

CEWP 22-06

**Split Personalities: The Changing Nature
of Technology Shocks***

Christopher Görtz
University of Birmingham

Christopher Gunn
Carleton University

June 12, 2022

CARLETON ECONOMICS WORKING PAPERS



Carleton
UNIVERSITY

Department of Economics

1125 Colonel By Drive
Ottawa, Ontario, Canada
K1S 5B6

Split Personalities: The Changing Nature of Technology Shocks *

Christoph Görtz
University of Birmingham

Christopher Gunn
Carleton University

June 2022

Abstract

This paper analyzes the nature of technology shocks and documents important changes in their propagation over time. We employ a vector-autoregression and identify a shock that explains the maximum variation in total factor productivity (TFP) at a long finite horizon. This agnostic identification suggests that the dominant shock driving TFP is not necessarily a surprise shock, but exhibits features consistent with a shock that is anticipated or diffuses over time: GDP and consumption rise prior to any significant increase in TFP. We further find that shock transmission has changed over time. In a sample that ends in the mid 1980s, the shock triggers a decline in hours-worked and inventories, and a rise in credit spreads. In a post-Great Inflation sample the response of these variables is reversed and the shock generates an outright expansion in hours, inventories, GDP and consumption that is accompanied by a decline in credit spreads. We find that the importance of technology shocks as a major driver of aggregate fluctuations has increased over time — they play a dominant role in the second subsample, but much less so in the first.

Keywords: technology shocks, total factor productivity, business cycles, shock transmission.

JEL Classification: E2, E3.

*Görtz: University of Birmingham, Department of Economics, Email: c.g.gortz@bham.ac.uk. Gunn: Carleton University, Department of Economics, Email: chris.gunn@carleton.ca.

1 Introduction

Since at least the onset of the era of modern macroeconomics, the idea of stochastic shifts in the technological frontier of the economy as driver of business cycles has had a very prominent role in macroeconomics. These so-called “technology shocks” have remained controversial since their inception, and in tandem with work studying their role in theoretical models, empirical researchers have sought out evidence about their potential prominence and role in the data. Do shifts in technology trigger a response that resembles business cycles? If so, how important are such disturbances for explaining aggregate fluctuations? Does it matter if these technological shifts are surprise shocks or anticipated in advance? Not surprisingly, the empirical literature addressing these questions has at times arrived at dramatically different answers.

While much of the literature has focused on exploring implications of different time-series treatments and identifications, in this paper we take a step back and explore the role of technological era. Using a standard VAR-based Max Share technology shock identification, we show that analyzing two separate subsamples created by splitting the data around the generally considered onset of the Great Moderation yields dramatically different results, and these results are remarkably robust to identification method and data treatment. We can characterize the general results over these two subsamples as follows: (i) the most relevant shock driving TFP is not necessarily a surprise shock as assumed in many models, but rather a news/diffusion shock; (ii) the importance of technology shocks has increased over time, and, (iii) the importance of technology shocks in terms of business cycles has increased over time.

With respect to our first general result above, following Kurmann and Sims (2021), our Max Share empirical shock identification seeks out the shock that “best explains” the variance in TFP at some long but finite horizon, and makes no attempt to impose any sort of additional short-run restriction in order to separately identify the surprise versus anticipated (“news”) component of the technology shocks. The identification thus remains agnostic about

the presence of surprise or news components in technology shocks, allowing us to address more generally the debate on the nature of technology shocks. Nevertheless, in line with the results of Kurmann and Sims (2021) over a single sample, in each of our subsamples TFP only rises with statistical significance after several periods, and then grows gradually beyond that, consistent with the idea of “anticipated/news” shocks, or technological diffusion. In this sense we conclude that anticipated/diffused technological growth is the dominant form of the technological shock over both samples.

With respect to our second general result above, the change in the transmission of technology shocks is best reflected in the striking difference in the response of hours-worked across the two subsamples: in the first subsample hours falls on impact; in the second subsample it rises. Yet consumption and stock prices rise consistently in both samples. Moreover, although hours responds differently in both samples, it co-moves positively with investment, inventories, the real wage and negatively with the BAA spread in both samples. As a group then, the response of hours, investment, inventories, the real wage and the BAA spread flips over the two samples relative to the consistent rise in consumption and stock prices over both subsamples. Interestingly, this connection between hours and inventories in particular is consistent with the literature that suggests a tight relationship between hours and inventories (and other variables, e.g. spreads) and argues for these to be assessed in conjunction.

Finally, with respect to our third general result above, about the increase in importance of TFP shocks over time, our forecast error variance decompositions show that while the identified shock explains a large and similar share of TFP over the two subsamples, the shock explains a substantially larger share of output variations in the second subsample than the first. This result is also related to our finding that the response of hours and other key variables in the second subsample, conditional on the identified shocks, is consistent with the unconditional correlations of those variables in the data. Said another way, the negative comovement of hours and consumption in the first subsample makes it difficult for the shock to account for a large proportion of business cycle activity when unconditionally hours and

consumption co-move positively.

To attempt to isolate the source of the change in the response over the sub-samples in the VAR, we perform a counterfactual exercise that “re-recovers” the technology shocks in the first subsample using the polynomial lag coefficients estimated from the first subsample but the variance-covariance matrix estimated from the second sample. Similarly, we re-recover the shocks in the second subsample using the polynomial lag coefficients estimated from the second subsample but variance-covariance matrix estimated from the first. The results are striking: the impulse response functions are largely unchanged from our core results, suggesting that potential structural change in the variance-covariance matrix is not driving change in the results over the sample. Rather, the exercise points toward changes in the polynomial lag coefficients of the underlying VAR.

Our work links to an ongoing literature that focuses on the importance of the long run to identify technology shocks in VARs. Galí (1999) employs long-run restrictions on labor productivity to identify technology shocks and finds a decline in hours-worked. Technology shocks account just for a very small part of total fluctuations in output and hours-worked at business cycles frequencies which is taken as evidence against the Real Business Cycle paradigm.¹ Others including Christiano et al. (2004) find the opposite result with regards to the response of hours-worked and the importance of technology shocks for aggregate fluctuations, which was attributed to a differences in the specification of hours in the VAR.² Following this debate, another strand of the literature emerged which focused on alternative identification. Francis et al. (2014) propose the so-called, Max Share identification which identifies a technology shock as the one that that maximizes the forecast-error variance of labor productivity at some long by finite horizon, and which addresses some of the shortcomings of long-run identification. In particular, Francis et al. (2014) show that the Max Share identification outperforms standard long-run restrictions by significantly reducing the bias in the short-run impulse responses and raising their estimation precision. They find —

¹See also Shea (1998), Ramey (2005), Pesavento and Rossi (2005) and Basu et al. (2006).

²See also Uhlig (2004) and Dedola and Neri (2007).

independently of the data specification considered for hours-worked — that hours respond negatively. Notably, Francis et al. (2014) derive their results from a single sample using 1948Q2-2009Q4.³ We build on the insights of this debate and employ the Max Share identification but focus on the analysis of two distinct subsamples for which the literature has documented differences in unconditional time series behavior. While we use TFP instead of labor productivity for our core analysis, we also show that our split-sample result for hours-worked holds using labor productivity instead of TFP.

Our work also connects with the ongoing empirical literature that studies anticipated shocks to technology, typically framed as TFP news shocks. This literature has for several years debated the response of key economics variables to TFP news, and as with the Galí (1999) debate discussed above, the response of hours-worked to the identified news shock has been a key feature of this debate. Some studies, e.g. Kurmann and Sims (2021) and Barsky and Sims (2011) (both with sample period 1960q1 to 2007q3), find that hours-worked do not co-move with output and consumption but decline in response to favorable anticipated technology shocks. Others document a broad-based expansion of macroeconomic aggregates — see e.g. Görtz et al. (2021) and Görtz et al. (2019) who consider 1984:Q1–2017:Q1 and 1983Q1-2018Q2 samples, respectively, which closely correspond to the second subsample in our paper. The differences in findings with regards to the response of hours-worked is important as it speaks to the notion of whether anticipated technology shocks are potentially important drivers of aggregate fluctuations. In relation to this, co-movement of macroeconomic aggregates has also been an important criterion for news-shock models.

We also speak to the large literature that documents differences in time series behavior across the Great Inflation/Great Moderation samples.⁴ While this literature documents the data unconditionally, we point to important changes conditional on technology shocks. This literature and our work has implications for the estimation of structural models —

³For a further contribution to the methodological debate on shock identification, see also Fève and Guay (2009) who document a decline in hours-worked over a 1948Q1–2003Q4 sample.

⁴We cannot do justice here to this extensive literature, see e.g. McCarthy and Zakrajsek (2007), Kahn et al. (2002) and Sarte et al. (2015).

in particular in relation to technology shocks. We speak to the relevance of subsample estimation or estimation with time varying parameters. See e.g. Fuentes-Albero (2019) who documents that contemporaneous to the Great Moderation there was a widespread increase in the volatility of financial variables. She comments on changes in the transmission of financial shocks. Cúrdia and Finocchiaro (2013) show that ignoring regime changes leads to spurious estimates.

The remainder of the paper proceeds as follows. In Section 2.1 we discuss our empirical methodology and the data. Sections 2.2-2.4 we take a break-date around the Great Moderation as given and analyze the two separate subsamples using a minimally specified VAR framework, explore robustness along a number of dimensions and and perform various empirical exercise to try and isolate and understand the source of this technology change in the role of technology. Section 3 concludes.

2 A tale of two eras

We begin by providing some VAR-based evidence about the importance of subsample era to the role and response of the macroeconomy to technology shocks. To make our point most clearly, we keep our analysis as simple and direct as possible, focusing on a small VAR with a relatively agnostic identification using two different subsamples. We then discuss the implications of these results, and provide an initial first-pass analysis of the source of the changes over subsample era.

2.1 Empirical Methodology and Data

Our identification objective is to isolate broadly-defined technology shocks and we want to be agnostic about whether technology instantaneously reacts to the shock or with a lag.

Like much of the literature, we focus on a identification condition at a long horizon based on the idea that a distinguishing feature of a technology shock is its ability to influence the behaviour of the macro-economy at long-horizons. As such, we identify the technology shock

using the Max Share methodology as suggested in Francis et al. (2014), who maximize the forecast error variance share of a productivity measure at a long but finite horizon.⁵ As in Francis et al. (2014), we consider this horizon h at which the forecast error variance is maximised to be 10 years. This approach is consistent with suggestions in Uhlig (2003) and in the spirit of Angeletos et al. (2020). Following Kurmann and Sims (2019), we use TFP as the target variable, such that identification isolates the shock that best-explains TFP at a long horizon. As in Kurmann and Sims (2019), we do not impose any additional restrictions intended to separate anticipated from surprise shocks to technology (such as a no-impact orthogonality restriction). As argued by Kurmann and Sims (2019), doing so helps to avoid measurement issues that may arise with a variable like TFP in the short-run. Moreover, it also allows us to put the least possible restrictions on our identification, thereby increasing the scope of our subsample dependence result. As such, the identification allows us to remain agnostic about the type of technology shock being identified (anticipated vs. surprise), and does not require us to make the strong assumption that TFP is completely exogenous at all horizons and comprised of just surprise and news shocks.

We include five variables in our baseline VAR model: TFP, GDP, consumption, hours-worked and the S&P500. A key measure to identify the shock that moves productivity is an observable for TFP. We use the TFP measure provided by Fernald (2014) which is based on the growth accounting methodology in Basu et al. (2006) and corrects for unobserved capacity utilization. GDP, consumption and hours-worked serve as our measures of economic activity, and the S&P500 serves as a forward-looking capturing information available to economic agents about future macroeconomic developments, helping to avoid non-invertibility issues. The GDP, consumption and hours-worked are all seasonally adjusted and in real per-capita terms (except for hours-worked which are not deflated). Appendix C provides details on the data sources and all used time series. The time series included in the VAR enter in levels, consistent with the practice in the empirical VAR literature (e.g. Barsky and

⁵Francis et al. (2014) show that in comparison to other long-run identification schemes, the Max Share approach's focus on a long and finite horizon helps reducing small-sample bias in VARs.

Sims (2011), Francis et al. (2014)). To estimate the VAR we use three lags with a Minnesota prior and compute confidence bands by drawing from the posterior.⁶

There is wide agreement in the literature that the structure of the US economy changed during the 1980s — what we now call the end of the Great Inflation and the onset of the Great Moderation — which resulted in substantial unconditional changes in time series behavior. We frame our investigation around the two subsamples on either side of the onset of the Great Moderation, estimating a VAR on U.S. data separately for each of two subsamples spanning the periods 1954Q2–1983Q4 and 1984Q1–2019Q4. This subsample horizon is guided by the literature that documents differences in cross-correlation patterns of several macro-aggregates in samples before and after the mid-1980s. In particular, McConnell and Perez-Quiros (2000) and Kim and Nelson (1999) document a structural break at the first quarter of 1984 (see also e.g. Galí and Gambetti (2009) and Stock and Watson (1999) for further evidence on this structural break).

2.2 Evidence from Two Eras

Figure 1 shows impulse response functions (IRFs) to our identified technology shock with the red and blue lines corresponding to the first and second subsamples respectively. There are several important points to note. First, while our agnostic shock identification does not exclude the possibility that TFP jumps on impact, in both subsamples, the dominant effect on TFP is one that grows over time. In particular, in both subsamples TFP only rises significantly with a lag of eleven quarters and after the other variables in the VAR. This is consistent with a diffusion-based or anticipated (news) technology shock. Second, there is a striking difference in the co-movement of the key aggregate variables between the two subsamples. Whereas in the more recent subsample we see a broad-based and positively co-moving expansion of GDP, consumption and hours-worked, in the earlier subsample hours-

⁶Further details about the VAR model, the Max Share identification and prior specifications are provided in Appendix A.

worked fall.⁷ Consumption rises also in the first subsample, yet its short- and medium-run expansion is less pronounced than that in the second subsample. For GDP this disparity is even more apparent as output rises in the first subsample significantly only after seven quarters. Finally, stock prices rise in both subsamples. This rise in stock prices along with that of consumption over the two subsamples is generally consistent with a “good news” technological expansion, despite the differential response of hours-worked between the subsamples. Overall, we observe for almost identical TFP responses a marked difference in the response of the other variables over the two episodes.⁸




Figure 1: **IRF to TFP shock.** First subsample 1954Q2-1983Q4 (red), second subsample 1984Q1-2019Q4 (blue). The solid line is the median and the shaded colored areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

Figure 2 shows the forecast error variance decompositions relating to the estimated VARs for the two subsamples. While the identified shock explains a substantial and very similar share of variation in TFP across the two episodes, in the first subsample the shock is of substantially lower importance for fluctuations in GDP at business cycle frequencies (red lines, approximately 10-55%) than in the second subsample (blue lines, approximately 70-85%). The rise in the shock’s importance for business cycle fluctuations in the second subsample is consistent with the IRF evidence from Figure 1, where we observed a stronger shock propagation and comovement across all macroeconomic aggregates, including hours-

⁷The qualitative differences across subsamples with respect to hours-worked is reflected in the labor market overall. Consistent with the decline in hours during the first subsample, Appendix B documents a decline in the labor force participation rate and a rise in the unemployment rate. In contrast, for the second subsample, the labor force participation rate increases and the unemployment rate declines.

⁸These impulse response functions are robust to using labor productivity as an alternative measure for productivity. Details are documented in Appendix B. Our results are also robust to alternating the number of lags and to variations in the Max Share horizon h . Results are available upon request.

worked. The opposite sign response of hours-worked – relative to GDP and consumption – is consistent with the notion that the technology shock in the first subsample is of lesser importance for business cycle fluctuations.




Figure 2: **Forecast Error Variance Decomposition — share explained by the TFP shock.** First subsample 1954Q2-1983Q4 (red), second subsample 1984Q1-2019Q4 (blue). The solid line is the median and the shaded colored areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

In summary, the above results suggest that: (1) The importance of technology shocks has increased over time — as a major driver of aggregate fluctuations they play a dominant role in the second subsample but less so in the first; (2) the transmission of technology shocks has changed over time, especially with regards to the qualitative response of hours-worked; (3) the most relevant shock driving TFP is not necessarily a surprise shock as assumed in many models, but rather a news/diffusion shock. We will discuss the implications of these findings further in the next section which investigates the shock transmission in more detail.

2.3 Digging Deeper: Subsample Differences in Shock Transmission

The above section documents differences in the transmission of TFP shocks over two subsample eras, most significantly manifested in the response of hours-worked. Developments in the labor market are often tightly linked to other key margins. In this section, we inspect these to gain a deeper understanding for differences in the shock transmission across the two subsamples.

Figure 3 shows responses of multiple variables of interest for the transmission of TFP shocks. Subplots in this figure are from a VAR with TFP, GDP, consumption, hours-worked,

the S&P 500 and one of the plotted variables of interest at a time. The plotted response for hours is from the VAR that includes inventories. The variables not shown are very similar to those in Figure 1.

In addition to the responses of hours-worked, a number of other variables also display remarkable differences across the two subsamples in their response to a TFP shock. In particular, inventories, investment and the real wage fall, and the BAA spread rises in the first subsample, whereas in the second subsample, the behaviour is reversed. In addition, there is a short-lived decline in inflation in both subsamples. The patterns of the remaining two variables are less certain: the federal funds rate doesn't respond significantly in either subsample, and capital utilization rises in the second subsample, but is insignificant in the first.

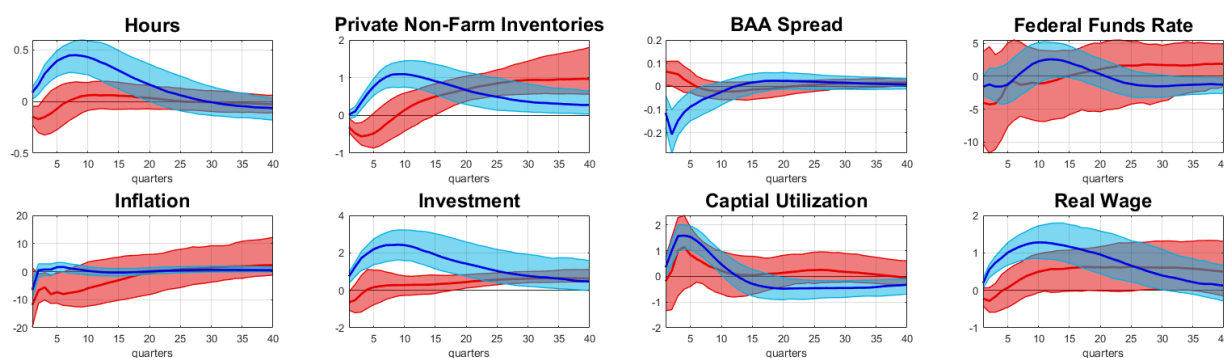


Figure 3: **IRF to TFP shock.** First subsample 1954Q2-1983Q4 (red), second subsample 1984Q1-2019Q4 (blue). The solid line is the median and the shaded colored areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations. Subplots are based on a VAR with TFP, GDP, consumption, hours-worked, the S&P 500 and one of the plotted variables at a time.

Taking together, the results from Figures 1 to 3 suggest the following with regards to the behaviour of the key variables in response to the technological shock. First, consumption and stock prices rise and inflation falls in both subsamples. This rise in consumption and stock prices in tandem with the delayed rise in TFP is consistent with the idea of “good news” associated with a rise in lifetime wealth due to expected TFP growth (see e.g. Beaudry and Portier (2006)). Moreover the short-lived decline in inflation is a widely reported response

to technology news shocks (see e.g. Barsky and Sims (2011) Kurmann and Sims (2019), Görtz et al. (2022)). Second, hours-worked, investment, inventories, the real wage and the BAA spread co-move in a consistent way *with each other* over both samples – and indeed, consistent with their unconditional correlations in the data – however, as a group, their response flips between the two subsamples. In particular, as a group, these variables respond in the short run in a “contractionary” way in the first subsample, and “expansionary” in the second subsample. This is also consistent with the somewhat more muted response of output in the first subsample relative to that in the second subsample, reported in Figure 1.

2.3.1 Group mentality: Labour, inventories, investment and credit spreads

The second observation made in the paragraph above is suggestive of a potential connection between developments on the labor market, inventories, investment and credit spreads. The close relationship between hours and inventories has been stressed for example by Macchini and Rossana (1984) and Galeotti et al. (2005), who point out the need for a joint understanding of the dynamics of inventories and hours-worked. Also Chang et al. (2009) emphasize this point and document the co-movement of inventories and employment conditional on (unanticipated) technology shocks. They further stress the connection between the sign of the employment response to technology shocks and the cost of holding inventories. Their notion that a positive response of hours-worked is more likely the less costly it is to hold inventories, is consistent with the patterns we document in Figure 3 on inventories, hours and credit spreads. Risk premia, such as credit spreads, have been recognised in the literature also as a measure for the opportunity cost of holding inventories. See for example Jones and Tuzel (2013) who document this relationship between risk premia and inventories unconditionally and Görtz et al. (2019) who stress the importance of credit spreads as opportunity cost for inventory holdings conditional on anticipated technology shocks. Hence, the decline (rise) in inventories shown in Figure 3 for the first (second) subsample is consistent with a rise (fall) in their opportunity cost captured by credit spreads.

A vast body of research finds that financial markets are characterized by frictions that

lead to credit spreads and hence affect the financing of investment projects.⁹ In particular, Görtz and Tsoukalas (2018) and Görtz et al. (2021) emphasize that the empirical relevance of technology news shocks hinges crucially on the shock’s transmission being amplified by frictions in financial markets. The responses of investment and the BAA spread shown in Figure 3 are consistent with this finding in so far as the response of the BAA spread indicates a much stronger transmission via financial markets in the second subsample. This and the relaxation of credit frictions, as indicated by the decline of the BAA spread, is consistent with the strong expansion in investment we document for the second subsample.¹⁰ In contrast, the somewhat muted rise of credit spreads in the first subsample is indicative of tighter lending conditions which is consistent with the somewhat less pronounced rise in investment.

Changes in the nature of US business cycles during the mid-1980s are a widely documented phenomenon. By considering two separate subsamples we take account of this finding and avoid masking differences in shock transmission across the two subsamples. Estimating the VAR over the entire sample (1954Q2-2019Q4) yields responses that are similar to those of the second subsample. Details are provided in Appendix B.

2.3.2 Conditional Evidence and Unconditional Dynamics in the Data

Our sample split coincides with the end of the Great Inflation and the literature has documented a number of structural changes in the economy that occurred around this time. Interestingly, these structural changes would be reflected in some of those variables that we find to depict the most substantial differences in responses across subsamples, i.e. inventories, hours-worked and credit spreads. McCarthy and Zakrajsek (2007) and Kahn et al. (2002) document that significant changes in inventory dynamics occur in the mid-1980s due to improvements in inventory management. Sarte et al. (2015) document that time-series properties of inventories and hours have changed with the onset of the Great Moderation

⁹See for example Philippon (2009) and Gilchrist and Zakrajsek (2012).

¹⁰Görtz et al. (2021) stress the importance of movements in credit spreads for the propagation of anticipated technology shocks. They show that such a favorable shock is amplified via financial markets since an endogenous strengthening of banks’ balance sheets relaxes lending conditions associated with a decline in credit spreads.

and attribute this, at least partly, to variations in credit market frictions. Adrian et al. (2010) and Jermann and Quadrini (2012) argue that the importance of the financial sector for the determination of credit and asset prices has risen significantly from the mid-1980s. Further, Jermann and Quadrini (2009) discuss a variety of financial innovations that were taking place or intensified in the 1980s — including banking liberalization, and flexibility in debt issuance through the introduction of the Asset Backed Securities market — and stress their role for a slowdown in output volatility. Fuentes-Albero (2019) documents that contemporaneous to the onset of the Great Moderation there was a widespread increase in the volatility of financial variables. This literature studies the unconditional dynamics of inventories, hours and credit spreads in relation to potential sources for the end of the Great Inflation. While our paper does not aspire to speak to the reasons for the onset of the Great Moderation, we note that there might potentially be a link between the sources of structural change — i.e. improvements in inventory management and developments in financial markets — that have been attributed to be potential sources of the Great Moderation and our documented changes in the transmission of technology shocks.¹¹ The following section builds on our econometric setup to provides some first insights on potential sources of the subsample differences conditional on technology shocks.

2.4 Exploring the Source of Subsample Differences: Impulse or Propagation?

Our results above suggest that not only have technology shocks played more of a role in accounting for aggregate fluctuations over time, but their impact on the macroeconomy has also changed. While the former effect on its own could simply reflect some change in a feature of the technology shock itself, the latter result however is more suggestive of a

¹¹Other factors that have been suggested to contribute to the end of the Great Inflation are changes in monetary policy making and smaller shocks. While this paper does not attempt to speak to this debate on unconditional changes in time series behavior, it is interesting to note that our results suggest that the transmission of technology shocks actually resulted in larger, rather than smaller, fluctuations in macroeconomic aggregates in response to technology shocks in the second subsample. The insignificant response of the federal funds rate is indicative for the limited role of changes in monetary policy in context of our conditional responses to technology shocks.

change in some underlying feature of the macroeconomy. We now take a first-pass at trying to understand the reason for this change within the context of our econometric setup.

As we show in detail in Appendix A, our econometric approach considers the following vector autoregression (VAR), which describes the joint evolution of an $n \times 1$ vector of variables y_t :

$$y_t = A(L)u_t.$$

$A(L) = I + A_1L + \dots + A_pL^p$ is a lag polynomial of order p over conformable coefficient matrices $\{A_p\}_{i=1}^p$. u_t is an error term with $n \times n$ covariance matrix Σ . We assume a linear mapping between the reduced form errors u_t and the structural errors ε_t :

$$u_t = B_0\varepsilon_t,$$

where B_0 is an identification matrix. We can then write the structural moving average representation of the VAR:

$$y_t = C(L)\varepsilon_t,$$

where $C(L) = A(L)B_0$, $\varepsilon_t = B_0^{-1}u_t$, and the matrix B_0 satisfies $B_0B_0' = \Sigma$. B_0 can also be written as $B_0 = \tilde{B}_0D$, where \tilde{B}_0 is any arbitrary orthogonalization of Σ and D is an orthonormal matrix such that $DD' = I$.

Thus through the lens of our structural moving-average representation in equation (3), the subsample differences can be driven by: (i) differences in the polynomial lag matrix $C(L)$, (ii) differences in the variance-covariance matrix associated with ε_t , which in turn results from differences in the estimates in the variance-covariance matrix Σ . We test for this as follows: We draw from the posterior coefficient matrix based on the reduced form VAR estimated for each of the two subsamples (we use the same seed for the random number generator). We then identify the TFP shock for the first subsample (as outlined in Section 2.1 and Appendix A.1) using the second-subsample polynomial-lag coefficients and the first-subsample variance-covariance matrix. Similarly, we identify a TFP shock for the second

subsample, using the first-subsample polynomial-lag coefficients and the second-subsample variance-covariance matrix.

Figure 4 shows the results of this exercise. The red shaded areas shown in the first row are the IRFs based on the first subsample. The blue shaded areas in the second row are the IRFs based on the second subsample. These shaded areas are congruent with those shown in Figure 1 and are used as a point of reference. The blue dashed and dotted lines in the first row show the median and posterior bands if the shock is identified using first-subsample polynomial-lag coefficients and second-subsample variance co-variance matrix. Similarly, the red lines in the second row of Figure 4 show the responses if the shock is identified using the second-subsample polynomial-lag coefficients and the first-subsample variance co-variance matrix. It is striking from the first row that if we identify the shock using polynomial-lag coefficients that are consistent with the first subsample and a second-subsample variance co-variance matrix, the resulting IRFs are extremely similar to the original first subsample responses. The same holds vice versa for the second row. This implies that the documented differences across subsamples are driven to a large extent by differences in the polynomial-lag coefficients, rather than differences in the variance co-variance matrices. This is indicative of a role for differences in the shock's transmission through the economy across the two subsamples.

3 Conclusion

While not as far-reaching as once advocated in the 1980s, technology shocks continue to play an important role in our understanding of aggregate fluctuations. Dis-satisfaction with the idea and plausibility of unexpected high-frequency technology shocks – especially negative shocks – lead researchers in the early 2000's to study whether technology could still play a role in the absence of surprise shocks and technological regress. Beaudry and Portier (2006) showed how a business cycle boom-bust could result in such an environment when the driving impulse was changes in expectations about future positive shifts in technology




Figure 4: **IRF to TFP shock.** First row: subsample 1954Q2-1983Q4 responses (red). Second row: second subsample 1984Q1-2019Q4 responses (blue). The solid line is the median and the shaded colored areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The blue (red) dashed and dotted lines in the first (second) row show the median and 16% and 84% posterior bands if the shock is identified using the beta coefficients implied by the first (second) subsample and the variance covariance matrix implied by the second (first) subsample.

rather than surprise changes in technology itself, and a vibrant literature was launched to study the importance and role of such “news shocks”.

In this paper we add to the empirical literature attempting to understand the role and importance of technology shocks. We take an agnostic view of the presence of surprise versus anticipated shocks, using a well-established empirical identification that seeks to best account for the variation in TFP at some far out but finite horizon. Rather than using a single sample as much of the work to date, we split our sample at the onset of the Great Moderation and study each sample independently. Our results suggest that the qualitative response of TFP is consistent with a dominant anticipated or diffused shock, that the importance of TFP shocks has increased over the sub-samples, and that the transmission of the shocks into the broader economy has changed.

This change in the transmission is manifested most clearly in the response of hours-worked: hours falls in the first subsample, but rises in the second, despite consumption and stock prices rising consistently in both subsamples. Moreover, despite its differential response over the two subsamples, hours co-varies in a consistent way with investment, inventories,

the real wage, and the credit spread over both subsamples. Studying these puzzles is an important next step both for understanding technology specifically and aggregate fluctuations more generally.

References

- Adrian, T., Moench, E., and Shin, H. S. (2010). Financial intermediation asset prices and macroeconomic fundamentals. *Federal Reserve Bank of New York Staff Report*, (422).
- Angeletos, G.-M., Collard, F., and Dellas, H. (2020). Business-cycle anatomy. *American Economic Review*, 110(10):3030–70.
- Barsky, R. B. and Sims, E. R. (2011). News shocks and business cycles. *Journal of Monetary Economics*, 58(3):273–289.
- Basu, S., Fernald, J., and Kimball, M. (2006). Are technology improvements contractionary? *American Economic Review*, 96(5):1418–1448.
- Beaudry, P. and Portier, F. (2006). News, stock prices and economic fluctuations. *The American Economic Review*, 96(4):1293–1307.
- Chang, Y., Hornstein, A., and Sarte, P.-D. (2009). On the employment effects of productivity shocks: The role of inventories, demand elasticity, and sticky prices. *Journal of Monetary Economics*, 56(3):328 – 343.
- Christiano, L. J., Eichenbaum, M., and Vigfusson, R. (2004). The Response of Hours to a Technology Shock: Evidence Based on Direct Measures of Technology. *Journal of the European Economic Association*, 2(2-3):381–395.
- Cúrdia, V. and Finocchiaro, D. (2013). Monetary regime change and business cycles. *Journal of Economic Dynamics and Control*, 37(4):756–773.
- Dedola, L. and Neri, S. (2007). What does a technology shock do? A VAR analysis with model-based sign restrictions. *Journal of Monetary Economics*, 54(2):512–549.
- Fernald, J. (2014). A quarterly, utilization–adjusted series on total factor productivity. *Working Paper*, (2012-19).

- Feve, P. and Guay, A. (2009). The response of hours to a technology shock: A two-step structural var approach. *Journal of Money, Credit and Banking*, 41(5):987–1013.
- Francis, N., Owyang, M., Roush, J., and DiCecio, R. (2014). A flexible finite-horizon alternative to long-run restrictions with an application to technology shocks. *Review of Economics and Statistics*, 96:638–647.
- Fuentes-Albero, C. (2019). Financial frictions, financial shocks, and aggregate volatility. *Journal of Money, Credit and Banking*, 51(6):1581–1621.
- Galeotti, M., Maccini, L. J., and Schiantarelli, F. (2005). Inventories, employment and hours. *Journal of Monetary Economics*, 52(3):575–600.
- Galí, J. (1999). Technology, employment, and the business cycle: Do technology shocks explain aggregate fluctuations?. *American Economic Review*, 89(1):249 – 271.
- Galí, J. and Gambetti, L. (2009). On the sources of the great moderation. *The American Economic Journal: Macroeconomics*, 1:26–57.
- Gilchrist, S. and Zakrajsek, E. (2012). Credit spreads and business cycle fluctuations. *American Economic Review*, 102(4):1692–1720.
- Görtz, C., Gunn, C., and Lubik, T. A. (2019). What Drives Inventory Accumulation? News on Rates of Return and Marginal Costs. Working Paper 19-18, Federal Reserve Bank of Richmond.
- Görtz, C. and Tsoukalas, J. (2018). Sectoral TFP news shocks. *Economics Letters*, 168(C):31–36.
- Görtz, C., Tsoukalas, J., and Zanetti, F. (2021). News Shocks under Financial Frictions. *American Economic Journal: Macroeconomics*, forthcoming.
- Görtz, C., Gunn, C., and Lubik, T. A. (2022). Is there news in inventories? *Journal of Monetary Economics*, 126:87–104.

- Jermann, U. and Quadrini, V. (2009). Financial innovations and macroeconomic volatility. *mimeo*.
- Jermann, U. and Quadrini, V. (2012). Macroeconomic effects of financial shocks. *American Economic Review*, 102(1):238–71.
- Jones, C. S. and Tuzel, S. (2013). New orders and asset prices. *Review of Financial Studies*, 26(1):115–157.
- Kahn, J. A., McConnell, M. M., and Perez-Quiros, G. (2002). On the causes of the increased stability of the U.S. economy. *Federal Reserve Bank of New York Economic Policy Review*, 8(1):183–206.
- Kim, C.-J. and Nelson, C. R. (1999). Has the u.s. economy become more stable? a bayesian approach based on a markov-switching model of the business cycle. *The Review of Economics and Statistics*, 81(4):608–616.
- Kurmann, A. and Sims, E. (2019). Revisions in utilization-adjusted tfp and robust identification of news shocks. *Review of Economics and Statistics*, forthcoming.
- Kurmann, A. and Sims, E. (2021). Revisions in utilization-adjusted tfp and robust identification of news shocks. *Review of Economics and Statistics*, 103(2):216–235.
- Maccini, L. J. and Rossana, R. J. (1984). Joint production, quasi-fixed factors of production, and investment in finished goods inventories. *Journal of Money, Credit and Banking*, 16(2):218–236.
- McCarthy, J. and Zakrajsek, E. (2007). Inventory Dynamics and Business Cycles: What Has Changed? *Journal of Money, Credit and Banking*, 39(2-3):591–613.
- McConnell, M. M. and Perez-Quiros, G. (2000). Output fluctuations in the United States: What has changed since the early 1980’s? *The American Economic Review*, 90(5):1464–1476.

- Pesavento, E. and Rossi, B. (2005). Do Technology Shocks Drive Hours Up Or Down? A Little Evidence From An Agnostic Procedure. *Macroeconomic Dynamics*, 9(4):478–488.
- Philippon, T. (2009). The bond market’s q. *Quarterly Journal of Economics*, 124:1011–1056.
- Ramey, Valerie Francis, N. (2005). Is the technology-driven real business cycle hypothesis dead? shocks and aggregate fluctuations revisited. *Journal of Monetary Economic*, 52:1379–99.
- Sarte, P.-D., Schwartzman, F., and Lubik, T. A. (2015). What inventory behavior tells us about how business cycles have changed. *Journal of Monetary Economics*, 76:264 – 283.
- Shea, J. (1998). What do technology shocks do? *NBER Macroeconomics Annual*, 13:275–310.
- Stock, J. H. and Watson, M. W. (1999). *Business Cycle Fluctuations in US Macroeconomic Time Series*. Elsevier, Amsterdam, Holland.
- Uhlig, H. (2003). What moves real gnp? Technical report, Humboldt University Mimeo.
- Uhlig, H. (2004). Do Technology Shocks Lead to a Fall in Total Hours Worked? *Journal of the European Economic Association*, 2(2-3):361–371.

Appendix

A Details on the VAR model

This appendix provides details on the VAR model, shock identification and prior specifications.

A.1 VAR-Based Identification of Technology Shocks

We consider the following vector autoregression (VAR), which describes the joint evolution of an $n \times 1$ vector of variables y_t :

$$y_t = A(L)u_t.$$

$A(L) = I + A_1L + \dots + A_pL^p$ is a lag polynomial of order p over conformable coefficient matrices $\{A_p\}_{i=1}^p$. u_t is an error term with $n \times n$ covariance matrix Σ . We assume a linear mapping between the reduced form errors u_t and the structural errors ε_t :

$$u_t = B_0\varepsilon_t,$$

where B_0 is an identification matrix. We can then write the structural moving average representation of the VAR:

$$y_t = C(L)\varepsilon_t,$$

where $C(L) = A(L)B_0$, $\varepsilon_t = B_0^{-1}u_t$, and the matrix B_0 satisfies $B_0B_0' = \Sigma$. B_0 can also be written as $B_0 = \tilde{B}_0D$, where \tilde{B}_0 is any arbitrary orthogonalization of Σ and D is an orthonormal matrix such that $DD' = I$.

We identify the technology shock using the Max Share methodology as suggested in Francis et al. (2014) who maximize the forecast error variance share of a productivity measure at a long but finite horizon. Following Kurmann and Sims (2019), we use TFP as the measure

for productivity. The Max Share methodology identifies productivity variations in the long run. The absence of any short run restrictions makes our applied identification robust to cyclical measurement issues of technology. Note that the methodology does not make an a priori assumption on whether technology reacts to the shock only with a lag or not.

Mechanically, we identify the technology shock by finding a rotation of the identification matrix \tilde{B}_0 , which maximizes the forecast error variance of the TFP series at some finite horizon. In this, we follow the Max Share approach of Francis et al. (2014). Specifically, the h -step ahead forecast error is given by:

$$y_{t+h} - E_{t-1}y_{t+h} = \sum_{\tau=0}^h A_{\tau} \tilde{B}_0 D \varepsilon_{t+h-\tau}.$$

The share of the forecast error variance of variable i attributable to shock j at horizon h is then:

$$V_{i,j}(h) = \frac{e_i' \left(\sum_{\tau=0}^h A_{\tau} \tilde{B}_0 D e_j e_j' D' \tilde{B}_0' A_{\tau}' \right) e_i}{e_i' \left(\sum_{\tau=0}^h A_{\tau} \Sigma A_{\tau}' \right) e_i} = \frac{\sum_{\tau=0}^h A_{i,\tau} \tilde{B}_0 \gamma \gamma' \tilde{B}_0' A_{i,\tau}'}{\sum_{\tau=0}^h A_{i,\tau} \Sigma A_{i,\tau}'},$$

where e_i denotes a selection vector with one in the i -th position and zeros everywhere else. The e_j vector picks out the j -th column of D , denoted by γ . $\tilde{B}_0 \gamma$ is therefore an $n \times 1$ vector corresponding to the j -th column of a possible orthogonalization and can be interpreted as an impulse response vector.

The Max Share approach chooses the elements of \tilde{B}_0 to make this restriction on forecast error variance share hold as closely as possible. This is equivalent to choosing the impact matrix so that contributions to $V_{1,2}(h)$ are maximized. Consequently, we choose the second column of the impact matrix to solve the following optimization problem:¹²

$$\arg \max_{\gamma} V_{1,2}(h) = \frac{\sum_{\tau=0}^h A_{i,\tau} \tilde{B}_0 \gamma \gamma' \tilde{B}_0' A_{i,\tau}'}{\sum_{\tau=0}^h A_{i,\tau} \Sigma A_{i,\tau}'}, \quad \text{s.t. } \gamma \gamma' = 1.$$

We restrict γ to have unit length to be a column vector of an orthonormal rotation matrix

¹²The optimization problem is written in terms of choosing γ conditional on any arbitrary orthogonalization \tilde{B}_0 to guarantee that the resulting identification belongs to the space of possible orthogonalizations of the reduced form.

of the Choleski decomposition of the reduced-form variance covariance matrix.

A.2 Specification for the Minnesota Prior in the VAR

We estimate the VAR using a Bayesian approach. The prior for the VAR coefficients A is a standard Minnesota prior as commonly used in the literature. It is of the form

$$vec(A) \sim N(\underline{\beta}, \underline{V}),$$

where $\underline{\beta}$ is one for variables in the baseline specification which are in log-levels, and zero for hours. The prior variance \underline{V} is diagonal with elements,

$$V_{i,jj} = \begin{cases} \frac{a_1}{p^2} & \text{for coefficients on own lags} \\ \frac{a_2 \sigma_{ii}}{p^2 \sigma_{jj}} & \text{for coefficients on lags of variable } j \neq i \\ a_3 \sigma_{ii} & \text{for intercepts} \end{cases}$$

where p denotes the number of lags. Here σ_{ii} is the residual variance from the unrestricted p -lag univariate autoregression for variable i . The degree of shrinkage depends on the hyperparameters $\underline{a}_1, \underline{a}_2, \underline{a}_3$. We set $\underline{a}_3 = 1$ and we choose $\underline{a}_1, \underline{a}_2$ by searching on a grid and selecting the prior that maximizes the in-sample fit of the VAR, as measured by the Bayesian Information Criterion.¹³

B Additional VAR Evidence

This section provides some additional empirical evidence that corroborates the results presented in the main body.

Labor Market Responses. Figure 5 shows that the subsample differences in hours-worked documented in Section 2.2 are also present if we replace total hours-worked with its components, the labor force participation rate and the unemployment rate. Consistent

¹³The grid of values we use is: $\underline{a}_1 = (1e-4:1e-4:9e-4, 0.001:0.001:0.009, 0.01:0.01:0.1, 0.1:0.1:1)$, $\underline{a}_2 = (0.01, 0.05, 0.1, 0.5, 1, 5)$. We consider all possible pairs of \underline{a}_1 and \underline{a}_2 in the above grids.

with the decline in hours-worked documented for the first subsample, Figure 5 documents a decline in the labor force participation rate and a rise in the unemployment rate. For the second subsample, the rise in hours-worked comes along with a rise in the labor force participation rate and a decline in the unemployment rate.




Figure 5: **IRF to TFP shock.** First subsample 1954Q2-1983Q4 (red), second subsample 1984Q1-2019Q4 (blue). The solid line is the median and the shaded colored areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

An Alternative Measure for Technology. Figure 6 shows impulse responses to a shock that maximizes the share of variance explained in labor productivity as in Francis et al. (2014). This shows that responses in Figure 1 are robust to using labor productivity instead of TFP as an alternative measure for productivity. In particular, also when using this measure for productivity we observe an expansion in GDP, consumption and stock prices that is more pronounced in the second subsample. Importantly hours work continue to decline in the first subsample and rise in the second subsample. An important difference between Figures 1 and 6 is that labor productivity responds strongly in the first subsample. This is consistent with findings in Francis et al. (2014) and Kurmann and Sims (2021) who flag this is due to a short-run capital deepening effect: the capital to labor ratio is driven up by the fall in hours-worked which in turn boosts labor productivity on impact relative to the more gradual rise in TFP documented in Figure 1.




Figure 6: **IRF to shock that maximizes variation in labor productivity.** First subsample 1954Q2-1983Q4 (red), second subsample 1984Q1-2019Q4 (blue). The solid line is the median and the shaded colored areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

Responses over the Entire Sample. Figure 7 shows the responses to a technology shock over the whole sample (1954Q2-2019Q4). All macroeconomic aggregates increase strongly and instantaneously in response to the shock. We also observe a rise in stock prices and a decline in credit spreads, so that these impulse responses resemble those documented in Figures 1 and 3 for the second subsample. Particularly the decline in hours-worked and inventories as well as the rise in credit spreads that we document for the first subsample is not evident when we estimate a VAR over the entire sample.

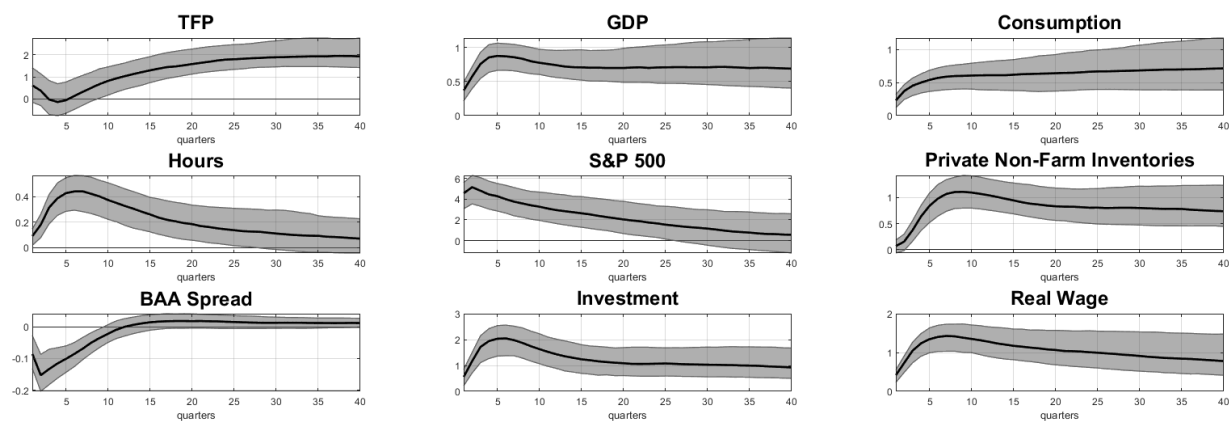


Figure 7: **IRF to TFP shock.** Entire sample 1954Q2-2019Q4. The solid line is the median and the shaded areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

C Data Sources and Time Series Construction

This section provides an overview of the data used to construct the observables. All the data transformations we have made in order to construct the dataset used for estimating the various VAR specifications and they enter in levels. The majority of the raw data described below were retrieved from the Federal Reserve of St. Louis FRED database. The exceptions are the TFP and utilization data series which is from Fernald (2014) at the Federal reserve bank of San Francisco, and the data on market yield and the BAA spread which are from the Federal reserve board and Bloomberg.

Data Sources. We describe the exact source of each data series below.

Gross domestic product, current prices: U.S. Bureau of Economic Analysis, Gross Domestic Product [GDP], retrieved from FRED, Federal Reserve Bank of St. Louis; *https://fred.stlouisfed.org/series/GDP*.

Gross Private Domestic Investment, current prices: U.S. Bureau of Economic Analysis, Gross Private Domestic Investment [GPDI], retrieved from FRED, Federal Reserve Bank of St. Louis; *https://fred.stlouisfed.org/series/GPDI*.

Real Gross Private Domestic Investment: U.S. Bureau of Economic Analysis, Real Gross Private Domestic Investment [GPDIC1], retrieved from FRED, Federal Reserve Bank of St. Louis; *https://fred.stlouisfed.org/series/GPDIC1*.

Personal Consumption Exp.: Durable Goods, current prices: U.S. Bureau of Economic Analysis, Personal Consumption Expenditures: Durable Goods [PCEDG], retrieved from FRED, Federal Reserve Bank of St. Louis; *https://fred.stlouisfed.org/series/PCEDG*.

Real Personal Consumption Exp.: Durable Goods: U.S. Bureau of Economic Analysis, Real Personal Consumption Expenditures: Durable Goods [PCEDGC96], retrieved from FRED, Federal Reserve Bank of St. Louis; *https://fred.stlouisfed.org/series/PCEDGC96*.

Personal Consumption Expenditures: Services, current prices: U.S. Bureau of Economic Analysis, Personal Consumption Expenditures: Services [PCES], retrieved from FRED, Fed-

eral Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/PCES>.

Real Personal Consumption Expenditures: Services: U.S. Bureau of Economic Analysis, Real Personal Consumption Expenditures: Services [PCESC96], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/PCESC96>.

Personal Consumption Exp.: Nondurable Goods, current prices: U.S. Bureau of Economic Analysis, Personal Consumption Expenditures: Nondurable Goods [PCEND], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/PCEND>.

Real Personal Consumption Exp.: Nondurable Goods: U.S. Bureau of Economic Analysis, Real Personal Consumption Expenditures: Nondurable Goods [PCENDC96], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/PCENDC96>.

Real Private Nonfarm Inventories: U.S. Bureau of Economic Analysis [A373RX1Q020SBEA], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/A373RX1Q020SBEA>.

Civilian Noninstitutional Population: U.S. Bureau of Labor Statistics, Population Level [CNP16OV], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/CNP16OV>.

Non-farm Business Sector: Compensation Per Hour: U.S. Bureau of Labor Statistics, Non-farm Business Sector: Compensation Per Hour [COMPNFB], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/COMPNFB>.

Non-farm Business Sector: Hours of All Persons: U.S. Bureau of Labor Statistics, Nonfarm Business Sector: Hours of All Persons [PRS85006031], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/PRS85006031>.

Effective Federal Funds Rate: Board of Governors of the Federal Reserve System (US), Effective Federal Funds Rate [FEDFUNDS], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/FEDFUNDS>.

Implicit GDP deflator: U.S. Bureau of Economic Analysis, Gross Domestic Product: Implicit Price Deflator [A191RI1Q225SBEA], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/A191RI1Q225SBEA>.

10 year treasury yield: The market yield on U.S. Treasury securities at 10-year constant maturity are available from the Federal Reserve Board H.15 database.

The BAA yield is Moody's Bond Indices Corporate BAA obtained from Bloomberg.

The real S&P 500 index is obtained from Robert Shiller's website (<http://www.econ.yale.edu/shiller/>)

The utilization adjusted TFP data and the series for capacity utilization can be accessed at www.frbsf.org/economic-research/economists/jfernalld/quarterly_tfp.xls.

The raw data are transformed as follows for the analysis. Consumption (in current prices) is defined as the sum of personal consumption expenditures on services and personal consumption expenditures on non-durable goods. The times series for real consumption is constructed as follows. First, we compute the shares of services and non-durable goods in total (current price) consumption. Then, total real consumption growth is obtained as the chained weighted (using the nominal shares above) growth rate of real services and growth rate of real non-durable goods. Using the growth rate of real consumption we construct a series for real consumption.

Real output is GDP derived by dividing current price GDP with the GDP deflator and the Civilian Noninstitutional Population measure. Similarly for hours-worked, consumption, investment and hourly wages (defined as total compensation per hour). All these series, as well as the real inventory measure are expressed in per capita terms using the series of non-institutional population, ages 16 and over. The nominal interest rate is the effective federal funds rate. The BAA spread series is the difference between the BAA yield and the 10 year treasury yield.