WORKING PAPERS

N° 1539

May 2024

"Believe it or not, it's all about Beliefs!"

Fabrice Collard, Patrick Fève and Alain Guay

Toulouse School of Economics

Believe it or not, it's all about Beliefs!

Fabrice Collard*
Toulouse School of Economics (CNRS, UTC) and CEPR

Patrick Fève
Toulouse School of Economics (UTC)

Alain Guay
Université du Québec à Montréal, CIREQ and CMP†

May 23, 2024

Abstract

This paper studies the impact of Higher Order Belief (HOB) shocks, representing shifts in agents' beliefs about others' beliefs, on macroeconomic outcomes. The dynamic causal effects of these shocks are identified by leveraging a combination of a proxy-VAR approach and DSGE-based instruments. Our findings suggest that HOB shocks are indeed a key driver of the business cycle and account for a sizeable share of the observed business cycle volatility. Finally, our identification approach ensures that these shocks are not confounded with other structural/fundamental shocks.

Keywords: Higher Order Beliefs, Business cycles, proxy-VAR, DSGE.

JEL Classification: C32, E32

*Collard: fabrice.collard@gmail.com, Fève: patrick.feve@tse-fr.eu, Guay: guay.alain@uqam.ca. We are thankful to M. Angeletos, G. Candian and R. Chahrour for insightful comments. Collard and Fève acknowledge funding from the French National Research Agency (ANR) under the Investments for the Future (Investissements d'Avenir) program, grant ANR-17-EURE-0010. Collard and Fève also acknowledge funding from the French National Research Agency (ANR) under the project ANR-23-CE26-0011-01.

†Chaire en Macroéconomie et Prévisions.

> "*To understand how economies work and how we can manage them and prosper, we must pay attention to the thought patterns that animate people's ideas and feelings, their animal spirits.*" (Akerlof and Shiller, 2009)

Introduction

While the theoretical literature has offered many insights on the role of animal spirits in driving economic dynamics either in models of indeterminate equilibria (see e.g. Cass and Shell, 1983; Diamond, 1982; Benhabib and Farmer, 1994, 1999), or more recently in models of incomplete information (see e.g. Acharya et al., 2021; Angeletos and La'O, 2013; Angeletos, 2018; Angeletos et al., 2018; Benhabib et al., 2015), it has mainly relied on highly specified structural models that makes them susceptible to misspecification. Very little is however known about the role of sentiment shocks or shocks to animal spirits in less constrained environments.[1] This paper contributes to the literature by unraveling the dynamic causal effects of higher order belief (HOB, *e.g.* beliefs about the beliefs of others) shocks on the economy in an unrestricted dynamic setting —*e.g.* a Vector AutoRegressive (VAR) model, relying on an instrument derived from a theoretical model.

The identification of sentiment shocks, or HOB shocks, poses a clear challenge to the applied econometrician as such shocks do not have any direct observable counterpart. Consumer or business confidence indices, as obtained from survey data, are, at best, proxies for such shocks, and likely contaminated by other structural shocks in so far they affect current business conditions and their perception by economic agents. Likewise standard identification techniques, as used in the Structural VAR (SVAR) literature, are likely to fail properly identify such shocks. Imposing short-run exclusion restrictions, by assuming for example that some variables do not react on impact to these shocks, would certainly be misleading and yield spurious results, as these shocks are precisely shocks to the expectations about short-run economic outcomes and should precisely exert most of their effect in the short-run. Imposing long-run restrictions may not be appropriate either. These shocks are indeed not expected to exert long run effects, and imposing such constraints would not provide any relevant information for their identification. Sign restrictions do not seem to offer a promising avenue either as, since they are usually interpreted as real demand shocks (see *e.g.* Angeletos, 2018), they could easily be confounded with other demand shocks. Finally max-share identification techniques may give rise to combination of shocks rather than provide a proper identification of a unique shock.

In order to circumvent all these difficulties, this paper adopts a proxy-VAR approach initially proposed by Beaudry and Saito (1998), further developed in Mertens and Ravn (2013), Stock and Watson (2018), Kilian and Lütkepohl (2018) and recently applied by Gertler and Karadi (2015), Lagerborg et al. (2023) or Baker et al. (2024) in various contexts. This approach

[1] One important exception is Lagerborg et al. (2023) who rely on a structural vector autoregression approach to identify the dynamic causal effect of a form of sentiment shocks.

proposes to use an external instrument —a proxy variable correlated with the shock of interest but with no other shocks— to unravel the dynamic causal effect of the shock under study. Finding such an instrument to identify higher belief shocks is challenging, as most candidates, such as confidence indices, are susceptible to being correlated with other fundamental shocks. Lagerborg et al. (2023) solve this problem by using the number of fatalities in mass shootings in the US to instrument consumer confidence. They indeed find a significant effect of the associated sentiment shocks on the economy. In this paper, we propose an alternative approach to identification. More precisely, we borrow from Stock and Watson (2012) and exploit exogenous variations in a proxy measure of expectations about short-run economic outcomes generated by a Dynamic Stochastic General Equilibrium (DSGE) model. In so doing, we maintain the unrestricted dynamic setting of the VAR model while relying on a theoretically sound external instrument. More precisely, we rely on the paper by Angeletos et al. (2018) that enriches a DSGE model with a tractable form of aggregate variations in higher-order beliefs. They estimate the model on quarterly US data, which in turn enables them to explore the quantitative implications of such variations for the business cycle. Their approach offers the possibility to use Kalman techniques to recover these exogenous variations and use them as a proxy in our VAR.[2] At this stage, it is worth mentioning some technical, although important, issues. First, none of the variable included in our benchmark VAR (*i.e.* expectation error, unemployment, industrial production, capacity utilization and and the first two business cycle factors identified by McCracken and Ng (2021)) is used to recover our proxy from the DSGE model.[3] Furthermore, when variables are added to this benchmark, we make sure that they are kept totally external to Angeletos et al. (2018) (*i.e.* macroeconomic uncertainty at various horizons, credit spreads, risk premium, confidence indices, ...). Second, our proxy suffers from a generated regressor problem that ought to bias inference. In order to accommodate this problem, we take advantage of the fact that Angeletos et al. (2018) estimated the model using Bayesian maximum likelihood techniques, randomly draw N vectors of parameters from the posterior distribution of the structural parameters and generate one realization of the proxy variable for each draw. Combining this procedure with standard bootstrapping techniques when drawing inference on the VAR accounts for the generated regressor problem. More importantly, this also permits to consider a variety of models, as multiple draws of structural parameters correspond to different models featuring various degrees of real and nominal rigidities and, more importantly, degrees of information frictions (*e.g.* HOB shocks).

We estimate a VAR model for the US economy using quarterly data from 1968Q4 to 2019Q4. Importantly for our identification the model includes the forecast error on the unemployment rate, which is instrumented by our DSGE-based proxy. The proxy-VAR approach is valid in so far as the model is invertible, and the proxy satisfies two conditions: exogeneity and relevance. We start by running the invertibility test developed by Plagborg-Møller and Wolf (2022), which

[2]To our knowledge, this is the sole quantitative medium-scale model featuring real, nominal and information frictions offering a reasonable fit of US data and allowing for the construction of HOB shocks.

[3]Although they display some correlation with the variables used by Angeletos et al. (2018), this correlation is far from being perfect (see e.g. Appendix A).

amounts to test whether the proxy jointly Granger causes the variables of the VAR, and find that the invertibility condition is satisfied across all the draws of the proxy. By construction, our proxy shall be exogenous with respect to the other shocks present in the DSGE model. Nevertheless, we also verify that, for any of the draw we consider, the proxy is not Granger caused by the variables of the VAR. The test clearly rejects Granger causality in more than 97.5% of the cases. Finally, we assess the relevance of the set of proxies by conducting a standard Montiel Olea and Pflueger (2013) weak instrument test. The test reveals that the instrument is indeed relevant in more than 98% of the draws of the proxy from the DSGE model.

The HOB shock induces an economic boom, characterized by a significant and persistent decline in the unemployment rate, an increase in industrial production and capacity utilization. Even though the effect of the shock on expectation errors is short lived, the response of the macro aggregates displays a clear hump-shaped pattern, which suggests a strong persistence mechanism at play in the transmission of these shocks. HOB shocks are also found to be key contributors to the volatility of aggregates (unemployment, industrial production, capacity utilization). It however leaves no footprint on industrial production in the long run. All these results are found to be robust to variations in the timing of the expectation error, substituting the error on unemployment by the error on output, variations in the estimation method used to recover the dynamic causal effect (Cholesky VAR, Local Projections). Using a placebo experiment, we also establish that they are also not the outcome of a divine coincidence.

These findings extend to other business cycle indicators, such as consumption expenditures of durable goods, gross private domestic investment or hours worked. In contrast, these shocks do not affect inflation. Using the technique developed by Angeletos et al. (2020), we show the HOB shocks are a key contributor of the main business cycle (MBC) shock —*i.e.* the shock that explains the bulk of various indicators volatility at business cycle frequencies (6-32 quarters). The response of the main aggregates to the MBC shock are essentially identical to the HOB shock. The history of aggregates conditional on both shocks are also found to be highly correlated, especially over the last 20 years.

We then examine additional implications of HOB shocks on the economy, starting with macroeconomic uncertainty. Our results indicate that HOB shocks lead to a drop in macroeconomic uncertainty, as measured by Jurado et al. (2015), and account for about 30% of its volatility in the short-run and quickly fades away. This can be interpreted as HOB shocks ease the short-run forecasting task of economic agents, which provides a candidate transmission channel to the rest of the economy. In line with this result, the shock leads to a reduction in the risk premium faced by economic agents and accounts for about 30% of its short-run volatility. A similar result obtains for the credit spread, which recedes in the short-run, indicating that HOB shocks do indeed improve the financial conditions in the economy, which in turn promotes the boom. Interestingly, while HOB shocks shall be interpreted as shocks pertaining to what agents think about the beliefs of others, they do not explain the evolution of confidence indices —the Michigan consumer confidence index— at any horizon. In fact, replacing the expectation error by the confidence index in our VAR to identify the dynamic causal effect of HOB shocks makes the proxy suffer from a weak instrument problem. In other words, the proxy does not provide useful information about confidence as measured by economic surveys.

One potential challenge lies in the possibility that HOB shocks be confounded with another source of shifts to expectations: technological news. We therefore consider a version of our VAR featuring, in the lines of Beaudry and Portier (2006), stock prices and total factor productivity (TFP). The HOB shock does not affect the evolution of TFP, explaining less than 2% of its volatility at any horizon in a Vector Error Correction Model. The HOB shock is therefore not confounded with a technological news shock. When we use the DSGE-based news shock to instrument the expectation error, this instrument is found to suffer a serious weak instrument problem, hence reinforcing the preceding conclusion. Similar conclusions hold with other structural shocks such as unanticipated technology, investment, discount factor or policy shocks.

The remaining of the paper is organized as follows. Section 1 provides an overview of our methodology, first outlining the proxy VAR identification strategy and then delving into the construction of our model-based instrument. Section 2 presents our main results. It starts by considering a baseline VAR and draws the connection to the main business cycle shock identified in Angeletos et al. (2020). It then provides additional insights on the effects of our HOB shock on key economic variables such as uncertainty and spreads. It also make precise the way in which our identification does not confound variations in HOB with other shifters of expectations. Section 3 finally assesses the robustness of our findings to variations in the econometrician's information set and identification strategy. A last section offers some concluding remarks.

1 Empirical Approach

In this section, we briefly review the proxy-VAR approach, as developed by Beaudry and Saito (1998), Mertens and Ravn (2013), Stock and Watson (2018), Kilian and Lütkepohl (2018) and applied by Gertler and Karadi (2015), Lagerborg et al. (2023) or Baker et al. (2024), to identify and estimate the dynamic causal effects of higher order beliefs (Hereafter HOB) shocks on aggregate activity. This approach relies on external instruments —a variable that is correlated with a shock of interest, but not with other shocks— to estimate dynamic causal effects of a structural shock on macroeconomic aggregates.[4] We then explain in detail the construction of our instrument and its main time series properties.

[4]Note that although we will focus on a proxy VAR approach to recover the dynamic causal effects of higher order belief shocks, we will also assess the robustness of our results to using a Cholesky decomposition (see Sims, 1980; Lagerborg et al., 2023; Plagborg-Møller and Wolf, 2021) and a local projection IV approach (see Ramey and Zubairy, 2018; Stock and Watson, 2018).

1.1 The Proxy-VAR Approach

Let Y_t be a $n_y \times 1$ vector of second order stationary endogenous variables, whose dynamics can be represented by the following Vector-AutoRegressive (VAR) process

$$A(L)Y_t = u_t \tag{1}$$

L denotes the lag operator ($L^i Y_t = Y_{t-i}$) and $A(L) = I - \sum_{i=1}^{p} A_i L^i$ is a matrix polynomial where A_i a $(n_y \times n_y)$ matrix, p denotes the number of lags in the VAR.[5] u_t is a $(n_y \times 1)$ vector of canonical innovations satisfying $\mathbb{E}[u_t] = 0$ and $\mathbb{E}[u_t u_t'] = \Sigma$ and $\mathbb{E}[u_t u_{t-j}'] = 0$ for any $j > 0$. These innovations are assumed to be linear combinations of n_y mutually orthogonal shocks, ε_t such that

$$u_t = S\varepsilon_t \tag{2}$$

where S is a non-singular $(n_y \times n_y)$ matrix. The orthogonal shocks satisfy $\mathbb{E}[\varepsilon_t] = 0$ and $\mathbb{E}[\varepsilon_t \varepsilon_t'] = \Omega$ where $\Omega_{ij} = 0$ for $i \neq j$. The Wold decomposition of the process is given by

$$Y_t = C(L)\varepsilon_t \tag{3}$$

where $C(L) = A(L)^{-1}S$. In this paper, we are interested in identifying a single structural shock —the HOB shocks. By convention, and for notation convenience, we will order this shock first in this vector, $\varepsilon_{1,t}$. Our approach to identifying this shock follows the instrumental strategy developed by Stock and Watson (2018).

Let Z_t be the (uni-dimensional) proxy variable —the instrument— for $\varepsilon_{1,t}$. This instrument has to satisfy

$$\mathbb{E}[\varepsilon_{1,t} Z_t] = \theta \neq 0 \tag{Relevance}$$

$$\mathbb{E}[\varepsilon_{j,t} Z_t] = 0 \;\; \forall j > 1 \tag{Exogeneity}$$

The first condition, relevance, ensures that the instrument exhibits a significant correlation with the shock of interest. The second condition, exogeneity, guarantees orthogonality between the instrument and the remaining shocks. Using (2) and the relevance and exogeneity conditions, we have

$$\mathbb{E}[u_t Z_t] = S\mathbb{E}[\varepsilon_t Z_t] = S \begin{pmatrix} \theta \\ 0 \end{pmatrix} = \begin{pmatrix} S_{1,1}\theta \\ S_{2,1}\theta \\ \vdots \\ S_{n_y,1}\theta \end{pmatrix}$$

When the effect of the shock of interest on a reference variable (say the first appearing in the vector Y_t) is normalized to unity ($S_{1,1} = 1$)[6], the last equality implies that in this case, $S_{i,1}$ can be directly recovered as:

$$S_{i,1} = \frac{\mathbb{E}[u_{i,t} Z_t]}{\mathbb{E}[u_{1,t} Z_t]}$$

[5] Without loss of generality we omit constant terms.

[6] In our case, the shock is the HOB shock and, as will become clear in the next section, the variable of interest is an expectation error.

which is readily obtained from the IV regression

$$u_{i,t} = S_{i,1}u_{1,t} + \sum_{j=2}^{n_y} \alpha_{i,j}\varepsilon_{j,t}.$$

Since the innovations u_t are unobservable, direct estimation of this regression is possible only at the cost of a generator regressor problem. In order to circumvent this problem, we follow the method proposed by Stock and Watson (2018) and exploit the fact that $u_{i,t} = Y_{i,t} - \mathbb{P}[Y_{i,t}|Y_{t-j}, j > 1]$, where $\mathbb{P}[Y|X]$ denotes the projection of Y onto X, to rewrite the previous regression as:

$$Y_{i,t} = S_{i,1}Y_{1,t} + \psi_i(L)Y_{t-1} + \sum_{j=2}^{n_y} \alpha_{i,j}\varepsilon_{j,t} \text{ for } i = 2, \ldots, n_y$$

where $\psi_i(L)$ collects all the coefficients of the projection of $Y_{i,t} - S_{i,1}Y_{1,t}$ onto the space spanned by the past history of Y_t. Subsequently, $S_{i,1}$ and $\psi_i(L)$ can then simply be estimated using a two stage least square method, employing Z_t as an instrument for $Y_{1,t}$. By applying this regression to each variable, we obtain the column vector $S_{\cdot,1}$ conditional on $S_{1,1} = 1$.

In this paper, we do not impose a unit effect of the shock on the variable of reference, but rather impose a unit shock. This requires a re-scaling of $S_{\cdot,1}$. This is achieved by noting that as long as (2) holds, we have $\varepsilon_{1,t} = \gamma u_t$ where $\gamma = S'_{\cdot,1}\Sigma^{-1}/(S'_{\cdot,1}\Sigma^{-1}S_{\cdot 1})$. Consequently, the volatility of $\varepsilon_{1,t}$ can be computed as $\sigma^2_{\varepsilon_1} = \gamma\Sigma\gamma' = (S'_{\cdot,1}\Sigma^{-1}S_{\cdot,1})^{-1}$, which can be used to re-scale S by $1/\sigma_{\varepsilon_1}$. In this scenario, the experiment effectively applies a unit shock to the system. Impulse responses and the associated forecast error variance decomposition can then be derived from (3). Note that, critically, this final step is embedded into the bootstrap procedure when drawing inference.

1.2 A Model Based Instrument

The previous section underscored the need for an instrument to identify HOB shocks. Lagerborg et al. (2023) employed mass shootings in the US as an instrument. In this paper, we follow Stock and Watson (2012) and leverage the latest advancements in modeling sentiment/HOB driven fluctuations within Dynamic Stochastic General Equilibrium (DSGE) models to construct such an instrument. Our approach hinges on the idea that a structural model can provide a theory-driven proxy for the instrument. We emphasize the use of the term 'proxy' to acknowledge the potential misspecification of the model, as it is likely calibrated/estimated to a specific economy and sample period. However, from a theoretical standpoint, if the model incorporates HOB shocks as a significant factor, it will impose useful constraints on its behavior. Furthermore, we will meticulously consider a variety of models and, consequently, instruments to assess the sensitivity of the results to a population of instrumental variables.

Standard macroeconomic models, in particular those within the DSGE framework, leave no room for the possibility of phenomena such as waves of optimism, pessimism, or sentiments. In these idealized economies, all agents possess the same information, share a common prior, have perfect knowledge of the beliefs of other players, and fully comprehend their behavior in response to exogenous shocks. Consequently, they reach a unanimous consensus regarding the current state and future economic outlook. However, relaxing these stringent assumptions opens the door to models that can accommodate phenomena akin to coordination failures and self-fulfilling fluctuations, despite the existence of a unique equilibrium (see e.g. Angeletos and La'O, 2013, among others).

In this paper, we will derive our instrument employing the framework developed in Angeletos et al. (2018), which offers a simple way to recover a higher order belief shock. Following Angeletos and La'O (2013), they consider a setup in which the economy consists of a continuum of islands and a mainland. Each island is populated by households and firms that interact on local input markets, where they produce differentiated intermediate goods. At this stage, agents on different islands are unable to communicate or coordinate their decisions, leading to incomplete information regarding the choices made on other islands. Subsequently, all island-specific intermediate goods are traded on the centralized mainland market, where they are combined to produce the final good for consumption and investment purposes.

Each period is split into two interim stages, Stage 1 and Stage 2. Stage 1 —the stage in which firms decide on their inputs for a given period t— deviates from the common knowledge assumption. Each island i observes only a vector of private signals, s_{it}, related to the vector of fundamental shocks, z_t. The fundamental shocks[7] are modeled as an autoregressive process of the form:

$$z_t = \Phi z_{t-1} + \varepsilon_{z,t} \tag{4}$$

where $\varepsilon_{z,t}$ follows a Gaussian distribution with mean 0 and diagonal covariance matrix Σ_z. The matrix Φ has all its eigenvalue inside the unit circle. The signal s_{it} is given by

$$s_{it} = z_t + \eta_{it} \tag{5}$$

where η_{it} is an island specific noise. This structure is then used to engineer variations in higher order beliefs by departing from the common prior assumption. While each island considers its own signal to be unbiased —*i.e.* $\eta_{it} \sim N(0, \sigma^2)$— it assumes that other islands' signals are systematically biased —*i.e.* $\eta_{jt} \sim \mathcal{N}(\xi_t, \sigma^2)\ \forall j \neq i$. The bias ξ_t represents the commonly known disagreement among agents in Stage 1 and is assumed to follow an AR(1) process of the form

$$\xi_t = \rho \xi_{t-1} + \varepsilon_{\xi,t} \tag{6}$$

where $\varepsilon_{\xi,t}$ follows a Gaussian distribution with mean 0 and standard deviation σ_ξ, and $\mathbb{E}[\varepsilon_z \varepsilon_\xi] = \mathbf{0}$. Finally, innovations in the gap between first- and higher-order beliefs can exhibit some persistence but eventually vanish, $|\rho| < 1$. In that setting, for all agents i, all other agents $j \neq i$,

[7] The model features 8 fundamental shocks: a news shock, a permanent and a transitory technology shock, a permanent and a transitory investment shock, a discount factor shock, a government expenditure shock and a monetary policy shock.

all periods t, and all states of nature, agent i's belief during stage 1 satisfy

$$\mathbb{E}^1_{it}[z_t] = s_{it},$$
$$\mathbb{E}^1_{it}[\mathbb{E}^1_{jt}[z_t]] = \mathbb{E}^1_{it}[s_{jt}] = s_{it} + \Lambda\xi_t,$$

where Λ is a matrix loading the belief shock on the fundamental shocks. We denote $\bar{s}_t$ the average signal in the economy and note that the "truth" is such that $s_{it} = \bar{s}_t = z_t$. However, this true state is not publicly revealed until Stage 2 of period t. In stage 1, instead, each island erroneously believes that

$$\mathbb{E}^1_{it}\bar{s}_t = s_{it} + \Lambda\xi_t.$$

In stage 2, the realizations of the fundamentals, z_t, of all the signals, s_{it}, and of all the Stage 1 decisions become common knowledge. Furthermore, the actual realization of the signals are such that $s_{it} = z_t$ for any period t and any agent i. This implies in particular that, in an equilibrium,

$$\mathbb{E}^2_{it-1}\bar{s}_t = \Phi z_{t-1} + \Lambda\rho\xi_{t-1}.$$

Angeletos et al. (2018) propose a tractable method for solving the model under this information structure. Specifically, they show that, as long as the cross-sectional volatility of signals $\sigma = 0$,[8] the solution of the (log-linearized) model can be expressed as:[9]

$$X_{t+1} = \Phi_x X_t + \Phi_s \bar{s}_t + \Phi_\xi \xi_t \tag{7}$$
$$Y_t = \Pi_x X_t + \Pi_s \bar{s}_t + \Pi_\xi \xi_t \tag{8}$$

where Y_t denotes the vector of forward-looking variables and X_t is the vector of predetermined variables. The matrices Φ_x, Φ_s, Π_x, Π_s correspond to those associated with the standard full information setting under rational expectations, Φ_ξ and Π_ξ are specific to the information structure.[10] Using the information structure and the form of the solution in an equilibrium, we can readily derive the stage-2 forecast error of Y_t, $\zeta^Y_{t|t-1} \equiv Y_t - \mathbb{E}^2_{it-1}[Y_t]$, as[11]

$$\zeta^Y_{t|t-1} = \underbrace{\Pi_s\varepsilon_{z,t}}_{\substack{\text{Fundamental}\\ \text{innovations}}} + \underbrace{\Pi_\xi\varepsilon_{\xi,t}}_{\substack{\text{HOB}\\ \text{innovations}}} - \Pi_s\Lambda\rho\xi_{t-1} \tag{9}$$

In other words, the expectation error is shifted both by the fundamentals and the higher order belief shock. Belief shock identification can therefore be achieved working directly with

[8] This essentially amounts to assume that all agents attribute the same bias, ξ_t, to the other agents' beliefs.

[9] The reader should keep in mind that, in an equilibrium, $\bar{s}_t = z_t$.

[10] Huo and Takayama (2023) also derive a state space representation of the solution of the model in the more general case where $\sigma > 0$ and agents share a common prior. The form of the solution is however more complicated as it involves higher order AR and MA components.

[11] This results derives from the fact that the Stage 1 forecast of island i about Y_t as $\mathbb{E}^1_{it}Y_t = \Pi_X X^b_t + \Pi_s s_{it} + (\Pi_\xi + \Pi_s\Lambda)\xi_t$. Keeping in mind that in Stage 2 shocks are perfectly observed, its Stage 2 forecast of Y_t in the previous period is given by $\mathbb{E}^2_{it-1}Y_t = \Pi_X X^b_t + \Pi_s\Phi z_{t-1} + (\Pi_\xi + \Pi_s\Lambda)\rho\xi_{t-1}$. The perfect observability of the shock in stage-2 implies that $\mathbb{E}^2_{it-1}[Y_t] = \mathbb{E}^2_{jt-1}[Y_t] = \mathbb{E}^2_{t-1}[Y_t]$. Accordingly, the expectation error is given by $\zeta^Y_{t|t-1} \equiv Y_t - \mathbb{E}^2_{it-1}[Y_t] = \Pi_s(\bar{s}_t - \Phi z_{t-1}) + \Pi_\xi(\xi_t - \rho\xi_{t-1}) - \Pi_s\Lambda Q\xi_{t-1}$. Using the fact that in equilibrium, $\bar{s}_t = s_{it} = z_t$, the result follows.

expectation errors. This principle will guide us in the design of our VAR in the next section. Since $\varepsilon_{\xi,t}$ is an innovation ($\mathbb{E}[\varepsilon_{\xi,t}\xi_{t-j}] = 0\ \forall j \in \mathbb{N}\backslash\{0\}$) and since $\mathbb{E}[\varepsilon_{z,t}\varepsilon_{\xi,t-j}] = 0\ \forall j \in \mathbb{Z}$, then the projection of $\zeta^Y_{t|t-1}$ onto $\varepsilon_{\xi,t}$ identifies the belief shock in the data. In other words, our instrument is the innovation to the belief shock process, as identified in the DSGE model.

In what follows, we use a New-Keynesian model, in the lines of Smets and Wouters (2007), as reported in Angeletos et al. (2018) to obtain realizations of the innovation of the belief shock. Specifically, we draw 200 parameter realizations from the posterior density of the structural parameters estimated in their study. For each parameter realization, we solve the model and employ the Kalman smoother to recover the structural shocks of the model from the observed quarterly data for real GDP, consumption (non-durables and services), investment (gross private domestic investment + durables), total hours worked, GDP deflator inflation, and the federal fund rate over the period 1968Q4-2019Q4. This procedure yields 200 time series for our instrument, each corresponding to a particular model. This approach allows us to achieve two important goals. First, from a purely statistical point of view, it takes model uncertainty into account when drawing inference in our proxy-VAR analysis. Second, it permits to obtain HOB instruments in a variety of models, each characterized by different degrees of nominal, real frictions and properties of the belief shock. In this sense, it allows us to assess the robustness of our instrument approach across various economic environments. Table 1 illustrates this point by presenting descriptive statistics for parameters related to real frictions (investment adjustment cost, IAC, and habit persistence in consumption, HPC), nominal rigidities (Calvo probability of price resetting) and the properties of the HOB shock. Real frictions display substantial

Table 1: Variations in Rigidities

	Frictions			Beliefs	
	IAC	HPC	Calvo	Persistence	Volatility
min	0.800	0.539	0.579	0.405	0.180
mean	3.427	0.752	0.730	0.826	0.671
max	8.378	0.902	0.851	0.979	2.057
s.d.	1.059	0.057	0.033	0.061	0.270

Note: IAC: investment adjustment cost and HPC: habit persistence in consumption.

variation, with the investment adjustment costs parameter ranging from 0.8 to 8.4, and the degree of habit persistence ranging from 0.5 to 0.9, Notably, the posterior distribution of the Calvo parameter implies an average length of price contracts ranging from 2.4 to 6.6 quarters. The HOB shock displays substantial variation in persistence—which can be interpreted as the persistence of coordination failures— ranging from 0.4 to 0.98, and volatility, from 0.18 to 2.06. Our analysis will therefore span a large spectrum of models.

Panel (a) of Figure 1 depicts the dynamics of the constructed HOB shock (left panel) and its associated innovation (right panel). The shaded band surrounding each series represents the

Figure 1: A Variety of Models

(a) Dynamics of Instruments

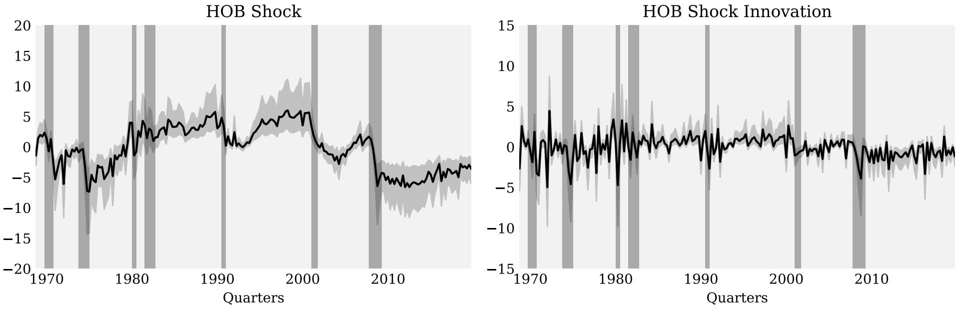

(b) Cross-Model Volatility

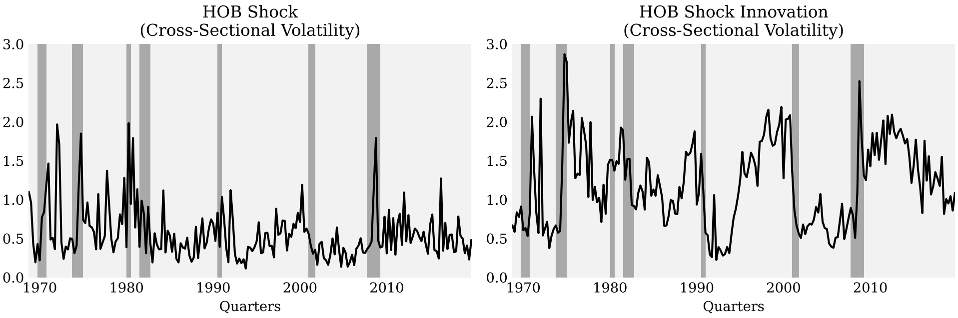

Panel (a): Plain Line: average across models, Shaded area: 95% of the across models distribution, Vertical shaded areas: NBER recessions.

95% cross-model variation for each period. The vertical shaded areas indicate NBER recession periods. The figure clearly demonstrates that beliefs tend to decline during recessions and rise during periods of economic expansion. Panel (b) of the figure illustrates the cross-sectional (model) variation of the belief shock and its innovation for each period. It reveals that drawing economies from the posterior distribution generates substantial and time-varying cross-sectional volatility in both the belief shock and its innovation. In other words, this procedure introduces significant variability into the instrument, which, as previously mentioned, accounts for model uncertainty and allows us to assess the robustness of our instrument.

Figure 2: Serial Correlation of the Instrument across Models

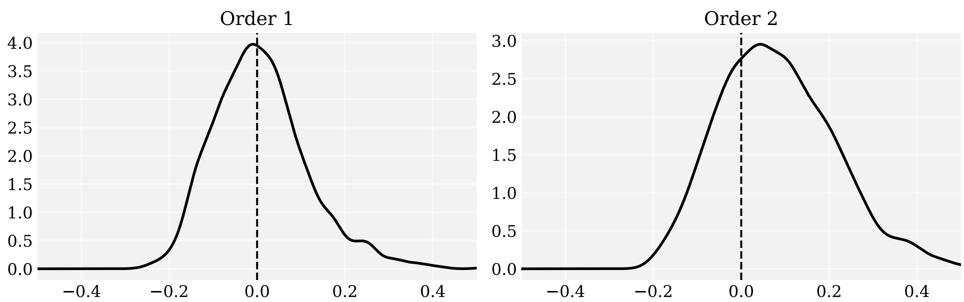

Figure 2 reports the cross-model distribution of the autocorrelation of the innovation of order 1 and 2. The graph suggests that the innovation exhibits minimal serial correlation, making it a suitable candidate as a shock instrument.[12]

2 Empirical Results

This section unveils our main findings, encompassing both the impulse response function (IRF) of key economic aggregates to a HOB shock and its contribution to the forecast error variance decomposition (FEVD). Subsequently, we delve into the question of whether this HOB shock truly represents a business cycle shock. Further, we explore the broader implications of this shock for the economy and its potential entanglement with another important expectation shifter: news shocks.

2.1 Baseline VAR

To capture the effects of HOB shocks on business cycle dynamics and their contribution to overall business cycle volatility, we begin by estimating a VAR model for the US economy using quarterly data from 1968Q4 to 2019Q4. The model includes a set of key business cycle variables: the unemployment rate (u), the industrial production index (IP, in log. difference) and capacity utilization (Util) obtained from the Federal Reserve Economic Database (`https://fred.stlouisfed.org/`). Additionally, we incorporate the first two business cycle factors identified by McCracken and Ng (2021) to capture any potential missing information about the business cycle, and hence avoid any nonfundamentalness representation issue (see Forni and Gambetti, 2014; Beaudry et al., 2019). Finally, and most importantly, the VAR includes the forecast error on the unemployment rate, $\zeta^u_{t|t-1}$, which will play a key role in our identification (see Equation 9). This variable is constructed by subtracting the one-step-ahead forecast of the unemployment rate, $u^e_{t|t-1}$, as reported in the survey of professional forecasters, from the first release of the unemployment rate for that period. The VAR features 4 lags, as selected by a standard Bayesian Information Criterion.

For our proxy-VAR analysis to be valid, the HOB shock must be invertible —*i.e.* spanned by past and current (but not future) values of the endogenous variables. As shown by Plagborg-Møller and Wolf (2022), this is equivalent to saying that our instrument does not Granger cause the variables included in the VAR.[13] Plagborg-Møller and Wolf (2022) propose a simple test consisting of expanding the VAR with the instrument and testing for the joint Granger causality of the instrument on the other variables. Figure 3(a) reports the distribution of the test statistics across our 200 models. The vertical dashed line corresponds to the threshold value of joint significance of the HOB in the VAR at the 95% confidence level, the plain vertical line is the average of the test across our 200 models. First and foremost, the average value of the test lies way below the threshold, indicating that, on average, the HOB shock is indeed invertible.

[12]Note that, at this stage we have establish neither the relevance of the instrument nor its exogeneity, which will be assessed in Section 2.

[13]Intuitively, if the MA(∞) of the shock is invertible, then the lags of the variables in the VAR capture all the forecasting power of lags of the HOB shock. In other words, the lags of the instrument do not help predict the variables in the VAR.

Additionally, the probability that the test statistics lies below the threshold is 1, indicating that the invertibility of the model cannot be rejected.

Figure 3: Preliminary Tests

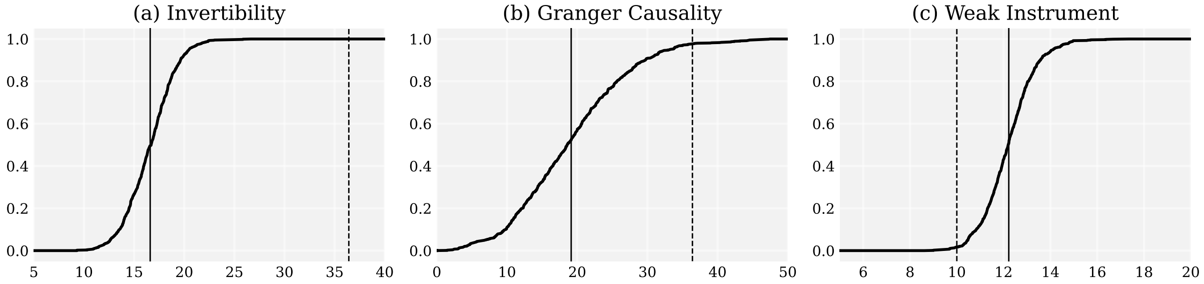

Note: Vertical line: Average value of the Test across 200 models , Vertical dashed line: Threshold value of the test (rule of thumb threshold of 10 for the weak instrument test).

A second key assumption of the approach is that our instruments be exogenous. Note that, by construction, in the DSGE model, the innovation of the HOB shock, our instrument, is orthogonal to any of the other fundamental shocks (technology, investment, discount rate, monetary policy, ...) and is structurally exogenous.[14] Nevertheless, we present the Granger causality test of the VAR variables onto our instrument across the 200 models in Figure 3(b). As for the invertibility test, the average value of the test (plain vertical line) lies way below the threshold (dashed vertical line), indicating that, on average, the HOB shock is not Granger caused by any of the macroeconomic aggregates we consider. Likewise, the probability that the test statistics lies below the threshold is greater than 0.975, indicating that the shock is not Granger caused by any of the macroeconomic aggregates across almost all models.

Finally, we assess their relevance by conducting a standard Montiel Olea and Pflueger (2013) weak instrument test across our 200 draws of the belief shocks from the model. This test evaluates the significance of the instrument in the projection of the unemployment forecast error, $\zeta^u_{t|t-1}$, onto lags of the variables in the VAR and our instrument. Figure 3 depicts the empirical cumulative distribution function of the test statistic for the information set utilized in our VAR. The average of the distribution (represented by the plain vertical line) is well above the rule-of-thumb threshold (dashed vertical line) of 10. In fact, the test values exceed this threshold in more than 98% of the draws, strongly suggesting that the weak instrument problem is unlikely to be an issue over all models. Figure 4 illustrates this point and reports the joint evolution of the forecast error on the unemployment rate and the average HOB instrument across models. [15] The two series are clearly highly correlated (-0.58), and the HOB instrument closely tracks the booms and recessions.

Figure 5 reports the impulse response functions (hereafter IRF) of the variables of the VAR to a 1 standard deviation HOB shock, along with the ± 1 standard deviation band (shaded

[14] We address the possibility, in sample, of a correlation between the HOB shock and other fundamental shocks in Sectionsec:corrxi.

[15] Appendix B reports the corresponding figure for the forecast error of output. In that case, the correlation is 0.46.

Figure 4: HOB instrument vs Forecast Error on Unemployment Rate

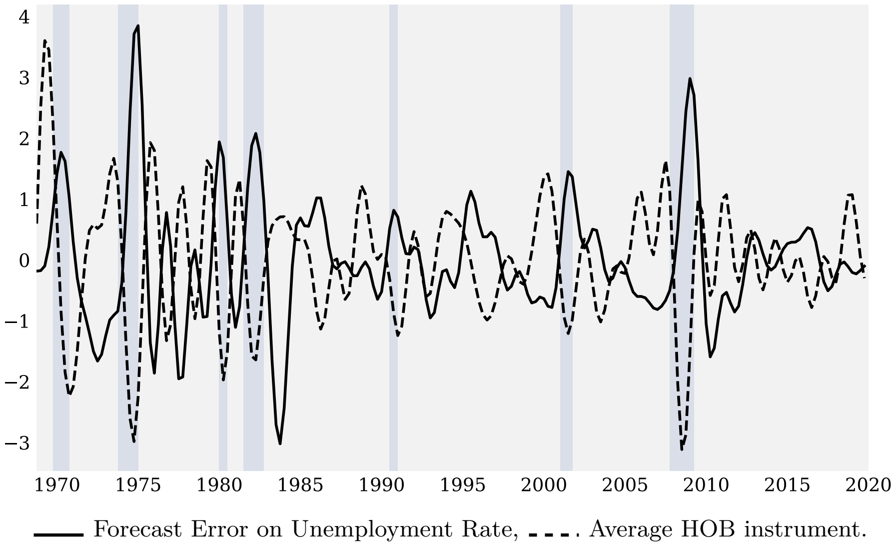

The HOB instrument corresponds to the average of the HOB instrument across models. Both series is filtered by applying the Christiano and Fitzgerald's 2003 filter (6-32 quarters frequencies).

area).[16] The shock induces an economic boom, characterized by a significant and persistent

Figure 5: Baseline VAR: Impulse Response to a HOB Shock

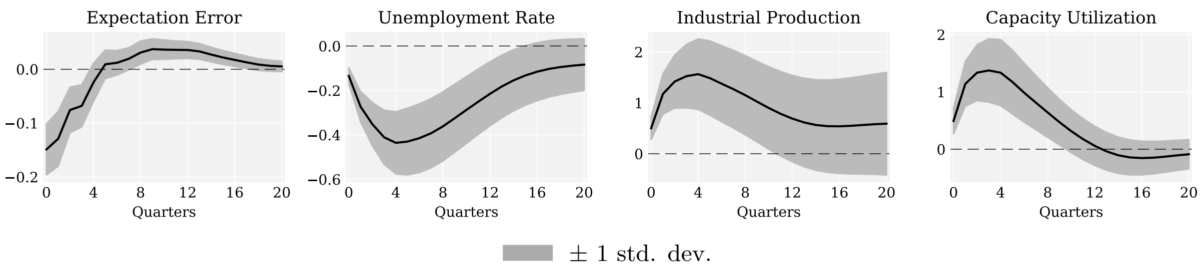

decline in the unemployment rate, an increase in the (log) level of industrial production, and a surge in the capacity utilization rate. Even though the effect of the shock on expectation errors is short lived, the IRFs point to the presence of a strong propagation mechanism, evident from the overshooting dynamics that suggests the existence of some form of learning behavior. Interestingly, this comes together with a reduction in the expectation error. In other words, following a positive HOB shock, agents tend to underestimate the actual decline in the unemployment rate. Table 2 reports the contribution of the HOB shock at various horizons. The shock is a major contributor to the overall business cycle volatility. It accounts for a substantial 55% of the expectation error and unemployment fluctuations, and over a third of the volatility in the (log) level industrial production and utilization rates upon impact. This influence persists, with the shock contributing close to 70% to unemployment volatility (55% for industrial production and 52% for utilization) after a year. Furthermore, the shock essentially leaves no footprint on business cycle volatility in the long run. For instance, the shock only accounts for 1% of

[16] We abstract from the IRFs of the 2 McCracken's factor as they have no clear economic interpretation *per se*.

Table 2: FEVD: Baseline VAR

	Impact	1-Quarter	1-Year	2-Years	5-Years	Long-Run
Exp. Error ($\zeta^u_{t\|t-1}$)	54.32	51.26	43.38	41.11	42.03	41.54
Unemployment Rate	54.79	72.66	69.78	63.68	47.05	28.04
Industrial Production	34.40	55.62	55.50	47.57	30.33	1.39
Capacity Utilization	33.19	53.37	52.15	46.11	39.96	12.76

the overall volatility of industrial production in the long run, which confirms that while HOB shocks may exert a persistent effect on the business cycle they do not leave a permanent mark on the economy.[17]

To further examine the business cycle implications of our HOB shock, we expand our VAR model by introducing, one at a time, investment indicators (consumption expenditures on durable goods, and gross domestic private investment), and hours worked. This allows us to assess the extent to which our HOB shock drives fluctuations in the economy. Figure 6 and Table 3 depict the IRFs and variance contribution of the HOB shock across these variables, respectively. These results suggest that our HOB shock effectively generates a business cycle by inducing simultaneous increases in durable goods consumption, investment, and hours worked. Additionally, Table 3 reveals that the HOB shock accounts for between 35 and 40% of the volatility of durable goods and investment at the one-year horizon and 75% of the volatility of hours worked. This highlights the significant role of HOB shocks in driving the business cycle. Interestingly, just as for industrial production, the HOB shock does not exert a long run effect on neither the consumption of durables nor the gross private domestic investment, while no long run restriction was imposed. This reinforces the view that these HOB shocks have purely transitory effects.

Figure 6: Impulse Response to a HOB Shock: Additional Business Cycle Variables

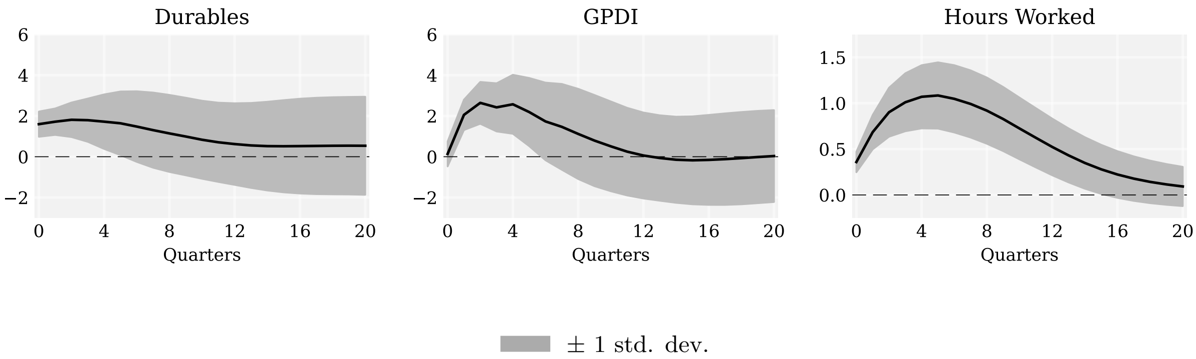

Note: GPDI: Gross Private Domestic Investment

In contrast, inflation remains largely unaffected, with its response insignificant even at the 68% confidence level. Accordingly, its impact on inflation volatility remains muted, with con-

[17] Note that no long-run restriction was imposed on the VAR. Furthermore, also note that industrial production was evaluated in first-difference, implying that, by construction, any shock is left free to exert a permanent effect on its level in the long run

Table 3: FEVD: Additional Business Cycle Variables

	Impact	1-Quarter	1-Year	2-Years	5-Years	Long-Run
Business Cycle Indicators						
Durables	39.83	46.84	43.41	33.72	20.23	4.22
Investment (GPDI)	0.82	27.51	34.71	24.42	13.12	2.99
Hours Worked	55.16	75.24	74.16	71.45	62.65	6.87
Nominal Variables						
Inflation Rate	4.12	4.13	3.88	5.90	4.60	1.56
Federal Fund Rate	65.91	76.70	73.00	55.40	26.75	4.53

Note: GPDI: Gross Private Domestic Investment, Inflation refers to CPI inflation.

Figure 7: IRF in the Nominal Side

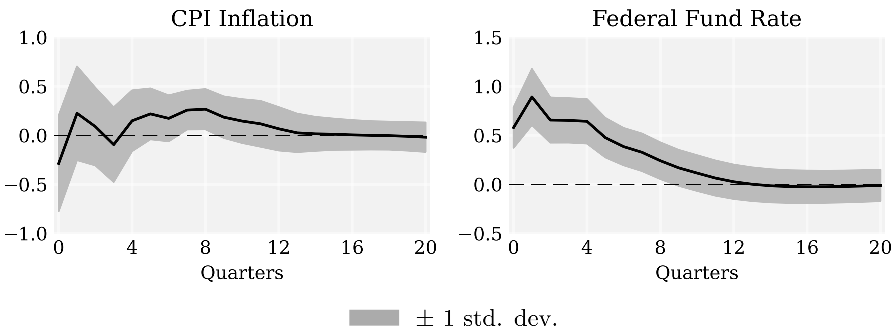

tributions below 5% across all horizons. The HOB shock is therefore non-inflationary and hence not confounded with a monetary policy shock. This finding is also reminiscent of the non-inflationary real demand shock interpretation of HOB shocks (see e.g. Angeletos, 2018).

A version of the VAR in the lines of Angeletos et al. (2020)[18] featuring the expectation error, GDP, consumption (non durables+services), investment (fixed private investment+durables), the unemployment rate, hours worked, the inflation rate and the federal fund rate yields very similar results. The shock effectively induces a business cycle, simultaneously raising GDP, consumption, investment, and hours worked, while reducing the unemployment rate and minimally impacting the inflation rate. The shock also explains a sizeable share of economic fluctuations. For instance, it accounts for about 40% of unemployment and GDP , 50% of hours worked, about 25% of investment and 70% of consumption at the one year horizon.

All these dynamics are a mix of the dynamics taking place at various frequencies. In the next section we explore the extent to which HOB shocks indeed account for business cycle fluctuations.

[18]The detailed results are presented in Appendix C. Although this version of the VAR model produces interesting and corroborating results, it was not chosen as our preferred specification due to its reliance on variables employed for constructing our instrumental variables.

2.2 The HOB Shock as a Main Business Cycle Shock

This section explores the relationship between our methodology and that proposed by Angeletos et al. (2020), who identified a main business cycle (MBC) shock —the shock that accounts for the bulk of macroeconomic aggregate fluctuations within business-cycle frequencies (6-32 quarters). More specifically, the main business cycle shock is identified as the shock that contributes the most to the volatility of a particular variable over the business cycle frequency band $(2\pi/32, 2\pi/6)$. This approach is similar to the "max-share" method developed in Faust (1998) and Uhlig (2004), but is formulated in the frequency domain.[19]

We employ this approach on our baseline VAR,[20] which encompasses the forecast error of unemployment ($\zeta^u_{t|t-1}$), the unemployment rate (u_t), the industrial production index (IP_t) and the rate of capacity utilization (Util_t). Panel (a) of Figure 8 compares the responses to an HOB shock (plain line) with the responses to an MBC shock when the latter is identified by targeting either the expectation error, the unemployment rate or the industrial production index. The outcomes align with those reported in Angeletos et al. (2020).

Figure 8: Extended VAR on MBC component: IRFs

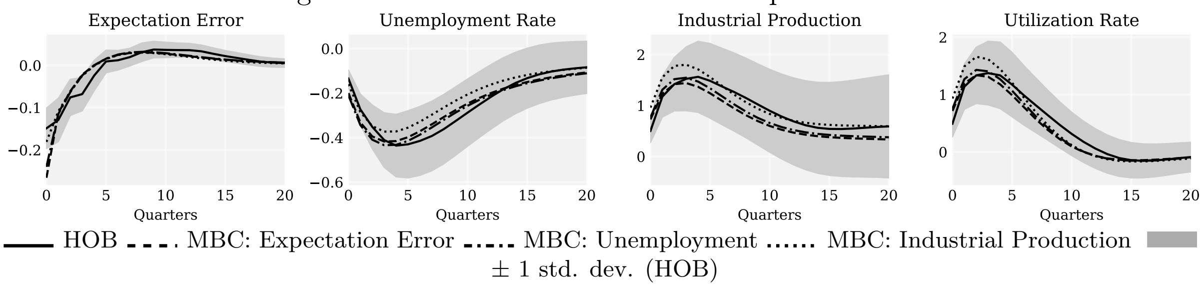

As in Angeletos et al. (2020), the MBC shocks, obtained by targeting the unemployment rate, industrial production or capacity utilization, are all interchangeable in that they all indeed induce the same business cycle, characterized by an increase in industrial production and the utilization rate and a decrease in the unemployment rate and the expectation error on unemployment rate. The figure also presents the MBC obtained by targeting the expectation error on the unemployment rate. The results are essentially identical. One may argue that this observation should not come as a surprise, as targeting either the unemployment rate or its forecast error should yield the same shock. This holds true in a rational expectations framework, but there is actually no a priori reason it should hold in the survey of professional forecasters. For instance, within our sample, the first-order (resp. second-order) autocorrelation of the expectation error $\zeta^u_{t|t-1}$ is 0.51 (resp.0.36). Similarly, the Granger causality tests yield p-values of 0.0 for the expectation error, the unemployment rate and first McCracken factor, indicating that these variables significantly Granger-cause the expectation error. These findings together suggest a

[19]See Appendix D.1 for details.

[20]Angeletos et al. (2020) consider all variables in level. We stick to this specification in this section, but also report results with the industrial production in difference in Appendix D.2. If anything, the difference specification reinforces our findings.

departure from rational expectations.

Table 4: FEVD: Baseline VAR (MBC Shocks)

	Impact	1-Quarter	1-Year	2-Years	5-Years	Long-Run
Exp. Error ($\zeta^u_{t\|t-1}$)						
Exp. Error ($\zeta^u_{t\|t-1}$)	90.84	90.25	88.97	88.74	88.71	88.67
Unemployment Rate	76.12	86.48	90.78	90.36	88.48	87.30
Industrial Production	51.08	58.77	62.09	60.66	53.34	34.33
Capacity Utilization	50.47	58.29	62.29	62.65	62.77	62.28
Unemployment Rate						
Exp. Error ($\zeta^u_{t\|t-1}$)	74.06	76.67	76.63	77.06	77.88	77.91
Unemployment Rate	84.00	93.40	98.10	97.84	95.68	94.42
Industrial Production	56.53	66.04	70.87	69.66	61.55	39.73
Capacity Utilization	55.85	65.48	71.04	71.73	71.76	71.17
Industrial Production						
Exp. Error ($\zeta^u_{t\|t-1}$)	42.36	48.47	50.12	51.22	52.30	52.29
Unemployment Rate	48.94	61.43	70.27	70.07	66.80	66.46
Industrial Production	88.54	95.56	98.07	96.47	89.29	66.91
Capacity Utilization	85.52	92.92	96.04	95.46	94.48	93.61

More importantly, the impulse responses to an MBC shock are almost identical to those to an HOB shock. This points to the existence of a strong connection between the MBC shock and the HOB shock identified by the proxy-VAR. It, however, does not imply that the HOB shock is the sole cause of the MBC shock, as the MBC shock likely arises from a combination of many structural shocks. Instead, these findings suggest that the HOB shock likely plays a significant role in driving the dynamics of the MBC shock. In order to investigate this issue further, Figure 9 reports the bandpass filtered (6-32 quarters) history of the variables conditional on the HOB shock only, and on the MBC (targeting the unemployment rate). As can readily be seen from the figure, the two variables comove positively and generate the same business cycle.

Figure 9: MBC vs HOB: Conditional History

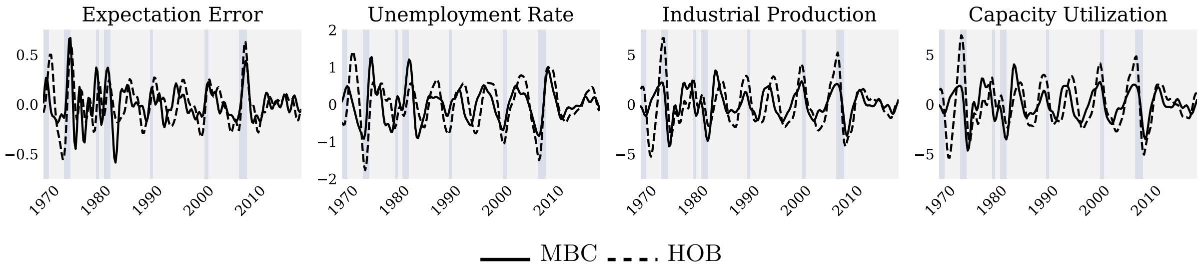

Table 5 indicates that the two conditional histories are quite strongly correlated over the whole sample. The correlation ranges from 0.53 for the expectation error to 0.60 for the industrial production index. This correlation evolved over the sample. In the seventies, the main business cycle shock was less connected to HOB shocks than in the later part of the sample. For instance, while the correlation between the two conditional histories of unemployment was 0.3

Table 5: HOB as MBC

Sample	$\zeta^u_{t\|t-1}$	u_t	IP_t	Util_t
Whole Sample	0.53	0.56	0.59	0.57
(1968Q4-2019Q4)	(0.27)	(0.31)	(0.34)	(0.33)
Pre-Volcker	0.53	0.31	0.36	0.35
	(0.26)	(0.07)	(0.11)	(0.10)
Post-Volcker	0.60	0.70	0.71	0.71
	(0.35)	(0.49)	(0.50)	(0.50)
1980Q1-2019Q4	0.56	0.69	0.70	0.69
	(0.31)	(0.48)	(0.48)	(0.47)
1990Q1-2019Q4	0.68	0.78	0.78	0.77
	(0.46)	(0.61)	(0.60)	(0.58)
2000Q1-2019Q4	0.83	0.92	0.91	0.89
	(0.68)	(0.85)	(0.82)	(0.80)

Note: MBC shock obtained by targeting the unemployment rate. The table reports the correlation between the history of each variable conditional on the MBC and the corresponding history conditional on our HOB shock. The number in parenthesis is the R^2 of the projection of the history conditional on the MBC shock onto the history conditional on the HOB shock.

in the pre-Volcker period, it reached 0.9 in the 2000Q1-2019Q4 period. The table also reports, in parenthesis, the R^2 of the projection of the history of each variable conditional on the MBC shock onto the corresponding history obtained from the HOB shock. In the last sub-sample, the R^2 is above 0.8 for the same business cycle variables, indicating a very strong connection between the MBC and the HOB shock.

2.3 Further Implications

In this section we derive additional implications of the HOB shock for the economy, delving into its effects on financial variables and uncertainty. We also go back to the relationship between HOB shocks and confidence. Finally, we investigate whether our shock can be confused with another shifter of expectations (news shock).

2.3.1 HOB Shocks and Uncertainty

We first examine the impact of the HOB shock on uncertainty. To do so, we augment our benchmark specification with the macroeconomic uncertainty measure developed by Jurado et al. (2015) (JLN hereafter). This measure captures the conditional volatility of the purely unforecastable component of the future value of a broad range of macroeconomic time series. It essentially gauges whether the economy has become more or less predictable. JLN constructed this measure for various forecasting horizons. In the sequel, we incorporate measures of uncertainty at the 1-month, 3-month, and 1-year forecast horizons into our VAR model.

Figure 10 depicts the IRF of the macroeconomic uncertainty measure for various forecasting horizon to a HOB shock, introduced one at a time in our benchmark VAR. To save space, we focus solely on the response of the uncertainty measures (and later FEVDs). However, it's worth noting that the response of the other VAR variables is qualitatively and quantitatively very similar to those obtained in the benchmark VAR presented in Section 2.1. A positive

Figure 10: IRF of Uncertainty Measures

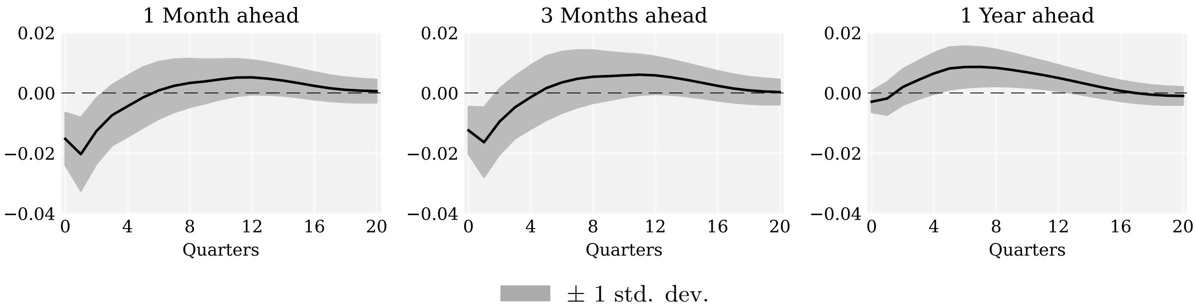

HOB shock triggers a decline in uncertainty, suggesting that it acts as a coordinating device, inducing agents to align their forecasts and consequently reduce the JLN's uncertainty measure. However, the effect varies across the horizon of the uncertainty measure. A positive shock elicits a stronger response for shorter-horizon uncertainty measures.

Table 6: FEVD: The Broader Economy

	Impact	1-Quarter	1-Year	2-Years	5-Years	Long-Run
Uncertainty						
JLN (1-Month ahead)	29.50	29.26	18.29	11.40	10.02	8.30
JLN (3-Months ahead)	23.07	20.60	10.80	7.32	7.18	5.80
JLN (1-Year ahead)	6.45	3.15	4.67	10.06	10.56	7.32
Risk Premium						
JLN (1-Month ahead)	29.46	29.67	19.96	15.35	16.47	16.06
Risk Premium	13.15	16.88	22.40	19.60	16.19	15.05
Credit Spread						
Credit Spread	49.52	49.81	51.08	48.27	43.09	18.63
Credit Spread + Labor Productivity						
Credit Spread	43.33	44.53	47.68	47.24	44.94	36.39
Labor Productivity	12.83	12.90	6.38	3.82	2.62	11.68
Confidence						
12-month Index	8.28	13.10	9.77	7.68	6.79	5.81
5-year Index	5.31	11.99	10.19	7.66	5.59	3.56
Synthetic Index	13.41	20.06	14.74	10.18	7.93	6.83

This is corroborated by the FEVD reported in Table 6, where the HOB shock accounts for about 30% of the volatility of the 1-month ahead uncertainty upon impact. This contribution drops to 23% for the 3-month ahead uncertainty and further drops to about 6.5% for the 12-month ahead uncertainty. This finding reinforces the notion that the HOB shock primarily affects expectations about short-run economic outcomes. The effect on uncertainty is short-lived and fades away rapidly, as reflected both by the IRFs and the FEVDs. For example, the contribution of the HOB shock to the volatility of the 1-month ahead uncertainty is one third its short-run contribution after 2 years (11%).

2.3.2 HOB Shocks and Financial Conditions

We first consider a variant of our benchmark VAR that incorporates the 1-month ahead macroeconomic uncertainty measure and a risk premium indicator,[21] in order to shed light on how financial markets value the drop in volatility that follows a positive HOB shock. Figure 11 illustrates the responses of uncertainty and the risk premium to the HOB shock. The uncertainty measure declines in response to the HOB shock and remains significantly influenced by the shock. Consequently, the risk premium falls upon impact and stays below its initial level for about two years. As reported in Table 6, the shock explains approximately 20% of the risk premium's volatility at the same two-year horizon. In the next section we delve further in the implications of the shock for financial conditions.

Figure 11: Uncertainty and the Risk Premium

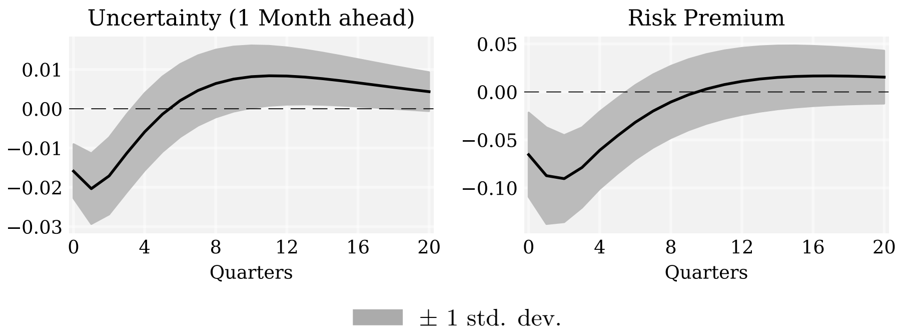

Financial conditions, particularly the severity of financial frictions, can be represented by the credit spread, measured as the gap between Moody's BAA corporate bond rate and the 10-year maturity bond. We augment our benchmark VAR by incorporating this credit spread. As depicted in Figure 12, a positive HOB shock induces a persistent decrease in the credit spread, indicating an easing of credit conditions, which aligns with the wave of optimism associated with the shock. The shock is a significant contributor to the credit spread's volatility, accounting for approximately 40% of its short-run volatility (see Table 6).

While these findings support the notion of a shock that stimulates economic activity, they pose a challenge, as they suggest that our HOB shock could be confused with a financial shock. To the extent that financial frictions and financial shocks impedes the efficient allocation of resources, a financial shock should have a non-negligible effect on productivity (see e.g. Moll, 2014). Therefore, we augment our latest VAR to incorporate labor productivity. Figure 12 illustrates the response of labor productivity to our shock and indicates that productivity remains largely unaffected. This is corroborated by Table 6. The shock explains approximately 13% of labor productivity volatility upon impact and only about 6% after one year. In other

[21] The risk premium is computed as the difference between the Moody's BAA and AAA corporate yield rates.

Figure 12: Credit Spread

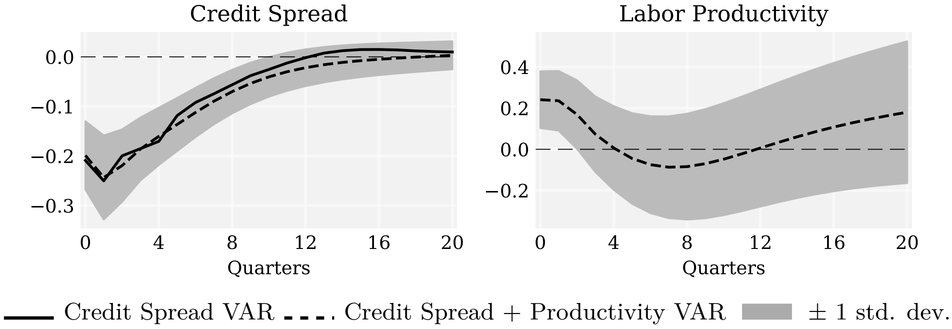

words, our shock is unlikely to be confounded with a financial shock[22] and instead triggers an improvement in financial conditions.

2.3.3 HOB is not Confidence

Previous studies exploring the role of belief and sentiment shocks (see *e.g.* Barsky and Sims, 2012; Forni et al., 2017; Fève and Guay, 2019; Lagerborg et al., 2023) heavily relied on confidence indices, such as the Michigan Consumer Sentiment Index, to facilitate their identification. However, we caution readers that HOB shocks may not provide substantial information about such indices, and relying solely on these indices for identification may not be appropriate.

We start by incorporating various measures of consumer sentiment into our benchmark VAR. The University of Michigan provides several indices that capture insights into consumer confidence. In this paper, we will utilize three of these indices. The first index indicates how consumers feel confident about short-run (12 months) economic outcomes, the second index relates to longer run (5 years) economic outcomes, the last one synthetizes the two indices.[23] Figure 13 depicts the IRF of the three confidence indices to a sentiment shock, introduced one at a time in our benchmark VAR. As in the previous section, we focus solely on the response of confidence measures (and later FEVDs). However, it's worth noting that the response of the other VAR variables is qualitatively and quantitatively very similar to those obtained in the benchmark VAR presented in Section 2.1. Inspection of the average response of household confidence, across all indices, indicates that the HOB shock triggers a surge of optimism, consistent with theoretical expectations. However, the magnitude of the effect varies depending on the index's horizon. A positive shock elicits a much stronger response for shorter-horizon confidence indices, further corroborating that the HOB shock measures sentiment about near-

[22]This is not the case for a transitory investment shock. For instance, when we use it as an instrument (See Section 2.3.5 for details about the construction of this shock) in place of our HOB shock, we indeed see a mild decrease in the credit spread but it is accompanied by a persistent increase in labor productivity.

[23]The first one, which we will refer to as the 12-month index, summarizes responses to the question (`BUS_12R`): "*Now turning to business conditions in the country as a whole, do you think that during the next twelve months we'll have good times financially, or bad times, or what?*" The corresponding question for the 5-year index (`BUS_5R`) reads: "*Looking ahead, which would you say is more likely–that in the country as a whole we'll have continuous good times during the next five years or so, or that we will have periods of widespread unemployment or depression, or what?*". The last index gathers information about 5 questions related to the overall economic situation. Greater details are available from https://data.sca.isr.umich.edu/fetchdoc.php?docid=75432.

Figure 13: IRF of Confidence Indices

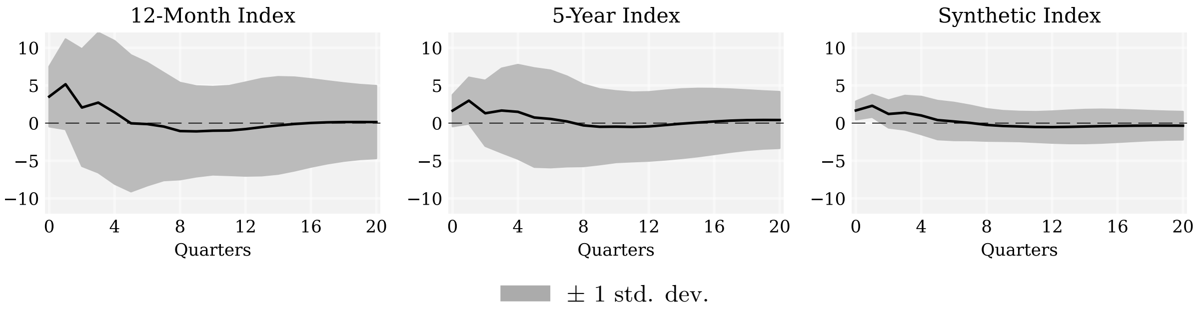

term economic outcomes. Importantly, these responses are imprecisely estimated, which, in this case, would lead to reject the existence of an effect at the 68% confidence level. As shown in lower panel of Table 6, the HOB shock's contribution to the overall volatility of confidence is small across all considered horizons. In other words, the HOB shock does not provide much information about the evolution of the confidence index, as measured by the Michigan index. This observation holds true for all the indices considered. Even at its peak, the HOB shock contributes less than 15% to the synthetic index volatility, and this contribution falls below 10% for both the 12-month and 5-year indices.

Since, the HOB shock is not informative about the Michigan index, there may not be enough information to properly identify the HOB shock from this variable. In order to evaluate this statement, we replace the expectation error by one of the three confidence measures. Then we conduct a standard Montiel Olea and Pflueger (2013) weak instrument test across our 200 draws of the HOB shocks from the model. Figure 14 reports the empirical cumulative distribution function of the test statistics across our 200 experiments. The results indicate that the all

Figure 14: Weak Instrument Test for Confidence Based Identification

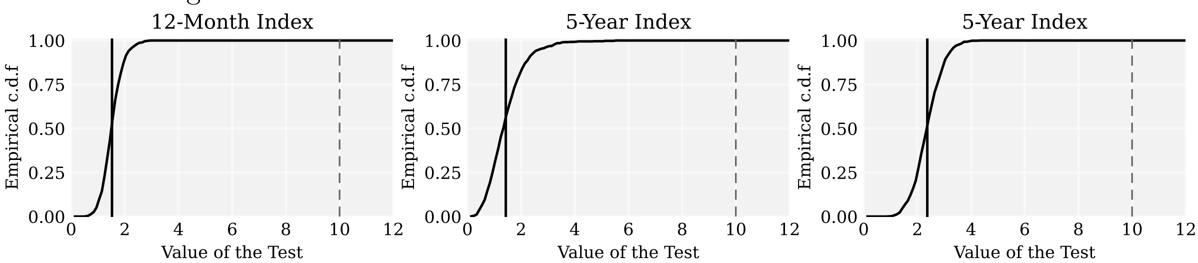

Vertical line: Average value of the Test , Vertical dashed line: rule of thumb threshold (10).

of the distribution's mass lies below the 10 rule-of-thumb threshold, suggesting that the null hypothesis of a weak instrument is unlikely to be rejected. In other words, our instrument cannot be effectively used in conjunction with the confidence index to elucidate the economic dynamics following a HOB shock.

2.3.4 Can it be news shock?

Thus far, we have maintained that our procedure identifies a shifter of beliefs about short-run economic outcomes. In this section, we examine whether it could be confounded with another shifter of beliefs about economic outcomes: news shocks about TFP (see Beaudry and Portier, 2014, for a review).

We begin by augmenting our benchmark VAR with Total Factor Productivity (TFP) and the real stock price index –the two variables utilized by Beaudry and Portier (2006) to identify news shocks.[24] We then investigate the degree to which our HOB shock explains TFP volatility. To account for the long-run behavior of both TFP and stock prices, we consider three distinct specifications. The first specification mirrors our benchmark specification, assuming that industrial production, TFP, and the stock price are all represented in log-differences. In the second specification, all variables are expressed in (log-)levels. The final specification, VECM, incorporates the first difference of (log-)TFP, the residual of the cointegration relationships between TFP and the log of the real stock price (resp. industrial production). Figure 15 reports the IRFs of TFP and the stock price index to a HOB shock in our three specifications.[25] TFP shifts up on impact and eventually seems to settle on a higher level. In the VECM case, the TFP IRF mirrors that reported in Beaudry and Portier (2006), barely responding on impact and eventually reaching a higher level in the longer run, suggesting some similarity to the response to a news shock. This resemblance is further accentuated by the immediate increase in the stock price index, which effectively front-loads the effects of a positive news shock.

Figure 15: IRF of TFP and the Stock Price Index

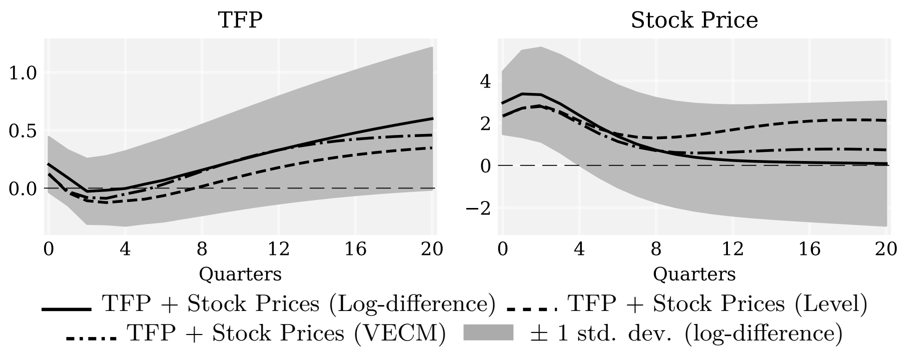

Does our shock effectively capture a news shock? Table 7 summarizes the variance contribution of our HOB shock to the volatility of TFP and the stock price index. The results unequivocally demonstrate that our shock does not contribute significantly to TFP volatility. This is particularly evident in the cointegration specification, where the contribution of the shock is less than 5% at any horizon and around 1% in the long run. Similarly, the shock's contribution to the volatility of stock prices fades away beyond the first quarter. Taken together, these observations effectively rule out the possibility of the shock being a news shock.

To further investigate the potential confounding effect of the news shock on our HOB shock, we replicate our identification procedure using the news shock, generated from the DSGE model,

[24] We use the measure of TFP corrected for utilization provided by Fernald (2012). The real stock price index is obtained from Shiller's website (`https://www.econ.yale.edu/~shiller/`).

[25] Like in the previous sections, the results pertaining to the variables present in our benchmark case survive. We therefore focus on TFP and the stock price index in this section.

Table 7: FEVD: News

	Impact	1-Quarter	1-Year	2-Years	5-Years	Long-Run
TFP and Stock Prices						
First Difference						
TFP	3.46	2.44	1.09	1.09	7.69	20.60
Stock Price	25.74	20.88	14.92	9.59	4.56	7.45
Level						
TFP	1.60	1.17	1.57	1.33	5.84	8.09
Stock Price	16.12	13.50	10.65	7.66	6.38	4.60
Cointegration						
TFP	1.44	1.00	0.92	1.01	8.54	1.55
Stock Price	16.29	13.46	10.02	6.48	3.29	3.68
News Shocks						
Exp. Error ($\zeta^u_{t\|t-1}$)	4.82	5.92	7.58	7.61	7.86	7.86
Unemployment Rate	4.62	6.81	9.82	11.17	13.97	23.56
Industrial Production	8.59	9.06	8.01	7.10	4.92	24.44
Capacity Utilization	8.65	9.25	8.82	9.24	10.81	25.79

as the instrument. Specifically, we follow the methodology outlined in Section 1.2. Employing the 200 parameter realizations from the posterior density of the structural parameters we used to derive the HOB shock, we utilize the Kalman smoother to recover 200 realizations of the news shock from the solution of the model. Subsequently, we apply our identical proxy-VAR approach, substituting the news shock for the expectation error as the instrument. Figure 16 illustrates the IRFs to the news shock. Strikingly, none of these IRFs is precisely estimated.

Figure 16: IRF to a (DSGE-) News Shock

Furthermore, the shock's contribution to the volatility of each variable, as shown in Table 7, is negligible, therefore suggesting that our HOB shock cannot be confused with news —the other belief shifter. One reason is that news shift beliefs about the longer run economic outcomes. This can actually be seen from Table 7 —the maximal contribution of the news shock is found in the very long run, although it remains limited.

A second reason is that this shock ought not to be informative about the expectation error. To assess this, we conduct the Montiel Olea and Pflueger (2013) weak instrument test across our 200 draws of the HOB shocks from the model. Figure 17 reports the empirical cumulative distribution function of the test statistics across our 200 experiments. The results indicate that

Figure 17: Weak Instrument Test: News

Vertical line: Average value of the Test , Vertical dashed line: rule of thumb threshold (10).

all of the mass of the distribution lies below the 10 rule-of-thumb threshold, suggesting that news shocks are indeed a weak instrument for the expectation error. All in all, our HOB shock does not capture information that could be confused for a news shock.

2.3.5 Other Fundamental Shocks?

Equation (9) in Section 1.2 suggests that the expectation error ought to be shifted by any of the fundamental shock driving the dynamics of the DSGE model. We therefore recover all the shocks from the DSGE model and use each of them as a potential instrument in our proxy-VAR. In particular, we run the Montiel Olea and Pflueger (2013) weak instrument test across our 200 draws of each of the shocks from the model.

Figure 18: Weak Instrument Test

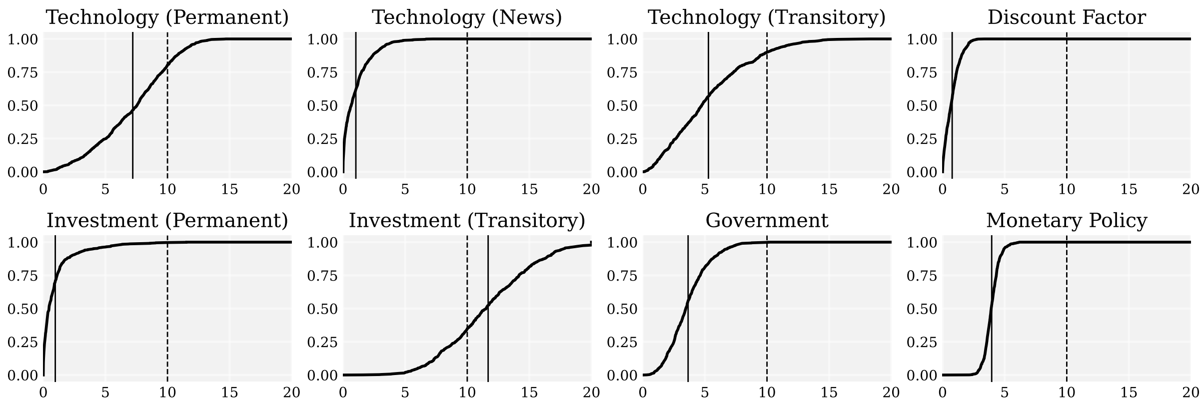

Vertical plain line: Average value of test, Vertical dashed line: rule of thumb threshold (10)

Figure 18 presents the cumulative distribution function of the weak instrument test across the range of models we consider. In all case but the investment transitory shock, the average of the test statistics falls way below the conservative rule of thumb threshold of 10, demonstrating that none of the remaining structural shocks are promising instruments. In the case of the transitory investment shock, although the average statistics lies above 10, there is a significant mass of models (about 30%) for which the weak instrument test is below the threshold.[26] One

[26]Assume that the transitory investment shock be recovered from a complete information version of the model

may however wonder to what extent there can be a potential confounding effect between the HOB shock and this investment shock. However, the two shocks are mutually orthogonal, with an average correlation of -0.04 across models. Taken together, all these observations lead us to conclude that none of the other shocks present in the DSGE model is a likely candidate for a relevant instrument.

3 Robustness

This section investigates the robustness of our findings to variations in the information and the identification approach employed to recover the HOB shock. Additionally, we conduct a placebo experiment to further examine the validity of our instrument strategy.

3.1 Alternative Expectation Errors

This section investigates the robustness of our main findings to changes in the expectation error. We delve into two distinct aspects of the expectation error's definition: the timing of the expectation and the variable for which we calculate the expectation error.

Timing: Figure 20 contrasts the impulse response functions of the principal variables in our benchmark VAR to those generated by a VAR in which the one-step-ahead expectation error is not calculated as $\zeta^u_{t|t-1} = u_t - u^e_{t|t-1}$ but rather as $\zeta^u_{t+1|t} = u_{t+1} - u^e_{t+1|t}$ (dashed line).[27] While these two measures would be identically distributed and convey identical information in a purely rational expectation stationary environment, this is not likely the case in survey data. This is illustrated in Figure 19 which depicts how two expectation errors are not identically distributed. The distribution of $\zeta^u_{t+1|t}$ is slightly shifted to the left, suggesting that individuals tend to overestimate unemployment rates in economic downturns to a greater extent when forming expectations in period t compared to their expectations formed in period $t-1$.

Figure 20 indicates that this does not markedly affect the dynamics of the economy following a HOB shock. The impulse response functions of the VAR variables are only marginally affected by this change in timing. Similarly, as indicated by the left panel of Table 8, the contribution of the HOB shock remains essentially unaffected. The shock accounts for about 70% of the unemployment volatility at the 1-year horizon, 50-55% of that of industrial production and 40-45% for capacity utilization.

We also consider the expectation error associated with the nowcast of the unemployment rate — namely $\zeta^u_{t|t} = u_t - u^e_{t|t}$. Figure 20 reveals that, consistent with the previous case, the dynamics of the VAR variables to a HOB shock are only marginally affected by the change in

—therefore ignoring HOB shocks. In that case, a sizeable part of the variation captured by HOBs is transmitted to this transitory investment shock, leaving the econometrician with the impression that this shock indeed matters a lot for business cycle volatility, while the data overwhelmingly favor the incomplete information version (see Section 5.4 in Angeletos et al. (2018)). *Why is the information captured by the transitory investment shock?* The investment decision trades-off the cost of investing in capital accumulation against the expected future benefits of this investment, which are precisely affected by the transitory investment shock as of the short-run. Therefore, it's natural that HOB shock variations are reflected in these investment shocks. See Section E in the Appendix.

[27] In the latter case, all variables of the VAR are shifted by one period to preserve the timing of expectations.

Figure 19: Distribution of Expectation Errors: Timing

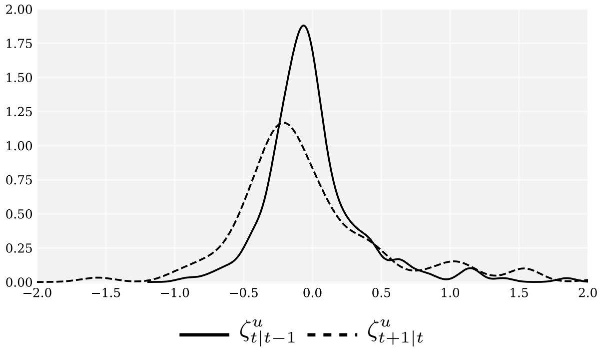

Figure 20: Impulse Response to a HOB Shock: Varying the timing of Expectations

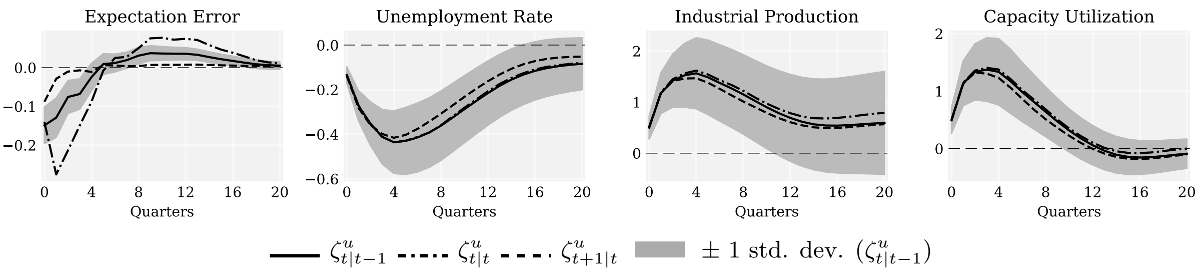

Table 8: FEVD: Varying Expectations

	Impact	1-Quarter	1-Year	2-Years	5-Years	Long-Run
Error on Unemployment ($\zeta^u_{t+1\|t}$)						
Exp. Error	51.60	69.98	53.70	48.77	49.04	48.23
Unemployment Rate	51.64	70.51	68.94	62.99	46.60	28.64
Industrial Production	35.49	56.90	57.53	49.65	33.01	0.50
Capacity Utilization	34.18	54.42	53.60	47.38	40.59	11.04
Error on Unemployment ($\zeta^u_{t\|t}$)						
Exp. Error	36.10	37.11	36.05	35.70	35.83	30.07
Unemployment Rate	53.20	74.24	71.86	58.92	40.10	23.06
Industrial Production	33.10	53.68	51.92	42.60	25.16	1.15
Capacity Utilization	31.82	51.22	47.95	40.43	33.92	12.33
Error on Output ($\zeta^y_{t\|t-1}$)						
Exp. Error	43.47	43.40	39.28	36.26	35.82	32.45
Unemployment Rate	57.34	72.46	72.50	66.08	50.06	31.73
Industrial Production	36.30	52.73	51.70	44.43	28.28	5.54
Capacity Utilization	35.14	50.64	48.29	42.60	37.60	16.66

the definition of the expectation error. Similarly, as evidenced in the right panel of Table 8, the HOB shock accounts for a substantial proportion of the short-run volatility of unemployment, industrial production, and capacity utilization.

Support Variable: We now explore the implications of utilizing the expectation error on real output instead of unemployment. One potential challenge of utilizing the real output expectation error stems from the multiple changes in the base year for computing the GDP deflator across the sample period, leading to artificial shifts in the expectation error. To address these jumps, we project the expectation error on a set of dummy variables that take the value 1 in years where a base year change occurred and 0 otherwise, and employ the residual of this regression as our expectation error. This approach completely eliminates the variation in those years, which may raise concerns. However, we also investigated alternative approaches to rectify these jumps, including replacing those observations with an average of the neighboring observations. The findings are only marginally affected by such alterations.

Figure 21: Impulse Response to a HOB Shock: Output vs Unemployment Expectation Error

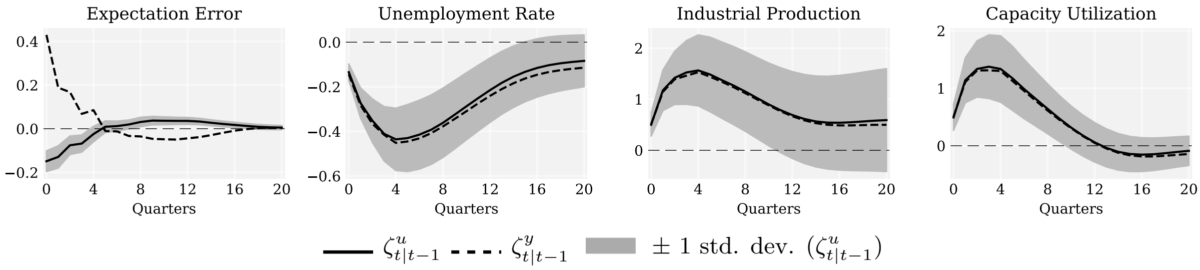

Figure 21 and Table 8 present, respectively, the impulse response functions (IRFs) of the VAR variables to a HOB shock and its contribution to the volatility of the same variables. The responses are remarkably similar regardless of whether the unemployment or the output expectation error is used to identify the HOB shock. The only distinction lies in the sign of the response of the expectation error itself. This is unsurprising, given the negative correlation between output and unemployment. In both cases, this suggests that agents tend to underestimate the response of output or the unemployment rate.

Table 8 indicates that the influence of the HOB shock on the volatility of the variables remains consistent when the definition of the expectation error shifts. If anything, the impact on the unemployment rate strengthens, reaching 72% at the one-year horizon, further reinforcing the role of HOB shocks in shaping the business cycle.

3.2 Addressing Potential Contamination of Shocks

Our identification strategy rests on the assumption that the innovation of the HOB shock, our instrument, is orthogonal to the innovation of the other fundamental shocks (technology, investment, discount rate, monetary policy, ...). This is, by construction, true in the DSGE model. However, one may be concerned that each of the 200 shocks is recovered applying

Kalman techniques to a small sample of macroeconomic time series. In that case, one cannot rule out the possibility that the estimated HOB shock be correlated with other fundamental shocks. In order to address this issue, for each of our 200 models, we project the innovation of the HOB shock onto the other fundamental shocks and use the residuals of that regression as our proxy variable. These residuals are then, by construction, orthogonal to the other shocks.

Figure 22: Impulse Response to a HOB Shock: Output vs Unemployment Expectation Error

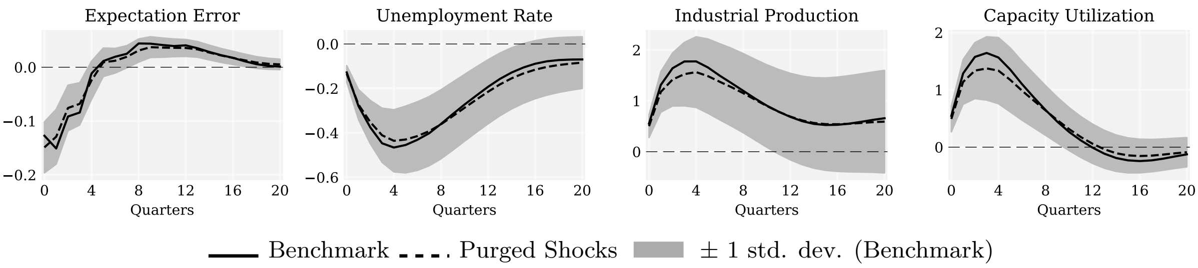

Table 9: FEVD: Purged Shocks

	Impact	1-Quarter	1-Year	2-Years	5-Years	Long-Run
Exp. Error ($\zeta^u_{t\|t-1}$)	40.22	51.89	48.15	46.76	48.22	47.58
Unemployment Rate	50.28	74.57	79.40	70.32	48.95	30.20
Industrial Production	42.59	69.37	72.75	59.57	36.47	2.14
Capacity Utilization	41.86	68.09	71.59	61.48	53.30	17.33

Figure 22 compares our benchmark IRF (plain line) to the IRF obtained from these "purged" HOB shocks (dashed line). Table 9 reports the corresponding variance decomposition. Inspection of the figure and the table clearly indicates that the results are not affected by applying this projection, which in turn suggests that our instrument does not suffer a contamination problem.

3.3 Alternative Ways to Recover Dynamic Causal Effects

To assess the robustness of our findings to the identification strategy, we first examine the responses to a HOB shock identified using a recursive timing approach inspired by Sims (1980). This approach imposes a sequential ordering on the variables' responses to the various shocks in the VAR. Specifically, we expand our benchmark VAR by directly incorporating our instrument into the vector of endogenous variables, ordering it first in the vector (see Plagborg-Møller and Wolf, 2021). The HOB shock then represents the first "structural" shock. The impulse response functions for the variables in response to the HOB shock are subsequently derived from the first column vector of the Cholesky decomposition of the VAR's covariance matrix. Panel (a) of Figure 23 compares the IRFs obtained from our benchmark VAR to those generated using the recursive timing approach.

The responses exhibit similarities to our benchmark in that the HOB shock triggers a decline in the expectation error and the unemployment rate, while simultaneously boosting industrial production and capacity utilization. However, some differences emerge. First, the recursive timing assumption tends to produce less persistent dynamics. For instance, the peak in the response of industrial production is reached after 2 quarters in the recursive timing identification, while the economy reaches its peak after 1 year in the proxy-VAR. Second, the responses of unemployment, industrial production, and capacity utilization eventually reverse direction under Cholesky identification. However, over all, both identification approaches recover very similar dynamics.

Figure 23: Alternative Identification Strategies: IRFs

(a) Cholesky Decomposition

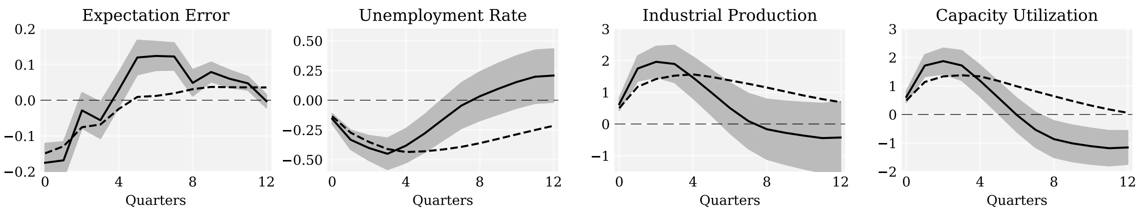

(b) LP-IV

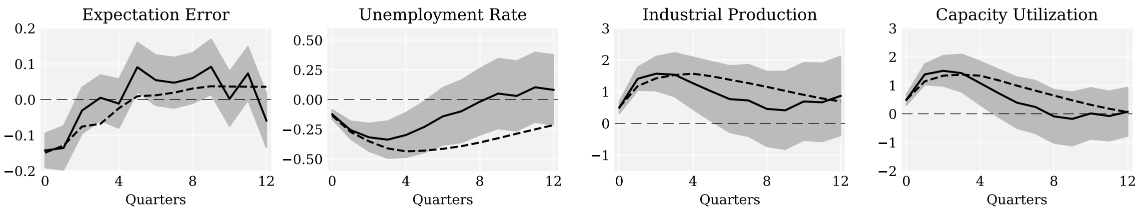

Alternative Idenfication - - - Benchmark VAR ± 1 std. dev. (Alternative)

For robustness analysis we also employed a local projection estimator to assess dynamic causal effects, which entails less stringent assumptions. We utilized a straightforward version of the local projection estimator, where the impulse responses at horizon h are derived from the estimates of $\{\varphi_h\}_{h=0}^{H}$ obtained from the regression

$$Y_{j,t+h} - \overline{Y}_{j,t-1} = \alpha_{i,h} + \varphi_{i,h}\zeta^u_{t|t-1} + \Gamma_{i,h}(L)Y_{t-1} + \nu^x_{i,j,t+h}$$

where $Y_{j,t}$ represents the j-th variable in vector Y_t; $\overline{Y}_{j,t-1}$ is a variable that takes the value $Y_{j,t-1}$ is the variable is growing (Industrial production in our benchmark VAR) and 0 otherwise; ν^x denotes the residual of the regression. In this regression, $\zeta^u_{t|t-1}$ is instrumented by the HOB shock drawn from the model. Since we utilize a variety of models – hence, a variety of instruments – we estimate one such equation for each model i we draw. Note that this regression makes no invertibility assumption. It however maintains, as in the VAR, the assumption of exogeneity of the instrument.[28]

Panel (b) of Figure 23 depicts the average IRF of the VAR variables across models (dashed line) in this regression alongside those in our benchmark VAR (plain curve). Despite the aforementioned flexibility, the impulse response functions presented in the figure exhibit remarkable similarities to those obtained from our benchmark VAR. The expectation error falls on impact,

[28]Note that we also relax the normalization assumption on the effect of the shock on the expectation error. The IRF is therefore rescaled by σ_{ε_1} as explained in Section 1.1.

along with the unemployment rate, while industrial production and capacity utilization rise. Inflation exhibits a muted response. These responses closely mirror those observed in the VAR, although some minor discrepancies emerge due to the reduced restrictions imposed in the local projection approach. Specifically, the response of the expectation error is strikingly similar throughout the entire trajectory in the local projection estimator compared to the proxy-VAR. But, the trough in the response of unemployment is reached earlier (after three quarters rather than four). Likewise, the peaks in the responses of industrial production and capacity utilization also occur earlier. Nevertheless, these minor variations aside, the responses remain remarkably similar to those obtained in the proxy-VAR, providing a reassuring validation.

3.4 A Placebo Experiment

This section introduces two placebo experiments that further demonstrate the validity of our instruments. In the first experiment, dubbed placebo I, we randomly reorder the values of our instrument across its realizations. In the second experiment, labeled placebo II, we randomly select a date T within the sample of each realization of our instrument. We then split the time series of the instrument into two subsequences: the first running from 1968Q4 to T and the second from T to 2019Q4. We subsequently swap these two blocks to create a new time series. We then replicate our identification using these artificial instruments, both under the placebo I and placebo II experiments, and compare the results with those of our benchmark VAR. The underlying idea is that should our instrument lack informativeness, it should not hold more information about the HOB shock than these entirely random instruments. Figure 24 compares

Figure 24: Placebo test

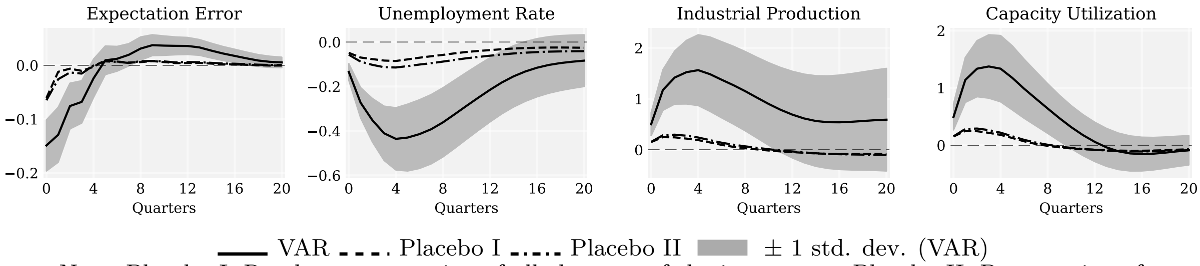

Note: Placebo I: Random permutation of all elements of the instrument, Placebo II: Permutation of 2 blocs of the instrument, cutoff randomly selected. The VAR includes the first two McCracken's factors.

the impulse response functions obtained from our benchmark VAR to those obtained both in the Placebo I and II experiment. The results are clear. The macroeconomic variables actually do not respond to the shock identified using the placebos. In fact, the placebo instruments are not sufficiently correlated with the true shocks to provide reliable estimates of their effects. Figure 25 depicts the empirical cumulative distribution function of the test statistics across draws of the placebo instruments. The figure shows that the distribution of the test statistics is heavily skewed to the left, with many of the values falling below the critical value, implying that the placebo instruments are weak instruments and lack sufficient information to effectively identify a meaningful HOB shock.[29]

Figure 25: Weak Instrument Test for Placebo

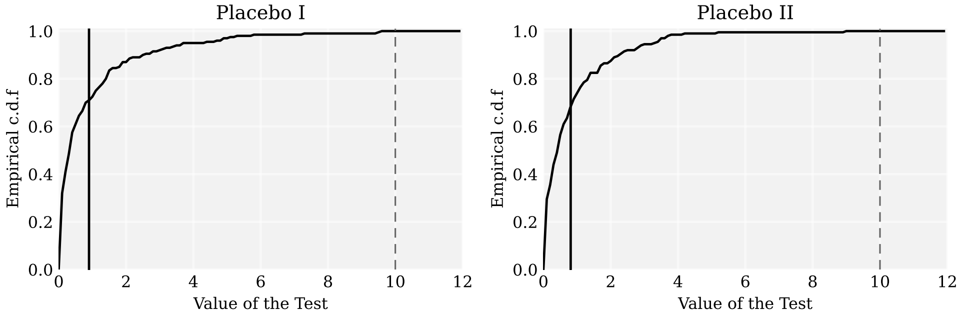

Dashed line: 95% threshold. Placebo I: Random permutation of all elements of the instrument, Placebo II: Permutation of 2 blocs of the instrument, cutoff randomly selected.

Therefore, our benchmark instrument captures significant variation associated with higher-order beliefs, eliciting aneconomic response that none of our placebo variations can explain.

4 Concluding Remarks

This paper examines the role of HOB shocks in driving business cycle fluctuations. The identification of such shocks is a challenging task for the econometrician as they do not have a direct observable counterpart in the data. In addition, popular identification schemes are not well-suited to identify such shocks. A notable exception are proxy-VARs, as exemplified by Lagerborg et al. (2023), which rely on an instrumental variable to identify the sentiment shock. In this paper, we employ a similar proxy-VAR approach but leverage exogenous variations in a variety of proxy measures of expectations generated by a DSGE model (see e.g. Angeletos et al., 2018).

Our findings highlight the significant impact of HOB shocks on the economy. Positive shocks trigger economic booms, characterized by a decrease in unemployment, a rise in industrial production and capacity utilization, and a persistent hump-shaped response in these variables. HOB shocks are also found to be a major contributor to the volatility of unemployment, industrial production, and capacity utilization, explaining a substantial part of their fluctuations at business cycle frequencies. This is confirmed by the fact that HOB shocks are a major contributor to the main business cycle shock as identified by Angeletos et al. (2020).

Furthermore, HOB shocks are shown to improve short-run economic forecasting by economic agents and decrease the risk premium and credit spread, fostering better financial conditions. Interestingly, HOB shocks do not seem to have a significant impact on inflation or explain a large share of the volatility in sentiment indices. We also address potential identification concerns by demonstrating that HOB shocks are not confounded with technological news or other structural/fundamental shocks like productivity, investment, discount factor, or policy shocks.

[29] This weakness is witnessed in the large width of the confidence bands for the impulse responses to our placebo shocks (reported in Figure 33 in Appendix F.)

References

Sushant Acharya, Jess Benhabib, and Zhen Huo (2021): "The anatomy of sentiment-driven fluctuations," *Journal of Economic Theory*, 195(C), S0022053121000971.

George A. Akerlof and Robert J. Shiller (2009): *Animal Spirits: How Human Psychology Drives the Economy, and Why It Matters for Global Capitalism.* Princeton University Press.

George-Marios Angeletos, Fabrice Collard, and Harris Dellas (2018): "Quantifying Confidence," *Econometrica*, 86(5), 1689–1726.

——— (2020): "Business-Cycle Anatomy," *American Economic Review*, 110(10), 3030–70.

George-Marios Angeletos (2018): "Frictional Coordination," *Journal of the European Economic Association*, 16(3), 563–603.

George-Marios Angeletos and Jennifer La'O (2013): "Sentiments," *Econometrica*, 81(2), 739–779.

Scott R Baker, Nicholas Bloom, and Stephen J Terry (2024): "Using Disasters to Estimate the Impact of Uncertainty," *The Review of Economic Studies*, 91(2), 720–747.

Robert Barsky and Eric Sims (2012): "Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence," *American Economic Review*, 102(4), 1343–77.

Paul Beaudry, Patrick Feve, Alain Guay, and Franck Portier (2019): "When is Nonfundamentalness in SVARs a Real Problem?," *Review of Economic Dynamics*, 34, 221–243.

Paul Beaudry and Franck Portier (2006): "Stock Prices, News, and Economic Fluctuations," *American Economic Review*, 96(4), 1293–1307.

——— (2014): "News-Driven Business Cycles: Insights and Challenges," *Journal of Economic Literature*, 52(4), 993–1074.

Paul Beaudry and Makoto Saito (1998): "Estimating the effects of monetary shocks: An evaluation of different approaches," *Journal of Monetary Economics*, 42(2), 241–260.

Jess Benhabib and Roger Farmer (1994): "Indeterminacy and Increasing Returns," *Journal of Economic Theory*, 63(1), 19–41.

Jess Benhabib and Roger E.A. Farmer (1999): "Indeterminacy and sunspots in macroeconomics," in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and M. Woodford, vol. 1 of *Handbook of Macroeconomics*, chap. 6, pp. 387–448. Elsevier.

Jess Benhabib, Pengfei Wang, and Yi Wen (2015): "Sentiments and Aggregate Demand Fluctuations," *Econometrica*, 83(2), 549–585.

David Cass and Karl Shell (1983): "Do Sunspots Matter?," *Journal of Political Economy*, 91(2), 193–227.

Lawrence J. Christiano and Terry J. Fitzgerald (2003): "The Band Pass Filter," *International Economic Review*, 44(2), 435–465.

Peter Diamond (1982): "Aggregate Demand Management in Search Equilibrium," *Journal of Political Economy*, 90(5), 881–94.

Jon Faust (1998): "The Robustness of Identified VAR Conclusions about Money," *Carnegie-Rochester Conference Series on Public Policy*, 49, 207–244.

John G. Fernald (2012): "A Quarterly, Utilization-adjusted Series on Total Factor Productivity," Working Paper Series 2012-19, Federal Reserve Bank of San Francisco.

Mario Forni and Luca Gambetti (2014): "Sufficient information in structural VARs," *Journal of Monetary Economics*, 66(C), 124–136.

Mario Forni, Luca Gambetti, Marco Lippi, and Luca Sala (2017): "Noisy News in Business Cycles," *American Economic Journal: Macroeconomics*, 9(4), 122–152.

Patrick Fève and Alain Guay (2019): "Sentiments in SVARs," *The Economic Journal*, 129(618), 877–896.

Mark Gertler and Peter Karadi (2015): "Monetary Policy Surprises, Credit Costs, and Economic Activity," *American Economic Journal: Macroeconomics*, 7(1), 44–76.

Zhen Huo and Naoki Takayama (2023): "Rational Expectations Models with Higher-Order Beliefs," miméo, Yale University.

Kyle Jurado, Sydney C. Ludvigson, and Serena Ng (2015): "Measuring Uncertainty," *American Economic Review*, 105(3), 1177–1216.

Lutz Kilian and Helmut Lütkepohl (2018): *Structural Vector Autoregressive Analysis*, no. 9781107196575 in Cambridge Books. Cambridge University Press.

Andresa Lagerborg, Evi Pappa, and Morten Ravn (2023): "Sentimental Business Cycles," *Review of Economic Studies*, 90(3), 1358–1393.

Michael W. McCracken and Serena Ng (2021): "FRED-QD: A Quarterly Database for Macroeconomic Research," *Review*, 103(1), 1–44.

Karel Mertens and Morten Ravn (2013): "The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States," *American Economic Review*, 103(4), 1212–47.

Benjamin Moll (2014): "Productivity Losses from Financial Frictions: Can Self-Financing Undo Capital Misallocation?," *American Economic Review*, 104(10), 3186–3221.

José Luis Montiel Olea and Carolin Pflueger (2013): "A Robust Test for Weak Instruments," *Journal of Business and Economic Statistics*, 31(3), 358–369.

Mikkel Plagborg-Møller and Christian K. Wolf (2022): "Instrumental Variable Identification of Dynamic Variance Decompositions," *Journal of Political Economy*, 130(8), 2164–2202.

Mikkel Plagborg-Møller and Christian K. Wolf (2021): "Local Projections and VARs Estimate the Same Impulse Responses," *Econometrica*, 89(2), 955–980.

Valerie Ramey and Sarah Zubairy (2018): "Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data," *Journal of Political Economy*, 126(2), 850 – 901.

Christopher Sims (1980): "Macroeconomics and Reality," *Econometrica*, 48(1), 1–48.

Frank Smets and Rafael Wouters (2007): "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, 97(3), 586–606.

James H. Stock and Mark W. Watson (2012): "Disentangling the Channels of the 2007-09 Recession," *Brookings Papers on Economic Activity*, 43(1 (Spring)), 81–156.

——— (2018): "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments," *Economic Journal*, 128(610), 917–948.

Harald Uhlig (2004): "What moves GNP?," Econometric Society 2004 North American Winter Meetings 636, Econometric Society.

— Appendix —

A Correlation Analysis

Figure 26 reports the correlation between each of the variable in our benchmark VAR and the variables used to recover the structural shocks in Angeletos et al.' (2018) model (*e.g.* GDP, Investment, Consumption, hours worked, inflation and the fed fund rate). Figure 26 reveals a

Figure 26: Correlation Analysis

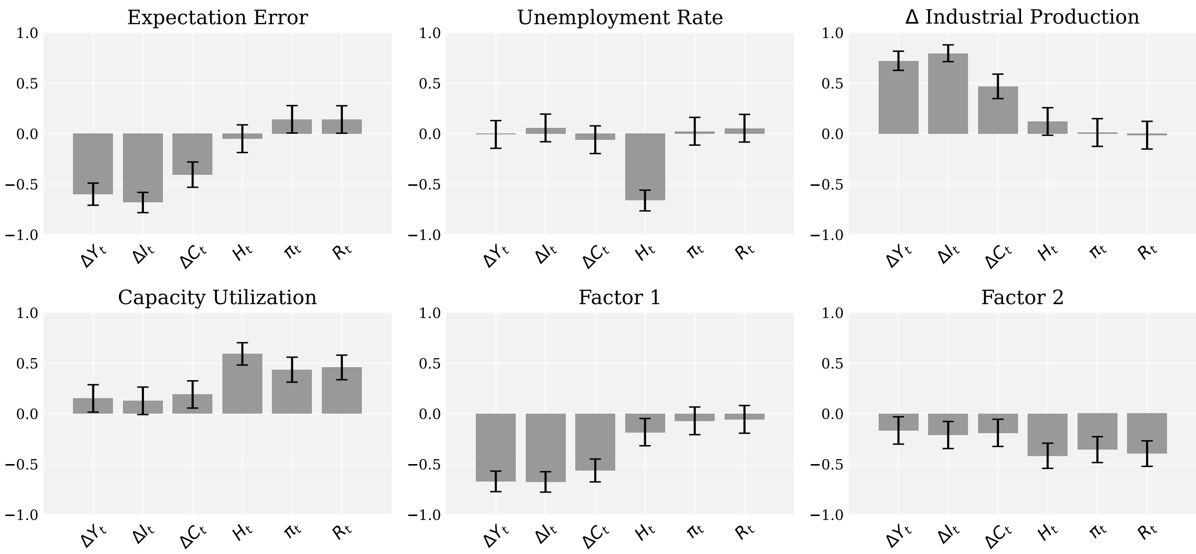

Note: Δ denotes the fist difference log operator, Y_t, C_t, I_t, H_t, π_t and R_t denote, respectively, GDP, consumption (non-durables+services), investment (Gross private investment + durable consumption), hours worked, CPI inflation and the Federal Fund Rate. Cap lines: 95% confidence bands.

non-zero correlation between several variable pairs in our study and those employed in Angeletos et al. (2018). While the strongest correlation, at 0.8, is observed between industrial production and investment, it remains far from perfect. In fact, most absolute correlation values fall below 0.5. This suggests a weak to moderate relationship between the variables used in Angeletos et al. (2018) and in our VAR, despite both studying macroeconomic aggregates. This implies that our VAR model incorporates significant information not exploited in the Angeletos et al. (2018) model's recovery of structural shocks.

B HOB vs Forecast Error on Output

Figure 27: HOB instrument vs Forecast Error on Output

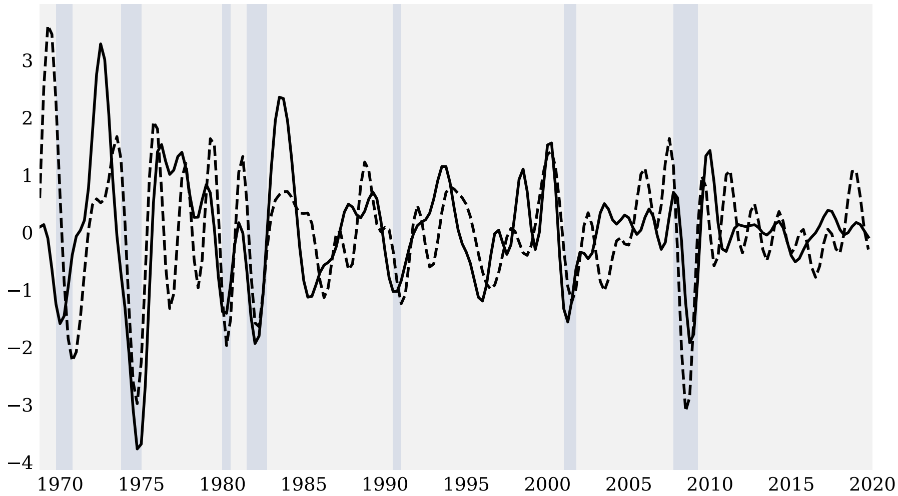

Forecast Error on Output, Average HOB instrument.

The HOB instrument corresponds to the average of the HOB instrument across models. Both series is filtered by applying the Christiano and Fitzgerald's 2003 filter (6-32 quarters frequencies).

C A Version of Angeletos et al. (2020)

Table 10: FEVD: ACD VAR

	Impact	1-Quarter	1-Year	2-Years	5-Years	Long-Run
Expectation Error	20.46	29.24	30.04	32.12	34.36	34.46
Unemployment Rate	20.86	30.92	39.80	36.62	28.42	39.88
Hours Worked	25.02	38.16	49.66	49.19	38.74	56.77
Output	21.81	32.81	39.26	34.35	19.39	56.56
Investment	7.75	17.01	24.11	19.39	14.54	25.04
Consumption	63.28	65.89	67.97	65.33	55.79	63.55
CPI Inflation	0.47	2.11	11.42	19.66	26.84	30.80
Fed. Fund Rate	33.84	42.51	55.61	61.31	64.20	62.45

Figure 28: ACD VAR

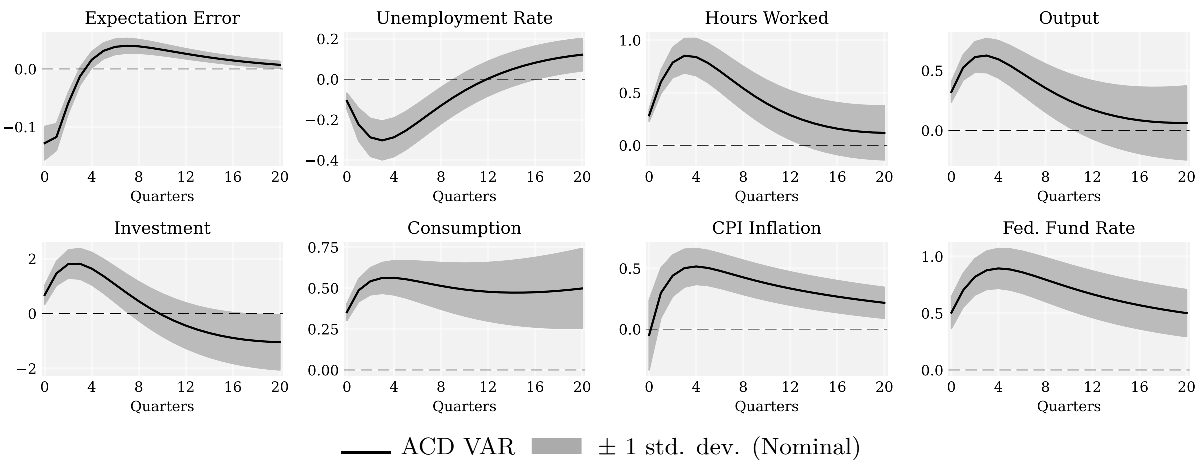

D The Main Business Cycle Shock

D.1 Methodology

This section gives technical details regading the construction of the MBC shock as identified in Angeletos et al. (2020). For exposition purposes, let us repeat the VAR reported in Equation (1)

$$A(L)Y_t = u_t,$$

where the definitions of $A(L)$, Y_t and u_t are the same and satisfy the same properties as in Section 1.1. The innovations, u_t, are assumed to be linear combinations of n_y mutually orthogonal shocks, η_t, such that

$$u_t = R\eta_t \tag{10}$$

where R is an invertible $n_y \times n_y$ matrix and η_t is i.i.d. over time, with $E(\eta_t \eta_t') = I$. We let $R = SQ$, where S is the Cholesky decomposition of the covariance matrix, Σ, of the VAR residuals, and Q is a rotation matrix (*i.e.* $QQ' = I$). As R is invertible, it follows that $\eta_t = R^{-1}u_t = Q'S^{-1}u_t$. This implies that each orthogonal shock η_t can be expressed as a linear combination of the Choleski shocks $S^{-1}u_t$. The coefficients of this linear combination are given by the corresponding column of matrix Q. The main business cycle shock is then obtained as the linear combination of the Choleski shocks —the column vector, q, of Q—that maximizes the variance of a particular variable at business cycle frequencies. The identification proceeds in three steps.

First, Angeletos et al. (2020) obtain the Wold representation of the VAR as

$$Y_t = B(L)u_t$$

where $B(L) = A(L)^{-1}$ is an infinite matrix polynomial, or $B(L) = \sum_{\tau=0}^{\infty} B_\tau L^\tau$. Replacing $u_t = SQ\eta_t$, this rewrites as

$$X_t = C(L)Q\eta_t = D(L)\eta_t,$$

The infinite matrix polynomials $C(L)$ and $D(L)$ are defined as $C(L) = \sum_{j=0}^{\infty} C_j L^j$ and $D(L) = \sum_{j=0}^{\infty} D_j L^j$, where $C_j \equiv B_j S$ and $D_j \equiv C_j Q$ for all $j \geqslant 0$. The sequence $\{D_j\}_{j=0}^{\infty}$ corresponds to the IRFs of the VAR variables to the shock of interest. These IRFs can be obtained as a linear combination of the IRFs of the variables to each of the Choleski shock, which are represented by the sequence $\{C_j\}_{j=0}^{\infty}$. The weights of the linear combination are then simply given by the column vector q. Henceforth, we will denote by D_j^k (respectively C_j^k) the j-horizon response of the k-th variable in vector Y.

Second, they build the contribution of the shock of interest to the volatility of the k-th variable of vector Y over a given band of frequency $[\underline{\omega};\overline{\omega}]$

$$V(q;k,\underline{\omega},\overline{\omega}) = q'\,\Omega(k,\underline{\omega},\overline{\omega})\,q$$

where[30]

$$\Omega(k,\underline{\omega},\overline{\omega}) \equiv \int_{\underline{\omega}}^{\overline{\omega}} \overline{C^k(e^{-i\omega})} C^k(e^{-i\omega})\mathrm{d}\omega$$

[30] For any vector $v \in \mathbb{C}$, $\overline{v}$ denotes the complex conjugate of v.

represents the volatility of the k-th variable in Y_t over the aforementioned frequency band, expressed in terms of the contributions of all the Cholesky shocks.

Finally, the main business cycle shock is identified by maximizing its contribution to the volatility of a particular variable over the business cycle frequency band $(2\pi/32, 2\pi/6)$ —*i.e.* $q \in \text{Argmax } V(q; k, 2\pi/32, 2\pi/6)$. It follows that q is the eigenvector associated to the largest eigenvalue of the matrix $\Omega(k, 2\pi/32, 2\pi/6)$. This approach is similar to the "max-share" method developed in Faust (1998) and Uhlig (2004), but is formulated in the frequency domain.

D.2 Difference Specification

Figure 29: Extended VAR on MBC component: IRFs (Difference Specification)

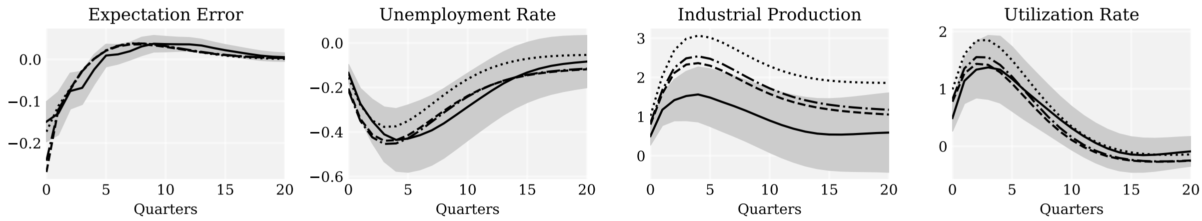

HOB MBC: Expectation Error MBC: Unemployment MBC: Industrial Production ± 1 std. dev. (HOB)

Table 11: FEVD: Baseline VAR (MBC Shocks, Difference Specification)

	Impact	1-Quarter	1-Year	2-Years	5-Years	Long-Run
Expectation Error						
Exp. Error ($\zeta^u_{t\|t-1}$)	89.98	90.45	89.24	88.58	88.22	88.08
Unemployment Rate	76.83	86.68	91.09	89.67	81.87	76.39
Industrial Production	56.31	56.31	58.59	56.24	46.37	16.32
Capacity Utilization	55.39	58.84	57.38	53.32	52.15	52.83
Unemployment Rate						
Exp. Error ($\zeta^u_{t\|t-1}$)	72.87	76.18	76.43	76.84	77.44	77.34
Unemployment Rate	83.30	92.02	96.99	95.47	86.18	79.53
Industrial Production	62.57	62.57	66.55	64.66	53.94	19.05
Capacity Utilization	61.68	66.86	67.24	63.50	61.29	61.02
Industrial Production						
Exp. Error ($\zeta^u_{t\|t-1}$)	37.35	44.29	46.57	48.67	50.19	50.08
Unemployment Rate	44.43	55.87	64.76	62.31	52.32	45.59
Industrial Production	92.50	92.50	96.70	96.67	90.22	50.75
Capacity Utilization	90.27	94.75	96.08	93.70	87.25	81.61

Figure 30: MBC vs HOB: Conditional History (Difference Specification)

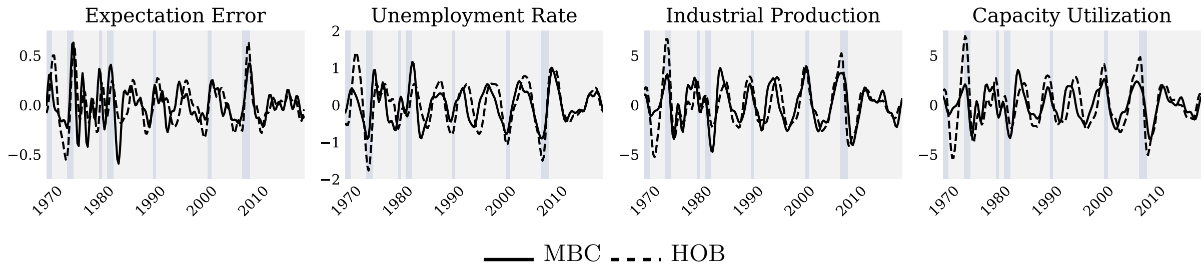

Table 12: HOB as MBC

Sample	$\zeta^u_{t\|t-1}$	u_t	IP_t	Util_t
Whole Sample	0.57	0.68	0.75	0.70
	(0.33)	(0.47)	(0.56)	(0.48)
Pre-Volcker	0.60	0.56	0.73	0.63
	(0.35)	(0.30)	(0.52)	(0.38)
Post-Volcker	0.62	0.75	0.80	0.77
	(0.39)	(0.56)	(0.63)	(0.58)
1980Q1-2019Q4	0.58	0.74	0.80	0.75
	(0.34)	(0.55)	(0.63)	(0.57)
1990Q1-2019Q4	0.69	0.80	0.82	0.79
	(0.47)	(0.64)	(0.67)	(0.63)
2000Q1-2019Q4	0.84	0.94	0.87	0.90
	(0.70)	(0.88)	(0.75)	(0.80)

Note: MBC shock obtained by targeting the unemployment rate. R^2 of the projection of the history conditional on the MBC shock onto the history conditional on the HOB shock.

E Alternative Shocks from a Complete Information Model

This section investigates whether alternative shocks, derived from a standard New Keynesian model (essentially Angeletos et al. (2018) without information frictions), can replicate the business cycle dynamics we uncovered using HOB shocks. We leverage the same approach employed to recover the belief shock in the main text. We extract 200 parameter sets from the posterior distribution estimated by Angeletos et al. (2018) for their perfect-information model (i.e., no role for HOB shocks). For each parameter set, we solve the model and utilize the Kalman techniques to extract structural shocks from observed quarterly data (1968Q4-2019Q4) on GDP, consumption, investment, hours worked, inflation, and the federal funds rate. This process yields 200 sets of instrumental time series, each corresponding to a specific model configuration, encompassing permanent and transitory technology and investment shocks, a discount factor shock, a government spending shock and a monetary policy shock.

We begin by winnowing down the candidate instruments for our expectation error within the benchmark VAR. Figure 31 depicts the distribution of the weak instrument test across all

Figure 31: Weak Instrument Test (No HOB shock model)

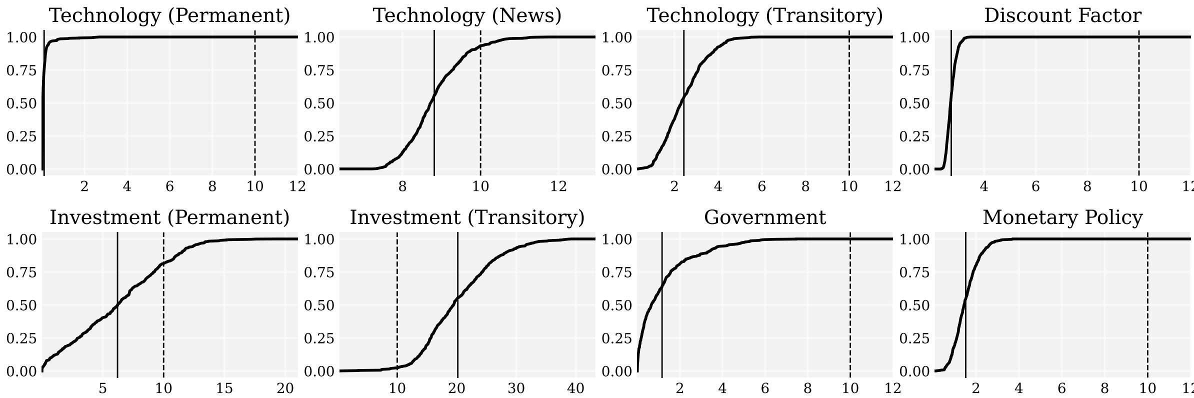

Vertical plain line: Average value of test, Vertical dashed line: rule of thumb threshold (10)

parameter draws. Except for the transitory investment shock, the test statistics decisively fall below the conventional threshold of 10, indicating that the remaining shocks are unsuitable instruments. Interestingly, this mirrors findings from the HOB shock model (Section 2.3.5). However, unlike the HOB case where 30% of models passed the test, less than 5% do in the perfect-information model. This stark reduction suggests the transitory investment shock likely captures the variation previously explained by HOB shocks.

To verify the latter assumption, Figure 32 presents impulse response functions along with 1 standard deviation bands for key variables in our VAR. These responses are generated using the transitory investment shock instrument extracted from both the complete and incomplete information models. Table 13 complements this analysis by reporting the corresponding forecast error variance decompositions. Two key findings emerge. First, the transitory investment shock instrument derived from the incomplete information model yields less precise estimates for the impulse responses. Notably, the effect of the shock on unemployment becomes statistically

Figure 32: Impulse Response to a Transitory Investment Shock

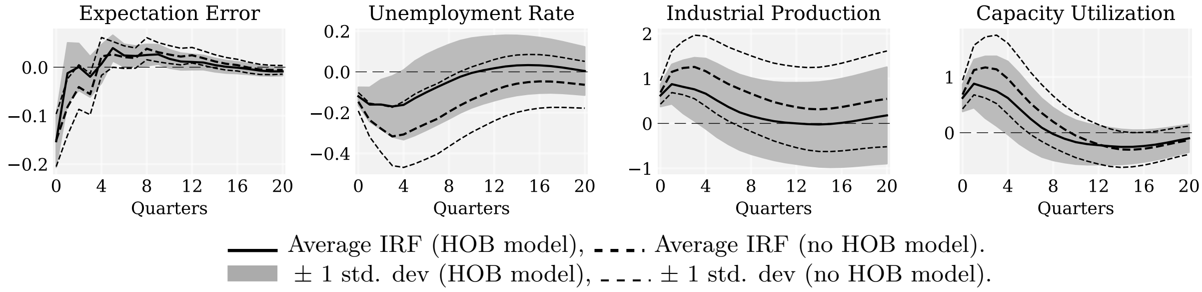

insignificant (at the 68% confidence level) after just four quarters, compared to two years with the complete information instrument. Second, the complete information model's transitory investment shock response closely resembles the HOB shock response in our VAR. Conversely, the incomplete information version exhibits significantly weaker responses, confirmed by FEVD analysis. For instance, the complete information shock explains 65% of unemployment volatility on impact (44% after one year) compared to a 20 percentage point (50% after one year) reduction in the incomplete version. Taken together, these findings suggest that HOB shock information is captured by the complete information model's transitory investment shock.

Table 13: FEVD: Baseline VAR (Transitory Investment Shock, w/ and w/o HOB)

	Impact	1-Quarter	1-Year	2-Years	5-Years	Long-Run
HOB Model						
Exp. Error	57.74	33.99	24.96	25.76	24.63	24.47
Unemployment Rate	44.16	34.39	20.28	14.46	9.76	7.10
Industrial Production	53.99	42.20	23.76	15.21	8.06	1.17
Capacity Utilization	54.57	43.31	25.97	19.07	17.95	6.41
No HOB Model						
Exp. Error	56.03	39.62	31.51	31.36	30.13	29.81
Unemployment Rate	65.37	61.53	44.02	32.73	22.06	18.39
Industrial Production	65.42	61.32	41.66	28.61	16.89	3.84
Capacity Utilization	64.59	60.18	40.01	28.37	25.68	13.17

Two intriguing questions remain. First, why isn't this information transmitted to news shocks, which are shocks to expectations? Unlike HOB shocks that target short-run outcomes, news shocks focus on longer-term productivity gains. Second, why is the information captured by the transitory investment shock? The investment decision trades-off the cost of investing in capital accumulation against the expected future benefits of this investment, which are precisely affected by the transitory investment shock as of the short-run. Therefore, it's natural that HOB shock variations are reflected in these investment shocks.

F Placebo

Figure 33: Placebo test

(a) Placebo I

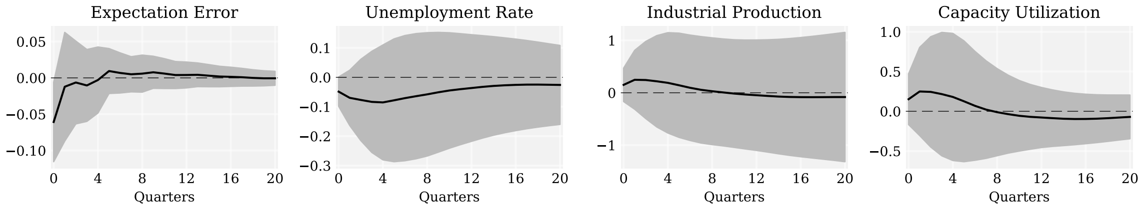

(b) Placebo II

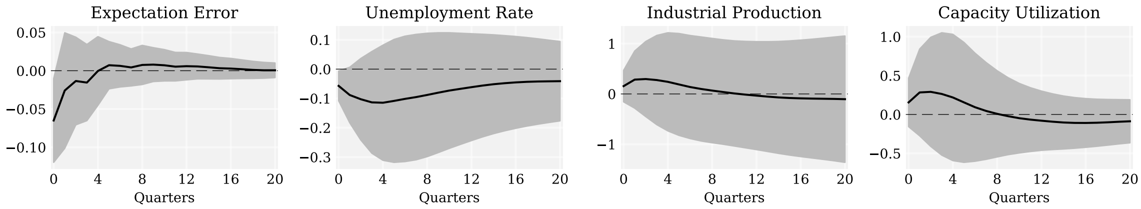

Placebo ± 1 std. dev.

Note: Placebo I: Random permutation of all elements of the instrument, Placebo II: Permutation of 2 blocs of the instrument, cutoff randomly selected. The VAR includes the first two McCracken's factors.