

Business Cycle Anatomy^{*}

George-Marios Angeletos
MIT

Fabrice Collard
Toulouse School of Economics
(CNRS)

Harris Dellas
University of Bern

August 7, 2019

Abstract

We propose a new strategy for dissecting the macroeconomic time series, provide a template for the propagation mechanism that best describes the observed business cycles, and use its properties to appraise models of both the parsimonious and the medium-scale variety. Our findings support the existence of a main business-cycle driver but rule out the following candidates for this role: technology or other shocks that map to TFP movements; news about future productivity; and inflationary demand shocks of the textbook type. Prominent members of the DSGE literature also lack the propagation mechanism seen in our anatomy of the data. Models that aim at accommodating demand-driven cycles under flexible prices appear promising.

^{*}We thank the editor, Mikhail Golosov, and three anonymous referees for extensive feedback. For useful comments, we also thank Larry Christiano, Patrick Feve, Francesco Furlanetto, Jordi Gali, Lars Hansen, Franck Portier, Juan Rubio-Ramirez and participants at various seminars and conferences. Angeletos acknowledges the financial support of the National Science Foundation (Award #1757198). Collard acknowledges funding from the French National Research Agency (ANR) under the Investments for the Future program (Investissements d'Avenir, grant ANR-17-EURE-0010).

“One is led by the facts to conclude that, with respect to the qualitative behavior of co-movements among series, business cycles are all alike. To theoretically inclined economists, this conclusion should be attractive and challenging, for it suggests the possibility of a unified explanation of business cycles.” Lucas (1977)

1 Introduction

In their quest to explain macroeconomic fluctuations, macroeconomists have often relied on models in which a single, recurrent shock acts as the main, or even the sole, driver of the business cycle.¹ This practice is grounded not only on the desire to offer a parsimonious, unifying explanation as suggested by Lucas, but also on the property that such a model may capture diverse business-cycle triggers if these share a common propagation mechanism: multiple shocks that produce the same impulse responses for *all* variables of interest can be considered as the same shock.²

Is there evidence of such a common propagation mechanism in macroeconomic data? And if yes, how does it look like?

We address these questions with the help of a new empirical strategy. The strategy involves taking multiple cuts of the data. Each cut corresponds to a VAR-based shock that accounts for the maximal volatility of a particular variable over a particular frequency band. Whether these empirical objects have a direct structural counterpart or not, their properties form a rich set of cross-variable, static and dynamic restrictions, which can inform macroeconomic theory. We call this set the “anatomy.”

A core subset of the anatomy is the collection of the five shocks obtained by targeting the main macroeconomic quantities (unemployment, output, hours worked, consumption and investment) over the business-cycle frequencies. These shocks turn out to be interchangeable in the sense of giving rise to nearly the same impulse response functions (IRFs) for all the variables, as well as being highly correlated with one another.

The interchangeability of these shocks supports the hypothesis of a main, unifying, propagation mechanism. Their shared impulse response functions provide an empirical template of it.

In combination with other elements of our anatomy, this template rules out the following candidates for the *main* driver of the business cycle: technology or other shocks that map to TFP movements; news about future productivity; and inflationary demand shocks. Prominent members of the DSGE literature also lack the propagation mechanism seen in the data. In contrast, models that allow for demand-driven cycles even in the absence of nominal rigidity or, equivalently, even when monetary policy replicates flexible-price allocations, seem to fit the provided template better.³

¹E.g., this is the monetary shock in Lucas (1975), the TFP shock in Kydland and Prescott (1982), the sunspot in Benhabib and Farmer (1994), the investment shock in Justiniano, Primiceri, and Tambalotti (2010), the risk shock in Christiano, Motto, and Rostagno (2014), and the confidence shock in Angeletos, Collard, and Dellas (2018).

²To echo Cochrane (1994): “The study of shocks and propagation mechanisms are of course not separate enterprises. Shocks are only visible if we specify something about how they propagate to observable variables.”

³Recent examples of such models include Angeletos and La’O (2010, 2013), Bai, Ríos-Rull, and Storesletten (2017), Beaudry and Portier (2018), Beaudry, Galizia, and Portier (2018), Benhabib, Wang, and Wen (2015), Eusepi and Preston (2015) Jaimovich and Rebelo (2009), Huo and Takayama (2015), and Ilut and Saijo (2018). Related is also the earlier literature on coordination failures (Diamond, 1982; Benhabib and Farmer, 1994; Guesnerie and Woodford, 1993).

The empirical strategy. We first estimate a VAR (or a VECM) on the following ten macroeconomic variables over the 1955-2017 period: the unemployment rate; the per-capita level of GDP, investment (inclusive of consumer durables), consumption (of non-durables and services), and total hours worked; labor productivity in the non-farm business sector; utilization-adjusted TFP; the labor share; the inflation rate (GDP deflator), and the federal funds rate. We next compile a collection of reduced-form shocks, each of which is identified by maximizing its contribution to the volatility of a particular variable over either business-cycle frequencies (6-32 quarters) or long-run frequencies (80- ∞). We finally inspect the empirical patterns encapsulated in each of these shocks, namely the implied IRFs and variance contributions.

This approach departs from standard practice in the SVAR literature, which aims at identifying empirical counterparts to specific theoretical mechanisms (for a review, see Ramey, 2016). Instead, it sheds light on dynamic comovements by taking multiple cuts of the data, one per targeted variable and frequency band. For example, one cut is obtained by targeting unemployment over the business-cycle frequencies, another by targeting TFP over the long-run frequencies, and so on. These cuts, which may or may not have a direct structural interpretation, comprise our “anatomy” of the data and form a rich set of empirical restrictions that can discipline theory.⁴

The Main Business Cycle Shock. Consider the shocks that target any of the following variables over the business-cycle frequencies: unemployment, hours worked, GDP, and investment. These shocks are interchangeable in terms of the dynamic comovements (IRFs) they produce. Furthermore, any one of them accounts for about three-quarters of the business-cycle volatility of the targeted variable and for more than one half of the business-cycle volatility in the remaining variables, and triggers strong positive co-movement in all variables. The shock that targets consumption is less tightly connected in terms of variance contributions, but still similar in terms of comovements/IRFs.

These findings offer support for theories featuring either a single, dominant, business-cycle shock, or multiple shocks that leave the same footprint because they share the same propagation mechanism. With this idea in mind, we use the term “Main Business Cycle shock” (henceforth, MBC shock) to refer to the common empirical footprint, in terms of IRFs, of the aforementioned reduced-forms shocks. This provides the sought-after template.⁵

A central feature of this template is the interchangeability property, namely all the aforementioned reduced-form shocks produce essentially the same IRFs, or the same propagation mechanism. Below, we describe a few additional features of the MBC shock and of the overall anatomy, and discuss their lessons for theory. At first, we draw lessons through the perspective of single-shock models. Later, we switch to multi-shock models and discuss the challenges and the use of our method in such models.

Disconnect from TFP and from the long run. The MBC shock is disconnected from TFP at *any* frequency. It also accounts for little of the long-term variation in output, investment, consumption,

⁴The basic idea of identifying a shock by maximizing its variance contribution to a variable is borrowed from Faust (1998) and Uhlig (2003); see also Barsky and Sims (2011) and Francis et al. (2014). What distinguishes our contribution is the multitude of such shocks considered, the empirical regularities recovered, and the lessons drawn for theory. Also, an early version of our method and results appeared in Section 2 of Angeletos, Collard, and Dellas (2015); the present paper subsumes this earlier work.

⁵Additional support for the existence of a main business-cycle driver is provided by recovering the first principle component of the business-cycle frequencies of the data. However, principal component analysis (PCA) does not allow for the construction of IRFs and therefore does not provide the template sought after.

and labor productivity. Symmetrically, the shocks identified by maximizing the long-term volatility in any of these variables make a negligible contribution to the business cycle.

These findings are inconsistent not only with the baseline RBC model but also with models that map other shocks, including financial, uncertainty and sunspot shocks into endogenous TFP fluctuations. In these models, the productivity movements over the business-cycle frequencies ought to be tightly tied to the MBC shock, which is not the case.⁶

These findings also challenge Beaudry and Portier (2006), who emphasize news of productivity and income in the future. If such news were the main driver of the business cycle, the MBC shock would be a sufficient statistic of the available information about future TFP movements, which is not the case. Instead, a semi-structural exercise based on our anatomy suggests that the contribution of TFP news to unemployment fluctuations is in the order of 10%, which is broadly consistent with the estimate provided by Barsky and Sims (2011).

The MBC shock fits better the notion of an aggregate demand shock unrelated to productivity and the long run, in line with Blanchard and Quah (1989) and Galí (1999). However, as discussed below, this shock ought to be non-inflationary, which may or may not fit the Keynesian framework.

Disconnect from inflation. The shock that targets unemployment accounts for 74% of the business-cycle variation in unemployment and only for 7% of the variation in inflation. And conversely, the shock that targets inflation explains 83% of the variation in inflation and only 4% of the variation in unemployment. Moreover, the magnitude of the inflation response to the MBC shock is close to zero.

These properties preclude the interpretation of the MBC shock as an *inflationary* demand shock, of the type contained in the textbook, New Keynesian model. Could it be that the MBC shock represents a mixture of an inflationary demand shock and a disinflationary supply shock? While this possibility cannot be ruled out in general, it is invalid insofar as we proxy the supply shock by the TFP shock in the data. It is also incompatible with estimated, New Keynesian, DSGE models.⁷ Such models attribute the bulk of the business cycle to demand shocks and, at the same time, make sure that demand shocks are nearly non-inflationary by assuming a large degree of nominal rigidity.

Another possibility is that the MBC shock represents a demand shock whose importance does not vanish when prices are flexible, or when monetary policy replicates flexible-price outcomes. This possibility, which is accommodated by the models cited in footnote 3, seems consistent with the fact that the MBC shock induces a strong countercyclical response in monetary policy without a sizable movement in inflation.

The anatomy of medium-scale DSGE models. Our empirical strategy was motivated by parsimonious models. Does it retain its probing power in medium-scale DSGE models?

Such models pose a challenge for the interpretation and use of the MBC shock identified in the data, as this may correspond to a combination of theoretical shocks, none of which individually has its properties.⁸ But at the same time, such models give rise to a larger set of cross-variable, static and

⁶Benhabib and Farmer (1994), Bloom et al. (2018) and Bai, Ríos-Rull, and Storesletten (2017) are notable examples of such models: the first generates procyclical TFP movements out of sunspots, the second out of uncertainty shocks, and the third out of demand shocks.

⁷Smets and Wouters (2007), Justiniano, Primiceri, and Tambalotti (2010), and Christiano, Motto, and Rostagno (2014).

⁸This difficulty is not specific to our approach. It concerns any approach that requires a single shock to drive some conditional variance in the data. For instance, Galí (1999) requires that a single shock drives productivity in the long run,

dynamic restrictions that can be confronted with our anatomy of the data.

We demonstrate these ideas in Section 6 using state-of-the-art, DSGE models. One is the sticky-price model of Justiniano, Primiceri, and Tambalotti (2010); this is essentially the same as that developed in Christiano, Eichenbaum, and Evans (2005) and Smets and Wouters (2007). Another one is the flexible-price model found in Angeletos, Collard, and Dellas (2018); this is an extension of the RBC model that disentangles the expectations of the short-term economic outlook from fundamentals such as TFP and allows business cycles to be driven by variation in the level of confidence about the behavior of others. We view the former as representative of the New Keynesian paradigm and the latter as an example of the literature cited in footnote 3, which aims at accommodating demand-driven business cycles outside the realm of nominal rigidities and Philips curves.

In each model, we perform an anatomy similar to that carried out in the data: we consider different linear combinations of the theoretical shocks, each one constructed by maximizing the business-cycle volatility of a different macroeconomic quantity. We then compare the model-based objects to their empirical counterparts.

Both models match the disconnect of the MBC shock from TFP and inflation. However, the first model does not match the interchangeability property of the MBC template: the reduced-form shocks obtained by targeting the key macroeconomic quantities are less similar in the model than their empirical counterparts. This is because this model—like many other members of the DSGE literature—attributes the business cycle to a fortuitous combination of specialized shocks, none of which generates the empirically relevant comovement patterns in the key macroeconomic quantities. By contrast, the second model fits the patterns seen in the data because it contains a dominant shock/mechanism that alone generates these patterns.

As an additional demonstration of the value of our method, we use it to evaluate the model of Christiano, Motto, and Rostagno (2014). This model is a leader in a new strand of the DSGE literature that includes financial frictions and uses financial (risk) shocks to drive the business cycle. We find that this model has difficulties, not only in the dimension discussed above, but also in satisfying the cross-variable, dynamic restrictions between the MBC shock and the new, financial variables it contains, in particular the credit spread and the level of credit.

To summarize: Although there is no presumption that the reduced-form objects comprising our anatomy can have a meaningful structural interpretation in the realm of arbitrary multi-shock models, they prove quite effective in the evaluation of the shock structure and propagation mechanisms in models actually used in the literature. And whether we interpret the evidence through the lenses of a parsimonious, single-shock perspective or through those of the state-of-the-art DSGE models considered, they point in the direction of theories that aim at making sense of demand-driven business cycles without the assistance of nominal rigidities and accommodative monetary policies.

Layout. The rest of the paper is organized as follows. Section 2 describes the empirical method. Section 3 reviews our empirical findings. Section 4 reports the various robustness exercises. Section 5 offers an interpretation of the main empirical findings. Section 6 contains the application to medium-scale models. Section 7 concludes.

an assumption inconsistent with the literature on news shocks.

2 Data and Method

The data used in our main specification consists of quarterly observations on the following ten, key macroeconomic variables: the unemployment rate (u); the real, per-capita levels of GDP (Y), investment (I), consumption (C); hours worked per person (h); labor productivity in the non-farm business sector (Y/h); the level of utilization-adjusted total factor productivity (TFP); the labor share ($\frac{Wh}{Y}$); the inflation rate (π), as measured by the rate of change in the GDP deflator; and the nominal interest rate (R), as measured by the federal funds rate. The sample starts in the first quarter of 1955, the earliest date of availability for the federal funds rate, and ends in the last quarter of 2017.

Following standard practice, and to ensure compatibility with the models used in Section 6, our measure of investment includes consumer expenditure on durables, while that of consumption consists of expenditure on non-durables and services. Both measures are herein deflated by the GDP deflator. Section 4 establishes the robustness of our results to the use of component-specific deflators; to different samples, such as the pre- and post-Volcker periods or excluding the Great Recession and the ZLB period; and to the incorporation of additional information, such as that contained in stock prices and financial variables. Appendix A contains the definitions and data sources.

We now turn to the description of the empirical method. As mentioned in the Introduction, the method involves running a VAR on the aforementioned ten variables and recovering certain “shocks.” As in the SVAR literature, any of the shocks constructed here represents a particular linear combination of the VAR residuals. What distinguishes our approach is the criterion used in the identification of such a linear combination.

Let the VAR take the form

$$A(L)X_t = \nu_t,$$

where the following definitions apply: X_t is a $N \times 1$ vector, containing the macroeconomic variables under consideration; $A(L) \equiv \sum_{\tau=0}^p A_\tau L^\tau$ is a matrix polynomials in the backshift operator L , with $A(0) = A_0 = I$; p is the number of lags included in the VAR; and ν_t is the vector of VAR residuals, with $E(\nu_t \nu_t') = \Sigma$ for some positive definite matrix Σ . Because of its large size, the VAR was estimated with Bayesian methods, using a Minnesota prior.⁹ Also, our baseline specification set $p = 2$, which was the number of lags suggested by standard Bayesian criteria. Section 4 shows the robustness of our main findings to the inclusion of additional lags and the use of a VECM instead of a VAR.¹⁰

We assume the existence of a linear mapping between the residuals, ν_t , and some mutually independent “structural” shocks, ε_t , that is, we let

$$\nu_t = S\varepsilon_t$$

where S is an invertible $N \times N$ matrix and ε_t is i.i.d. over time, with $\mathbb{E}(\varepsilon_t \varepsilon_t') = I$. These “structural” shocks may or may not correspond to the kind of structural shocks featured in theoretical models; they are transformations of the VAR residuals, whose interpretation is inherently delicate.

⁹The posterior distributions were then obtained using Gibbs sampling with 50,000 draws, and the reported highest posterior density intervals (HPDI) were obtained by the approach described in Koop (2003).

¹⁰A VECM may be recommended if the analyst believes, perhaps on the basis of theory, that certain variables are co-integrated. But a VECM is also sensitive to the assumed co-integration relations, which explains why we, as much of the related empirical literature, use the VAR as our baseline specification.

We can always write $S = \tilde{S}Q$, where \tilde{S} is the Cholesky decomposition of Σ , the covariance matrix of the VAR residuals, and Q is an orthonormal matrix, namely a matrix such that $Q^{-1} = Q'$. We then have that $\varepsilon_t = S^{-1}\nu_t = Q'\tilde{S}^{-1}\nu_t$, which means that each one of the shocks in ε_t corresponds to a column of the matrix Q . Furthermore, Q satisfies $QQ' = I$ by construction, which is equivalent to S satisfying $SS' = \Sigma$. But this by itself does not suffice to identify any of the underlying shocks: additional restrictions must be imposed on Q in order to identify any of them. The typical SVAR exercise in the literature employs exclusion or sign restrictions motivated by specific theories. We instead identify a shock by the requirement that it contains the maximal share of all the information in the data about the volatility of a particular variable in a particular frequency band.

Let us fill in the details. The Wold representation of the VAR is given by

$$X_t = B(L)\nu_t$$

where $B(L) = A(L)^{-1}$ is an infinite matrix polynomial of the form $B(L) = \sum_{\tau=0}^{\infty} B_{\tau}L^{\tau}$. Replacing $\nu_t = \tilde{S}Q\varepsilon_t$, we can rewrite the above as follows:

$$X_t = C(L)Q\varepsilon_t = \Gamma(L)\varepsilon_t,$$

where $C(L)$ and $\Gamma(L)$ are infinite matrix polynomials of the form $C(L) = \sum_{\tau=0}^{\infty} C_{\tau}L^{\tau}$ and $\Gamma(L) = \sum_{\tau=0}^{\infty} \Gamma_{\tau}L^{\tau}$, with $C_{\tau} \equiv B_{\tau}\tilde{S}$ and $\Gamma_{\tau} \equiv C_{\tau}Q$ for all $\tau \in \{0, 1, 2, \dots\}$. The sequence $\{\Gamma_{\tau}\}_{\tau=0}^{\infty}$ represents the IRFs of the variables to the structural shocks. This is obtained from the sequence $\{C_{\tau}\}_{\tau=0}^{\infty}$, which encapsulates the Cholesky transformation of the VAR residuals.

For any pair $(k, j) \in \{1, \dots, N\}^2$, take the k -th variable in X_t and the j -th shock in ε_t . As already noted, this shock corresponds to the j -th column of the matrix Q . Let this column be the vector q . For any $\tau \in \{0, 1, \dots\}$, the effect of this shock on the aforementioned variable at horizon τ is given by the (k, j) element of the matrix $\Gamma_{\tau} \equiv C_{\tau}Q$, or equivalently by the number $C_{\tau}^{[k]}q$, where $C_{\tau}^{[k]}$ henceforth denotes the k -th row of the matrix C_{τ} . Similarly, the contribution of this shock to the spectral density of this variable over the frequency band $[\underline{\omega}, \bar{\omega}]$ is given by

$$\Upsilon(q; k, \underline{\omega}, \bar{\omega}) \equiv \int_{\omega \in [\underline{\omega}, \bar{\omega}]} \left(\overline{C^{[k]}(e^{-i\omega})q} C^{[k]}(e^{-i\omega})q \right) d\omega = q' \left(\int_{\omega \in [\underline{\omega}, \bar{\omega}]} \overline{C^{[k]}(e^{-i\omega})} C^{[k]}(e^{-i\omega}) d\omega \right) q$$

where, for any vector v , \bar{v} denotes its complex conjugate transpose.

Consider the matrix

$$\Theta(k, \underline{\omega}, \bar{\omega}) \equiv \int_{\omega \in [\underline{\omega}, \bar{\omega}]} \overline{C^{[k]}(e^{-i\omega})} C^{[k]}(e^{-i\omega}) d\omega$$

This matrix captures the entire volatility of variable k over the aforementioned frequency band, expressed in terms of the contributions of all the Cholesky-transformed residuals. It can be obtained directly from the data (i.e., from the estimated VAR), without any assumption about Q . The contribution of any structural shock can then be re-written as

$$\Upsilon(q; k, \underline{\omega}, \bar{\omega}) = q'\Theta(k, \underline{\omega}, \bar{\omega})q, \tag{1}$$

where, as already explained, q is the column vector of Q corresponding that shock.

The above is true for any shock, no matter how it is identified. Our approach is to identify a shock by maximizing its contribution to the volatility of a particular variable over a particular frequency band, that is, to choose q so as to maximize the number given in (1). It follows that q is the eigenvector associated to the largest eigenvalue of the matrix $\Theta(k, \underline{\omega}, \bar{\omega})$.

This identification strategy reminds principle component analysis; we expand on this point in Section 3.3. It is also similar to the “max-share” method employed in Faust (1998), Uhlig (2003), Barsky and Sims (2011) and Francis et al. (2014), except for two differences. The first, and mostly technical, difference is that we work on the frequency domain rather than the time domain.¹¹ The second, and more substantial, difference is that we systematically vary the targeted variable and/or the targeted frequency band instead of committing to a specific such choice.

In the next section, we start by targeting unemployment and setting $[\underline{\omega}, \bar{\omega}] = [2\pi/32, 2\pi/6]$, which is the frequency band typically associated with the business cycle (e.g., Stock and Watson, 1999). We then proceed to vary both the targeted variable and the targeted frequency band. This produces many different cuts of the data, the collection of which comprises the “anatomy” offered in this paper and forms the basis of the lessons we draw for theory.

3 Empirical findings

This section presents the main findings of our empirical method and discusses a few tentative lessons for macroeconomic theory. These lessons are sharpest under our preferred perspective, namely, when seeking to understand the business cycle as the product of a single, dominant shock/mechanism. This is the perspective adopted in this section. Its relaxation in subsequent sections reveals the broader usefulness of our findings.

3.1 The Main Business Cycle Shock: Targeting Unemployment

A key finding in this paper is that the shocks that target the aggregate quantities over the business-cycle frequencies can be thought of as interchangeable facets of (what we call) the MBC shock. But as our anatomy consists of individual cuts of the data, we need to start with one of these shocks. We choose the shock that targets unemployment, rather than any of its “sister” shocks, because unemployment is the most widely recognized indicator of the state of the economy.

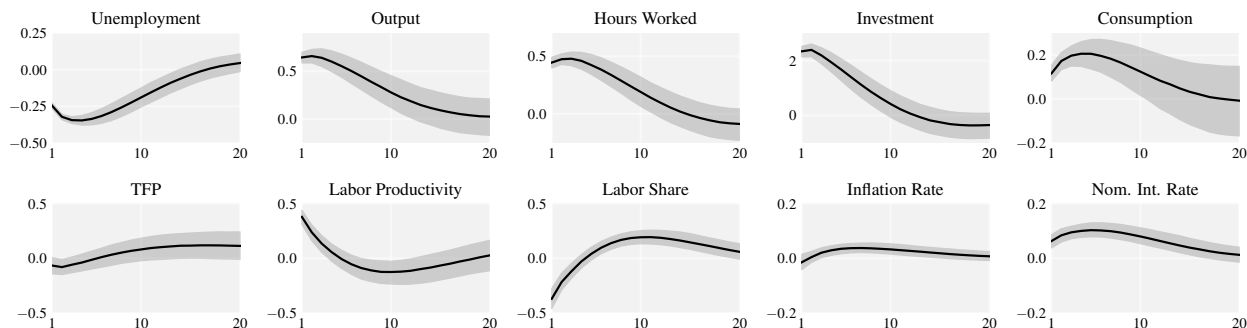
Figure 1 reports the impulse response functions (IRFs) of all the variables to this shock. As very similar IRFs are produced by the shocks that target the other key macroeconomic quantities, this figure plays a crucial role in our analysis: it serves as the empirical template for the propagation mechanism of models that contain a single or dominant business-cycle driver.

Table 1 adds more information about the identified shock by reporting its contribution to the volatility of all the variables over two frequency bands: the one used to construct it, which corresponds to the range between 6 and 32 quarters and is referred to as “Short Run” in the table; and a different band, which is referred to as “Long Run” and corresponds to the range between 80 quarters and ∞ . This helps assess whether the identified shock can indeed account for the bulk of the business-cycle

¹¹As shown in Online Appendix D, similar results obtain if we repeat our exercises on the time domain.

fluctuations in the key macroeconomic quantities, as well as how large its footprint is on inflation or the long run.¹²

Figure 1: Impulse Response Functions to the MBC Shock



Note: Impulse Response Functions of all the variables in our VAR to the identified MBC shock. Horizontal axis: time horizon in quarters. Shaded area : 68% Highest Posterior Density Interval (HPDI henceforth).

Table 1: Variance Contributions

	u	Y	h	I	C
Short Run (6-32 quarters)	73.71	58.51	47.72	62.09	20.38
	[66.80,79.94]	[50.65,65.07]	[40.77,54.45]	[54.09,68.46]	[13.61,27.53]
Long Run (80- ∞ quarters)	20.83	4.64	5.45	5.16	4.13
	[8.37,38.94]	[0.52,15.85]	[1.25,15.40]	[0.79,16.81]	[0.38,14.93]
	TFP	Y/h	wh/Y	π	R
Short Run (6-32 quarters)	5.86	23.91	27.02	6.96	22.27
	[2.44,10.96]	[17.27,31.22]	[18.39,35.93]	[3.24,12.28]	[14.22,30.97]
Long Run (80- ∞ quarters)	4.09	3.88	3.12	5.77	9.12
	[0.41,14.48]	[0.37,14.19]	[0.78,10.16]	[1.70,13.54]	[2.68,20.00]

Note: Variance contributions of the MBC shock at two frequency bands. The first row (Short Run) corresponds to the range between 6 and 32 quarters, the second row (Long Run) to the range between 80 quarters and ∞ . The shock is constructed by targeting unemployment over the 6-32 range. The notation used for the variables is the same as that introduced Section 2. 68% HPDI into brackets.

What are the main properties of the identified shock?

First, over the business-cycle frequencies, it explains about 75% of the volatility in unemployment, 60% of that in investment and output, and 50% of that in hours. It also gives rise to a realistic business cycle, with all the aforementioned variables, as well as consumption, moving in tandem. These properties together with those reported in the next subsection justify labeling the identified shock as the “main business cycle shock.”

Second, the identified shock contains little statistical information about the business-cycle variation in either TFP or labor productivity. This is *prima facie* inconsistent, not only with the baseline

¹²Figure 12 in Online Appendix D contains similar information as Table 1, but applied to the time domain: it reports the contributions of the identified shock to the forecast error variances (FEV) of the variables at different horizons.

RBC model, but also with a class of models that let financial or other shocks trigger business cycles only, or primarily, by causing endogenous movements in productivity. We expand on this point in Section 3.4.¹³

Third, the contribution of the shock to economic activity peaks within a year of the occurrence of the shock, fades out before long, and leaves a negligible footprint on the long run. This finding extends and reinforces the key message of Blanchard and Quah (1989): what drives the business cycle appears to be distinct from what drives productivity and output in the longer term. This point is further corroborated later.

Fourth, the shock triggers a small, almost negligible, and delayed movement in inflation. This precludes the interpretation of the identified shock as an inflationary demand shock of the textbook variety, but allows its interpretation as either a demand shock of the DSGE type, where the use of sufficiently flat wage and price Philips curves guarantee that demand shocks can generate sizable departures from flexible-price outcomes without commensurate inflationary pressures, or a demand shock operating outside the realm of nominal rigidities and Philips curves, as in the models cited in footnote 3. We revisit this point in Sections 3.5 and 6.

Fifth, the shock triggers a strong, procyclical movement in the nominal interest rate. And because inflation hardly moves, this translates to a strong, procyclical movement in the real interest rate. At face value, this seems consistent with a monetary policy that raises the nominal interest in response to the boom triggered by the identified shock, stabilizes inflation, and perhaps even closes the gap from flexible-price outcome (or, equivalently, tracks the natural rate of interest). This scenario is ruled out by the prevailing New Keynesian paradigm, because a gap from flexible-price outcomes is needed in order to accommodate demand-driven business cycles. But there is no way to verify or reject this assumption on purely empirical grounds, because the natural rate of interest and the flexible-price outcomes are not directly observable (and not even defined outside specific models).

Finally, the shock triggers a countercyclical response in the labor share for the first few quarters, which is reversed later on. Relatedly, when looking at the response of the real wage, as inferred by the difference between the response of the labor share and that of labor productivity, we see that the real wage remains relatively flat in response to the identified shock. This is consistent with the well-known, unconditional fact that real wages display very weak procyclicality, which is typically interpreted as being due to some form of real-wage rigidity.

3.2 The Main Business Cycle Shock: Targeting Other Quantities

Figure 2 compares the IRFs of the shock that targets the business-cycle volatility of the unemployment rate (black line) to the IRFs of the shocks that are identified by targeting the business-cycle volatility of a some other key macroeconomic quantities: GDP (red line), hours (green line), investment (blue line), and consumption (gray line).

As is evident from the figure, these shocks are nearly indistinguishable: targeting any one of

¹³Also, the mild and short-lived, procyclical response of labor productivity to the identified shock could reflect the impact of the latter on capacity utilization. This hypothesis is corroborated by the evidence in Online Appendix G.2, where the MBC shock is shown to lead to a similarly short-lived but significant increase in capacity utilization while accounting for more than 50% of its volatility at the business-cycle frequencies.

the aforementioned variables seems to give rise to the same dynamic comovement properties. This explains the rationale of interpreting these reduced-form shocks as interchangeable facets of the empirical footprint of the same propagation mechanism, or of what we have called the MBC shock.¹⁴

Figure 2: The Various Facets of the MBC Shock, IRFs

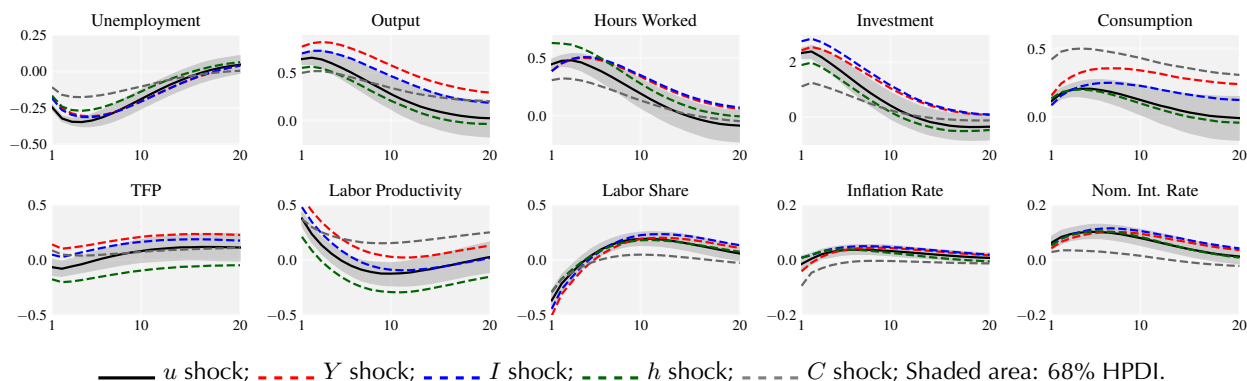


Table 2 paints a complementary picture in terms of the variance contributions: the shock that targets *any* one of unemployment, GDP, hours and investment explains the bulk of the business-cycle volatility in all of these variables. The following caveat applies to consumption: the shock that targets consumption explains less than one quarter of the fluctuations in unemployment, hours, or investment; and symmetrically, the other shocks that make up our MBC template account for less than one quarter of the fluctuations in consumption.¹⁵ Nonetheless, the consumption shock conforms quite closely to the other shocks when it comes to the IRFs it induces, as well as in terms of the disconnect from TFP and inflation. That is, it shares roughly the same propagation mechanism.

Finally, the interchangeability property extends from the IRFs to the times series produced by the different facets of the MBC shock. This is shown in Table 3. The table reports, for any of the variables of interest, the correlations between the times series of that variable produced by the unemployment shock and that produced by any of its sister shocks. The nearly perfect correlations seen in this table mean that that recovered shocks are essentially the same, not only in terms of IRFs, but also in terms of realizations, as manifested in the times series they produce for the main variables of interest.¹⁶

¹⁴Recall that, for our purposes, different shocks that are observationally equivalent in terms of IRFs are essentially one and the same shocks. This perspective is consistent with standard practice in both the SVAR and the DSGE literatures: as echoed in the quote from Cochrane cited in footnote 2, shocks are visible—and hence distinguishable—only through the dynamic comovement patterns they induce in the variables of interest.

¹⁵Recall that our measure of consumption excludes spending on durables, which is instead included in investment.

¹⁶Let $X \in \{u, Y, C, I, h\}$ denote any one of the variables of interest. Next, let X_s denote the bandpass-filtered time series of the predicted value of that variable produced by the shock that targets the variable $s \in \{u, Y, C, I, h\}$ (where s may or may not coincide with X). The typical cell in Table 3 reports, for a variable X (across rows) and a shock $s \neq u$ (across columns), the correlation of X_s and X_u . This summarizes the information seen in Figure 9 in Appendix B, which depicts the full scatterplots of the series X_s against the series X_u , for all X and s . The similarity is also present in terms of the *innovations* that correspond to the different shocks. But these innovations are not meaningful per se, nor are the corresponding column vectors of the matrix Q . Rather, what matters is how these innovations propagate over time and across variables, which is what the IRFs seen in Figure 2 reveal, or how they manifest themselves in terms of the predicted time series X_s , which explains the focus of Table 3 and Figure 9.

Table 2: The Various Facets of the MBC Shock, Variance Contributions

Targeted Variable	u	Y	h	I	C
Unemployment	73.71 [66.80,79.94]	58.51 [50.65,65.07]	47.72 [40.77,54.45]	62.09 [54.09,68.46]	20.38 [13.61,27.53]
Output	56.24 [48.94,61.93]	80.13 [72.80,86.44]	44.73 [37.36,51.68]	67.13 [60.72,72.82]	33.03 [25.04,40.44]
Hours Worked	49.84 [42.43,56.53]	47.54 [38.20,55.67]	70.45 [64.25,77.04]	47.99 [38.49,55.96]	21.78 [15.30,29.22]
Investment	59.03 [51.73,64.55]	66.60 [60.40,72.21]	45.20 [37.93,51.98]	80.29 [72.82,86.97]	19.01 [12.27,27.34]
Consumption	19.19 [12.12,27.73]	31.59 [21.81,40.90]	20.15 [13.60,27.68]	17.10 [9.96,25.94]	68.30 [60.61,75.53]
Targeted Variable	TFP	Y/h	wh/Y	π	R
Unemployment	5.86 [2.44,10.96]	23.91 [17.27,31.22]	27.02 [18.39,35.93]	6.96 [3.24,12.28]	22.27 [14.22,30.97]
Output	4.24 [1.76, 8.32]	41.31 [35.29,47.43]	40.20 [32.75,47.40]	10.47 [5.97,16.75]	16.89 [11.00,26.08]
Hours Worked	11.62 [6.14,18.14]	22.61 [15.58,29.66]	19.47 [11.73,29.24]	7.23 [3.32,13.31]	22.38 [15.09,31.87]
Investment	3.81 [1.38, 7.83]	33.74 [27.72,40.30]	36.44 [29.21,44.21]	7.69 [3.65,12.96]	21.51 [13.91,30.28]
Consumption	1.57 [0.59, 3.57]	12.93 [7.40,20.54]	10.31 [5.08,17.88]	9.93 [4.70,17.05]	4.50 [1.38,10.63]

Note: The rows correspond to different targets in the construction of the shock. The columns give the contributions of the constructed shock to the business-cycle volatility of the variables. In this and in all following tables, square brackets contain the 68% HPDI.

Table 3: Correlations of Conditional Times Series

	Y shock	I shock	C shock	h shock
Unemployment	0.973	0.982	0.931	0.941
Output	0.997	0.997	0.991	0.992
Investment	0.990	0.996	0.938	0.989
Consumption	0.987	0.983	0.739	0.964
Hours Worked	0.973	0.982	0.931	0.941

Note: Each row reports the correlation between each bandpass-filtered variable as predicted by the unemployment shock and that predicted by the other facets of the MBC shock.

3.3 The MBC Shock and Principal Component Analysis

The finding that there is a single force that drives the bulk of the fluctuations in various measures of economic activity naturally invites a comparison to principal component analysis (PCA). Is our main business cycle shock similar to the first principal component of the data over business cycle frequencies? And if yes, are there any reasons to favor employing our method over PCA in pursuing an anatomy of the business cycle?¹⁷

To address the first question, we perform PCA in the frequency domain. For each variable $X_j \in \{u, Y, h, I, \dots\}$, we construct the bandpass-filtered variable X_j^{bc} that isolates its business cycle frequencies (6-32 quarters). We then use the covariance matrix of all the filtered variables to construct the first principal component, denoted by $PC1^{\text{bc}}$. We finally project each X_j^{bc} on $PC1^{\text{bc}}$ and compute the R-square of the projection. This gives the percentage of the business-cycle volatility in variable j accounted for by the principal component.¹⁸

Four different versions of this exercise are carried out. In the first version, X^{bc} is derived by applying the bandpass filter directly on the raw data, variable by variable. In the second version, we first run a VAR on all the variables jointly, use it to estimate the cross-spectrum of the data, and then construct the band passed variables X_j^{bc} . Hence, the bandpass filter is the ideal one in the latter case, whereas it is only an approximate one in the former.

In the third and fourth version, the filtered variables are normalized by their respective standard deviations before extracting the first principal component. Such a normalization is often employed in the PCA literature in order to cope with scaling issues and/or to focus on the co-movements in the data. But it also reduces the role played by the more volatile variables (e.g., investment), which may or may not be desirable depending on the context. As we do not have a strong prior on how to properly weight the variables, we carry the exercise on both normalized and non-normalized data.

The results are reported in Table 4. In all cases, the first principal component accounts for the bulk of the business-cycle volatility in unemployment, hours, output, and investment but only a small fraction of the business-cycle volatility in either TFP or inflation.

Table 4: The First Principal Component, Business Cycle Frequencies

	u	Y	h	I	C	TFP	Y/h	wh/Y	π	R
Raw Data	75.33	92.26	81.24	99.80	60.19	6.10	17.73	3.02	2.33	12.27
VAR-Based	63.31	87.33	62.47	99.72	26.67	1.22	29.19	14.16	0.68	8.10
Normalized Raw	91.50	86.76	91.26	80.59	76.75	17.32	2.59	0.33	19.22	38.21
Normalized VAR	82.87	93.86	78.12	82.59	54.86	1.81	19.36	5.28	2.09	19.63

This is reassuring: the picture obtained here is similar to that obtained in Table 2 about the various facets of the MBC shock. As shown in Online Appendix F, a similarly reassuring connection holds between the main long-run shock obtained by our method in the next section and the principal component obtained by applying PCA to the long-run components of the data. But are there any

¹⁷We thank an anonymous referee for suggesting the exploration of these questions.

¹⁸Recall that the first principal component is constructed by taking the eigenvector corresponding to the largest eigenvalue of the covariance matrix. It is thus designed to account for as much as possible of the volatility and the co-movement of all the (filtered) variables at once.

reasons to favor our method over PCA for performing the type of business cycle analysis carried out in this paper? We think there are.

First, note that the two principal components that account for, respectively, the business-cycle and long-run movements in the data are orthogonal to each other *by construction*, because the two frequency bands do not overlap. Consequently, PCA is not useful for addressing the question of whether the forces that drive the business cycle and long run are related. That is, PCA cannot generate the type of information contained in the second row of Table 1, or its mirror image regarding the short term effects of the main long run shock shown in the sequel.

Second, PCA does not contain information about how the variables respond on impact and over time to a shock. That is, PCA does not accommodate the construction of IRFs, which are of paramount importance for our purposes.¹⁹

And third, by targeting each time an individual variable but also systematically varying the possible targets, our method avoids the difficulties associated with having to choose the “best” weights in PCA and, more importantly, helps reveal patterns that may prove useful in the validation of existing models or in the construction of new ones. Consider, in particular, the interchangeability pattern documented in Figure 2. This finding offers, not only hope for parsimonious models that aspire to account for the majority of business cycles with a single shock/mechanism, but also a test for models that employ a multitude of shocks/mechanisms. This will become clear in Section 6.

3.4 The Long Run and the Short Run

In the preceding analysis we recovered a MBC shock by targeting the business cycle frequencies. We now document the existence of an analogous object for the long run frequencies. We also discuss the implications of our results for theories that link the business cycle to technology and news shocks.

Consider the shocks constructed by targeting GDP, investment, consumption, TFP, or labor productivity at the frequencies corresponding to $80\text{-}\infty$ quarters.²⁰ Figure 3 and Table 5 show that these shocks are nearly indistinguishable in terms of either IRFs or variance contributions.²¹ Hence, one may advance the concept of the “main long-run shock” as the main driver of long-run movements in a manner analogous to that of the MBC.²²

This finding also motivates us to repeat our exercises using a VECM in which the aforementioned quantities share a common stochastic trend, while the remaining variables are stationary. The use of such a VECM instead of our baseline VAR is recommended if the analyst has a strong prior that the aforementioned quantities are cointegrated—a prior that is not only imposed in standard models

¹⁹A version of Dynamic Factor Analysis appropriately adapted to the frequency domain could address this caveat, thus serving as a closer substitute to our approach. However, the point we make next about PCA would still apply: such an approach would not necessarily reproduce the richness of information contained in our *multiple* cuts of the data and could miss critical properties such as the interchangeability of the various facets of the MBC shock.

²⁰Here, we omit the shocks that target the unemployment rate and hours worked per person because these variables do not have a long-run trend, at least in the context of most models.

²¹A similar picture emerges from inspection of the properties of the first principal component over these long term data; see Table 15 in Appendix F.

²²We have verified that the shocks considered here are nearly identical to those identified by targeting the frequency exactly at ∞ , which amounts to imposing a set of long-run restrictions as in Blanchard and Quah (1989) and Galí (1999).

Figure 3: Long-Run Shocks

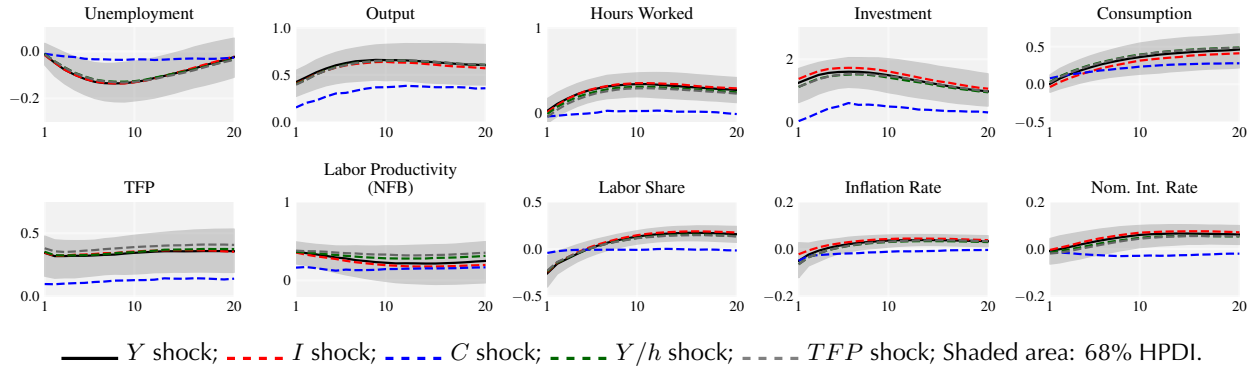


Table 5: Long-Run Shocks, Contributions at Long-Run Frequencies ($80-\infty$ q)

Targeted Variable	Y	I	C	TFP	Y/h
Output	99.59 [98.53,99.92]	95.94 [89.26,98.93]	99.47 [98.33,99.86]	95.66 [88.38,98.87]	96.92 [90.68,99.13]
Investment	96.88 [88.35,99.39]	97.83 [93.39,99.39]	96.41 [87.05,99.31]	91.62 [74.88,97.83]	91.75 [72.74,97.94]
Consumption	99.34 [97.57,99.85]	95.63 [87.91,98.81]	99.53 [98.23,99.90]	95.39 [87.38,98.81]	96.69 [90.51,99.12]
TFP	97.39 [88.33,99.48]	92.55 [76.41,98.11]	97.40 [88.33,99.49]	98.43 [94.49,99.70]	98.43 [93.92,99.67]
Labor Productivity	98.30 [91.73,99.60]	93.23 [77.39,98.28]	98.55 [92.92,99.66]	97.60 [91.37,99.50]	98.97 [95.10,99.84]

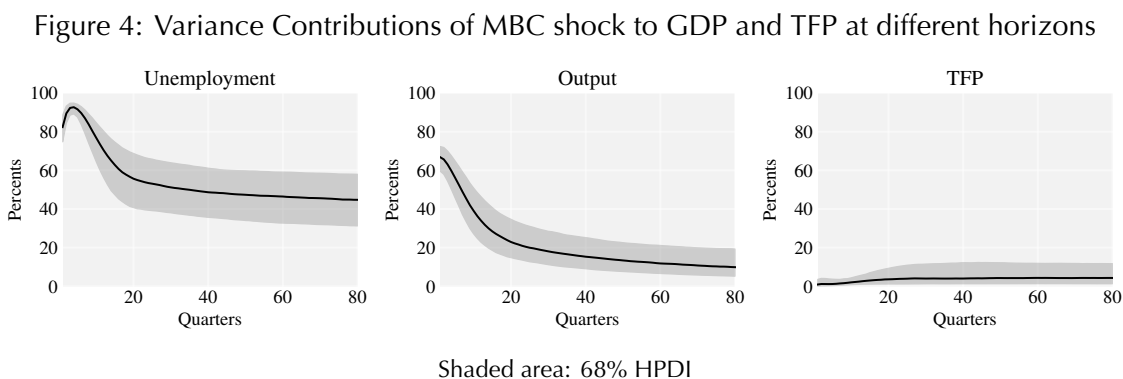
Table 6: VECM, Long-Run TFP Shock, Contributions at Business-Cycle Frequencies

u	Y	h	I	C
9.63 [3.46,18.43]	24.78 [11.41,40.32]	11.01 [4.99,19.60]	17.56 [7.31,29.53]	15.58 [5.71,27.20]
TFP	Y/h	wh/Y	π	R
22.01 [5.95,42.17]	21.89 [10.96,35.27]	10.19 [2.75,21.70]	12.59 [4.64,28.59]	7.26 [2.52,16.84]

but also corroborated by the evidence presented above as well as by familiar cointegration tests. For robustness, we also consider a variant VECM in which we add a second stochastic trend that drives inflation and the nominal interest rate. This helps capture the familiar indeterminacy of the long-run values of these variables in theoretical models and their high persistence in the actual data.

These VECMs produce essentially the same empirical regularities as those presented above, for both the short and the long run shocks. An example of this robustness is provided in Table 6. This table reports the contribution of the main long run shock, represented by the shock that targets TFP over the $80-\infty$ range, to the volatilities of all the variables over the 6-32 range. The emerging picture is essentially the mirror image of that contained in the second row of Table 1. There, we reported that the MBC shock has a small contribution to the long run. Here, we see that the shock that accounts for the long run has a small footprint on the business cycle.

The disconnect between the short and the long run can be seen in a more continuous manner when moving to the time domain. Figure 4 shows the contribution of the MBC shock to the volatility (FEV) of unemployment, output and TFP at different time horizons.²³ The MBC shock explains more than 60% of unemployment and output movements during the first two years, but less than 7% of the TFP movements at *any* horizon; and conversely, the main long run shock explains nearly all the long-run variation in investment and TFP, but less than 10% of the unemployment and investment movements over the first two year.²⁴



How do these findings compare to related ones in the existing literature?

First, consider Blanchard and Quah (1989). They seek to represent the data in terms of two shocks, a “supply shock” and a “demand shock.” To this goal, they run a VAR on two variables, GDP and unemployment; identify the supply shock as the shock that accounts for GDP movements in the very long run (at ∞) and the demand shock as the residual shock; and document that each one of these shocks accounts for nearly one half of business-cycle volatility in GDP. The additional information contained in our larger VAR reduces the contribution of—our various proxies of—the supply shock to between one tenth and one fifth.

Second, consider Uhlig (2003). This work, too, pursues a two-shock representation of the data.

²³The MBC shock is still identified in the frequency domain and used in the VAR to generate forecast errors at horizons 1-100 in the time domain. The same picture emerges when the MBC is identified in the time domain; see Appendix E.

²⁴It is worth noting that the disconnect between the short and the long run extends from neutral technology, as measured by TFP, to investment-specific technology, as measured by the relative price of investment; see Appendix G.2.

The main difference from Blanchard and Quah (1989) is that it uses a different identification scheme: it identifies the two shocks that *jointly* maximize the prediction error variances in real GNP for horizons between 0 and 5 years. Uhlig offers a tentative interpretation of one shock as being a productivity shock of the RBC type and the other as a cost-push shock of the New Keynesian type. This interpretation finds little support in our anatomy, especially due to our finding of a disconnect between our main business-cycle shock and TFP at all horizons.²⁵

Third, consider Galí (1999) and Francis et al. (2014). Our long-run TFP shock is essentially the same as the technology shock identified in those papers. Tables 5 and 6 confirm their finding that this shock has a small contribution to the business cycle. This property extends to the robustness exercises reviewed in Section 4. It also connects to the point made next.

Finally, consider Beaudry and Portier (2006). The first part of that paper uses a two-variable VAR with TFP and the SP500 index to identify a shock—interpreted as TFP news—that accounts for the bulk of both the short-run movements in stock prices and the long-run movements in TFP. The second part proceeds to argue, with the help of three- to five-variable VARs and more delicate identifying assumptions, that TFP news account for about 50% of the short-run volatility in hours and total private spending, about 80% of that in consumption, and about 80% the long-run movements in private spending. In short, TFP news emerges as the main driver of *both* the business cycle and the long run.

This picture is hard to reconcile not only with the one painted here but also with the earlier results of Galí (1999) and Francis et al. (2014). If TFP news was the main driver of *both* the business cycle and the long run, one would expect to see a strong connection between the two. But as seen in Table 6, the main long-run shock identified here accounts for only 10% of the short-run volatility in unemployment and hours and 17% of that in investment. The identification strategies employed in Galí (1999) and Francis et al. (2014) use slightly different but produce similar numbers for the short-run contributions of the shock that drives the long run. And perhaps most tellingly, Figure 4 above shows that the MBC shock accounts for nearly zero of the FEV of TFP at any horizon. That is, the MBC shock itself contains no news about future TFP.²⁶

We believe that, while TFP news may be a non-trivial contributor to macroeconomic fluctuations, the numbers reported by Beaudry and Portier (2006) are biased upwards due to: (i) the use of relatively small VARs; and (ii) the reliance on delicate identifying assumptions. We elaborate on these points in Section 5 and Appendix C. There, we use a semi-structural exercise, based on our anatomy of the data, to shed new light on the business-cycle effects of technology and news shocks. Our explorations suggest that the contribution of news shocks to unemployment fluctuations is 11-13%, which is quite more modest than that suggested by Beaudry and Portier (2006) but still non-trivial.

²⁵We emphasize that the interpretation offered in Uhlig (2003) was tentative as that paper was not completed. Also note that the approach adopted in that paper allows for the identification of the two shocks *together* but does not separate one shock from the other, so the aforementioned interpretation relied on particular orthogonalizations. Finally, because the VAR considered in that paper did not contain TFP, the disconnect documented here could not have been detected.

²⁶Moreover, these findings are not due to the absence or presence of Stock Prices in the VARs. As can be seen in row 9 of Tables 8 and 9, which appear in the sequel and report results from various robustness exercises, the inclusion of Stock Prices is inconsequential for the properties of the MBC shock as well as for those of the short and long run TFP shocks.

3.5 Inflation and the Business Cycle

We now turn attention to the nexus of real economic activity and inflation. Seen through the lens of our method, the link is weak. First, as shown in the first row of Table 7 (which repeats a portion of the first row of Table 1), the identified MBC shock accounts for only 7% of the business-cycle variation in inflation, which is as low as the corresponding number for TFP. Second, the shock that targets inflation explains 83% of the business-cycle volatility in inflation and only 4 to 8% of that in unemployment, output, and investment. Finally, the shock that targets inflation explains only 2% of the labor share, a commonly used proxy of the real marginal cost, or the driving force of inflation according to the New Keynesian Philips Curve (Galí and Gertler, 1999); and symmetrically, the shock that targets the labor share explains 86% of the labor share itself but only 4% of inflation.

Table 7: Inflation and the Business Cycle

Targeted Variable	u	Y	π	Wh/Y
Unemployment	73.71 [66.80,79.94]	58.51 [50.65,65.07]	6.96 [3.24,12.28]	27.02 [18.39,35.93]
Inflation	4.24 [1.62, 8.20]	7.88 [3.77,12.87]	83.03 [76.11,88.46]	1.96 [0.66, 4.60]
Labor Share	26.01 [18.13,33.99]	35.33 [27.88,43.68]	4.03 [1.45, 7.94]	85.59 [80.04,90.02]

These findings pose a challenge for the prevailing formalization of demand-driven business cycles. Because of its transitory nature and its disconnect from TFP, it is tempting to interpret the MBC shock in the data as a demand shock in the New Keynesian model. In that model, demand shocks generate business cycles only by inducing positive output gaps from flexible-price outcomes. Furthermore, because replicating flexible-price outcomes is equivalent to stabilizing inflation, such gaps are the main “fundamental” driving inflation.²⁷ In particular, the baseline New Keynesian model imposes that inflation is the best predictor of future output gaps, or real marginal costs, similarly to how the baseline asset-pricing model imposes that the asset prices is the best predictor of future earnings. From this perspective, Table 7 suggests that the failure of the two models are comparable: the link between inflation and real economic activity is as strong as the link between asset prices and earnings.

Another challenge emerges from contrasting the magnitude of the actual inflation response to the identified MBC shock to that predicted the textbook version of the New Keynesian model under the interpretation of this shock as an aggregate demand shock in the theory: as illustrated in Figure 17 in Online Appendix H.1, the predicted response is over ten times larger than the observed one.

These challenges are familiar, albeit through other lenses,²⁸ and have already shaped the existing DSGE literature. This literature has sought to address them by making the Phillips curve much flatter

²⁷This is true even in versions of the model that feature markup shocks and the Hybrid NKPC. However, inflation can also be driven by other factors, including shifts in market expectations about the central bank’s long-run inflation target.

²⁸For instance, the weak comovement of inflation and real economic activity is also evident in the unconditional moments, although it is less pronounced than that seen in Table 7. See also the survey by Mavroeidis, Plagborg-Møller, and Stock (2014) on the large empirical literature on the various incarnations of the Phillips curve.

than, not only its textbook version, but also that implied by menu-cost models calibrated to micro-economic evidence; and by attributing the bulk of inflation fluctuations to large, mysterious, cost-push shocks.²⁹ The empirical foundations of these and various other add-ons that help improve the fit of DSGE models remains a contested issue, which is certainly beyond the scope of this paper. Instead, we only wish to raise the possibility that the MBC shock in the data represents an aggregate demand shock of a different kind that that presently formalized in the New Keynesian framework, namely one that operates *inside* its flexible-price core rather than outside it.

4 Robustness

In this section we report results from an extensive battery of robustness exercises we have conducted. The main exercises are described below, the rest are delegated to the Appendix.

Tables 8 and 9 describe the variance contribution of the MBC shock over business cycle and longer term frequencies, respectively, and across many alternative specifications (different samples, statistical models estimated, set of variables, numbers of lags). As in Table 1, we use the shock that targets unemployment as the measure of the MBC shock. Appendix G reports similar tables for the shocks that target GDP, hours, etc...The first row in Tables 8 and 9 corresponds to our baseline specification, that is, it repeats the information from Table 1. The remaining rows correspond to ten alternative specifications.

Row 2 corresponds to a VAR with four lags instead of two; the results with six or eight lags are almost the same and are thus omitted. Rows 3 and 4 correspond to two VECMs: the first allows for a single unit root that drives the real quantities, while the second allows inflation and the nominal interest rate to be driven by the first, “real” root as well as by a second, “nominal” root. Row 5 extends the sample backwards to 1948, by replacing the Federal Reserve Rate with the 1-month T-bill rate. Row 6 constrains the sample to 1960-2007, leaving out the Great Recession and the ZLB; this is also the period used in the estimation and validation of the two DSGE models considered in the next section. Rows 7 and 8 split the sample to two sub-samples, pre- and post-Volcker. Row 9 adds the following three variables to the VAR: the SP500 index, the relative price of investment, and capital utilization. Row 10 adds the credit spread, a common measure of the severity of financial frictions. Finally, row 11 considers a version where consumption and investment are deflated by their respective, chained-type price indices rather than the GDP deflator, as a way to take relative-price effects into account.³⁰

The results speak for themselves. As we move across specifications (rows), the contribution of the identified shock to the variance of the key macroeconomic quantities remains almost unchanged.³¹

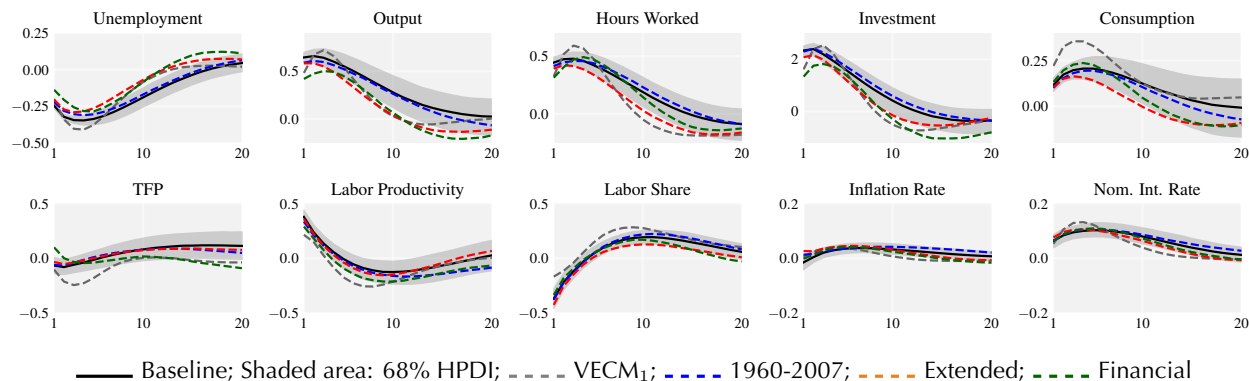
²⁹Note that, if taken literally, the assumed flatness of the Philips curve requires that inflation be unresponsive, not only to the output gap, but also to any cost-push shock, which further strains the offered explanation of the data.

³⁰Given that consumption is the sum of non durables and services, and investment is the sum of gross private domestic investment and durables, some care must be take to build the corresponding chained type price indices. The construction of the indices is detailed in Appendix G.5.

³¹The only sensitivities worth mentioning are the following. First, the VECMs raise slightly the long-run footprint of the MBC shock and more noticeably its short-run co-movement with consumption. And second, the pre-Volcker sample features a smaller disconnect between real economic activity and inflation than the post-Volcker one. These findings are

Similar results obtain in additional robustness exercises which we have undertaken but omit here for the sake of saving space.³²

Figure 5: Robustness, IRFs



More importantly, the same robustness is present when considering the IRFs. We illustrate this in Figure 5 for the shock that targets unemployment for a select subset of the eleven specifications under consideration.³³ This is re-assuring as the properties of the IRFs, and in particular the interchangeability of the various facets of the MBC shock, represent the key criterion for judging the empirical plausibility of a model's propagation mechanism.³⁴

hardly surprising and, in any case, do not change the main picture.

³²For instance, we have verified that the properties of the MBC shock remain largely the same if we drop any one of the variables in our baseline VAR, or if we add labor market indicators such as vacancies. The results become sensitive only when the size of the VAR becomes very small. See Appendix C for an illustration. This is not surprising given the well-known fragility of small VARs. To the contrary, this fact along with the already reported robustness to the addition of stock prices and other variables suggests that our baseline VAR has the “right” size in order to reveal robust properties.

³³The remaining specifications are also similar. They are omitted only because they would have over-crowded the figure.

³⁴As can be seen by comparing the baseline and the 1960-2007 cases in Figure 5, the interchangeability property and the profile of the MBC shock are not sensitive to the inclusion or exclusion of the ZLB period. This fact may seem puzzling when viewed through the lenses of a model in which the ZLB constraint is binding and dramatically changes the propagation of the main driver(s) of the business cycle. But if this constraint is largely bypassed by the effective use of other policy tools, the main propagation mechanism seen in the data need not change as one moves between ZLB and non-ZLB samples; see Debortoli, Galí, and Gambetti (2019) for corroborating evidence. Yet another possibility is that the ZLB constraint matters for the amplitude of the business cycle but not for the propagation dynamics.

Table 8: Robustness, Short-Run Variance Contributions

	u	Y	h	I	C	TFP	Y/h	Wh/Y	π	R
[1] Benchmark	73.71	58.51	47.72	62.09	20.38	5.86	23.91	27.02	6.96	22.27
	[66.80,79.94]	[50.65,65.07]	[40.77,54.45]	[54.09,68.46]	[13.61,27.53]	[2.44,10.96]	[17.27,31.22]	[18.39,35.93]	[3.24,12.28]	[14.22,30.97]
[2] 4 lags	74.49	58.23	49.16	62.42	21.20	6.28	23.10	27.87	6.91	24.75
	[67.98,80.77]	[50.51,65.05]	[42.24,56.10]	[55.15,69.04]	[14.13,28.78]	[2.82,11.74]	[16.83,31.02]	[18.93,37.34]	[3.23,12.15]	[16.20,33.77]
[3] VECM(1)	62.43	50.27	48.81	53.39	34.88	18.13	23.80	24.11	10.46	33.37
	[56.47,68.44]	[43.46,57.44]	[42.14,55.91]	[47.05,60.01]	[26.27,44.47]	[9.03,29.45]	[17.14,32.73]	[16.36,34.17]	[4.39,20.13]	[19.07,48.60]
[3] VECM(2)	64.85	54.99	48.82	53.78	44.93	12.17	19.51	29.71	11.29	19.51
	[57.60,71.25]	[46.53,62.59]	[42.52,55.66]	[46.37,60.86]	[33.73,55.68]	[6.00,19.88]	[13.11,27.14]	[20.04,39.49]	[5.09,19.32]	[10.94,32.92]
[5] 1948-2017	78.98	65.32	49.61	63.76	19.52	6.14	26.53	29.62	5.16	16.94
	[72.86,84.10]	[59.25,71.33]	[43.55,55.83]	[57.87,70.19]	[13.70,26.91]	[2.51,11.05]	[19.68,33.57]	[22.10,37.53]	[2.28,10.00]	[10.37,24.31]
[6] 1960-2007	68.15	59.93	55.99	65.02	20.67	6.02	25.04	29.96	10.70	27.03
	[61.82,73.98]	[48.14,68.85]	[47.10,63.10]	[55.39,72.59]	[13.52,31.01]	[2.24,13.76]	[16.29,36.15]	[19.57,43.29]	[5.49,18.89]	[16.86,37.53]
[7] pre-Volcker	74.23	56.75	43.21	61.50	23.43	6.82	30.69	28.43	17.45	27.60
	[64.05,82.35]	[45.87,66.62]	[32.38,53.49]	[51.63,70.37]	[13.58,35.24]	[2.45,15.11]	[20.09,42.11]	[16.92,42.01]	[9.39,28.74]	[16.81,40.08]
[8] post-Volcker	73.39	50.37	50.65	58.44	20.23	7.94	18.46	23.01	4.65	15.05
	[65.47,80.53]	[41.45,58.81]	[42.60,59.01]	[50.17,66.23]	[12.46,28.65]	[3.67,14.49]	[11.61,26.94]	[14.23,33.51]	[1.74,10.06]	[7.48,25.22]
[9] Extended	59.33	50.61	45.50	52.91	21.83	4.81	26.69	27.82	12.12	28.99
	[53.73,65.69]	[43.05,57.99]	[39.71,51.26]	[44.97,60.17]	[14.87,31.14]	[1.95,10.39]	[19.36,34.75]	[14.05,44.15]	[6.57,19.70]	[17.38,42.75]
[10] Financial	68.57	57.56	46.84	59.95	25.94	7.04	27.20	26.86	8.42	26.59
	[62.38,74.87]	[49.74,64.87]	[39.39,54.03]	[52.26,66.82]	[17.80,34.98]	[3.10,12.97]	[19.45,35.96]	[18.53,37.07]	[3.77,14.98]	[16.82,36.24]
[11] Chained-Type C&I	81.41	59.04	45.96	61.52	17.36	4.03	20.35	20.19	5.82	23.17
	[75.30,86.36]	[52.45,64.82]	[39.33,52.36]	[54.39,67.49]	[12.10,23.41]	[1.56, 7.51]	[14.80,26.64]	[13.97,26.72]	[2.62,10.41]	[16.31,30.38]

Table 9: Robustness, Long-Run Variance Contributions

	u	Y	h	I	C	TFP	Y/h	Wh/Y	π	R
[1] Benchmark	20.83 [8.37,38.94]	4.64 [0.52,15.85]	5.45 [1.25,15.40]	5.16 [0.79,16.81]	4.13 [0.38,14.93]	4.09 [0.41,14.48]	3.88 [0.37,14.19]	3.12 [0.78,10.16]	5.77 [1.70,13.54]	9.12 [2.68,20.00]
[2] 4 lags	18.22 [7.27,34.06]	4.39 [0.61,14.67]	5.19 [1.32,15.39]	4.94 [0.89,15.34]	3.98 [0.47,13.45]	3.66 [0.41,13.53]	3.67 [0.40,13.01]	2.93 [0.57, 9.08]	5.44 [1.59,13.00]	9.81 [3.07,20.60]
[3] VECM(1)	12.97 [4.50,29.34]	14.07 [2.53,29.11]	8.06 [2.67,18.73]	14.07 [2.53,29.11]	14.07 [2.53,29.11]	14.07 [2.53,29.11]	14.07 [2.53,29.11]	13.91 [3.26,29.03]	7.50 [2.76,17.41]	13.82 [4.77,26.70]
[4] VECM(2)	23.29 [8.05,47.79]	16.70 [3.31,37.32]	9.22 [3.13,20.76]	16.70 [3.31,37.32]	16.70 [3.31,37.32]	16.70 [3.31,37.32]	16.70 [3.31,37.32]	10.55 [2.62,26.43]	8.66 [2.08,22.19]	8.66 [2.08,22.19]
[5] 1948-2017	31.82 [15.20,50.79]	7.44 [1.22,19.37]	4.43 [1.34,14.13]	7.80 [1.52,20.10]	6.66 [0.96,18.27]	7.20 [1.12,19.01]	6.72 [0.98,17.50]	4.85 [1.47,11.57]	3.37 [0.80, 9.24]	4.91 [1.30,12.70]
[6] 1960-2007	11.85 [5.40,22.31]	4.17 [0.52,16.00]	8.83 [3.25,18.36]	4.84 [0.72,16.69]	3.96 [0.43,15.18]	4.11 [0.73,14.05]	5.29 [1.45,16.02]	5.63 [1.52,15.31]	12.48 [4.85,23.35]	21.09 [8.45,35.63]
[7] pre-Volcker	29.37 [9.83,55.35]	8.15 [1.21,26.52]	9.33 [2.55,23.58]	8.23 [1.49,25.68]	7.10 [1.06,24.84]	7.31 [0.96,25.64]	7.55 [0.93,25.89]	7.17 [1.78,22.43]	8.82 [2.07,24.81]	18.60 [5.74,41.73]
[8] post-Volcker	19.30 [6.59,38.92]	3.58 [0.80,12.17]	9.96 [3.93,20.78]	6.07 [2.06,15.37]	3.04 [0.49,12.11]	3.41 [0.55,11.59]	3.03 [0.47,12.04]	5.05 [1.40,13.61]	9.54 [2.63,25.64]	14.30 [4.49,32.05]
[9] Extended	9.49 [3.03,24.04]	4.52 [0.45,17.60]	3.96 [1.11,11.23]	4.58 [0.78,18.25]	4.43 [0.40,16.92]	4.39 [0.59,17.66]	4.59 [0.52,17.40]	4.36 [0.79,14.99]	7.03 [2.20,16.45]	11.23 [2.88,24.32]
[10] Financial	16.97 [5.77,34.68]	4.85 [0.54,15.56]	4.85 [1.04,14.40]	5.20 [0.74,16.24]	4.40 [0.53,14.80]	4.26 [0.59,14.78]	3.98 [0.48,13.92]	3.40 [0.75,10.82]	5.06 [1.60,12.87]	8.35 [2.38,18.61]
[11] Chained-Type C&I	13.94 [5.61,27.35]	3.79 [0.49,14.58]	5.24 [1.00,15.39]	3.73 [0.55,13.92]	3.63 [0.46,14.16]	3.67 [0.54,13.27]	3.20 [0.42,12.74]	3.88 [1.11,11.46]	7.41 [2.22,17.54]	11.91 [4.11,25.74]

5 Interpretation

In this section, we first summarize what can be learned from the properties of our anatomy if one views them from a parsimonious, single-shock perspective. We then discuss the robustness of such lessons and the use of our anatomy outside the realm of single-shock models.

5.1 The Lesson for Parsimonious, Single-Shock Models

In the beginning of the Introduction, we asked: Is it possible to account for the bulk of the business cycle with a parsimonious, single-shock model? And if so, how should this shock look like? Our empirical findings provide the following answer.

Tentative lesson. *It is possible to account for the bulk of the business-cycle fluctuations in unemployment, hours, GDP, investment, and, to a somewhat lesser extent, consumption using a parsimonious, one-shock model. This shock must have the following key properties:*

- *it causes strong, positive, and transient comovements in the aforementioned quantities;*
- *it is an indicator of the short-run economic outlook and not of the medium- and long-run prospects;*
- *it is essentially orthogonal to both TFP and inflation at all horizons;*

As already discussed, these properties are hard to reconcile with the baseline RBC model, as well as with models that attribute the bulk of the business cycle to news about productivity and income in the medium to long run. They also speak against models in which financial, uncertainty, or other shocks matter primarily by triggering endogenous procyclical movements in aggregate TFP.³⁵ In contrast, the evidence seems consistent with a shock that triggers transitory movements in the labor wedge—but only insofar as these movements occur without commensurate movements in aggregate TFP and without opposite movements in the real wage. This rules out shocks to labor supply, as well as productivity shocks intermediated by labor-market frictions, but leaves room for other possibilities.³⁶ The evidence is also consistent with the Keynesian narrative that the bulk of the business cycle is due to shifts in aggregate demand—but only insofar as these shifts do not trigger significant movements in inflation. This, in turn, requires either a very flat Philips curve, as in state-of-the-art DSGE models, or demand shocks operating outside the realm of sticky prices and Philips curves, as in the literature cited in footnote 3 in the Introduction.

5.2 The Anatomy of Multi-Shock Models

So far, we have attempted to give structural meaning to the identified MBC shock through the lenses of models that aspire to explain the bulk of the observed business cycles with a single shock/propagation

³⁵Benhabib and Farmer (1994) and Bloom et al. (2018) are notable examples of such models: the former generates procyclical TFP movements out of animal spirits, the latter out of uncertainty shocks.

³⁶For example, in Angeletos, Collard, and Dellas (2018) the requisite movements in the measured labor wedge are the byproduct of higher-order uncertainty about the short-term economic outlook; in Arellano, Bai, and Kehoe (2018) these movements are attributed to the interaction of financial frictions and firm-level uncertainty shocks; and in Golosov and Menzio (2015) they obtain from animal spirits in frictional labor markets.

mechanism. This choice reflects, in part, a “philosophical” attitude. But this choice can be consequential for the interpretation of the MBC shock and more generally for the use of our anatomy. As suggested in the Introduction, the reason is that, in principle, any of the reduced-form objects contained in our anatomy may map into a un-interpretable combination of multiple theoretical shocks, none of which possesses the properties of the empirical object.

In this section, we use two examples to illustrate both this challenge and a resolution offered by our method. By design, our anatomy contains not only the reduced-form shock that targets unemployment over the business-cycle frequencies but also the other reduced-form shocks we have discussed in the previous section. This additional information comes into play when there is more than one shock in the model and holds the key for the effectiveness of our anatomy in multi-shock contexts. It turns out, at least within the set of semi-structural and fully-structural exercises consider in this and the next section, that this extra information suffices to pin down the nature of the main driving force of the business cycle, corroborating the main claim from the previous section, namely, that this force corresponds to a non-inflationary, demand shock.³⁷

Our first pedagogical example revisits the disconnect between the MBC shock and inflation within the textbook AD-AS paradigm. Let the AD and AS equations be given by, respectively,

$$y_t = -\pi_t + v_t^s \quad \text{and} \quad \pi_t = y_t + v_t^s, \quad (2)$$

where y_t denotes output, π_t denotes inflation, and v_t^d and v_t^s are the structural shocks to aggregate demand and aggregate supply, respectively. Imposing equilibrium gives

$$y_t = \frac{1}{2}(v_t^d + v_t^s) \quad \text{and} \quad \pi_t = \frac{1}{2}(v_t^d - v_t^s).$$

Assume now that v_t^d and v_t^s follow independent AR(1) processes, with the same persistence and variance. This implies (i) that each structural shock drives 50% of the volatility of both output and inflation and (ii) that output and inflation are orthogonal to each other. As a result, our “output shock,” which is here given by output itself, accounts for 100% of the fluctuations in output and 0% of those in inflation. This matches the MBC shock seen in the data, but rather than representing a single, dominant, non-inflationary, business-cycle shock, it is the sum of two distinct structural shocks, an inflationary and a dis-inflationary one.

Our second example demonstrates that a similar problem may plague the interpretation of the finding that the short and the long run factors are disconnected. Consider a model that contains two types of TFP shocks, namely, unanticipated and anticipated (news) shocks. Suppose further that each shock contributes 50% of the long-run volatility in TFP and 50% of the short-run volatility in unemployment. Finally, let the two shocks have symmetrically opposite effects on unemployment, one increasing it and the other decreasing it. The constructed “unemployment shock” then accounts for 100% of the short-run fluctuations in unemployment and 0% of the long-run fluctuations in TFP, which matches the disconnect of the short run and the long run seen in the data. Yet, the business cycle is not driven by a single, dominant, transitory shock. Instead, it is driven by two unit-root shocks, which have the same long-run effect on TFP but opposite short-run effects on unemployment.

³⁷Needless to say, this particular conclusion need not extend to *arbitrary* multi-shock models, because any structural interpretation is ultimately model-specific. But the use of our anatomy does extend, because the panoply of empirical restrictions contained can help model evaluation regardless of the model structure and the associated interpretation.

In both of these examples the basic challenge is the same: a reduced-form shock identified via our method does not map into a “true” structural shock. Clearly, this problem is not unique to our method. For instance, the second example also invalidates the interpretation of the “demand shock” identified in Blanchard and Quah (1989), or the “technology shock” identified in Galí (1999).³⁸ Nevertheless, additional, pertinent information can often remove this kind of challenge. Our approach provides ample such information in the form a panoply of conditional, cross-variable, static and dynamic restrictions, which can be deployed in both semi-structural and fully-structural endeavors.

To illustrate the use of our method in a semi-structural context, consider the second example. We used this example to argue that the disconnect between the short and the long run does not suffice to rule out technology, or news about it, as an important business-cycle driver. But this disconnect is not the only restriction contained in the anatomy. Another key restriction is that the MBC shock accounts for essentially zero of the TFP fluctuations at *any* horizon, including the short run. This helps reject the story proposed above: if that story were correct, the MBC shock would have been strongly correlated with current TFP, which is not the case.

We expand on this point in Appendix C. There, we impose no structure other than the assumption that TFP is driven by exactly two shocks, an unanticipated, permanent technology shock that has an immediate effect on TFP, and a news shock that has a delayed effect. We then show how two elements of our anatomy, namely the reduced-form shocks that target TFP in the short and the long run, provide an estimate of the contribution of the news shock to the unemployment fluctuations. This estimate turns out to be 13% in our baseline VAR and a bit lower in extended VARs that add stock prices.³⁹

In Online Appendix H, we carry out a similar semi-structural exercise in the context of the first example: we show that the simple story of offsetting demand and supply shocks does not work insofar as the supply shock can be proxied by the reduced-form shock that captures the bulk of the TFP movements in the data. To put it differently, the supply shock has to be a mysterious markup shock. We then proceed to conduct a second, fully structural yet relatively parsimonious, exercise: we revisit the example through the lenses of a two-variable, two-shock, New Keynesian model and ask what it takes for this model to match the relevant elements of our anatomy, namely the dynamic responses of output and inflation to our identified output and inflation shocks. The answer turns out to be consistent with the interpretation of the output shock in the data as a dominant, non-inflationary demand shock in the model (and of the inflation shock as the markup shock).

All in all, these simple exercises illustrate how one can utilize additional elements of our anatomy and/or additional theoretical structure to extend the use of our method to multi-shock environments. They also serve as a prelude for the more compelling analysis in the next section, which demonstrates the effectiveness of our method in the context of three state-of-the-art DSGE models. Relative to the exercises discussed above, those in the next section make use of both more elaborate theoretical

³⁸More generally, for any “structural” shock identified in the existing SVAR literature, one can always concoct examples that deconstruct it into a combination of two or more distinct shocks, none of which resembles the object identified in the data. Whether the problem is more severe in our case depends on whether one finds the premise of a dominant business-cycle shock less defensible than those other identifying assumptions in the literature.

³⁹Another function of Appendix C is to show how the estimated contribution of the news shock depends on the number of variables included in the VAR. This corroborates a point made in Section 3.4, that our conclusions about the importance of news shocks differ from those of Beaudry and Portier (2006) in large part due to the amount of data used.

structures and a broader set of elements from our anatomy, which helps keep the balance between degrees of freedom and empirical restrictions.

6 An Application to Medium-Scale DSGE Models

In the previous section we argued that our method can be of use in multi-shock environments thanks to the rich set of cross-variable, dynamic restrictions it contains. In this section, we put this argument on trial by applying our method to three off-the-shelf, state-of-the-art DSGE models. This application corroborates the structural interpretation of the MBC shock suggested on the basis of single-shock models. Most importantly, it demonstrates the probing power of our method, in the sense that the conditional moments comprising the anatomy help identify flaws in the propagation mechanism of models that may have gone unnoticed before.

We first study the properties of the sticky-price model in Justiniano, Primiceri, and Tambalotti (2010) and the flexible-price model in Angeletos, Collard, and Dellas (2018), henceforth referred to as JPT and ACD, respectively. The first is a representative of the New Keynesian, DSGE paradigm.⁴⁰ The second is an example of a recent literature that aims at disentangling demand-driven fluctuations from nominal rigidities and Philips curves (see the references in footnote 3).

Both models have been estimated and evaluated in the respective papers using familiar, pre-existing methods.⁴¹ The value added here is to revisit their performance through lenses of our new method. We thus take each model as is and use it to construct the linear combinations of the theoretical shocks that maximize the business-cycle volatility of GDP, investment, consumption or hours worked in the model. These objects are the theoretical counterparts to the reduced-form shocks that were previously identified in the data via our method. To avoid confusion between these objects and the primitive theoretical shocks, we henceforth refer to the former as “factors” and reserve the term “shocks” for the latter.⁴²

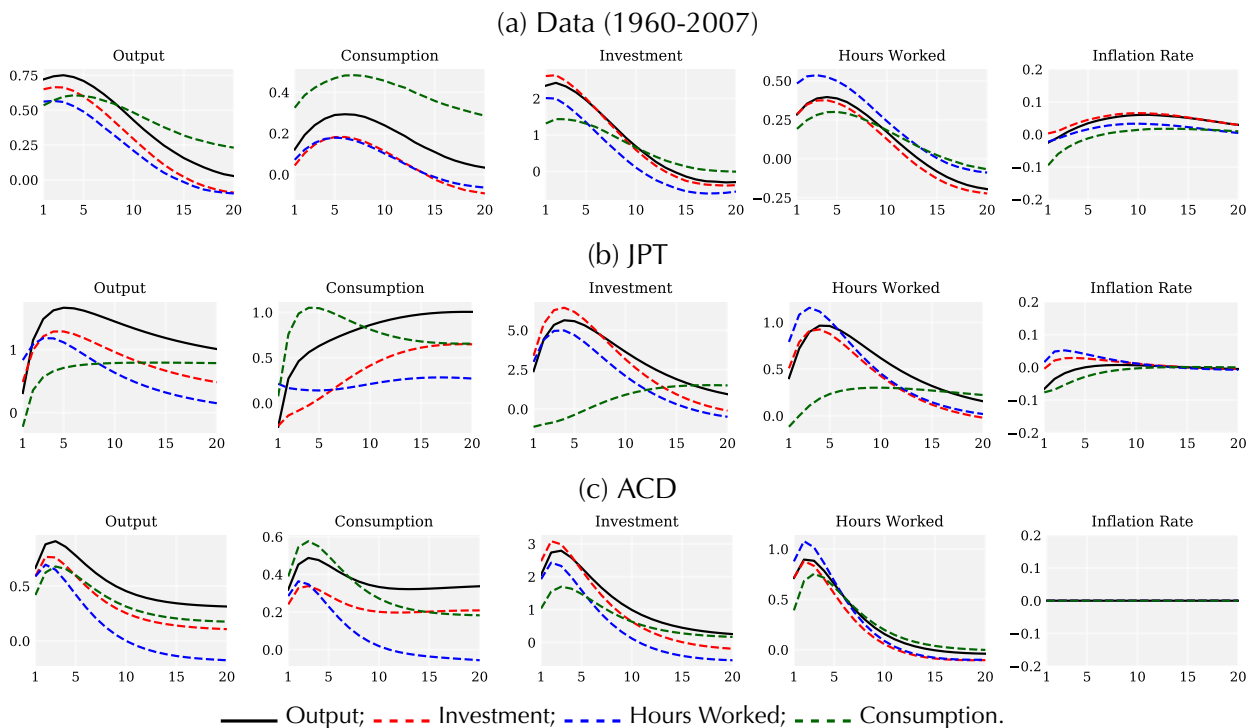
Figure 6 reports the IRFs of the various factors in the data (top panel) and in the two models (middle panel for JPT, bottom for ACD). As seen in this figure, the various factors are highly interchangeable in ACD, as they are in the data, whereas they are quite distinct in JPT. This is most evident in the responses

⁴⁰Indeed, it is essentially the same model as that in Smets and Wouters (2007), but with more appropriate mapping to the data. The measure of consumption used in Smets and Wouters (2007) includes expenditure on durables, which is at odds with the specification in the model. Justiniano, Primiceri, and Tambalotti (2010) fix this problem by including such expenditure to the measure of investment, just as we have done both here and in Angeletos, Collard, and Dellas (2018).

⁴¹In particular, both JPT and ACD have been estimated with Bayesian maximum likelihood. But whereas ACD has been estimated on the frequency domain using the levels of all variables, JPT has been estimated on the time domain using the growth rates of output, investment, and consumption. Another difference concerns the sample used: 1954Q3 to 2004Q4 in JPT vs 1960Q1-2007Q4 in ACD. As discussed later on and shown in Online Appendix I.2, re-estimating the JPT in the exact same way as ACD does not change the take-home lesson of this section. With this in mind, and to make sure that the two models are evaluated on the basis of the same sample period as that used in their estimation, the data underlying the top panels of Figure 6 refer to the VAR that appeared earlier as row [6] in Table 8, namely the one that spans the 1960Q1-2007Q4 period; as already emphasized, this makes little difference from our baseline specification.

⁴²Our “factors” should not be confused with those in dynamic factor analysis. Also, the construction of the factors in the models abstracts from small-sample issues, because this seems ideal for revealing the theoretical mechanisms of these models. As shown in Online Appendix I.1, however, the lessons drawn below are robust to a Monte Carlo exercise that accounts for sampling uncertainty.

Figure 6: The MBC Shock in the Data and the Models



of output and consumption to the various factors, as well as in the comparison of the consumption factor to the other factors.⁴³

We can offer a quantitative measure of these differences by constructing a metric of the interchangeability of factors in the data and in each of the models. Let $Z_{v,k}^f$ denote the impulse response function of variable $v \in V$ to factor $f \in F$, where $k \geq 0$ indexes the horizon, V is the set of the four key macroeconomic quantities (output, hours, consumption, and investment), and F is the set of the corresponding four factors. Next, let $\bar{Z}_{v,k} \equiv \frac{1}{4} \sum_{f \in F} Z_{v,k}^f$ and consider the following object:

$$D_v = \frac{1}{4} \sum_{f \in F} \sqrt{\sum_{k=0}^{20} (Z_{v,k}^f - \bar{Z}_{v,k})^2}$$

This is a measure of the dispersion of the IRFs of variable v across the factors. The closer D_v is to zero, the greater the degree of interchangeability. Conversely, a large value for D_v indicates low interchangeability vis-a-vis that particular variable. Finally, let $\bar{D} \equiv \frac{1}{4} \sum_{v \in V} D_v$. This gives a metric of how interchangeable the factors are over all the variables of interest.

⁴³Another noticeable feature is the magnitude of the responses, which are roughly twice as large as in JPT than the corresponding ones in either the data or ACD. This is because the original estimation of JPT, which is based on growth rates, produces excess volatility in the levels. As can be seen in Figure 19 in Online Appendix I.2, re-estimating JPT on levels, and in the same way as in ACD, fixes this excess-volatility problem but does not overcome the interchangeability challenge. Finally, the response of inflation appears to be much more sluggish in the data than in JPT, despite the inclusion of the hybrid versions of the price and wage Philips curves. This seems interesting, although it may not be directly related to the main point we wish to make here regarding the interchangeability of factors.

Table 10 reports the results of these calculations for the data and the two models (first row for the data, second row for JPT, third row for ACD). In each case, we report both the variable-specific metrics D_v (columns named “Y” through “h”) and the average metric \bar{D} (column named “Average”). It is evident that ACD produces nearly the same interchangeability as that observed in the data, while JPT produces much less.

Table 10: Interchangeability of Factors

	<i>Y</i>	<i>C</i>	<i>I</i>	<i>h</i>	Average
Data	0.47	0.52	1.28	0.28	0.64
JPT	2.90	2.21	6.29	1.35	3.19
ACD	0.56	0.49	1.61	0.30	0.74

Note: This table reports the distance of factors, measured in the way described in the main text. A number closer to zero indicates a larger degree of interchangeability.

We now shed light on this result and the mechanics of the models by doing a decomposition of their factors in terms of the underlying theoretical shocks.

Let us first consider JPT. In this model, the four macroeconomic quantities, and hence also the factors that target them, are driven by different mixtures of three distinct theoretical shocks: the investment-specific shock, the discount-factor shock, and the technology shock. As evident in the top panel of Figure 7, none of these shocks looks like the MBC shock in the data. In particular, both the investment-specific shock and consumption-specific shock induce *negative* co-movement between investment and consumption.⁴⁴ And because each of these shocks contribute differentially to the model’s factors, the latter are less interchangeable than the empirical counterparts.⁴⁵

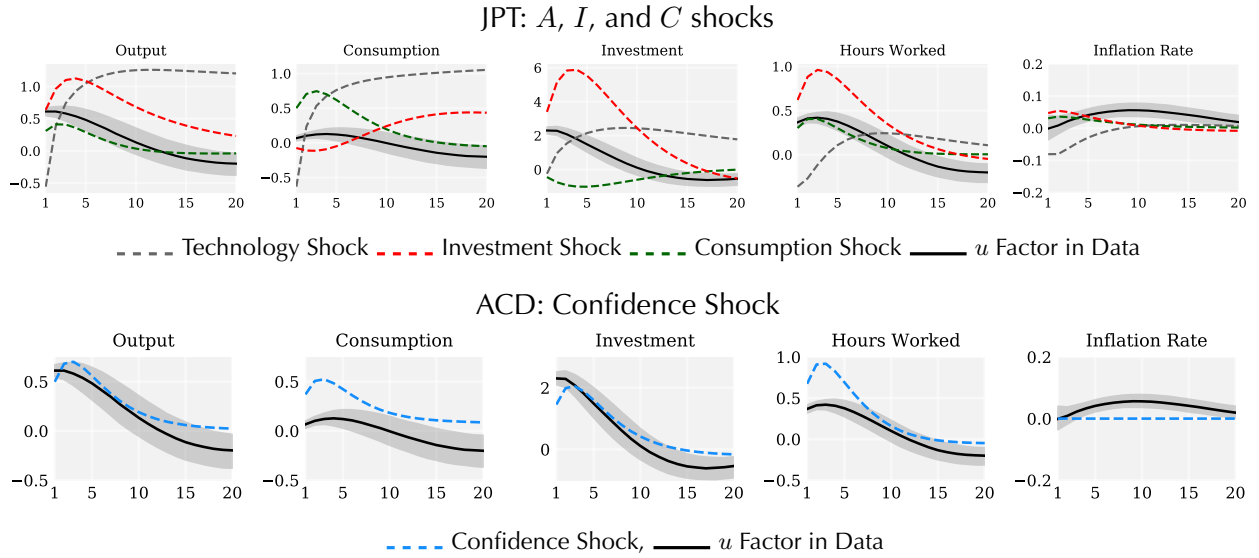
Consider next ACD. In this model, all variables are driven, to a large extent, by the same shock, the confidence shock. As explained in more detail in Angeletos, Collard, and Dellas (2018), this shock represents a shift in higher-order beliefs, or the expectations of the behavior of other firms and consumers, and helps capture waves of optimism and pessimism about the short-term economic outlook without commensurate shifts in the expectations of the long run. What is key for the present purposes is the observation, evident in the bottom panel of Figure 7, that this shock is quite similar to the MBC shock in the data, in terms of co-movements and relative volatilities. This helps explain why the factors in ACD are almost as interchangeable as those in the data.⁴⁶

⁴⁴This negative co-movement originates in the problem first highlighted by Barro and King (1984) and would have been even sharper if were not for the assumptions of time-non-separable preferences, sticky prices, and a monetary policy that induces an expansion relative to flexible prices.

⁴⁵Although the anatomy of JPT offered here is new, the basic property that the investment-specific shock produces negative co-movement between consumption and investment is known. Most of the existing attempts to fix this property maintain the reliance on non-separable preferences, sticky prices, and a monetary policy that induces an expansion relative to flexible-price outcomes (Furlanetto, Natvik, and Seneca, 2013; Furlanetto and Seneca, 2014; Ascari, Phaneuf, and Sims, 2016). An evaluation of the relative merits of these works vis-a-vis ACD, whose good comovement properties do not rely on any of these features, or any other member of the literature cited in footnote 3 is clearly beyond the scope of this paper.

⁴⁶Clearly, a bare-bones version of ACD, which shuts down all shocks except the confidence shock, achieves perfect

Figure 7: MBC Shock in Data vs Key Theoretical Shocks in JPT and ACD



These lessons are robust to two additional exercises, which are reported in Online Appendix I.2. In the first, we re-estimate JPT with same frequency-domain method as that used in the estimation of ACD. In the second exercise, we re-estimate both JPT and ACD on the basis of our anatomy, namely by minimizing the distance of each model from the data in terms of the impulse responses of the output, consumption, investment, and hours to the four factors that target the same quantities. Both exercises help JPT produce more interchangeability, but the model still falls far short of that found in the data as well as of that produced by the ACD model. That said, the goal of these exercises is not to argue that ACD is superior to JPT, but rather to illustrate the probing power of our empirical method.

In the same vein, we have applied our method to another important DSGE model, that of Christiano, Motto, and Rostagno (2014), henceforth CMR. This model is on the forefront of a new strand of the DSGE literature that pays close attention to the real-financial nexus. Its main differences from the model used in Christiano, Eichenbaum, and Evans (2005) and Justiniano, Primiceri, and Tambalotti (2010) are the following three. First, it includes a financial friction that constrains investment, the latter been broadly defined to include consumer durables. Second, it contains a new structural shock (“risk shock”) that determines the severity of the financial friction.⁴⁷ And third, it uses financial variables, most notably the credit spread between the gross nominal interest rate on debt and the risk free rate and the level of credit to such firms in the estimation and validation of the model.

The anatomy of this model involves not only the behavior of the macroeconomic quantities we have focused on so far, but also that of the new, financial variables. We have thus extended our anatomy of the data in Online Appendix G.3 to include information about these variables.⁴⁸

interchangeability without a big sacrifice in terms of matching the MBC shock in the data. This is not true for any single-shock restriction of

⁴⁷To be precise, this shock comes in nine flavors, depending on whether it hits the idiosyncratic volatility of firm returns with a lag of 0, 1, 2, ..., 8 quarters.

⁴⁸This is done in Online Appendix G.3 using three complementary VARs. The first one is obtained by adding only the credit spread to our baseline VAR. This allows us to keep the original sample size and corresponds to what is reported

Figure 8: Comparing Business-Cycle Factors

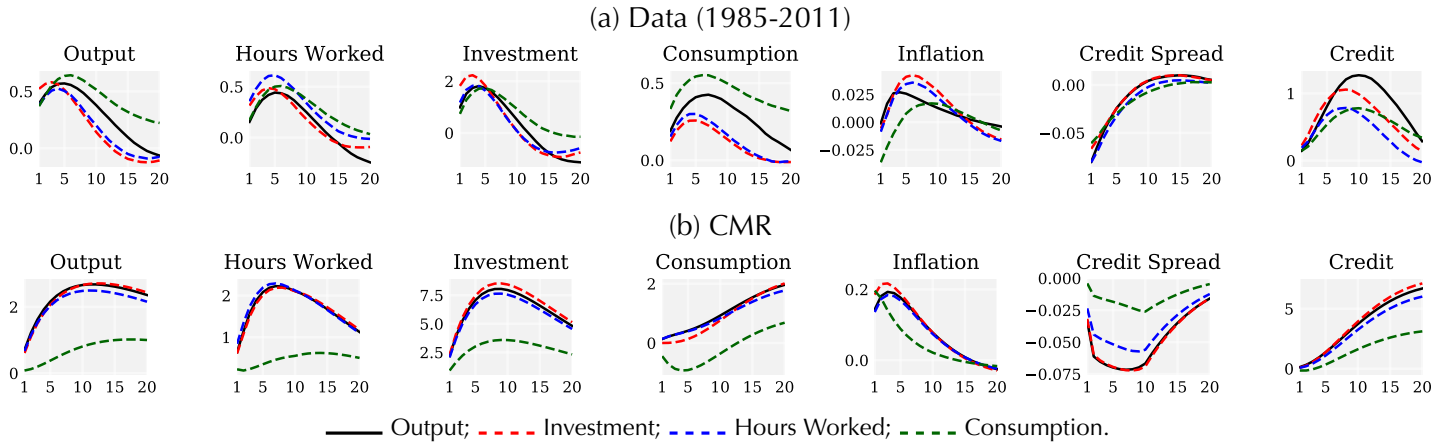


Figure 8 conducts a similar exercise as Figure 6. The top panel reports the IRFs of a few key variables to the output, hours, investment and consumption factors. The bottom panel reports the corresponding objects in the model. The only changes are the use of CMR instead of JPT or ACD; the focus on the sub-sample used in the estimation of that model;⁴⁹ and the addition of the impulse responses of the credit spread and the level of credit.

The following four patterns emerge. First, CMR improves upon JPT in terms of featuring more interchangeability between the output, hours, and investment factors—actually too much of it. Second, CMR does worse than JPT in terms of missing the business-cycle properties of consumption. This is evident both in the response of consumption to the aforementioned factors and in the response of all variables to the consumption factor. Third, CRM produces too much volatility and persistence compared to the data. Fourth, despite its use of a very flat Philips curve and very sticky wages, CMR produces a much steeper relation between inflation and real economic activity than that seen in the data, underscoring its reliance on nominal rigidity. Last but not least, the model fails to capture the dynamics of the response of the credit spread to all of these factors: while in the data the credit spread appears to lead the MBC shock, in the sense that it peaks before the macroeconomic quantities, it does the opposite in the model.

Whether these patterns represent critical failures for the CMR’s ability to capture the propagation of business cycles or easily fixable weaknesses is an open question beyond the scope of our paper.⁵⁰

as row 10 in Tables 8 and 16–19. The second is obtained by adding all the four financial variables used in CMR. In this case, data limitations force a shorter sample, 1971Q1-2014Q4. The third is obtained by restricting the second VAR to 1985Q1-2010Q4, which is the sample period used in the original estimation of CMR. The three VARs produce similar results, underscoring the robustness not only of our main findings but also of the additional findings reported in Figure 8 regarding the real-financial nexus.

⁴⁹That is, the empirical IRFs are obtained by using the last of the three VARs mentioned in footnote 48 above. Similarly to what we did in the case of JPT and ACD, this ensures that the model is evaluate on the basis of the period used in its estimation. But as already mentioned, the empirical patterns themselves are robust to the longer period spanned by our baseline specification.

⁵⁰The excessive persistence appears to be the product of the model’s reliance on an unusually high adjustment cost for investment as well as on very persistent shocks. The property that the business cycle leads, rather than lags, the credit spread

As already noted, the main goal of all the exercises conducted in this section is to illustrate the probing power of our method in the context of state-of-the-art, medium-scale models.

7 Conclusion

We have proposed a new strategy for dissecting macroeconomic time series and have used its findings to guide macroeconomic theory. The strategy involves employing a VAR to construct a variety of reduced-form shocks, each of which maximizes the volatility of a particular individual variable at particular frequencies. The constructed shocks, which may or may not have direct theoretical counterparts, represent a rich set of one-dimensional cuts of the data, or conditional comovements, which we call the anatomy of the data.

Prominent among the shocks constructed are those that target the main macroeconomic quantities (unemployment, output, hours worked, investment and consumption) at the business-cycle frequencies. The near interchangeability of these objects in terms of IRFs motivates the concept of the MBC shock: we use this term to refer to the dynamic comovement patterns that are common to all these cuts of the data. These include a strong, positive, and transient comovement between the aforementioned quantities; little relation with both inflation and TFP at any horizon; and a disconnect between the short run and the long run.

We have argued that these patterns speak against theories that seek to attribute the bulk of the business cycle to any of the following forces: technology shocks; financial, uncertainty and other shocks that matter primarily by affecting the concurrent level of aggregate TFP; news about medium-to long-run productivity prospects of the economy; and demand shocks that give rise to procyclical movements in inflation. In contrast, models that contain a non-inflationary, demand shock as the main driver of the business cycle seem a priori consistent with this evidence.

This conclusion is based on the premise that the bulk of the business cycle can be attributed to a single shock or, equivalently, by a dominant propagation mechanism shared by multiple shocks. But even if this is not the case, the rich set of the dynamic comovement patterns that comprise our MBC shock, or more generally, our anatomy, serve as useful yardstick for model evaluation: models of any size and complexity can be readily confronted with these patterns. Prominent members of the DSGE literature, such as Justiniano, Primiceri, and Tambalotti (2010) have difficulty passing this test, despite the inclusion of a dominant business-cycle driver and a flat Philips curve.

All in all, our findings suggest that the flexible-price core of the mainstream, New Keynesian framework is problematic to start with, in the sense that this core is itself unable to accommodate the kind of non-inflationary demand shock we have documented in the data. Encouragingly, there is now a growing literature that attempts to accommodate demand-driven business cycles even without nominal rigidities and Phillips curves; see the references in footnote 3. We hope that the characterization of the data performed in the present paper will stimulate further research in this direction and help guide macroeconomic theory.

appears to be driven by the model's reliance on a number of news shocks, which have a relatively more pronounced and front-loaded effect on investment, hours and output than on the credit spread. We do not know how changing these features would impact on the empirical performance of the model along other dimensions and also on the structural interpretation offered to the data.

Finally, while our anatomy is quite comprehensive, it could be further enriched by more refined cuts of the data. Consider, in particular, the following enrichment. For each variable $X \in \{u, Y, h, I, C\}$, first filter out the effect of the shock that accounts for most of the business-cycle volatility in that variable (i.e., the kind of shock we have focused in this paper) and then construct the shock that accounts for most of the *residual* volatility in the same variable. This procedure gives the “second largest shocks” corresponding to the main macroeconomic quantities. We find that these shocks display a similar interchangeability property as the one that motivated our concept of the MBC shock. They can thus be thought of as different facets of the same, secondary, business-cycle shock (subject, of course, to the same kind of interpretation challenges as those discussed in Section 5). Online Appendix J details the empirical profile of this shock, compares it to that of the MBC shock, and illustrates how this could serve as an additional model-validation tool.

References

- Angeletos, George-Marios, Fabrice Collard, and Harris Dellas. 2015. "Quantifying Confidence." *NBER Working Paper No. 20807* .
- . 2018. "Quantifying Confidence." *Econometrica* 86 (5):1689–1726.
- Angeletos, George-Marios and Jennifer La'O. 2010. "Noisy Business Cycles." In *NBER Macroeconomics Annual 2009, Volume 24*. University of Chicago Press, 319–378.
- . 2013. "Sentiments." *Econometrica* 81 (2):739–779.
- Arellano, Cristina, Yan Bai, and Patrick Kehoe. 2018. "Financial Frictions and Fluctuations in Volatility." *Journal of Political Economy* Forthcoming.
- Ascari, Guido, Louis Phaneuf, and Eric Sims. 2016. "Business Cycles, Investment Shocks, and the "Barro-King" Curse." *NBER Working Paper* (22941).
- Bai, Yan, José-Victor Ríos-Rull, and Kjetil Storesletten. 2017. "Demand Shocks as Productivity Shocks." .
- Barro, Robert J. and Robert G. King. 1984. "Time-Separable Preferences and Intertemporal Substitution Models of Business Cycles." *The Quarterly Journal of Economics* 99 (4):817–839.
- Barsky, Robert B. and Eric R. Sims. 2011. "News Shocks and Business Cycles." *Journal of Monetary Economics* 58 (3):273 – 289.
- Beaudry, Paul, Dana Galizia, and Franck Portier. 2018. "Reconciling Hayek's and Keynes' Views of Recessions." *The Review of Economic Studies* 85 (1):119–156.
- Beaudry, Paul and Franck Portier. 2006. "Stock Prices, News, and Economic Fluctuations." *American Economic Review* 96 (4).
- . 2018. "Real Keynesian Models and Sticky Prices." miméo, UBC/TSE.
- Benhabib, Jess and Roger E. A. Farmer. 1994. "Indeterminacy and Increasing Returns." *Journal of Economic Theory* 63 (1):19–41.
- Benhabib, Jess, Pengfei Wang, and Yi Wen. 2015. "Sentiments and Aggregate Demand Fluctuations." *Econometrica* 83 (2):549–585.
- Blanchard, Olivier J. and Danny Quah. 1989. "The Dynamic Effects Of Aggregate Demand And Supply Disturbances." *American Economic Review* 79 (4):655–673.
- Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J. Terry. 2018. "Really Uncertain Business Cycles." *Econometrica* forthcoming.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy* 113 (1):1–45.

- Christiano, Lawrence J., Roberto Motto, and Massimo Rostagno. 2014. "Risk Shocks." *American Economic Review* 104 (1):27–65.
- Christiano, Lawrence J. and Robert J. Vigfusson. 2003. "Maximum Likelihood in the Frequency Domain: the Importance of Time-to-Plan." *Journal of Monetary Economics* 50 (4):789–815.
- Cochrane, John H. 1994. "Shocks." *Carnegie-Rochester Conference Series on Public Policy* 41:295 – 364.
- Debortoli, Davide, Jordi Galí, and Luca Gambetti. 2019. "Business Cycles, Investment Shocks, and the "Barro-King" Curse." *NBER Working Paper* (25820).
- Diamond, Peter A. 1982. "Aggregate Demand Management in Search Equilibrium." *The Journal of Political Economy* :881–894.
- Eusepi, Stefano and Bruce Preston. 2015. "Consumption Heterogeneity, Employment Dynamics and Macroeconomic Co-movement." *Journal of Monetary Economics* 71 (C):13–32.
- Faust, Jon. 1998. "The robustness of identified VAR conclusions about money." *Carnegie-Rochester Conference Series on Public Policy* 49:207 – 244.
- Fernald, John G. 2012. "A Quarterly, Utilization-adjusted Series on Total Factor Productivity." Working Paper Series 2012-19, Federal Reserve Bank of San Francisco.
- Forni, Mario, Luca Gambetti, and Luca Sala. 2019. "Structural VARs and Non-invertible Macroeconomic Models." *Journal of Applied Econometrics* 34 (2):221–246.
- Francis, Neville, Michael T. Owyang, Jennifer E. Roush, and Riccardo DiCecio. 2014. "A Flexible Finite-Horizon Alternative to Long-Run Restrictions with an Application to Technology Shocks." *The Review of Economics and Statistics* 96 (4):638–647.
- Furlanetto, Francesco, Gisle J. Natvik, and Martin Seneca. 2013. "Investment shocks and macroeconomic co-movement." *Journal of Macroeconomics* 37:208 – 216.
- Furlanetto, Francesco and Martin Seneca. 2014. "Investment shocks and consumption." *European Economic Review* 66:111 – 126.
- Galí, Jordi. 1999. "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review* 89 (1):249–271.
- Galí, Jordi and Mark Gertler. 1999. "Inflation Dynamics: A Structural Econometric Analysis." *Journal of Monetary Economics* 44 (2):195–222.
- Golosov, Mikhail and Guido Menzio. 2015. "Agency Business Cycles." Working Paper 21743, National Bureau of Economic Research.
- Guesnerie, Roger and Michael Woodford. 1993. "Endogenous Fluctuations." In *Advances in Economic Theory*, vol. 2, edited by Jean-Jacques Laffont. Cambridge University Press, 289–412.

- Huo, Zen and Naoki Takayama. 2015. "Higher Order Beliefs, Confidence, and Business Cycles." miméo, Yale University.
- Ilut, Cosmin and Hikaru Saijo. 2018. "Learning, Confidence, and Business Cycles." miméo, Duke University.
- Jaimovich, Nir and Sergio Rebelo. 2009. "Can News about the Future Drive the Business Cycle?" *American Economic Review* 99 (4):1097–1118.
- Johansen, S. and K. Juselius. 1990. "Maximum Likelihood Estimation and Inference on Cointegration with Applications to Demand for Money." *Oxford Bulletin of Economics and Statistics* 52 (2):169–210.
- Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti. 2010. "Investment Shocks and Business Cycles." *Journal of Monetary Economics* 57 (2):132–145.
- Koop, Gary. 2003. *Bayesian Econometrics*. John Wiley & Sons.
- Kydland, Finn E and Edward C Prescott. 1982. "Time to Build and Aggregate Fluctuations." *Econometrica* 50 (6):1345–1370.
- Lucas, Robert E. Jr. 1975. "An Equilibrium Model of the Business Cycle." *Journal of Political Economy* 83 (6):1113–44.
- . 1977. "Understanding Business Cycles." *Carnegie-Rochester Conference Series on Public Policy* 5:7 – 29.
- Mavroeidis, Sophocles, Mikkel Plagborg-Møller, and James H. Stock. 2014. "Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve." *Journal of Economic Literature* 52 (1):124–88.
- Ramey, Valerie A. 2016. *Macroeconomic Shocks and Their Propagation, Handbook of Macroeconomics*, vol. 2, chap. 2. Elsevier, 71–162.
- Sala, Luca. 2015. "DSGE Models in the Frequency Domains." *Journal of Applied Econometrics* 30 (2):219–240.
- Smets, Frank and Rafael Wouters. 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review* 97 (3):586–606.
- Stock, James H. and Mark W. Watson. 1999. "Business Cycle Fluctuations in US Macroeconomic Time Series." In *Handbook of Macroeconomics, Handbook of Macroeconomics*, vol. 1, edited by J. B. Taylor and M. Woodford, chap. 1. Elsevier, 3–64.
- Uhlig, Harald. 2003. "What Moves Real GNP?" mimeo.

APPENDICES

A The Data

The data is from the Federal Reserve Economic Database (FRED). TFP corresponds to the TFP time series corrected for utilization produced by Fernald (2012) (downloaded 2016). Tables 11 and 12 describe the original data and the transformations used in our VARs. Table 13 reports the raw (unconditional) correlations over the business-cycle frequencies.

Table 11: Description of Data

Data	Mnemonic	Freq.	Transform
Real gross domestic product per capita	A939RX0Q048SBEA	Q	–
Gross Domestic Product	GDP	Q	–
Gross Domestic Product: Implicit Price Deflator	GDPDEF	Q	–
Personal Consumption Expenditures: Nondurable Goods	PCND	Q	–
Personal Consumption Expenditures: Services	PCESV	Q	–
Personal Consumption Expenditures: Goods	PCDG	Q	–
Gross Private Domestic Investment	GPDI	Q	–
Nonfarm Business Sector: Real Output Per Hour of All Persons	OPHNFB	Q	–
Nonfarm Business Sector: Labor Share	PRS85006173	Q	–
Nonfarm Business Sector: Average Weekly Hours	PRS85006023	Q	–
Civilian Noninstitutional Population	CNP160V	M	EoP
Civilian Unemployment Rate	UNRATE	M	Ave
Effective Federal Funds Rate	FEDFUNDS	M	Ave
Total Factor Productivity (Growth rate)	DTFPu	Q	–

Note: Q: Quarterly, M: Monthly, EoP: end of period, Ave: quarterly average.

Table 12: Variables in the VARs

Real GDP per capital	$Y = \log(A939RX0Q048SBEA)$
Real consumption per capita	$C = \log((PCND + PCESV) * A939RX0Q048SBEA / GDP)$
Real investment per capita	$I = \log((PCDG + GPDI) * A939RX0Q048SBEA / GDP)$
Hours worked	$H = \log(PRS85006023 * CE160V / CNP160V)$
Inflation Rate	$\pi = \log(GDPDEF / GDPDEF(-1))$
Interest Rate	$R = FEDFUNDS / 400$
Productivity (NFB)	$YSHnfb = OPHNFB$
Labor Share	$wh/y = \log(PRS85006173)$
TFP	$TFP = \log(\text{cumulative sum}(DTFPu / 400))$

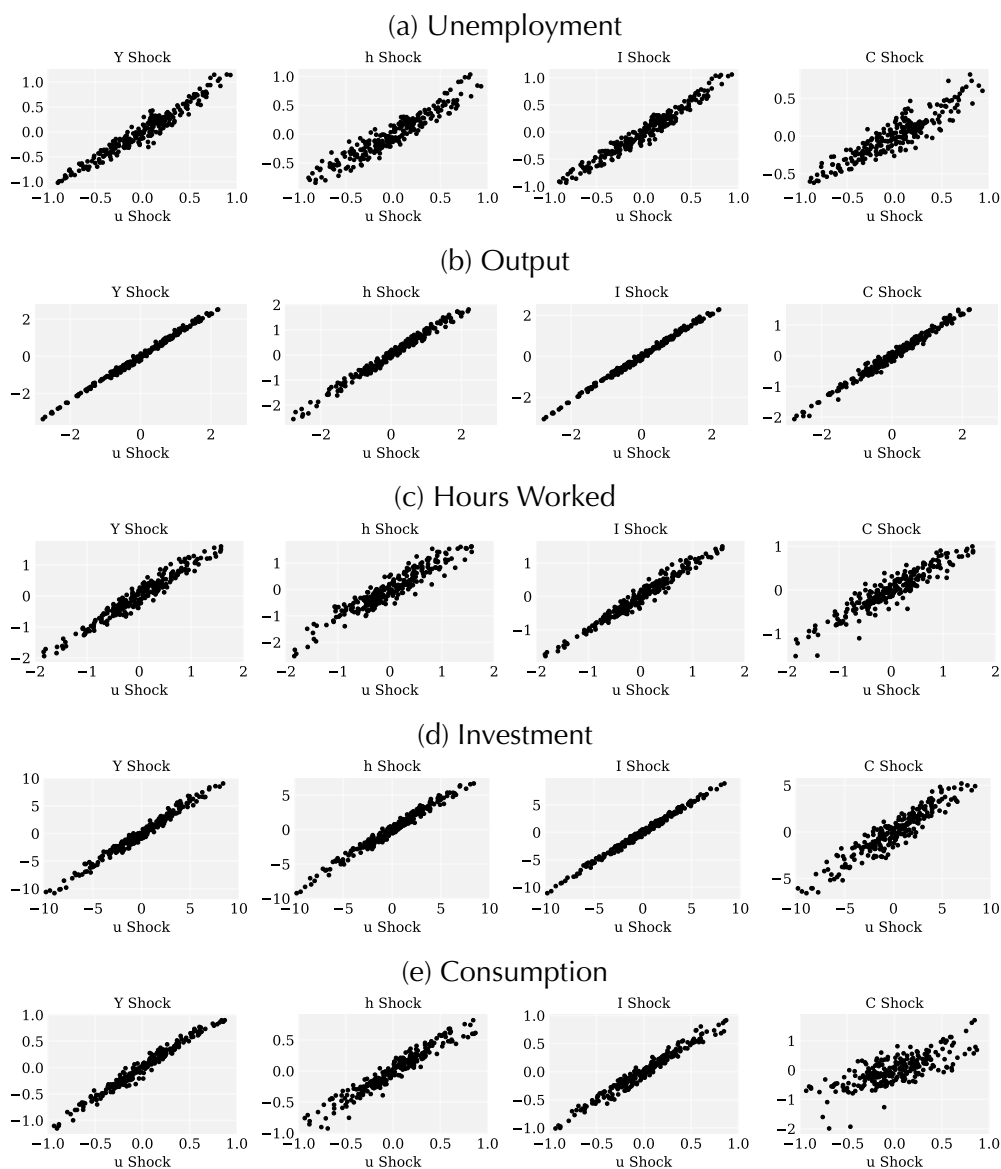
Table 13: Correlations (Bandpass filtered, 6-32 Quarters)

	Y_t	C_t	I_t	h_t	u_t	TFP_t	$(Y/h)_t$	$(Wh/Y)_t$	π_t	R_t
Y_t	1.00	0.84	0.95	0.89	-0.88	-0.19	0.47	-0.15	0.21	0.40
C_t	0.84	1.00	0.76	0.82	-0.78	-0.28	0.24	0.05	0.31	0.42
I_t	0.95	0.76	1.00	0.89	-0.85	-0.24	0.44	-0.18	0.13	0.33
h_t	0.89	0.82	0.89	1.00	-0.93	-0.46	0.11	0.06	0.29	0.47
u_t	-0.88	-0.78	-0.85	-0.93	1.00	0.41	-0.06	-0.16	-0.37	-0.59
TFP_t	-0.19	-0.28	-0.24	-0.46	0.41	1.00	0.45	-0.23	-0.27	-0.34
$(Y/h)_t$	0.47	0.24	0.44	0.11	-0.06	0.45	1.00	-0.56	-0.30	-0.31
$(Wh/Y)_t$	-0.15	0.05	-0.18	0.06	-0.16	-0.23	-0.56	1.00	0.31	0.23
π_t	0.21	0.31	0.13	0.29	-0.37	-0.27	-0.30	0.31	1.00	0.72
R_t	0.40	0.42	0.33	0.47	-0.59	-0.34	-0.31	0.23	0.72	1.00

B Interchangeability in the Time Series

In the main text we emphasized the interchangeability of the various facets of the MBC shock in terms of IRFs. Figure 9 shows that a similar interchangeability property is present in terms of the time series generated by the reduced-form shocks. Each row in this figure reports, for each one of the key macroeconomic quantities, the scatterplot of that variable as predicted by the Y , I , C , and h shocks against its value as predicted by the unemployment shock. Table 3 in the main text summarizes the information contained in this figure in terms of correlations.

Figure 9: The Various Facets of the MBC Shock, Scatterplots



C Application to New Shocks

In this Appendix, we use our method to identify news shocks and examine how their properties, in particular their contribution to business cycles, vary with the size of the VAR used to identify the shocks. This serves two purposes. It sheds light on the source of the difference reported in the main text between our findings and those of Beaudry and Portier (2006). And it provides yet another example of the usefulness of our method outside the realm of one-shock representations of the business cycle, in particular, in the context of semi-structural explorations.

The exercise conducted here is based on the premise that the vast majority, if not all, of the TFP fluctuations at all frequencies can be accounted by two structural shocks: an unanticipated, permanent shock and a news shock. The former affects TFP both in the short and the long run, while the latter does not have an effect on impact.⁵¹

As explained in Section 5, the accommodation of these two structural shocks complicates the interpretation of the empirical MBC shock and in particular of its disconnect from the long run: this disconnect is consistent with models in which the two structural shocks under consideration have significant but offsetting effects on unemployment in the short run. Still, insofar as only these two shocks drive TFP, and regardless of how many other shocks may drive unemployment, we can identify the news shock and its business-cycle contribution as follows.

We first construct, via our method, the two empirical shocks that have the maximal contribution to the volatility of TFP in the long-run and the business-cycle frequencies ($80 - \infty$ and $6 - 32$ quarters, respectively). Denote these by s_t^1 and s_t^2 , respectively. These shocks do not have a structural interpretation but are linear combinations of the two “true” structural shocks, the unanticipated technology shock, s_t^{tech} , and the news shock, s_t^{news} . The two sets of shocks are related as follows:

$$\begin{bmatrix} s_t^1 \\ s_t^2 \end{bmatrix} = A \begin{bmatrix} s_t^{tech} \\ s_t^{news} \end{bmatrix}$$

for some matrix A . As long as both s_t^1 and s_t^2 have a non-zero impact effect on TFP (which is true for all the specifications considered below), one can construct their unique (up to rescaling) linear combination that has a zero impact effect on TFP. This combination recovers the news shock.

We have implemented this identification strategy in our baseline VAR, as well as in several other smaller and larger VARs. We report results below for seven nested specifications, denoted as VAR₁ through VAR₇. The smallest one, VAR₁, contains only the main two variables of interest, TFP and unemployment. VAR₂ adds investment. VAR₃, adds GDP, consumption and hours, giving the “real core” of our baseline VAR. The latter is herein denoted by VAR₄; this contains all the 10 variables described in Section 2. VAR₅ adds the SP500 index. VAR₆ adds capacity utilization. VAR₇ adds the credit spread.

⁵¹One may object to the assumption of only two TFP shocks, on the basis, for instance, that the “right” model features multiple news shocks, each one corresponding to different horizons at which TFP is expected to change. But this is a slippery road that ultimately leads one to give up hope on “a-theoretic” endeavors and, instead, commit to a particular, fully-specified model. Clearly, each approach has its strengths and limitations. We follow the one approach here and the other in Section 5.

In all of the VARs, the two empirical shocks, s_t^1 and s_t^2 , together account for over 95% of the volatility of TFP at the long-run frequencies and for over 85% of that at the business-cycle frequencies. In our baseline specification, in particular, these numbers are 99% and 92%, respectively. In this regard, our two-shock representation of TFP works well. Moreover, the effect of the identified news shock on the dynamics of TFP is quite similar across the VARs: see the left panel of Figure 10. Such robustness, however, is absent in the relationship between news shocks and unemployment fluctuations; see the right panel of Figure 10. In particular, the news shock switches from being strongly expansionary in the smallest VAR to being slightly contractionary in the largest VAR.

Figure 10: IRF of TFP and Unemployment to News Shock

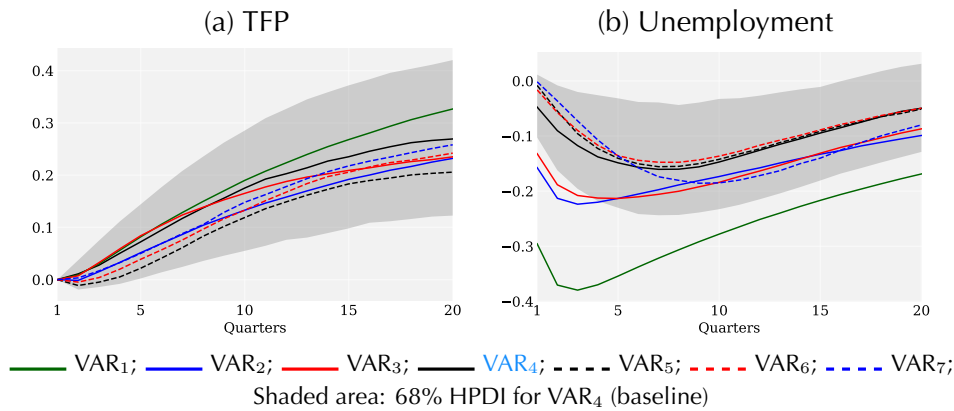
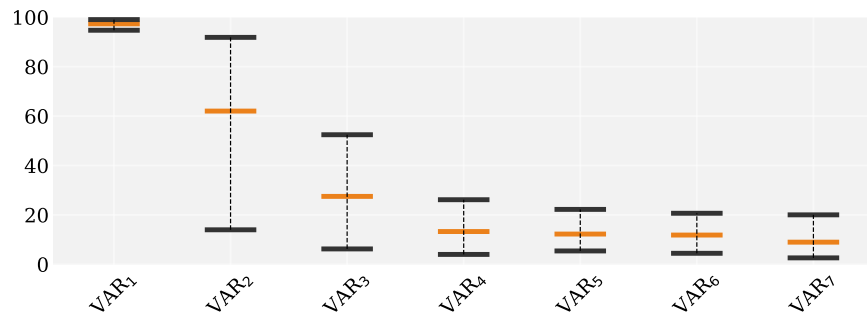


Figure 11 presents this sensitivity in terms of the contribution of the identified news shock to the volatility of unemployment at the business-cycle frequencies. On the horizontal axis, we vary the size of the VAR used in the construction of s_t^1 and s_t^2 and, thereby, of the news shock: as we move from left to right, we progressively add more data and, accordingly, increase the size of the VAR from 2 variables to a total of 13.

Figure 11: Variance Contribution of News Shock to Unemployment



Note: Contribution of news shock to unemployment at business-cycle frequencies. Red line gives median, upper and lower black lines give 68% HPDI. $VAR_1 = \{u, TFP\}$, $VAR_2 = VAR_1 \cup \{I\}$, $VAR_3 = VAR_2 \cup \{Y, C, h\}$, $VAR_4 =$ Baseline VAR, $VAR_5 = VAR_4 \cup \{SP500\}$, $VAR_6 = VAR_5 \cup \{\text{utilization}\}$, $VAR_7 = VAR_6 \cup \{\text{credit spread}\}$.

The figure speaks for itself: as more information (in the form of the additional variables) is incor-

porated, the estimated contribution of the news shock declines dramatically, stabilizing at around 11% in the last four specifications. In our baseline specification, the number is 13%.

Due to the well-known potential fragility of results from small VARs (Forni, Gambetti, and Sala, 2019), we trust more the results from the medium and larger ones, specially because size seizes to matter after a certain size. Larger VARs contain more information, while smaller ones may mechanically attribute a larger share of the business cycle to the news shock.

To illustrate the latter point, consider VAR₁. In this specification, the news shock accounts for 97% of the short-run fluctuations in unemployment. Why? In a two variables-two shocks specification, s_t^{tech} and s_t^{news} must together account for all of the fluctuations in unemployment. Due to the assumption that s_t^{tech} is the only shock that has an immediate, impact effect on TFP, s_t^{tech} is closely associated with actual TFP in the short run. But as we have established, TFP is nearly orthogonal to unemployment at the business-cycle frequencies (and beyond). It then follows that s_t^{tech} can account for only a trivial fraction of the unemployment fluctuations—which leaves s_t^{news} as the only shock to explain unemployment fluctuations. In short, this VAR mechanically attributes a large fraction of the business cycle to the news shock, simply because the only other allowed shock is a “dead horse” to start with.

As we move to larger VARs, we add more data but also more shocks that can contribute to the fluctuations in unemployment. So the role of news is bound to wither. Figure 11 shows that the decline is precipitous at first, but stabilizes once we reach the baseline specification.

This helps shed light on one of the reasons why our results differ from those in Beaudry and Portier: we use larger VARs than they do. Another part of the difference comes from using different identifying assumptions. In this context, note that our identification strategy remains the same as we move from smaller to larger VARs. The same is true for the strategy employed in Barsky and Sims, which, reassuringly, leads to a similar conclusion as ours.⁵² By contrast, the one employed in Beaudry and Portier requires the introduction of progressively more delicate exclusion restrictions as more variables are added.

The exercise conducted here also serves another important purpose. Namely, it helps showcase the usefulness of our approach in the realm of multi-shock models without a need for the explicit intermediation of a particular, fully-specified model. The key is to drop the exclusive focus on the MBC shock and include other features of the anatomy—here for instance the shocks that target TFP in the short and the long run—and to utilize the cross-equation restrictions associated with them. As shown in Section 6, the same procedure also proves very effective in the context of fully-structural endeavors.

⁵²Barsky and Sims emphasize that their identified news shocks is contractionary, but what we take as a more robust conclusion from the combination of their explorations and ours is the small contribution of TFP news to the business cycle.

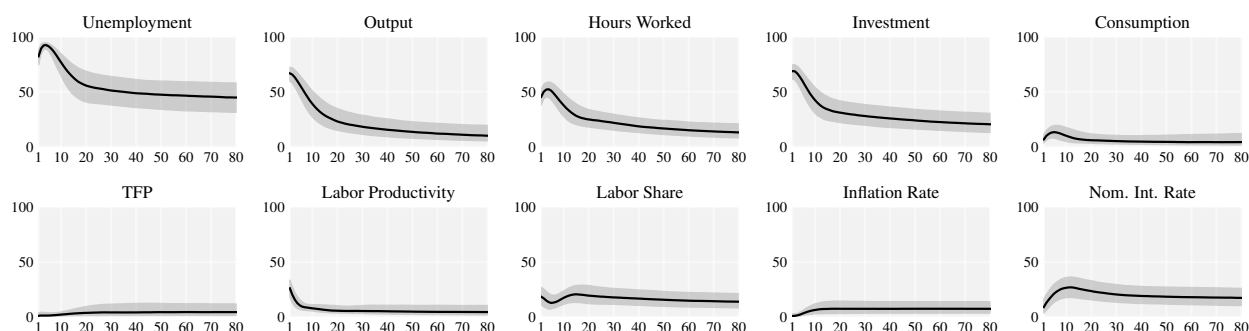
ONLINE APPENDICES

D Variance Contributions on the Time Domain

Figure 12 complements Table 1 in the main text by reporting the contribution of the identified MBC shock to the FEV of the variables at different horizons. To avoid any confusion, let us emphasize that the shock is still identified in the frequency domain, by targeting the volatility of unemployment over the band of the business-cycle frequencies. The time domain is used only in the calculation of variance contributions.

The picture that emerges is fully consistent with that painted in the main text: the identified shock explain the bulk of the short-run variation in the key macroeconomic quantities, and has a negligible footprint to TFP and inflation at all horizons. The only subtlety worth noting here is that “short run” in the time domain maps to a horizon of about 4 to 8 quarters. This is evident not only in the FEV contributions reported here but also in the IRFs shown in the main text, which pick within the first few quarters. And it anticipates the choice of the horizon targeted in a variant, time-domain identification considered in Online Appendix E.

Figure 12: Variance Contributions at Different Horizons



Note: Variance contributions of the MBC shock in the time domain. Horizontal axis: time horizon in quarters. Shaded area : 68% HPDI.

E Identification in the Frequency vs Time Domain

In this Appendix we illustrate the robustness of our findings to employing time-domain instead of frequency-domain methods. In particular, we identify the relevant shocks by maximizing their contribution to the FEV of the corresponding variable at horizons of 1, 4, and 8 quarters (which are the time horizons typically associated with the business cycle). As is evident in Figure 13 and Table 14, this change does not affect our findings. To save space we only report the results for the shocks that target unemployment and GDP. But the same picture obtains for the other elements of our anatomy.

Figure 1 in the main text and Figure 12 in Online Appendix D paint a complementary picture in terms of the *time*-domain properties of the shock identified on the *frequency* domain: its IRFs and

the FEV contributions peak within 1 to 4 quarters. Together, these results clarify the following point: targeting the business-cycle frequencies in the frequency domain is essentially the same as targeting a horizon of a few (say 1, 4 or 8) quarters in the time domain.

Still, one can ask why to construct shocks by targeting the chosen frequency band rather than at the corresponding time horizon. We do so for three interrelated reasons. First, the mapping between the frequency band of interest and the appropriate time horizon is not *a priori* known. Second, it is a standard convention in the literature to equate the business cycle with the chosen band of frequencies; see, e.g., Stock and Watson (1999) for the use of the 6-32 bandpass filter in the construction of stylized business-cycle moments and its near coincidence with the HP filter. By contrast, there is less agreement on whether the volatility induced by a shock at a particular horizon is or is not part of the business cycle. Finally, our approach to the data is more directly comparable to a DSGE practice that estimates business-cycle models on the frequency domain (Christiano and Vigfusson, 2003; Sala, 2015). This in turn seems a good practice, because such models are not meant to account for the productivity slowdown, structural reforms, demographic changes and other slow-moving forces that may “distort” the usual, time-domain estimation practice.

Figure 13: Frequency-Domain vs Time-Domain Identification (IRFs)

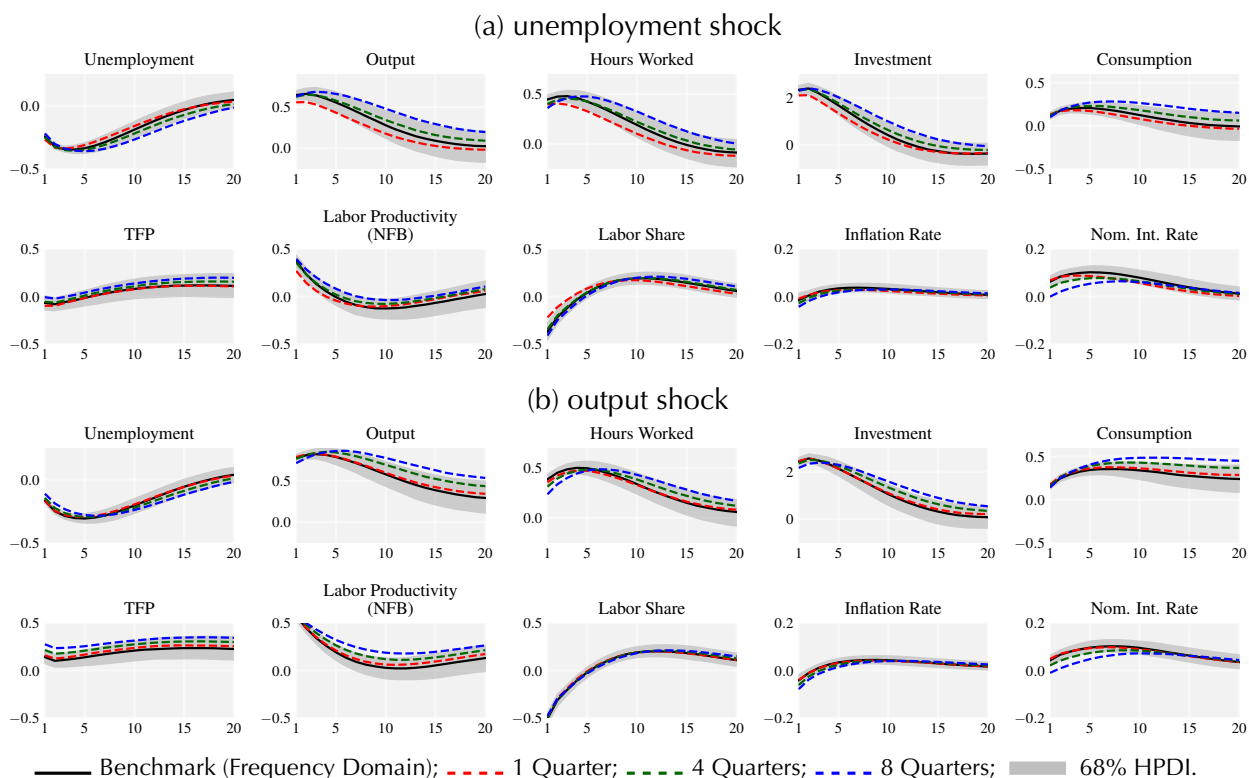


Table 14: Frequency-Domain vs Time-Domain Identification (Variance Contributions)

	u	Y	h	I	C	TFP	Y/h	Wh/Y	π	R
<i>Unemployment Shock</i>										
Benchmark	73.71 [66.80,79.94]	58.51 [50.65,65.07]	47.72 [40.77,54.45]	62.09 [54.09,68.46]	20.38 [13.61,27.53]	5.86 [2.44,10.96]	23.91 [17.27,31.22]	27.02 [18.39,35.93]	6.96 [3.24,12.28]	22.27 [14.22,30.97]
1 Qrt	66.07 [57.97,73.47]	40.79 [34.52,46.67]	34.15 [27.88,39.68]	46.32 [40.16,52.47]	15.79 [11.04,21.59]	6.07 [2.91,10.55]	13.03 [8.91,17.47]	14.21 [10.03,19.22]	3.83 [1.70, 7.22]	18.18 [12.68,23.37]
4 Qrts	71.18 [62.71,77.79]	53.14 [46.19,59.10]	41.02 [33.52,47.92]	57.29 [50.43,63.96]	19.43 [13.25,26.17]	5.14 [2.15, 9.59]	19.40 [14.14,25.08]	24.43 [18.27,30.77]	5.62 [2.59, 9.95]	12.80 [7.50,19.49]
8 Qrts	67.39 [58.03,75.33]	56.44 [48.28,63.11]	41.53 [32.86,50.02]	59.41 [51.33,67.17]	21.81 [14.52,29.32]	4.03 [1.62, 8.32]	22.41 [16.49,29.26]	31.49 [23.32,38.90]	7.41 [3.49,12.75]	8.73 [4.34,14.91]
<i>Output Shock</i>										
Benchmark	56.24 [48.94,61.93]	80.13 [72.80,86.44]	44.73 [37.36,51.68]	67.13 [60.72,72.82]	33.03 [25.04,40.44]	4.24 [1.76, 8.32]	41.31 [35.29,47.43]	40.20 [32.75,47.40]	10.47 [5.97,16.75]	16.89 [11.00,26.08]
1 Qrt	51.17 [44.04,57.04]	77.54 [69.45,84.83]	38.13 [31.31,45.26]	63.88 [56.57,70.14]	32.83 [25.25,40.11]	4.60 [2.13, 8.73]	38.34 [32.88,44.22]	36.75 [30.29,42.53]	8.90 [5.07,14.03]	15.88 [10.55,22.61]
4 Qrts	47.35 [39.89,53.93]	77.31 [68.65,84.43]	36.58 [29.28,44.34]	61.44 [54.05,68.39]	34.87 [27.06,42.04]	6.73 [3.28,11.52]	40.73 [34.40,46.94]	38.00 [31.73,44.45]	11.84 [7.51,17.55]	10.59 [6.13,16.46]
8 Qrts	40.15 [32.02,47.83]	71.69 [61.49,80.46]	32.35 [24.00,40.56]	54.78 [45.85,63.46]	35.75 [27.81,43.20]	9.66 [5.38,16.43]	40.72 [33.65,47.58]	36.28 [28.79,43.36]	15.20 [9.70,21.96]	8.14 [4.33,13.26]

Note: The two parts of the table correspond to different targeted variables, unemployment or GDP. In each part, the first row correspond to our benchmark, frequency-domain identification of the shock, while the other rows correspond to time-domain identification. In particular, three cases are reported, depending on whether the shock is constructed by maximizing its contribution to the FEV of the respective variable at horizons of 1, 4, or 8 quarters. The columns report the contributions of the thus-identified shocks to the business-cycle volatilities of all the variables.

F Long Run PCA

Table 4 in Section 3.3 reported the first principal component over the business-cycle frequencies (the band corresponding to 6 – 32 quarters). For completeness, Table 15 here reports the corresponding object over the long-run frequencies (the band corresponding to 80 – ∞ quarters). The picture that emerges corroborates the existence of a single unit-root force driving almost the entirety of the long-run fluctuations in TFP and the key macroeconomic quantities.⁵³

Table 15: First Principal Component, Long Term, 1955-2017

	u	Y	h	I	C	π	R	r	TFP	Y/h	w	wh/Y
Raw Data	10.43	99.93	64.93	98.11	99.66	6.20	6.97	2.44	98.33	99.32	99.74	73.89
VAR-Based	12.20	97.88	5.82	95.08	96.32	3.94	6.66	9.99	88.97	98.18	96.59	32.74
Normalized Data	10.38	99.18	62.59	95.57	99.83	9.85	10.33	3.52	96.69	98.96	98.72	78.28
VAR Normalized	29.44	90.64	17.49	86.44	89.46	11.89	20.37	19.00	88.36	89.67	94.16	49.18

G Robustness of Empirical Findings

In Section 4, we established the robustness of the empirical properties of the shock that targets unemployment across eleven specifications. In the first subsection of this appendix, we first show that the same robustness property characterizes the other shocks that form our anatomy. In the next two subsections, we expand on some additional findings from the two extended VARs that show up as rows 9 and 10 in these tables. In the last two subsections, we finally fill in a few details regarding the VECM specifications and measurement of the relative prices of investment.

G.1 Beyond the unemployment shock: other elements of the anatomy

Table 8 in the main text reported the variance contributions of the shock that targets unemployment across eleven specifications. Table 16 through Table 19 here repeat the exercise of a select subset of the other elements comprising our anatomy: the shocks that target GDP, hours, investment, and inflation. Although omitted here for the sake of saving space, the same robustness property is also present in terms of IRFs.

⁵³The cells in this table that appear in gray color correspond to the variables that, at least according to most theories, ought to be stationary, in which case the reported numbers are meaningless.

Table 16: The MBC Shock, Targeting Output, Variance Contributions (6-32 Quarters)

	u	Y	h	I	C	TFP	Y/h	Wh/Y	π	R
[1] Benchmark	56.24 [48.94,61.93]	80.13 [72.80,86.44]	44.73 [37.36,51.68]	67.13 [60.72,72.82]	33.03 [25.04,40.44]	4.24 [1.76, 8.32]	41.31 [35.29,47.43]	40.20 [32.75,47.40]	10.47 [5.97,16.75]	16.89 [11.00,26.08]
[2] 4 lags	56.48 [50.18,63.14]	79.38 [71.95,85.64]	44.56 [37.14,52.69]	67.35 [61.08,73.31]	33.20 [26.42,40.63]	5.49 [2.44,10.40]	40.56 [34.22,46.76]	41.06 [33.38,48.23]	11.35 [6.31,17.09]	17.71 [9.90,26.49]
[3] VECM(1)	51.21 [43.96,57.68]	62.37 [56.11,69.41]	43.05 [35.30,50.83]	54.74 [48.66,61.51]	44.17 [36.00,54.01]	9.71 [5.27,17.85]	30.54 [24.50,37.65]	35.49 [26.32,44.21]	9.37 [4.40,17.63]	21.55 [9.85,39.01]
[4] VECM(2)	52.31 [45.04,59.90]	68.59 [60.91,76.13]	43.52 [36.30,50.88]	55.54 [48.61,62.06]	56.07 [46.24,64.66]	7.65 [4.38,12.83]	33.22 [26.99,40.03]	37.57 [29.75,45.14]	9.14 [3.75,16.24]	15.80 [8.90,25.36]
[5] 1948-2017	62.00 [56.59,67.36]	86.39 [80.69,91.04]	52.46 [46.51,58.63]	70.81 [65.86,75.73]	34.79 [27.48,42.14]	3.17 [1.37, 6.49]	43.83 [38.37,49.88]	41.02 [34.62,47.68]	5.32 [2.49, 9.78]	14.96 [9.05,22.02]
[6] 1960-2007	55.40 [48.07,62.00]	78.24 [71.45,84.76]	48.87 [41.66,56.51]	70.64 [64.26,75.98]	36.65 [27.53,44.92]	15.65 [8.55,24.41]	44.61 [37.30,52.12]	42.96 [35.77,50.99]	12.49 [6.78,20.65]	16.21 [8.36,25.16]
[7] pre-Volcker	60.57 [50.61,68.94]	71.01 [61.45,80.34]	45.61 [34.80,56.13]	61.91 [51.71,70.91]	39.59 [28.04,50.75]	5.58 [2.11,14.16]	45.38 [36.13,55.02]	43.92 [32.53,54.58]	19.53 [11.25,30.88]	23.52 [13.15,37.61]
[8] post-Volcker	46.34 [37.67,54.73]	77.66 [68.56,84.52]	40.88 [32.34,50.00]	66.18 [57.96,73.11]	35.62 [25.20,45.83]	7.63 [3.30,14.45]	26.34 [19.91,33.98]	27.27 [19.66,35.55]	3.59 [1.36, 8.36]	17.45 [8.49,27.94]
[9] Extended	47.56 [41.35,54.06]	65.28 [58.72,72.44]	40.18 [33.19,46.69]	56.71 [50.57,63.11]	31.43 [24.19,38.98]	4.73 [2.11, 9.17]	40.33 [33.83,46.75]	42.69 [33.03,51.32]	10.89 [6.26,17.07]	17.55 [9.74,28.63]
[10] Financial	53.90 [47.10,60.67]	75.33 [68.02,82.18]	43.57 [36.40,50.77]	62.44 [55.85,68.60]	35.42 [27.88,43.94]	5.19 [2.82, 9.31]	41.43 [34.82,47.79]	38.42 [31.20,45.65]	11.54 [6.56,17.79]	19.98 [12.54,29.42]
[11] Chained-type C&I	57.80 [51.26,63.00]	85.61 [79.50,90.50]	43.46 [36.44,50.46]	69.68 [64.03,74.42]	32.40 [25.23,40.53]	2.76 [1.43, 5.07]	39.00 [33.69,45.44]	31.36 [24.98,37.66]	8.85 [4.82,14.15]	18.31 [11.45,26.07]

Table 17: The MBC Shock, Targeting Hours Worked, Variance Contributions (6-32 Quarters)

	u	Y	h	I	C	TFP	Y/h	Wh/Y	π	R
[1] Benchmark	49.84 [42.43,56.53]	47.54 [38.20,55.67]	70.45 [64.25,77.04]	47.99 [38.49,55.96]	21.78 [15.30,29.22]	11.62 [6.14,18.14]	22.61 [15.58,29.66]	19.47 [11.73,29.24]	7.23 [3.32,13.31]	22.38 [15.09,31.87]
[2] VECM(1)	52.16 [45.43,58.79]	46.09 [38.60,54.00]	58.32 [53.32,63.44]	48.52 [41.43,55.97]	32.81 [23.69,43.81]	28.64 [15.87,40.14]	23.63 [16.48,32.70]	18.58 [11.07,29.94]	13.87 [5.96,25.56]	39.95 [25.71,53.96]
[3] VECM(2)	53.91 [45.99,61.44]	50.41 [39.92,59.50]	57.82 [52.81,62.97]	49.65 [41.34,57.77]	41.91 [26.35,55.14]	16.99 [7.95,28.24]	18.34 [11.83,26.80]	25.72 [13.88,40.00]	10.93 [4.98,18.61]	23.69 [12.51,44.29]
[4] 4 lags	51.82 [44.30,58.55]	46.53 [37.75,56.07]	70.17 [63.67,76.61]	45.99 [36.73,54.81]	23.11 [16.46,30.73]	10.22 [5.22,18.05]	19.54 [13.51,26.97]	19.25 [10.70,28.70]	6.80 [3.25,11.93]	24.55 [15.81,33.91]
[5] 1948-2017	51.98 [45.78,57.75]	57.31 [50.34,63.96]	76.44 [70.91,81.81]	56.45 [48.96,63.94]	23.48 [16.93,30.48]	8.49 [4.35,14.47]	23.93 [17.81,30.80]	25.26 [17.60,34.06]	7.85 [4.09,13.28]	16.43 [10.29,23.33]
[6] 1960-2007	53.21 [46.03,60.23]	50.95 [42.71,59.85]	70.91 [63.83,77.35]	52.51 [44.58,60.62]	21.39 [13.62,30.77]	5.83 [2.43,10.92]	18.52 [11.48,27.04]	26.91 [17.88,37.23]	7.75 [3.22,15.89]	18.67 [10.74,29.52]
[7] pre-Volcker	45.56 [33.61,56.43]	47.14 [36.05,58.16]	67.93 [58.71,76.98]	50.35 [37.67,61.28]	23.45 [14.49,35.17]	19.40 [9.36,30.19]	27.09 [17.01,39.19]	21.50 [11.24,35.79]	17.76 [10.48,29.26]	24.53 [13.88,40.91]
[8] post-Volcker	50.25 [42.13,58.72]	44.09 [35.21,53.20]	72.21 [63.07,80.11]	44.75 [35.54,54.40]	19.96 [12.81,28.05]	6.93 [2.91,13.12]	16.02 [9.41,23.70]	14.80 [8.02,24.30]	3.61 [1.32, 8.35]	13.01 [5.95,22.97]
[9] Extended	43.09 [36.17,49.37]	41.15 [33.92,48.68]	61.33 [55.15,67.24]	43.02 [35.24,50.56]	23.81 [17.36,31.14]	10.31 [4.86,17.26]	22.64 [16.06,29.80]	14.07 [7.21,24.92]	12.55 [7.13,19.79]	26.84 [17.00,38.82]
[10] Financial	50.45 [42.91,57.94]	49.94 [39.97,58.12]	63.65 [57.66,69.85]	50.13 [39.73,59.48]	27.20 [18.83,36.31]	11.29 [5.35,18.69]	25.81 [17.25,35.19]	22.27 [12.36,34.58]	8.77 [4.29,16.04]	26.53 [17.47,35.91]
[11] Chained-type C&I	48.43 [41.28,54.30]	46.76 [39.48,53.30]	78.87 [72.70,85.02]	46.11 [38.74,52.65]	20.37 [14.80,26.59]	10.92 [6.25,16.74]	19.51 [14.11,26.13]	13.41 [7.98,20.55]	5.76 [2.67,10.49]	20.27 [13.40,27.48]

Table 18: The MBC Shock, Targeting Investment, Variance Contributions (6-32 Quarters)

	u	Y	h	I	C	TFP	Y/h	Wh/Y	π	R
[1] Benchmark	59.03 [51.73,64.55]	66.60 [60.40,72.21]	45.20 [37.93,51.98]	80.29 [72.82,86.97]	19.01 [12.27,27.34]	3.81 [1.38, 7.83]	33.74 [27.72,40.30]	36.44 [29.21,44.21]	7.69 [3.65,12.96]	21.51 [13.91,30.28]
[2] VECM(1)	54.47 [47.86,60.60]	55.01 [48.96,61.65]	45.49 [38.05,53.16]	61.58 [55.78,68.31]	34.54 [25.35,45.08]	12.29 [5.84,22.09]	26.98 [20.34,34.17]	32.02 [22.71,41.00]	9.54 [4.00,18.54]	29.65 [16.48,45.86]
[3] VECM(2)	55.79 [49.03,62.87]	60.32 [53.38,67.58]	46.08 [39.48,53.54]	63.02 [56.30,69.67]	44.57 [32.28,55.14]	8.59 [4.23,15.06]	27.15 [20.32,34.38]	37.96 [28.53,46.61]	9.59 [3.90,17.23]	20.51 [11.51,33.76]
[4] 4 lags	59.99 [53.25,66.00]	66.75 [60.22,72.56]	43.60 [36.01,51.36]	79.98 [72.18,86.39]	20.51 [13.93,28.34]	5.22 [1.99,10.09]	32.41 [26.04,39.20]	37.29 [29.53,44.75]	7.29 [3.68,12.94]	21.25 [13.48,30.63]
[5] 1948-2017	61.66 [56.29,67.03]	72.01 [67.21,76.62]	53.31 [46.78,59.21]	85.20 [79.20,90.07]	21.44 [14.54,29.61]	2.98 [1.19, 6.60]	36.88 [30.74,43.40]	36.80 [30.54,43.51]	7.46 [3.92,13.31]	18.81 [12.01,26.03]
[6] 1960-2007	56.94 [50.22,63.46]	67.79 [60.98,73.81]	48.22 [40.67,55.65]	81.22 [74.33,87.11]	23.69 [15.10,32.48]	11.53 [5.03,20.50]	36.28 [28.74,43.88]	37.39 [29.88,45.86]	11.20 [5.71,19.37]	22.37 [13.64,31.05]
[7] pre-Volcker	62.79 [53.55,70.93]	60.25 [49.47,69.59]	48.49 [37.33,58.33]	72.75 [62.21,81.58]	24.92 [13.48,37.86]	7.25 [2.49,15.90]	36.32 [25.65,47.21]	32.97 [21.26,45.81]	17.94 [9.66,29.65]	29.75 [17.67,44.22]
[8] post-Volcker	51.27 [42.22,59.14]	62.59 [54.28,69.59]	40.40 [31.31,49.31]	82.79 [73.94,89.33]	21.88 [14.03,31.04]	5.89 [2.18,11.48]	19.01 [13.31,26.96]	25.19 [17.22,33.15]	3.72 [1.42, 7.89]	17.72 [9.66,27.00]
[9] Extended	49.51 [43.52,55.92]	56.64 [50.63,62.73]	42.79 [35.92,48.65]	65.72 [58.67,72.73]	20.17 [13.41,27.74]	3.91 [1.54, 8.04]	34.47 [28.44,41.41]	41.46 [31.04,50.82]	10.87 [6.03,16.56]	21.42 [12.92,32.65]
[10] Financial	57.04 [50.63,63.29]	63.64 [57.22,69.74]	44.94 [37.75,52.18]	74.05 [66.67,80.32]	23.94 [15.73,32.62]	4.92 [2.40, 9.55]	35.15 [28.22,41.96]	35.00 [27.81,42.40]	8.54 [4.11,14.77]	24.44 [16.05,33.52]
[11] Chained-type C&I	59.34 [53.12,64.87]	69.12 [63.69,74.05]	42.24 [34.89,49.01]	86.02 [79.25,90.69]	18.43 [12.48,25.84]	2.42 [0.93, 4.86]	31.03 [25.76,37.27]	27.74 [21.75,34.11]	6.49 [3.22,11.60]	22.05 [15.23,29.99]

Table 19: The Inflation Shock, Variance Contributions (6-32 Quarters)

	u	Y	h	I	C	TFP	Y/h	Wh/Y	π	R
[1] Benchmark	4.24	7.88	3.32	3.01	15.14	3.55	7.37	1.96	83.03	7.61
	[1.62, 8.20]	[3.77,12.87]	[1.21, 6.92]	[1.12, 6.60]	[10.00,21.93]	[1.75, 7.08]	[4.11,12.31]	[0.66, 4.60]	[76.11,88.46]	[3.36,14.61]
[2] VECM(1)	11.81	14.22	11.98	9.92	21.13	12.05	17.10	6.59	86.63	18.65
	[5.47,19.24]	[7.93,22.41]	[5.34,19.78]	[4.24,16.89]	[12.53,30.13]	[6.91,18.10]	[9.85,24.32]	[2.74,12.27]	[80.27,91.16]	[10.45,27.75]
[3] VECM(2)	4.03	2.00	4.46	3.11	1.84	11.15	3.37	4.22	85.90	5.17
	[1.31, 8.55]	[0.64, 5.43]	[1.77, 8.71]	[1.05, 7.13]	[0.47, 5.11]	[6.55,16.98]	[1.34, 7.04]	[2.01, 7.65]	[78.72,91.04]	[2.38, 9.46]
[4] 4 lags	5.08	9.21	3.87	3.49	15.77	3.70	9.85	2.30	82.22	6.89
	[2.14, 9.53]	[4.82,15.07]	[1.49, 8.12]	[1.18, 7.46]	[10.29,22.23]	[1.89, 6.83]	[5.48,15.73]	[0.81, 5.66]	[76.14,87.42]	[2.84,13.13]
[5] 1948-2017	2.71	2.53	4.60	5.90	12.50	7.25	6.62	2.03	86.62	6.52
	[0.95, 5.85]	[0.88, 5.31]	[2.00, 7.99]	[3.24, 9.79]	[7.19,19.13]	[3.47,12.15]	[3.57,10.92]	[0.65, 4.87]	[81.29,90.86]	[2.54,12.23]
[6] 1960-2007	8.86	8.93	10.01	5.84	19.06	3.47	10.74	4.70	80.78	11.71
	[4.33,15.49]	[4.25,16.27]	[4.63,17.43]	[2.52,11.75]	[12.21,27.47]	[1.68, 7.16]	[5.63,17.61]	[1.95, 9.68]	[73.48,86.89]	[5.21,20.70]
[7] pre-Volcker	10.46	14.57	6.81	11.29	21.23	12.30	17.25	8.99	66.39	9.26
	[3.59,22.60]	[6.74,27.14]	[2.18,17.67]	[4.00,22.56]	[12.76,32.51]	[5.03,24.28]	[9.03,28.81]	[3.32,20.32]	[55.30,77.59]	[3.22,23.14]
[8] post-Volcker	6.76	9.02	7.02	5.40	14.74	2.34	7.96	2.51	87.67	22.97
	[2.78,13.21]	[4.46,16.22]	[2.70,13.10]	[2.18,10.68]	[8.25,23.75]	[0.85, 6.05]	[3.50,14.84]	[0.95, 5.88]	[81.23,92.33]	[12.99,33.79]
[9] Extended	8.24	9.45	7.13	5.22	14.13	5.30	11.37	3.68	75.28	13.59
	[3.68,14.72]	[4.90,15.69]	[3.06,13.46]	[1.95,10.61]	[8.24,21.01]	[2.67, 9.34]	[6.50,17.81]	[1.43, 8.28]	[67.59,81.92]	[7.09,22.06]
[10] Financial	4.85	7.93	3.88	3.69	14.06	3.92	7.89	2.07	80.61	8.49
	[2.03, 9.37]	[3.94,13.34]	[1.46, 8.31]	[1.32, 7.50]	[8.20,20.25]	[1.88, 7.15]	[4.36,12.52]	[0.83, 4.70]	[73.22,86.65]	[3.55,15.36]
[11] Chained type C&I	1.88	4.64	1.54	2.11	6.80	3.23	6.25	1.75	80.18	6.92
	[0.57, 5.11]	[1.86, 9.17]	[0.56, 3.91]	[0.68, 4.69]	[3.03,11.99]	[1.41, 6.36]	[2.97,10.39]	[0.64, 3.88]	[73.40,85.71]	[2.40,13.24]

G.2 Stock Prices, Relative Price of Investment, and Utilization

Here, we describe additional properties of the specification in row 9 (“Extended”) of Tables 8-9 and 16-19. Recall that this specification contains three additional variables: stock prices (SP); the relative price of investment (P_i/P_c); and capital utilization (z). Our measure of stock prices is in real terms, is the same as that used by Beaudry and Portier, and is taken from Robert Shiller’s website (http://www.econ.yale.edu/~shiller/data/ie_data.xls). The relative price of investment is the ratio of the price of Gross Private Domestic Investment and Durables to the price of Non Durables and Services; its computation is detailed in Online Appendix G.5. Finally the capacity utilization rate variable corresponds to the Capacity Utilization in Manufacturing (SIC), $CUMFNS$ in the Federal Reserve Economic Database.

The inclusion of stock prices and the relative price of investment is motivated by works that uses these variable in the identification of, respectively news shocks and investment-specific technology shocks. The inclusion of capacity utilization, on the other hand, helps shed light on why labor productivity moves with the MBC shock while TFP does not. Last but not least, the inclusion of all three variables at once helps illustrate the robustness of our main findings to the addition of more information—a point already made in Tables 8-9 and 16-19.

Here, Tables 20-21 and Figure 14 complete the picture by reporting the contribution of the MBC shock to the short-run and long-run volatility of the aforementioned three variables, as well as the properties of the shock that targets the business-cycle volatility of stock prices.⁵⁴ The most noteworthy new findings are the following.

First, the disconnect between the business cycle and technology applies to both TFP and investment-specific technology, as measured by the relative price of investment. For instance, the MBC shock explains less than 5% of the volatility of either of these variables at either the business-cycle or the long-run frequencies.

Second, the shock that targets Stock Prices accounts for 21 to 24% of the business-cycle volatility in unemployment, output and investment, and 15 to 22% of the long-run volatility in TFP, output and investment. In this regard, the fluctuations in stock prices appear to be disconnected from current technology and to contain non-trivial statistical information about both the business cycle and the long-term prospects of the economy. The extent to which these patterns reflect the presence of a news shock is explored further in Appendix C.

Finally, the shock that targets utilization at the business-cycle frequencies is similar to the MBC shock in terms of both variance contributions and IRFs (Figure 14). This helps understand why labor productivity increases in response to the MBC shock, while TFP does not move.

⁵⁴The shocks that target business cycle volatility in TFP and the relative price of investment lack novelty as they contribute negligibly to the volatility of the macroeconomic quantities.

Table 20: Extended VAR, Business Cycle Variance Contributions

	u	Y	h	I	C	z	TFP	Y/h
MBC shock	59.33 [53.73,65.69]	50.61 [43.05,57.99]	45.50 [39.71,51.26]	52.91 [44.97,60.17]	21.83 [14.87,31.14]	51.71 [45.55,57.66]	4.81 [1.95,10.39]	26.69 [19.36,34.75]
SP shock	24.14 [18.31,31.23]	23.05 [16.99,29.55]	15.75 [10.45,22.24]	21.65 [15.75,28.29]	24.63 [18.47,31.05]	18.10 [12.64,24.36]	4.37 [2.46, 7.30]	10.81 [6.55,16.04]
	P_i/P_c	SP	wh/Y	π	R	GDP/h	w	r
MBC shock	4.42 [1.69, 9.62]	11.54 [5.16,22.75]	27.82 [14.05,44.15]	12.12 [6.57,19.70]	28.99 [17.38,42.75]	10.70 [5.36,19.24]	4.48 [1.93,10.10]	12.52 [5.56,21.67]
SP shock	3.39 [1.32, 7.33]	82.82 [76.59,87.93]	11.29 [6.25,17.22]	9.27 [4.28,14.73]	5.48 [2.40,10.26]	12.39 [7.59,18.64]	13.19 [7.93,19.43]	2.40 [0.87, 5.04]

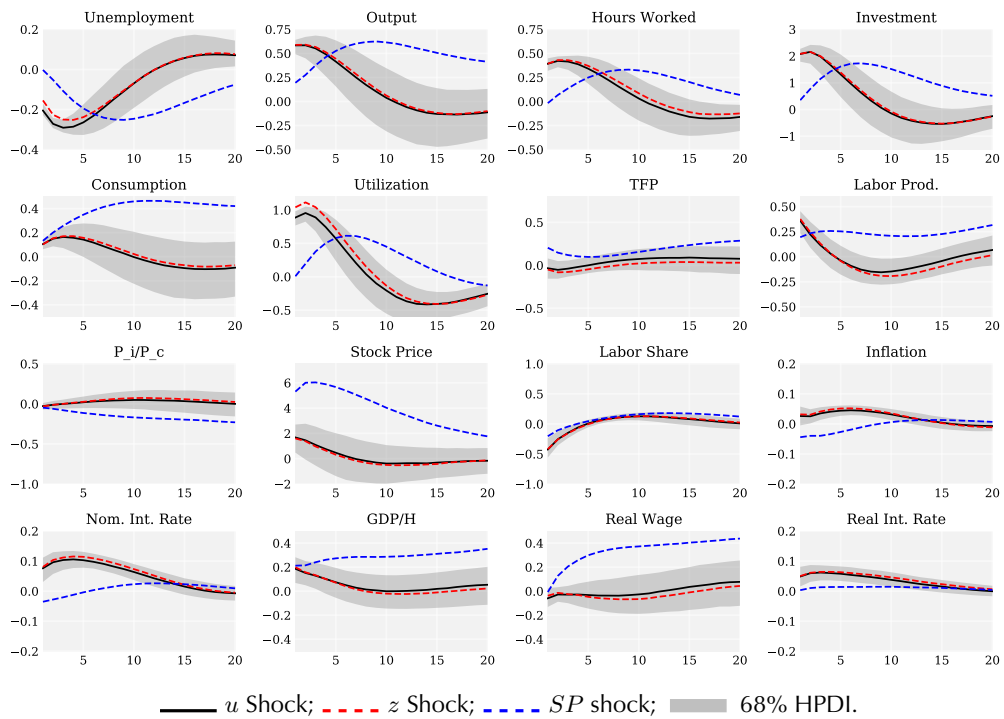
Note: The rows correspond to the shocks targeting business-cycle variation in unemployment (MBC shock) and Stock Prices (SP shock), respectively. The columns correspond to the 13 variables in the VAR. These are the 10 variables from our baseline specification, and also capacity utilization z , the Relative Price of Investment P_i/P_c and stock prices SP .

Table 21: Extended VAR, Long-Run Variance Contributions (80- ∞ Quarters)

	u	Y	h	I	C	z	TFP	Y/h
MBC shock	9.49 [3.03,24.04]	4.52 [0.45,17.60]	3.96 [1.11,11.23]	4.58 [0.78,18.25]	4.43 [0.40,16.92]	6.36 [2.19,15.41]	4.39 [0.59,17.66]	4.59 [0.52,17.40]
SP shock	30.39 [14.89,47.64]	14.55 [2.96,38.30]	8.95 [2.29,25.66]	14.85 [3.49,38.37]	14.76 [2.87,38.64]	17.35 [7.85,32.71]	21.67 [5.53,43.52]	21.88 [5.59,44.09]
	P_i/P_c	SP	wh/Y	π	R	GDP/h	w	r
u	4.60 [0.59,17.02]	5.23 [1.13,16.97]	4.36 [0.79,14.99]	7.03 [2.20,16.45]	11.23 [2.88,24.32]	4.55 [0.50,17.35]	4.58 [0.50,17.62]	8.19 [2.05,19.88]
SP	26.99 [8.96,46.73]	34.63 [16.95,52.70]	24.51 [9.89,42.83]	9.16 [3.08,21.50]	12.68 [3.21,32.22]	20.31 [4.72,42.71]	20.96 [5.00,43.46]	18.88 [6.11,37.94]

Note: The rows correspond to the shocks targeting business-cycle frequencies variation in unemployment (MBC shock) and Stock Prices (SP shock) respectively. The columns correspond to the 13 variables in the VAR. These are the 10 variables from our baseline specification, plus capacity utilization (z), the Relative Price of Investment (P_i/P_c) and stock prices (SP).

Figure 14: Extended VAR, IRFs



G.3 Financial Variables

Here we provide additional information on the VAR that adds the credit spread (CS) and appears as row 10 (“Financial”) of Tables 8-9 and 16-19. We also consider a more comprehensive specification, called “Financial-Full,” that contains three additional financial variables at the expense of a shorter sample period. The additional variables are the slope of the term structure (TS), the level of credit to non-financial firms (Cr), and the net worth of such firms (WS).

Our measurement of all these variables follows Christiano, Motto, and Rostagno (2014). The credit spread (CS) is the difference between the interest rate on BAA-rated corporate bonds and the 10 year US government bond rate. The slope of the term structure (TS) is the difference between the 10-year constant maturity US government bond yield and the Federal Funds rate. The level of credit (Cr) is taken from the Flow of Funds of the US Federal Reserve Board. Finally, net worth (WS) is measured by the Dow Jones Wilshire 5000 index.⁵⁵ Because this index only starts in 1971 and the measure of credit is only available until 2014, the VAR that contains all four financial variables (“Financial-Full”) is estimated for the period running from 1971Q1 to 2014Q4. By contrast, the VAR that contains only the credit spread (“Financial”, or row 10 of the aforementioned tables) spans the entire 1955Q1-2017Q4 period.

For the purposes of the model evaluation done in Section 6, we have also considered a third specification, which is obtained by restricting the second specification to 1985Q1-2010Q4. This is the period used in the original estimation of the model in Christiano, Motto, and Rostagno (2014). We refer to this specification as “Financial-CMR.”

Figure 15 reports the IRFs of the various facets of the MBC shock obtained from these three specifications. Although there are some differences,⁵⁶ the main picture remains the same: the reduced-form shocks obtained by targeting unemployment, hours, output, investment and consumption are highly interchangeable.

Perhaps more interestingly, we can now detect the empirical footprint of the MBC shock on the new, financial variables. In particular, we see that the credit spread spikes on impact, while output and the other key macroeconomic quantities respond with a delay, in a hump-shaped manner. From this perspective, the credit spread leads the business cycle. As discussed in Section 6, this property, which is presumably informative about the real-financial nexus, is unfortunately not captured by the model of Christiano, Motto, and Rostagno (2014).⁵⁷

⁵⁵Note that the measure of net worth is a stock-market valuation, which differs from that used in the previous subsection (SP500) because the present specification aims at replicating the data used in CMR, while the previous one followed Beaudry and Portier. In any case, it makes little difference which one of these two measures is used as their business-cycle behavior is nearly identical.

⁵⁶Most notably, consumption appears to more closely connected to the MBC shock in the third specification.

⁵⁷Although we have omitted it here, we have also looked at the shock that targets the credit spread itself. This shock is similar to the MBC shock in terms of IRFs (comovements), although less so with regard to variance contributions. Importantly, this shock, too, gives rise to pattern mentioned above, with the credit spread itself moving before the key macroeconomic quantities.

Figure 15: Comparing Business-Cycle Factors

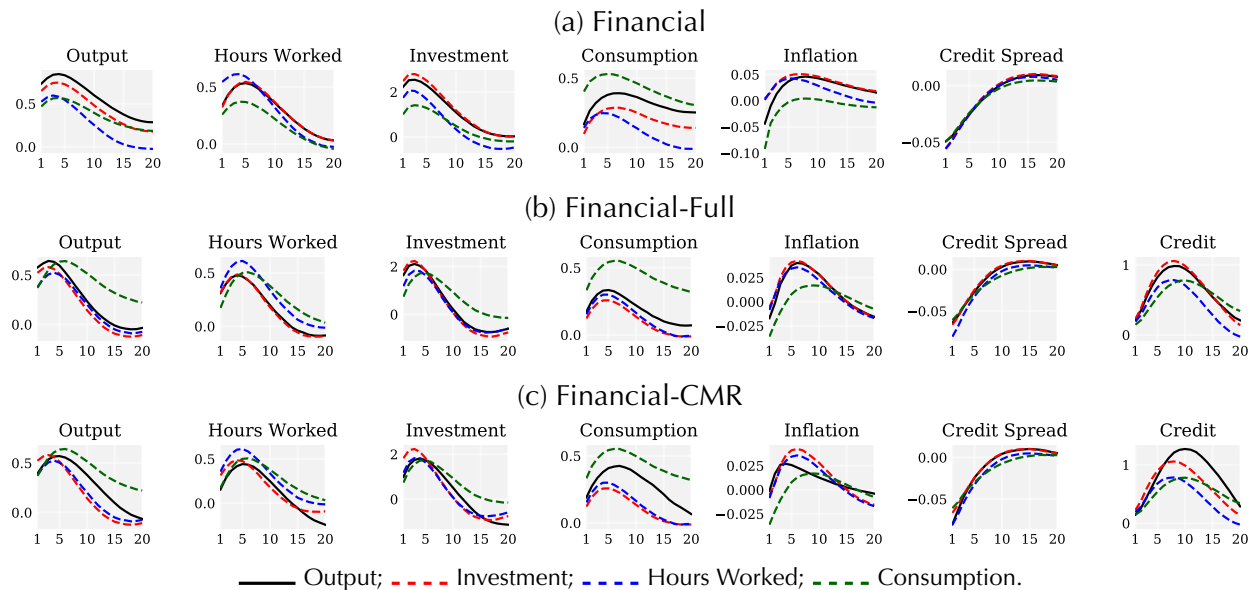


Table 22: Financial VARs, Short-Run Contributions of MBC Shock

	u	Y	h	I	C	π	CS	Cr
Financial	68.57 [62.38,74.87]	57.56 [49.74,64.87]	46.84 [39.39,54.03]	59.95 [52.26,66.82]	25.94 [17.80,34.98]	8.42 [3.77,14.98]	41.56 [30.02,54.08]	
Financial-Full	60.47 [54.39,67.41]	51.65 [43.81,59.37]	53.32 [45.45,61.18]	54.63 [47.11,62.53]	33.84 [22.66,46.48]	13.29 [6.12,24.44]	49.68 [29.51,62.90]	39.69 [28.46,51.23]
Financial-CMR	64.76 [56.31,73.66]	53.26 [40.61,64.00]	59.60 [48.45,69.36]	55.90 [45.01,66.30]	35.93 [21.80,51.92]	15.83 [6.79,30.26]	56.05 [36.50,70.72]	46.16 [29.83,61.78]

Note: The rows correspond to the shocks targeting business-cycle frequencies variation in unemployment (MBC shock) for the various financial VARs described in the text. CS denotes the Credit Spread, Cr the measure of credit.

G.4 Description of VECMs

We now fill in the details of the VECMs reported in rows 3 and 4 of Tables 8-9 and 16-19. Both of these VECMs are nested in the following form:

$$\Delta X_t = \Gamma_0 \Theta X_{t-1} + \sum_{i=1}^p \Gamma_i \Delta X_{t-i} + \nu_t$$

where Θ is the matrix of co-integration coefficients and Γ_0 is the matrix of loadings of these co-integration relationships. The difference between the two VECMs is the specification of the number of unit roots and the co-integration relations.

In $VECM_1$, we assume that the real quantities (Y, C, I, APL) and TFP share a single stochastic trend, while the remaining variables are assumed to be stationary. The co-integrating relationship is of the type $x_t = \alpha_x + \beta_x TFP_t$ for each variable $x \in \{Y, C, I, APL\}$.

In $VECM_2$, the real quantities (Y, C, I, APL) and TFP share one stochastic trend; the nominal variables, π and R , share another stochastic trend; and the remaining variables (the unemployment, hours, and the labor share) are stationary. The co-integration relationships are of the type $x_t = \alpha_x + \beta_x TFP_t$ for $x \in \{Y, C, I, APL\}$ and $R_t = \delta + \gamma \pi_t$.

We have also considered a third specification that allows the number of stochastic trends and the co-integration relationships to be determined completely a-theoretically, by means of the standard maximum eigenvalue and trace tests proposed by Johansen and Juselius (1990). Relative to the aforementioned two specifications, this “unrestricted” VECM marginally reinforces the disconnect between the short run and the long run;⁵⁸ but it also produces six (!) unit roots, which makes little sense from the perspective of theory.

G.5 Measuring the Relative Price of Investment

We now describe the measure of the relative price of investment that is used in one of our robustness exercises, the one appearing as row 9 (“Extended”) of Tables 8-9 and 16-19.

Let P_t^x denote the chained price index of aggregate x at time t , and similarly Q_t^x the quantity of aggregate x at time t , where x can denote either gross domestic private investment (GPDI), durable consumption (D), non durable consumption (ND) or services (S). The change in investment ($I=GPDI+D$) price, is then given by

$$\Delta P_t^i = \sqrt{\Delta P_t^i(Q_{t-1}^i) \Delta P_t^i(Q_t^i)} - 1$$

where

$$\Delta P_t^i(Q_{t-1}^i) = \frac{P_t^{\text{gpdi}} Q_{t-1}^{\text{gpdi}} + P_t^{\text{d}} Q_{t-1}^{\text{d}}}{P_{t-1}^{\text{gpdi}} Q_{t-1}^{\text{gpdi}} + P_{t-1}^{\text{d}} Q_{t-1}^{\text{d}}} \quad \text{and} \quad \Delta P_t^i(Q_t^i) = \frac{P_t^{\text{gpdi}} Q_t^{\text{gpdi}} + P_t^{\text{d}} Q_t^{\text{d}}}{P_{t-1}^{\text{gpdi}} Q_t^{\text{gpdi}} + P_{t-1}^{\text{d}} Q_t^{\text{d}}}$$

Similarly, we define the change in the consumption ($C=ND+S$) price as

$$\Delta P_t^c = \sqrt{\Delta P_t^c(Q_{t-1}^c) \Delta P_t^c(Q_t^c)} - 1$$

⁵⁸In particular, the unemployment shock accounts 10% of the long-run volatility in output and TFP, compared to 14% in $VECM_1$ or $VECM_2$.

where

$$\Delta P_t^c(Q_{t-1}^c) = \frac{P_t^{\text{nd}}Q_{t-1}^{\text{nd}} + P_t^sQ_{t-1}^s}{P_{t-1}^{\text{nd}}Q_{t-1}^{\text{nd}} + P_{t-1}^sQ_{t-1}^s} \text{ and } \Delta P_t^c(Q_t^c) = \frac{P_t^{\text{nd}}Q_t^{\text{nd}} + P_t^sQ_t^s}{P_{t-1}^{\text{nd}}Q_t^{\text{nd}} + P_{t-1}^sQ_t^s}$$

Let us denote by Q_t the relative price of investment as $Q_t = P_t^i/P_t^c$, then Q_t satisfied

$$Q_t = (1 + \Delta P_t^i - \Delta P_t^c)Q_{t-1}$$

H Pushing the AD-AS Example

In this appendix we conduct two “pedagogical” exercises motivated by the AD-AS example mentioned in Section 5. In the first, which is semi-structural in nature, we show that the narrative of offsetting demand and supply shocks does not work insofar as the supply shock is proxied by the productivity shock identified via our method. In the second exercise, which is fully structural, we show that this story is also inconsistent with a textbook New Keynesian model calibrated to the relevant elements of our anatomy.

H.1 Proxying the AS shock with the TFP shock

Our first, semi-structural exercise is based on the following simple idea. If the MBC shock is a mixture of an inflationary demand shock and a disinflationary supply shock, and if the supply shock reflects movements in productivity, then the documented disconnect between the MBC shock and inflation should be weakened, and the role of the demand shock be revealed, if we control for the effect of productivity. This in turn can be done by purging from the data the reduced-form shock that targets TFP over the business-cycle frequencies.⁵⁹ We thus repeat our identification of the shocks that target unemployment, GDP, and inflation after this purging and ask whether this reduces the disconnect between the MBC shock and inflation.

As evident in Table 23 and Figure 16, the answer is clearly negative. Whether we look at original reduced-form shocks or the ones obtained after purging the effects of productivity, the aforementioned disconnect and indeed the shocks themselves remain almost unchanged.

Furthermore, insofar as one accepts the interpretation of the MBC shock identified in the data as the AD shock in the theory, the challenge for the theory is twofold: not only does the MBC shock accounts for a small fraction of volatility in inflation, but it has such a small impact on inflation that the theory can make sense of only if the AS curve is extremely flat.

We illustrate this point in Figure 17. The solid black line shows the actual response of inflation to the MBC shock in the data. The dashed red line shows the response predicted by the New Keynesian Philips Curve, under a textbook calibration and with the real marginal cost proxied by the response

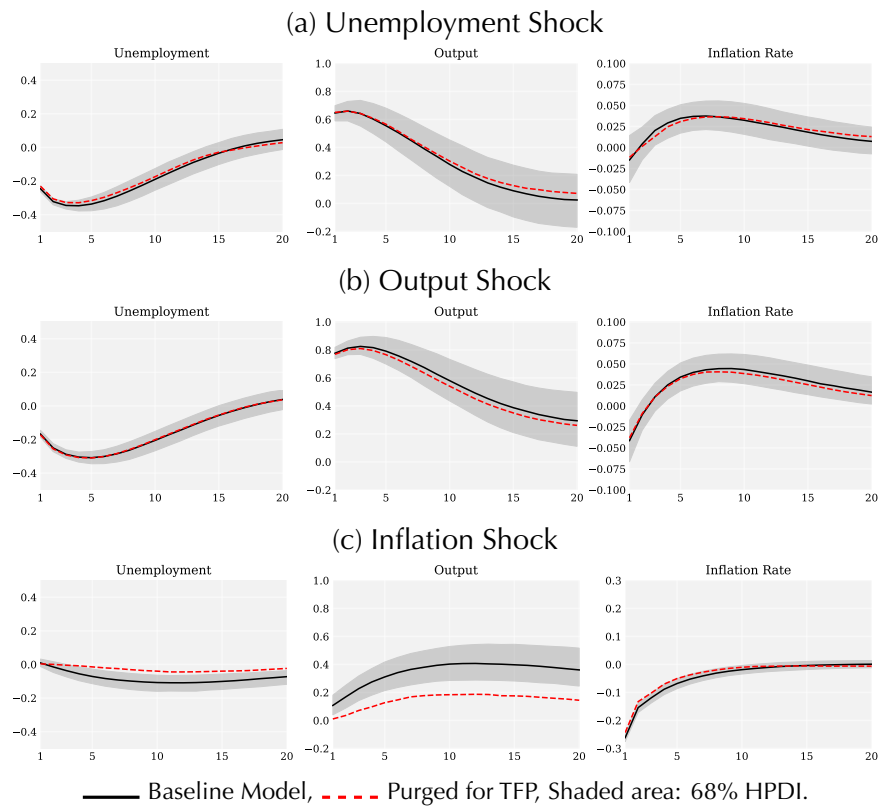
⁵⁹We have obtained almost identical results with a variant specification that proxies the supply shock with the technology shock identified as in Galí (1999), as well as with one that purges both the short-run and the long-run TFP shocks identified via our method. These alternatives, however, seem less appropriate for the present purposes, because they amount to purging also the effects of news about future productivity, which in standard models maps do a demand rather than a supply shock.

Table 23: Variance Contributions

	u	Y	π
<i>Unemployment Shock</i>			
Baseline	73.71	58.51	6.96
Purged	70.98	61.10	8.05
<i>Output Shock</i>			
Baseline	56.24	80.13	10.47
Purged	57.48	78.29	9.55
<i>Inflation Shock</i>			
Baseline	4.24	7.88	83.03
Purged	3.78	6.04	79.97

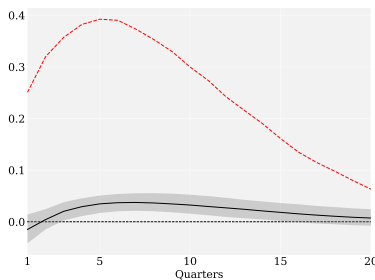
68% HPDI into brackets

Figure 16: Impulse Response Functions



of the labor share to the MBC shock.⁶⁰ The large gap between the two lines illustrates that, even after controlling for the possible sluggishness in the response of the real marginal cost due to wage rigidities, the predicted response of inflation is over 10 times larger than the actual one. Conversely, the Phillips curve has to be very flat for the theory to match the observed inflation response. A similar picture is painted in the next subsection, which takes a fully-structural approach to both the MBC shock and the shock that accounts for the volatility in inflation.

Figure 17: The MBC Shock and the NKPC



— Actual inflation response; Shaded area: 68% HPDI; - - - Predicted response.

H.2 A 2x2 New Keynesian model

We now turn to second, fully-structural exercise: we employ a two-shock, two-variable version of the New Keynesian model and ask what it takes for this model to account for the relevant elements our anatomy.

In particular, we estimate both the shock processes and the main parameters of the model—those that govern the slopes of the AS and AD curves and the sluggishness of the inflation and output dynamics—by minimizing the distance between four empirical IRFs and their theoretical counterparts. These are the IRFs of output and inflation to the output shock and to the inflation shock, as identified by our method. We focus on these objects because the simple, textbook-style model considered here is meant to speak to the only dynamics of output and inflation.⁶¹

We then use the estimated model to answer two questions. First, what parameter values (for instance, the slope of the Phillips curve) does the model need in order to achieve maximum fit vis-a-vis our facts? And second, does the MBC shock identified via our method correspond to a single structural shock in the model or to a mixture of structural shocks, as suggested by the AD-AS example used in Section 5?

Like the textbook version of the New Keynesian model, the version considered here reduces to two equations in the (y, π) space, one representing aggregate demand (AD) and the other representing

⁶⁰To construct this line, we proceed as follows. First, we take the New Keynesian Philips Curve: $\pi_t = \kappa x_t + \beta \mathbb{E}_t[\pi_{t+1}]$, where $\beta \in (0, 1)$ is the discount factor, $\kappa = (1 - \theta)(1 - \beta\theta)/\theta$, and θ is the Calvo parameter. Next, we set $\theta = 2/3$ (prices are, on average, reset every 3 quarters) and $\beta = 0.99$ (an annual discount rate of 4%). Finally, we feed x_t with the response of the labor share to the MBC shock.

⁶¹The empirical IRFs are obtained from our VAR by targeting the inflation rate or output (see Figure 2 for example). The theoretical IRFs are constructed in an analogous manner, treating the model as the DGP.

aggregate supply (AS). At the same time, our version mimics richer DSGE versions by allowing for a flat Phillips curve, habit persistence and price indexation. These enhancements may lack empirical micro-foundations but are customarily used in the literature in order to improve the model's empirical performance.

Let us start with the textbook version of the New Keynesian model, which can be expressed by the following equations:

$$y_t = -\sigma (R_t - \mathbb{E}_t[\pi_{t+1}]) + \mathbb{E}_t[y_{t+1}] + \sigma \xi_t \quad (3)$$

$$\pi_t = \lambda mc_t + \beta \mathbb{E}_t[\pi_{t+1}] + \lambda \mu_t \quad (4)$$

$$mc_t = \kappa y_t - \frac{1+\nu}{\alpha} a_t + \varsigma_t \quad (5)$$

$$R_t = \varphi \pi_t + \psi y_t + m_t \quad (6)$$

The interpretation is familiar: (3) is the Dynamic IS curve, (4) is the NKPC, (5) describes the real marginal cost as a function of output and productivity, and (6) specifies monetary policy. The notation is also standard: y_t is output, π_t is inflation, mc_t is the real marginal cost, R_t is the nominal interest rate, \mathbb{E}_t is the rational expectations operator, a_t is the productivity shock, ξ_t is the discount-rate shock, μ_t is the markup shock, ς_t is the cost-push shock, m_t is the monetary-policy shock, $\sigma > 0$ is the elasticity of intertemporal substitution, $\beta \in (0, 1)$ is the steady-state discount factor, $\lambda \equiv \frac{(1-\theta)(1-\beta\theta)}{\theta}$ is the slope of the NKPC with respect to the real marginal cost (and to the markup shock, too), θ is the Calvo parameter (the probability of a firm's not being able to reset its price), $\kappa \equiv \frac{1+\nu}{\alpha} + \frac{1-\sigma}{\sigma} > 0$ is the slope of the real marginal cost with respect to output, $\nu \geq 0$ is the Frisch elasticity of labor supply, $\alpha \in (0, 1]$ is the short-run elasticity of output with respect to labor, and $\varphi > 1$ and $\psi \geq 0$ parameterize the responsiveness of monetary policy to, respectively, inflation and output.

To simplify the exposition of the AD and AS curves below, we set $\psi = 0$.⁶² For the reported experiments, we also interpret a period as a quarter and set $\beta = .99$, $\varphi = 2$, $\alpha = 1$, and $\nu = 0$.⁶³ More crucially, the parameters λ and σ , which govern the slopes of the two curves, and two additional parameters, which are introduced momentarily and which govern the endogenous persistence in the model, are left free to be estimated in one of the experiments.

Substituting (6) in (3) and (5) in (4), we can reduce the model to the following two equations in output and inflation alone:

$$y_t = -\sigma\varphi\pi_t + \sigma\mathbb{E}_t[\pi_{t+1}] + \mathbb{E}_t[y_{t+1}] + u_t^d \quad (7)$$

$$\pi_t = \lambda\kappa y_t + \beta\mathbb{E}_t[\pi_{t+1}] - u_t^s \quad (8)$$

where $u_t^d \equiv \sigma\xi_t - \sigma m_t$ and $u_t^s \equiv \lambda\kappa a_t - \lambda\kappa\varsigma_t - \lambda\mu_t$. Condition (7) represents aggregate demand, AD, (8) represents aggregate supply, AS. Accordingly, u_t^d and u_t^s are the (composite) demand and supply shocks. We assume that these shocks follow independent $AR(1)$ process and let (σ_d, σ_s) denote their standard deviations and (ρ_d, ρ_s) their autocorrelations.

⁶²Since the experiments conducted here do not utilize data on the interest rate, the effect of a positive ψ on the dynamics of output and inflation can be proxied by appropriately adjusted values for other model parameters. Accordingly, we have verified that our findings about the model's performance remain essentially unchanged if we let, for example, $\psi = 0.5$.

⁶³The values of β and φ are standard, while those for α and ν help reduce the sensitivity of the real marginal cost to output (intuitively, a high value for α mimics variable utilization and a low value for ν mimics real wage rigidity), which in turn helps improve the empirical performance of the model (and makes our own job harder)

This completes the description of the baseline version of the New Keynesian model, which is the building block for the enhanced, DSGE-like variant used here. This variant is obtained by including habit persistence in the Dynamic IS curve and by replacing the standard NKPC with the hybrid one. The modified equations are given by

$$\begin{aligned}
 y_t &= -\sigma \frac{1-h}{1+h} (\varphi \pi_t - \mathbb{E}_t \pi_{t+1}) + \frac{1}{1+h} \mathbb{E}_t y_{t+1} + \frac{h}{1+h} y_{t-1} + u_t^d \\
 \pi_t &= \lambda \left(\kappa y_t + \frac{h}{\sigma(1-h)} (y_t - y_{t-1}) \right) + \frac{\beta \theta}{\theta + \omega(1-\theta(1-\beta))} \mathbb{E}_t \pi_{t+1} + \frac{\omega}{\theta + \omega(1-\theta(1-\beta))} \pi_{t-1} - u_t^s
 \end{aligned}$$

for some $h \in [0, 1)$ and $\omega \in [0, 1)$. These capture the inertia added to the aggregate demand and aggregate supply equations, respectively.⁶⁴ Finally, λ is allowed to take low enough values so as to accommodate a relatively weak positive co-movement between inflation and output in response to demand shocks.

Let $\Theta \equiv (\sigma_d, \sigma_s, \rho_d, \rho_s; \lambda, \sigma, h, \omega)$ collect the parameters that regulate the shock processes and the internal propagation, namely the slopes of the AS and AD curves and the corresponding sources of sluggishness. We estimate Θ by minimizing the distance between the IRFs of output and inflation to the output and inflation shocks identified in the data via our method and the corresponding objects in the model.

Table 24 reports the estimated parameter values. Table 25 reports the variance contributions of the model's two structural shocks. The most notable features are that λ is nearly zero, that the output fluctuations are dominated by a non-inflationary demand shock, and that the inflation fluctuations are dominated by a disinflationary supply shock. That is, confronted with the relevant elements of our anatomy, the model demands a very flat AS (or Philips) curve and specialized structural shocks, a picture consistent with that painted in Section 3.⁶⁵

Table 24: Parameters

σ_s	σ_d	ρ_s	ρ_d	h	ω	λ	σ
0.0789	0.0316	0.7016	0.9540	0.1979	0.0000	0.0004	0.2764

The purpose of this—pedagogical—exercise was to illustrate how the combination of our anatomy with a model can help discipline the AD-AS narrative offered in Section. The same strategy is applied to, and works well for, the three state-of-the-art DSGE models considered in Section 6. Naturally, while all of these exercises support the interpretation of the empirical MBC shock as a non-inflationary demand shock, they cannot establish its universality.

⁶⁴The standard interpretation of h is as the degree of habit persistence in consumption. But as there is no capital in the model, h represents all the adjustment frictions in aggregate demand. On the other hand, ω corresponds to the fraction of irrational, backward-looking firms in Galí and Gertler (1999), or the degree of automatic past-price indexation in Christiano, Eichenbaum, and Evans (2005). These model enhancements lack solid empirical micro-foundations but are customarily used in the DSGE literature.

⁶⁵Another interesting finding, which is though not particular relevant for the present purposes, is that the estimation of the model based on our anatomy yields $\omega = 0$, that is, no past-price indexation or backward-looking element in the Philips curve. This appears to be driven by the absence of sluggishness in the response of inflation to the inflation shock and suggests that the “right” model is one that somehow allows for such sluggishness in the response of inflation to the main driver of the real quantities without however introducing such sluggishness in the overall inflation dynamics.

Table 25: Variance Contributions

	Output	Inflation
Supply Shock	7.62	98.90
Demand Shock	92.38	1.10

I Robustness of Model Evaluations

This appendix assesses the robustness of the lessons drawn in Section 6 regarding the evaluation of the JPT and ACD models under the lenses of our method.

I.1 Running the Same VAR on Data and Models

In the main text, we evaluated the ability of JPT and ACD to account for the MBC shock in the data using the theoretical, asymptotic properties of the two models. We now explore the robustness of our findings to a Monte Carlo exercise that runs the same, small-size VAR on artificial data from each model and on the actual US data.

Because both models have a stochastic dimension smaller than that of our benchmark VAR, first rerun our empirical specification on a restricted VAR featuring Output, Consumption, Investment, Hours worked, Fernald’s measure of Total Factor Productivity (corrected for utilization), the nominal interest rate and the inflation rate. As can be seen in the first row of Figure 18, this smaller VAR gives rise to the same picture as our baseline VAR: the shocks that target output, hours, investment and consumption are essentially indistinguishable from one another.

Because the smaller VAR run here has exactly the same stochastic dimension as the JPT model, it can be readily run on artificial data generated by that model. By contrast, the ACD model has one dimension less: being a flexible-price, no-monetary model, it makes no prediction about inflation (and nominal variables). To be able run the same VAR on artificial data from that model, we augment it with the simplest model of inflation we could think of: an exogenous AR(1) process.⁶⁶ Clearly, this add-on has no effect on the model’s predictions regarding any of the real variables. It only permits us to run the same VAR on the two models under consideration.

Each model is then simulated 1000 times to generate artificial time series for the aforementioned set of variables. Each artificial time series has the same length as in the data (192 quarters from 1960Q1 to 2007Q4). Note that, in order to avoid any dependence on initial conditions, we actually simulated 292 observations and discarded the first 100. Then, for each set of simulated data, we estimated the same VAR as in actual data and applied our methodology to extract the various VAR-based shocks, or “factors,” and build their IRFs.

The second and the third row of Figure 18 show the median of the so-obtained distribution of IRFs for the JPT and ACD models, respectively. The comparison of these rows to one another and with the

⁶⁶We estimated this process using inflation data alone. This gave an estimate of 0.89 for the persistence parameter and 0.27% for the standard deviation of the innovation. All the other (real) parameters of the model were fixed at their values in the original article. Finally, the nominal interest rate was obtained directly from the Fisher equation, using the AR(1) process for inflation and the model’s prediction about the real rate.

first row (the data) corroborates the lesson obtained in the main text on the basis of the theoretical state-space representation of the two models: the factors in JPT are less interchangeable than their counterparts either in ACD or the data. The visual impression is corroborated by Table 26, which reports the metric discussed in the main text.

Figure 18: The MBC Shock

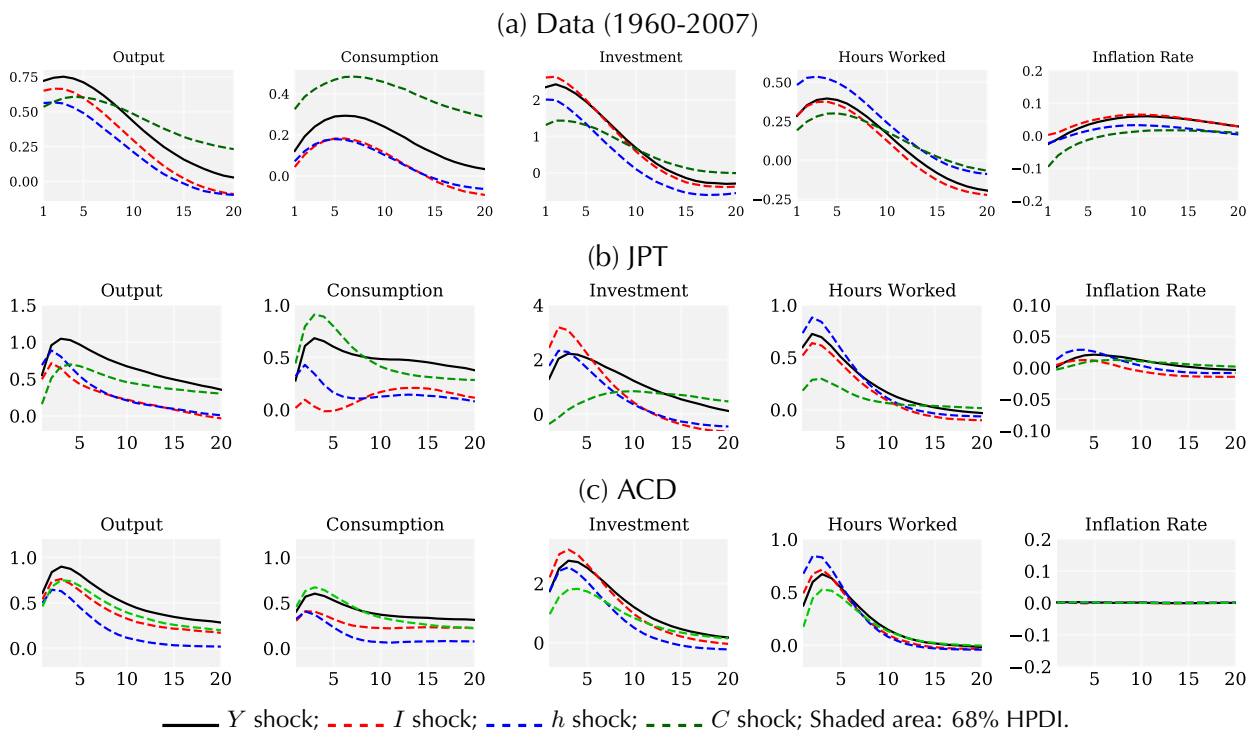


Table 26: Interchangeability of Factors, Simulated VARs

	Y	C	I	h	Average
Data	0.47	0.52	1.28	0.28	0.64
JPT	0.80	0.90	2.58	0.42	1.17
ACD	0.42	0.47	1.34	0.25	0.62

Note: The metric is the same as that in Table 10. A number closer to zero indicates a larger degree of interchangeability.

1.2 Re-estimating JPT/ACD

We now turn to the remaining two robustness exercises mentioned in Section 6.

First, in order to offer a proper comparison between JPT and ACD, we re-estimated the JPT model the same frequency-domain Bayesian technique used to estimate ACD. More precisely, the model is estimated over the business-cycle band of frequencies (6-32 quarters), using the levels of all variables,

and using the 1960-2007 data. This set of results is labeled *JPT - Freq. Domain* in the tables and figures that follow.

Second, we re-estimated both models using a minimum-distance estimation technique, with the parameters selected in order to minimize the distance between IRFs of output, consumption, investment and hours worked to the output, consumption, investment and hours worked factors over the horizon of 20 quarters (a set of 320 moments). Denoting by $IRF_{j,h}^i$ (resp. $\widetilde{IRF}_{j,h}^i(\Theta)$) the response of variable j to factor i at horizon h found in the data (resp. in the model) and $\sigma_{j,h}^i$ the variance of $IRF_{j,h}^i$, the vector of structural parameters Θ is found by solving the problem

$$\min_{\Theta} \sum_{i=1}^4 \sum_{j=1}^4 \sum_{h=1}^{20} \frac{(\widetilde{IRF}_{j,h}^i(\Theta) - IRF_{j,h}^i)^2}{\sigma_{j,h}^i}$$

Given our focus on the real IRFs, the parameters pertaining to the nominal part of JPT (Calvo probabilities, indexation parameters, parameters of nominal shocks) are not identified. We therefore set the values of these parameters to those estimated by JPT and re-estimated the parameters pertaining to the real side of the model (preferences, technology, adjustment costs, parameters of real shock processes). The relevant set of results is labeled *JPT - Matching Factors* and *ACD - Matching Factors*.

Figure 19 and Table 27, which extend Figure 6 and Table 10 from the main text, provide a comprehensive comparison of the dynamic properties of the two models under alternative specifications. The main findings are as follows. Re-estimating the JPT model in the frequency domain has a significant but still insufficient impact on the model's ability to reproduce the interchangeability of factors in the data. Re-estimating it by targeting the factors helps the model even more, but it still falls short of that in the data. Re-estimating the ACD by targeting the factors does not upset its already good performance, but it overshoots in the direction of producing too much interchangeability. All in all, the metric of how different the factors are is systematically greater for JPT than ACD, irrespective of the estimation method.

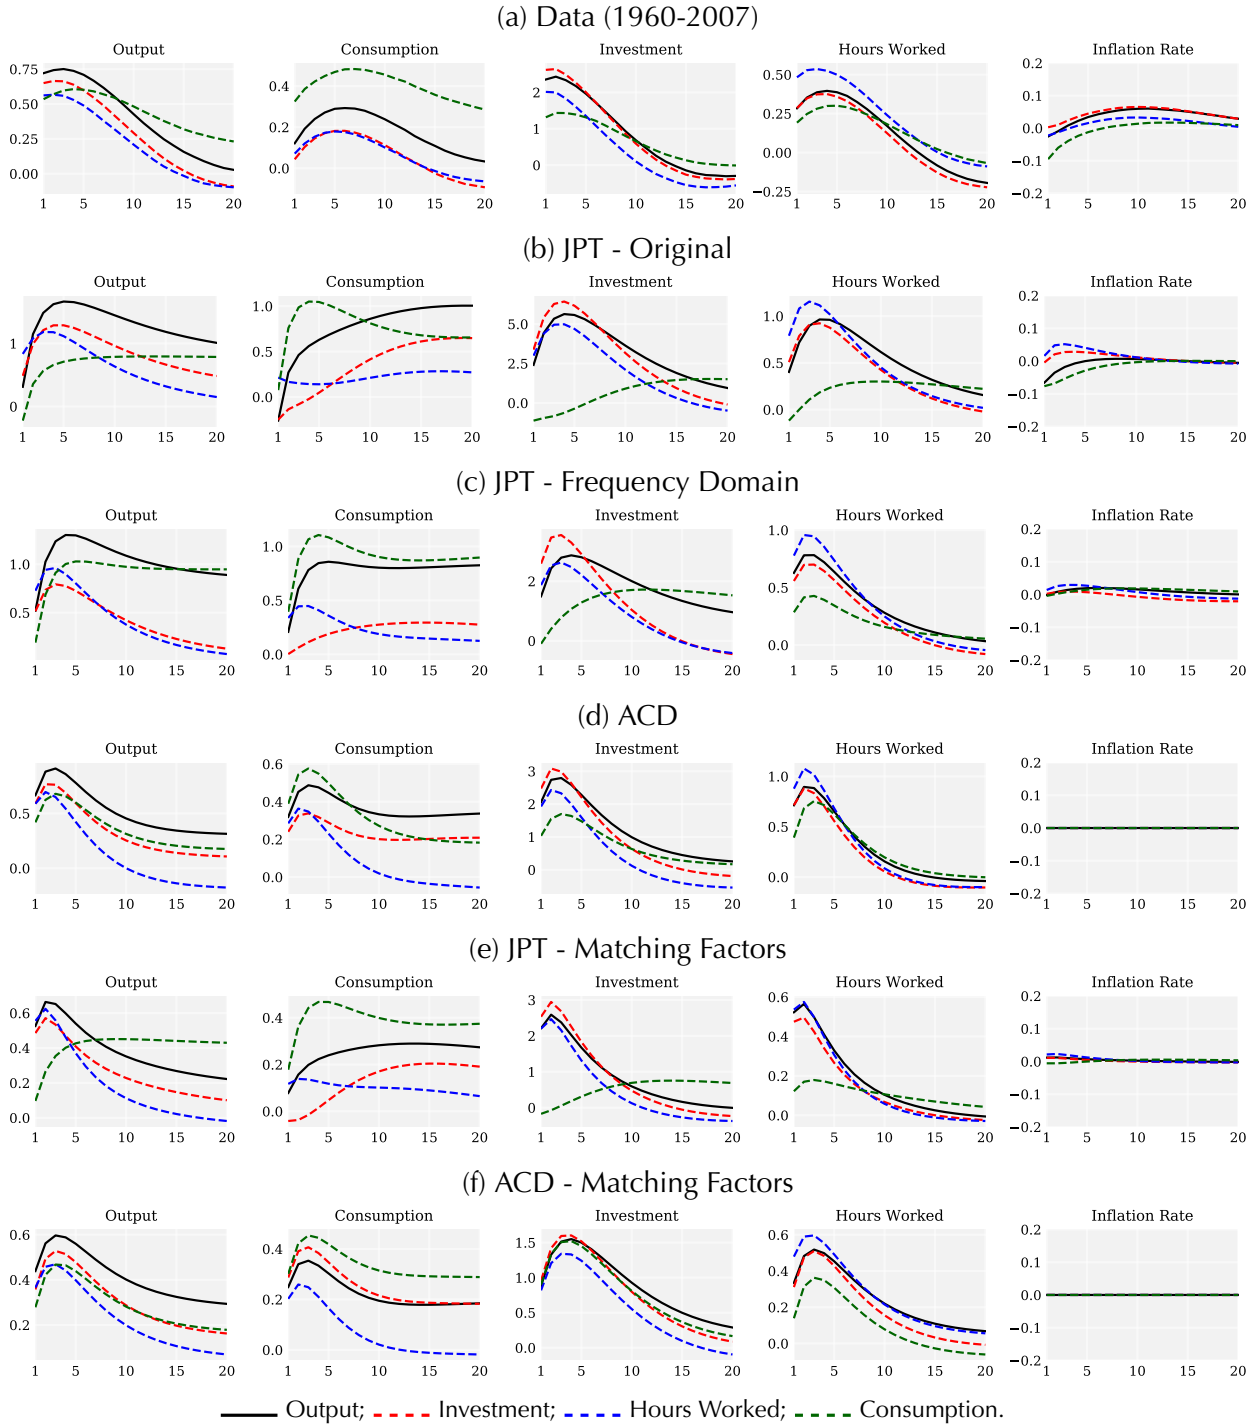
Table 27: Interchangeability of Factors

	Y	C	I	h	Average
Data (1960-2007)	0.47	0.52	1.28	0.28	0.64
JPT - Original	2.90	2.21	6.29	1.35	3.19
JPT - Freq. Domain	1.41	1.42	3.24	0.42	1.62
ACD	0.56	0.49	1.61	0.30	0.74
JPT - Matching Factors	0.56	0.51	2.26	0.27	0.90
ACD - Matching Factors	0.26	0.36	0.49	0.26	0.34

Note: The metric is the same as that in Table 10. A smaller number indicates greater interchangeability.

In conclusion, let us reiterate that the main goal of the application of our method to ACD and JPT is not to judge the superiority of one model over the other, but rather to illustrate the probing power of our method in the context of existing, medium-scale, DSGE models that have already been estimated and evaluated via other methods. This is best exemplified by the exercise conducted in the main text.

Figure 19: Comparing Business-Cycle Factors



The second robustness exercise in this appendix serves a complementary objective, namely to inform on whether is at all possible for these models to be replicate the propagation mechanism we observe in the data and, if so, what this requires in terms of their parameters. In short, the two exercises illustrate two different ways in which our anatomy of the data can inform theory.

J The Secondary Business Cycle Shock

For each of the five macroeconomic quantities, $X \in \{u, Y, h, I, C\}$, we now identify *two* shocks. The first shock is the one already reported in the main text: it is obtained by maximizing its contribution to the business-cycle volatility of that variable. The second shock is obtained by maximizing its contribution to the residual, business-cycle volatility of the targeted variable after filtering out the effect of the first shock. This procedure produces a collection of five new shocks, one for each of the macroeconomic quantities of interest.

Figure 20 reports the IRFs to these shocks and Table 28 their variance contributions. The IRFs are nearly the same, suggesting that these shocks, too, represent interchangeable facets of one shock—the “secondary” business cycle shock, or SBC for short.

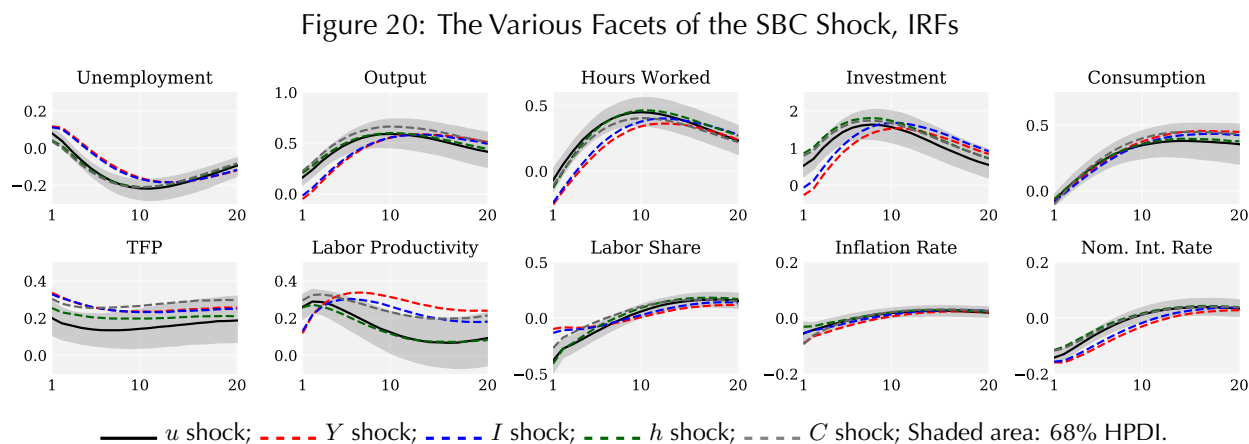


Figure 20 also reveals that the impact of the SBC shock on the economy builds up slowly over time, peaking after several quarters. By contrast, the impact of the MBC shock peaks within a year and fades shortly after. Furthermore, the SBC shock contains relatively more information about TFP and the prospects of the economy in the medium to long run. In this sense, whereas the MBC shock fits the profile of a quick-moving demand shock, the SBC shock fits the profile of a slow-moving supply shock. Both shocks, however, have a similar, zero to weakly positive, effect on inflation. Neither of them therefore fits easily in the traditional AD-AS framework.⁶⁷

In the main text, we focused on the MBC shock as the main probing tool of our anatomy and treated the SBC shock as part of the residual. While the SBC shock represents subsidiary rather than

⁶⁷Of course, any attempt to offer a structural interpretation of the MBC and SBC shocks, either jointly or in isolation, faces the basic challenge discussed in detail in Section 5 that such objects could be different combinations of multiple theoretical shocks, none of which fits the profile of either of these empirical objects. The aforementioned interpretation is therefore possible but not necessary.

Table 28: Variance Contributions of Second Largest Shocks

Target	u	Y	h	I	C
Unemployment	24.50 [18.61,31.29]	21.49 [15.33,28.89]	23.76 [17.41,30.40]	19.73 [13.66,27.51]	21.41 [15.01,29.22]
Output	23.94 [17.80,30.83]	18.54 [12.56,25.71]	24.36 [17.10,32.15]	17.63 [11.55,24.49]	23.08 [15.98,31.39]
Hours Worked	21.11 [14.99,28.76]	22.59 [15.45,30.75]	28.07 [21.65,34.19]	24.56 [17.12,33.76]	21.92 [15.33,29.96]
Investment	24.00 [17.82,30.91]	17.99 [11.98,25.21]	26.08 [18.77,33.67]	18.41 [12.18,25.47]	22.23 [15.44,30.75]
Consumption	20.90 [14.06,28.17]	26.19 [18.58,35.60]	22.23 [15.77,29.44]	22.90 [14.63,31.78]	28.00 [21.31,34.84]
	TFP	Y/h	Wh/Y	π	R
Unemployment	5.71 [2.16,11.87]	19.84 [14.26,26.43]	26.44 [17.47,36.69]	11.08 [4.38,20.46]	38.49 [28.16,48.05]
Output	17.79 [9.10,29.06]	17.16 [11.33,24.25]	6.46 [2.16,13.23]	20.55 [10.19,33.00]	53.85 [43.53,62.99]
Hours Worked	10.15 [4.72,18.29]	17.60 [11.30,24.42]	29.44 [19.92,39.71]	8.10 [3.27,16.17]	28.54 [18.83,39.20]
Investment	17.52 [8.56,28.84]	16.37 [10.67,23.05]	8.59 [3.25,16.75]	12.14 [4.31,23.20]	51.56 [40.34,61.00]
Consumption	13.48 [6.43,22.85]	21.38 [15.09,28.84]	15.68 [7.04,25.63]	21.52 [12.81,31.48]	30.63 [19.76,41.06]

primary variation in the data, it can still serve as an additional or complementary validation tool in exercises like those conducted in Section 6. For instance, consider Figure 21. This figure redoes Figure 6 for the SBC shock in place of the MBC shock. That is, it compares the various second largest shocks to the corresponding objects in JPT and ACD, the models considered in Section 6. Clearly, JPT does rather poorly vis-a-vis the SBC shock, too. But now this failure is shared by ACD.

On the one hand, these findings confirm the validation prowess of our empirical strategy. On the other hand, they serve as an additional warning that the postulated propagation mechanisms of state-of-the-art models, even of the most successful ones, remain crude representations of the propagation mechanisms that best characterize the data.

Figure 21: The SBC in the Data and the Models

