Split Personalities: The Changing Nature of Technology Shocks*

Christoph Görtz
University of Birmingham

Christopher Gunn
Carleton University

Thomas Lubik
Federal Reserve Bank of Richmond

March 2023

***** Preliminary and Incomplete *****

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. We then turn to a rich structural model to study potential causes of the changing impact of technology. Using an IRF-matching feature, we find that a change in the stance of monetary policy and the nature of intangible capital accumulation both played dominant roles in accounting for the changing impact of technology over time.

Keywords: technology shocks, TFP, business cycles, shock transmission.

JEL Classification: E2, E3.

*We thank Frank Schorfheide, Felipe Schwartzman and Michael Weber for helpful comments and suggestions. The views expressed in this paper are those of the authors and not necessarily those of the Federal Reserve Bank of Richmond or the Federal Reserve System.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. Lubik: Federal Reserve Bank of Richmond, Research Department, Email: thomas.lubik@rich.frb.org

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 transmission of technology shocks has changed 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 subsampple 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.

Next we turn to a rich structural model to study the potential causes of the changing impact of technology. In order to consider various different possibilities for the sources of the change, we augment a New Keynesian framework with a banking sector and financial frictions based on Gertler and Karadi (2011), inventory holding by firms as in Lubik and Teo (2012), intangible capital as an additional input into production, which we refer to as knowledge capital, as in Gunn and Johri (2011a), as well as various other standard real rigidities. Using the same data subsample split as in our VAR analysis, we then estimate values for key parameters seprately for each subsample using an impulse response matching procedure which minimizes a function of the distance between the empirical VAR and model-generated impulse response functions, conditional on an anticipated shock to the growth rate of non-stationary Total Factor Productivity. Using the results of this procedure and a series of model-based experiments, we then evaluate various candidate hypothesises for the source of the change through the lens of the model. Our results suggest that the change in the response of technology over time was likely some combination of a change in the stance of monetary policy, a change in the nature of knowledge capital accumulation, and a change in the cost of utilizing capital.

With respect to the change in the stance of monetary policy, our results suggest a move towards tighter monetary policy in response to inflation and looser response to the output gap in a Taylor-type rate-setting rule. This change in stance then works through real-interest rate effects on labour and inventory decisions, and a powerful channel discussed in Christiano et al. (2007) through which under nominal wage rigidities, an inflation-targeting central bank influences the path of the real wage. With respect to the change in the nature of knowledge capital accumulation, our results suggest a change in the elasticity in which hours-worked in production contributes to the accumulation of new knowledge capital, causing firms increase labour demand and lower markups as they seek to acquire valuable new knowledge capital in the face of expanding future technology. Finally, with respect to the change in the cost of utilizing capital, our results suggest an increase in this cost working through general equilibrium effects though the credit sector by the influence of this cost on on the return to capital and associated increase in demand for new capital.

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 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 short-comings 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

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

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

bias in the short-run impulse responses and raising their estimation precision. They find — 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.

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

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

This literature and our work has implications for the estimation of structural models — 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 analyze impulse response functions based on two separate subsamples using a minimally specified VAR framework, explore robustness along a number of dimensions and perform various empirical exercise to try and isolate and understand the source of this technology change in the role of technology. Section 3.1 outlines a structural model. This model points in Section 4 to potential channels that may drive the differences across subsamples. In Sections 5 and 7 we employ an IRF matching procedure to investigate the empirical relevance of these channels. Section 8 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 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 (2021), 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 (2021), 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 (2021), 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 identification 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

[5]Francis 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.

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 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

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

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-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]

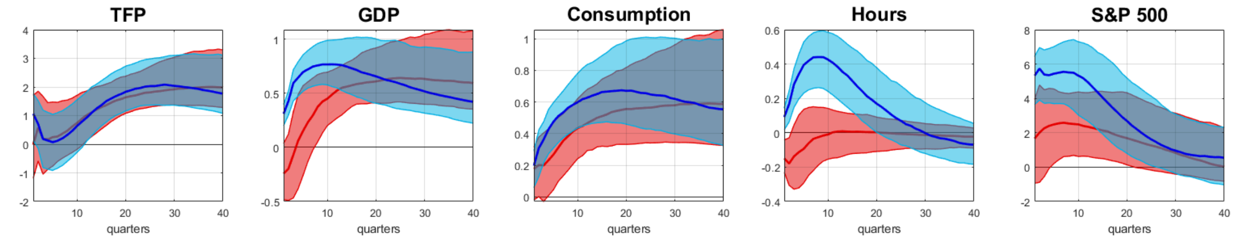

Figure 1: **IRF to TFP shock.** First subsample 1954Q2-1983Q4 (red), second subsample 1984Q1-2019Q4 (blue). The solid line is the median and the shaded colored areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage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

[7]The 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.

[8]These 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.

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-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.

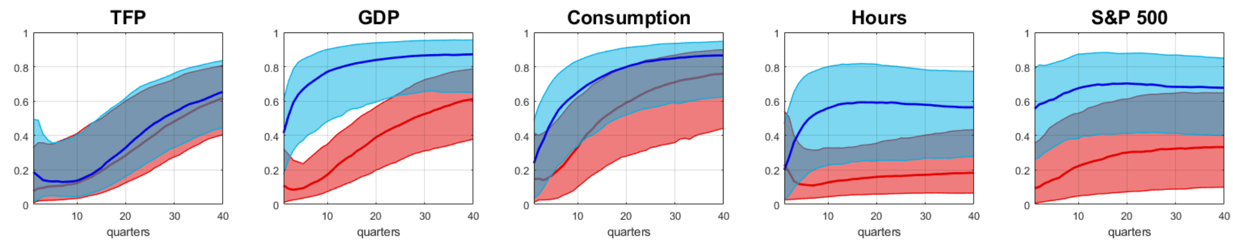

Figure 2: **Forecast Error Variance Decomposition — share explained by the TFP shock.** First subsample 1954Q2-1983Q4 (red), second subsample 1984Q1-2019Q4 (blue). The solid line is the median and the shaded colored areas are the 16% and 84% posterior bands generated from the posterior distribution of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

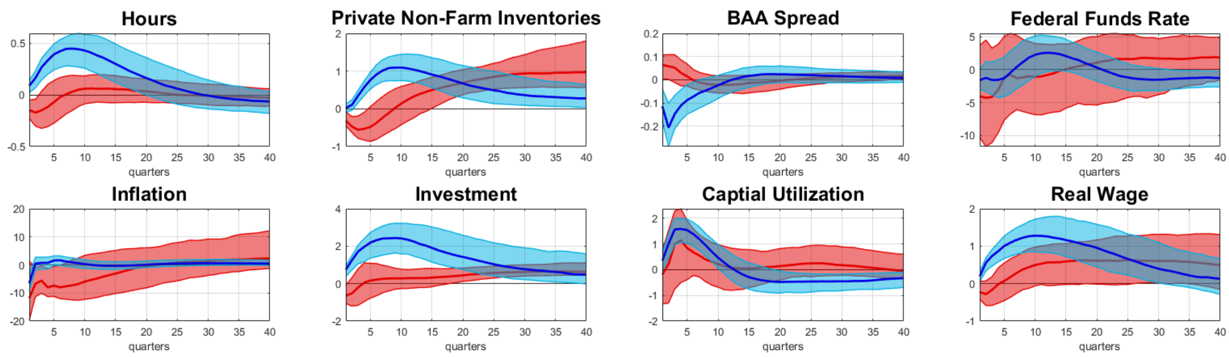

Figure 3: **IRF to TFP shock.** First subsample 1954Q2-1983Q4 (red), second subsample 1984Q1-2019Q4 (blue). The solid line is the median and the shaded colored areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage 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.

Taken 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 (2021), 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 Maccini 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

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

[10]Görtz 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.

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 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.

Before we do so, we want to shed some more light into the timing of when impulse responses flip sign. For this purpose, we estimate our VAR model over rolling windows of 119 quarters. This choice implies that the first rolling window is consistent with our first subsample and we shift this window forward until its end corresponds to the end of the second subsample. Figure 4 displays the maximum or minimum (whichever is larger in absolute terms) of IRF responses within the first ten quarters.[12] For hours worked, it is evident that

[11] Other 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.

[12] In Appendix B, we report corresponding statistics for the impact responses of IRFs. Results are consistent

six quarters after the rolling window shifts beyond the end of the first subsample, the largest (in absolute terms) median response within the first ten quarters turns positive. Once the end of the rolling estimation window includes the year 2000 the positive response is almost always significant. For inventories the picture is similar, although here the response flips into positive significant territory already for a sample end at around 1994. Also investment moves very quickly from a negative response to an insignificant one before it becomes significant and positive once the sample end includes 1999. Also the response of the BAA spread becomes insignificant very soon once estimation windows move away from the first subsample. The spread response remains insignificant somewhat longer than those of the other variables and flips to be negative and significant once the sample includes observations after the financial crisis.

Overall, Figure 4 shows that as soon as the sample includes observations that are considered part of the Great Moderation period, the negative response of hours, investment and inventories, and the positive response of the spread become insignificant. Once samples include more post 1984 observations the IRFs flip sign and remain in this territory. This rolling window exercise illustrates that the findings discussed in relation to Figure 1 are not solely related to the two specific subsamples under consideration but reflect a broader feature in the data. It also shows that the transmission of technology shocks has been affected for some variables by significant events — such as the financial crisis for the BAA spread and investment, the Great Moderation for inventories, investment and hours, and the late 1990's technology booom for investment and hours.

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

with those of Figure 4. The same holds for corresponding figures with a shorter window length, which are also provided in this appendix.

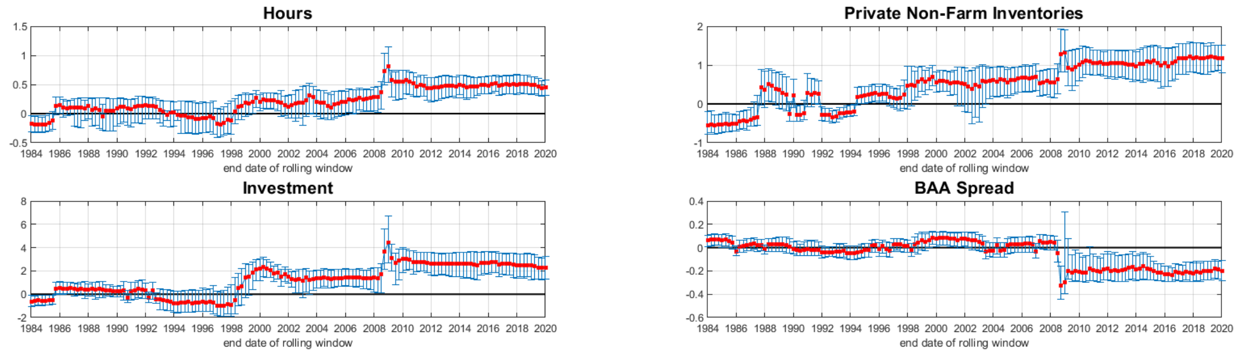

Figure 4: **Maximum/minimum (whichever is largest in absolute terms) IRF response within the first ten quarters to a TFP shock for rolling window.** First rolling window sample is 1954Q2-1983Q4 (119 quarters). The window is shifted up to 2019Q4. We display the median (red dot) and the 16% and 84% posterior 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.

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 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:

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

$A(L) = I + A_1 L + ... + A_p L^p$ 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 shocks 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 $B_0 = \widetilde{B}_0 D$, where $\widetilde{B}_0$ 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 coefficents and the second-subsample variance-covariance matrix.

Figure 5 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 VAR is run on the second subsample but the shock is identified using the first-subsample polynomial-lag coefficients and the second-subsample variance co-variance matrix. Similarly, the red lines in the second row of Figure 5 show the responses if the VAR is run on the first subsample and 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.

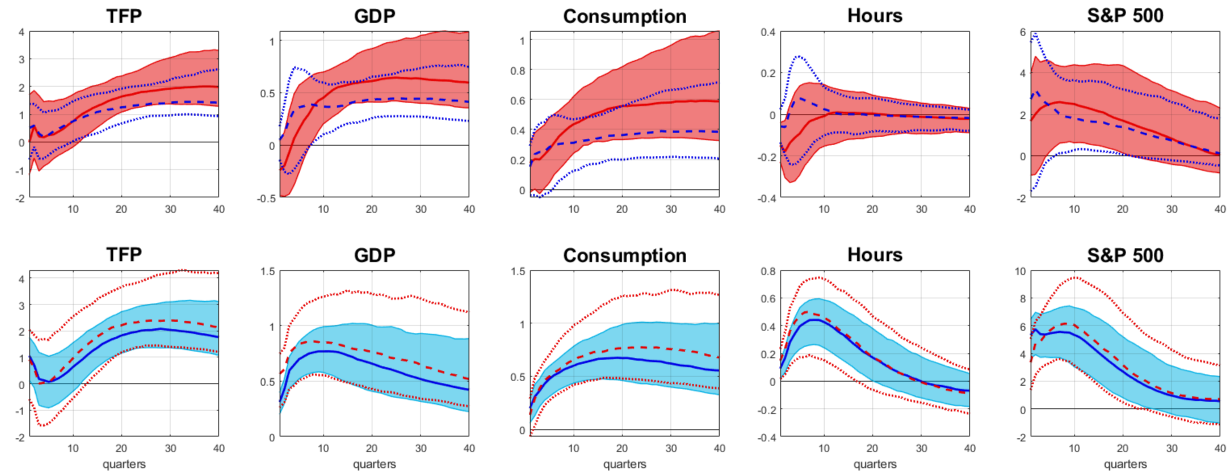

Figure 5: **IRF to TFP shock.** The solid red (blue) line is the median and the shaded red (blue) areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters on the first (second) subsample. First subsample is subsample 1954Q2-1983Q4, second subsample is 1984Q1-2019Q4. The blue (red) dashed and dotted lines in the subplots in row one (two) are the median and posterior bands when running the VAR on the second (first) subsample, but identifying the shock using the polynomial lag coefficients implied by the first (second) subsample and the variance co-variance matrix implied by the second (first) subsample.

3 Differences in Shock Transmission through the Lens of a Structural Model

Our empirical results above document the changes in the response of the economy to technological shocks over time, yet the analysis remains agnostic about the underlying source of these changes. We now use use a structural model to provide some interpretation to potential underlying causes. As we discussed above, the extensive literatures studying changes in the structure of the macroeconomy over the 1970's – 1990's have suggested several important changes over this period, including: (1) changes in inventory management (2) changes in labour market rigidities (3) changes in monetary policy (4) emergence of the information and communications technology (ICT) era. While the focus of these studies has not been the changing effect of technology shocks, our rich structural model allows us to provide some insight into whether these potential underlying changes in the economy could also be behind the changing impact of technology.

Our structural framework is a medium scale New Keynesian model of augmented with inventories, a financial sector with financial frictions, and knowledge capital accumulation by firms. We model inventories as in Lubik and Teo (2012), based on the stock-elastic demand model of Bils and Kahn (2000), where finished goods inventories are sales-enhancing. The financial side of the model uses the setup of Gertler and Karadi (2011). Finally, knowledge capital accumulation by firms follows the approach of Gunn and Johri (2011b), Chang et al. (2002) and Cooper and Johri (2002) whereby firms accumulate productivity-enhancing knowledge through an internalized learning-by-doing process in labor.

Our results above focus on the response of the economy to an identified exogenous technology shock, and thus our core analysis in our theoretical model focuses on the conditional response of the model economy to an exogenous technological shock. Additionally, our results above suggest that dominant technological shock in both subsamples is an anticipated or diffused shock where the 16% lower posterior band of TFP impulse response only rises above zero in the range of 12 periods out, and thus our exogenous technological shock takes the form of a ”news” or anticipated shock to TFP received 12 periods in advance of the actual change in TFP. Nevertheless, in documenting our model, we include a full suite of shocks to facilitate additional analysis.

3.1 Model description

The model consists of a large number of identical infinitely-lived households, a competitive intermediate goods-producing firm, a continuum of monopolistically competitive distributors, a competitive final goods producer, a continuum of competitive financial intermediaries, a competitive capital services firm, a competitive capital goods producer, a continuum of monopolistically competitive labour unions, a competitive employment agency and a monetary policy authority. The intermediate goods firm produces a homogeneous good that it sells to distributors. This good is then differentiated by the distributors into distributor-specific varieties that are sold to the final-goods firm. The varieties are aggregated into final output, which then becomes available for consumption or investment. Households are comprised of a fraction $1 - f$ of workers and f of bankers. Workers supply labor, bankers manage financial intermediaries, and both return their earnings to the household. Since this particular decentralization of wage stickiness implies that choices on consumption and hours-worked are identical across households, for simplicity we will refer to a stand-in representative household. The model economy contains several stationary stochastic shock processes as well as non-stationary TFP and IST shocks and a suite of shocks that are standard in the literature to facilitate estimation.

3.2 Households and Government

The stand-in household's lifetime utility is defined over sequences of consumption C_t and hours worked N_t and is given by

$$E_0 \sum_{t=0}^{\infty} \beta^t \Gamma_t \frac{\left(V_t^{1-\sigma} - 1\right)}{1-\sigma}, \tag{1}$$

where $0 < \beta < 1$, $\sigma > 0$, and where Γ_t is a stationary stochastic preference shock process. The argument V_t is given by

$$V_t = C_t - bC_{t-1} - \psi N_t^{\xi} F_t, \quad \text{where} \quad F_t = (C_t - bC_{t-1})^{\gamma_f} F_{t-1}^{1-\gamma_f}, \tag{2}$$

is a preference component that makes consumption and labor non-time-separable and is consistent with the balanced-growth path in a growing economy. This preference structure, which follows Schmitt-Grohe and Uribe (2012) and is based on Jaimovich and Rebelo (2009), nests the no-income effect structure of Greenwood et al. (1988) in the limit as the parameter $0 < \gamma_f \leq 1$ tends toward zero. The parameter $0 \leq b < 1$ allows for habits in consumption;

and $\xi > 1$ is related to the Frisch elasticity of labour supply.

The household enters each period with real financial securities, B_t, which serve as deposits with the financial intermediaries, and nominal bonds, B_t^n, earning risk-free gross real rate of return, R_t, and risk-free gross nominal rate of return, R_t^n, respectively, receiving nominal wage, W_t^h, for supplying hours, N_t^h, to the labour union, and receiving a share of real profits from the various other entities in the model, denoted collectively as Π_t. At the end of the period, the household chooses its consumption C_t, its holdings of financial deposits B_{t+1} and nominal nominal bonds B_{t+1}^n. The household's period t budget constraint is given by

$$C_t + B_{t+1} + \frac{B_{t+1}^n}{P_t} + T_t = R_t B_t + R_t^n \frac{B_t^n}{P_t} + \frac{W_t^h}{P_t} N_t^h + \Pi_t, \tag{3}$$

where P_t is the price of the final good in terms of the nominal unit under the control of the central bank and T_t denotes lump-sum taxes. The household's problem is to choose sequences of C_t, N_t^h, B_{t+1} and B_{t+1}^n to maximize equation (1) subject to equations (2)–(3), resulting in standard first-order conditions.

Government spending follows the process $G_t = \left(1 - \frac{1}{\varepsilon_t}\right) Y_t$, where ε_t is a stationary stochastic government spending shock.

3.3 Financial intermediaries

Our financial intermediary framework follows that of Gertler and Karadi (2011) and so we show only the core elements here. In period t, the jth financial intermediary obtains deposit funds, B_{jt+1}, from households. The intermediary uses those funds and their own net-worth, N_{jt}, to make state-contingent loans, S_{jt}, to non-financial capital services firms, such that the intermediary's financing satisfies the balance sheet identity

$$Q_t S_{jt} = N_{jt} + B_{jt+1},$$

where Q_t is the price of the state-contingent loan. The intermediary's net-worth then evolves as

$$N_{jt+1} = R_{kt+1} Q_t S_{jt} - R_{t+1} B_{jt+1},$$

where R_{t+1} is the non-contingent rate paid on household deposits (determined in t), and R_{kt+1} is the state-contingent return on loans.

An intermediary this period stays an intermediary next period with exogenous probability θ_b, at which point they will have been deemed to exit, a feature that insures an intermediary will not grow so large as to be able to self-finance all its loans. The intermediary will continue to operate and build wealth until exiting as long as the risk adjusted premium on making loans over borrowing is positive. The intermediary thus maximizes expected terminal wealth, given by

$$V_{jt} = \max_{\{S_{jt+i}, N_{jt+1+i}\}} E_t \sum_{i=0}^{\infty} (1 - \theta_b) \theta_b^i \Lambda_{t,t+1+i} N_{jt+1+i},$$

where $\Lambda_{t,t+i} = \beta^i \frac{\lambda_{t+i}}{\lambda_t}$ is the household's stochastic discount factor.

The financial friction takes the form of a moral hazard/cost enforcement problem where each period, the intermediary can divert the fraction λ_b of assets back to the household, at which point the intermediary is forced into bankruptcy and the depositors recover the fraction $1 - \lambda_b$. Thus for depositors to be willing to supply funds to the intermediary, the following enforcement constraint

$$V_{it} \geq \lambda_b Q_t S_{jt}$$

must hold, such that the value to the intermediary of continuing to operate is at least as large as the value of absconding funds. Conjecturing and subsequently verifying that the solution is linear in its balance sheets components in the form

$$V_{jt} = \nu_{bt} Q_t S_{jt} + \eta_{bt} N_{jt},$$

and that leverage, ϕ_{bt}, defined as

$$\phi_{bt} = \frac{Q_t S_{jt}}{N_{jt}},$$

is not dependent on intermediary-specific factors, we can then solve for ν_{bt} and η_{bt}. Under the case that the enforcement constraint is binding (as in Gertler and Karadi (2011)),

$$\phi_{bt} = \frac{\eta_t}{\lambda_b - \nu_{bt}},$$

and ν_{bt} and η_{bt} are then given by

$$\nu_{bt} = E_t \beta \frac{\lambda_{t+1}}{\lambda_t} \Gamma_{t+1} \delta_{bt+1} \qquad \text{and} \qquad \eta_{bt} = E_t \beta \frac{\lambda_{t+1}}{\lambda_t} \Gamma_{t+1} R_{t+1},$$

where $\Gamma_t = 1 - \theta_b + \theta_b(\nu_{bt}\phi_{bt} + \eta_{bt})$ and $\delta_{bt} = R_t^k - R_t$.

3.4 Employment unions and employment agency

Our sticky-wage framework follows the decentralization of Schmitt-Grohe and Uribe (2012) and Smets and Wouters (2007) with a continuum of monopolistically competitive labor unions on a unit mass indexed by $j \in [0,1]$, and a competitive employment agency. Monopolistic unions buy homogeneous labor from households, transform it into differentiated labor inputs, and sell it to the employment agency who aggregates the differentiated labor into a composite which it then sells to the intermediate goods producer. The unions face frictions in setting wages for each labor type. The unions face Calvo frictions in setting their wages for each labour type, and re-set their wage according to an indexation rule when unable to reoptimize.

Labor unions acquire homogeneous labor, N_t^h, from the household at wage W_t^h, differentiate it into labor types N_{jt}, $j \in [0,1]$, and then sell the differentiated labor it to the employment agency for wage, W_{jt}. The unions have market power, and can thus choose the wage for each labor type subject to the labor demand curve for that labor type. In particular, the unions face Calvo frictions in setting their wages, such that each period they can re-optimize wages with probability $1-\zeta_w$. A union that is unable to re-optimize wages re-sets it according to the indexation rule $W_{jt} = W_{jt-1}\pi_{t-1}^{\iota_w}\pi^{1-\iota_w}, \quad 0 \leq \iota_w \leq 1$, where $\pi_t = P_t/P_{t-1}$ and π is its steady state, and where $0 \leq \iota_w \leq 1$. A union that can re-optimize its wage in period t chooses its wage W_{jt}^* to maximize

$$E_t \sum_{s=0}^{\infty} \zeta_w^s \beta^s \frac{\lambda_{t+s} P_t}{\lambda_t P_{t+1}} \left[W_{jt}^* (\Pi_{k=0}^s \pi_{t+k-1}^{\iota_w} \pi^{1-\iota_w}) - W_{t+s}^h \right] n_{jt+s},$$

subject to the demand curve for N_{jt}.

The employment agency acquires each jth intermediate labor type N_{jt}, $j \in [0,1]$, at wage W_{jt} from the labor unions, and combines the differentiated labor into a composite n_t according to

$$n_t = \left[\int_0^1 n_{jt}^{\nu_w} dj \right]^{\frac{1}{\nu_w}}, \quad 0 < \nu_w \leq 1.$$

The agency sells the composite labor to the intermediate goods producers for wage W_t. The agency chooses n_{jt} $\forall j$ to maximize profits $W_t n_t - \int_0^1 W_{jt} n_{jt} dj$, yielding a demand function n_{jt} for the jth labor type,

$$N_{jt} = \left[\frac{W_{jt}}{W_t} \right]^{\frac{1}{\nu_w - 1}} N_t,$$

and wage index W_t, given respectively by

$$W_t = \left[\int_0^1 W_{jt}^{\nu_w/(\nu_w-1)} dj \right]^{\frac{(\nu_w-1)}{\nu_w}}.$$

The sticky wage framework results in a time-varying markup μ_t^w between the wage W_t paid by the intermediate goods firm and the wage W_t^h paid to the household, such that

$$\mu_t^w = \frac{w_t}{w_t^h}, \tag{4}$$

where $w_t = \frac{W_t}{P_t}$ and $w_t^h = \frac{W_t^h}{P_t}$. The dynamics of μ_t^w is captured by a resulting equilibrium wage Phillips curve derived from imposing equilibrium on the combination of the employment agency and union's problem.

3.5 Intermediate goods firm

The competitive intermediate goods firm produces the homogeneous good Y_t with technology

$$Y_t = z_t \left(\Omega_t N_t\right)^{\alpha_n} \widetilde{K}_t^{\alpha_k} \left(\Omega_t H_t\right)^{1-\alpha_n-\alpha_k},$$

where z_t is a stationary exogenous stochastic productivity process, Ω_t is a non-stationary exogenous stochastic productivity process, and H_t is stock of intangible capital that resides within the firm and that we refer to as knowledge capital.

Following Gunn and Johri (2011a), Chang et al. (2002) and Cooper and Johri (2002) we assume that the stock of knowledge capital, H_t, evolves as an internalized learning-by-doing process to capture the idea that agents acquire new technological knowledge through their experiences in engaging labor in the production process.[13] Accordingly, H_t evolves as

$$H_{t+1} = (1-\delta_h)H_t + H_t^{\gamma_h} N_t^{1-\gamma_h}, \qquad \text{where} \qquad 0 \leq \delta_h \leq 1, \quad 0 \leq \gamma_h < 1, \quad \nu_h > 0. \tag{5}$$

The accumulation (5) nests a log-linear specification for $\delta_h = 1$ common in the literature such as in Chang et al. (2002), Cooper and Johri (2002) and d'Alessandro et al. (2019), but also allows for a more general linear formulation for $0 < \delta_h < 1$.

Each period, the firm acquires labor, N_t, at wage, w_t, from the labor market, and capital services, $\widetilde{K}_t$, at rental rate r_t from the capital services market. It then sells its output, Y_t, at real price, τ_t, to the distributors. Additionally, we find it convenient to define the marginal cost of production for intermediate goods, $mc_t = \frac{w_t}{MPN t} = \frac{w_t}{\alpha_n Y_t / N_t}$, where $MPN_t = F_{Nt}$ is the marginal product of labor. It then follows that the output price, τ_t, is equal to the marginal cost of production, mc_t.

The firm's optimization problem involves choosing N_t, $\widetilde{K_t}$ and H_{t+1} to maximize its stream of profits, $E_0 \sum_{t=0}^{\infty} \frac{\beta^t \lambda_t}{\lambda_0} \Pi_t^y$, subject to the production function and knowledge capital accumulation equation, where $\Pi_t^y = \tau_t Y_t - w_t N_t - r_t \widetilde{K}_t$.

[13] See Görtz et al. (2022), Gunn and Johri (2011a), Chang et al. (2002) and Cooper and Johri (2002) for detailed discussions of learning-by-doing as a business cycle mechanism.

3.6 Capital services firm

At the end of each period, the competitive capital services firm buys capital, K_{t+1}, from the capital producer at price q_t^k, financing it with loans from the financial intermediaries in the form of state-contingent claims, S_t^b, equal to the number of units of capital, and pricing each claim at the price of a unit of capital. At the beginning of $t+1$, the firm rents services of the capital, $\tilde{K}_{t+1} = u_t K_{t+1}$, to intermediate goods firms at price r_t. At the end of the period, the firm incurs utilization costs of $a(u_{t+1})K_{t+1}\Upsilon_{t+1}$, sells the undepreciated capital back to capital goods producers at price q_{t+1}^k, and pays out state-contingent profits Π_{t+1}^k to financial intermediaries, where

$$\Pi_{t+1}^k = r_{t+1}u_{t+1}K_{t+1} - a(u_{t+1})K_{t+1}\Upsilon_{t+1} + q_{t+1}^k K_{t+1}.$$

After observing the aggregate state in $t+1$, the firm the faces the problem to choose u_{t+1} to maximize Π_{t+1}^k, yielding the optimality condition $a'(u_{t+1})\Upsilon_{t+1} = r_{t+1}$.

Letting R_{t+1}^k be the state-contingent gross real return on the on claims issued in t, then $\Pi_{t+1}^k = R_{t+1}^k q_t^k S_t^b = R_{t+1}^k q_t^k K_{t+1}$, such that using the firm's optimality conditions for u_{t+1},

$$R_{t+1}^k = \frac{r_{t+1}u_{t+1} - a(u_{t+1})\Upsilon_{t+1} + q_{t+1}^k}{q_t^k}.$$

3.7 Capital-producer

The competitive capital-goods producer operates a technology that combines existing capital with new investment goods to create new installed capital. At the end of each period it purchases existing capital K_t^k from entrepreneurs at price $\bar{q}_t$, combining it with investment I_t to yield new capital stock K_t^{nk}, which it sells back to entrepreneurs in the same period at price q_t. The capital-producer faces investment adjustment costs in the creation of new capital, and incurs depreciation in the process, so that

$$K_t^{nk} = (1-\delta)K_t^k + I_t\left[1 - S\left(\frac{m_t I_t}{I_{t-1}}\right)\right]. \tag{6}$$

where m_t is a stationary exogenous stochastic process for marginal efficiency of investment, and $S(x)$ is an investment adjustment cost function with the properties $S(x) = 0$, $S'(x) = 0$, and $S''(x) = s''$, where s'' is a parameter. The capital producer's period t profits are given by $\Pi_t^k = q_t K_t^{nk} - \bar{q}_t K_t^k - I_t/\Upsilon_t$, where Υ_t captures non-stationary exogenous stochastic investment-specific technological change. Since the capital producer faces intertemporal investment adjustment costs, it solves a dynamic problem, choosing K_t^{nk}, K_t^k and I_t to maximize $E_0 \sum_{t=0}^{\infty} \frac{\beta^t \lambda_t}{\lambda_0} \Pi_t^k$ subject to equation (6).

3.8 Final goods firm

The competitive final goods firm produces goods for sale, S_t, by combining distributor-specific varieties S_{it}, $i \in [0, 1]$, according to the technology

$$S_t = \left[\int_0^1 \nu_{it}^{\frac{1}{\theta}} S_{it}^{\frac{\theta-1}{\theta}} di \right]^{\frac{\theta}{\theta-1}}, \qquad \text{with} \qquad \nu_{it} = \left(\frac{A_{it}}{A_t} \right)^{\zeta}, \qquad \text{and} \qquad \theta > 1,\ \zeta > 0.$$

where ν_{it} is a taste shifter that depends on the stock of goods available for sale A_{it}. The latter is composed of current production and the stock of goods held in inventory.[14] We assume that ν_{it} is taken as given by the final goods producer and A_t is the economy-wide average stock of goods for sale, given by $A_t = \int_0^1 A_{it} di$. The parameters θ and ζ capture, respectively, the elasticity of substitution between differentiated goods and the elasticity of demand with respect to the relative stock of goods.

The firm acquires each variety i from the distributors at relative price $p_{it} = P_{it}/P_t$, where $P_t = \left[\int_0^1 v_{it} P_{it}^{1-\theta} di \right]^{\frac{1}{1-\theta}}$ is the aggregate price index. It sells the final good for use in consumption or as an input into the production of investment goods. The firm maximizes the profit function $\Pi_t^s = S_t - \int_0^1 \frac{P_{it}}{P_t} S_{it} di$ by choosing S_{it}, $\forall i$. This results in demand for S_{it} for the ith variety

$$S_{it} = \nu_{it} p_{it}^{-\theta} S_t. \tag{7}$$

An increase in ν_{it} shifts the demand for variety i outwards. This preference shift is influenced by the availability of goods for sale of variety i, which thereby provides an incentive for firms to maintain inventory to drive customer demand and avoid stockouts.

[14]This structure follows Bils and Kahn (2000) and is standard in modeling demand for goods drawn from inventories. It also supports a convenient decentralization of production.

3.9 Distributors

We follow Bils and Kahn (2000) in modeling inventories as a mechanism that helps generate sales, while at the same time implying a target inventory-sales ratio that captures the idea of stockout avoidance. Distributors acquire the homogeneous good Y_t from the intermediate goods firms at real price τ_t. They differentiate Y_t into goods variety Y_{it} at zero cost, with a transformation rate of one-to-one. Goods available for sale are the sum of the differentiated output and the previous period's inventories subject to depreciation

$$A_{it} = (1 - \delta_x) X_{it-1} + Y_{it}, \tag{8}$$

where the stock of inventories X_{it} are the goods remaining at the end of the period

$$X_{it} = A_{it} - S_{it}, \tag{9}$$

and $0 < \delta_x < 1$ is the rate of depreciation of the inventory stock.

The distributors have market power over the sales of their differentiated varieties. The ith distributor sets price p_{it} for sales S_{it} of its variety subject to its demand curve (7). Distributors face frictions in setting their prices, and as in Lubik and Teo (2012) , we assume that the ith distributor faces convex adjustments costs in the form $\frac{\kappa}{2}\left[\frac{P_{it+k}}{\pi_{t-1}^{\iota_p}\pi^{1-\iota_p}P_{it+k-1}} - 1\right]^2 S_t$. Each period, a distributor faces the problem of choosing p_{it}, S_{it}, Y_{it}, and A_{it} to maximize profits

$$E_t \sum_{k=0}^{\infty} \beta^k \frac{\lambda_{t+k}}{\lambda_t} \left\{ \frac{P_{it+k}}{P_{t+k}} S_{it+k} - \tau_t Y_{it+k} - \frac{\kappa}{2}\left[\frac{P_{it+k}}{\pi_{t-1}^{\iota_p}\pi^{1-\iota_p}P_{it+k-1}} - 1\right]^2 S_t \right\},$$

subject to the demand curve (7), the law of motion for goods available for sale (8), and the definition of the inventory stock (9), and where λ_t is household's marginal utility of wealth.

3.10 Monetary policy

We close the model with a standard monetary policy rule where the interest rate, R_{t+1}^n, is set by the monetary authority according to a feedback rule,

$$\frac{R_{t+1}^n}{R^n} = \left(\frac{R_t^n}{R^n}\right)^{\rho_r} \left(\left(\frac{\Pi_t}{\Pi}\right)^{\phi_\pi} \left(\frac{Y_t}{Y_t^*}\right)^{\phi_y}\right)^{(1-\rho_r)} e^{\eta_t},$$

where Π_t is the gross inflation rate, η_t is a monetary policy shock, and Y_t^* is level of output that would preside under flexible prices and without wage or price markup shocks.

3.11 Non-stationary TFP Stochastic Process

The non-stationary exogenous stochastic TFP process Ω_t, with growth rate g_t^Ω is given by:[15]

$$\ln\left(\frac{g_t^\Omega}{g^\Omega}\right) = \rho_{g^\Omega}\ln\left(\frac{g_{t-1}^\Omega}{g^\Omega}\right) + u_t^{g^\Omega}, \qquad \text{with} \qquad u_t^{g^\Omega} = \epsilon^0_{g^\Omega t} + \epsilon^4_{g^\Omega t-4} + \epsilon^8_{g^\Omega t-8} + \epsilon^{12}_{g^\Omega t-12},$$

where $\epsilon^0_{g^\Omega t}$ is an unanticipated shock and $\epsilon^p_{g^\Omega t-p}$ is a news shock that agents receive in period t about the innovation in time $t+p$.

4 Understanding the response of hours

The primary qualitative change in the response to a technology shock evidenced in our empirical analysis was the change in the co-movement of hours-worked with consumption. In addition, inventories maintained its comovement with hours over the subsamples. Moreover, in both subsamples, productivity evolved in a diffused manner, consistent with the interpretation of the technology shock as an anticipated or news shock. Before we confront the model with the data to study these features, we first highlight some key mechanisms of the model to understand the response of hours to such news shocks and frame our subsequent analysis.

We examine the key equations of labour market equilibrium to develop an expression the characterizes the response of hours-worked. We work with the linearizatons of the stationary transformations of the underlying non-stationary system, and introduce wedges into the model as stand-ins for several of the structural mechanisms in the model, where the wedges can be interpreted as endogenous equilibrium objects that represent deviations of some from a reference model. Additionally, in our linearizations we focus on a "news phase" where the

[15]We discuss details of the other shock processes in the appendix.

model model economy has received a news-shock about an increase in future TFP, but where the TFP shock has not yet materialized, and thus the linearized shocks are all zero.

We begin with the labor-supply equation

$$\xi\psi\Gamma_t v_t^{-\sigma} n_t^{\xi-1}\frac{f_t}{\bar{\lambda}_t} = \frac{\bar{w}_t}{\phi_t^{ls}}, \tag{10}$$

where we define ϕ_t^{ls} as a labour supply wedge between the marginal rate of substitution on the left-hand side and the real wage, and which in this model is equal to the wage markup term resulting from sticky wages which we discuss more below. Next we write the labor demand equation as

$$\bar{w}_t\phi_t^{ld} = \alpha\tau_t\frac{y_t}{n_t}, \tag{11}$$

where ϕ_t^{ld} is a labor demand wedge, equal to the knowledge capital markup wedge which we will discuss more below, and τ_t is the relative price of output, which itself acts as a wedge through its link to the inventory stocking equation. Finally we write the production function as

$$y_t = (n_t)^{\alpha}\left(u_t\frac{k_t}{g_t^k}\right)^{1-\alpha}\phi_t^e, \tag{12}$$

where ϕ_t^e is an efficiency wedge, equal to the input of knowledge capital in production, $h_t^{(1-\alpha_n-\alpha_k)}$. Linearizing (10), (11) and (12), eliminating the real wage and isolating hours-worked gives

$$\hat{n}_t = \frac{1}{(\xi-\alpha_n)}\left[\hat{\psi}_t^n - \hat{\phi}_t^{ls}\right] + \frac{1}{(\xi-\alpha_n)}\left[\hat{\tau}_t + \hat{\phi}_t^e - \hat{\phi}_t^{ld} + \alpha_k\hat{u}_t + \alpha_k\hat{k}_t\right], \tag{13}$$

where "hat's" denote percent-deviations of the transformed stationary variables from steady state, and $\hat{\psi}_t^n = \hat{\lambda}_t + \sigma\hat{v}_t - \hat{f}_t$ is a stand-in for the preference elements from the Jaimovich-Rebelo class of preferences.

Equation (13) describes the response of hours during the news-phase when no shocks other than the news shock are present. The terms in the first set of square brackets on the right hand side are labor supply shifters, and those in the second set of square brackets are labor demand shifters. Movements in the former and latter that are associated with an increase in the response of hours will tend to lower and raise the real wage respectively.

The coefficient term on both sets of brackets is a function of $\xi - \alpha_n$ which contribute to the relative slopes of the linearized labor supply and demand curves. Through the lens of the model, any change in the response of hours over the subsamples must show up somehow in the elements of this equation, and thus we use it to summarize the main possibilities for a change in the response of hours over the two subsamples.

Preference hypothesis. In principle, a change in preferences over time could account for the change in the response of hours over time, either through the parameter ξ which parameterizes the Frisch elasticity of labor and thus the amount of labour households are willing to supply for a given wage, or through the stand-in variable $\hat{\psi}_t^n$, which itself depends on the "wealth effect" parameter γ_f and consumption habits parameter b. The parameter γ_f is a particularly strong potential channel given the strong link between the wealth effect of expanding technology on consumption and the comovement of hours and consumption. As has been studied extensively in the literature, when γ_f is large, the standard income-effect on leisure means that while consumption rises in response to the the increase in lifetime wealth from the increase in technology, leisure also rises, and thus consumption and hours negatively co-move. When γ_f is near zero on the other hand, the income-effect on labor is minimal, such that consumption can increase in response to an increase in wealth without implying a corresponding drop in hours.[16] While changes in preferences over time are possible, we find large changes unlikely. Moreover, our empirical results suggest that consumption rises even more in the second subsample that the first, suggesting an even larger wealth effect in the second subsample than the first: if preferences where to change to account for this, they would more to overcome in the second subsample than the first to increase the response of hours. Nevertheless, in our quantitative analysis, we will allow for changes in ξ and γ_f, but we will limit their range within common limits in the literature to limit the possibility that large jumps in preferences alone explaining the change in hours.

Labor market frictions hypothesis. The direct effect of changes in labor market

[16]To see this most clearly, note that with $\gamma_f = 0$ and no consumption habits ($b = 0$), the stand-in variable $\hat{\psi}_t^n = 0$ and thus, it drops out of equation (13).

frictions in equation (13) work through the labor supply wedge ϕ_t^{ls}, which equals the wage markup from the wage Phillip's curve. This occurs through the changes in the parameter ω in the model, which measures the Calvo probability of not being able to optimally re-set household wages in a given period. An increase in ω would imply a more sluggish response of the real wage w_t^h facing the household, and thus a larger drop in ϕ_t^{ls} putting upward pressure on hours-worked.

Monetary policy hypothesis. A change in the stand of monetary policy in the model impacts the response of hours through at least two main channels: the first through interaction sticky wages, and the second through the real interest rate. For the first channel, as discussed extensively by Christiano et al. (2007), under sticky nominal wages, an inflation-targeting central bank directly impacts the real wage through the impact of inflation. Like the change in labor market frictions discussed above, this would manifest itself directly in equation (13) through the labor supply wedge ϕ_t^{ls}. The second channel impacts equation (13) indirectly through the general equilibrium impacts of the real interest rate on the variables in this equation, such as the preference term ψ_t^c, capacity utilization u_t, and the marginal cost of output τ_t through the impact of the real interest rate on inventory. For our quantitative analysis we allow for changes in the parameters ρ_r, ϕ_π, and ϕ_y of the monetary policy rule.

Credit market hypothesis. Like the second channel of monetary policy above, changes in credit market frictions in the model manifest themselves in equation (13) indirectly through the general equilibrium impacts of the real interest rate on variables in this equation, as well as the impacts of the credit spread capacity utilization u_t and choice of capital k_t. We note however that with capital predetermined on impact and sluggish in subsequent periods relative to the other variables such as capacity utilization, variation in k_t isn't likely to be a dominant factor in the response of hours in the initial few periods. Changes in credit market frictions in the model occur through changes in the parameter λ_b, which captures the proportion of capital a financial intermediary threatens to abscond. For our quantitative analysis, we allow for changes in λ_b by estimating steady-state leverage, ϕ_b, which based on our model solution and partial calibration maps directly to λ_b.[17]

Inventory hypothesis. The equilibrium optimal stocking condition in the model implies that the inventory sales ratio $\frac{X_t}{S_t}$ is given by

$$\frac{X_t}{S_t} = \chi(\tau_t, \mu_t^x), \tag{14}$$

where $\chi_\tau(t) = \frac{\partial \chi(\tau_t, \mu_t^x)}{\partial \tau_t} < 0$ and $\chi_{\mu^x}(t) = \frac{\partial \chi(\tau_t, \mu_t^x)}{\partial \mu_t^x} < 0$, and where μ_t^x is equal to the expected discounted value of future marginal costs, $\mu_t^x = (1-\delta_x)\,\beta E_t \frac{\lambda_{t+1}}{\lambda_t}\tau_{t+1}$. Faced with an increase in demand for sales triggered by the TFP news shock, distributors can satisfy the demand by some combination of running down inventories or purchasing new output at real price τ_t. The function $\chi(\cdot)$ depends on the parameter ζ, which measures the elasticity of demand for sales with respect to the relative stock of good varieties, and thus changes in ζ will impact the equilibrium response of τ_t and the associated response of hours in equation (13).

We note here also our empirical results suggesting that in both subsamples, hours and inventory co-move positively in both subsamples, no matter how hours and consumption comove. In terms of equation (13), we see that all else equal, hours varies positively with τ_t. A change in inventory management which implies meeting any increase in sales demand with relatively more new production relative to existing stocks of inventory would drive up the real price of output τ_t and thus hours-worked, implying upward force on both hours and inventories.

Knowledge capital hypothesis. The Intermediate Goods Firm's optimal labour choice is given by

$$w_t = \tau_t \alpha \frac{Y_t}{N_t} + q_t^h (1-\gamma_h) \frac{H_t^{\gamma_h} N_t^{1-\gamma_h}}{N_t}, \tag{15}$$

where q_t^h as the Lagrange multiplier on equation the knowledge capital accumulation equation (5) and has the interpretation as the marginal value of acquiring new knowledge capital in terms of expected future lifetime profits. q_t^h is in turn given by

$$q_t^h = \beta E_t \frac{\lambda_{t+1}}{\lambda_t} \left\{ (1-\alpha_n-\alpha_h)\tau_{t+1} \frac{Y_{t+1}}{H_{t+1}} + q_{t+1}^h \left(1 - \delta_h + \gamma_h \frac{H_{t+1}^{\gamma_h} N_{t+1}^{1-\gamma_h}}{H_t} \right) \right\}. \tag{16}$$

[17] Due to non-linearities in the steady-state relatons, estimating ϕ_b instead of λ_b reduces computational complexity.

The presence of internalized knowledge capital in the firm's technology adds an additional term into the firm's hours-worked first order condition (15) that shifts labor demand. News about future TFP increases the value of future knowledge through the impact of knowledge in future production according to (16), increasing the value of knowledge today, q_t^h, shifting out the firm's labor demand.

This manifests itself directly in equation (13) as a decrease in the labor demand wedge, ϕ_t^{ld} as the firm increases hours today in order to increase future knowledge capital, thus lowering its markup and current profits in the present in order to increase profits in the future. In our quantitative analysis we study changes in knowledge capital accumulation in the model occur through changes in the parameters δ_h, the depreciation rate of knowledge capital, and $1-\gamma_h$, the elasticity of labor in the production of new knowledge.

Other channels A change in the cost of adjusting capacity utilization, ϵ_u, directly impacts equation (13) through its impact on u_t. We can gain additional insight by using the Capacity Utilization firm's first-order condition for untilization, $\bar{r}_t = \delta\prime(u_t)$ and the Intermediate Goods Firm's first-order condition for capital services $\bar{r}_t = (1-\alpha)\tau_t \frac{y_t}{u_t \frac{k_t}{g_t^k}}$ to eliminate u_t in equation (13), resulting in

$$\hat{n}_t = \frac{1}{(\xi - \chi_u \alpha_n)} \left[\hat{\psi}_t^n - \hat{\phi}_t^{ls}\right] + \frac{1}{(\xi - \chi_u \alpha_n)} \left[\chi_u \hat{\tau}_t + \chi_u \hat{\phi}_t^e - \hat{\phi}_t^{ld} + (1 + \chi_u(\alpha_k - 1))\alpha_k \hat{k}_t\right], \tag{17}$$

where $\chi_u = \frac{1-\epsilon_u}{1-\epsilon_u-\alpha_k}$, where we can see that the primary role of capacity utilization in the model is to increase the elasticity of the other components of labour demand.[18] In addition to the direct effect, a change in the cost of utilizaton can impact equation (17) indirectly through the general equilibrium impacts of the real interest rate due to the influence of the cost of utilization on the return to capital and thus real interest rate.

In additional to the cost of utilization, a change in the cost of adjusting investment

[18]We note that the relative price of capital, q_{kt} does not play a role here as an independent shift factor for utilization, as in Greenwood et al. (1988) or Jaimovich and Rebelo (2009). Unlike in those models where the cost of utilization is incurred within the capital accumulation equation and this in terms of units of capital, in this model decentralization with separate financial, capital services and production sectors that subdivide this overall capital accumulation process, the utilization cost is incurred in terms of consumption units as in Christiano et al. (2005).

or capital, s'', impacts equation (17) indirectly through the general equilibrium impacts of the real interest rate on variables in this equation, working through the credit sector by influencing the price of capital and thus the return on capital.

5 Quantitative approach

We now detail our approach for quantitatively studying how the model might account for the changes in the response to technology shocks that we documented earlier in the empirical section. The approach is a hybrid of calibration, econometric and counterfactural exercises designed to illuminate candidate channels that could explain the changes in response across subsamples.

As in the empirical section, we break the sample data into two subsamples, 1954Q2-1983Q4 and 1984Q1-2019Q4, as an approximation for gradual or more abrupt structural change over that period. We then estimate a subset of the model parameters independently over the two subsamples, allowing the remaining subset of the parameters to remain fixed over both subsample. Our estimation for each subsample follows the approach of Christiano et al. (2005) such that the parameters are estimated by minimizing a measure of the distance between the model and empirical impulse response functions, conditional on a single structural shock in the model that corresponds to the shock identified in the empirical VAR. We use the results of the estimation over the two subsamples to highlight the key parameter changes, and then perform counterfactual exercises to explore the potential role of each key parameter in the change in impulse responses over the two subsamples.

5.1 Fixed parameters

The parameters that we hold fixed at calibrated value over both subsamples are detailed in Table 1 along with their calibrated values. Our choice values for this subset of parameters is guided by the existing literature, where we maintain comparability with Jaimovich and Rebelo (2009) and Schmitt-Grohe and Uribe (2012) for the aspects of the news shock mechanism and Lubik and Teo (2012) for the inventory component.

We set the household's discount factor β to 0.9957, which is implied by the real interest rate computed from average inflation and the federal funds rate over our sample period. We set the elasticity of intertemporal substitution as in Jaimovich and Rebelo (2009), $\sigma = 1$, and the consumption habits parameter b to 0.7.

On the firm side, we set the elasticity parameter in the production function to $\alpha = 0.64$ as in Jaimovich and Rebelo (2009), and the degree of decreasing-returns-to-scale (DRS) to labor and capital in production, $1 - \alpha_n - \alpha_k$, to 0.1, following Jaimovich and Rebelo (2009) and Schmitt-Grohe and Uribe (2012). For the parameters related to physical capital, we fix steady-state physical capital depreciation at $\delta = 0.025$.

The parameters related to inventories are based on the empirical estimates in Lubik and Teo (2012). The goods aggregator curvature parameter θ is set to 6.8, which results in a steady-state goods markup of 10%.

For the parameters related to the banking sector, following Gertler and Karadi (2011), we set θ_b, the determinant of a banker's life horizon, to 0.972. We then set w_b, the proportional transfer to enter bankers to 0.0038, such that for a steady-state leverage ratio of 4 (which we will estimate), the annualized steady state credit spread is about 100 basis points.

For the parameters related to the nominal side of the economy, we choose values consistent with the literature, setting the steady state wage markup to 10%, and wage and price indexation to 0.5.

Finally, a number of steady-state parameter values are implied by average values in the data, such as the (quarterly) steady-state growth rates of GDP g^y and the relative price of investment (RPI) g^{RPI}, which we find to be 0.43 and -0.58, respectively. We also set the steady-state government-spending ratio to output to $g/y = 0.18$ following Smets and Wouters (2007) and target a level of hours in steady state of 0.2, while steady-state capacity utilization is targeted at one.

Table 1: Calibrated model parameters

Description	Parameter	Value
Subjective discount factor	β	0.9975
Household elasticity of intertemporal substitution	σ	1
Habit persistence in consumption	b	0.70
DRS to N and K in production	$1-\alpha_n-\alpha_k$	0.1
Labor elasticity in production	α_n	0.64
Capital depreciation	δ_k	0.025
Goods aggregator curvature	θ	6.8
Price indexation	ι_p	0.5
Wage indexation	ι_w	0.5
Inventory depreciation	δ_x	0.05
Proportional transfers to entering bankers	w_b	0.0038
Survival rate of bankers	θ_b	0.972
Steady state government spending over output	g/y	0.18
Steady state hours	n	0.2
Steady state capacity utilization	u	1
Steady state wage markup	λ_w	1.1
Steady state GDP growth rate (in %)	g^y	0.42545
Steady state RPI growth rate (in %)	g^{rpi}	-.58203

5.2 Estimated parameters

The parameters that we estimate and that will be the focus of our analysis correspond to the parameters highlighted in our analytical analysis in Section 4 as being potentially important for the response of hours-worked. Let $\vartheta = (\xi, \gamma_f, \epsilon_u, s'', \zeta, \delta_h, \nu_h, \rho_\Omega, \phi_\pi, \phi_y, \rho_r, \kappa, \zeta_w, \phi_b)$ be the vector of these parameters. Then let $\psi(\vartheta)$ be the mapping from ϑ to the first ten elements of the model impulse response functions for the particular target variables under consideration, and $\hat{\psi}$ be the median of the estimated posterior distribution of the corresponding empirical impulse response functions. For each of the two subsamples, we then estimate ϑ as the solution to the problem

$$J_i = \min_{\vartheta_i} \left[\hat{\psi}_i - \psi(\vartheta_i)\right]' V_i^{-1} \left[\hat{\psi}_i - \psi(\vartheta_i)\right], \tag{18}$$

where V_i is a weighting matrix, and $i = 1, 2$ denotes the first or second subsample. We construct V_i using the variances of the posterior distribution of empirical impulse response functions along the diagonal for each subsample.

Due to small size of the empirical VARs, we are somewhat limited in the set of target variables. To best avoid issues resulting from arbitrary selection of these target variables, we perform the matching exercise described above for three separate sets of target variables: (i) consumption (C), output (Y), hours-worked (N) and investment (I)) (ii) consumption (C), output (Y), hours-worked (N) and inventory (X), and (iii) consumption (C), output (Y), hours-worked (N), inventory (X) and investment (I).

Table 2: Estimated model parameters: **IRF-match target variables C,Y,N,I**

Description	Parameter	Range	1954-1983	1984-2019
Determinant of Frisch elasticity of labor supply	ξ	$[1.1, 6]$	2.02	1.1
Wealth elasticity parameter (GHH/KPR pref)	γ_f	$[0.001, 0.5]$	0.0089	0.031
Elasticity of capacity utilization	ϵ_u	$[0.05, 5]$	0.054	0.26
Investment adjustment cost	s''	$[0.01, 2]$	0.22	0.47
Inventory taste shifter curvature	ζ	$[0.55, 0.7]$	0.62	0.70
Knowledge capital depreciation	δ_h	$[0.001, 0.999]$	0.16	0.16
Labor elasticity in knowledge capital	$1 - \gamma_h$	$[0.001, 0.9]$	0.90	0.45
TFP growth process persistence	ρ_Ω	$[0.001, 0.999]$	0.25	0.48
Taylor rule inflation	ϕ_π	$[1.1, 2.5]$	1.1	2.5
Taylor rule output	ϕ_y	$[0.05, 0.1]$	0.1	0.05
Taylor rule smoothing	ρ_r	$[0.5, 0.95]$	0.5	0.5
Price adjustment costs	κ	$[75, 300]$	100.40	100.5
Calvo wage parameter	ζ_w	$[0.5, 0.95]$	0.95	0.95
Steady state leverage	ϕ_b	$[3.5, 6]$	6	3.5

6 Quantitative Results

We now presents the results of the impulse response function matching as well as a series of counterfactual experiments.

6.1 Parameter Estimates

Tables 2, 3 and 4 respectively show the estimated values from the impulse response function matching exercise for the three sets of target variables, (i) C,Y,N,I, (ii) C,Y,N,X

Table 3: Estimated model parameters: **IRF-match target variables C,Y,N,X**

Description	Parameter	Range	1954-1983	1984-2019
Determinant of Frisch elasticity of labor supply	ξ	$[1.1, 6]$	1.58	1.36
Wealth elasticity parameter (GHH/KPR pref)	γ_f	$[0.001, 0.5]$	0.001	0.001
Elasticity of capacity utilization	ϵ_u	$[0.05, 5]$	0.064	0.22
Investment adjustment cost	s''	$[0.01, 2]$	0.17	0.01
Inventory taste shifter curvature	ζ	$[0.55, 0.7]$	0.70	0.70
Knowledge capital depreciation	δ_h	$[0.001, 0.999]$	0.26	0.999
Labor elasticity in knowledge capital	$1 - \gamma_h$	$[0.001, 0.9]$	0.90	0.16
TFP growth process persistence	ρ_Ω	$[0.001, 0.999]$	0.001	0.001
Taylor rule inflation	ϕ_π	$[1.1, 2.5]$	1.1	1.22
Taylor rule output	ϕ_y	$[0.05, 0.1]$	0.1	0.05
Taylor rule smoothing	ρ_r	$[0.5, 0.95]$	0.5	0.95
Price adjustment costs	κ	$[75, 300]$	99.7	99.2
Calvo wage parameter	ζ_w	$[0.5, 0.95]$	0.95	0.95
Steady state leverage	ϕ_b	$[3.5, 6]$	3.5	6

Table 4: Estimated model parameters: **IRF-match target variables C,Y,N,X,I**

Description	Parameter	Range	1954-1983	1984-2019
Determinant of Frisch elasticity of labor supply	ξ	$[1.1, 6]$	1.66	1.16
Wealth elasticity parameter (GHH/KPR pref)	γ_f	$[0.001, 0.5]$	0.0016	0.001
Elasticity of capacity utilization	ϵ_u	$[0.05, 5]$	0.07	3.35
Investment adjustment cost	s''	$[0.01, 2]$	0.22	0.01
Inventory taste shifter curvature	ζ	$[0.55, 0.7]$	0.70	0.70
Knowledge capital depreciation	δ_h	$[0.001, 0.999]$	0.31	0.999
Labor elasticity in knowledge capital	$1 - \gamma_h$	$[0.001, 0.9]$	0.90	0.13
TFP growth process persistence	ρ_Ω	$[0.001, 0.999]$	0.001	0.001
Taylor rule inflation	ϕ_π	$[1.1, 2.5]$	1.1	2.5
Taylor rule output	ϕ_y	$[0.05, 0.1]$	0.1	0.05
Taylor rule smoothing	ρ_r	$[0.5, 0.95]$	0.5	0.5
Price adjustment costs	κ	$[75, 300]$	99.9	99.6
Calvo wage parameter	ζ_w	$[0.5, 0.95]$	0.95	0.95
Steady state leverage	ϕ_b	$[3.5, 6]$	3.5	6

and (iii) C,Y,N,X,I[19]. In each table, Column 3 shows the search domain for each parameter of the minimization procedure, and columns 4 and 5 show the estimated values for the first and second subsamples respectively. While the estimates vary over the three tables, the three exercises suggest a generally consistent pattern of results. We draw attention to some key insights from the three tables.

First, several key parameter estimates change very little over the two substantials. Both the wage and price rigidity parameters ω_w and κ are nearly constant over the subsamples in all three tables, implying a moderate degree of price stickiness and high degree of wage stickiness. The very small value of γ_j over both samples in all three tables implies nearly "zero income effect" on labour supply, and is consistent with very small values found in studies such as in Bayesian estimation in Schmitt-Grohe and Uribe (2011) and Görtz et al. (2022). Small values of this parameter are typically important in models where comovement of hours-worked and consumption is important, as in many news-shock models, which is interesting given that in our empirical results hours-worked and consumption positively comove in the second sample but negatively comove in the first. Taken together, the above results suggest that a change in goods or labour market frictions, or changes in the wealth effect component of preferences were not likely a factor in the change in the response to technology shocks over the two samples.

Second, the estimates of several parameters changed substantially over the two subsamples, and the direction of change over the subsamples was consistent in all three tables. These include the disutility of working parameter ξ, knowledge capital depreciation δ_h, labor elasticity in knowledge capital $1-\gamma_h$, the Taylor rule inflation parameter ϕ_π, the Taylor rule output parameter ϕ_π elasticity of utilization parameter ϵ_u. For each table, while the disutility of working parameter ξ decreases in the second subsample implying a higher labor supply elasticity, the values nevertheless imply a high labor supply elasticity in both subsamples. Additionally, Tables 3 and 4 that include inventory in the target variable set show higher labor supply elasticities in the first subsample than in table 2 that does not contain,

[19] In this current current draft we are missing the standard errors in Tables 2, 3 and 4.

potentially capturing the empirical regularity of positive inventory and hours co-movement over both subsamples. The knowledge capital parameters δ_h and ν_h change substantially over the two samples, implying a change in the dynamics of knowledge capital accumulation in the second sample. The Taylor rule inflation parameter ϕ_y increases substantially in the second subsample, implying tighter monetary policy, and the Taylor rule output parameter ϕ_y decreases in the second subsample.

Finally, the pattern of estimates of a few parameters was inconsistent over the three tables. Tables 3 and 4 that include inventory as a target variable suggest that the investment adjustment cost parameter s'' decreased and steady state leverage ϕ_b increased in the second subsample, whereas Table 2 that excludes inventory suggests the reverse. Both the TFP growth process persistence parameter ρ_Ω and inventory taste shifter parameter ζ were constant over the two subsamples in two of the tables, but increased in the other.

6.2 Impulse Response Functions

The first and second panels in each of Figures 6, 7 and 8 show the impulse response functions exercise corresponding to the matching procedure for the first and second subsamples respectively. In the top row of each figure, the red solid line and shading indicate the median and 16% and 84% posterior bands respectively for the posterior distribution of VAR parameters for the first subsample. The red dashed line is the model IRF obtained from the IRF matching procedure over this period. In the second row of each figure, the blue solid line and shading indicate the median and 16% and 84% posterior bands respectively for the posterior distribution of VAR parameters for the second subsample. The blue dash-dotted line is the model IRF obtained from the IRF matching procedure over this period.

The figures show that the estimation procedure over the two subsamples captures the primary nature of the change in the empirical response over the two subsamples: in the model IRFs like as in the empirical IRFs, consumption rises in both samples, whereas hours and inventory fall in the first subsample and rise in the second. Implicit within this is that the model captures the tendency discussed in our empirical analysis for hours to co-move with inventories, even when hours doesn't co-move with consumption.

Figure 9 shows the IRFs to a larger set of variables using both the parameters estimated from the first and second subsample, for each of the three target variable cases. The top panel of the figure shows the C,Y,N,I target variable case, the middle panel the C,Y,N,X target variable case, and the bottom panel the C,Y,N,X,I target variable case. For each case, the red dotted line and blue dash-dottled line show the IRFs using parameters estimated on the first and second subsamples respectively.

The top panel shows that the C,Y,N,I targeting procedure is able to capture the general pattern of an initial fall in hours and investment and rise in credit spreads in the first subsample, and the reswerve in the second subsample, as in our empirical results. The model also captures the larger positive response of the real wage in the second subsample compared to the first as in our empirical results, but not an initial drop in wages in the first subsample. Without inventory as a target, the model is unable to match the rise in inventory in the second subsample.

The middle panel shows that the C,Y,N,X targeting procedure is able to capture the general pattern of a fall in inventory along with hours in the first subsample, and rise in the second, but now investment falls and spreads rise in both subsamples, contrary to our empirical results. The case does not capture the overall pattern of wages.

The lower panel shows the C,Y,N,X,I targeting procedure captures the general pattern of hours, inventory, investment and spreads from the empirical results. Like the middle panel however, the case does not captures the overall pattern of wages.

The above discussion reveals a tension in the model between matching the IRFs of inventory and investment. In general, including inventory in the target variable set makes it more difficult to capture the investment response (and thus spreads response also), and vice versa. Adjustment costs to investment/capital and the interaction with the banking sector play a key role in this. In general, higher adjustment costs increase net-worth during a boom, increasing the demand for capital, and putting upward pressure on the real interest rate. For investment, the optimal stocking equation depends on current marginal costs but also expected future marginal costs, discounted using the real interest rate. With marginal costs growing during the boom, a higher real interest rate means a lower weight is placed on expected future marginal costs in the inventory decision, providing less incentive to increase inventories in the present to avoid having to build them up in a higher cost future.

Subsample 1: 1954Q2-1983Q4.

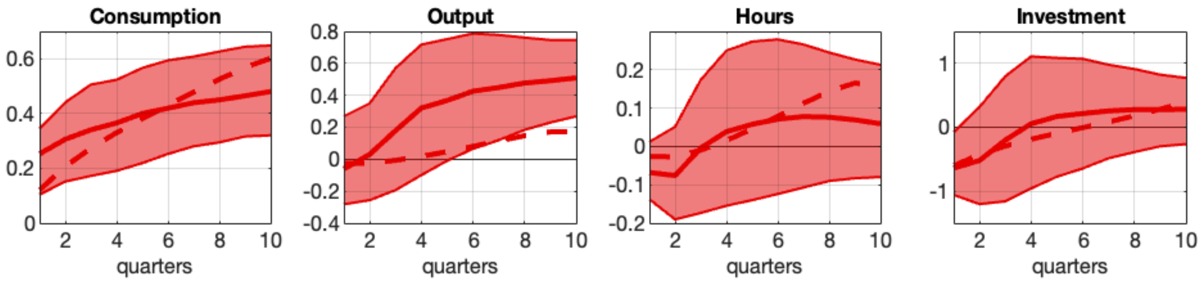

Subsample 2: 1984Q1-2019Q4.

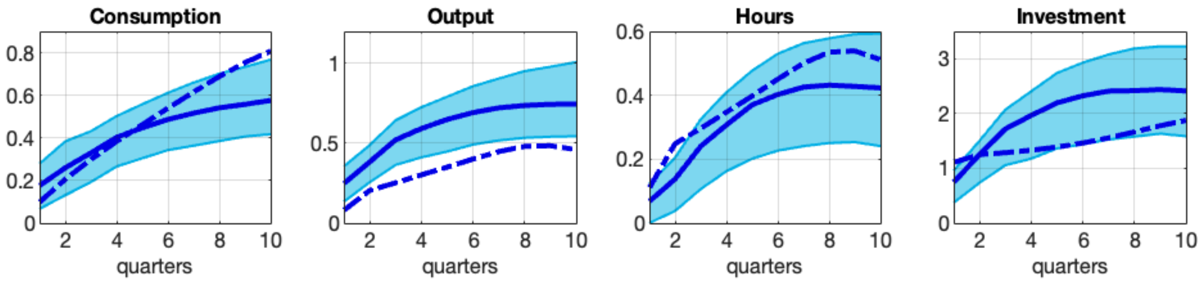

Figure 6: **VAR and model based IRF to permanent productivity shock: IRF match target variables C,Y,N,I.** Red (blue) solid line is the median VAR estimate on the first (second) subsample and the shaded areas are the 16% and 84% posterior bands generated from the posterior distribution of empirical VAR parameters. Red (blue) dashed (dash-dotted) line is model IRF to 12 period ahead news shock using IRF-matching procedure on the first (second) subsample. The units of the vertical axes are percentage deviations.

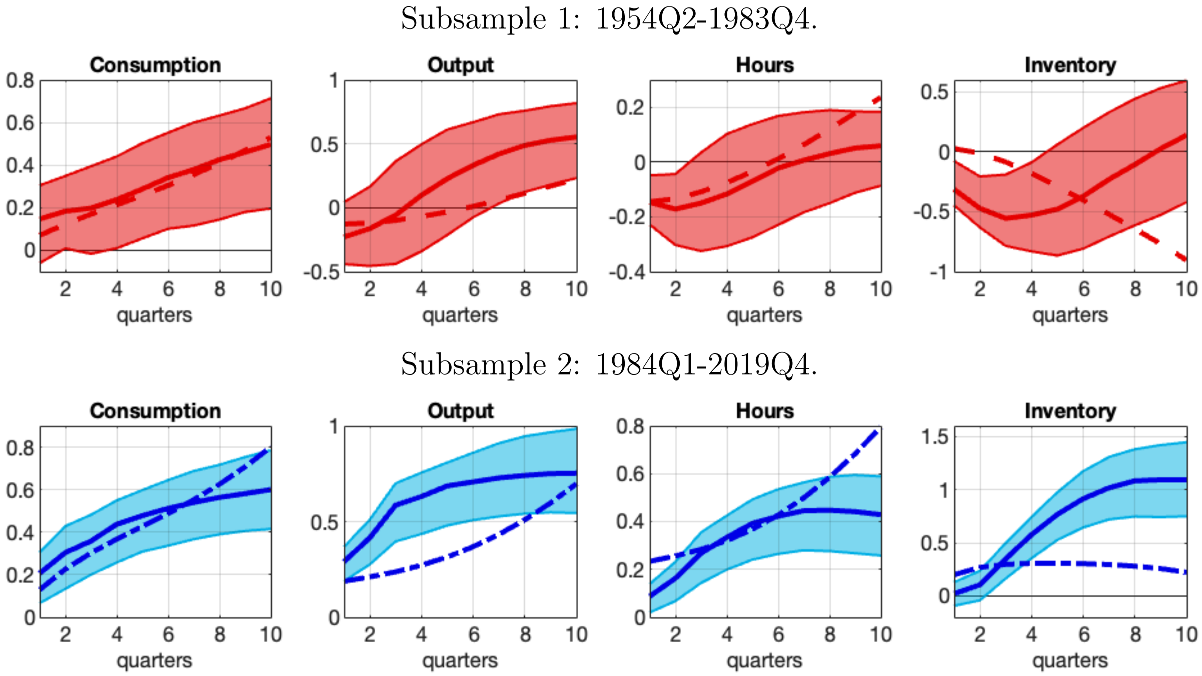

Figure 7: **VAR and model based IRF to permanent productivity shock: IRF match target variables C,Y,N,X.** Red (blue) solid line is the median VAR estimate on the first (second) subsample and the shaded areas are the 16% and 84% posterior bands generated from the posterior distribution of empirical VAR parameters. Red (blue) dashed (dash-dotted) line is model IRF to 12 period ahead news shock using IRF-matching procedure on the first (second) subsample. The units of the vertical axes are percentage deviations.

Subsample 1: 1954Q2-1983Q4.

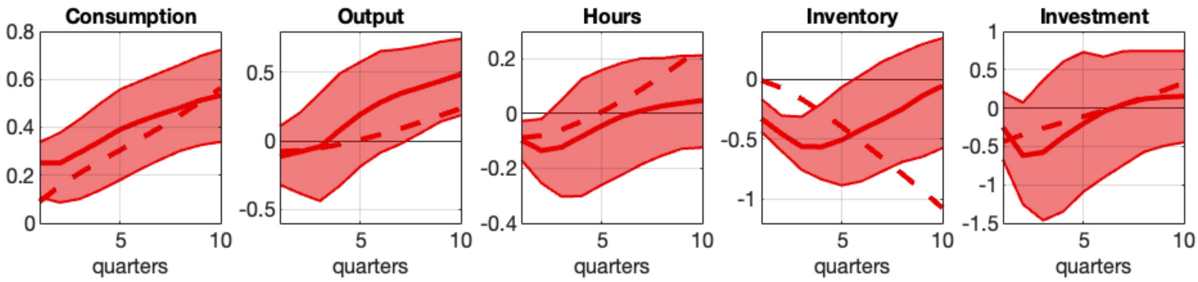

Subsample 2: 1984Q1-2019Q4.

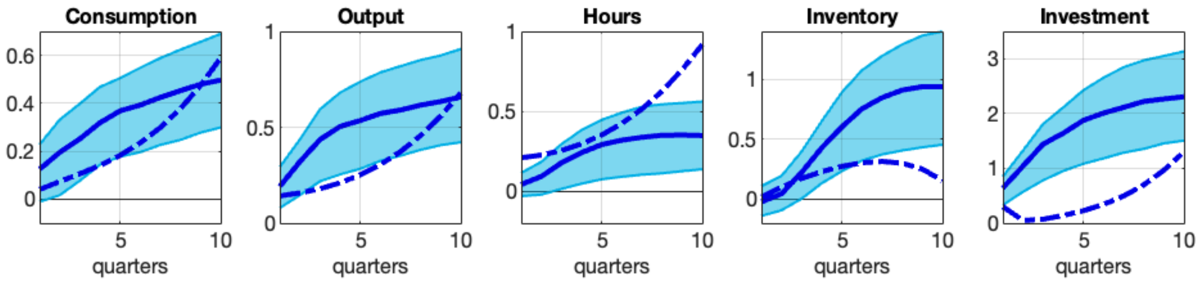

Figure 8: **VAR and model based IRF to permanent productivity shock: IRF match target variables C,Y,N,X,I.** Red (blue) solid line is the median VAR estimate on the first (second) subsample and the shaded areas are the 16% and 84% posterior bands generated from the posterior distribution of empirical VAR parameters. Red (blue) dashed (dash-dotted) line is model IRF to 12 period ahead news shock using IRF-matching procedure on the first (second) subsample. The units of the vertical axes are percentage deviations.

IRF-match target variables: C,Y,N,I

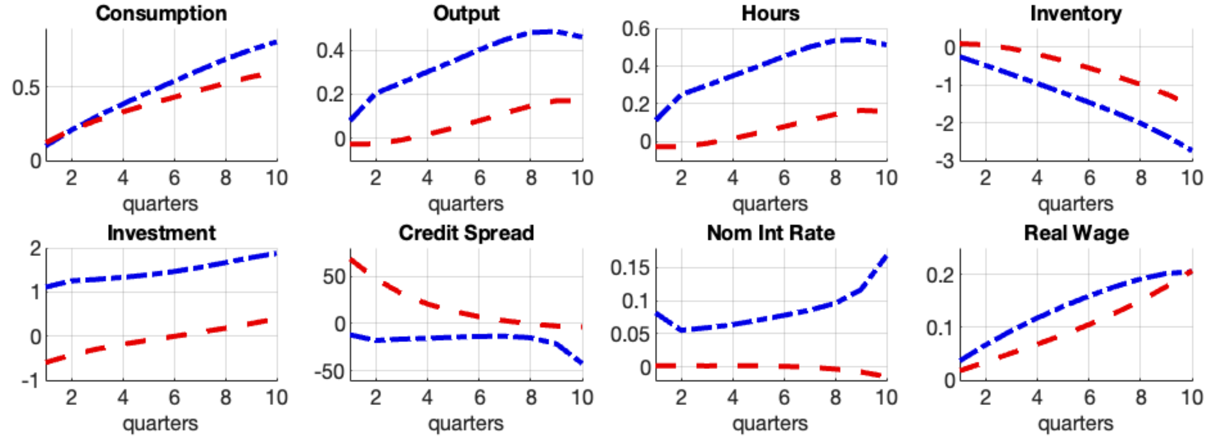

IRF-match target variables: C,Y,N,X

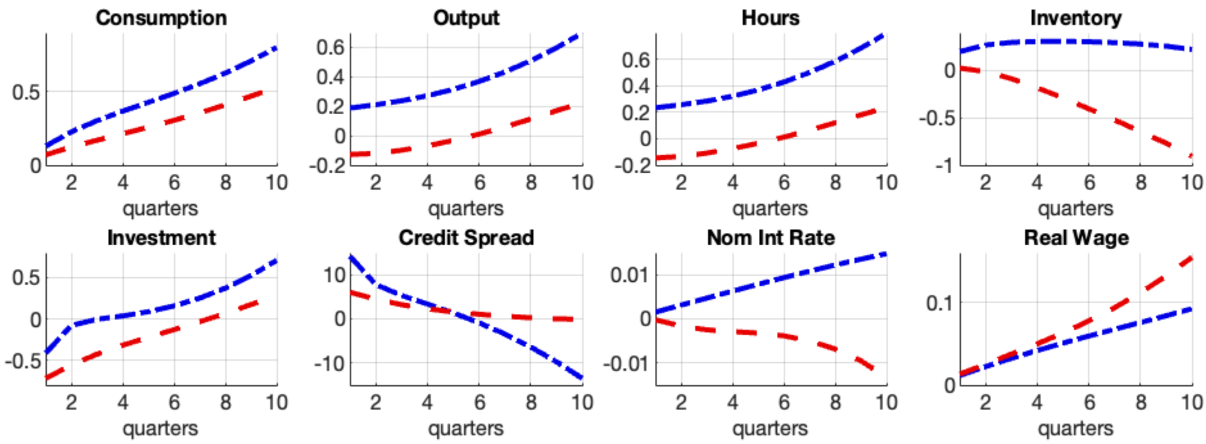

IRF-match target variables: C,Y,N,X,I

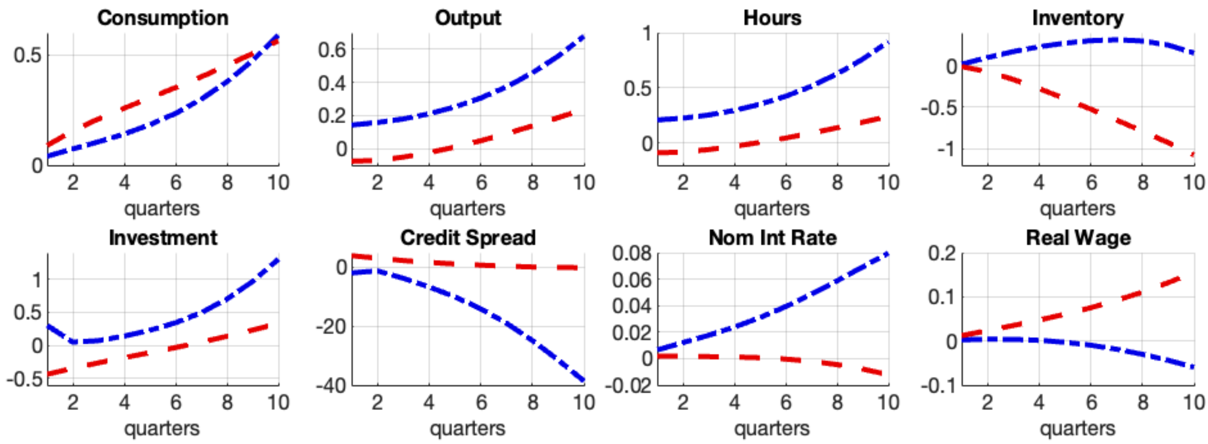

Figure 9: **Model based IRF to permanent productivity shock: IRF match target variables as indicated for each of three panels.** Red (blue) dashed (dash-dotted) line is model IRF to 12 period ahead news shock using IRF-matching procedure on the first (second) subsample. The units of the vertical axes are percentage deviations.

6.3 Counterfactual Experiments

The above results reveal several parameter changes that together account for the general nature of the change in model IRFs over the two subsamples. Was the empirical change in the response of the IRFs due to change in a combination of underlying factors - captured in the model by a change in multiple parameters - or could a few factors (parameters) have been dominant? We now qualitatively explore this question through the lens of the model, examining the role of each of the key parameter changes through counterfactual experiments to determine if one or a small subset of these factors were dominant. To do this, we calibrate the model to the parameterization obtained from the estimation results for the first subsample, and then one-by-one we vary the parameters to study their role delivering the change in IRF in the second subsample. In the process, we highlight several important structural features of the model that are important for accounting for the changes.

Panels A to J in Figures 10 to 11 show this exercise for the group of parameters identified in Section 6.1 whose estimates changed materially over the two subsamples in at least one of the three cases: $\xi, s'', \delta_h, 1-\gamma_h, \phi_\pi, \phi_y, \epsilon_y, s'', \phi_b, \rho_\omega, \zeta$. We note that the panels ordered A to J are the parameters identified in 6.1 with consistent changes in the direction of estimatates over the two subsamples over all three cases. In each panel, the blue solid line and shading indicate the median and 16% and 84% posterior bands respectively for the posterior distribution of VAR parameters for the second subsample, and the blue dash-dotted line is the model IRF obtained from the IRF matching procedure over this same period (i.e. the blue shading and blue dash-dotted line in each panel reproduces the blue shading and dash-dotted line of Figure 8). The red dashed line in each panel is the model IRF obtained from the IRF matching procedure over the first subsample (i.e. the red dashed line in each of these figures reproduces the red dashed line of Figure 8). The black dotted line in each panel is then the counterfactual model IRF using all parameters values obtained from the IRF matching procedure over the *first* subsample, except for the parameter in question where the value is set to the value obtained form the IRF matching procedure over the *second* subsample (ie with no counterfactual change in the parameter of question, the black dotted line would coincide exactly with the red dashed line). In other words, the dotted black line shows the extent to which the single parameter change in question can shift the model IRF from the red dashed line to the blue dash-dot line.

Panel A of Figure 10 shows that the model IRFs from changing ξ to the value estimated over the second subsample actually moves the black dotted line for all four variables further below the red dash line, and thus further away from the targeted empirical IRFs. Panel B of the same figure suggests the same for δ_h. These results suggest that while these parameter changes may help the fit of the overall model over the second subsample when all variables change also, on their own and given the direction of their change, they are not likely the dominant factors in accounting for the change in the IRFs over the two samples.

Panel C Figure 10 show that in changes in ν_h moves the response of hours and inventory substantially toward or beyond the blue dash-dotted line. Panels D and E show that changes in ϕ_π and ϕ_y also move the response of hours in this same direction, but have little impact on inventory. These result suggest changes in ν_h, ϕ_π and δ_b may be dominant factors in accounting for the change in the IRFs over the two samples.

Panel F of Figure 11 shows that a change in the utilization cost moves the response of both hours and investment towards second subsample response, but moves the response of inventory slightly away from it. We note that the change here represents an *increase* in the cost of utilization, and thus dominant channel of this changes is likely through the general equilibrium effects in the banking sector whereby the increased cost of utilization drives up the return to capital and increases investment demand, similar to the effect of an increase in adjustment costs discussed earlier.

Panel G of Figure 11 shows mixed results for the investment adjustment cost parameter s''. While it moves the response of consumption, output and hours further from target, it moves inventory closer, via the mechanism discussed earlier whereby adjustment costs work through the banking sector and create a tension between inventory and investment. Thus this factor may be important for helping to maintain the co-movement between hours and inventories over the two samples. We note however that while the estimates of this parameter decreased in the second subsample in the IRF matching exercises that included inventory (as in the figure here), it increased in the matching exercise that excluded inventory. Our results are thus somewhat inconclusive about the role of this parameter.

Panels H, I and J of Figure 11 show relatively small impact of changes in ϕ_b, ρ_Ω and and ζ respectively over the two subsamples. This result suggests that on their own changes in these three parameters are not likely dominant factors in accounting for the change in the IRFs over the two samples.

Overall, our results from this exercise suggest that changes in the nature of knowledge capital accumulation ν_h, tighter monetary policy in response to inflation (ϕ_π), looser monetary policy in response to output (ϕ_y), an increase in the cost of utilization (ϵ_u) all potentially contributed to the change in the response of technology shocks over the two subsamples.

7 Re-visiting the model-based hypotheses

With our evidence in hand from the various model-based exercises above, we can now circle back to the potential model-based hypotheses concerning the sources of the change in the response to technology shocks that we outlined in Section 4.

Preference hypothesis Our estimates suggest that the wealth effect parameter γ_f did not change materially over the subsamples, and while the disutility or working parameter ξ changed over the subsamples, our experiments suggest that change did not move the response in the correct direction. Thus our results suggest that a change in preferences was not likely a dominant source of the change.

Labor Market Frictions Hypothesis Our estimates suggest a high degree of wage rigidities that did materially change over both subsamples. While this rigidity is a key propagation mechanism, the lack of a change suggests that a change in labor market frictions where likely not a source of the change in the impact of technology.

Monetary Policy Hypothesis Our results suggest that changes in the stance of monetary policy between the two subsamples was a likely contributor to the change in the impact

Panel A: Frisch elasticity parameter, ξ, counterfactual.

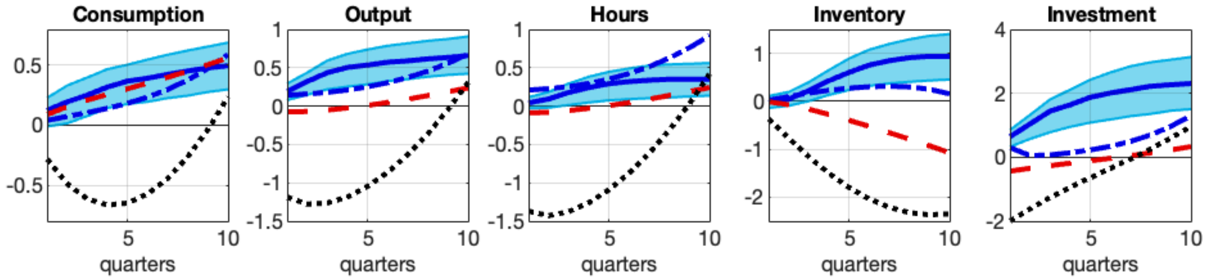

Panel B: Knowledge capital depreciation, δ_h, counterfactual.

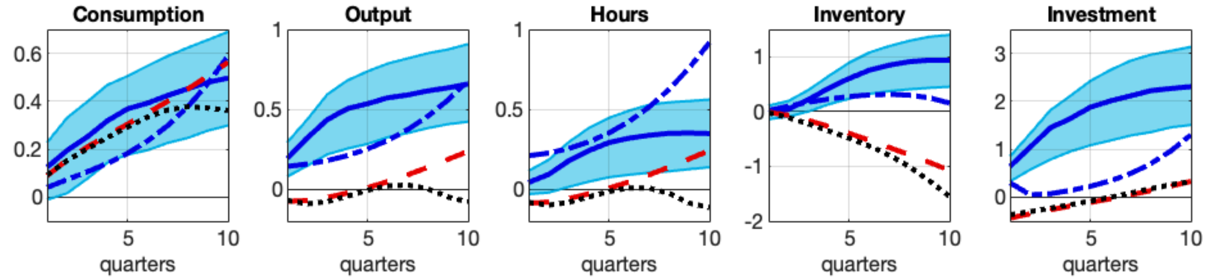

Panel C Labor elasticity in knowledge capital, ν_h, counterfactual.

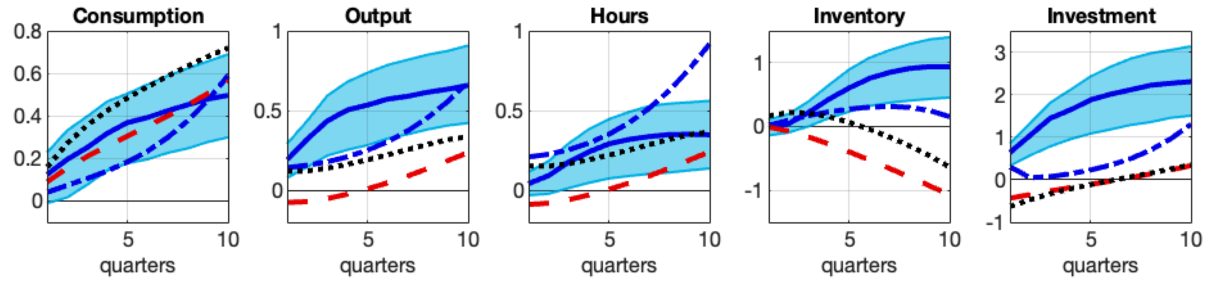

Panel D: Taylor rule inflation, ϕ_π, counterfactual.

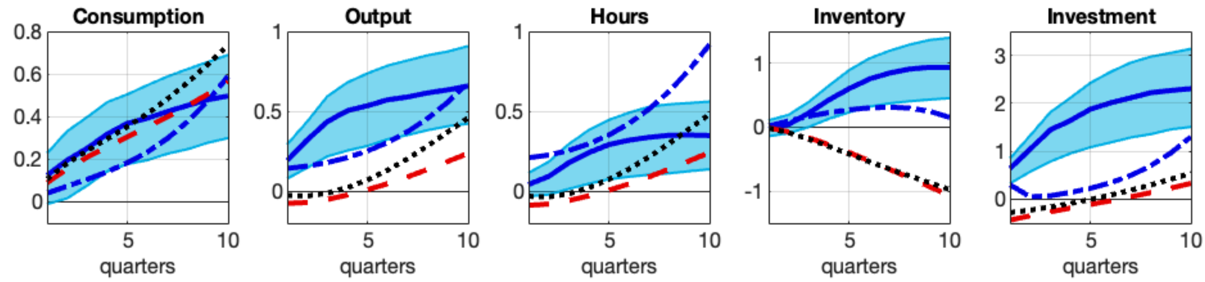

Panel E: Taylor rule output, ϕ_y, counterfactual.

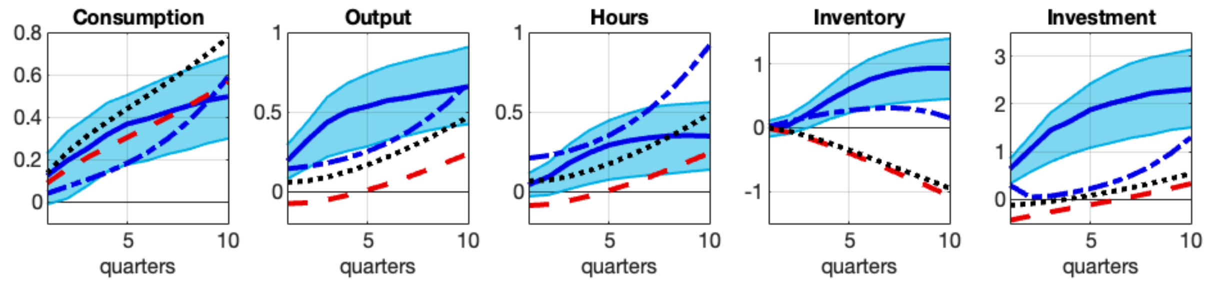

Figure 10: **IRFs to 12 period out permanent TFP news shock.** Blue solid line (shaded blue areas) is the median (the 16% and 84% posterior bands generated from the posterior distribution) of empirical VAR parameters using data on 1984Q1-2019Q4. Blue dash-dotted line is model IRF using IRF-matching procedure on data sample 1984Q1-2019Q4. Red dashed line is model IRF using IRF-matching procedure on data sample 1954Q2-1983Q4. Black dotted line is counterfactual model IRF using all parameter values obtained from IRF matching procedure on first subsample except for parameter in question which is set subsample 2 value or otherwise indicated. The units of the vertical axes are percentage

Panel F: Utilization cost, ϵ_u, counterfactual.

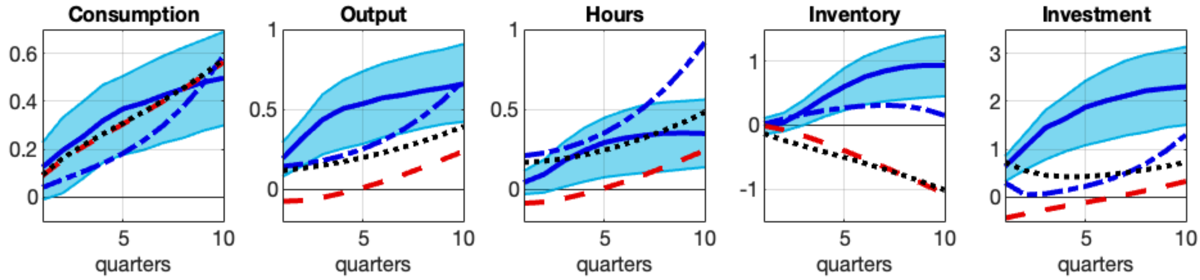

Panel G: Investment adjustment cost, s'', counterfactual.

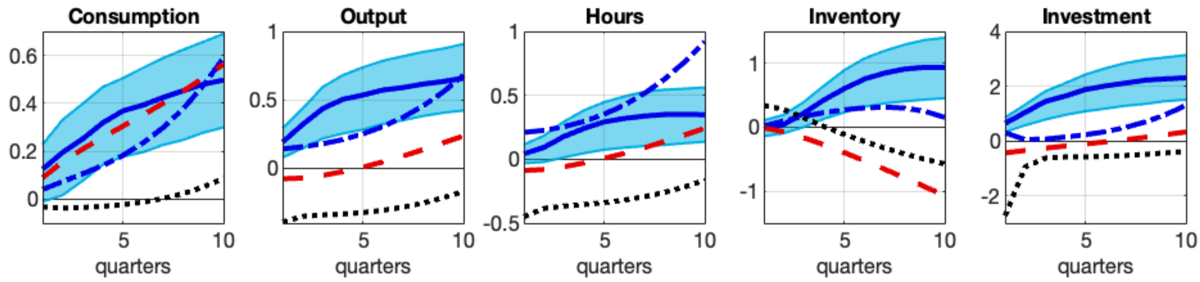

Panel H: Steady state leverage, ϕ_b, counterfactual.

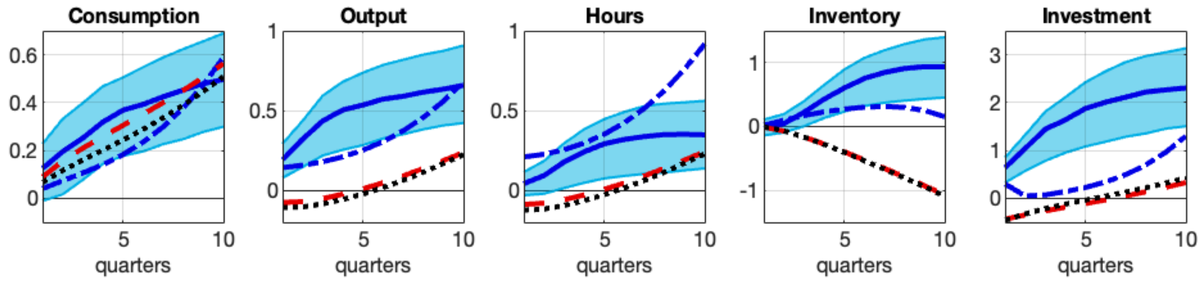

Panel I: TFP growth process persistence, ρ_Ω, counterfactual: $0.001 \rightarrow 0.23$.

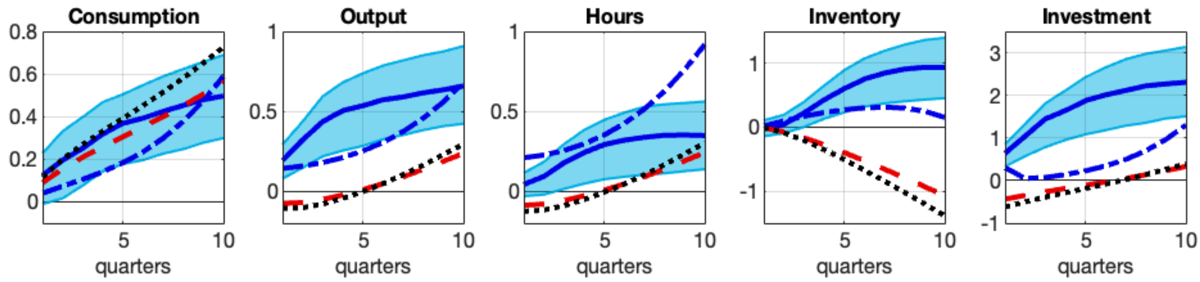

Panel J: Inventory taste shifter, ζ, counterfactual: $0.7 \rightarrow 0.73$

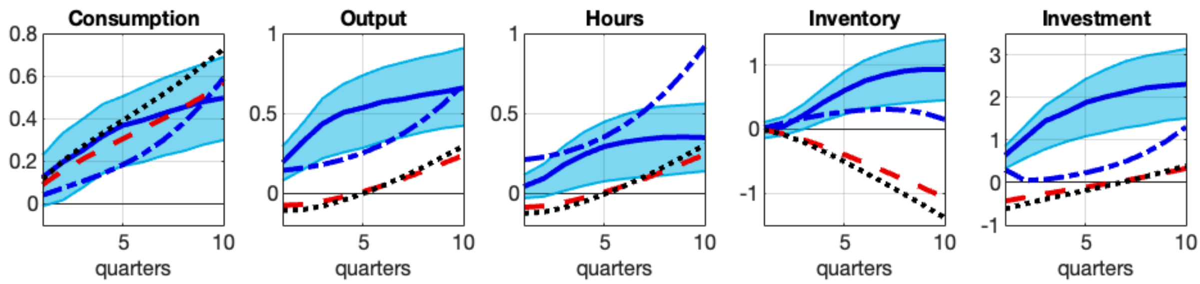

Figure 11: **IRFs to 12 period out permanent TFP news shock.** Blue solid line (shaded blue areas) is the median (the 16% and 84% posterior bands generated from the posterior distribution) of empirical VAR parameters using data on 1984Q1-2019Q4. Blue dash-dotted line is model IRF using IRF-matching procedure on data sample 1984Q1-2019Q4. Red dashed line is model IRF using IRF-matching procedure on data sample 1954Q2-1983Q4. Black dotted line is counterfactual model IRF using all parameter values obtained from IRF matching procedure on first subsample except for parameter in question which is set subsample 2 value or otherwise indicated. The units of the vertical axes are percentage

of technology. This includes estimates implying tighter monetary policy in response to inflation (ϕ_π) and looser monetary policy in response to output (ϕ_y) over all three IRF matching exercises, and a evidence from our experiments suggesting these changes were associated with the correct pattern of change of hours-worked in the second subsample.

Credit Hypothesis Our estimates suggest an inconsistent direction of change in the steady state leverage parameter ϕ_b over the three IRF matching exercises, and our experiments suggest the magnitude of change in the response of the IRFs was not material. Thus a change in credit market frictions over time was not likely a source of the change in the impact of technology. Our evidence does however suggest that the banking sector was a powerful propagation mechanism of the change in other parameters.

Inventory Hypothesis Our estimates suggest an inconsistent direction of change in the inventory taste shifting parameter ζ over the three IRF matching exercises, and our experiments suggest the magnitude of change in the response of the IRFs was not material. Thus a change in inventory management practices as evidence through this parameter was not likely a source of the change in the impact of technology.

Knowledge Capital Hypothesis Our results suggest that changes in the nature of intangible capital accumulation was a likely contributor to the change in the impact of technology over time in matching the response of both hours and inventories, in particular due to the change in the elasticity of labor in knowledge capital accumulation. Taking this modeling mechanism literally, one interpretation is a change in the way firms learned about organizing the inputs into production as the processes of production changed rapidly heading into the information and technology revolution of the 1980's and 1990's. Taken more broadly, the effect can be seen to be symptomatic of the emergence of a labor-demand side wedge that feeds into an efficiency wedge in production.

Other channels Our results suggest that an increase in the cost of utilizing capital contributed to the change in the the response of technology. Our estimates suggested an increase in this cost over all three matching exercises, and our experiments showed that the magnitude of magnitude of change in the response of the IRFs was material. The likely mechanisim for this change in the model was through the credit sector by influencing the return to capital.

8 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 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 literature attempting to understand the role and importance of technology shocks. We take an agnostic view of the presence of surprise verss 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.

We when add to the theoretical literature to study the source of the changes in the response of technology through the lens of a rich structural model. We use both an IRF matching procedure and model experiments to evaluate various different hypotheses for the change. Our results suggest that the change in the response of technology over time was likely some combination of a change in the stance of monetary policy, a change in the nature of intangible capital accumulation, and a change in the cost of utilizing capital.

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 Kimpball, 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.

Bils, M. and Kahn, J. A. (2000). What inventory behavior tells us about business cycles. *American Economic Review*, 90(3):458–481.

Chang, Y., Gomes, J. F., and Schorfheide, F. (2002). Learning-by-Doing as a Propagation Mechanism. *American Economic Review*, 92(5):1498–1520.

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., Ilut, C., Motto, R., and Rostagno, M. (2007). Monetary policy and stock market boom-bust cycles. *manuscript, Northwestern University.*

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

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.

Cooper, R. and Johri, A. (2002). Learning-by-doing and aggregate fluctuations. *Journal of Monetary Economics*, 49(8):1539–1566.

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.

d'Alessandro, A., Fella, G., and Melosi, L. (2019). Fiscal Stimulus with Learning-By-Doing. *International Economic Review*, 60(3):1413–1432.

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.

Gertler, M. and Karadi, P. (2011). A model of unconventional monetary policy. *Journal of Monetary Economics*, 58(1):17 – 34.

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.

Greenwood, J., Hercowitz, Z., and Huffman, G. (1988). Investment, capacity utilization, and the real business cycle. *The American Economic Review*, 78:402–217.

Gunn, C. and Johri, A. (2011a). News and knowledge capital. *Review of Economic Dynamics*, 14(1):92–101.

Gunn, C. and Johri, A. (2011b). News and knowledge capital. *Review of Economic Dynamics*, 14(1):92–101.

Görtz, C., Gunn, C., and Lubik, T. A. (2022). Is there news in inventories? *Journal of Monetary Economics*, 126:87–104.

Jaimovich, N. and Rebelo, S. (2009). Can news about the future drive the business cycle? *American Economic Review*, 99(4):1097–1118.

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. (2021). Revisions in utilization-adjusted tfp and robust identification of news shocks. *Review of Economics and Statistics*, 103(2):216–235.

Lubik, T. A. and Teo, W. L. (2012). Inventories, inflation dynamics and the new keynesian phillips curve. *European Economic Review*, 56(3):327–346.

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.

Schmitt-Grohe, S. and Uribe, M. (2011). Business cycles with a common trend in neutral and investment-specific productivity. *Review of Economic Dynamics*, 14:122–135.

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

Shea, J. (1998). What do technology shocks do? *NBER Macroeconomics Annual*, 13:275–310.

Smets, F. and Wouters, R. (2007). Shocks and frictions in US business cycles: A Bayesian DSGE approach. *American Economic Review*, 97(3):586–606.

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:

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

$A(L) = I + A_1 L + ... + A_p L^p$ 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 $B_0 B_0' = \Sigma$. B_0 can also be written as $B_0 = \widetilde{B}_0 D$, where $\widetilde{B}_0$ 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 (2021), 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 prior 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 $\widetilde{B}_0$, 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:

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

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

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

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. $\widetilde{B}_0\gamma$ 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 $\widetilde{B}_0$ 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:[20]

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

We restrict γ to have unit length to be a column vector of an orthonormal rotation matrix of the Choleski decomposition of the reduced-form variance covariance matrix.

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

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 a standard Minnesota prior as commonly used in the literature. It is of the form

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

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,

$$\underline{V}_{i,jj} = \begin{cases} \frac{\underline{a}_1}{p^2} \text{ for coefficients on own lags} \\ \frac{\underline{a}_2 \sigma_{ii}}{p^2 \sigma_{jj}} \text{ for coefficients on lags of variable } j \neq i \\ \underline{a}_3 \sigma_{ii} \text{ for intercepts} \end{cases}$$

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 hyper-parameters $\underline{a}_1, \underline{a}_2, \underline{a}_3$. 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.[21]

B Additional VAR Evidence

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

Labor Market Responses. Figure 12 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 with the decline in hours-worked documented for the first subsample, Figure 12 documents a decline in the labor force participation rate and a rise in the unemployment rate. For

[21] The 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.

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.

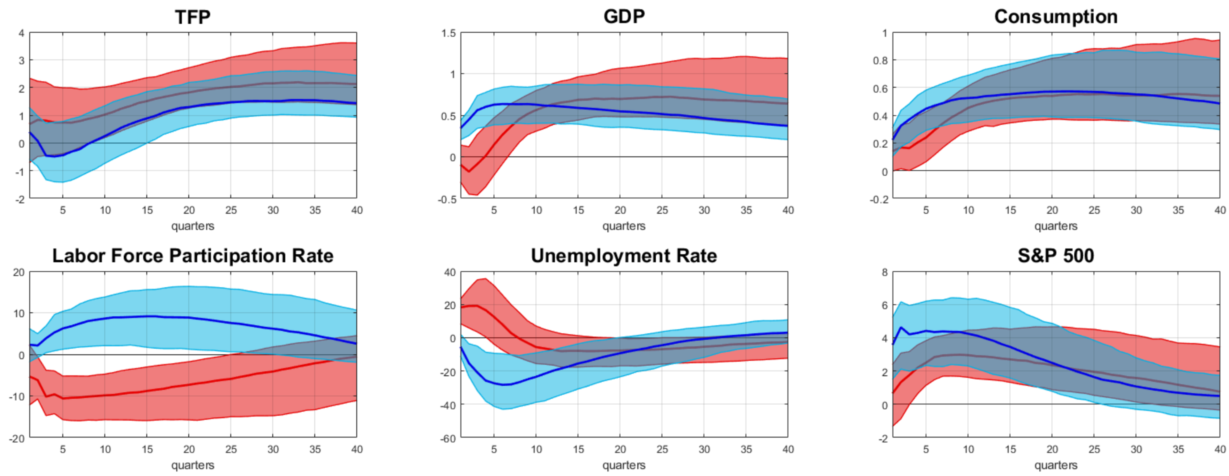

Figure 12: **IRF to TFP shock.** First subsample 1954Q2-1983Q4 (red), second subsample 1984Q1-2019Q4 (blue). The solid line is the median and the shaded colored areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

An Alternative Measure for Technology. Figure 13 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 13 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.

Responses over the Entire Sample. Figure 14 shows the responses to a technology shock over the whole sample (1954Q2-2019Q4). All macroeconomic aggregates increase

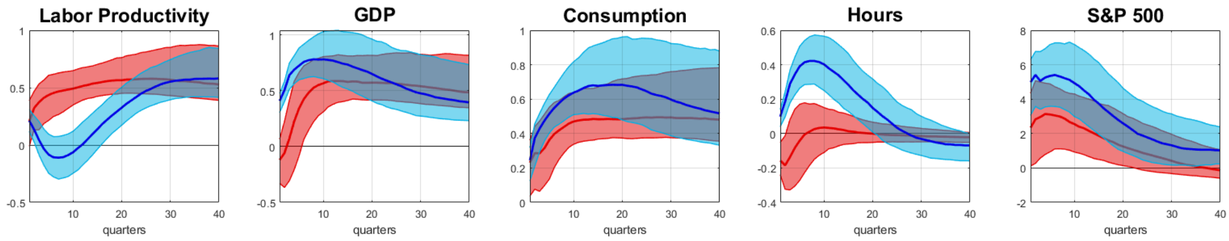

Figure 13: **IRF to shock that maximizes variation in labor productivity.** First subsample 1954Q2-1983Q4 (red), second subsample 1984Q1-2019Q4 (blue). The solid line is the median and the shaded colored areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

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.

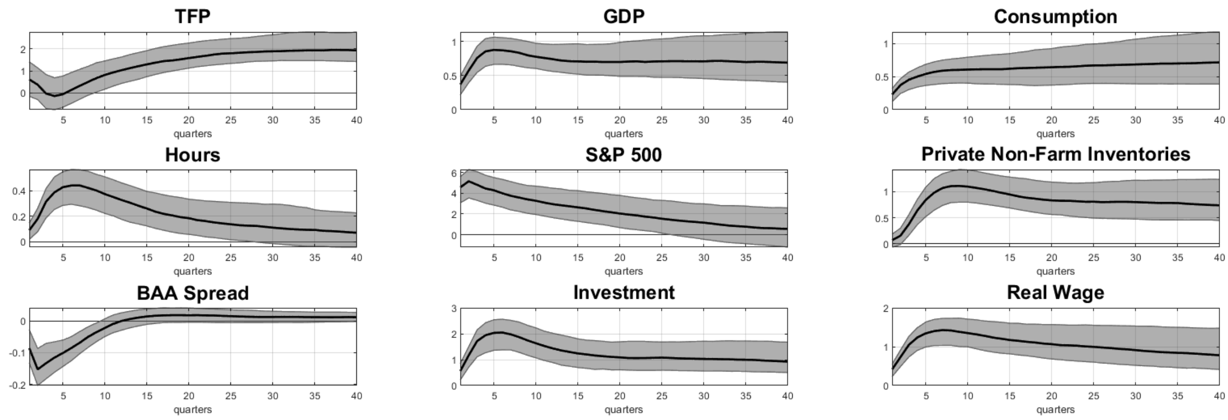

Figure 14: **IRF to TFP shock.** Entire sample 1954Q2-2019Q4 . The solid line is the median and the shaded areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

Robustness for Rolling Window Analysis. Figure 15 shows the median and posterior bands of impact responses for selected variables over different samples. Results are consistent with those in Figure 4 in the main body which shows the most extreme response within the first ten quarters. For hours, inventories and investment, it is evident that impact responses move away from the negative territory over the rolling window analysis. On impact the BAA becomes negative particularly once the window includes the time around the financial crisis.

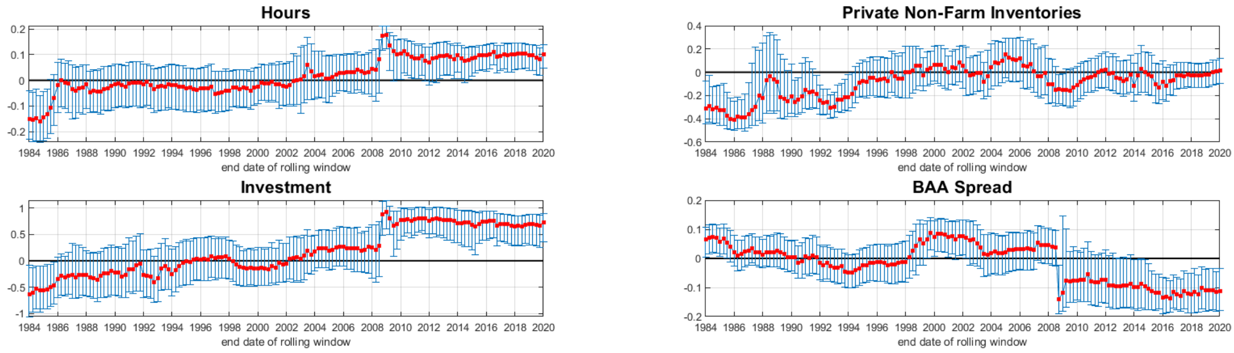

Figure 15: **Impact IRF responses to a TFP shock for rolling windows.** First rolling window sample is 1954Q2-1983Q4. The window is shifted up to 2019Q4. We display the median (red dot) and the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical 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.

Figures 17 and 16 show statistics corresponding to those in Figure 4 in the main body. They shows for each rolling window the maximum or minimum IRF (whichever is largest in absolute terms) within the first ten quarters to a TFP shock for rolling window. Figures 17 and 16 differ from the one depicted in the main body in that they consider a shorter rolling window of 90 and 100 quarters, respectively, instead of 119 quarters.

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.Luis 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.

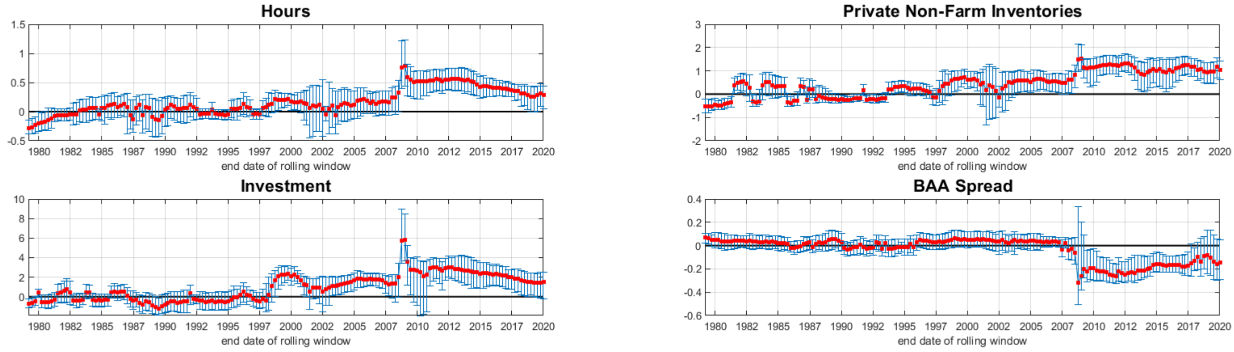

Figure 16: **Maximum/minimum (whichever is largest in absolute terms) IRF response within the first ten quarters to a TFP shock for rolling window.** First rolling window sample is 1954Q2-1979Q1 (100 quarters). The window is shifted up to 2019Q4. We display the median (red dot) and the 16% and 84% posterior 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.

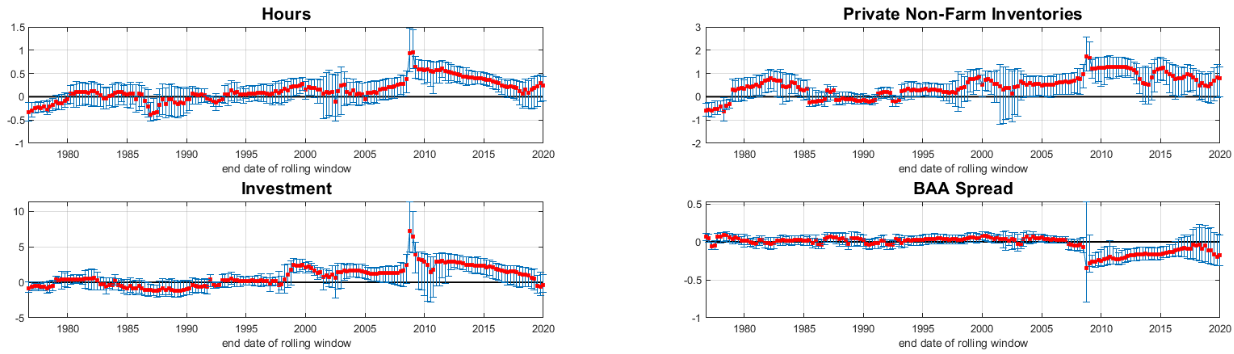

Figure 17: **Maximum/Minimum (whichever is largest in absolute terms) IRF response within the first ten quarters to a TFP shock for rolling window.** First rolling window sample is 1954Q2-1976Q3 (90 quarters). The window is shifted up to 2019Q4. We display the median (red dot) and the 16% and 84% posterior 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.

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, Federal 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; $https://fred.stlouisfed.org/series/A373RX$
Civilian Noninstitutional Population: U.S. Bureau of Labor Statistics, Population Level [CNP16OV], retrieved from FRED, Federal Reserve Bank of St. Louis; $https://fred.stlouisfed.org/se$
Non-farm Business Sector: Compensation Per Hour: U.S. Bureau of Labor Statistics, Nonfarm 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 ($http://www.econ.yale.edu/\ shiller/d$
The utilization adjusted TFP data and the series for capacity utilization can be accessed at $www.frbsf.org/economic-research/economists/jfernald/quarterly\_tfp.xls$.

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.