	0.0500	0.100	0.143
AR(1) Coefficient of $u_t^{2,w}$		B	0.4000	0.100	0.674
Core PCE					
Loading 2	$\gamma^{\pi,1}$	N	0.0000	1.000	0.045
Standard Deviation of $u_t^{1,p}$		I	0.0500	0.100	0.046
AR(1) Coefficient of $u_t^{1,p}$		B	0.2000	0.100	0.108
Core CPI					
Loading 1	$\beta^{\pi,2}$	N	1.0000	0.100	0.808
Loading 2	$\gamma^{\pi,2}$	N	0.0000	1.000	0.087
Standard Deviation of $u_t^{2,p}$		I	0.1000	0.100	0.077
AR(1) Coefficient of $u_t^{2,p}$		B	0.4000	0.200	0.586
Market-Based Core PCE					
Constant	π_*^3	N	-0.1000	0.100	-0.037
Loading 1	$\beta^{\pi,3}$	N	1.0000	0.100	1.121
Loading 2	$\gamma^{\pi,3}$	N	0.0000	1.000	0.015
Standard Deviation of $u_t^{3,p}$		I	0.0500	0.100	0.035
AR(1) Coefficient of $u_t^{3,p}$		B	0.2000	0.100	0.144
1Q Ahead Expected PCE					
Standard Deviation of $u_t^{1,1,\pi}$		I	0.0500	0.100	0.026
AR(1) Coefficient of $u_t^{1,1,\pi}$		B	0.2000	0.100	0.196
1Q Ahead Expected CPI					
Loading	$\beta^{1,2}$	N	1.0000	0.100	0.980
Standard Deviation of $u_t^{1,2,\pi}$		I	0.0500	0.100	0.062
AR(1) Coefficient of $u_t^{1,2,\pi}$		B	0.2000	0.100	0.198
10Y Ahead Expected PCE					
AR(1) Coefficient of $u_t^{40,1,\pi}$		B	0.2000	0.100	0.271
10Y Ahead Expected CPI					
Loading	$\beta^{40,2}$	N	1.0000	0.100	1.021
AR(1) Coefficient of $u_t^{40,2,\pi}$		B	0.2000	0.100	0.213

Notes: Distributions (**N**) Normal, (**G**) Gamma, (**B**) Beta, (**I**) Inverse-gamma-1, (**U**) Uniform

First Sample Estimated Parameters (Continued)

Parameter	**Symbol**	**Density**	**Prior Mean**	**Std.Dev**	**Posterior Mode**
PCE Durable Goods Inflation					
Constant	π_*^d	N	-0.3500	0.100	-0.360
Standard Deviation of u_t^d		I	0.2000	2.000	0.286

Notes: Distributions (**N**) Normal, (**G**) Gamma, (**B**) Beta, (**I**) Inverse-gamma-1, (**U**) Uniform

First Sample Estimated Parameters (Continued)

Parameter	Symbol	Density	Prior Mean	Std.Dev	Posterior Mode
Inflation in Relative Price of Government,					
Inventories and Net Exports to Consumption					
Constant	π^g_*	N	0.1980	1.000	-0.666
Standard Deviation of u^g_t		I	0.5000	2.000	1.861
Factor A					
Loading 1	α_1	N	0.6839	0.200	1.305
Loading 2	α_2	N	0.5224	0.200	0.877
Loading 3	α_3	N	0.4314	0.200	0.306
Loading 4	α_4	N	0.3243	0.200	-0.012
Standard Deviation	σ_α	N	0.1000	0.0750	0.040
Factor B					
Loading 1	β_1	N	0.3310	0.200	0.656
Loading 2	β_2	N	0.6525	0.200	1.104
Loading 3	β_3	N	0.8059	0.200	1.162
Standard Deviation	σ_β	N	0.1000	0.0750	0.078

Notes: Distributions (**N**) Normal, (**G**) Gamma, (**B**) Beta, (**I**) Inverse-gamma-1, (**U**) Uniform

Table 5: Second Sample Calibration Targets (Different from First Sample)

Description	Expression	Value
Fixed Interest Rate (quarterly, gross)	R^*	1.007
Per-Capita Steady-State Output Growth Rate (quarterly)	Y_{t+1}/Y_t	1.003
Growth Rate of Relative Price of Consumption to Investment	P_C/P_I	0.171

Table 6: Second Sample Calibrated Parameters (Different from First Sample)

Parameter	Symbol	Value
Steady-State Measured TFP Growth (quarterly)	z_*	0.489
Investment-Specific Technology Growth Rate	ω_*	0.171
Steady-State Marginal Depreciation Cost	δ_1	0.038
Core CPI, 1Q Ahead and 10Y Ahead Expected CPI		
Constant	$\pi_*^2, \pi_*^{l,2}$	0.122
10Y Ahead Expected CPI and PCE		
Standard Deviation of $u_t^{40,j,\pi}$		0.020
PCE Durable Goods Inflation		
1st Lag Coefficient	$\beta_{1,1}$	0.000
2nd Lag Coefficient	$\beta_{1,2}$	0.000
Inflation in Relative Price of Government, Inventories and Net Exports to Consumption		
1st Lag Coefficient	$\beta_{2,1}$	0.320
2nd Lag Coefficient	$\beta_{2,2}$	-0.240
Compensation Loading	$\beta^{w,1}$	1.000
Loading 5 Factor A	α_5	0.932
Loading 8 Factor B	β_8	0.210
Loading 10 Factor B	β_{10}	0.000

Table 7: Second Sample Estimated Parameters

Parameter	Symbol	Prior Mean	Std.Dev	Posterior Mode
Compensation				
Constant	w_*^1	-0.2023	0.000	-0.2023
Standard Deviation of $u_t^{1,w}$		0.1941	0.100	0.284
AR(1) Coefficient of $u_t^{1,w}$		0.4579	0.000	0.4579
Earnings				
Constant	w_*^2	-0.2370	0.000	-0.237
Loading 1	$\beta^{w,2}$	0.9039	0.000	0.9039
Standard Deviation of $u_t^{2,w}$		0.1434	0.100	0.304
AR(1) Coefficient of $u_t^{2,w}$		0.6741	0.000	0.6741
Core PCE				
Loading 2	$\gamma^{\pi,1}$	0.0449	0.000	0.0449
Standard Deviation of $u_t^{1,p}$		0.0457	0.100	0.274
AR(1) Coefficient of $u_t^{1,p}$		0.1081	0.000	0.1801
Core CPI				
Loading 1	$\beta^{\pi,2}$	0.8083	0.00	0.8083
Loading 2	$\gamma^{\pi,2}$	0.0868	0.000	0.0868
Standard Deviation of $u_t^{2,p}$		0.0770	0.100	0.2517
AR(1) Coefficient of $u_t^{2,p}$		0.5856	0.000	0.5856
Market PCE				
Constant	π_*^3	-0.0367	0.000	-0.0367
Loading 1	$\beta^{\pi,3}$	1.1213	0.000	1.1213
Loading 2	$\gamma^{\pi,3}$	0.0153	0.000	0.0153
Standard Deviation of $u_t^{3,p}$		0.0349	0.100	0.2553
AR(1) Coefficient of $u_t^{3,p}$		0.1436	0.000	0.1436
1Q Ahead Expected PCE				
Standard Deviation of $u_t^{1,1,\pi}$		0.0259	0.020	0.0412
AR(1) Coefficient of $u_t^{1,1,\pi}$		0.1960	0.050	0.1832
2Q Ahead Expected PCE				
Standard Deviation of $u_t^{2,1,\pi}$		0.0259	0.020	0.0175
AR(1) Coefficient of $u_t^{2,1,\pi}$		0.1960	0.050	0.2140
3Q Ahead Expected PCE				
Standard Deviation of $u_t^{3,1,\pi}$		0.0259	0.020	0.0193
AR(1) Coefficient of $u_t^{3,1,\pi}$		0.1960	0.050	0.2202
4Q Ahead Expected PCE				
Standard Deviation of $u_t^{4,1,\pi}$		0.0259	0.020	0.0156
AR(1) Coefficient of $u_t^{4,1,\pi}$		0.1960	0.050	0.2075
1Q Ahead Expected CPI				

Second Sample Estimated Parameters (Continued)

Parameter	Symbol	Prior Mean	Std.Dev	Posterior Mode
Loading	$\beta^{1,2}$	0.9803	0.080	1.0022
Standard Deviation of $u_t^{1,2,\pi}$		0.0622	0.020	0.095
AR(1) Coefficient of $u_t^{1,2,\pi}$		0.1982	0.050	0.206
2Q Ahead Expected CPI				
Loading	$\beta^{1,2}$	0.9803	0.080	1.2433
Standard Deviation of $u_t^{2,2,\pi}$		0.0622	0.020	0.0411
AR(1) Coefficient of $u_t^{2,2,\pi}$		0.1982	0.050	0.2532
3Q Ahead Expected CPI				
Loading	$\beta^{1,2}$	0.9803	0.080	1.2662
Standard Deviation of $u_t^{3,2,\pi}$		0.0622	0.020	0.0399
AR(1) Coefficient of $u_t^{3,2,\pi}$		0.1982	0.050	0.2607
4Q Ahead Expected CPI				
Loading	$\beta^{1,2}$	0.9803	0.080	1.2354
Standard Deviation of $u_t^{4,2,\pi}$		0.0622	0.020	0.0406
AR(1) Coefficient of $u_t^{4,2,\pi}$		0.1982	0.050	0.2782
10Y Ahead Expected PCE				
AR(1) Coefficient of $u_t^{40,1,\pi}$		0.2711	0.000	0.2711
10Y Ahead Expected CPI				
Loading	$\beta^{40,2}$	1.0207	0.000	1.0207
AR(1) Coefficient of $u_t^{40,2,\pi}$		0.2133	0.000	0.2133
1Q Ahead Expected GDP				
Standard Deviation of $u_t^{1,1,Y}$		0.10	0.100	0.9827
AR(1) Coefficient of $u_t^{1,1,Y}$		0.20	0.100	0.1300
2Q Ahead Expected GDP				
Standard Deviation of $u_t^{2,1,Y}$		0.10	0.100	0.6263
AR(1) Coefficient of $u_t^{2,1,Y}$		0.20	0.100	0.1825
3Q Ahead Expected GDP				
Standard Deviation of $u_t^{3,1,Y}$		0.10	0.100	0.9779
AR(1) Coefficient of $u_t^{3,1,Y}$		0.20	0.100	0.1767
4Q Ahead Expected GDP				
Standard Deviation of $u_t^{4,1,Y}$		0.10	0.100	0.3664
AR(1) Coefficient of $u_t^{4,1,Y}$		0.20	0.100	0.2747
PCE Durable Goods Inflation				
Constant	π_*^d	-0.4500	0.200	-0.4858
Standard Deviation of u_t^d		0.5000	0.150	0.325
Inflation in Relative Price of Government, Inventories and Net Exports to Consumption				
Constant	π_*^g	0.8900	0.400	-0.1177

Second Sample Estimated Parameters (Continued)

Parameter	**Symbol**	**Prior Mean**	**Std.Dev**	**Posterior Mode**
Standard Deviation of u_t^g		0.8143	0.150	1.267
Factor A				
Loading 0	α_0	0.0180	0.250	0.099
Loading 1	α_1	0.0574	0.250	0.202
Loading 2	α_2	0.1941	0.250	0.397
Loading 3	α_3	0.3996	0.250	0.591
Loading 4	α_4	0.6520	0.250	0.792
Loading 6	α_6	1.2266	0.250	1.116
Loading 7	α_7	1.5237	0.250	1.281
Loading 8	α_8	1.8139	0.250	1.406
Loading 9	α_9	2.0914	0.250	1.517
Loading 10	α_{10}	2.3523	0.250	2.851
Standard Deviation	σ_α	0.0442	0.100	0.056
Factor B				
Loading 0	β_0	-0.0181	0.300	0.051
Loading 1	β_1	0.2211	0.300	0.083
Loading 2	β_2	0.3679	0.300	0.125
Loading 3	β_3	0.4424	0.300	0.152
Loading 4	β_4	0.4612	0.300	0.167
Loading 5	β_5	0.4370	0.300	0.181
Loading 6	β_6	0.3817	0.300	0.192
Loading 7	β_7	0.3032	0.300	0.203
Loading 9	β_9	0.1074	0.300	0.210
Standard Deviation	σ_β	0.0334	0.100	0.449
Standard Deviations of Monetary Policy Innovations				
Unanticipated	σ_{η_0}	0.0061	0.005	0.011
1Q Ahead	σ_{η_1}	0.0021	0.005	0.010
2Q Ahead	σ_{η_2}	0.0004	0.005	0.010
3Q Ahead	σ_{η_3}	0.0019	0.005	0.010
4Q Ahead	σ_{η_4}	0.0001	0.005	0.010
5Q Ahead	σ_{η_5}	0.0025	0.005	0.009
6Q Ahead	σ_{η_6}	0.0019	0.005	0.010
7Q Ahead	σ_{η_7}	0.0011	0.005	0.009
8Q Ahead	σ_{η_8}	0.0001	0.005	0.009
9Q Ahead	σ_{η_9}	0.0014	0.005	0.010
10Q Ahead	$\sigma_{\eta_{10}}$	0.0028	0.005	0.0001

G. Covid parameter estimates

Table 8 displays our estimates of the Covid shock's parameters and of the Covid factor, f. The first three columns show the estimated ϕ_is and λ_js. The final column shows the λ_is that replicate the perfectly anticipated propagation of the Covid corresponding to the Delta scenario. The parameters ϕ_s, ϕ_b, ϕ_i, ϕ_ν, and ϕ_p, denote the loadings of the Covid shock on the liquidity preference, discount rate, MEI, neutral technology, and price markup components of Covid, respectively. These parameters are estimated using 2020q2 data only and so their values do not change across the periods.

Table 8: Parameter Estimates

	2020Q2	2020Q3	2020Q4	2021Q3
ϕ_s	1	1	1	1
ϕ_b	0.0038	0.0038	0.0038	0.0038
ϕ_i	0.1696	0.1696	0.1696	0.1696
ϕ_ν	-0.444	-0.444	-0.444	-0.444
ϕ_p	0.0103	0.0103	0.0103	0.0103
λ_0	1	1	1	1
λ_1	-0.1298	-0.3922	-0.0109	-0.5074
λ_2	-0.1044	-0.3535	-0.1051	-0.1193
λ_3	-0.1405	-0.225	-0.2956	-0.0052
λ_4	-0.1068	0.1724	0.2786	0.1021
σ_f	11.6863	11.6737	9.7799	NA
$\hat{f}$	30.453	-11.639	-5.039	0.879

Note: The parameters ϕ_s, ϕ_b, ϕ_i, ϕ_ν, and ϕ_p, denote the loadings of the Covid shock onto its components which include the liquidity preference, discount rate, MEI, permanent neutral technology, and the cost-push shocks. The entries in the rightmost column replicate the perfect anticipation of the path of the Covid shock over the horizon of the SPF forecasts. Since this is a deterministic path the variance is not applicable. The estimates are modal values of the posterior distribution based on normally distributed loose priors. The parameters ψ_s and λ_0 are normalized to one.

The estimates of the Covid factor seem broadly consistent with the historical narrative at least as it pertains to the mandates on social distancing. They indicate that the dramatic effects in 2020q2 were somewhat reversed in the following two quarters. Whole swaths of the economy were shut down in 2020q2 and they opened back up as the year progressed.

The relative magnitudes of the ϕs do not provide a complete picture of the contributions of each wedge to the effects of the Covid shock. The λs influence the contributions as anticipated movements in the wedges influence current decisions. The sizes of the contributions also depend on the degree of amplification within the model. Still, the relatively large loadings on the liquidity preference and neutral technology wedges are broadly consistent with the widely held view that the public and private response to the pandemic had both demand and supply side aspects. A shock to the liquidity preference wedge is a demand shock — it moves output and prices in the same direction. A shock to the neutral technology wedge is a supply shock as it moves output and prices in opposite directions.

The loadings λ of the common factor onto the Covid news are re-estimated each quarter. Recall from equation (3), that news about the Covid shock j-steps ahead equals one-step-ahead forecast errors of the Covid shock. In 2020q2 the Covid shock is entirely unanticipated by construction. The negative values of the λs in that quarter indicate expectations of the Covid shock going forward were a persistent reversal from the unanticipated shock. The table shows there are similar revisions to expectations in the remaining quarters of 2020.

We can gain insight into the sign reversals by studying the one-step-ahead forecast errors and revisions to expectations for GDP growth and core PCE inflation in 2020q2 (these model forecasts are very close to the SPF forecasts). These are displayed in Figure 10. The red lines corresponds to the forecast conditioned on 2020q1 data and the black line is the forecast conditioned on 2020q2 data. The bars, which sum to the forecast error and revisions, are discussed below. Note that the forecasts for GDP growth and core PCE growth in 2020q2 are close to the SPF data in that quarter.[28] In 2020q2 output collapsed and prices fell, but forecasters expected a quick rebound of both in 2020q3. The expected rebound explains the sign reversal in our estimates of the λs in 2020q2.

H. The role of the prior in our results

The prior does not play a large role in driving the results. We look at two statistics to show this: the prior sensitivity on impulse response functions and the difference between the log posterior kernel and the log likelihood. In the first case, we run a prior sensitivity check on the responses of output after an unusual shock to see if the prior could constrain the shape of the impulse response distribution. This conditional moment is crucial to capture the extent to which the unusual shock is able to characterize the economic dynamics during the pandemic period. To this aim, we randomly draw 1000 values from the prior distribution of the unusual shock parameters and for each draw computed the implied impulse response of output. Figure 11 reports the prior-implied impulse responses of output in percentage points; in particular the gray area indicates the 90% prior confidence set and the dashed-dotted black line the min and the max across all simulations.[29] Clearly the prior is not constraining the size nor the shape of the dynamic transmission of the unusual shock to output and *a priori* the shock could capture a large array of GDP growth outcomes.

The second statistic is meant to measure the influence of the prior on the posterior mode, in particular how much the likelihood and the posterior kernel differ in a neighborhood around

[28] This is primarily because of the reduction in the measurement error on the SPF forecasts discussed in Section 4.2.

[29] Impulse responses are normalized so that the unusual shock has a negative impact on impact on output.

Figure 10: One-step-ahead forecast error decomposition in 2020q2

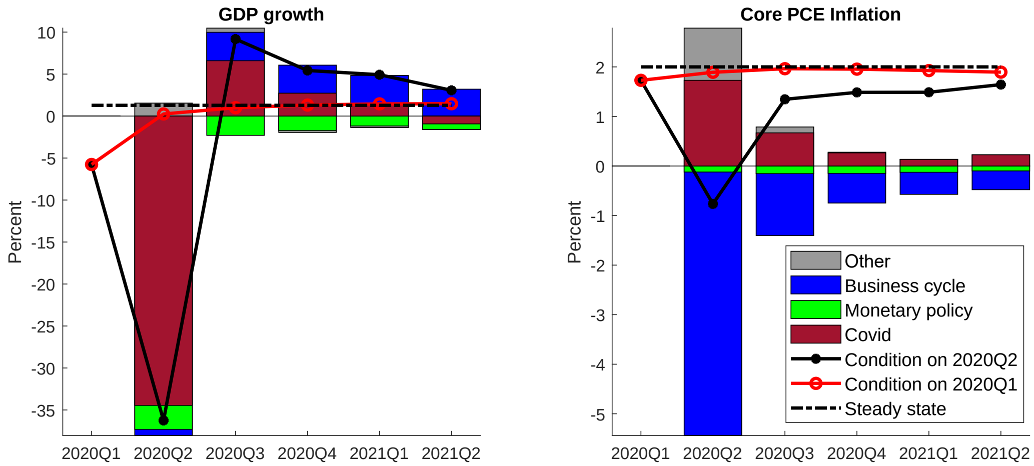

Note: This figure features the decomposition of one-step-ahead forecast error of output and inflation into the parts attributed to the Covid shock, usual business cycle shocks, surprise and news monetary policy shocks, and other shocks which include variables from outside the model used to measure model-consistent GDP and core PCE inflation. The black lines indicate 2020q2 output and inflation and SPF forecasts of these variables one to four quarters ahead. The red line is the forecast conditioned on 2020q1 data. The colored bars show the contribution to the forecast error of the indicated shocks. Source: Authors' calculations, SPF, and Haver Analytics.

the posterior mode when using 2020Q2 data. The difference between the likelihood and the posterior kernel is entirely due to the prior. If the latter is not influencing the results, this difference should be roughly constant in a neighborhood around the posterior mode. If this is not the case and if we see the posterior kernel and the likelihood decoupling, then we should worry about the prior having an influence on the posterior mode. To check that this was not the case, we randomly draw the unusual shock parameter values from a normal distribution centered on the posterior kernel mode with the covariance given by the inverse of the Hessian evaluated at the mode. For each draw, we computed the log posterior kernel value in deviation from the posterior mode value and the log likelihood value in deviation from the likelihood mode value and took the difference between these two statistics. In this way we removed the contribution of the prior at the mode and isolated the relative importance of the prior in a neighborhood of it. In figure 12, we report this statistic across two hundred draws and find that the difference between posterior and likelihood is strictly less than three log points in absolute terms suggesting an insignificant role of the prior in influencing the posterior.

Figure 11: Prior IRFs. 90 % percent bands and min and max. Number of draws form the prior 1000.

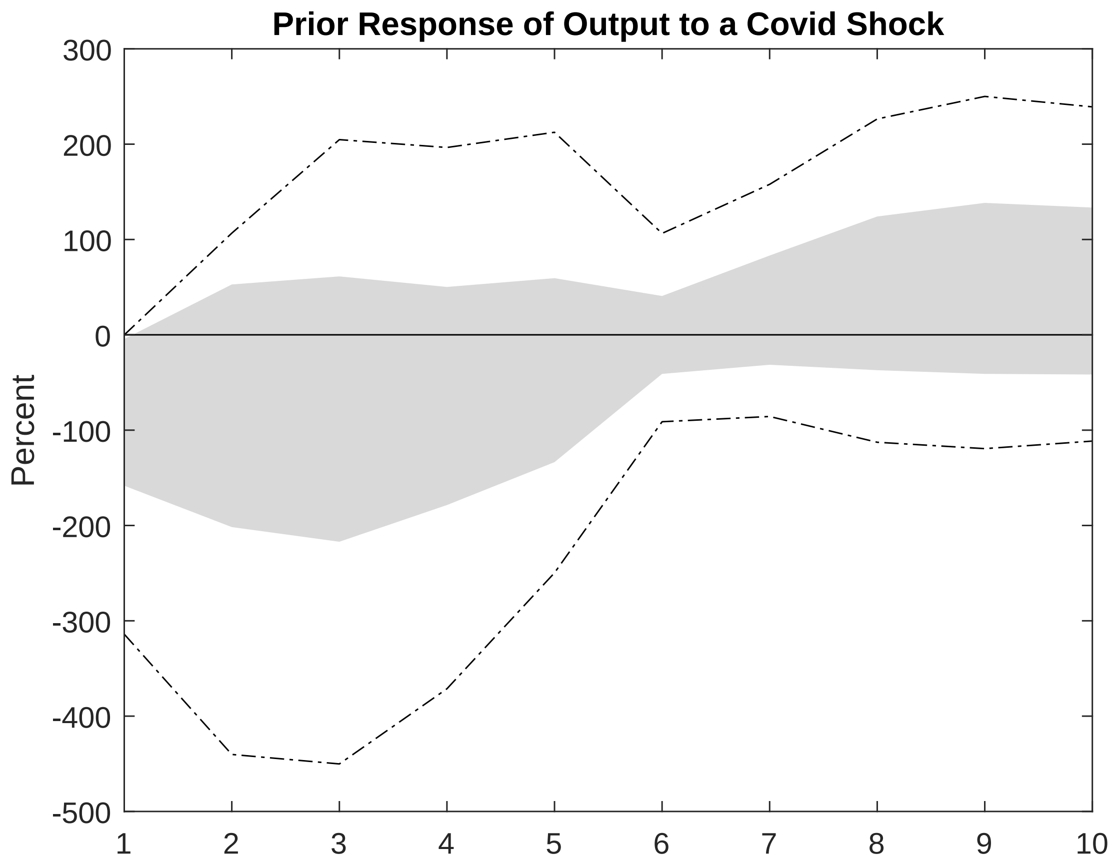

I. The effect of adding the Covid shock on the realizations of the usual shocks

The usual shocks can account for the magnitude and co-movement, just the shocks have to be far outside the historical distribution. This is shown in Figure 13. We add just one shock which makes the usual shocks more normal-sized. This can be seen by comparing Figure 14 to Figure 13).

J. A placebo experiment

What if we assumed the Covid shocks arrive in 2019q1 instead of when they actually arrived? Would the model ignore these degrees of freedom or, again, lean very heavily on them to match the data? Figures 15 and 16 show the smoothed shock decomposition of hours and inflation in terms of the structural shocks over the period 2017q4-2019q4 in the style of the main text. We assume the unusual shocks arrive in 2019q1, 2019q2 and 2019q3 and reduce

Figure 12: Difference between the log posterior kernel value in deviation from the mode and the log likelihood value in deviation from the mode value in a random neighbor of the posterior mode. Values larger than three in absolute terms indicate significant influence of the prior to the posterior.

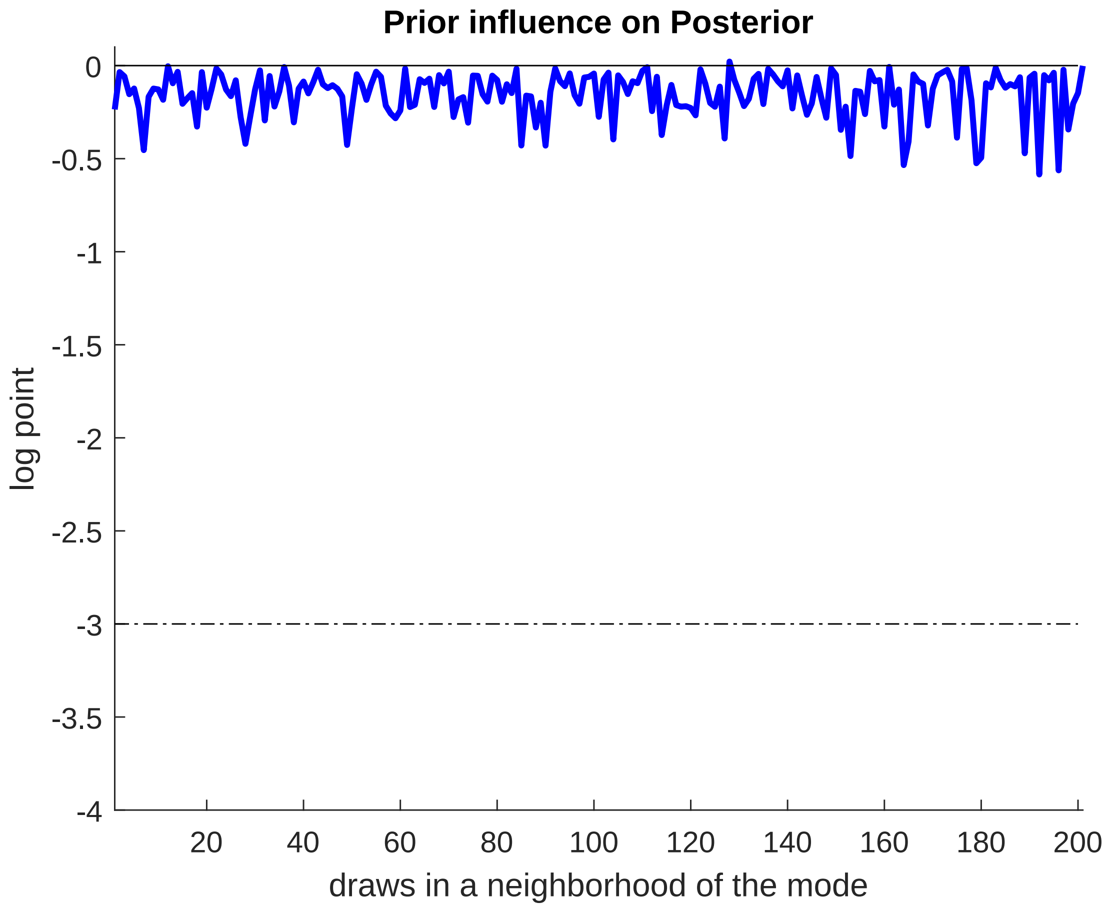

the measurement error on the forecasts to estimate the unusual shock's parameters just as we do for the Covid period. The figures show that the model does not lean on the unusual shock very much to explain the data.

Figure 13: Smoothed realizations of the usual shocks without the Covid shocks

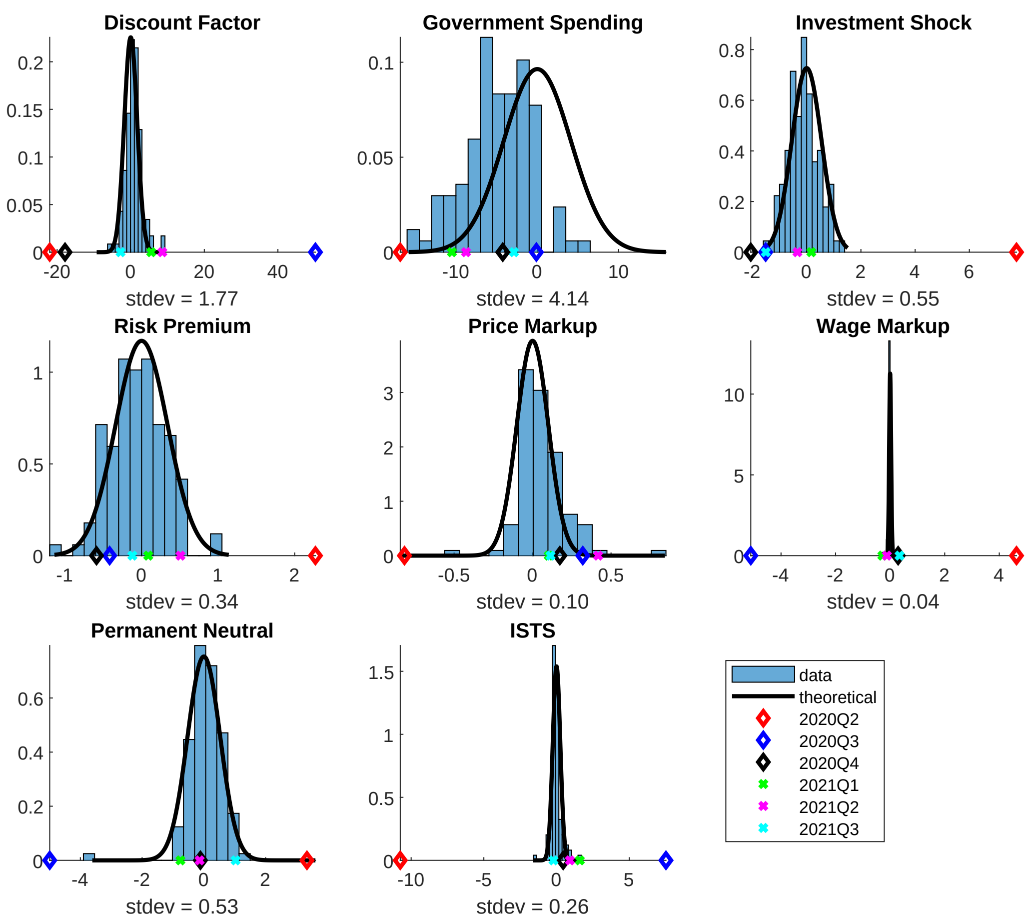

Figure 14: Smoothed realizations of the usual shocks with the Covid shocks

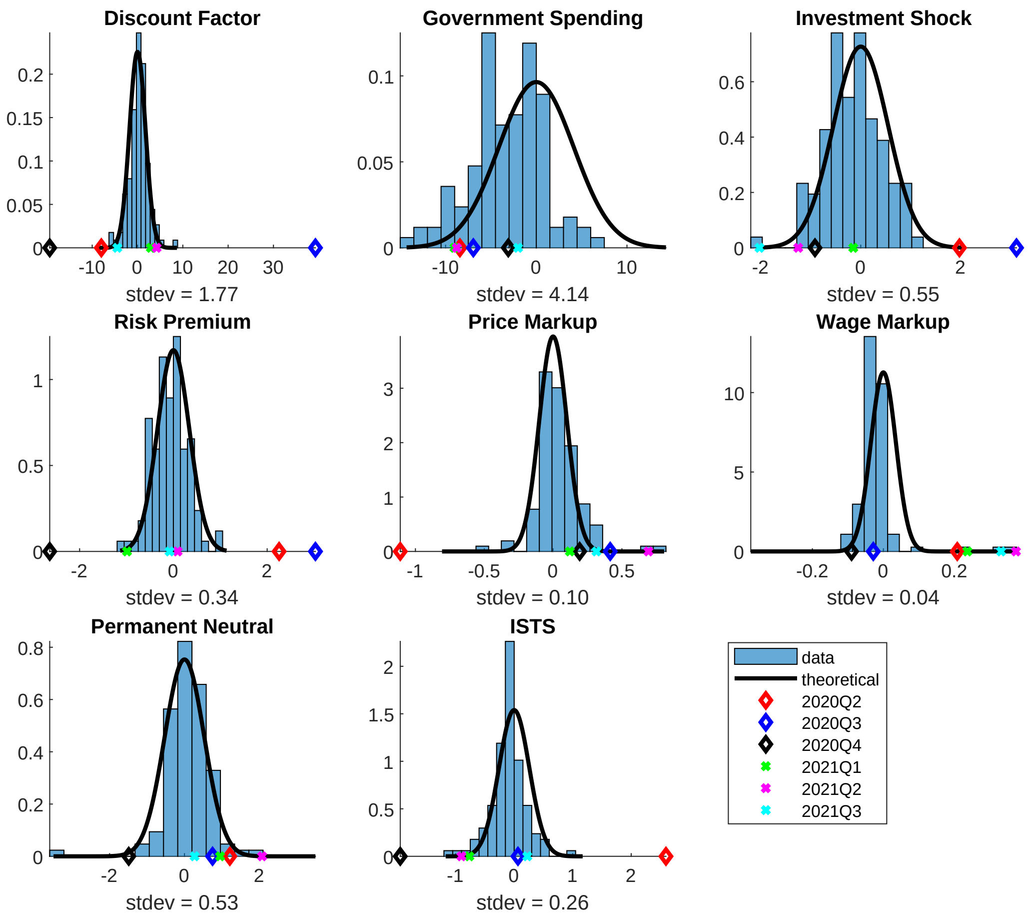

Figure 15: Smoothed decomposition of hours in the placebo experiment.

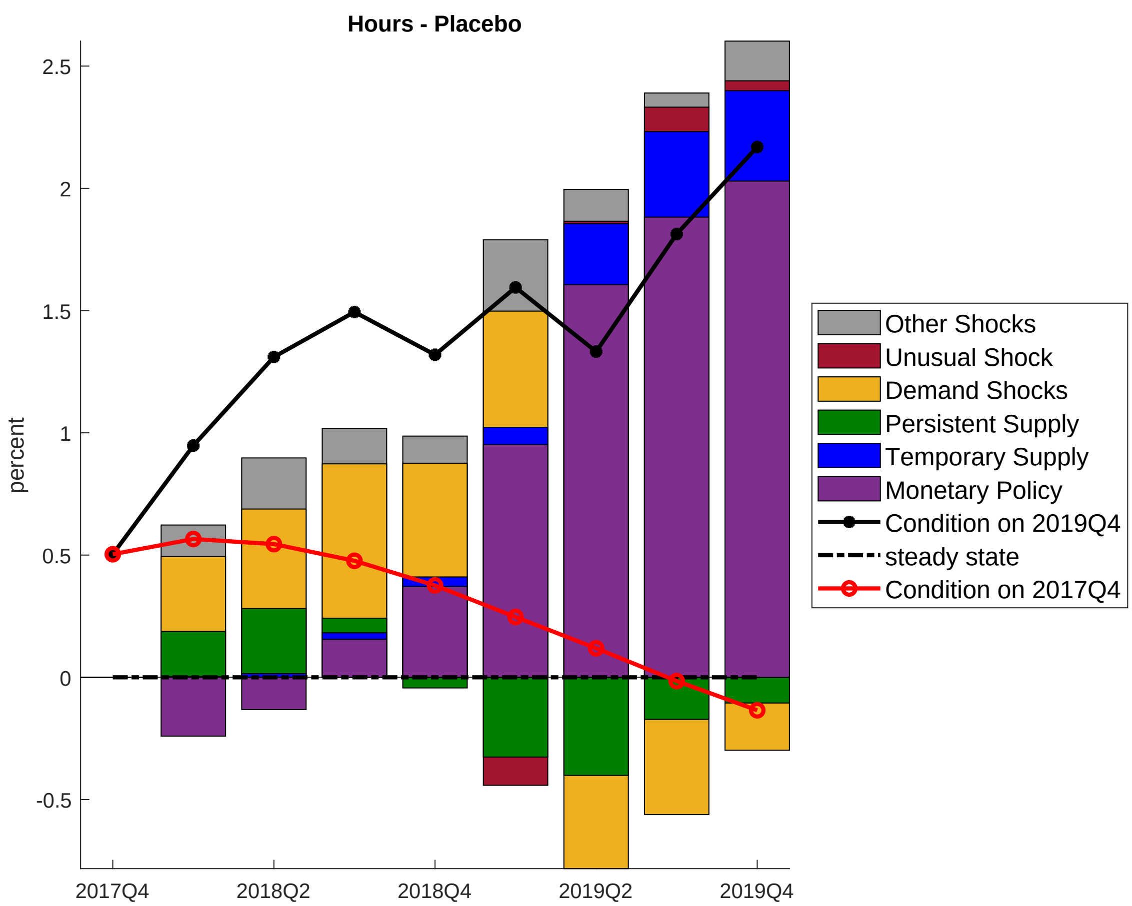

Figure 16: Smoothed decomposition of inflation in the placebo experiment.