Mainly Employment: Survey-Based News and the Business Cycle

Riccardo M. Masolo[†]

November 29, 2021

Abstract

Surprises in survey responses on perceived business conditions produce strong comovement in unemployment, consumption, investment, and output, and a muted response of inflation and measured Total Factor Productivity (TFP). This suggests that news play an important role in explaining business cycle fluctuations, but also that attention should not be limited to TFP news. Employment news are the main driver of the overall index of reported business conditions.

Vector Autoregression impulse-responses can be matched by a New-Keynesian model in which individual risk, a positive supply of liquid funds, and complementarity between labor and capital inputs are modeled explicitly and the assumption of free entry of vacancies is done away with.

JEL Codes: C30, E31, E32.

Keywords: News, Unemployment, Business Cycles, Search Frictions, Individual Risk.

1 Introduction

News shocks have long been considered a potential driver of the business cycle (Pigou, 1927 and Beaudry and Portier, 2006, among others). For a time, now, the news literature has largely focused on Total Factor Productivity (TFP) news,[1] as discussed in Beaudry and Portier (2014).

However, influential work by Angeletos, Collard, and Dellas (2020), Barsky, Basu, and Lee (2014), and Shimer (2005), among others, challenges the idea that TFP news have the potential to explain business cycles, in particular due to the weak comovement between TFP and labor market variables.

In this paper, I take a step back and let the data suggest what sources of news, if any, have the potential to explain business cycle fluctuations, in the spirit of Cochrane (1994). Cochrane famously studied *consumption shocks*, which he defines as ”news consumers see but we do not see” (Cochrane, 1994, p. 296).

[†]Bank of England and Centre for Macroeconomics. Contact: rmmasolo@gmail.com

I am grateful to Tommaso Aquilante, Federico Di Pace, Michele Piffer, Ricardo Reis, and Vincent Sterk for their comments and suggestions.

[1]With the notable exception of a relatively small number of papers on tax news (e.g. Mertens and Ravn, 2012), on monetary policy news (Milani and Treadwell, 2012), and on investment news (Ben Zeev and Khan, 2015).

I subscribe to that general definition of news but I take advantage of survey information to peer into what an econometrician cannot usually see using only macro outcomes, such as consumption.

The University of Michigan Survey of Consumers (UMSC) question on news heard about changes in business conditions exhibits at at least two key features that makes it extremely valuable in this regard.[2] It asks consumers whether they have heard of any favorable of unfavorable changes in business conditions *during the last few months*. The timing qualification is critical. It makes the variable naturally predetermined to current realizations of macro aggregates. So a business-conditions news shock can be identified with a simple timing restriction, in an otherwise flexible Vector Autoregression (VAR) model.[3] Moreover, the backward looking nature of the question makes it easier to control for relevant information available to survey respondents. This news shock captures, by definition, information that consumers have seen but econometricians have not, as it is orthogonal to past realizations of the economic variables of interest.

A second key feature of questions on business-conditions news, is that the survey inquires about the *source* of news. Respondents can specify if they heard news relative to prices, demand conditions, employment, credit conditions and so on. This enables me to verify if the propagation of different sources of news displays significant differences. Also, it allows me to make some progress, relative to the exercise in Cochrane (1994), towards understanding which sources of news are important.

My first finding is that the economy responds differently to different types of reported news, in a way consistent with economic intuition. Reports of price news induce a swift increase in both consumers' inflation expectations as well as realized inflation. In response to this shock, unemployment and prices move in the same direction. The VAR responses look like the textbook model responses to a cost-push shock. Credit conditions shocks lead to a countercyclical response of interest rates. Demand news appear to generate negligible inflationary pressure, as in New-Keynesian models with a flat Phillips Curve.

Overall, consumers have a good grasp of economic conditions and of the drivers of future developments: these indices appear to be capturing specific sources of news and so they can be used to discriminate among them.

As for the aggregate business conditions index, it induces a muted response of inflation, a strong positive comovement among output, consumption and investment, and is orthogonal to measured TFP (Fernald, 2014). It is consistent with the so-called *Business-Cycle Anatomy*, as described by Angeletos, Collard, and Dellas (2020). The shock explains a large share of the forecast-error variance of real variables, such as unemployment, output, investment and consumption.

The third empirical contribution amounts to narrowing down the main driver of business conditions to labor-market news. From an economic standpoint, using the labor-market conditions index as opposed

[2]This series has been used, in the macroeconomic literature, by Barsky and Sims (2012) to study how it associates with consumer confidence. I retain the benefit of using the rich set of questions about business conditions, but I will use the data differently: they focus on how responses to news about business conditions correlate with innovations from their VAR specifications. I will confirm that business conditions news relate to consumer confidence, but also show how they appear to convey more information and in which sense they differ from a confidence shock.

[3]The use of survey data also increases the information content of the VAR, as advocated by Beaudry and Portier (2014) to guard against possible fundamentalness issues.

to the overall index, produces IRFs that are very close those to surprises to the aggregate index. From a statistical perspective, a principal component analysis of the the eight sub-indices provided with the Survey of Consumers shows that the first principal component, which explains more than two thirds of the variation in the aggregate index, is almost perfectly correlated (correlation coefficient of .998) with the employment news index. The second displays a very strong correlation with price news (coefficient of .867). Close to 80 percent of the variance of the overall business conditions index is explained by these two components which, in turn, directly relate to two separate sources of news. This complements the findings in Barsky and Sims (2012) that relate the prominence of these two sub indices to their popularity: they are the most often reported by consumers.

It is also important to note that the response of the economy to a business-conditions shock is clearly different relative to that to a consumer confidence shock. Positive surprises in confidence strongly associate with expectations of increases in TFP and, consequently, tend to be deflationary. Barsky, Basu, and Lee (2014), for instance, use the popular measure of *consumer sentiment* from the UMSC and find that ”to a large extent a news shock is a consumer confidence shock” (Barsky, Basu, and Lee, 2014, p. 239), where news refers to TFP news.[4] Business conditions shocks, however, leave both inflation and measured TFP unchanged. Consumer sentiment responds strongly to business conditions shocks. However, making the shock orthogonal to contemporaneous values of the consumer sentiment shock, does not significantly alter the peculiar characteristics of the responses of macro variables to a surprise in the business conditions index.

Overall, my empirical analysis supports the idea that news, in the general sense described above, play an important role in explaining business cycle fluctuations. News about the labor market are central to the aggregate business-conditions index.

The second part of the paper builds on these empirical findings by presenting a model in which labor-market shocks with a news component reproduce the comovement pattern identified in the data.

Angeletos, Collard, and Dellas (2020) show that state-of-the-art New Keynesian DSGEs can only do so by a constellation of shocks, each of which would struggle to generate that pattern in isolation. In a companion paper, Angeletos, Collard, and Dellas (2018) propose a model with autonomous variations in higher-order beliefs that could match some of these key features, which do not however directly include unemployment and other labor-market aggregates.

Mapping *labor-market news* from the VAR into structural shocks in a DSGE poses a challenge. This is inevitable when dealing with endogenous variables (Cochrane, 1994; Beaudry and Portier, 2014). Beaudry and Portier (2013) faced a similar problem as they start discussing a generic demand shock, but then they had to specify one to study its properties in a model.

[4]Barsky, Basu, and Lee (2014) summarizes the key properties of a TFP news shock. First, TFP news anticipate future changes in TFP. In other words, measured TFP responds to TFP news. Second, positive TFP news cause a fall in inflation, in line with New-Keynesian logic: positive TFP news decreases future expected marginal costs, which in turn affect the pricing decision of forward-looking firms. Third, oftentimes identified TFP news shocks cause investment and consumption impact responses to have opposite signs (in line with the classic result in Barro and King, 1984.

Fortunately the empirical analysis provides some clear indications, that limit the set of candidate shocks. The shock has to be directly related to labor-market conditions, it has not to comove with measured TFP, and it has to have a limited impact on inflation. Moreover it should induce positive comovement between consumption and investment, and a negative relationship between unemployment and vacancies, tracing out a *Beveridge* curve.

Shimer (2005) sets the standard for search-and-matching models of the labor market (Mortensen and Pissarides, 1994) aiming to explain business cycle fluctuations in unemployment. It considers two shocks, to TFP and to job-separation. The question then is, can a DSGE model reproduce the comovement observed in the data in response to a job-separation shock with a news component? And, is there any other shock that fits the bill?

As for the first question, a job-separation news shock in a standard representative-agent New-Keynesian model (e.g. Monacelli, Perotti, and Trigari, 2010, which in turn builds on Gertler and Trigari, 2009) is strongly deflationary and causes positive comovement between unemployment and vacancies (a shift in the *Beveridge* curve). The positive comovement between consumption and investment only emerges in the presence of sufficiently strong real frictions – in the form of investment-adjustment costs and consumption habits. Theodoridis and Zanetti (2016) also show, in a quantitative representative-agent model with search-frictions and labor-market news, that labor-market news do not produce realistic comovement among macro variables.

So I augment the model along two key dimensions. First, I separately model capitalists, who finance investment, from workers. The latter are subject to uninsurable unemployment risk. This type of Tractable Heterogeneous-Agent New-Keynesian model (THANK) is in line with recent work by Ravn and Sterk (2021) and Cui and Sterk (2021). Positive labor-market related news induce capitalists to invest more, due to the expected increase in capital productivity, and workers to consume more, due to a reduction in the risk of unemployment and the related precautionary-saving motive. The resulting increase in demand generates inflationary pressures.

Recent work by Coles and Moghaddasi Kelishomi (2018) and Broer, Druedahl, and Harmenberg (2021) shows that the counterfactual response of unemployment and vacancies to labor-market shocks depends on vacancies being infinitely elastic to economic conditions in standard business cycle models with search. I follow their insight and model a vacancy-opening cost, which implies that a vacancy is opened only if the expected return from it exceeds a randomly-drawn cost. This reduces the elasticity of vacancy creation with respect to the state of the economy and induces *negative* comovement between unemployment and vacancies in response to labor-market news shocks, in line with empirical findings.

The model I just outlined captures the key features of the data in response to a job-separation shock with a news component. All this in the absence of many of the frictions commonly employed in quantitative New Keynesian models, such as consumption habits and large investment-adjustment costs,[5] and with a reasonable slope of the Phillips curve. Wage stickiness improves the quantitative fit of the model but are

[5] In my baseline impulse-response matching exercise I estimate investment-adjustment costs. The estimate is tiny compared to macro literature and removing investment adjustment costs altogether does not affect my results.

not necessary to reproduce the comovement among macro variables. The same is true for the degree of complementarity between labor and capital inputs, which I model in line with the literature (Gechert et al., 2021), as well as for the degree of liquidity in the economy (Cui and Sterk, 2021).

I put the model to the test by conducting an impulse-response matching exercise, to have full control on the set of shocks the model can use to explain the observed relationship among macro variables - see discussion in Chahrour, Chugh, and Potter (2021). As mentioned above, I restrict my attention to shocks that directly relate to the labor market, that do not affect measured TFP, and that match up to the VAR evidence even in isolation. In practice this means considering shocks to job-destruction, and a shock to the cost of opening vacancies. The latter produces qualitatively similar responses when taken in isolation so I cannot rule it out a priori. On the other hand, the matching-efficiency shock, the other obvious candidate shock, produces a counterfactual positive correlation between vacancy and unemployment (Furlanetto and Groshenny, 2016), so I do not consider it in my baseline matching exercise.

Ultimately, the answers to the questions above are positive. A shock to job-separation augmented with a news component, matches the patterns we observe in response to a business-conditions shock. A shock to the cost of opening vacancies produces qualitatively similar effects.

Related Literature. This paper relates and contributes to a number of different strands in the macroeconomic literature. For one thing, it contributes the quest for *which* type of news shock is most important, addressed by Schmitt-Grohé and Uribe (2012) and Miyamoto and Nguyen (2020) in the context of a DSGE, while not abandoning the more traditional VAR-based news literature (Cochrane, 1994; Beaudry and Portier, 2006; Barsky and Sims, 2011; Barsky and Sims, 2012; Beaudry and Portier, 2014; Kurmann and Sims, 2021; Ramey, 2016). I allow for multiple sources of news, which is typically not the case in the VAR literature, while taking full advantage of the flexibility of a time-series specification. Interestingly, Schmitt-Grohé and Uribe (2012) find an important role for anticipated wage markup shocks. Shimer (2009) drives the comparison between wage markups and the labor wedge and makes a convincing argument for why search frictions may provide a more exhaustive explanation for the observed labor wedge. My empirical findings reflect the centrality of news related to the labor market, which, in a model with search friction, play a very important role.

The extent to which TFP news can drive the business cycle is still the subject of debate. Angeletos, Collard, and Dellas (2020) make a strong case for why it is not plausible that TFP news be an important driver of the business cycle. Chahrour, Chugh, and Potter (2021) argue that the max-share variance approach to identifying the *Main Business Cycle* shock, proposed by Angeletos, Collard, and Dellas (2020), produces long-term variations in measured TFP if lower frequencies are also considered and the covariance of output and hours is targeted, as opposed to an individual variable. They interpret their findings as supportive of TFP news being an important driver of the business cycle. Faccini and Melosi (2021) also make a case for noisy TFP news in an estimated DSGE.

I approach the problem from a different angle. My VAR analysis, does not set off to necessarily explain a large share of the variance in macro aggregates. Responses to business-conditions shocks turn out to be more in line with those in Angeletos, Collard, and Dellas (2020).[6] My modeling approach reflects this. Clearly my analysis cannot completely rule out that the sources of business conditions shocks are technological in nature, at least to an extent. If they are, though, they are not reflected in the first moment of measured TFP (Fernald, 2014), present or future. It could well be that they reflect technological advancements not easily captured in available measures of TFP, or they might reflect changes in the cross sectional distribution of TFP. It is possible, for instance, that the observed countercyclicality in TFP dispersion (Kehrig, 2015) may result in changes in reservation productivity (Mortensen and Pissarides, 1994) and ultimately in changes in the job-destruction rate. In this sense, a shock to job-destruction could be interpreted as related to the dispersion in idiosyncratic productivity, as modeled in Christiano, Motto, and Rostagno (2014).[7]

There have been various noteworthy attempts to present models consistent with the observed responses of the economy to news shocks. Jaimovich and Rebelo (2009) discuss the importance of capital utilization, investment adjustment costs, and and non-separable preferences in consumption and labor for a representative-agent DSGE model to produce accurate responses to news shocks. Den Haan and Kaltenbrunner (2009), as well as more recent related work by Chahrour, Chugh, and Potter (2021), focus on the propagation of news shocks in a model with search frictions, which implies a forward-looking behavior of labor-market participants. All these papers consider real models which, as such, have no implication for the behavior of inflation and nominal rates.

Beaudry and Portier (2013) are the first, to my knowledge, to move away from the representative-agent paradigm to explain news-driven, non-inflationary cycles. Recent years have witnessed an increasing number in papers that model individual unemployment risk and agent heterogeneity, Den Haan, Rendahl, and Riegler (2018), Mckay and Reis (2020), Challe (2020) being prominent examples, when restricting the attention to those including a search-and-matching mechanism (HANK-SAM). Broer, Druedahl, and Harmenberg (2021) specifically focus on the importance of unemployment risk when modeling business-cycle fluctuations.

The model I employ is closest to those in Ravn and Sterk (2021) and Broer, Druedahl, and Harmenberg (2021). In my model, capital accumulation plays a prominent role so that I can study the comovement between consumption and investment.[8] More importantly, I allow for positive levels of savings, along the lines of the *limited liquidity* approach proposed by Cui and Sterk (2021), so as to increase the quantitative accuracy of my model, while keeping it tractable enough so that standard impulse-response matching procedures apply.

Wage-determination is central in these models, all the more in the presence of wealth heterogene-

[6]Even if I extend the IRF out to 60 quarters, the response of TFP is - with a slight abuse of terminology given the Bayesian setting - not significantly different from zero, even at the 68 percent level.

[7]Interestingly, Christiano, Motto, and Rostagno (2014) show that news to productivity dispersion (or risk) play a more important role in explaining business cycle fluctuations than news to the level of TFP.

[8]See Monacelli, Perotti, and Trigari (2010) for a representative-agent model with search-and-matching frictions and capital. Ravn and Sterk (2021) present an extension with capital. Here I allow for more quantitative features such as capital utilization, investment-adjustment costs, input complementarity and I allow for different risk aversion on the part of capitalists and workers.

ity. Early discussions of the effects of wage determination and wealth heterogeneity go back to Gomes, Greenwood, and Rebelo (1997), and, more recently, Krusell, Mukoyama, and Şahin (2010). I want to put my model to the test with the most-standard wage-setting mechanism: Nash bargaining in its simplest form, i.e. independent of individual wealth. I thus assume that workers need the intermediation of an employment agency to search for a job, against the payment of a fee proportional to the aggregate wage (or home production when unemployed). This results in bargaining conditions comparable to those in representative-agent models like Monacelli, Perotti, and Trigari (2010).

I also keep the preferences deliberately simple, of the constant elasticity of substitution form with no habits. Equally standard is the pricing friction á la Rotemberg (1982). The way I model the cost of opening vacancies, as mentioned above, follows Coles and Moghaddasi Kelishomi (2018) and Broer, Druedahl, and Harmenberg (2021), which makes for the observed procyclical response of vacancies.

2 Empirical Analysis

I estimate a series of Bayesian Structural VARs of the form:

$$Y_t = A(L)\, Y_{t-1} + u_t \tag{1}$$

where Y_t is a vector of variables, $A(L)$ a matrix polynomial in the lag operator L, $\mathbb{E}[u_t] = 0$, $\mathbb{E}\left[u_t u_t'\right] = \Sigma$, $\mathbb{E}\left[u_t u_{t-j}'\right] = 0\ \forall j \geq 1$.

If the business conditions measure of choice is the first entry in Y_t, I need to identify the first column of C, $u_t = C\varepsilon_t$, such that $\mathbb{E}\left[\varepsilon_t \varepsilon_t'\right] = I$, $CC' = \Sigma$. I do so by exploiting the structure of the question I consider, which asks: "During the last few months, have you heard of any favorable or unfavorable changes in business conditions?". As the question inquires about the "last few months" it naturally makes the variable predetermined. The identification assumption is that the business-conditions series does not respond contemporaneously to any of the time-t macro variables. This translates into the first column of C being full, and the first row of C being all zeros except the first entry, a standard Cholesky identification scheme.

Clearly a timing identification assumption is more restrictive the longer the period under consideration. I will thus establish all the relevant results in the context of monthly VARs and use quarterly specifications primarily as a comparison to the business-cycle literature (Angeletos, Collard, and Dellas, 2020) and as a benchmark for the theoretical model. The quarterly specifications will also allow me to consider a larger set of macro variables, not available at a monthly frequency.

The survey question of interest becomes available at a monthly frequency from 1978 onwards. Monthly VARs are thus estimated over the 1978-2019 sample. Quarterly specifications can extend back to 1965. I also consider mixed-frequency VARs in which a core of monthly variables (the index of business conditions, and measures of unemployment, industrial production, aggregate prices, and interest rates) are

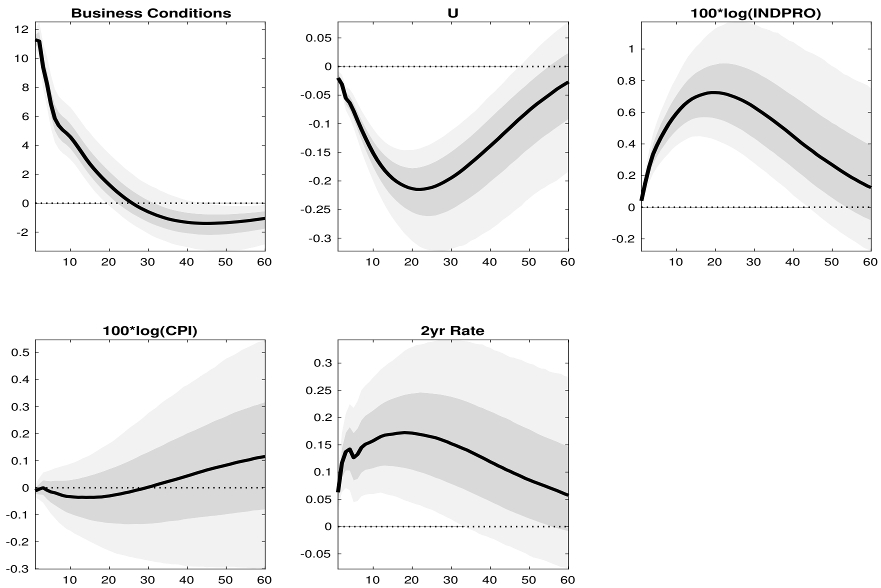

Figure 1: Monthly, 12-lag, Bayesian VAR including the Business Conditions index, unemployment, industrial production, CPI and the 2-year interest rate. Median response in black with 68 and 95 percent credible sets in gray.

complemented with some quarterly series. Mixed-frequency specifications have the advantage of not imposing the identifying timing restriction for the duration of an entire quarter, while enabling to study the responses of key variables such as investment and consumption.

Monthly VARs are estimated in log-levels – except rates which are in levels – and include 12 lags of the observables. For the quarterly specification I primarily follow the variable definitions in Angeletos, Collard, and Dellas, 2020 and include 4 lags. Bayesian estimation, as implemented by Ferroni and Canova (2020), employs Minnesota priors with hyperparameters optimized as in Giannone, Lenza, and Primiceri (2015).

2.1 Monthly Specifications

It includes the UMSC index of reported business conditions alongside unemployment, industrial production (in log levels), the CPI price index (in log levels) and the the 2-year interest rate.[9]

Figure 1 shows that, in response to a one standard-deviation shock to business conditions, the economy displays a large, persistent, and significant reduction in unemployment, peaking at a horizon of about 2 years. Industrial production increases and so do interest rates. This pattern would paint the picture of a standard demand shock, except for the fact that prices hardly move. So a business-conditions shock

[9] I use the 2-year rate in my baseline specification in line with Gertler and Karadi (2015) and Swanson and Williams (2014), who suggest using one or two-year rates as a better indicator of the monetary policy stance. My results do not depend on this choice as will become clear in specifications with Fed Funds Rates.

is strongly procyclical, explains about half of the variance in unemployment and industrial production (Figure 2), but is largely orthogonal to inflation.

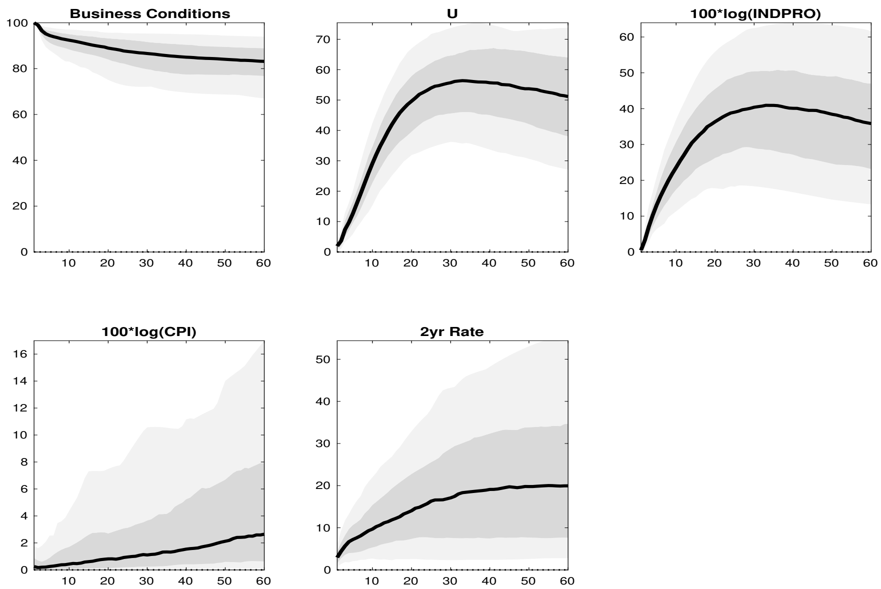

Figure 2: Baseline specification of the monthly VAR, share of forecast-error variance explained by shock to Business Conditions.

The response pattern is robust to estimating the VAR on shorter samples – either the post-1984 sample or the Great Moderation sample, presented in the Appendix.[10]

The formulation of the survey question suggests a clear timing sequence. When the survey is administered, figures for the current-period macro variables are yet to be released (in fact they are yet to fully materialize), so clearly they do not form part of the responders' information set. To restrict a business conditions shock to be orthogonal to current-period variables is thus not consistent with the timing of the survey and the nature of the question. However, in keeping with the macroeconomic literature, as a robustness check, I also present, in the Appendix, a specification in which the Cholesky ordering is reversed, and the business conditions shock is made orthogonal to current realizations of macro variables. The patterns emerged from the baseline specification survive under this alternative identification scheme. Which, as Barsky and Sims (2012) point out, is a strong indication that the business conditions index genuinely Granger causes the other macro variables.

More relevant is another robustness check that relies on intra-monthly information. I will illustrate below that the business conditions index is strongly related to labor-market conditions. As such it is sensible to expect that it might respond to the intra-monthly (weekly) releases of unemployment claims. Making the business-conditions shock orthogonal to unemployment claims, however, does not significantly affect

[10]The response pattern does not depend on the choice of priors either, as shown in the the Appendix, where I overlay IRFs from an OLS counterpart to the Minnesota prior baseline specification in Figure B.8.

the responses of the macro variables. As expected, the responses of both initial and continued claims are strongly countercyclical (Figure B.11, in the Appendix).

Sentiment and Inflation Expectations. A key and popular measure of consumer expectations is the UMSC consumer sentiment series. I thus consider a specification in which consumer sentiment is added, alongside the UMSC measure of inflation expectations and a measure of stock prices, in line with the recommendation of Beaudry and Portier (2014) to combine stock prices and survey information to increase the informational content of a VAR specification. Figure 3 shows that both the measure of sentiment and stock prices display a strong contemporaneous response to the business conditions shock, while inflation expectations hardly move. The responses of the main macro variables are largely unaffected by the introduction of the these extra controls.

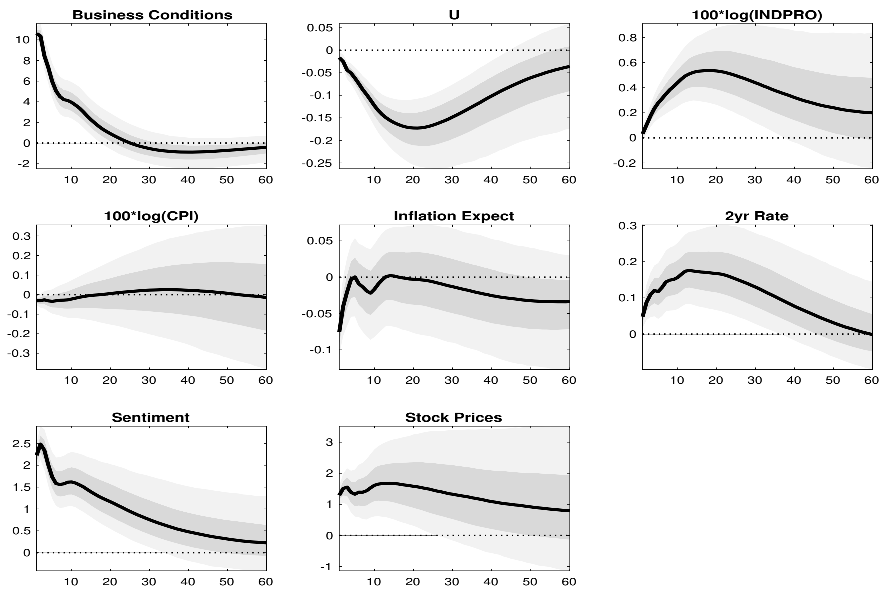

Figure 3: Specification of the monthly VAR that includes Consumer Sentiment inflation expectations and stock prices.

2.1.1 Sub Indices

The business conditions index aggregates eight sub indices. When respondents report hearing positive or negative news they are asked "What did you hear?". Responses are categorized as referring to: employment/unemployment, prices, demand conditions, government, credit conditions, stock markets, international trade, and energy crisis.

Studying each of them in turn, helps shed light on two key aspects:

i. do respondents understand the broad macroeconomic implications of the information they heard about?

ii. Is there a main driver of the results obtained with the headline index?

Answering the first question is important in that it gives me confidence that respondents can actually anticipate what the news they hear is bound to translate into. The second question is critical to build a structural model that explains the observed behavior.

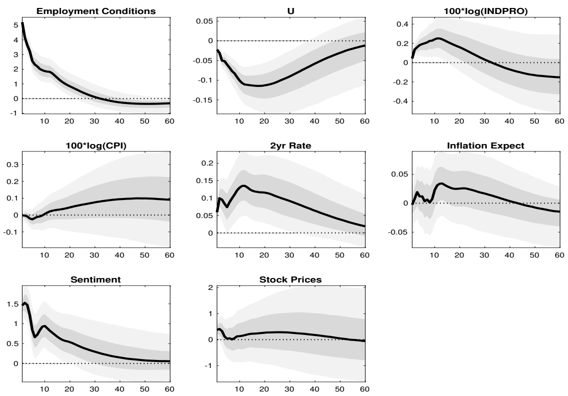

(a) Employment conditions

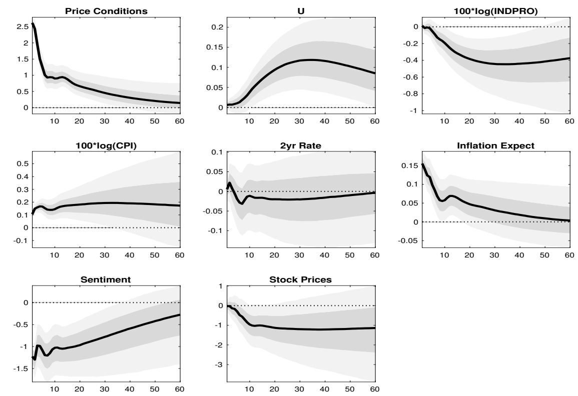

(b) Price conditions

Figure 4: Impulse responses to a shock to the employment conditions index (left) and to the price index (right).

The answer to the first question is affirmative. Figure 4, shows the IRFs to a shock to employment-conditions index, and to price conditions.[11]

The responses to an employment-related surprise change in business conditions are close to that to the overall index, with the possible exception of the response of stock prices, which is more muted. Contrasting them to the response to price news immediately shows that agents have a clear understanding of the driver of the news and the effects on the economy. When consumers hear of price news they correctly revise up their inflation expectations: the initial jump in inflation expectations is about exactly the same as the jump in the price index. News of high prices also associate with an increase in unemployment and a fall in industrial production, stock prices, and consumer sentiment. News of surprisingly high prices produce responses that look a lot like those to a standard cost-push shock.

I defer the IRFs to the other indices to the Appendix (Figure B.17) in the interest of space. Not all produce very large effects. However, taken together, the responses confirm the idea that consumers, at least on average, have a very good grasp of the news they report and of the ensuing macroeconomic developments. For instance, a reported easing in credit conditions results in an increase in industrial production and a fall in unemployment. Contrary to a shock to labor conditions, however, these associate to a countercyclical response of the interest rate, reflecting the reduced cost of borrowing.

[11] I focus on these two indices as they turn out to be the two most important, but in the Appendix I report IRFs for each of the 8 indices in Figure B.17.

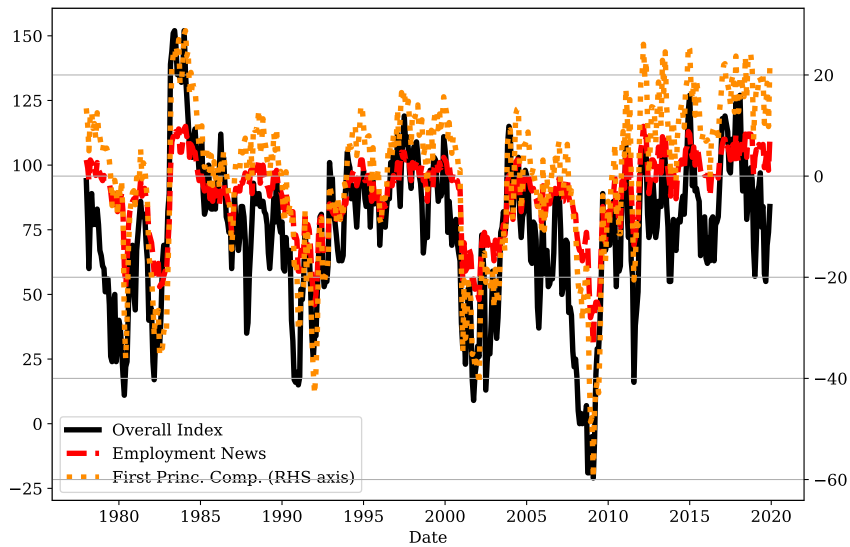

Figure 5: Overall news index (black solid), employment news index (red dashed), first principal component (orange dotted).

Principal Component Analysis. The responses of the macroeconomic variables to employment news appear fairly similar to those to the overall business-conditions index (Figure B.17). Indeed, the overall news index displays a correlation of about .8 with the employment news index – Table C.1.

More formally, a simple principal-component analysis of the set of news indices reveals that the first principal component is almost perfectly correlated with the employment news index (correlation coefficient of .998), explains two thirds of the overall variance (66.5 percent) and strongly correlates (.816) with the overall news index. Figure 5, reports the time series plots of the overall index, the employment index and the first principal component. It visually confirms the strong comovement between the series.

The second principal component is primarily correlated with price news (correlation coefficient of .867). The first two principle components explain close to 80 percent (77.9) of the overall variation in the business conditions index. None of the other components explains more than 10 percent of the overall variation.

The headline business-conditions index can thus be thought of as reflecting a main driver relating to employment conditions and a secondary, much less important, driver that correlates with price news. The association of these two sub-indices with different principal components, which are orthogonal by design, explains why they generate distinctly different effects in the economy.

2.1.2 Informational content: integrating information from the Survey of Professional Forecasters

Professionals' forecasts are widely considered to produce the most accurate and timely forecasts (Ang, Bekaert, and Wei, 2007). I thus want to control for the informational content of forecasts from the Survey of Professional forecasters. In particular, I build a series that captures the revision over a quarter, of the one-year ahead forecast for unemployment, as a benchmark for news about employment.

Controlling for these news requires a mixed-frequency VAR specification, as SPF data is available at quarterly frequency.[12]

Reports of positive developments in business conditions associate with downward revisions in professionals' forecasts for unemployment.[13] This reinforces the idea that business conditions shocks genuinely capture important developments on the labor market. The responses of the other variables are in line with those I estimate in the other monthly specifications.[14] So, ultimately, reports in business conditions, while correlated with professional forecasters' information, do not overlap completely. Indeed, the component that is orthogonal to information reflected in publicly available professional forecasts, retains a lot of explanatory power.

2.1.3 Consumption and Investment

I estimate a mixed-frequency specification which includes measures of consumption and investment.[15] Investment and consumption comove very strongly with each other and with industrial production in response to a shock to business conditions (Figure B.16). Expected is also the fact that the investment response is larger in magnitude.

Overall, these impulse responses paint a picture that is difficult to reconcile with "standard" news shocks which (Barsky, Basu, and Lee, 2014) typically imply deflationary pressures and negative comovement between investment and consumption.

2.2 Quarterly Specification and the Main Business Cycle shock

A quarterly specification allows me to compare my results to a wider literature, to consider a larger set of variables, and to produce a benchmark for the theoretical model.

[12] For mixed frequency VARs, I do not optimize over the prior hyperparameters. Rather I set them based on the optimized values for similar specifications which do not include quarterly variables.

[13] Figure B.10, in the Appendix, reports impulse-responses for this specification.

[14] In the Appendix I show that these findings are robust even when I focus exclusively on employment news (Figure ??). Making the monthly series for business conditions orthogonal to a quarterly series from the SPF survey is clearly unrealistic given the monthly nature of the former. However, for the sake of showing the robustness of the covariance between the business conditions index and the macro variables, I also report a specification in which business conditions are made orthogonal to contemporaneous realizations of the SPF forecast revisions (Figure B.13). Though smaller in magnitude, and less precisely estimated, responses of macro variables line up well with my baseline specification.

[15] The quarterly specification in the next section will allow me to study the comovement among a much larger set of variables. Given the centrality of the consumption-investment relationship, however, I want to establish that it hold even in a monthly specification, augmented with these two quarterly variables.

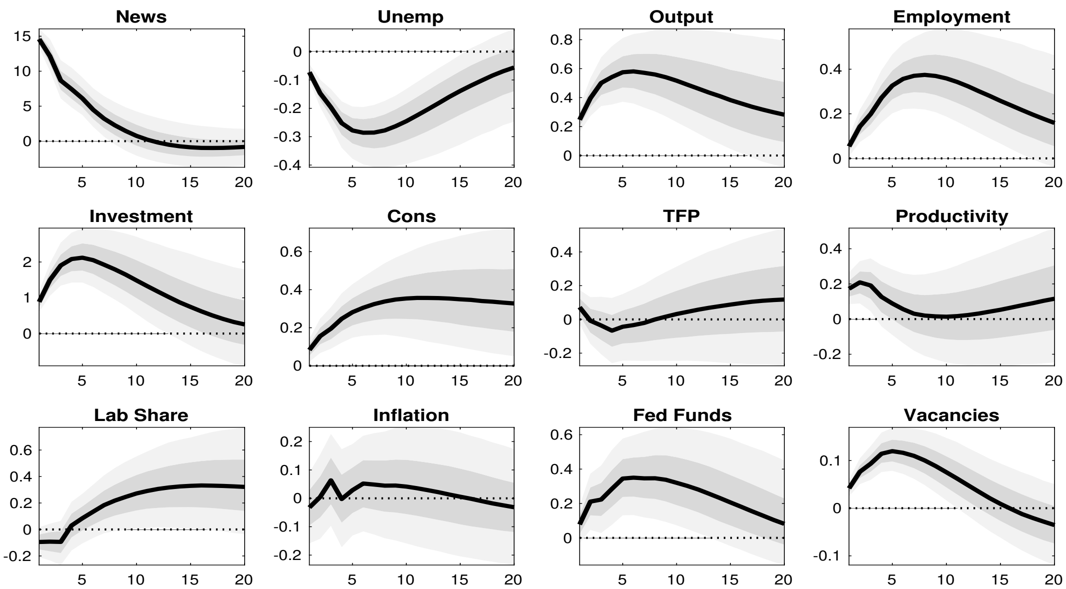

Figure 6: Quarterly VAR specification for IRF matching.

I will follow Angeletos, Collard, and Dellas (2020) in defining the variables to be included in my quarterly specification, so that I will be able to study the effects of the business-conditions shock on variables such as measured TFP (Fernald, 2014) and ultimately to compare its effects to those of the so-called *Main Business Cycle* shock.[16] Angeletos, Collard, and Dellas (2020) include unemployment, output, hours, investment, consumption, TFP, labor productivity, the labor share, inflation and the Fed Funds rates. I replace employment for hours, which is immaterial for the results but more in line with the specification of my model. Moreover, I add a measure of vacancies based on Barnichon (2010).[17]

Figure 6, reports the responses to a surprise in the business-conditions index. Consistent with the findings from the monthly specifications, output, hours, investment and consumption all increases significantly and persistently, while unemployment falls. The Fed Funds Rate increases in response to the shock, while measured TFP and inflation do not.

These responses bear a striking similarity to those identified by Angeletos, Collard, and Dellas (2020)[18] not only in terms of the signs of the responses, but also with regards to magnitudes and even the timing of the peak effect.

The responses to a surprise in the index of reported business conditions, meet all the characteristics of the *business cycle anatomy* described by Angeletos, Collard, and Dellas (2020). Namely, they display strong comovement of real macro variables and of interest rates, while being all but orthogonal to inflation and TFP. The key difference is that while Angeletos, Collard, and Dellas (2020) emphasize the fact that their

[16]In the Appendix I also report the quarterly counterpart to my baseline monthly specification for comparison.

[17]This limits the sample to 2016, but is immaterial for the estimates given the long sample starting in 1965.

[18]In the Appendix I present the exact comparison.

identification scheme is entirely agnostic by design, the use of survey information gives me an indication as to the economic nature of the shock.[19]

In the Appendix, I show that these patterns hold if I limit myself to the employment conditions sub index, and if I only consider business-conditions news reported in the survey administered in the first month of the quarter (available only starting in 1978), so that the identifying timing assumption is no more restrictive than in a monthly specification.

In sum, my empirical analysis shows that surprises in the business conditions index are mainly driven by employment news. They produce strong comovement between unemployment, vacancies, output, hours, investment, consumption and interest rates but do not affect inflation, nor TFP. The impulse responses bear striking similarity to those that describe the *business cycle anatomy*, in the words of Angeletos, Collard, and Dellas (2020), while providing a more precise indication with regards to the source of the variation.

3 Model

The empirical analysis restricts the set of structural shocks that can explain the responses of the macro variables at hand.

Responses to a job-separation shock from popular representative-agent New-Keynesian models with search frictions – Monacelli, Perotti, and Trigari (2010), which in turn relates to Gertler and Trigari (2009) –, are at odds with those identified in the VAR presented above. First, they are deflationary. Second, they can result in a negative comovement between investment and consumption (when a news component is considered), unless strong enough investment-adjustment costs are introduced. Third, a job-separation shock causes unemployment and vacancies to move in the same direction, or, equivalently, a shift in the *Beveridge* curve. Matching up to the empirical IRFs would require a combination of shocks and a host of real frictions, as discussed in Angeletos, Collard, and Dellas (2018).

The dynamics of the model responses to a job-separation shock described above are due primarily to two features of the model: the representative-agent assumption and the modeling of the cost of vacancies. I move away from a representative-agent setup, by modeling workers and capitalists separately. Capitalists own firms, and thus capital. Workers are subject to uninsurable unemployment risk, which I model as in Tractable HANK (THANK) models (Ravn and Sterk, 2021 and Cui and Sterk, 2021). In this environment, an improvement in labor market conditions reduces the risk of unemployment and thus precautionary

[19] Figure B.20 illustrates that the business conditions shock explains more than 50 percent of the variability of unemployment at business cycle frequency and somewhere between 25 and 40 percent of the forecast-error variance of output, hours, investment and consumption. With the noteworthy exception of unemployment, these shares are somewhat smaller than those found by Angeletos, Collard, and Dellas (2020). This is to be expected considering that my identification scheme is not deliberately targeting the highest share of the forecast-error variance. Consistent with Angeletos, Collard, and Dellas (2020), consumption is the variable for which this shock has the lowest explanatory power, though the business conditions shock still explains about a quarter of its business cycle variation.

savings, boosting demand and producing an upward pressure on prices that balances out the otherwise deflationary effects of a job-separation shock.

Free entry into vacancies, or vacancies being infinitely elastic to business conditions, is key to the counterfactual comovement of unemployment and vacancies to job-destructions and matching efficiency shocks (Shimer, 2005). Coles and Moghaddasi Kelishomi (2018) propose a novel way to model vacancy creation, also adopted by Broer, Druedahl, and Harmenberg (2021). It amounts to assuming that firms open a vacancy if the expected benefit is larger than a vacancy-opening cost. As a result, in the wake of a reduction in job-separations, vacancies do not necessarily fall.

A model with these features, similar to that in Broer, Druedahl, and Harmenberg (2021) but augmented with capital and investment, qualitatively replicates the observed comovement in the data. A better quantitative match obtains, however, when two more extensions are considered. First, I relax the commonly held assumption in THANK models that the supply of liquid assets is zero, so that every agent consumes her income. I allow for so-called *moderate liquidity*, as in Cui and Sterk (2021). That is to say, the supply of liquid assets is strictly positive, yet small enough to induce newly unemployed agents to consume all of their savings in the first period of unemployment. This assumption is realistic given the average level liquid wealth held by workers and it dramatically improves tractability. In this version of the model, an improvement of labor-market conditions materially reduces unemployment risk for workers who are building up their saving buffer, resulting in a marked increase in the level of their consumption.

In this economy, investment is financed by retained earnings, thus, indirectly, by a reduction in the consumption of the capitalists, who own the firms. The anticipation of a period of favorable labor-market conditions will increase capital productivity and boost investment. The higher the degree of complementarity between investment and labor, the more pronounced the investment response. I find that by modeling capital and labor as complements, with a degree of complementarity in line with the literature (Gechert et al., 2021), the response of investment in the model aligns well with that in the data.

3.1 Setup

There is a continuum of measure one of households. A fraction $\mathfrak{w}$ are workers, the rest are capitalists. Workers, can be employed or unemployed. The wealth heterogeneity in the model implies that I need to model H separate cohorts of employed workers. As for unemployed workers, newly unemployed start the period with liquid wealth but do not benefit from home production. As such, their problem is different relative to that of agents unemployed for two periods or longer.

3.1.1 Capitalists.

Households $i \in [\mathfrak{w}, 1]$ are capitalists. They do not supply labor[20] but own a differentiated portfolio of all the firms in the economy. Just like all the other agents in the economy, they face a borrowing constraint,

[20] I maintain that despite not being employed their home production is zero. One can rationalize this by assuming that their time is devoted to the management of the firms. This assumption is not central to the main implications of the model.

which is set to zero.

Provided all capitalists start off with the same level of wealth, they will all face the same problem:

$$\max_{C_{i,t+j},H_{i,t+j+1},b_{i,t+j+1}} \quad \mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \frac{C_{i,t+j}^{1-\sigma_{\mathrm{k}}}}{1-\sigma_{\mathrm{k}}}, \tag{2}$$

$$\text{s.t.} \quad C_{i,t} + p_t^H H_{i,t+1} + b_{i,t+1} = \frac{R_{t-1}}{\Pi_t} b_{i,t} + \left(p_t^H + d_t\right) H_{i,t} + T_{i,t}, \tag{3}$$

$$b_{i,t+1} \geq 0, \tag{4}$$

where $C_{i,t}$ is consumption for agent i in period t and σ_{k} governs the elasticity of intertemporal substitution for capitalists. $H_{i,t}$ are stock holdings,[21] $b_{i,t}$ bond holding, $T_{i,t}$ government transfers, p_t^H stock prices (in real terms, i.e. in units of consumption); d_t are dividends, R_t the nominal short-term rate set the by the Central Bank, and Π_t consumption price inflation.

I maintain that each capitalist starts off with equal holdings of stocks $H_{i,t} = \frac{1}{1-\varpi}$ and bonds $b_{i,t} = 0$. They have no incentives to trade stocks so their holdings will remain constant.

As for bonds, the maintained conjecture – to be verified later – is that the real rate of interest in and around the steady state will be below $1/\beta$, driven down by the desire of employed workers to save against unemployment risk. Under this conjecture capitalists would like to borrow, so their borrowing constraint will be binding.

The capitalists' problem can be characterized by the following two equations:

$$p_t^H = \beta \mathbb{E}_t \frac{C_{\mathrm{k},t+1}^{-\sigma_{\mathrm{k}}}}{C_{\mathrm{k},t}^{-\sigma_{\mathrm{k}}}} \left(p_{t+1}^H + d_{t+1}\right), \tag{5}$$

$$C_{\mathrm{k},t} = \frac{d_t}{1-\varpi} + T_{\mathrm{k},t}, \tag{6}$$

where $C_{\mathrm{k},t}$ is the consumption of the representative capitalist. Equation (5) is a simple pricing equation for stocks, while equation (6) pins down the consumption of capitalists, which equals their income in each period, as they are credit constrained.

3.1.2 Workers

Workers can be either employed or unemployed. Workers also need to make a decision regarding their participation in the labor market. In this models workers have differing levels of wealth. The level of wealth would, in principle, affect their bargaining power. I assume that workers can only participate in the labor market via an employment agency owned by capitalists, so that the standard wage-setting problem of representative-agent models obtains.

[21] One could think of this as shares in a mutual funds that owns all the firms in the economy.

Workers' participation decision. At the start of the period workers need to decide whether they want to participate in the labor market or live off their home production (autarky). To participate in the labor market they have to strike a deal with an employment agency that will cost them a fraction τ^A of wage when employed and of their home production when unemployed.

A sufficient conditions for all workers to participate in the labor market is:

$$\mathfrak{p}_t \mathbf{V}_{\mathbf{e},t} \left(\left(1 - \tau^a\right) w_t \right) + \left(1 - \mathfrak{p}_t\right) \mathbf{V}_{\mathbf{u},t} \left(\left(1 - \tau^a\right) \vartheta \right) \geq \mathbf{V}_{\mathbf{a},t} \left(\vartheta\right). \tag{7}$$

The left-hand side represents the expected utility of an unemployed worker at the start of period t who knows that, if he enters the contract, will find employment with probability $\mathfrak{p}_t$ and have income $\left(1 - \tau^a\right) w_t$, while with probability $1 - \mathfrak{p}_t$ will remain unemployed with income $\left(1 - \tau^a\right) \vartheta$. If this value exceeds that of autarky $\mathbf{V}_{\mathbf{a},t} \left(\vartheta\right) = \frac{1}{1-\beta} \frac{\vartheta^{1-\sigma_{\mathrm{w}}}}{1-\sigma_{\mathrm{w}}}$, the unemployed worker will opt in or remain in the contract. Notice that, if that condition is satisfied for unemployed workers, it will always hold for employed workers as well, since $\mathbf{V}_{\mathbf{e},t} \left(\left(1 - \tau^a\right) w_t \right) \geq \mathfrak{p}_t \mathbf{V}_{\mathbf{e},t} \left(\left(1 - \tau^a\right) w_t \right) + \left(1 - \mathfrak{p}_t\right) \mathbf{V}_{\mathbf{u},t} \left(\left(1 - \tau^a\right) \vartheta \right)$.

I will focus on the case in which $\tau^a \longrightarrow 0$ for two reasons. Firstly, this rules out transfers of resources between agents. Secondly, it ensures that the inequality in equation (7) will always be satisfied, so long as these three sufficient conditions hold: $\mathfrak{p}_t > 0$, $w_t > \vartheta$ and $w_t > \frac{R_{t-1}}{\Pi_t} b_{H,t}$. The first condition trivially holds, the second holds except in the limit case in which all the bargaining power rests with firms. The third condition refers to newly unemployed and states that consuming all their savings still does not afford them the consumption level they could afford by working.[22]

These sufficient conditions are stronger than needed as they disregard the benefit of future employment spells but are very easy to verify.

Employed Workers' consumption-saving decision. The consumption-saving decision is made after matches have occurred. Employed workers supply labor $N_{i,t} = 1$, their indivisible supply labor normalized to 1. They face uninsurable unemployment risk. So their optimization reflects the possibility that, with probability $\left(1 - \rho_t\right)$, their current employment contract will terminate at the end of period t. If so, with probability $\mathfrak{p}_{t+1}$ they will be matched to a new employer at the start of period $t + 1$, else they will become unemployed. I define $\mu_{t+1} \equiv \rho_t + \left(1 - \rho_t\right) \mathfrak{p}_{t+1}$, the probability of being employed next period, conditional on being employed in the current period.

The *moderate liquidity* assumption (Cui and Sterk, 2021) implies that unemployed workers consumer all of their savings in the first period out of work. As a consequence, the consumption of employed workers depends only on the number of periods of uninterrupted employment. I refer to new hires out of unemployment as cohort $h = 0$. If they remain in employment in the following period, they will be part of cohort $h = 1$, and so on. The number of periods of continued employment is the only source of heterogeneity.

[22] $b_{H,t}$ are the bond-holdings for cohort H, i.e. the highest level of wealth of all cohorts as will be discussed later. Notice that this condition is much stronger than needed. As autarky is a permanent condition in this model, even if savings would allow a newly unemployed to consume more than an employed worker, the lower continuation value would likely tilt the balance in favor of remaining in the contract with the agency.

After a sufficiently long period of continued employment, workers have built up their saving buffer and the levels of consumption and saving are all but the same across cohorts. As a result, I can truncate the number of cohorts I keep track of to H, with the understanding that cohort H includes all those that have been continually employed for H periods or longer.[23]

Defining $\mathbf{V_e}\left(b_{i,t}\right)$ the value function of an employed worker entering period t with wealth $b_{i,t}$, and $\mathbf{V}_{\mathbf{u},0}\left(b_{i,t}\right)$ the corresponding value function of a *newly* unemployed worker, I can write down the employed worker's problem as:

$$\begin{aligned}
\mathbf{V_e}\left(b_{i,t}\right) &= \max_{C_{i,t},b_{i,t+1}} \frac{C_{i,t}^{1-\sigma_{\mathrm{w}}}}{1-\sigma_{\mathrm{w}}} + \beta\mathbb{E}_t\left\{\mu_{t+1}\mathbf{V_e}\left(b_{i,t+1}\right) + \left(1-\mu_{t+1}\right)\mathbf{V}_{\mathbf{u},0}\left(b_{i,t+1}\right)\right\} \\
\text{s.t.}\quad & C_{i,t} + b_{i,t+1} = \frac{R_{t-1}}{\Pi_t}b_{i,t} + w_t - T_{i,t} \text{ if employed} \qquad (8)\\
& C_{i,t} + b_{i,t+1} = \frac{R_{t-1}}{\Pi_t}b_{i,t} - T_{i,t} \text{ if } \textit{newly} \text{ unemployed} \qquad (9)\\
& b_{i,t+1} \geq 0 \qquad (10)
\end{aligned}$$

where w_t is the real wage and $T_{i,t}$ transfers to employed workers to be defined later - I will maintain that all workers will be subject to the same transfers. Given the envelope condition $\frac{\partial \mathbf{V_e}\left(b_{i,t}\right)}{\partial b_{i,t}} = \frac{R_{t-1}}{\Pi_t}C_{i,t}^{-\sigma_{\mathrm{w}}}$, and the fact that all workers in the same cohort make the same decision (so that I can index by the cohort as opposed to the individual agent), the Euler condition reads:

$$C_{\mathbf{e},h,t}^{-\sigma_{\mathrm{w}}} = \beta\mathbb{E}_t\frac{R_t}{\Pi_{t+1}}\left\{\mu_{t+1}C_{\mathbf{e},h+1,t+1}^{-\sigma_{\mathrm{w}}} + \left(1-\mu_{t+1}\right)C_{\mathbf{u},0,h+1,t+1}^{-\sigma_{\mathrm{w}}}\right\} \quad \forall\, h = 0,\dots,H-1 \qquad (11)$$

$$C_{\mathbf{e},H,t}^{-\sigma_{\mathrm{w}}} = \beta\mathbb{E}_t\frac{R_t}{\Pi_{t+1}}\left\{\mu_{t+1}C_{\mathbf{e},H,t+1}^{-\sigma_{\mathrm{w}}} + \left(1-\mu_{t+1}\right)C_{\mathbf{u},0,H,t+1}^{-\sigma_{\mathrm{w}}}\right\} \quad h = H \qquad (12)$$

reflecting the fact that in period $t+1$ a cohort h agent in period t will be a cohort $h+1$ agent, either employed or newly unemployed.[24] By the argument presented above, an H-cohort agent will remain in that cohort until her occupational status does not change. The limited liquidity assumption implies that $b_{-1,t} = 0, \forall\, t$. The Euler equations, alongside the budget constraints, pin down the consumption level and bond holdings of each employed worker.

Unemployed workers' consumption-saving decision. When deciding consumption for period t an unemployed agent is no longer able to search for work in the current period and knows that $N_{i,t} = 0$. Next period he will find occupation with probability $\mathfrak{p}_{t+1}$. In anticipation of higher future income, he would like to borrow. His borrowing constraint will thus be binding and his consumption will be determined by his financial wealth, home production and any transfer.

There are two key differences between *newly* unemployed and all other unemployed workers. The

[23]The determination of H is ultimately a numerical problem. In my baseline specification keeping track of 40 cohorts is enough to obtain a well defined problem, in line with the findings of Cui and Sterk (2021).

[24]The borrowing constraint will not be binding for employed agents who want to save against the risk of losing their job.

former start the period with liquid wealth $b_{h,t}$ but cannot rely on home production for the first period of unemployment. This can be rationalized by thinking that they need to set themselves up to be productive at home and/or by considering that they incur one-off expenses upon losing their job (e.g. relocation). I maintain that these effect combine to the level of home production in the first period of unemployment. The latter group, by the limited-liquidity assumption,[25] have no bond-holdings left but have set themselves up to operate the home production technology, which yields ϑ units of consumption good in each period.

A *newly* unemployed worker, that would have been part of cohort h if he was still employed (indexing by cohort is necessary to pin down the level of wealth), faces the following problem:

$$\mathbf{V}_{\mathbf{u},0,h}\left(b_{i,t}\right) = \max_{C_{i,t},b_{i,t+1}} \frac{C_{i,t}^{1-\sigma}}{1-\sigma} + \beta\mathbb{E}_t\left\{\mathfrak{p}_{t+1}\mathbf{V}_{\mathbf{e},0}\left(b_{i,t+1}\right) + \left(1-\mathfrak{p}_{t+1}\right)\mathbf{V}_{\mathbf{u}}\left(b_{i,t+1}\right)\right\}$$

$$\text{s.t.} \quad C_{i,t} + b_{i,t+1} = \frac{R_{t-1}}{\Pi_t}b_{i,t} + w_t - T_{i,t} \text{ if employed} \tag{13}$$

$$C_{i,t} + b_{i,t+1} = \frac{R_{t-1}}{\Pi_t}b_{i,t} - T_{i,t} \text{ if newly unemployed} \tag{14}$$

$$C_{i,t} + b_{i,t+1} = \frac{R_{t-1}}{\Pi_t}b_{i,t} + \vartheta - T_{i,t} \text{ if unemployed for more than one period} \tag{15}$$

$$b_{i,t+1} \geq \overline{0} \tag{16}$$

The *limited liquidity* assumption hinges on the borrowing constraint being binding for all newly unemployed. Since the level of bond-holdings in non-decreasing in h, and the transfers are the same to all workers, newly unemployed of cohort H enjoy the highest level of consumption and, consequently, the lowest level of marginal utility $C_{\mathbf{u},0,H,t}^{-\sigma_{\mathrm{w}}}$. It is thus sufficient to verify that they are not on the Euler equation, that is:

$$C_{\mathbf{u},0,H,t}^{-\sigma_{\mathrm{w}}} > \beta\mathbb{E}_t\frac{R_t}{\Pi_{t+1}}\left\{\mathfrak{p}_{t+1}C_{\mathbf{e},0,t+1}^{-\sigma_{\mathrm{w}}} + \left(1-\mathfrak{p}_{t+1}\right)C_{\mathbf{u},t+1}^{-\sigma_{\mathrm{w}}}\right\} \tag{17}$$

where $C_{\mathbf{u},t}$ is the consumption level of an agent unemployed for more than one period.[26]
If (17) holds as an inequality, which will be maintained, then newly unemployed consumption will equal:

$$C_{\mathbf{u},0,h,t} = \frac{R_{t-1}}{\Pi_t}b_{\mathbf{e},h-1,t} - T_{\mathbf{w},t} \tag{18}$$

Since newly unemployed do not carry any savings to the following period, the problem of those unemployed for more than one period is straightforward and implies that:

$$C_{\mathbf{u},t} = \vartheta - T_{\mathbf{w},t} \tag{19}$$

[25]Sufficient conditions for this will be described below.
[26]It is easy to verify whether this condition is satisfied in a steady state and thus in its neighborhood.

3.1.3 Firms

All firms are owned by capitalists, hence they will discount their cash flows by the capitalists' marginal utility of consumption.

Labor-service providers. A unit measure of labor service providers decide whether to open vacancies which, when matched, will result in the production of 1 unit of labor services to be sold to wholesale good firms. Unfilled vacancies get destroyed at the same rate $1-\rho_t$ as are jobs (Coles and Moghaddasi Kelishomi, 2018).
The real value of an unfulfilled vacancy v_t^0 is (I drop the firm-specific index for simplicity):

$$v_t^0 \quad = \quad -\iota + \beta \mathbb{E}_t \frac{C_{\mathrm{k},t+1}^{-\sigma_{\mathrm{k}}}}{C_{\mathrm{k},t}^{-\sigma_{\mathrm{k}}}} \rho_t \left[\mathfrak{q}_t v_{t+1}^1 + (1-\mathfrak{q}_t) v_{t+1}^0 \right] \tag{20}$$

where $\mathfrak{q}_t = \frac{M_t}{V_t}$, the proportion of vacancies filled in period t, v_t^1 is the real value of a matched vacancy, and ι is the flow cost of keeping a vacancy open. The value of opening a vacancy derives from the possibility of it becoming a productive match.
The real value of a filled vacancy is:

$$v_t^1 \quad = \quad p_t^N - w_t + \beta \mathbb{E}_t \frac{C_{\mathrm{k},t+1}^{-\sigma_{\mathrm{k}}}}{C_{\mathrm{k},t}^{-\sigma_{\mathrm{k}}}} \rho_t v_{t+1}^1 \tag{21}$$

where p_t^N is the price (in units of the final consumption good) that a unit of labor services sells for. Unfilled vacancies are created by firms who draw an opportunity cost $c \leq v_t^0$, where $c = F_t x,\ x \in [0,1],\ Pr\{x \leq x_0\} = x_0^{\psi_V}$, a Power-law distribution, scaled by F_t. This distribution of the vacancy-creating cost, delivers a law of motion for vacancies in line with the literature (Coles and Moghaddasi Kelishomi, 2018, and Broer, Druedahl, and Harmenberg, 2021):

$$V_t \quad = \quad \rho_{t-1}(V_{t-1} - M_{t-1}) + \tilde{F}_t v_t^{0\psi_V} \quad \tilde{F}_t \equiv F_t^{-\psi_V} \tag{22}$$

the first term represents the unmatched vacancies from the previous period that are not destroyed, the second term represent the flow of new vacancies.
The law of motion for the aggregate labor-services supply is given by:

$$N_t \quad = \quad \rho_{t-1} N_{t-1} + M_t = \rho_{t-1} N_{t-1} + \mathfrak{q}_t V_t \tag{23}$$

Labor-service providers will bargain over the marginal surplus of filling a vacancy as opposed to keeping it unmatched $v_t = v_t^1 - v_t^0$. Real dividends from labor-service providers are $d_t^l = -\iota V_t + \left(p_t^N - w_t\right) N_t - \int_0^{v_t^0} c\, df(c)$, where $\int_0^{v_t^0} c\, df(c) = \frac{\psi_V}{\psi_V+1}\left(\frac{v_t^0}{F_t}\right)^{\psi_V+1}$.

Employment agency. On the labor-supply side operate employment agencies. If workers want to access the labor market they have to go through it, as described above. The agency internalizes the search friction as it takes into account the fact that it will never be convenient for a worker to opt for autarky.
If I define $\Lambda_{i,t}$ as the number of workers under contract with agency i and $\tilde{N}_{i,t}$, the number of workers under contract with agency i that are currently employed, the cash flow for agency i is
$\tau^{a}\mathbb{E}_{t}\sum_{j=0}^{\infty}\beta^{j}\frac{C_{\mathrm{k},t+1}^{-\sigma_{\mathrm{k}}}}{C_{\mathrm{k},t}^{-\sigma_{\mathrm{k}}}}\left(w_{t+j}\tilde{N}_{i,t}+\vartheta\left(\Lambda_{i,t}-\tilde{N}_{i,t}\right)\right)$, subject to the law of motion for employed workers, determined by the search friction, $\tilde{N}_{i,t}=\rho\tilde{N}_{i,t-1}+\mathfrak{p}_{t}\left(\Lambda_{i,t}-\rho_{t-1}\tilde{N}_{i,t-1}\right)$.[27]
The agency will send a representative to the bargaining table with the aim of maximizing the value of having an extra employed worker:

$$\eta_{t} \quad = \quad (w_{t}-\vartheta)+\beta\mathbb{E}_{t}\frac{C_{\mathrm{k},t+1}^{-\sigma_{\mathrm{k}}}}{C_{\mathrm{k},t}^{-\sigma_{\mathrm{k}}}}\rho\left(1-\mathfrak{p}_{t+1}\right)\eta_{t+1} \tag{24}$$

where η_t is the Lagrange multiplier on the law of motion for employment. The surplus takes the conventional form in this class of models (Monacelli, Perotti, and Trigari, 2010), which is important in that I want to keep wage determination comparable to that of representative-agent models. Also, it is independent of τ^a.

As discussed above, it is convenient to focus on the limit case in with $\tau^a \longrightarrow 0$. As a result, all workers will elect to participate and $\Lambda_{i,t}=\mathfrak{w},\ \forall\, t$.

Wholesale Good Firm. A representative competitive wholesale good firm combines labor services and capital to produce good Y_t with the following CES technology $Y_{t}=A_{t}\left[\alpha\left(\kappa_{t}K_{t}\right)^{\frac{\nu-1}{\nu}}+(1-\alpha)\,N_{t}^{\frac{\nu-1}{\nu}}\right]^{\frac{\nu}{\nu-1}}$.[28]

It also accumulates capital, $K_{t+1}=(1-\delta)\,K_{t}+\left(1-S\left(\frac{I_{t}}{I_{t-1}}\right)\right)I_{t}$, subject to an investment adjustment cost and decides the level of capital utilization κ_t subject to cost $a\left(\kappa_{t}\right)=\frac{r^{K}}{\overline{a}}\left[e^{\overline{a}(\kappa_{t}-1)}-1\right]$ (Christiano, Motto, and Rostagno, 2014).
The discounted sum of dividends $\mathbb{E}_{t}\sum_{j=0}^{\infty}\beta^{j}\frac{C_{\mathrm{k},t+j}^{-\sigma_{\mathrm{k}}}}{C_{\mathrm{k},t}^{-\sigma_{\mathrm{k}}}}\left(p_{t+j}^{Y}Y_{t+j}-p_{t+j}^{N}N_{t+j}-I_{t+j}-a\left(\kappa_{t}\right)K_{t}\right)$ is maximized subject to the production function, and the law of motion for capital.

[27] It is convenient to assume that the fee is proportional to the level of home production even for the newly unemployed who live off their savings as opposed to home production. This assumption simplifies the derivation and allows me to obtain a level of the surpluses in line with the literature, and is immaterial insofar as I will consider the solution for $\tau^a \longrightarrow 0$.

[28] Clearly A_t can only enter this way because I am not interested in assuming it to be trending. Indeed it is treated as constant in the impulse-response matching exercise.

Firm decisions are characterized by the following first-order conditions:

$$\varphi_t = \beta \mathbb{E}_t \frac{C_{\mathrm{k},t+1}^{-\sigma_{\mathrm{k}}}}{C_{\mathrm{k},t}^{-\sigma_{\mathrm{k}}}} \left(r_{t+1}^K - a(\kappa_{t+1}) + (1-\delta)\varphi_{t+1} \right) \tag{25}$$

$$p_t^N = \varrho_t \tag{26}$$

$$1 = \varphi_t \left[\left(1 - S\left(\frac{I_t}{I_{t-1}} \right) \right) - S'\left(\frac{I_t}{I_{t-1}} \right) \frac{I_t}{I_{t-1}} \right] + \beta \mathbb{E}_t \frac{C_{\mathrm{k},t+1}^{-\sigma_{\mathrm{k}}}}{C_{\mathrm{k},t}^{-\sigma_{\mathrm{k}}}} S'\left(\frac{I_{t+1}}{I_t} \right) \frac{I_{t+1}^2}{I_t^2} \varphi_{t+1} \tag{27}$$

$$r_t^K = a'(\kappa_t)\kappa_t, \tag{28}$$

where φ_{t+j} is the Lagrange multiplier on the law of motion for capital, which represents the relative price of capital (in the sense of Christiano, Ilut, et al., 2010), $r_t^K = p_t^Y \alpha (A_t \kappa_t)^{\frac{\nu-1}{\nu}} \left(\frac{Y_t}{K_t} \right)^{\frac{1}{\nu}}$, the marginal product of capital, and $\varrho_t = p_t^Y (1-\alpha) A_t^{\frac{\nu-1}{\nu}} \left(\frac{Y_t}{N_t} \right)^{\frac{1}{\nu}}$ the marginal product of labor.

Intermediate Good Firms. A continuum of firms i buy the wholesale good Y_t and differentiate it into the good $Z_t(i)$ according to a simple production function $Z_t(i) = Y_t^d(i)$. The receive a subsidy τ so their real net marginal cost is $MC_t = (1-\tau) p_t^Y$. They face a decreasing demand-function $Z_t(i) = \left(\frac{P_t(i)}{P_t} \right)^{-\epsilon} Z_t$ in a monopolistically competitive market, and are subject to a nominal friction á la Rotemberg (1982) and are also subject to indexation $\overline{\Pi}^{1-\xi} \Pi_{t-1}^{\xi}$. They maximize the discounted flow of future dividends, which results in the following Phillips-curve relationship:

$$1 - \psi \left(\frac{\Pi_t}{\overline{\Pi}^{1-\xi} \Pi_{t-1}^{\xi}} - 1 \right) \frac{\Pi_t}{\overline{\Pi}^{1-\xi} \Pi_{t-1}^{\xi}} + \psi \beta \mathbb{E}_t \frac{C_{\mathrm{k},t+1}^{-\sigma_{\mathrm{k}}} Z_{t+1}}{C_{\mathrm{k},t}^{-\sigma_{\mathrm{k}}} Z_t} \left(\frac{\Pi_{t+1}}{\overline{\Pi}^{1-\xi} \Pi_t^{\xi}} - 1 \right) \frac{\Pi_{t+1}}{\overline{\Pi}^{1-\xi} \Pi_t^{\xi}} = \epsilon (1 - MC_t) \tag{29}$$

Final Good Firm. A competitive firm buys intermediate goods and bundles them together with CES technology $Z_t = \left[\int_0^1 Z_t(i)^{\frac{\epsilon-1}{\epsilon}} \, di \right]^{\frac{\epsilon}{\epsilon-1}}$. The cost-minimization problem together with the zero-profit condition results in the demand function $Z_t(i) = \left(\frac{P_t(i)}{P_t} \right)^{-\epsilon} Z_t$ and the price index $P_t = \left[\int_0^1 P_t(i)^{1-\epsilon} \, di \right]^{\frac{1}{1-\epsilon}}$. The final good can then be converted one for one into investment and consumption.

3.1.4 Matching

Employment and Unemployment. Matches are formed at the start of each period and resolved at the end, randomly, at an exogenous rate $1-\rho_t$. Aggregate employment thus evolves as $N_t = \rho_{t-1} N_{t-1} + \mathfrak{p}_t J_t$ where J_t is the number of job seekers. Unemployment is defined as $U_t = \frac{\mathfrak{w} - N_t}{\mathfrak{w}}$.

Surplus and wage. The total surplus is defined as $S_t = v_t + \eta_t$. The maximum acceptable wage is $\overline{w}_t = p_t^N + \iota + \beta \mathbb{E}_t \frac{C_{\mathrm{k},t+1}^{-\sigma_{\mathrm{k}}}}{C_{\mathrm{k},t}^{-\sigma_{\mathrm{k}}}} \rho (1 - \mathfrak{q}_t) v_{t+1}$. The minimum acceptable wage reads $\underline{w}_t = \vartheta - \beta \mathbb{E}_t \frac{C_{\mathrm{k},t+1}^{-\sigma_{\mathrm{k}}}}{C_{\mathrm{k},t}^{-\sigma_{\mathrm{k}}}} \rho (1 - \mathfrak{p}_{t+1}) \eta_{t+1}$, so that $S_t = \overline{w}_t - \underline{w}_t$.

The Nash-bargained wage is $w_t^* = \omega \overline{w}_t + (1-\omega)\, \underline{w}_t$. In my impulse-response matching exercise I allow for a more flexible solution that nests both pure Nash-bargaining, as well as fixed wages, proposed by Hall (2005) and adopted by Ravn and Sterk (2017): $w_t = \psi_w w_{ss} + (1-\psi_w)\, w_t^*$.

Matching. Matches are formed according to the following technology: $M_t = \Xi_t J_t^{\gamma} V_t^{1-\gamma} = \Xi_t V_t \theta_t^{-\gamma}$, where $\theta \equiv \frac{V_t}{J_t}$ is labor-market tightness. It follows then that $\mathfrak{q}_t = \frac{M_t}{V_t} = \Xi_t \theta_t^{-\gamma}$.

3.1.5 Government and Central Bank

For simplicity, I assume away government spending, and denote with a superscript s the supply of bonds. The government budget-constraint (in real terms) is:

$$b_{t+1}^s = \frac{R_{t-1}}{\Pi_t} b_t^s + \tau p_t^Y Z_t - T_t \tag{30}$$

I maintain that the government will set T_t so as to maintain a constant level of debt, which implies:

$$T_t = \left(\frac{R_{t-1}}{\Pi_t} - 1\right) b^s + \tau p_t^Y Z_t \tag{31}$$

I assume that capitalist bear the financing of the production subsidy as they are the only beneficiaries of it, while all agents participate in the financing of interest-rate government expenditures, as they all have access to the bond market:

$$T_{\mathrm{w},t} = \left(\frac{R_{t-1}}{\Pi_t} - 1\right) b^s \tag{32}$$

$$T_{\mathrm{k},t} = \frac{\tau p_t^Y Z_t}{1-\mathfrak{w}} + \left(\frac{R_{t-1}}{\Pi_t} - 1\right) b^s \tag{33}$$

The central bank follows a simple monetary policy rule:

$$R_t = R_{t-1}^{1-\rho_{MP}} \left(R_t^{\text{flex}} \Pi_t^{1+\phi} \left(\frac{U_t}{U_t^{\text{flex}}}\right)^{-\phi_u} \right)^{1-\rho_{MP}} . \tag{34}$$

where it is maintained that the inflation target equals steady state inflation and both are zero (in logs). The *flex* superscript refers to the flexible-price counterpart to the corresponding variable. In the numerical exercise I will consider variants of this policy rule in which the output gap replaces the unemployment gap, or in which the central bank does not respond to the flex-price level of the interest rate.

3.1.6 Market Clearing

Clearing on the stock market, where the supply of stocks is normalized to one, simply implies that $H_{\mathrm{k},t} = \frac{1}{1-\mathfrak{w}}$.

For what concerns the wholesale goods market it has to be that $\int_0^1 Y^d(i)\ di = Y_t$, which implies $Z_t = Y_t$, in an equilibrium without price dispersion - I focus on a symmetric solution to the Rotemberg-pricing problem.

Clearing on the labor market is trivially verified by noting that $\mathfrak{w}U_t 0 + \mathfrak{w}(1-U_t)1 + (1-\mathfrak{w})0 = N_t$, where the left-hand side represents the labor supply of workers (both employed and unemployed) and capitalists and the left-hand side the labor demand.

Clearing on the bonds market requires $\int b_{i,t+1}\ di = b^s$, a market in which demand schedules from different cohorts of workers are pooled together. Considering that only employed workers will hold positive amounts of bonds, the integral can be expressed as:

$$\mathfrak{w}(1-U_t)\sum_h \varkappa_{h,t} b_{\mathbf{e},h,t+1} = b^s \tag{35}$$

where $\varkappa_{h,t}$ is the share of cohort-h agents, within the the population of employed workers.

Shares are defined by computing the probability that an individual worker has been in a continued employment spell for the last h periods:

$$\varkappa_{0,t} = \mathfrak{p}_t \frac{U_{t-1}}{(1-U_t)} \tag{36}$$

$$\varkappa_{h,t} = \mu_t \varkappa_{h-1,t-1} \quad \forall\, 1 \leq h \leq H-1 \tag{37}$$

$$\varkappa_{H,t} = 1 - \sum_{h=0}^{H-1} \varkappa_{h,t} \tag{38}$$

The definition of $\varkappa_{H,t}$ is just a normalization. As H increases, its value will mechanically decrease. But this is irrelevant so long as the bond-holdings of agents in neighboring cohorts are practically identical for agents having been continuously employed for a sufficiently long number of periods.

It is convenient to define the average level of consumption by employed agents as $C_{\mathbf{e},h,t} = \sum_{h=0}^{H} \varkappa_{h,t} C_{\mathbf{e},h,t}$, and by newly unemployed agents $C_{\mathbf{u},0,t} = \sum_{h=1}^{H+1} \varkappa_{h-1,t-1} C_{\mathbf{u},h,0,t}$.[29] Then total consumption equals:

$$C_t = \mathfrak{w}U_t\left(\varkappa_{\mathbf{u},0,t} C_{\mathbf{u},0,t} + \left(1-\varkappa_{\mathbf{u},0,t}\right) C_{\mathbf{u},t}\right) + \mathfrak{w}(1-U_t)\, C_{\mathbf{e},t} + (1-\mathfrak{w})\, C_{\mathbf{k},t} \tag{39}$$

where $\varkappa_{\mathbf{u},0,t} = \frac{(1-\rho)(1-\mathfrak{p}_t)N_{t-1}}{\mathfrak{w}U_t}$ is the share of newly unemployed, within the population of all unemployed workers. Finally the resource constraint is given by:

$$Z_t + \mathfrak{w}U_t\left(1-\varkappa_{\mathbf{u},0,t}\right)\vartheta = C_t + I_t + \iota V_t + \frac{\psi_V}{\psi_V+1}\left(\frac{v_t^0}{F_t}\right)^{\psi_V+1} + a(\kappa_t)\,K_t + \frac{\psi}{2}\left(\frac{\Pi_t}{\overline{\Pi}^{1-\xi}\Pi_{t-1}^{\xi}} - 1\right)^2 Z_t. \tag{40}$$

[29]Note that the there cannot be cohort-0 unemployed, as that means that would have never left the unemployment pool in the first place.

On the left-hand side is the production of goods by firms and, in the form of home production, by the unemployed. The final good is used for consumption and investment purposes as well as to pay the various costs described above.

3.2 Calibration

I calibrate the parameters that affect the model's steady state. I set the share of workers to .9, as in Lansing (2015). The relative risk-aversion coefficient for workers $\sigma_{\mathbf{w}}$ is set to 1 (log preferences), and the discount factor β to .99, a commonly used value also employed by Cui and Sterk (2021). These calibrated parameters alongside the effects of individual risk, which pushes down on the equilibrium interest rate, result in a steady state value of the real interest rate of about 1 percent (.98 percent).

I set $\nu = .5$ which makes labor and capital gross complements, a moderate degree of complementarity according to estimates in the literature (Gechert et al., 2021; Klump, McAdam, and Willman, 2012; Cantore et al., 2015; Di Pace and Villa, 2016). A Cobb-Douglas production function would still preserve the qualitative properties of the responses, but complementarity improves the quantitative performance of the model. The steady state level of A_t is just a scaling parameter. Setting it to 1.6, while setting the depreciation rate $\delta = .025$, and $\alpha = .4$, yields a labor share of two thirds (66.7 percent).

The bargaining power of workers (ω) and the elasticity parameter of the matching function (γ) are both set to .5, as in Monacelli, Perotti, and Trigari (2010). I set $\psi_V = .3$, the parameter which governs the elasticity of vacancy creation to the value of vacancies, and $\iota = 0$, the flow cost of keeping a vacancy open. Both are in the line with the values proposed by Coles and Moghaddasi Kelishomi (2018).

I calibrate the steady-state value of the job-continuation rate ρ to .9, which implies that the average match lasts two and a half years (Shimer, 2005). I target a value of labor-market tightness of .5 (Monacelli, Perotti, and Trigari, 2010), and a labor-finding probability $\mathfrak{p} = .7$.[30] This value corresponds to a monthly job-finding probability of about a third, in between the value of 25.2 percent used in Ravn and Sterk (2021) and the value of 45 percent reported in Shimer (2005). Importantly, it implies a realistic value for steady state unemployment of about 4.1 percent.

The calibration of the supply of bonds b^s and of home production ϑ is key to determine the level of individual risk, that of the real rate of interest, and whether the *moderate liquidity* condition in equation (17) is satisfied. I set them so that the fall in consumption is of the order of 20 percent for a cohort-H worker losing her job, in line with Chodorow-Reich and Karabarbounis (2016), Cui and Sterk (2021), and Ravn and Sterk (2021).[31] These assumptions imply that, in equilibrium, the bond holdings of a cohort-H employed worker amount to about 10 weeks of his labor income, in line with the evidence that liquid savings are of the order of a few weeks of income even for wealthier households (Cui and Sterk, 2021).

[30] These two restrictions pin down the steady state values of F and Ξ.

[31] I set $b^s = 1$, $\vartheta = 1.3$, which imply that a cohort-H worker losing her job will see her consumption fall to 79.7 percent of the steady-state consumption level of a cohort-H worker in the first period – when agents cannot rely on home production – and to 80.4 percent in the subsequent periods when savings have been exhausted but unemployed agents can rely on home production. Under this parametrization, the left-hand side of equation (17) is larger than the right-hand side, ensuring that newly unemployed agents will consume all their accumulated savings in the first period.

Finally I calibrate the elasticity of substitution between intermediate-good varieties to 11 implying a markup of 10 percent,[32] and gross steady-state inflation $\Pi = 1$.

3.3 Impulse-Response Matching

Given the steady state described above, I match the model's impulse responses to those from my quarterly VAR, with a Bayesian impulse response matching procedure based on Christiano, Trabandt, and Walentin (2011).[33]

Impulse response matching gives full control as to the set of shocks used to match up to the empirically observed patterns. The remaining question is which structural shocks capture the labor-market news reported by survey respondents. This parallels the problem faced by Beaudry and Portier (2013) when they had to translate a discussion about a generic "demand shock" into well-defined model shocks.

I base my selection on the literature, my VAR evidence, and the qualitative features of responses to different shocks in my model.

At least since Shimer (2005), job-separation shocks have played a prominent role in the literature on the business cycle properties of search-and-matching models. That is the one shock Shimer (2005) considers, other than a TFP shock.

My empirical analysis rules out a shock with a direct impact on the first-moment of measured TFP. Moreover the shock under consideration has to have a limited impact on inflation, induce positive co-movement between consumption and investment, and a negative relationship between unemployment and vacancies, tracing out a *Beveridge* curve.

There are two-labor market related shocks that meet these criteria, in the context of my model economy. One is the aforementioned job-separation shock, the other is a shock to the cost of opening vacancies. They both produce impulse responses of the same sign as VAR responses. This is not true, however, for matching-efficiency shocks, as they drive a positive comovement between vacancies and unemployment.[34]

Based on this consideration, my baseline impulse-response matching exercise will consider these two shocks as possible drivers of labor-conditions news.[35]

I set up the IRF-matching procedure by assuming there exists an underlying process:

$$x_t = \rho_x x_{t-1} + (1 - \theta_x)\, u_t^x + \theta_x u_{t-1}^x \tag{41}$$

where $u_t^x \sim \mathcal{N}\left(0, \sigma_{u^x}^2\right)$ and independent over time. The nature of the survey question I use in my VAR is such that the news-shock I am considering is orthogonal to all *past* realizations of macro variables, but not current ones. The specification in (41) allows for the possibility that the shocks in question start producing

[32] The production subsidy is calibrated to the optimal value of $\tau = 1/\epsilon$.

[33] As implemented in Dynare by Gauthier, 2021.

[34] Figures D.30 and D.29, in the Appendix, confirm my claim, reporting the outcome of a matching experiments in which each of the two shocks is considered in isolation. Figure D.31 reports the matching that obtains when a matching efficiency shock is added to the mix.

[35] In the Appendix, I will show how each of them fares when taken in isolation, and how the three of them combine together.

their effects in the current period, or only in the following one - the parameter θ_x will be estimated.

I then posit that the two processes in question evolve as:

$$\log(\rho_t) = (1-\rho_\rho)\log(\rho_{ss}) + \rho_\rho \log(\rho_{t-1}) + +u_t^\rho + \theta_\rho x_t \quad (42)$$
$$\log(\tilde{F}_t) = (1-\rho_F)\log\left(\tilde{F}\right) + \rho_F \log(\tilde{F}_{t-1}) + u_t^F + \theta_F x_t. \quad (43)$$

I will estimate both θ_ρ and θ_F so that the impulse-response matching procedure can attach different weights to the two shocks.

In my baseline, I estimate the parameters pertaining to the forcing process (θ_x, ρ_x, θ_ρ, ρ_ρ, θ_F, ρ_F),[36] the parameters of the monetary policy rule (ϕ, ϕ_u, ρ_R), and those governing the risk-aversion of capitalists (σ_{k}), the utilization, price and investment adjustment costs ($\overline{a}$, ψ, S'') and the degree of price indexation (ζ). Finally I estimate ψ_w, which governs the degree of wage stickiness.

Figure 7 presents the resulting model responses.[37] The model responses meet the key features of the empirical IRFs: a strong negative comovement between unemployment and vacancies, the positive co-movement between consumption and investment, the negligible response of inflation and the pro-cyclical response of Fed Funds rates.

The matching procedure is also informative with regards to the parameter estimates. Capitalists are estimated to be almost risk-neutral, an assumption often made in the literature (Ravn and Sterk, 2021). Investment-adjustment costs are estimated to be basically zero. This contrasts with comparable, representative-agent models, such as Monacelli, Perotti, and Trigari (2010), in which investment-adjustment costs play a prominent role. The model does not require an excessively flat Phillips curve. The prior for ψ is centered around a price-adjustment cost level, that translates into a price duration of 4 quarters in a Calvo setting, and it is diffuse (standard deviation of 30). The posterior mean would correspond to an average price duration of less that 4-and-a-half quarters. The monetary policy rule parameter estimates imply that nominal rates respond 1.7 times the deviation of inflation from target, in line with Taylor (1993) and the extensive literature that followed.

Importantly, the posterior mean for ψ_w is .87, i.e. the matching procedure favors a wage process that does not depart much from its steady state value. That of constant wages is quite a common assumption in this literature (Broer, Harbo Hansen, et al., 2020) and dates back all the way to Hall (2005). Gertler, Huckfeldt, and Trigari (2020) provides new empirical evidence for why this may be a reasonable assumption in macro models, while Christiano, Eichenbaum, and Trabandt (2016) provides a possible microfoundation for wage stickiness.

Wage rigidity, while certainly preferred by the data, is not strictly necessary to match the VAR responses, at least qualitatively. Neither is a particularly flat Phillips curve – a criticism leveled by Angeletos, Collard, and Dellas (2020) to state-of-the-art representative agents models –, or the presence of R_t^{flex} in the

[36] σ_{u^x} cannot be identified separately from θ_F and θ_ρ. I thus normalize it to .02.

[37] Priors and posterior estimates are presented in Table D.2.

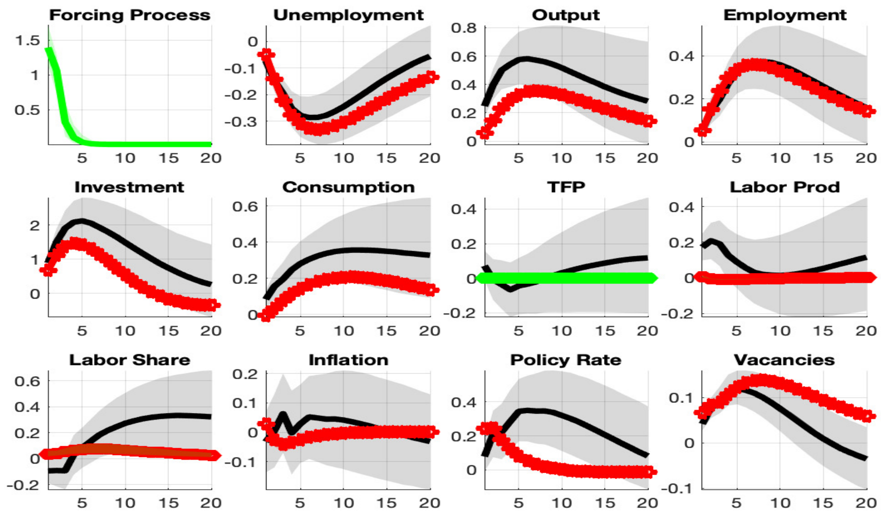

Figure 7: Impulse response matching to a quarterly VAR with employment. Red lines represent impulse-responses from the model which are matched to the VAR responses, while green lines represent impulse responses that are left unrestricted.

policy rule.[38]

Overall, this exercise does a good job at matching up to my VAR evidence.[39] This shows the potential for a fairly simple DSGE model to match up to empirical evidence without falling into Angeletos, Collard, and Dellas (2020) critique, i.e. using combinations of shocks, real frictions like investment-adjustment costs and consumption habits, and overly flat Phillips curves. All this without resorting to shocks to current or future expected TFP.

But what if TFP was actually somehow looming behind consumer reports of labor-market news despite this not showing up in the response of measured TFP? If I take my baseline parameter estimates and feed a TFP shock[40] into the model I get that it does an overall very good job at matching labor market variables. In part that is by design as I scale its size so that the peak response of unemployment is similar to that we observe in the data. The muted response of inflation and the procyclical response of the policy rate are also prima facie consistent with the VAR responses – Figure D.32.

However, the labor productivity response to such shock would be about ten times as large as that I

[38] Figure D.26 reports the matched IRFs for a model in which none of these features are present. For a good measure, I also set $\rho^x = 0$ to avoid the potential for excessive persistence stemming from the forcing process. I finally completely turn off any investment-adjustment cost.

[39] My battery of robustness exercises is completed by Figure D.28 in which I replace the unemployment gap with a measure of the output gap, and Figure D.27 in which I match the model responses up to a VAR that includes only news to the labor-conditions sub index.

[40] Scaling it so that the unemployment response peaks at about -.3 as in the data and with a news component captured by the estimate of θ_x.

observe in the data. This is another way of getting to a point Shimer (2005) makes (section E). The outsized labor-productivity response is the mirror image of an implausibly large response of output, investment and consumption. Finally, while the VAR identifies a small and largely procyclical response of the labor share, a TFP shock would generate a sharp fall in it.[41] This shock clearly would not make the cut if it came down to matching the VAR IRFs in isolation. The scale of the responses of aggregate demand components, and the countercyclical labor share are not consistent with the data.

4 Conclusion

I use information from the University of Michigan Survey of Consumers to re-examine the role of news in business cycle fluctuations, in the spirit of Cochrane (1994). Surprise reports of changes in business conditions explain a large share of the variation of real macro variables, but are essentially orthogonal to measured TFP and inflation. This is in line with the so-called *business-cycle anatomy* described by Angeletos, Collard, and Dellas (2020).

News related to the labor market are the main driver of these effects. This guides my modeling exercise, in which I extend a standard New-Keynesian economy with search frictions and physical capital by modeling explicitly individual risk, and the limited elasticity of vacancy creation to the state of the economy. Moreover, I account for complementarity between capital and labor inputs and strictly positive levels of liquidity in the economy. The resulting THANK-SAM model does a good job at reproducing the VAR-based impulse responses using a job-separation shock and a shock to the cost of opening vacancies (or each of them in isolation). In particular, this model generates positive comovement between consumption and investment, negative covariance between vacancies and unemployment, and a muted response of inflation, without resorting to combinations of shocks, an excessively flat Phillips curve, or real frictions such as consumption habits and investment-adjustment costs.

References

Ang, Andrew, Geert Bekaert, and Min Wei (May 2007). "Do Macro Variables, Asset Markets, or Surveys Forecast Inflation Better?" In: *Journal of Monetary Economics* 54.4, pp. 1163–1212.

Angeletos, George-Marios, Fabrice Collard, and Harris Dellas (2018). "Quantifying Confidence". In: *Econometrica* 86.5, pp. 1689–1726.

— (2020). "Business-Cycle Anatomy". In: *American Economic Review* 110.10, pp. 3030–3070.

Barnichon, Regis (2010). "Building a Composite Help-Wanted Index". In: *Economics Letters* 109.3, pp. 175–178.

Barro, Robert J. and Robert G. King (1984). "Time-Separable Preferences and Intertemporal-Substitution Models of Business Cycles". In: *Quarterly Journal of Economics*, pp. 154–177.

[41] See Chahrour, Chugh, and Potter (2021) for a recent discussion on the response of the labor share to a TFP news shock in a search-and-matching framework.

Barsky, Robert B and Eric R Sims (2012). “Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence”. In: *American Economic Review* 102.4, pp. 1343–1377.

Barsky, Robert B., Susanto Basu, and Keyoung Lee (2014). “Whither News Shocks?” In: *NBER Macroeconomics Annual* 29.1.

Barsky, Robert B. and Eric R. Sims (2011). “News Shocks and Business Cycles”. In: *Journal of Monetary Economics* 58.3, pp. 273–289.

Beaudry, Paul and Franck Portier (2006). “Stock Prices, News, and Economic Fluctuations”. In: *American Economic Review* 96.4, pp. 1293–1307.

— (2013). “Understanding Non-Inflationary Demand Driven Business Cycles”. In: *NBER Macroeconomics Annual* 28.

— (2014). “News Driven Business Cycles: Insights and Challenges”. In: *Journal of Economic Literature* 52.4, pp. 993–1074.

Ben Zeev, Nadav and Hashmat Khan (Oct. 2015). “Investment-Specific News Shocks and U.S. Business Cycles”. In: *Journal of Money, Credit and Banking* 47.7, pp. 1443–1464.

Broer, Tobias, Jeppe Druedahl, and Karl Harmenberg (2021). “Labor-Market Hysteresis and Persistent Paradox-of-Thrift Recessions”. In: *mimeo*.

Broer, Tobias, Niels-Jakob Harbo Hansen, et al. (2020). “The New Keynesian Transmission Mechanism: A Heterogeneous-Agent Perspective”. In: *The Review of Economic Studies* 87.1, pp. 77–101.

Cantore, Cristiano et al. (2015). “CES Technology and Business Cycle Fluctuations”. In: *Journal of Economic Dynamics and Control* 61, pp. 133–151.

Chahrour, Ryan, Sanjay K Chugh, and Tristan Potter (2021). “Anticipated Productivity and the Labor Market”. In: *mimeo*.

Challe, Edouard (2020). “Uninsured Unemployment Risk and Optimal Monetary Policy in a Zero-Liquidity Economy”. In: *American Economic Journal: Macroeconomics* 12.2, pp. 241–283.

Chodorow-Reich, Gabriel and Loukas Karabarbounis (2016). “The Cyclicality of the Opportunity Cost of Employment”. In: *Journal of Political Economy* 124.6.

Christiano, Lawrence, Cosmin L Ilut, et al. (2010). “Monetary Policy and Stock Market Booms”. In: *National Bureau of Economic Research Working Paper Series* No. 16402.

Christiano, Lawrence J, Roberto Motto, and Massimo Rostagno (2014). “Risk Shocks”. In: *American Economic Review* 104.1, pp. 27–65.

Christiano, Lawrence J., Martin S. Eichenbaum, and Mathias Trabandt (2016). “Unemployment and Business Cycles”. In: *Econometrica* 84.4, pp. 1523–1569.

Christiano, Lawrence J., Mathias Trabandt, and Karl Walentin (2011). “DSGE Models for Monetary Policy Analysis”. In: *Handbook of Monetary Economics* 3.C.

Cochrane, John H. (1994). “Shocks”. In: *Carnegie-Rochester Conference Series on Public Policy* 41, pp. 295–364.

Coles, Melvyn G. and Ali Moghaddasi Kelishomi (2018). "Do Job Destruction Shocks Matter in the Theory of Unemployment". In: *American Economic Journal: Macroeconomics* 10.3, pp. 118–136.

Cui, Wei and Vincent Sterk (2021). "Quantitative Easing". In: *Journal of Monetary Economics* 123 (October), pp. 68–90.

Den Haan, Wouter J. and Georg Kaltenbrunner (2009). "Anticipated Growth and Business Cycles in Matching Models". In: *Journal of Monetary Economics* 56.3, pp. 309–327.

Den Haan, Wouter J., Pontus Rendahl, and Markus Riegler (2018). "Unemployment (Fears) and Deflationary Spirals". In: *Journal of the European Economic Association* 16.5, pp. 1281–1349.

Di Pace, Federico and Stefania Villa (2016). "Factor Complementarity and Labour Market Dynamics". In: *European Economic Review* 82, pp. 70–112.

Faccini, Renato and Leonardo Melosi (2021). "Pigouvian Cycles". In: *American Economic Journal: Macroeconomics* forthcoming.

Fernald, John G. (2014). "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity". In: *Federal Reserve Bank of San Francisco, Working Paper Series*, pp. 01–28.

Ferroni, Filippo and Fabio Canova (2020). "A Hitchhiker Guide to Empirical Macro Models". In: *mimeo*.

Furlanetto, Francesco and Nicolas Groshenny (2016). "Mismatch Shocks and Unemployment During the Great Recession". In: *J. Appl. Econ.*, p. 18.

Gauthier, David (2021). "Bank Shocks and the Debt Structure". In: *SSRN Electronic Journal*.

Gechert, Sebastian et al. (2021). "Measuring Capital-Labor Substitution: The Importance of Method Choices and Publication Bias". In: *Review of Economic Dynamics*.

Gertler, Mark, Christopher Huckfeldt, and Antonella Trigari (2020). "Unemployment Fluctuations, Match Quality, and the Wage Cyclicality of New Hires". In: *Review of Economic Studies* 87.4, pp. 1876–1914.

Gertler, Mark and Peter Karadi (2015). "Monetary Policy Surprises, Credit Costs, and Economic Activity". In: *American Economic Journal: Macroeconomics* 7.1.

Gertler, Mark and Antonella Trigari (2009). "Unemployment Fluctuations with Staggered Nash Wage Bargaining". In: *Journal of Political Economy* 117.1, pp. 38–86.

Giannone, Domenico, Michele Lenza, and Giorgio E Primiceri (2015). "Prior Selection for Vector Autoregressions". In: *The Review of Economic Studies*.

Gomes, Joao, Jeremy Greenwood, and Sergio Rebelo (1997). "Equilibrium Unemployment". In: *NBER WORKING PAPER SERIES*.

Hall, Robert E (2005). "Employment Fluctuations with Equilibrium Wage Stickiness". In: *American Economic Review* 95.1, pp. 50–65.

Jaimovich, Nir and Sergio Rebelo (2009). "Can News about the Future Drive the Business Cycle?" In: *American Economic Review* 99.4, pp. 1097–1118.

Kehrig, Matthias (2015). "The Cyclical Nature of the Productivity Distribution". In: Working Paper.

Klump, Rainer, Peter McAdam, and Alpo Willman (2012). "The Normalized CES Production Function: Theory and Empirics". In: *Journal of Economic Surveys* 26.5, pp. 769–799.

Krusell, Per, Toshihiko Mukoyama, and Ayşegül Şahin (Oct. 2010). “Labour-Market Matching with Precautionary Savings and Aggregate Fluctuations”. In: *Review of Economic Studies* 77.4, pp. 1477–1507.

Kurmann, André and Eric Sims (May 14, 2021). “Revisions in Utilization-Adjusted TFP and Robust Identification of News Shocks”. In: *The Review of Economics and Statistics* 103.2, pp. 216–235.

Lansing, Kevin J (2015). “Asset Pricing with Concentrated Ownership of Capital and Distribution Shocks”. In: *American Economic Journal: Macroeconomics* 7.4, pp. 67–103.

Mckay, Alisdair and Ricardo Reis (2020). “Optimal Automatic Stabilizers”. In: *Nber Working Paper Series* 22359 (June), pp. 1–65.

Mertens, Karel and Morten Ravn (2012). “Empirical Evidence on the Aggregate Effects of Anticipated and Unanticipated US Tax Policy Shocks”. In: *American Economic Journal: Economic Policy* 4.2, pp. 145–181.

Milani, Fabio and John Treadwell (2012). “The Effects of Monetary Policy ”News” and ”Surprises””. In: *Journal of Money, Credit and Banking* 44.8, pp. 1667–1692.

Miyamoto, Wataru and Thuy Lan Nguyen (2020). “The Expectational Effects of News in Business Cycles: Evidence from Forecast Data”. In: *Journal of Monetary Economics* 116, pp. 184–200.

Monacelli, Tommaso, Roberto Perotti, and Antonella Trigari (2010). “Unemployment Fiscal Multipliers”. In: *Journal of Monetary Economics* 57.5, pp. 531–553.

Mortensen, D. T. and C. A. Pissarides (July 1, 1994). “Job Creation and Job Destruction in the Theory of Unemployment”. In: *The Review of Economic Studies* 61.3, pp. 397–415.

Pigou, A. C. (1927). *Industrial Fluctuations*. Mamillan & Co., Ltd.

Ramey, Valerie A (2016). “Macroeconomic Shocks and Their Propagation”. In: *NBER*.

Ravn, Morten O. and Vincent Sterk (2017). “Job Uncertainty and Deep Recessions”. In: *Journal of Monetary Economics* 90, pp. 125–141.

— (2021). “Macroeconomic Fluctuations with HANK & SAM: An Analytical Approach”. In: *Journal of the European Economic Association* 19.2, pp. 1162–1202.

Rotemberg, Julio J. (1982). “Monopolistic Price Adjustment and Aggregate Output”. In: *The Review of Economic Studies* 49.4, p. 517.

Schmitt-Grohé, Stephanie and Martín Uribe (2012). “What’s News in Business Cycles”. In: *Econometrica* 80.6, pp. 2733–2764.

Shimer, Robert (2005). “The Cyclical Behavior of Equilibrium Unemployment and Vacancies”. In: *The American Economic Review* 95.1, pp. 25–49.

— (2009). “Convergence in Macroeconomics : The Labor Wedge”. In: *American Economic Journal: Macroeconomics* 1.1, pp. 280–297.

Swanson, Eric T and John C Williams (2014). “Measuring the Effect of the Zero Lower Bound on Monetary Policy”. In: *American Economic Review* 104.10, pp. 3154–3185.

Taylor, John B (1993). “Discretion Practice versus Policy Rules in Practice”. In: *Carnegie-Rochester Conference Series on Public Policy* 39, pp. 195–214.

Theodoridis, Konstantinos and Francesco Zanetti (2016). “News Shocks and Labour Market Dynamics in Matching Models”. In: *Canadian Journal of Economics* 49.3, pp. 906–930.

A Data

The key data series I use is question A6 in the University of Michigan Survey of Consumers, which asks: *During the last few months, have you heard of any favorable or unfavorable changes in business conditions?* This has two implications:

1. It *directly* inquires about news instead of having to infer it from other variables. Also, it helps to focus on the news that are perceived as such by *consumers*.

2. It makes for a naturally *predetermined* variable because it asks about past months.

This makes for a straightforward identification scheme by providing a timing restriction. Moreover, Table 24 provides a classification about various kinds of news.

Michigan Survey of Consumer data is available monthly since 1978. I can thus estimate monthly SVARs starting in 1978 (ending in December 2019) including the log of industrial production, the log of CPI, the civilian unemployment rate, the two-year interest rate and the Michigan Survey measure of inflation expectations. These series are taken from the Federal Reserve Bank of St. Louis' FRED database. To these I add series from Tables 23 and 24 of the Survey of Consumers.

For the quarterly specification, I largely rely on the series definitions in Angeletos, Collard, and Dellas (2020) with all data taken from FRED, with the exception of the vacancies series which is based on Barnichon (2010).

B VAR Robustness checks

B.1 Monthly and Mixed-Frequency Specifications

Figure B.8 reports the estimates of my baseline monthly specification, with OLS estimates overlaid, to show that the priors, while making the estimation more amenable to short samples and large sets of variables, do not affect my findings.

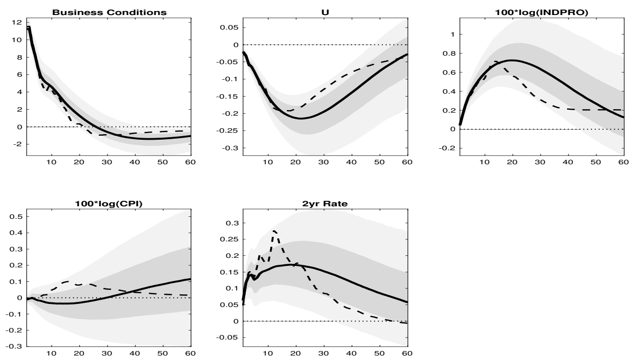

Figure B.8: Baseline specification of the monthly VAR, with OLS estimates overlaid.

Figure B.9 reports the estimates of my baseline monthly specification estimated over the post 1984 and over the 1984-2007 samples respectively, to verify that my key results do not depend neither on the ”Volcker recession” of the early 1980s, nor on the Great Recession and the Zero-Lower Bound period.

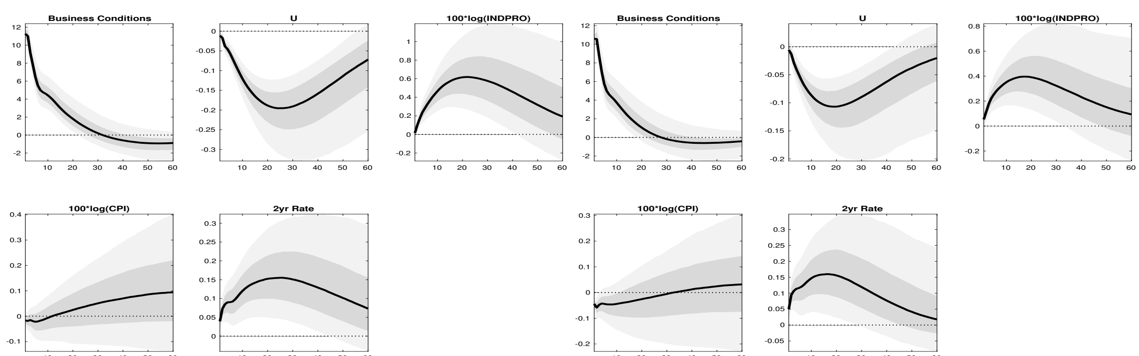

Figure B.9: Baseline specification of the monthly VAR estimated over the post-1984 sample (left pane) and the Great Moderation sample (1984-2007, right pane).

The informational content of the business conditions series is critical for the economic interpretation of my identified impulse responses. I will thus consider more restrictive specifications in which, information from the Survey of Professional Forecasters, and weekly information about unemployment claims are added to the VAR.

Figure B.10 presents IRFs from two specifications in which unemployment-expectations from the Survey of Professional Forecasters are added - one includes the overall business conditions index, the other only the labor-market index. Including forward-looking information from professionals controls for key information regarding the labor market that may not be reflected in past values of macro variables. The VAR has to be estimated with mixed-frequency techniques as the SPF survey is quarterly. Key properties of the responses are not affected. Moreover, this specification shows that in the wake of positive surprises in the business conditions index they revise down the forecasts for unemployment, which reinforces the idea that the business conditions index picks up relevant economic developments.

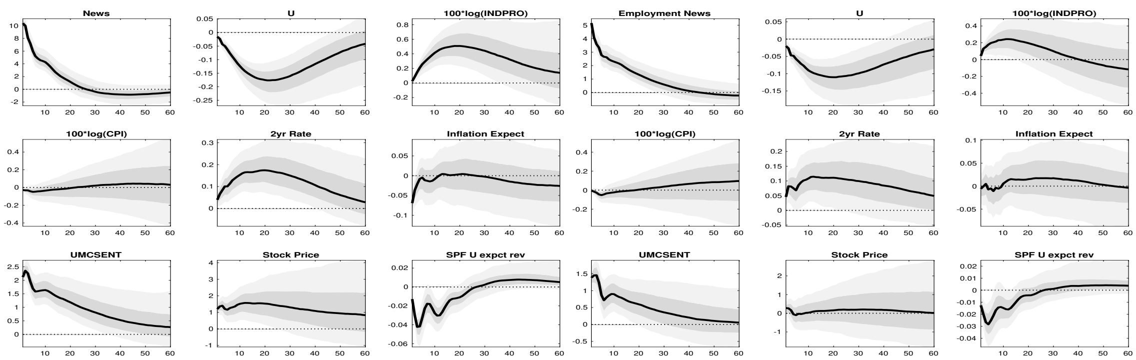

Figure B.10: Impulse Responses from a mixed-frequency VAR that includes revisions of SPF unemployment forecasts. The shock is given by surprises in the overall business conditions index (left), and the labor-market conditions index only (right).

Figure B.11 reports IRFs from a VAR in which the monthly business conditions index is made orthogonal to the key weekly unemployment indicators (for the second week of the month), initial and continued unemployment claims. The responses of unemployment, industrial production, CPI and the interest rate are hardly affected, the measures of claims respond as expected, negatively, to a positive business conditions shock.

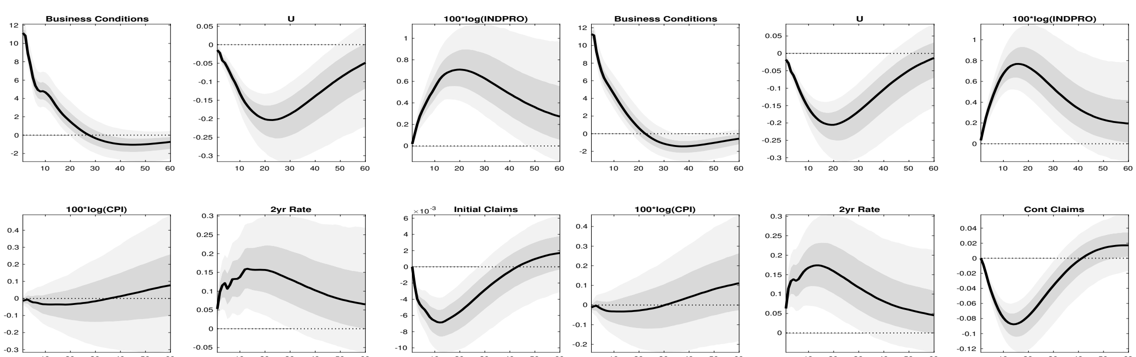

Figure B.11: Monthly VAR in which the shock is made orthogonal to the weekly series for Initial Unemployment Claims (left) and Continued Unemployment Claims (right).

The business conditions series is logically predetermined relative to the realization of current-period macro variables, as well as relative to forward-looking survey questions, such as inflation expectations or any question about the outlook for the future. The best way to control for intra-monthly information is to control for higher-frequency indicators such as unemployment claims, as presented above. However, it is common in the VAR literature to check the responses under an inversion of the Cholesky ordering, i.e. a situation in which shocks to the business conditions index are made orthogonal not only to past but also to current realizations of the other variables included in the VAR.

Figures B.12, B.13, B.14 reports the IRFs under an "inverse" Cholesky ordering for the specifications presented above. For a good measure, Figure B.15 displays results for a specification in which the measure of 5-year outlook from the UMSC is added to the mix (Barsky and Sims, 2012). The fact that the responses are qualitatively the same and remain significant, and always strongly so for unemployment, is a strong indication that the business conditions index genuinely Granger causes the other series (Barsky and Sims, 2012).

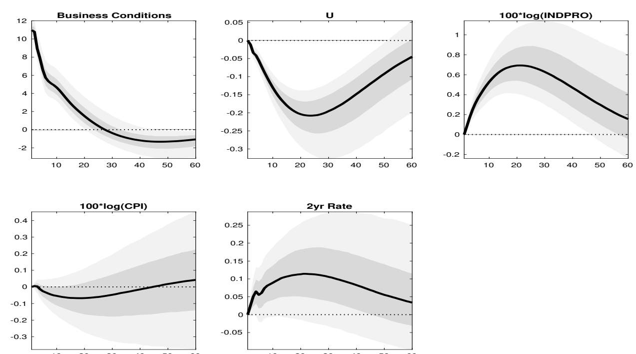

Figure B.12: Baseline specification of the monthly VAR, Business Conditions index ordered last in the Cholesky decomposition.

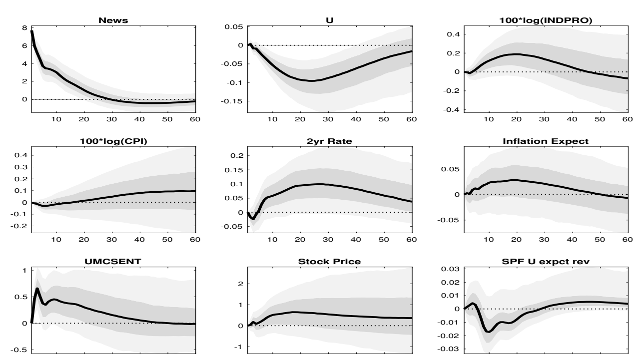

Figure B.13: Mixed-frequency specification with SPF revisions of unemployment forecasts. Inverse Cholesky ordering.

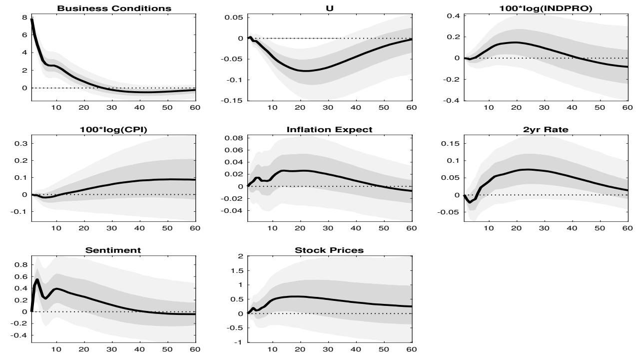

Figure B.14: Specification of the monthly VAR that includes Consumer Sentiment, inflation expectations and stock prices; Business Conditions index ordered last in the Cholesky decomposition.

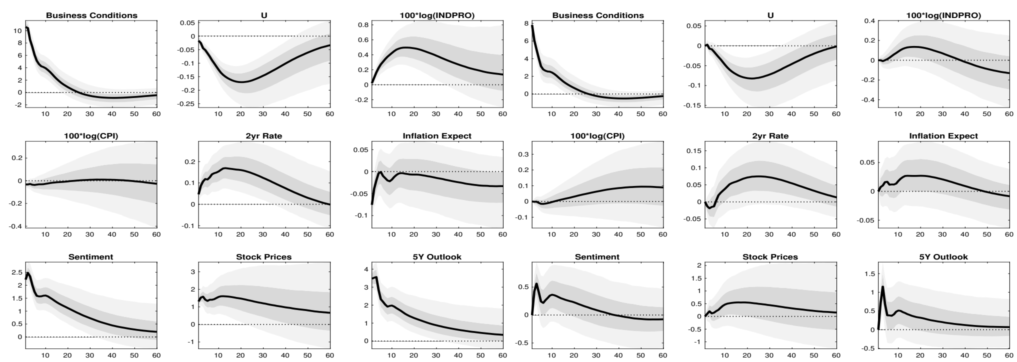

Figure B.15: Specification of the monthly VAR that also includes the 5-year Outlook variable (left); inverse Cholesky ordering (right).

The comovement of consumption and investment in response to the business conditions shock is critical to the understanding of its properties and in guiding the development of a suitable model. A mixed-frequency specification allows me to maintain a weaker timing restriction, while observing the responses of consumption and investment, which are only observed at a quarterly frequency. A strong positive comovement can be observed in Figure B.16, with investment responding more than consumption, in line with quarterly specifications and long-established business-cycle regularities.

Finally, Figure B.17 reports the IRFs for VAR specifications including the individual indices, one at a time.

B.2 Quarterly

For ease of comparison, I start off by estimating the quarterly counterpart to the baseline monthly specification presented above. Figure B.18 presents the impulse responses and the FEVD decomposition. Both are in line with the findings from the monthly VAR.

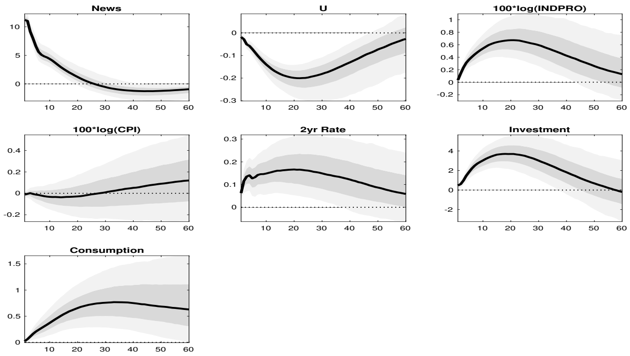

Figure B.16: Impulse Responses from a mixed-frequency VAR that includes investment and consumption.

A timing-identification scheme poses an inherently stronger restriction in a quarterly specification, than in a monthly one. For instance, it could be that survey responses in the third month of a quarter depended strongly on, say, unemployment in the first two months. To verify this phenomenon is not affecting my results, I present the IRFs to a VAR in which the quarterly series for the business conditions index is given by the level of the index in the first month of the quarter, as opposed to all three months as in the baseline, in Figure B.19. Responses are remarkably similar to the baseline in Figure B.18, also in terms of magnitudes.

Turning to the 12-variable quarterly VAR specification, Figure B.20 presents the Forecast-Error Variance Decomposition, corresponding to the shock presented in Figure 6 in the main body of the text. The shock explains between a third and a half of the variance in unemployment, output, employment, investment, and vacancies. The share of variance of consumption explained is around a quarter, for the Fed Funds rates of the order of 15 percent, and less than 10 percent for productivity, the labor share, inflation, and TFP.

The charts below present a series of robustness checks that parallel those reported for the monthly specifications. Inverting the Cholesky ordering of the variables does not affect the key results, Figure B.21; neither does restricting the attention to the business conditions news from the first month of the quarter, Figure B.22; or to news specific to the labor market, Figure B.23.

As explained in the main text, my 12-variable VAR specification uses primarily the variables definitions from Angeletos, Collard, and Dellas (2020), with three exceptions. I add the business conditions index and a measure of vacancies, and I replace hours with employment, as I do not explicitly model the intensive margin of labor. The specification presented in Figure B.24 does not include vacancies, and replaces employment with hours. The dashed lines represent the responses to a shock – in a 10-variable specification that excludes the business conditions index as well to be as close as possible to Angeletos, Collard, and Dellas (2020) – identified with an agnostic max-share approach, as the shock explaining the maximum variance share of unemployment. The similarities in the responses are striking in terms of sign, magnitude and timing.

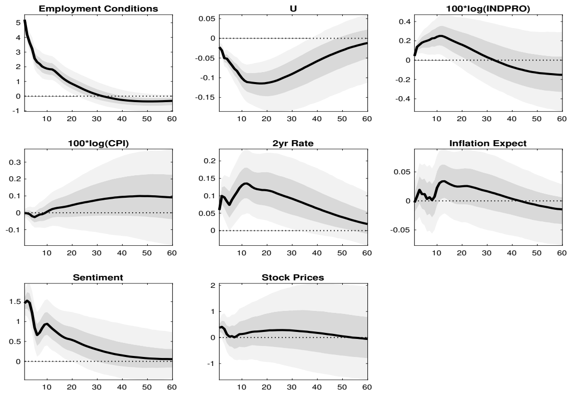

(a) Employment conditions

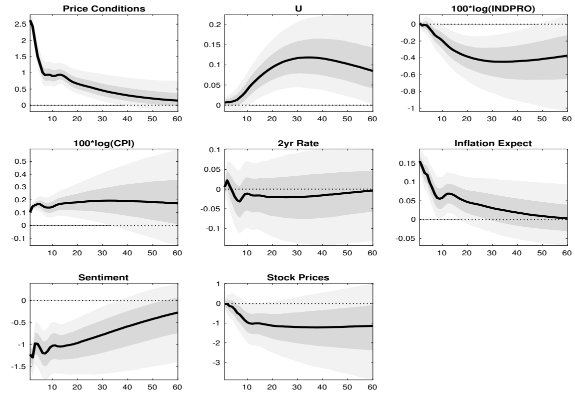

(b) Price conditions

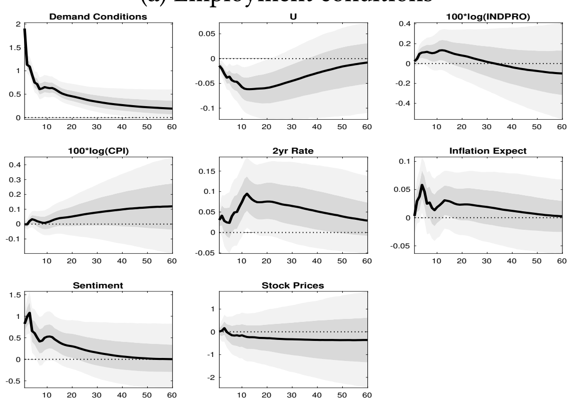

(c) Demand conditions

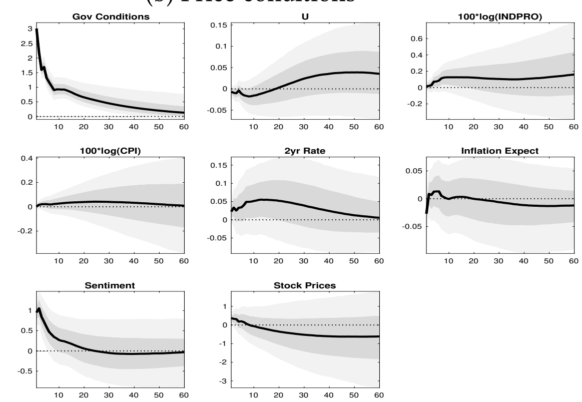

(d) Government conditions

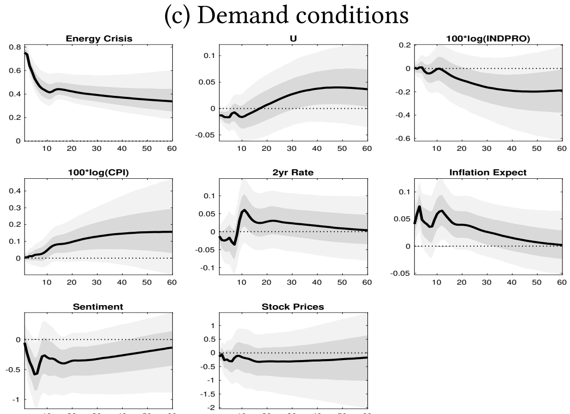

(e) Energy Crisis

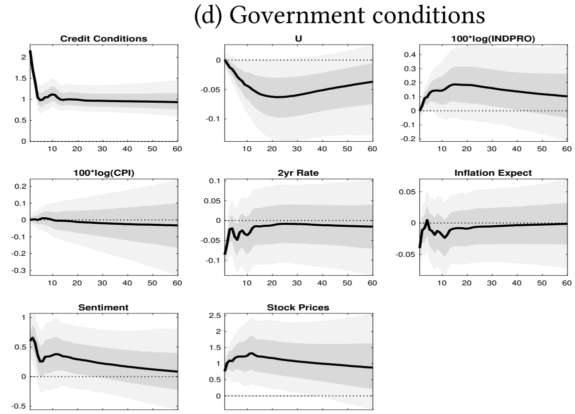

(f) Credit conditions

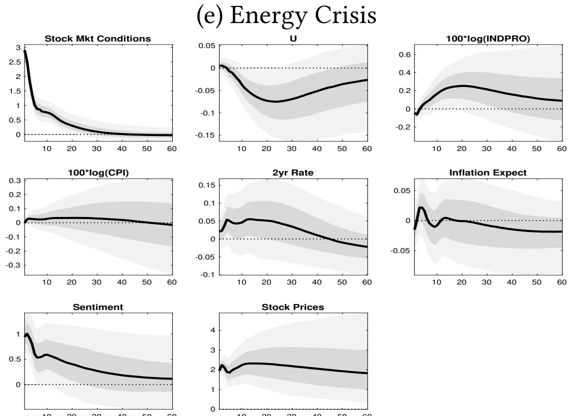

(g) Stock market

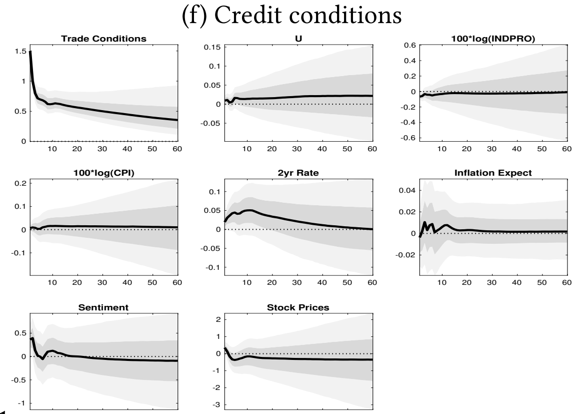

(h) Trade

Figure B.17: Impulse responses to shocks to the individual components of the business-conditions index.

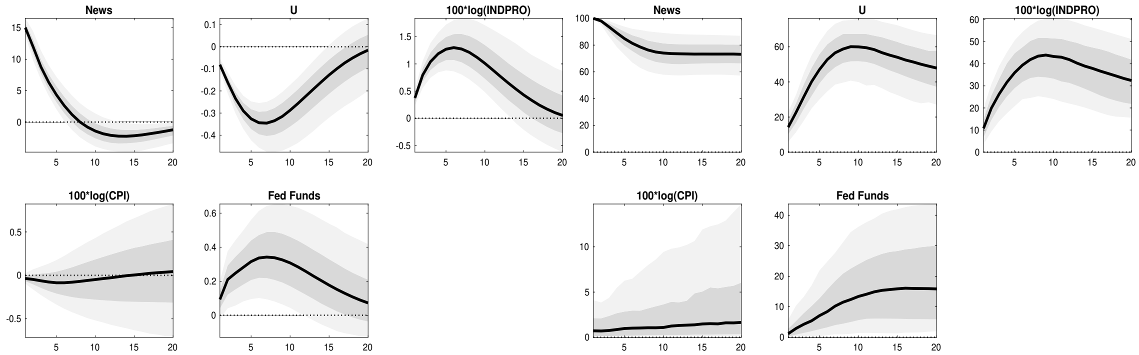

Figure B.18: Quarterly VAR, baseline specification (left), and corresponding forecast error variance decomposition (right).

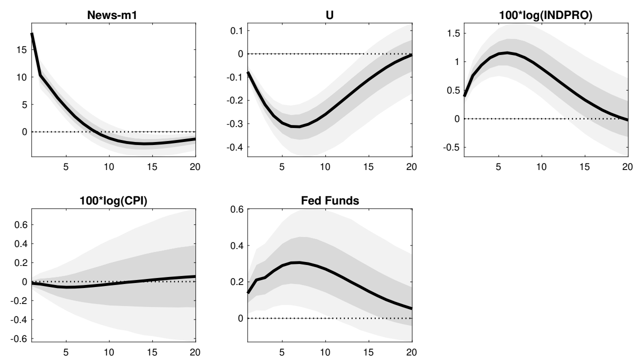

Figure B.19: Quarterly VAR, baseline specification, except for only using survey information from the first month of the quarter.

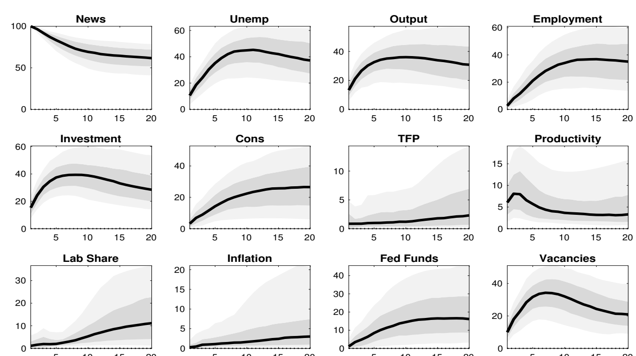

Figure B.20: Quarterly 12-variable specification, Forecast-Error Variance Decomposition.

Finally, Figure B.25 shows that including 4 lags in my quarterly specifications, as opposed to 2 in the baseline specification in Angeletos, Collard, and Dellas (2020), is effectively immaterial.

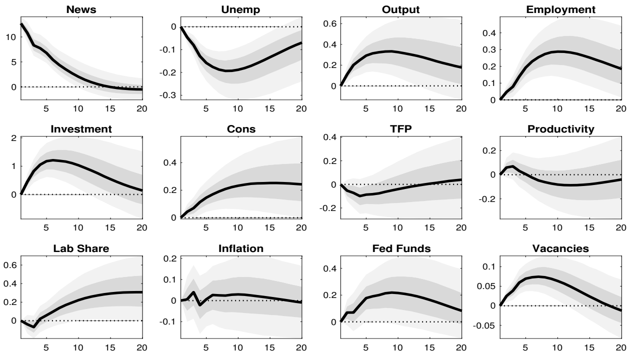

Figure B.21: Quarterly 12-variable specification, inverse Cholesky ordering.

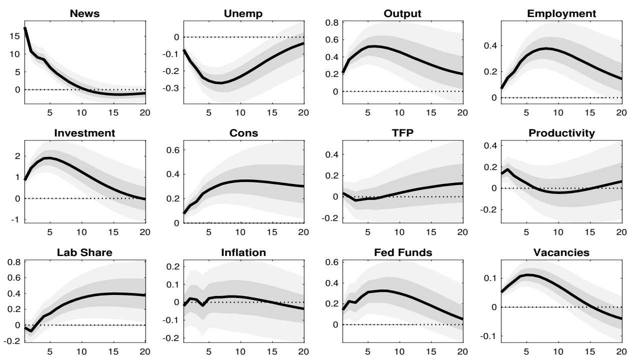

Figure B.22: Quarterly 12-variable specification, using only the news from the first month of the quarter.

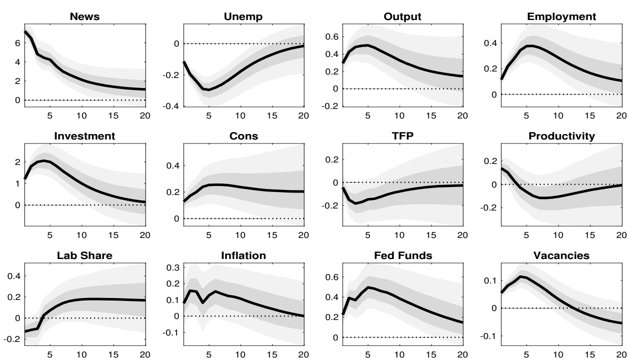

Figure B.23: Quarterly 12-variable specification, employment news only.

C Principal Component Analysis

Table C.1 reports the correlation between the eight the overall business conditions index, the eight sub-indices, and the eight principal components (PC1 through PC8). The principal components are listed in decreasing order of importance (i.e. share of variance explained) and are identified up to sign.

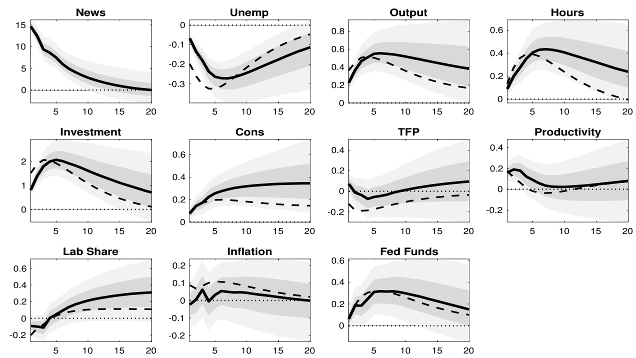

Figure B.24: VAR specification as in Angeletos, Collard, and Dellas (2020) with the addition of the business conditions index. The dashed lines are responses to the shock that maximizes the share of variance explained for unemployment.

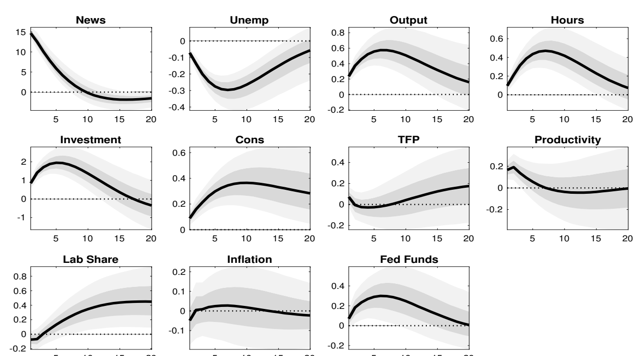

Figure B.25: Quarterly 12-variable specification including only 2 lags in line with Angeletos, Collard, and Dellas (2020).

Table C.1: Correlation coefficients between news indices and principal components. The sign of the principal components is not identified, so the correlation coefficients are identified up to sign as well.

	Overall	PriceNews	EmpNews	DemNews	GovNews	CredNews	StockNews	TradeNews	EnergyCrisisNews
Overall	1.000	-0.484	0.791	0.797	0.177	0.231	0.554	0.013	-0.178
PriceNews	-0.484	1.000	-0.139	-0.342	0.095	-0.390	-0.127	0.039	0.350
EmpNews	0.791	-0.139	1.000	0.668	-0.049	-0.126	0.326	-0.203	0.033
DemNews	0.797	-0.342	0.668	1.000	0.039	0.188	0.263	-0.063	-0.117
GovNews	0.177	0.095	-0.049	0.039	1.000	-0.039	0.002	0.169	-0.002
CredNews	0.231	-0.390	-0.126	0.188	-0.039	1.000	0.111	-0.072	-0.194
StockNews	0.554	-0.127	0.326	0.263	0.002	0.111	1.000	0.183	-0.007
TradeNews	0.013	0.039	-0.203	-0.063	0.169	-0.072	0.183	1.000	-0.000
EnergyCrisisNews	-0.178	0.350	0.033	-0.117	-0.002	-0.194	-0.007	-0.000	1.000
PC1	-0.816	0.171	-0.998	-0.695	0.055	0.108	-0.362	0.193	-0.020
PC2	-0.370	0.867	0.038	-0.298	0.295	-0.721	-0.228	0.034	0.344
PC3	-0.369	0.140	0.010	-0.193	-0.945	-0.191	-0.197	-0.212	0.096
PC4	-0.118	-0.242	0.025	0.087	0.116	-0.043	-0.870	-0.273	-0.114
PC5	-0.020	0.370	0.014	0.134	-0.018	0.645	-0.135	-0.202	0.098
PC6	0.111	0.056	-0.026	0.597	-0.043	-0.111	-0.008	0.293	-0.038
PC7	-0.043	0.001	-0.010	0.098	0.013	-0.051	0.051	-0.846	-0.022
PC8	-0.043	-0.021	-0.001	0.013	0.002	0.002	-0.001	-0.007	0.921

D Impulse-Response Matching

I complement the analysis in the main text reporting more information regarding the impulse-response matching exercise.

Table D.2 reports the prior and posterior information for the parameters estimated in my baseline setup, the impulse-responses for which are reported in the main text.

	prior mean	post. mean	10th	90th	prior	prior stdev
ρ_ρ	0.500	0.8011	0.7437	0.8461	beta	0.2000
ρ_F	0.500	0.2960	0.1151	0.4733	beta	0.2000
θ_ρ	0.500	0.0630	0.0444	0.0816	beta	0.2000
θ_F	0.500	0.5043	0.3779	0.6150	beta	0.2000
θ_x	0.500	0.2943	0.1256	0.4590	beta	0.2000
ρ_x	0.500	0.3062	0.1194	0.4961	beta	0.2000
σ_k	0.500	0.0308	0.0068	0.0530	beta	0.2000
$\overline{a}$	1.000	1.0948	0.7163	1.5290	gamm	0.2500
ψ_w	0.500	0.8743	0.7980	0.9528	beta	0.2000
ψ	120.000	143.3507	96.7798	182.8451	norm	30.0000
ζ	0.500	0.4615	0.1670	0.8214	beta	0.2000
S''	3.000	0.0992	0.0454	0.1554	gamm	2.0000
ϕ_u	0.010	0.0131	0.0044	0.0255	gamm	0.0050
ϕ	0.500	0.7114	0.1753	1.1689	gamm	0.3000
ρ_R	0.500	0.4171	0.2224	0.6030	beta	0.2000

Table D.2: Prior and posterior statistics. Baseline IRF-matching estimation. The columns represent the prior and posterior mean, the 10th and 90th percentile of the posterior draws, the prior distribution, and the prior standard deviation.

The Bayesian matching procedure attributes about a third of the weight to the component of the shock that will only directly affect the economy from period 2 onwards. Capitalists are estimated to be almost risk neutral, as often assumed in the literature. Wages are estimated to be sticky, which again is common. The price adjustment coefficient is not estimated to be too large (relative to a prior corresponding to a price duration of 4 quarters). The corresponding price duration, in a Calvo setting, would be of about 4.4 quarters. The investment-adjustment costs are estimated to be practically zero. The policy rule coefficient on inflation implies that policy rates would respond (in the long-run) 1.7 times the deviation of inflation from target, well within commonly estimated values. The coefficient on unemployment is very small because it scales *percent* variations in the unemployment rate. A .3 percentage point variation in unemployment (like that observed in my IRFs) would amount to a 7 percent variation (relative to the the 4.1 percent steady state value for unemployment). $\phi_u = .0131$ implies that, in response to that variation, rates would be increased by about 30 basis points (annualized).

Figure D.26 reports the IRFs estimated when I restrict the process x_t is not autoregressive, the investment-

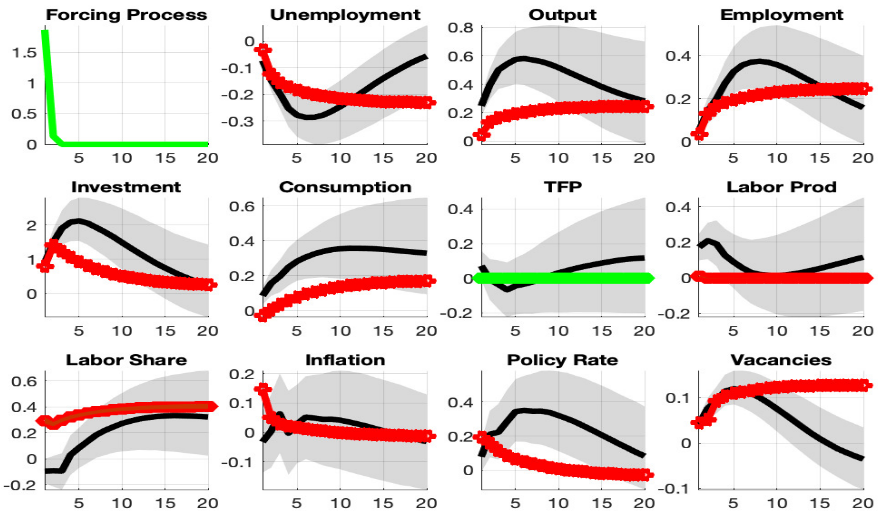

Figure D.26: Impulse response matching to the quarterly VAR impulse responses. R_t^{flex} is not included in the policy rule and only, $\rho^x = 0$, wages are fully flexible, and Phillips Curve Parameters are not estimated. Red lines represent impulse-responses from the model which are matched to the VAR responses, while green lines represent impulse responses that are left unrestricted.

adjustment costs to be identically zero, and I do not include R_t^{flex} in the policy rule. While the fit obviously deteriorates, the key comovement features remain the same. Table D.3 reports the parameter estimates. The main difference is observed for the response coefficient to inflation. Nominal rates are estimated to respond by about 2.5 percent for every 1 percent deviation of inflation from target.

	prior mean	post. mean	10th	90th	prior	prior stdev
ρ_ρ	0.500	0.1554	0.0408	0.2663	beta	0.2000
ρ_F	0.500	0.9960	0.9928	0.9996	beta	0.2000
θ_ρ	0.500	0.0632	0.0463	0.0772	beta	0.2000
θ_F	0.500	0.4101	0.3856	0.4324	beta	0.2000
θ_x	0.500	0.0671	0.0172	0.1172	beta	0.2000
$\sigma_{\mathbf{k}}$	0.500	0.2131	0.0661	0.4043	beta	0.2000
$\overline{a}$	1.000	0.9652	0.5541	1.2763	gamm	0.2500
ϕ_u	0.010	0.0102	0.0023	0.0176	gamm	0.0050
ϕ	0.500	1.5571	1.0628	2.1494	gamm	0.3000
ρ_R	0.500	0.4761	0.3113	0.6378	beta	0.2000

Table D.3: Prior and posterior statistics. Estimation with simpler policy rule and fixed Phillips Curve slope. The columns represent the prior and posterior mean, the 10th and 90th percentile of the posterior draws, the prior distribution, and the prior standard deviation.

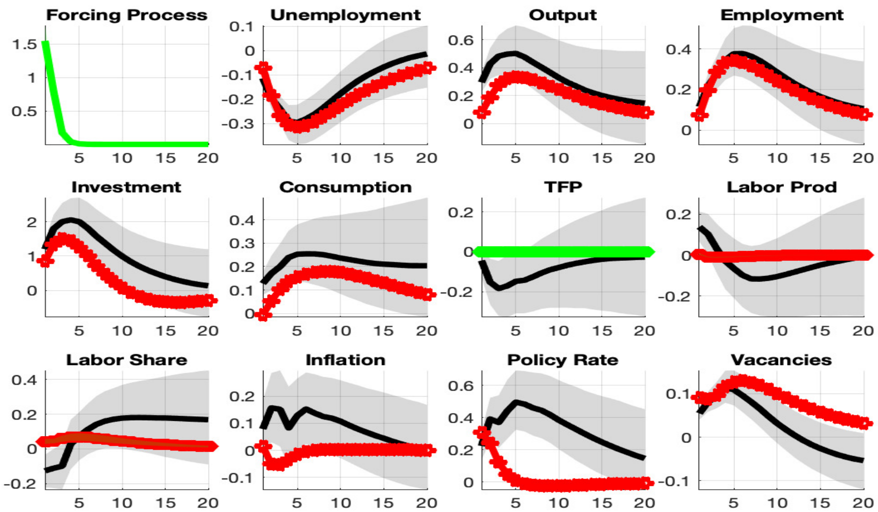

Figure D.27: Impulse response matching to the quarterly VAR impulse responses. The VAR includes only the labor-market component of the Business conditions index. Red lines represent impulse-responses from the model which are matched to the VAR responses, while green lines represent impulse responses that are left unrestricted.

Figure D.27 reports the matched IRFs when I only include labor-market news in the VAR. The fit is better for real variables and a bit worse for inflation. Table D.4 reports the parameter estimates. Posterior estimates are in line with those found above in the baseline case. However, while I maintain the prior mean the same, I somewhat reduce the prior standard deviation for some parameters. With the original prior standard deviations the mode finding step would work very well but the MCMC procedure would not.

Figure D.28 reports the IRFs in the case in which I replace the unemployment-gap in the monetary policy rule with the commonly-used output gap. Table D.5, reports the estimates. A coefficient on the output gap of .11 is in line with the literature. The fit is comparable to my baseline case.

Figures D.29, D.30, and D.31 report the IRF matching results when I only consider the shock to the cost of vacancy creation, when I only consider the shock to job destruction, and when, on top of these two shocks, I also allow for a shock to matching efficiency. The most important observation is that, as mentioned in the main text, both the shocks I use in my baseline setup can explain the comovement among the variables of interest even taken in isolation. This is important to reinforce the point that this model does not need a combination of shocks to reproduce the responses of the data to business-conditions.
Correspondingly, allowing for an extra shock, to matching efficiency, improves the quantitative fit of the model but does not change the general properties of the model.

Figure D.32 reports impulse responses to a TFP shock, from a model calibrated based on the posterior

	prior mean	post. mean	10th	90th	prior	prior stdev
ρ_ρ	0.500	0.6299	0.5503	0.6963	beta	0.2000
ρ_F	0.500	0.2273	0.0589	0.3539	beta	0.2000
θ_ρ	0.500	0.0913	0.0682	0.1108	beta	0.2000
θ_F	0.500	0.6540	0.5488	0.7829	beta	0.2000
θ_x	0.500	0.2166	0.0681	0.3413	beta	0.2000
ρ_x	0.500	0.2231	0.0739	0.3514	beta	0.2000
$\sigma_\mathbf{k}$	0.500	0.0310	0.0053	0.0538	beta	0.2000
$\overline{a}$	1.000	1.1330	0.7360	1.6214	gamm	0.2500
ψ_w	0.500	0.8810	0.8256	0.9309	beta	0.1000
ψ	120.000	127.7170	93.3148	156.6024	norm	20.0000
ζ	0.500	0.4857	0.3461	0.6390	beta	0.1000
S''	3.000	0.0441	0.0108	0.0751	gamm	2.0000
ϕ_u	0.010	0.0155	0.0030	0.0272	gamm	0.0050
ϕ	0.500	0.5862	0.4151	0.7732	gamm	0.1000
ρ_R	0.500	0.3877	0.2773	0.5086	beta	0.1000

Table D.4: Prior and posterior statistics. Labor news only, with slightly tighter priors. The columns represent the prior and posterior mean, the 10th and 90th percentile of the posterior draws, the prior distribution, and the prior standard deviation.

	prior mean	post. mean	10th	90th	prior	prior stdev
ρ_ρ	0.500	0.7777	0.6915	0.8617	beta	0.2000
ρ_F	0.500	0.2817	0.0466	0.5068	beta	0.2000
θ_ρ	0.500	0.0624	0.0442	0.0834	beta	0.2000
θ_F	0.500	0.4642	0.3627	0.5734	beta	0.2000
θ_x	0.500	0.2981	0.1337	0.4775	beta	0.2000
ρ_x	0.500	0.3414	0.1080	0.5973	beta	0.2000
$\sigma_\mathbf{k}$	0.500	0.0321	0.0065	0.0567	beta	0.2000
$\overline{a}$	1.000	1.0871	0.6039	1.4392	gamm	0.2500
ψ_w	0.500	0.8875	0.8159	0.9655	beta	0.2000
ψ	120.000	136.4337	92.4434	185.2885	norm	30.0000
ζ	0.500	0.5110	0.1831	0.8654	beta	0.2000
S''	3.000	0.1027	0.0129	0.1594	gamm	2.0000
ϕ_y	0.100	0.1125	0.0438	0.2065	gamm	0.0500
ϕ	0.500	0.8936	0.2473	1.4878	gamm	0.3000
ρ_R	0.500	0.4305	0.2249	0.6868	beta	0.2000

Table D.5: Prior and posterior statistics. Using the output gap as opposed to the unemployment gap in the policy rule.

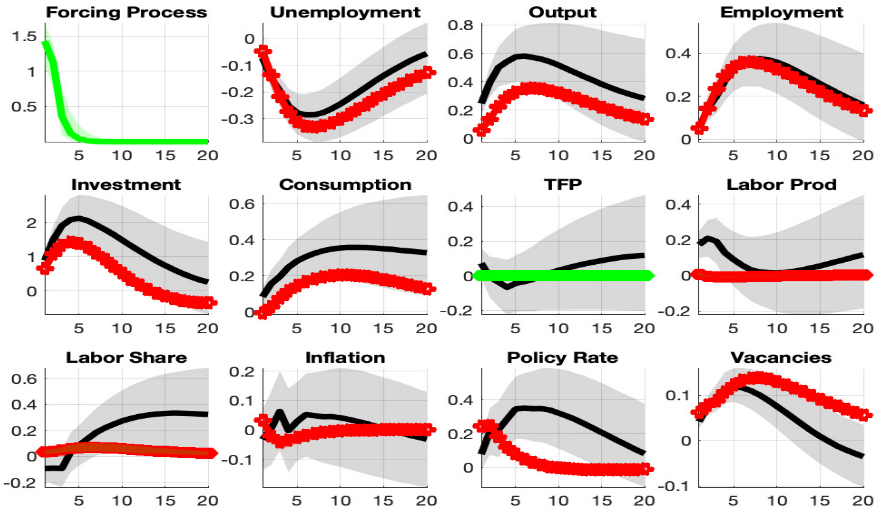

Figure D.28: Impulse response matching to the quarterly VAR impulse responses. In this version of the model, the unemployment gap is replace by the output gap in the policy rule. Red lines represent impulse-responses from the model which are matched to the VAR responses, while green lines represent impulse responses that are left unrestricted.

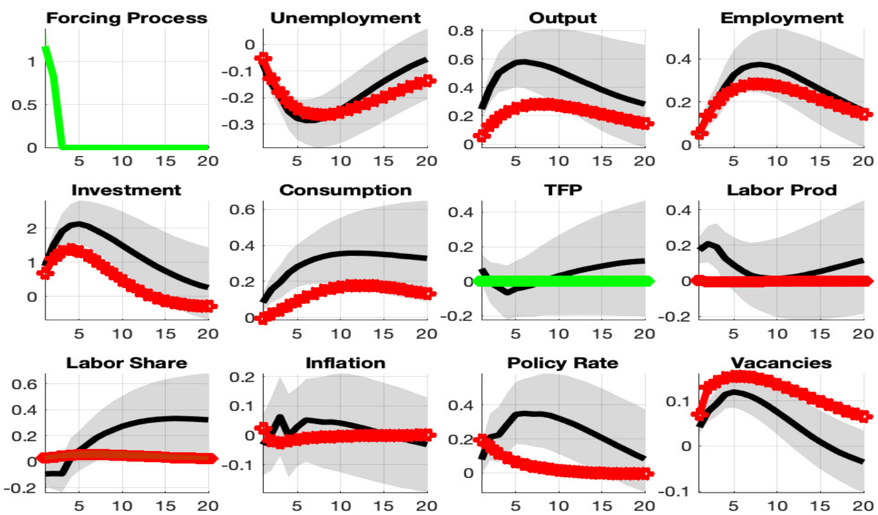

Figure D.29: Impulse response matching to the quarterly VAR impulse responses, using only shocks to the cost of opening new vacancies. Red lines represent impulse-responses from the model which are matched to the VAR responses, while green lines represent impulse responses that are left unrestricted.

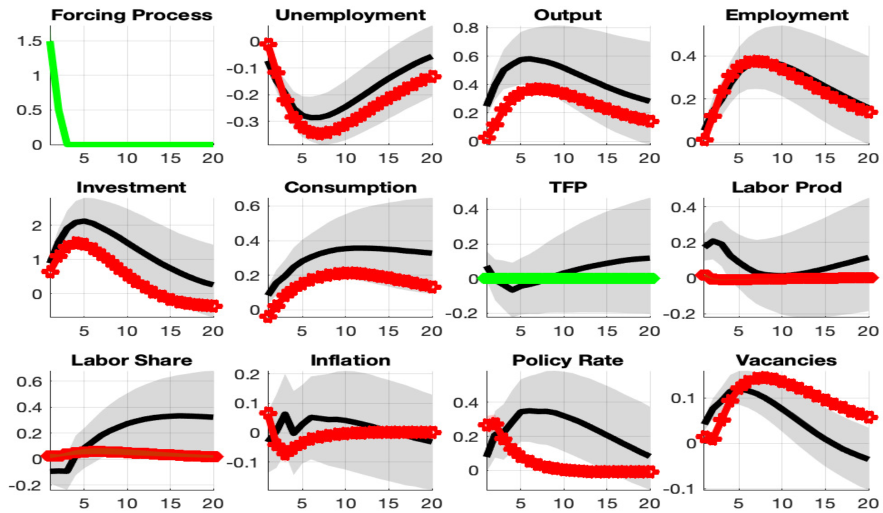

Figure D.30: Impulse response matching to the quarterly VAR impulse responses, using only shocks to the job-separation rate. Red lines represent impulse-responses from the model which are matched to the VAR responses, while green lines represent impulse responses that are left unrestricted.

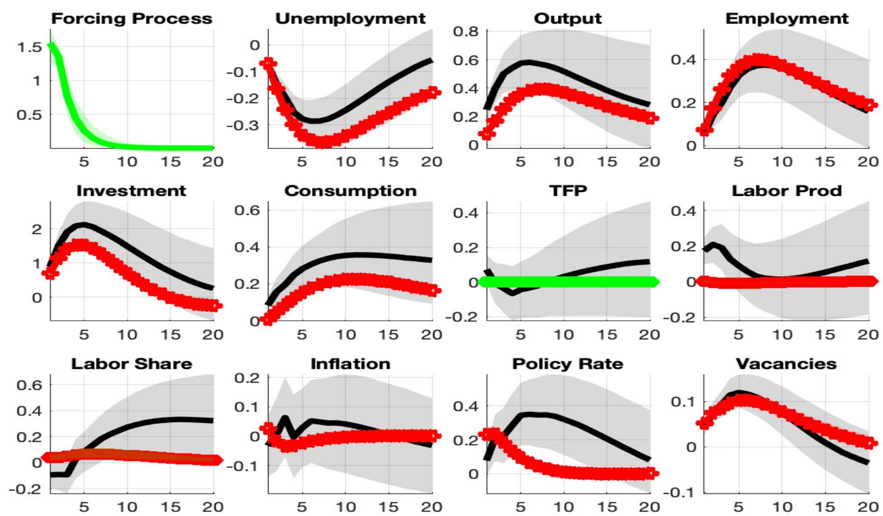

Figure D.31: Impulse response matching to the quarterly VAR impulse responses, in which I also allow for shocks to matching efficiency. Red lines represent impulse-responses from the model which are matched to the VAR responses, while green lines represent impulse responses that are left unrestricted.

mean estimates for my baseline matching exercise. The size of the shock is calibrated so that the peak response of unemployment is comparable to that in the VAR. The responses of most key macro variables

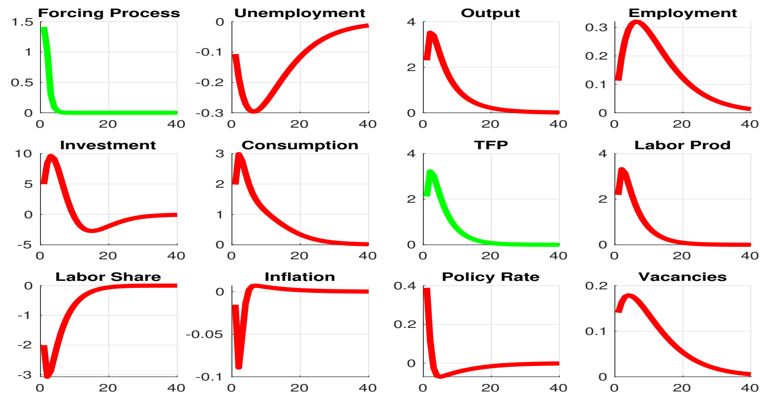

Figure D.32: I take the parameter estimates from by baseline IRF matching exercise. I then calibrate the size of a TFP shock (inclusive of news with the same process as x) to match the peak response of unemployment of about negative .3 percent - while calibrating the persistence of the process to .8 like I estimated for the job-destruction rate. The responses are reported here.

is in line, sign-wise, with the empirical evidence. However, the size of the responses of aggregate demand components is about 5 times as large as that observed in the VAR. Finally the labor share responds strongly and negatively, at odds with observed evidence.