Anticipated Productivity and the Labor Market*

Ryan Chahrour† Boston College

Sanjay K. Chugh‡ The Ohio State University

Tristan Potter§ Drexel University

December 18, 2019

Abstract

We identify the main shock driving the covariance of the labor market and output. The shock drives strong business cycle comovement among output, consumption, investment, hours, and stock prices but is essentially orthogonal to business cycle fluctuations in TFP. Yet, the shock is associated with future persistent TFP fluctuations, consistent with theories of technology news. A standard labor search model in which wages are determined by a cash flow sharing rule, rather than the net present value of match surplus, matches the observed responses to TFP news. The response of the wage implied by this rule is consistent with the empirical responses of a broad panel of wage series.

Keywords: News Shocks, Wages, Search and Matching, Business Cycles

JEL Classification: E32, E24

*We thank Susanto Basu, Fabio Schiantarelli, John Seater, Rosen Valchev, and seminar participants at Boston College, Boston University, CREI, the Society for Economic Dynamics Meetings (2017), the Econometric Society North American Summer meetings (2017), the Toulouse School of Economics, and Simon Fraser University for many helpful comments and suggestions.

†Department of Economics, Boston College, Chestnut Hill, MA 02467. Telephone: 617-552-0902. Email: ryan.chahrour@bc.edu.

‡Department of Economics, The Ohio State University, Columbus, OH 43210. Telephone: 614-247-8718. Email address: chugh.14@osu.edu.

§*Corresponding author.* School of Economics, LeBow College of Business, Drexel University, Philadelphia, PA 19103. Telephone: 215-895-2540. Email: tristan.l.potter@drexel.edu.

1 Introduction

Over the last two decades, macroeconomists have become increasingly skeptical that technology could be the primary driver of business cycle fluctuations, particularly in the labor market. From a theoretical perspective, this skepticism is rooted in the critique of Shimer (2005) that theories of flexibly-bargained wages cannot give rise to fluctuations in firm vacancy postings and therefore in employment. The alternative of fixed or extremely sticky wages, proposed by Hall (2005) and explored quantitatively by Gertler et al. (2008), could resolve the conundrum, save for mounting evidence of substantial wage flexibility in practice (Haefke et al., 2013; Kudlyak, 2014; Basu and House, 2016). Proponents of technology shocks have also faced the rather compelling counterpoint that fluctuations in the labor market do not seem to coincide with contemporaneous measures of productivity (Angeletos et al., 2019).

In this paper, we revisit the arguments against TFP as a primary driver of labor market fluctuations and conclude that the evidence against TFP may be less conclusive than initially appears. We begin our paper with a simple empirical VAR exercise, in which we identify in an agnostic way the "main" shock driving joint fluctuations in output and hours. The shock we uncover drives roughly half of business-cycle frequency fluctuations in standard macroeconomic aggregates, including hours, but only small and statistically insignificant fluctuations in contemporaneous TFP. Yet, at horizons beyond the typical business cycle, we find the shock is associated with a strong and extremely persistent increase in productivity: our identified shock strongly resembles a classic "news" shock.

We then ask if a standard search and matching model of the labor market can match the observed responses of business cycle quantities to a TFP news shock of the type we observe in the data. To this end, we perform an impulse response matching estimation exercise matching a theoretical model to our estimated impulse responses. Our model for this exercise is a real business cycle model with a matching friction in the labor market and two standard extensions: capital adjustment costs and consumption habit formation.

The ability of labor matching models to generate fluctuations in the labor market depends heavily on how wages are determined. Given the considerable disagreement about how to best model wage setting, and the open debate about the best empirical measures of wages, we perform our initial impulse response matching exercise without imposing any structure on the wage ex ante. Our exercise thus chooses the process for the wage that is most consistent with the observed responses of quantities.

Our estimation exercise delivers two main results. First, we find that the model — which has very few free parameters apart from the flexible wage process — does an excellent job at matching the impulse responses we find in the data. Second, we find that the implied wage process follows a distinctive pattern: Wages fall modestly during the anticipation period ahead of the TFP shock, and then rise quickly when the shock is realized. The estimated wage process is thus inconsistent with a model of extremely sticky real wages, but also hard to align with a model of constant-share Nash bargaining, which cannot cause wages to fall significantly in response to higher expected future product of labor.

A natural question then arises: What sort of wage determination mechanism would be consistent with our estimated "agnostic" wage process? It turns out that our estimated wage process is consistent with a model in which wages are driven primarily by current cash flows, rather than the net present value of match surplus. We thus propose a simple model of wage determination according to which workers receive a pro rata share of firms' available cash flow after accounting for payments to capital and the costs of hiring. This model of wage setting closely resembles the model studied by den Haan and Kaltenbrunner (2009), and entails only a single free parameter. We re-estimate our model using the flow-based wage determination mechanism and show that the model fit, the model-implied impulse responses, and the implied wage are all virtually identical to the results from the fully agnostic wage specification we originally estimated.

Our model of wage determination has two key elements. First, the wage splits current-period cash flows, rather than the present discounted value of match surplus as in Nash bargaining. This feature is essential for matching the large observed response of employment and output to news about the future before the shock is actually realized: In our model, good news about the future stimulates hiring today via the frictional matching process, which in turn increases employment, reduces labor's marginal product and cash flows per worker, and so reduces wages. When the shock is finally realized and labor becomes more productive, the wage rises in response to the increased revenue flows associated with higher productivity.

Because it is based on a present value calculation, a Nash bargained wage could never support a similar expectations-driven boom. For, any potential boom in employment and consumption today would lower future consumption growth, raising the present value of future cash flows and hence the Nash bargained wage itself. This negative feedback precludes a model with simple Nash bargaining from generating a boom in output and employment ahead of the realization of the shock.

Second, for our model of wage determination to match the data, the fraction of flow surplus accruing to households must be relatively high. This feature is closely related to the observation of Hagedorn and Manovskii (2008) that when firms receive a small fraction of flow surplus, small changes in productivity translate to large (in percentage terms) changes in flow profits and thus have an outsize effect on vacancy posting incentives. This effect is capable of generating large booms in response to *anticipated* changes in productivity because matching frictions pull forward the benefits of hiring, but only when those benefits are not offset by a forward-looking wage process such as Nash bargaining.

We conclude our main results by showing that the wage process we estimate is consistent with a variety of existing measures of the aggregate wage. To do this, we consider a panel of 19 commonly-used wage measures collected from various sources. Our first (and preferred) measure of the wage is aggregate wage and salary payments to labor in the private sector (compiled by the BEA) divided by total private sector hours worked. The response of this wage to our identified shock is in fact very similar to what our model predicts: The wage falls on impact and then eventually rises following TFP.

In addition to this series, we present a set of aggregate and sector-level wage series prepared by the BLS, and the new-hire wage series generated by Basu and House (2017). The responses of these variables to our identified shock differ substantially, but two patterns emerge. First, of the 19 series, all but 4 fall on impact according to our point estimates, and none is significantly positive. Second, virtually all of the wage series exhibit upward-sloping patterns in the period after the identified shock. In these respects, our panel of wage data is quite consistent with the wage process we estimate; indeed, our estimated wage process lies within the range of estimated responses in the panel for at least 10 years after the shock.

Our results are robust to a wide range of specifications of the empirical VAR, including different lag lengths, VECM estimation with one or more trends, and including additional variables in our VAR. We also show that the main features of the shock could also be recovered via an identification procedure that seeks to explain the maximal variation in future forecast revisions of TFP. In this respect, our agnostic identification procedure recovers the same shock that is recovered by a common approach that has been specifically designed to isolate news shocks.

Our results lead us to the conclusion that news about technology could well play an important role in driving the business cycle, including for the labor market. This result contrasts with some recent findings in the literature. In particular, though our methodology

is similar to Angeletos et al. (2019), they find that TFP cannot be the "main business cycle" shock. The crucial difference between our respective approaches is that, in their identification procedure, Angeletos et al. (2019) specifically target narrow portions of the spectrum, focusing on business cycle fluctuations between 6 and 32 quarters in the frequency domain, while we consider fluctuations at horizons of up to 500 quarters.[1]

Our empirical results are also related to a long literature seeking to identify news shocks in VARs, notably Barsky and Sims (2011) and more recently Kurmann and Sims (2017) and Bouakez et al. (2019). In fact, our main empirical results are recovered if we employ the indentation approach of Kurmann and Sims (2017), who seeks to identify the shock that best explains forecast revision in TFP in the distant future. Nevertheless, our empirical exercises give somewhat different results and our focus on the cyclicality of real wages and theories of labor market search is quite different from theirs.[2]

Recently, Faccini and Melosi (2019) have estimated a structural labor search model with sticky wages, and also find that expectations shocks play a crucial role in driving the labor market. Our semi-structural empirical approach reinforces these findings, and allows us to easily incorporate additional evidence on how wages respond to news shocks.

From a theoretical perspective, our paper is most related to den Haan and Kaltenbrunner (2009), which motivates our choice of a structural wage-setting mechanism. That paper was among the first to demonstrate that news shocks can, in principle, drive an immediate expansion in employment. We build on that paper by providing new empirical evidence in support of news shocks and showing that a model with capital adjustment costs and habit formation can quantitatively match the empirical responses of macroeconomic aggregates generated by such shocks, particularly measured investment. Theodoridis and Zanetti (2016) consider a search and matching model with Nash bargaining and several shocks, including news about TFP. While they find that news shocks are important for explaining consumption and investment dynamics, their model requires both job destruction shocks and shocks to the matching function to account for labor market dynamics. We provide quantitative evidence that news shocks alone can provide a compelling account of business cycles—including labor

[1] Our results do not depend on going this far out and are essentially identical so long as we consider frequencies corresponding to periods of at least 100 quarters. Narrowing the range beyond this gradually changes our results to look more like those in Angeletos et al. (2019).

[2] When we use the shock identified by Kurmann and Sims (2017) on our sample period (i.e. through 2018Q4), hours rise. However, when we instead use their sample period (i.e. through 2007Q3), hours do not always rise on impact. Our approach consistently gives a positive impact for hours across sample periods and specifications.

markets—with the right wage-setting mechanism.

Finally, the paper is related to Christiano et al. (2016), who also do an impulse response matching exercise with a labor search model. However, they do not consider the possibility of news shocks, which appear to be crucial in our data. Hall (2017) has argued that the data support a strong connection between stock market valuation and labor markets, a finding which our empirical and theoretical exercise supports.

This paper proceeds as follows. In Section 2, we describe our main empirical exercise aimed at identifying the shock that drives the covariance of output and hours. Section 3 describes a theoretical labor search model with an anticipated productivity shock. Section 4 estimates the parameters of the model needed to match our estimated impulse responses and discusses the implications of the estimation exercise for a plausible model of wage determination, as well as the relation with empirical measures of the wage. Section 5 argues that a flow-based surplus sharing model of the wage can fit the data well. Section 6 concludes.

2 Empirical Exercise

Our baseline empirical specification consists of a vector-autoregression of the form

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

where Y_t is a vector of observed variables, $B(L)$ contains the weights on past realizations of Y_t, ϵ_t is a vector of structural economic shocks, and A is the structural matrix that our procedure seeks to identify from the set of estimated residuals, $\mu_t \equiv A\epsilon_t$.

We take as our baseline set of variables $Y_t \equiv [TFP_t, GDP_t, C_t, H_t, SP_t]'$, which includes utilization-adjusted TFP from Fernald (2014), real per-capita GDP, real per-capita consumption, real per-capita investment, per-capita hours, and the real stock price. We estimate the VAR in levels via OLS and include four lags in the polynomial $B(L)$. Our sample ranges from 1966Q1 to 2018Q4. Additional details on data construction are provided in Appendix C.

We also consider a set of auxiliary variables, W_t, that includes 19 measures of the hourly wage drawn from several sources. The wage series are related to current and past observations of Y_t according to

$$W_t = \Gamma(L)Y_t + v_t \tag{2}$$

where the coefficient matrix $\Gamma(L)$ includes the same number of lags (four in our baseline) as the VAR in (1) and is estimated via OLS. We can thus construct impulse responses for any auxiliary wage measure in W_t using the responses of the variables Y_t and the estimated values of $\Gamma(L)$.

2.1 Identification Approach

We employ an approach to identifying the matrix A in the family of "max-share" approaches first introduced by Uhlig (2003). These approaches identify the shock which explains the largest portion of some covariance matrix implied by the model in (1), and have been used widely in the structural VAR literature.[3]

Like Angeletos et al. (2019), we target a moment extracted from the frequency domain, but our target is different in two respects. First, we target the covariance matrix of output and hours, rather than targeting a single variable (e.g. output only) at a time. In this respect, we are explicitly directing our procedure to identify the shock that drives comovement between output and the labor market. Second, we target a wider band of frequencies than do Angeletos et al. (2019), who attempt to isolate the sources of business cycle comovements from the sources of comovement at longer frequencies.

Specifically, define

$$\phi(z) \equiv (I - B(z))^{-1} A \tag{3}$$

as the z transfer-function associated with the MA-infinity representation of equation (1). Further, let s be a matrix selecting the target variables of interest. In our baseline case, $s = [e_2, e_5]'$, where e_i is the ith column basis vector. The covariance associated with spectra of periodicity $p \equiv [p_1, p_2]$ is given by

$$\Sigma_p^s \equiv \frac{1}{2\pi} \int_{2\pi/p_2}^{2\pi/p_1} \left[s\phi(e^{-i\lambda})\right] \left[s\phi(e^{i\lambda})\right]' d\lambda. \tag{4}$$

Conversely, the contribution of each shock to the variance in the same range is given by

$$\Omega_p^s \equiv \frac{1}{2\pi} \int_{2\pi/p_2}^{2\pi/p_1} \left[s\phi(e^{-i\lambda})\right]' \left[s\phi(e^{i\lambda})\right] d\lambda. \tag{5}$$

We can then find the shock that explains the most of Σ_p^s by computing q_1, the eigenvector

[3] See Francis et al. (2014), who build on Uhlig (2003, 2004) and Faust (1998).

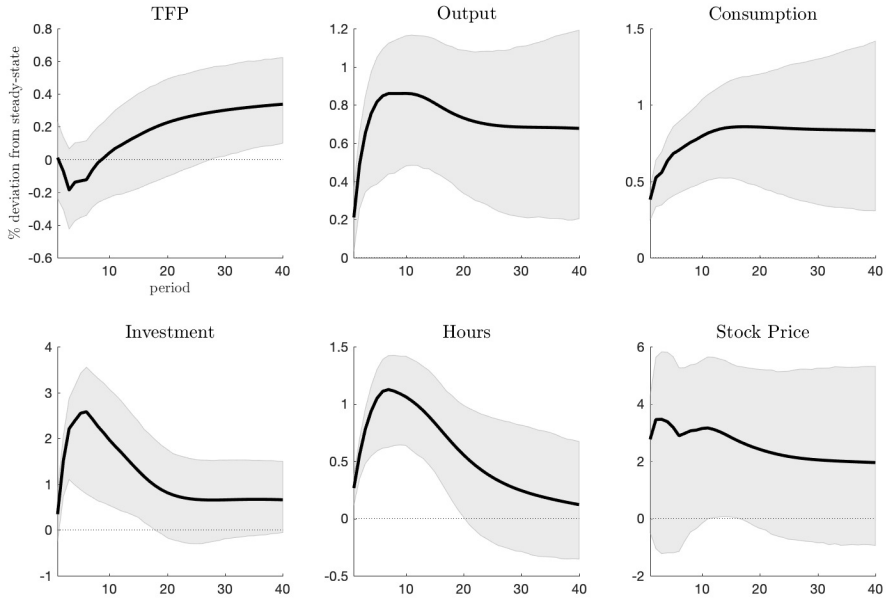

Figure 1: Impulse responses to identified shock.

associated with the largest eigenvalue of Ω_p^s and setting

$$A = \hat{A} q_1, \tag{6}$$

where $\hat{A}$ is the Cholesky decomposition of the matrix $\Sigma_u \equiv \text{cov}(\mu_t)$.

In contrast to Angeletos et al. (2019), we consider a wider range of spectra, with periodicity $p = [6, 500]$ quarters, because we do not want to impose an *ex ante* separation between the shocks that drive the business cycle and those that drive longer-run fluctuations. In practice, this is little different than considering unconditional covariances, but it has the advantages of (i) excluding extremely short-range fluctuations that might be associated with, e.g., measurement error, and (ii) remaining feasible even if the estimated process has a unit root and unconditional variances are not defined, as would occur should we estimated (1) as a VECM.

2.2 Results

Figure 1 presents the impulse responses to our identified shock selected to explain the covariances of output and hours, along with 80% confidence bands from a bias-corrected bootstrap. The Figure shows that our shock drives large and significant immediate fluctuations in out-

Table 1: Variance decompositions of VAR variables

Frequency (Quarters)	**TFP**	**Y**	**C**	**I**	**N**	**S&P**
Business Cycle (6-32)	0.149	0.433	0.658	0.384	0.539	0.166
Medium run (32-100)	0.148	0.647	0.784	0.488	0.644	0.129
Long run (100-500)	0.710	0.803	0.802	0.758	0.735	0.396

put, consumption, investment and employment, as well as a substantial though marginally significant response in stock prices. Moreover, while the responses of most of these variables are larger in the short run, they are extremely persistent, with both output and consumption significantly positive ten years after the shock.

Central for our insights in this paper is the response of utilization-adjusted TFP, our preferred measure of production technology. On impact of the shock, TFP is unchanged and then falls modestly (and insignificantly) for several quarters, before it begins a gradual rise that becomes statistically significant after roughly seven years. This pattern of productivity is consistent with the idea of a "news" shock and it is this interpretation that we explore in the following sections.

Table 1 presents the variance decomposition for our shock across three portions of the spectrum, corresponding to business cycle frequencies (6-32 quarters), medium run (32-100 quarters) and long run (frequencies greater than 100 quarters). The table shows that all of the quantity variables are substantially explained by the identified shock, with the contribution of the shock rising to well over 50% at longer horizons for variables other than the stock market.

Crucially (and consistent with the finding of Angeletos et al., 2019), we find that TFP fluctuations are essentially orthogonal to the effects of this shock at business cycle and even at medium run frequencies. It is only at periodicities of over 100 periods that the strong connection between our shock and productivity appears. These results are precisely consistent with the idea that expectations about very long-run productivity are playing a central role in driving fluctuations at shorter horizons, echoing the theories and structural estimation results of Blanchard et al. (2013) and Chahrour and Jurado (2018).

To understand the effect that the shock has on wages, we produce impulse responses for a number of empirical wage measures to our identified shock. These wage responses are displayed in Figure 2. Our preferred measure of the aggregate wage (aggregate wage and salary payments to labor in the private sector divided by total private sector hours worked)

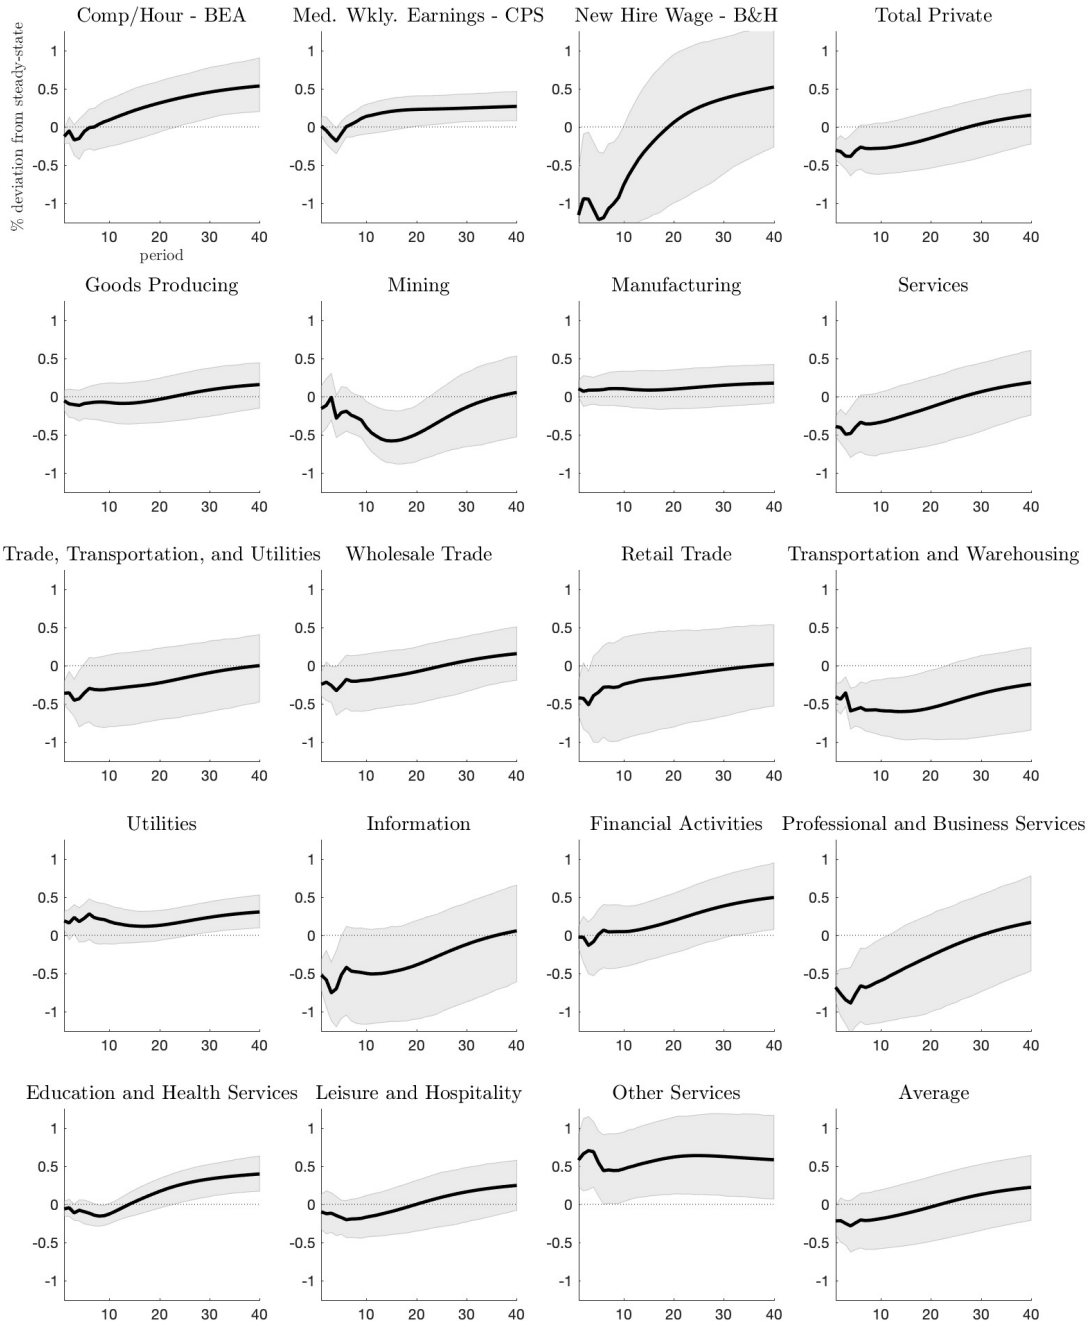

Figure 2: Wages responses to labor market shock.

displays two distinctive features. First, the wage falls modestly on impact. Second, the wage grows quickly as TFP begins to rise. The responses of the other wage series exhibit considerable heterogeneity but generally reflect similar patters: At their point estimates, all but four fall on impact, none is significantly positive, and nearly all of the series appear to grow over the horizon of the response.

Apparently, the data suggest that a productivity news shock could potentially play a central role in driving business cycle fluctuations, as well as longer-term changes in macroeconomic aggregates. Standard models, however, will have trouble rationalizing these patterns. Real models with flexible prices will not generally be able to produce an expansion during the anticipation period. By contrast, new-Keynesian models with sticky prices can explain the expansion during the anticipation period, but will generally lead TFP improvements to be contractionary for labor when they are realized (Basu et al., 2006). In the next sections, we explore the ability of a real model with a non-standard specification of the wage to account for our empirical results.

2.2.1 Relationship to News Shock procedure

One advantage of our approach to identifying a shock using the comovement of output and hours is that we do not commit *ex ante* to any particular interpretation of the shock. Given our finding that our identified shock closely resembles a news shock, however, it is natural to ask if we would have found the same impulse responses had we used a more standard approach to identifying news.

The answer, it turns out, is yes. In Appendix D, we present impulse responses for our baseline shock and for the long-horizon identification procedure used by Kurmann and Sims (2017). As the figure in the appendix shows, we find impulse responses that are extremely similar to the news responses that their alternative procedure identifies. We read this evidence as corroborating our interpretation that the shock we identify is, in fact, a news shock.

3 Model

The economy consists of a representative household and a representative firm who each trade in markets for consumption, labor and capital. Consumption and capital markets are competitive, while transactions in labor markets are subject to search and matching frictions

in the spirit of Mortensen and Pissarides (1994).

3.1 Households

The representative household consists of a continuum of *ex ante* identical members who are either employed or searching for work. The household derives utility at time t from consumption according to the period utility function $u(c_t; C_{t-1})$, where c_t is household consumption and C_{t-1} is lagged aggregate consumption capturing the habit stock in the economy.[4] Each period, non-employed workers search for a match in the labor market. Searching members match with probability p_t. Moreover, newly-created matches become productive within the period, so that unemployment is given by the measure of searchers who failed to match

$$1 - n_t = (1 - p_t)\, s_t, \tag{7}$$

where n_t denotes the measure of currently matched workers and $(1 - p_t)\, s_t$ denotes the measure of searchers who failed to find a match in period t.[5] Each period, previously productive matches dissolve with exogenous probability λ, so that employment evolves according to

$$n_t = (1 - \lambda) n_{t-1} + p_t s_t. \tag{8}$$

In addition to choosing its consumption, the household also chooses a level of investment subject to a capital adjustment cost. The law of motion for the stock of capital is given by

$$k_{t+1} = (1 - \delta) k_t + i_t, \tag{9}$$

where δ is the depreciation rate of the capital stock.

The household budget constraint is given by

$$c_t + i_t + \tau_t = R_t k_t + W_t n_t + (1 - p_t) s_t \kappa_t + d_t - k_t \Phi_k \left(\frac{i_t}{k_t} \right). \tag{10}$$

The household takes the rental rate of capital, the wage rate of labor, and benefits paid to unemployed workers (R_t, W_t and κ_t respectively), as given. It also receives d_t, lump-sum dividends from firms, and pays τ_t, a lump-sum tax used to finance any exogenous stream

[4] We assume external habits as is common in the literature, and suppress dependence of $u()$ on C_{t-1}.

[5] This timing convention is consistent with the evidence on labor market flows at quarterly frequency. See Davis et al. (2006).

of government expenditures and unemployment benefits. The benefit paid to unemployed workers is assumed to be a fixed fraction of the current wage rate, $\kappa_t = \kappa W_t$.

The representative household's problem may thus be expressed as

$$\max_{c_t, i_t, u_t, n_t, k_{t+1}} E_0 \sum_{t=0}^{\infty} \beta^t u(c_t) \qquad \text{s.t. (9) and (10).} \tag{11}$$

The first-order conditions for investment, i_t, and capital next period, k_{t+1}, are given by

$$\mu_t^K = u_{c,t}\left[1 + \Phi_k'\left(\frac{i_t}{k_t}\right)\right] \tag{12}$$

$$\mu_t^K = E_t\left\{\beta\left[(1-\delta)\mu_{t+1}^K + u_{c,t+1}\left[R_{t+1} - \Phi_k\left(\frac{i_{t+1}}{k_{t+1}}\right) + \frac{i_{t+1}}{k_{t+1}}\Phi_k'\left(\frac{i_{t+1}}{k_{t+1}}\right)\right]\right]\right\} \tag{13}$$

where μ_t^K denotes the Lagrange multiplier on the law of motion for capital in (9).

3.2 Firms

The representative firm chooses labor, capital and vacancy postings to maximize the present value of real dividends, discounted according to the consumer's stochastic discount factor. The firm produces output with a production function of the form

$$y_t = F(k_t, X_t n_t), \tag{14}$$

where X_t is a non-stationary labor-augmenting technology shock.

Our main shock is a news shock about future X_t. Define the growth rage of productivity $\gamma_{x,t} \equiv X_t/X_{t-1}$, and the long-run growth rate γ_x. We assume that productivity growth follows an AR(1) process with news,

$$\log(\gamma_{x,t}/\gamma_x) = \rho_x \log(\gamma_{x,t-1}/\gamma_x) + \epsilon_{x,t-h}. \tag{15}$$

In equation (15), the shock $\epsilon_{x,t-h}$ first influences productivity at time t but is observed by agents at time $t-h$. We refer to h as the *time horizon* of the news shock.

The law of motion of employed labor from the firm's perspective is given by

$$n_t = (1-\lambda)n_{t-1} + q_t v_t \tag{16}$$

where v_t denotes vacancies posted in the labor market and q_t denotes the probability of a vacancy returning a match. The firm's profit maximization problem is thus

$$\max_{v_t, n_t, k_t} E_0 \sum_{t=0}^{\infty} \beta^t u_{c,t} \left[y_t - W_t n_t - R_t k_t - a_n v_t \right] \qquad \text{s.t. (14) and (16)} \tag{17}$$

where a_n is the cost of posting a vacancy. The first-order condition for capital is given by

$$F_{k,t} = R_t. \tag{18}$$

The first-order conditions for vacancies and employment, respectively, are

$$\mu_t^N = \frac{a_n}{q_t} \tag{19}$$

$$\mu_t^N = F_{n,t} - W_t + (1-\lambda) E_t \left\{ \Omega_{t,t+1} \mu_{t+1}^N \right\}. \tag{20}$$

3.3 Government

The government runs a balanced budget, financing an exogenous stream of aggregate purchases G_t through lump-sum taxes τ_t net of unemployment benefit transfers $(1-p_t)s_t\kappa_t$:

$$G_t = \tau_t - (1-p_t) s_t \kappa_t. \tag{21}$$

To maintain balanced growth, we follow Schmitt-Grohé and Uribe (2012) in assuming that government spending gradually adjusts to restore the long-run share of government spending in the economy,

$$G_t = G_{t-1} \left(\frac{G_{t-1}}{\bar{g}} \right)^{\phi_g}. \tag{22}$$

3.4 Wages

In search and matching models such as the one described above, the presence of matching frictions gives rise to positive match surplus that is split by the wage. Any wage yielding weakly positive surplus for the firm and the worker is consistent with equilibrium. The basic theory thus provides little guidance on how to model wage setting. Furthermore, there is considerable disagreement about the best empirical measure of wages, making it difficult to elicit direct empirical guidance regarding the best model of wage setting.

Table 2: Calibrated Parameters

Parameter	Concept	Value
β	Discount factor	0.990
σ	Risk Aversion	1.500
κ	Replacement rate	0.600
α	Capital share	0.320
δ	Depreciation rate	0.030
λ	Separation rate	0.080
ρ_x	TFP growth (persistence)	0.900
σ_x	TFP growth (SD)	0.050
a_n	Vacancy posting cost	0.301
ϕ_x	Wage error-correction	0.050
γ_x	long-run growth	1.004

We therefore specify for our baseline an "agnostic" wage, which places essentially no *a priori* structure on how wages can respond to shocks. In particular, we model wage growth as an MA(H) process augmented with an error-correcting term designed to ensure that the wage eventually returns to its long run level. Specifically, we assume

$$\Delta w_t = \gamma(L)\epsilon_t - \phi_x(w_{t-1} - x_{t-1}), \tag{23}$$

where $w_t \equiv \log(W_t)$, $x_t \equiv \log(X_t)$, and $\Delta w_t \equiv \log(w_t) - \log(w_{t-1})$.

Accordingly, our wage process admits $H + 2$ free parameters ($H + 1$ associated with the polynomial $\gamma(L)$ plus ϕ_x). In Section 5, we propose a cash flow-based structural description of wages that contains only a single free parameter, and explore how well it can reproduce our agnostic estimates of the wage process described above.

4 Estimation and Results

We calibrate a large set of parameters, since most of the structural (non-wage) parameters in our simple model are naturally pinned down by long-run averages in the data. Our calibration choices are summarized in Table 2.

We select the discount factor $\beta = 1.04^{1/4}$ to be consistent with an annual real interest rate of 4%. The value $\sigma = 1.5$ corresponds to a standard calibration of the intertemporal elasticity. The replacement rate of unemployment benefits is set to $\kappa = 0.6$, in the middle

of the range typically used in the search and matching literature. We fix capital's share of output and the capital depreciation rate to standard values of $\alpha = 0.32$ and $\delta = 0.03$, respectively. The quarterly job separation rate is set to 0.08, consistent with a summary of the evidence in Yashiv (2008).

We fix the size and persistence of the productivity news shock, as well as the horizon of its arrival, to be consistent with our point estimates of the TFP response to our identified shock. As depicted in Figure 1, there is a gradual build up of productivity after agents learn of the change. The implied values are $\sigma_x = 0.05$, $\rho_x = 0.90$, and $h = 9$.

As a final parameter, we need to fix a_n, the cost of vacancy postings. The choice of this parameter is important, as it implicitly (in conjunction with our other targets) pins down the share of output that goes to workers' wages. We follow Fujita and Ramey (2012), who draw on survey evidence on employer recruitment behavior cited in Barron et al. (1997) and Barron and Bishop (1985) to arrive at an estimate that vacancy posting costs constitute 17% of the marginal product of labor. This corresponds to a value of 0.301 for a_n, which implies that 1.2% of gross firm revenues accrue to firms as surplus from labor relationships.

4.1 Estimation Procedure

We estimate our model parameters using via a standard impulse response matching exercise, where the target impulse responses correspond to the responses for all six of the variables in our baseline VAR (that is, Y_t) for horizons of up to 40 periods.[6] Because we aim to match 40 periods, we fix the horizon of the MA terms in the wage process at $H = 40$.

Let $\hat{\psi}$ denote the column vector stacking our point estimates of each of these impulse responses. Then our target objective function corresponds to

$$\mathcal{L}(\theta) = (\hat{\psi} - \psi(\theta))' W (\hat{\psi} - \psi(\theta)) \tag{24}$$

where $\theta \equiv \{\theta, \phi, \epsilon, \gamma_0, \gamma_1, ...\gamma_H\}$ is the vector of parameters we seek to estimate and W is a diagonal matrix consisting of the inverse of the bootstrapped variances of each entry in $\hat{\psi}$.[7]

[6] Because we calibrate the parameters of the exogenous process for TFP in Table 2, including the impulse response of TFP in the target moments $\hat{p}si$ is irrelevant for our results.

[7] Estimating the agnostic wage process occasionally delivers "jagged" responses near the end of the impulse response horizon. For this reason, we augment the loss function (24) with a small penalty for acceleration (changes in the growth rate) of the wage. This penalty accounts for less than 1% of the loss function at the optimum and does not affect our results in any qualitative way.

Table 3: Parameter Estimates (Agnostic Wage)

Parameter	Concept	Estimate	Std. Err.
θ	External Habit	0.510	0.101
ϕ	Capital Adj. Cost	1.824	1.618
ϵ	Matching function elasticity	0.857	0.025

4.2 Results

Table 3 reports our baseline estimates for the first three elements of $\hat{\theta} \equiv \arg\min \mathcal{L}(\theta)$ along with standard errors generated from the asymptotic delta method following Guerron-Quintana et al. (2017). For brevity, we refrain from reporting numerical values for each of the 41 parameters of $\hat{\gamma}(L)$, and refer the reader instead to the figures below for their implications for the wage. Our structural parameter estimates are largely in line with existing literature.[8]

Figure 3 plots the impulse responses from our empirical identification procedure against those implied by our estimated model in response to a news shock. The model fit is excellent. The model almost exactly matches the impact effect, as well as the subsequent trajectories, of output and hours in the data. The response of the stock price is somewhat muted on impact relative to the data, but quickly catches up as investment and employment rise. Investment, in turn, rises less quickly in the model than in the data, but subsequent dynamics are similar.

Figure 4 plots the impulse response of our estimated agnostic wage process following a news shock. We overlay empirical impulse responses from 19 commonly used aggregate and industry-level wage series to our identified shock. The latter serve to highlight both the considerable heterogeneity in responses across various measures of the real wage, but also the presence of several systematic components of how wages respond to news about future movements in TFP. In particular, all but four of the series fall on impact, and all but two of the series eventually rise above their initial levels in response to the shock. As it happens, these features are precisely what we find in our agnostic wage process. Furthermore, it bears emphasizing that the model wage was in no way constrained to match these empirical patterns: Our estimation procedure relied on the six series in Y_t alone.

[8]Estimates of the elasticity of the matching function with respect to vacancies, ϵ, tend to vary depending on the methodology and data. Our estimate is relatively high, but within the range reported in Petrongolo and Pissarides (2001) and similar to the value estimated by Yashiv (2000) (0.87). Our estimate of the external habit, θ, is similar to values commonly found in the DSGE estimation literature, such as Christiano et al. (2005) (0.6) and Gertler et al. (2008) (0.7).

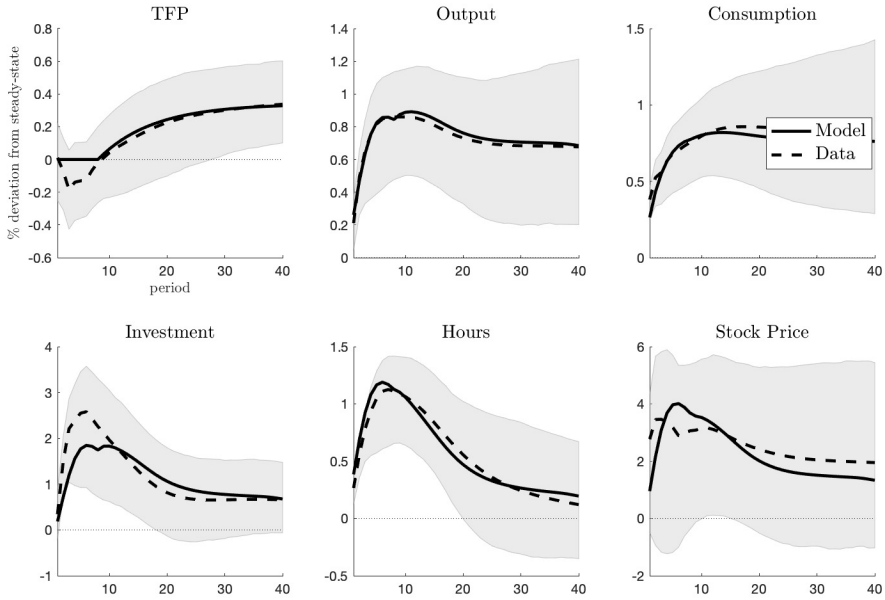

Figure 3: Quantity responses: Model (agnostic wage) and data

There are three principal take-aways from our estimation exercise: (i) A parsimoniously-specified labor search model with an entirely agnostic wage process can replicate the economy's dynamics response to a news shock; (ii) in order to do so, wages must fall on impact, remain low throughout the anticipation period, and rise quickly when the shock is realized; and (iii) such a wage response lies squarely within the range of empirical responses of wages to our identified shock, and is thus empirically plausible.

Of course, any reduced-form specification of wages, however well it may fit the data, is of limited utility beyond its ability to ultimately inform a structural theory of wage determination. We therefore take up the task of searching for a structural theory of wage determination that reflects what we have learned from our semi-structural exercise.

5 A Flow-Based Model of Wage Determination

Is there a structural model of wage determination that is consistent with our estimated agnostic wage process—and thus consistent with the economy's dynamic response to a news shock?

Models in which wages depend explicitly on the *present discounted value* of match surplus, such as Nash bargaining, will generally struggle to generate sizable anticipatory responses to

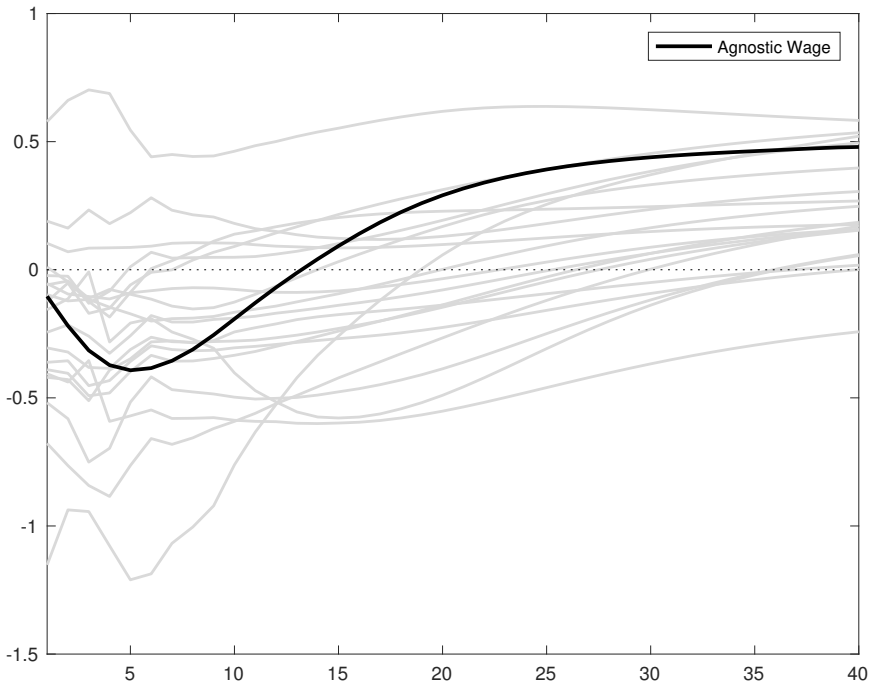

Figure 4: Wage response: Model (agnostic wage) and data

news shocks such as those we observe in the data. This is because, in the presence of matching frictions, a forward-looking wage throttles the benefits of hiring today in anticipation of higher future productivity. Indeed, as discussed above, most standard measures of real wages fall on impact in response to our identified shock, suggesting that such models will be poor candidates for explaining the data, both in terms of quantities and wages.

In light of this intuition, a natural alternative is a sharing rule based on the *flow* match surplus, rather than the present discounted value of match surplus. We thus propose a simple model of wage determination according to which workers receive a share of firms' available cash flow after accounting for payments to capital and the costs of hiring. This model, and the intuition underlying it, is closely related to the model studied by den Haan and Kaltenbrunner (2009). In particular, we consider a model in which wages are given by

$$W_t = \omega_0 P_t \tag{25}$$

where

$$P_t \equiv \frac{Y_t - R_t K_t - a_n V_t}{N_t}. \tag{26}$$

This process embodies the qualitative features towards which our estimated agnostic wage process pointed. Namely, it allows for wages to fall in response to expectations of a future

Table 4: Parameter Estimates (Flow wage)

Parameter	Concept	Estimate	Std. Err.
θ	External Habit	0.230	0.181
ϕ	Capital Adj. Cost	5.673	1.658
ϵ	Matching function elasticity	0.827	0.003
ω^F	Flow term	0.682	0.004

increase in productivity, and then rise when that increased productivity is finally realized. Why does the wage fall in response to expectations of future productivity? In the world we consider, a strong labor market is high is one in which (i) employment is high, so the average product of labor is relatively low, and (ii) expenditures on vacancies are relatively high, so cash flows after accounting for posting and capital costs are relatively low. Thus the wage can fall and the economy can boom when good news about future productivity arrives.[9]

5.1 Estimation

We next re-estimate the structural model described in the preceding section, replacing the flexible agnostic wage process with a version of (25) intended to allow the data to choose between our model and a simple inertial wage rule:

$$W_t = \omega_0 P_t^{\omega^F} W_{t-1}^{1-\omega^F} \tag{27}$$

where ω_0 is calibrated to match labor's share of income and ω^F is directly estimated. The model is otherwise identical to the model described in Section 3, and the estimation procedure is likewise unchanged. Since we are no longer estimating the 41 parameters associated with the reduced-form MA(40) wage process, and are instead estimating a single parameter governing the importance of the flow wage component, we are estimating 40 fewer parameters.

Table 4 reports results from our estimated model with the flow-based model of wage determination. Parameter estimates are broadly in line with the estimates from the model with the agnostic wage in Table 3. We estimate a somewhat lower degree of habit formation, a somewhat higher capital adjustment cost, and a nearly identical value for the matching function elasticity.

[9]Our main results continue to hold if we instead define P_t as per-worker cash flows net of capital costs only, $P_t \equiv (Y_t - R_t K_t)/N_t$.

Importantly, we estimate a value of 0.68 for ω^F, the weight on the flow wage process. This suggests that the flow wage process we describe is quantitatively important for accounting for quantity responses to a news shock. It is instructive to reflect on why this number might not be even higher. In particular, consider the case when there is no inertial component, so that $\omega^F = 1$. Momentarily neglecting capital, in this scenario it can be shown that the wage is proportional to the marginal product of labor.[10] It follows that the model will be unable to generate any significant volatility in hiring and thus hours. On the other hand, if $\omega^F = 0$, wages are entirely fixed; this case, wages again cannot fall on impact to the shock and hours will not move significantly prior to the realization of the shock. The value that we find for ω^F of 0.68 balances these forces, allowing employment to rise and hours to fall in anticipation of the shock.

Figure 5 plots the impulse responses from our empirical identification procedure against those implied by our simple model of wage setting. Despite the fact that we now have 40 fewer degrees of freedom, the model fit remains excellent. Observationally, there is little difference between Figures 3 and 5, while the minimized value of the criterion has increased only marginally despite the parsimony of the model—both in terms of the underlying search and matching structure and the single-parameter structural model of a flow-based wage.

To assess the correspondence between our estimated agnostic wage and our tightly-parameterized structural wage setting model, Figure 6 plots the response of both estimated wages processes as well as our panel of empirical wage responses. We include the flexible Nash bargained wage as a point of comparison.

Our estimated flow-based model of the wage is nearly identical to the 41-parameter reduced-form wage process we estimated in the previous section. Furthermore, both series lie well within the range of the empirical wage responses to our identified shock, falling on impact and eventually rising when the shock is realized. The Nash wage, by contrast, adjusts by a negligible amount on impact, only rising once the productivity improvement is realized. The inability of the Nash wage to fall on impact makes it incapable of generating the large anticipation effects on output and employment that we observe in the data.

5.2 Flow Wage: Critical Features

The empirical success of our simple flow-based wage mechanism rests on two key elements. The first critical feature is that the wage depends on contemporaneous cash flows, not the

[10] Precisely, $W_t = \Phi^N \cdot \frac{Y_t}{N_t}$ where Φ^N is the steady state share of compensation accruing to labor.

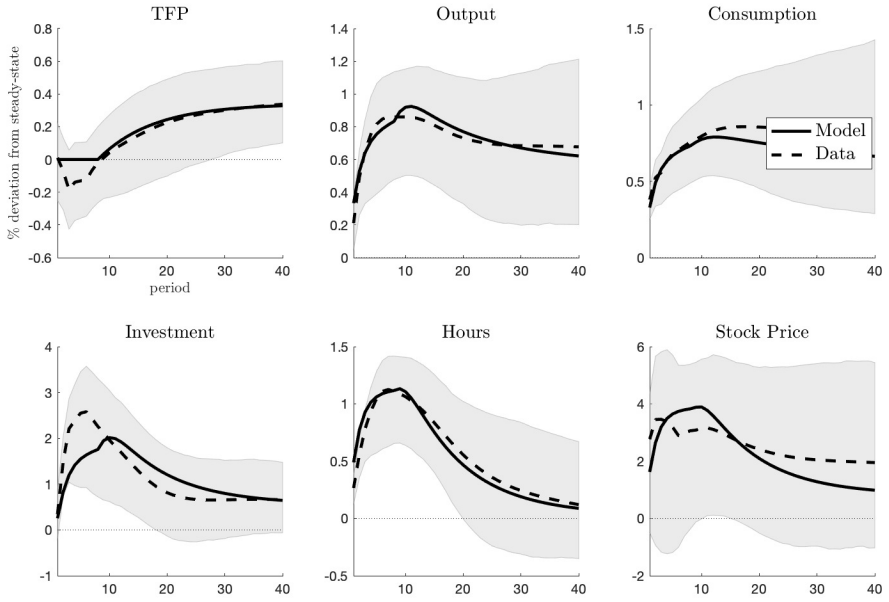

Figure 5: Quantity responses: Model (flow wage) and data

present discounted value of match surplus. For a given real interest rate and wage, expectations about higher future productivity increase the incentive for firms to post vacancies ahead of the realization of the shock. The question for the different wage mechanisms then becomes: how do real interest rates adjust in equilibrium?

If wages split the current period's surplus—cash flows net of non-wage costs—news about future productivity does not directly feed back into the wage, so that current hiring, output, and consumption can rise in anticipation of the shock. Higher current output, in turn, drives down the current marginal product of labor and wages while higher current consumption mitigates an increase the real interest rate, thus sustaining the increase in vacancy posting incentives.

In contrast, if wages are forward-looking and depend on the present-discounted value of match surplus, then the conjectured increase in the value of vacancies will increase wages, decreasing current hiring and consumption, forcing up the real interest rate which further decreasing the incentive to post vacancies. Through this negative feedback loop, the Nash bargained wage prevents a boom in anticipation of the future productivity, providing a result somewhat akin to the classic finding of Barro and King (1984), now adapted to the search economy.

This qualitative description of the mechanism that allows employment to rise in response

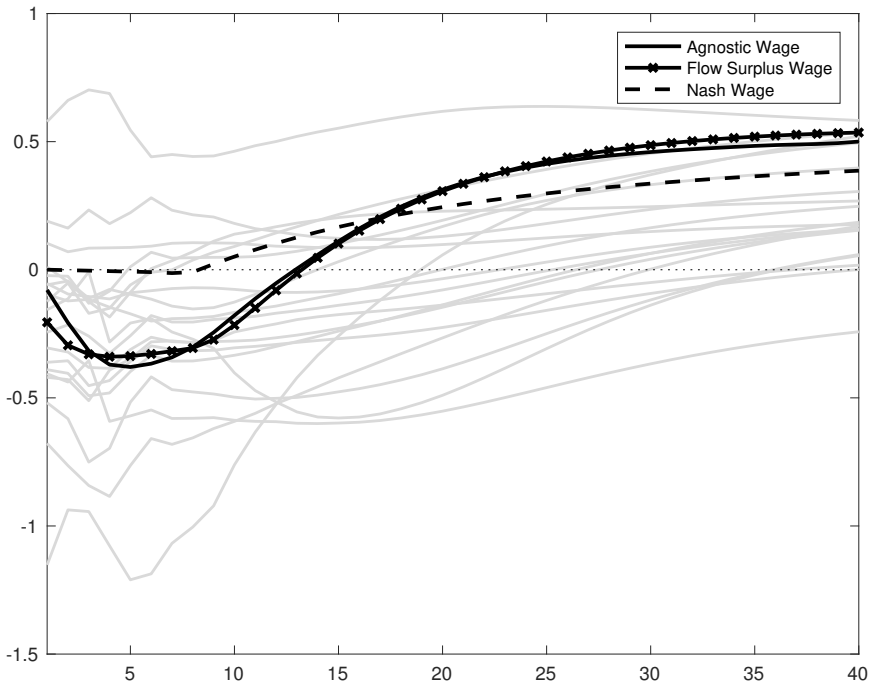

Figure 6: Estimated wage responses in data and model.

to an anticipated shock is general, but the quantitative power of the mechanism relies on a second key feature: the avereage share of match surplus captured by firms must be relatively small. This is precisely the observation of Hagedorn and Manovskii (2008) applied to the present context of anticipated shocks to productivity: When firms receive a small fraction of the total surplus, small changes in productivity translate to large (in percentage terms) changes in the value of a match, and leading to high volatility of firms' vacancy posting choices.

5.3 Non-targeted Moments

We next consider our model's implications for two widely studied features of labor markets that were not targeted in the course of estimation: the Beveridge curve and labor's share of income.

5.3.1 Beveridge Curve

Is our model consistent with the observed negative relationship between unemployment and vacancies in the data? To evaluate this, we compute the Beveridge curve using the most recent vacancy data (1996Q1-2016Q4) constructed as in Barnichon (2010), following the

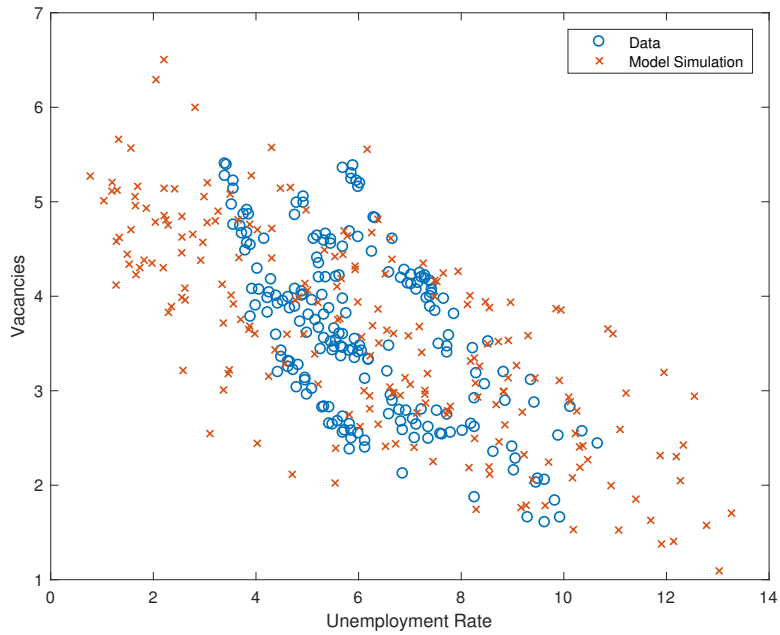

Figure 7: Beveridge curve

procedure of Elsby et al. (2015).

We estimate the slope of the Beveridge curve via OLS using the data described above and find a value of -0.32. Performing the same analysis on data simulated from our model, we find a value of -0.25. Figure 7 plots the curve for historical data and a single simulation of the same length. We take the ability of model to match these this out-of-sample target as further evidence that our estimated model provides a plausible account of labor market fluctuations.

5.3.2 Labor Share

What does our identified shock imply for labor's share of income, and is it consistent with our model? Because we prefer not to take a strong stance on the "right" wage series, to answer the first question, we construct labor's share of income using the average response of our panel of hourly wages, together with our series for hours and real GDP. Figure 8 depicts the empirical response of labor's share of income to our identified shock from Section 2. To answer the second question, we superimpose our model's implied response of labor's share of income to the news shock.

In the model, we find that labor's share falls very modestly in the model in response to the shock, and remains negative throughout the anticipation period and after the shock is realized. In the data (using the average of our wage series), labor's share likewise falls on

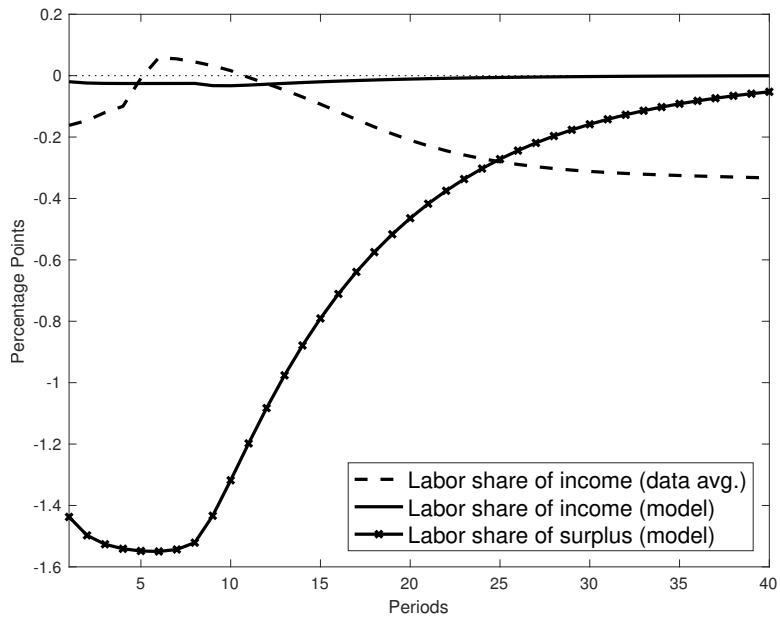

Figure 8: Labor's share of income and surplus

impact, briefly rises above its initial value when the shock is realized, before again falling to a lower level. None of these movements, however, are statistically distinguishable from a zero response. These results for the labor share of income contrast with the findings of Ríos-Rull and Santaeulàlia-Llopis (2010) for identified *surprise* technology shocks.

Despite the fact that the labor share is (close to) acyclical in both the model and the estimated data, we find that the model's implication for labor's share of surplus is much more pronounced. Labor's share of surplus is strongly counter-cyclical, falling by as much as 1.5 percentage points in response to the productivity shock. These endogenous fluctuations of labor's share of surplus are essential for the model's ability to deliver large swings in employment in response to our anticipated productivity shock.

6 Conclusion

This paper revisits a set of negative conclusions regarding the potential of productivity to be a main driver of labor market fluctuations. We show that both the data and a simple model of labor search are consistent with an important role for anticipated productivity shocks. The key requirement is a process for wages that falls modestly in response to good news about the future. Simple empirical wage measures provide tentative support for wages responding in this manner and a simple and plausible model of structural wage determination based on

cash flows delivers realistic responses from the theory.

One implication of our findings is that contemporaneous correlations, for example regarding the cyclicality of wages or the correlation of TFP with other variables, can mask important dynamic relationships that suggest tighter relationships between these objects. Similarly, focusing on only a portion of the spectrum in considering covariances can lead researchers to miss relationships that exist across frequencies in the data. While these insights are not new for macroeconomists, our results suggest that business cycle researchers should continue to take these features of the data into account as they search for the origins of fluctuations.

References

Angeletos, G.-M., F. Collard, and H. Dellas (2019). Business cycle anatomy. *Manuscript*.

Barnichon, R. (2010). Building a composite help-wanted index. *Economics Letters 109*(3), 175 – 178.

Barro, R. J. and R. G. King (1984). Time-separable Preferences and Intertemporal-Substitution Models of Business Cycles. *Quarterly Journal of Economics 99*(4), 817 – 839.

Barron, J., M. Berger, and D. Black (1997). Employer search, training, and vacancy duration. *Economic Inquiry 35*(1).

Barron, J. and J. Bishop (1985). Extensive search, intensive search, and hiring costs: New evidence on employer hiring activity. *Economic Inquiry 23*(3).

Barsky, R. B. and E. R. Sims (2011). News Shocks and Business Cycles. *Journal of Monetary Economics 58*(3), 273 – 289.

Basu, S., J. G. Fernald, and M. S. Kimball (2006). Are Technology Improvements Contractionary? *The American Economic Review 96*(5), pp. 1418–1448.

Basu, S. and C. House (2016). Allocative and remitted wages: New facts and challenges for keynesian models. Volume 2 of *Handbook of Macroeconomics*, pp. 297 – 354. Elsevier.

Basu, S. and C. House (2017). Allocative and remitted wages: New facts and challenges for keynesian models. *NBER Working Paper*.

Blanchard, O. J., J.-P. L'Huillier, and G. Lorenzoni (2013). News, Noise, and Fluctuations: An Empirical Exploration. *American Economic Review 103*(7), 3045–70.

Bouakez, H., L. Kemoe, et al. (2019). News shocks, business cycles, and the disinflation puzzle. Working paper.

Chahrour, R. and K. Jurado (2018). News or Noise? The Missing Link. *American Economic Review 108*(7), 1702–36.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (2005). Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy. *Journal of Political Economy 113*(1), 1–45.

Christiano, L. J., M. S. Eichenbaum, and M. Trabandt (2016). Unemployment and business cycles. *Econometrica 84*(4), 1523–1569.

Davis, S. J., R. J. Faberman, and J. Haltiwanger (2006, April). The flow approach to labor markets: New data sources and micro-macro links. Working Paper 12167, National Bureau of Economic Research.

den Haan, W. J. and G. Kaltenbrunner (2009). Anticipated growth and business cycles in matching models. *Journal of Monetary Economics*.

Edge, R. M. and R. S. Gürkaynak (2010). How useful are estimated dsge model forecasts for central bankers? [with comments and discussion]. *Brookings Papers on Economic Activity*, 209–259.

Elsby, M. W. L., R. Michaels, and D. Ratner (2015, September). The beveridge curve: A survey. *Journal of Economic Literature 53*(3), 571–630.

Faccini, R. and L. Melosi (2019). Pigouvian cycles. *Manuscript*.

Faust, J. (1998). The robustness of identified var conclusions about money. *Carnegie-Rochester Conference Series on Public Policy*.

Fernald, J. G. (2014). A quarterly, utilization-adjsuted series on total factor productivity. *Federal Reserve Bank of San Francisco Working Paper Series*.

Francis, N., M. Owyang, J. Roush, and R. DiDecio (2014). A flexible finite-horizon alternative to long-run restrictions with an application to technology shocks. *Review of Economics and Statistics 96*(4).

Fujita, S. and G. Ramey (2012). Exogenous versus endogenous separation. *American Economic Journal: Macroeconomics 4*(4).

Gertler, M., L. Sala, and A. Trigari (2008). An estimated monetary dsge model with unemployment and staggered nominal wage bargaining. *Journal of Money, Credit and Banking 40*(8), 1713–1764.

Guerron-Quintana, P., A. Inoue, and L. Kilian (2017). Impulse response matching estimators for dsge models. *Journal of Econometrics 196*(1), 144 – 155.

Haefke, C., M. Sonntag, and T. van Rens (2013). Wage rigidity and job creation. *Journal of Monetary Economics 60*(8), 887 – 899.

Hagedorn, M. and I. Manovskii (2008, September). The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited. *American Economic Review 98*(4), 1692–1706.

Hall, R. E. (2005). Employment fluctuations with equilibrium wage stickiness. *The American Economic Review 95*(1), 50–65.

Hall, R. E. (2017). High discounts and high unemployment. *American Economic Review 107*(2).

Kudlyak, M. (2014). The cyclicality of the user cost of labor. *Journal of Monetary Economics* (0), –.

Kurmann, A. and E. Sims (2017, February). Revisions in utilization-adjusted tfp and robust identification of news shocks. Working Paper 23142, National Bureau of Economic Research.

Mortensen, D. T. and C. A. Pissarides (1994, July). Job Creation and Job Destruction in the Theory of Unemployment. *Review of Economic Studies 61*(3), 397–415.

Petrongolo, B. and C. Pissarides (2001). Looking into the black box: A survey of the matching function. *Journal of Economic Literature*.

Ríos-Rull, J.-V. and R. Santaeulàlia-Llopis (2010). Redistributive shocks and productivity shocks. *Journal of Monetary Economics 57*(8), 931 – 948.

Schmitt-Grohé, S. and M. Uribe (2012). What's News in Business Cycles. *Econometrica 80*, 2733–2764.

Shimer, R. (2005). The cyclical behavior of equilibrium unemployment and vacancies. *The American Economic Review 95*(1), pp. 25–49.

Theodoridis, K. and F. Zanetti (2016). News shocks and labour market dynamics in matching models. *Canadian Journal of Economics*.

Uhlig, H. (2003). What moves real GNP? *Unpublished manuscript*.

Uhlig, H. (2004). Do technology shocks lead to a fall in total hours worked? *Journal of the European Economic Association*.

Yashiv, E. (2000). The determinants of equilibrium unemployment. *American Economic Review 90*(5).

Yashiv, E. (2008). U.s. labor market dynamics revisited. *Scandinavian Journal of Economics 109*(4).

A Stationary Representation

The model described in the body of the text is trend stationary with respect to labor-augmenting technological progress, X_t. Denoting by tildes the stationary counterparts to non-stationary variables, we can write the model in terms of only stationary variables:

$$\tilde{Y}_t = A_t \left(\tilde{K}_t\right)^{\alpha} (\gamma_{x,t} N_t)^{1-\alpha} \tag{28}$$

$$S_t = \frac{1 - N_t}{1 - p_t} \tag{29}$$

$$N_t = (1-\lambda) N_{t-1} + M_t \tag{30}$$

$$\tilde{K}_{t+1} = \gamma_{x,t}^{-1} \left[(1-\delta)\tilde{K}_t + \tilde{I}_t\right] \tag{31}$$

$$\tilde{Y}_t = \tilde{C}_t + \tilde{I}_t + \tilde{G}_t + \tilde{K}_t \Phi_k \left(\frac{\tilde{I}_t}{\tilde{K}_t}\right) + a_n V_t \tag{32}$$

$$\tilde{D}_t = \tilde{Y}_t - \tilde{W}_t N_t - \tilde{R}_t \tilde{K}_t - a_n V_t \tag{33}$$

$$\frac{\tilde{\mu}_t^K}{\tilde{\mu}_t} = 1 + \Phi_k' \left(\frac{\tilde{I}_t}{\tilde{K}_t}\right) \tag{34}$$

$$\frac{\tilde{\mu}_t^K}{\tilde{\mu}_t} = E_t \left\{ \Omega_{t,t+1} \left[(1-\delta) \frac{\tilde{\mu}_{t+1}^K}{\tilde{\mu}_{t+1}} + R_{t+1} - \Phi_k \left(\frac{\tilde{I}_{t+1}}{\tilde{K}_{t+1}}\right) + \frac{\tilde{I}_{t+1}}{\tilde{K}_{t+1}} \Phi_k' \left(\frac{\tilde{I}_{t+1}}{\tilde{K}_{t+1}}\right) \right] \right\} \tag{35}$$

$$\frac{a_n}{q_t} = (1-\alpha) A_t \left(\frac{\tilde{K}_t}{\gamma_{x,t} N_t}\right)^{\alpha} \gamma_{x,t} - \tilde{W}_t + (1-\lambda) E_t \left\{ \Omega_{t,t+1} \gamma_{x,t} \frac{a_n}{q_{t+1}} \right\} \tag{36}$$

$$\tilde{\mu}_t^N = \frac{a_n}{q_t} \tag{37}$$

$$\tilde{R}_t = \alpha A_t \left(\frac{\tilde{K}_t}{\gamma_{x,t} N_t}\right)^{\alpha-1} \tag{38}$$

where $\Omega_{t,t+1} \equiv \beta \frac{\tilde{\mu}_{t+1}}{\tilde{\mu}_t} \gamma_{x,t}^{-\sigma}$ and non-stationary variables are detrended according to $\tilde{\Delta}_t \equiv \frac{\Delta_t}{X_{t-1}}$ for $\Delta_t \in \{Y_t, C_t, D_t, W_t, K_t, I_t, \mu_t^N\}$ and $\tilde{\Delta}_t \equiv \frac{\Delta_t}{X_{t-1}^{-\sigma}}$ for $\Delta_t \in \{\mu_t, \mu_t^K\}$.

B Steady State and Calibration

We use the restrictions imposed by the deterministic steady state of the model, together with long-run values of $\bar{N}$, $\bar{\Phi}^N$ and $\bar{q}$ taken from the data, to analytically solve for all remaining endogenous variables, as well as χ and a_n. Below, all variables are detrended, and variables with bars denote long-run values taken from the data.

The first-order condition for investment implies that $\frac{\mu^K}{\mu} = 1$. This, together with the first-order condition for period-ahead capital imply

$$R = \frac{1}{\Omega} - 1 + \delta.$$

Solving the capital demand equation gives

$$K = \bar{N}\gamma_x \left(\frac{R}{\alpha}\right)^{1/(\alpha-1)}$$

which allows us to solve for output

$$Y = K^{\alpha}(\gamma_x \bar{N})^{1-\alpha}.$$

The timing of the labor market, together with the requirement that $M = \lambda\bar{N}$, implies that the number of searching workers is

$$S = 1 - (1-\lambda)\bar{N}$$

from which we obtain p and V using the definition of match probabilities for firms and workers, respectively

$$p = M/S$$
$$V = M/\bar{q}.$$

With values of S, V and M, we use the matching function to solve for match efficiency χ

$$\chi = M/(V^{\epsilon}S^{1-\epsilon}).$$

We next use the long-run value of labor's share of income $\bar{\Phi}^N \equiv \frac{W\bar{N}}{Y}$ to pin down the steady-state wage

$$W = \bar{\Phi}^N Y/\bar{N}$$

and the vacancy posting condition to solve for the vacancy posting cost a_n

$$a_n = \frac{\bar{q}}{1-(1-\lambda)\Omega\gamma_x}\left[(1-\alpha)A\left(\frac{K}{\bar{N}\gamma_x}\right)^{\alpha} - W\right].$$

Finally, the law of motion for capital and the aggregate resource constraint imply

$$I = K(\gamma_x - 1 + \delta_0)$$

and

$$C = Y - G - a_n V - I.$$

C Data Sources and Construction

Our main VAR specification consists of TFP, output, consumption, investment, hours, and the stock price. Except when otherwise noted, we download these series from the FRED database of the St. Louis Federal Reserve Bank.

For TFP, we use the capacity utilization adjusted measure described by Basu et al. (2006) and downloaded from https://www.frbsf.org/economic-research/indicators-data/ on June 1, 2019. To compute the level of TFP we cumulate the growth rates starting from the initial observation in 1947Q2.

Quantity variables are provided in real per-capita terms. Our population series is the civilian non-institutional population ages 16 and over, produced by the BLS. We convert this series to quarterly frequency using a three-month average and smooth it using an HP-filter with penalty parameter $\lambda = 1600$ to account for occasional jumps in the series that occur after census years and CPS rebasings (see Edge and Gürkaynak (2010)). Our deflator series is the GDP deflator produced by the BEA national accounts.

For output, we use seasonally adjusted nominal output produced by the BEA divided by the population and the GDP deflator. For investment, we take the sum of nominal gross private domestic investment and personal expenditure on durable goods, again divided by the population and the GDP deflator. Consumption consists of nominal personal consumption expenditures on non-durables and services, also divided by the GDP deflator and population. Our measure of total private sector hours come from the BLS Labor Productivity and

Table 5: CES Sectoral Wage Series

Sector	**Code**
Total Private	AHETPI
Goods Producing	CES0600000008
Mining	CES1000000008
Manufacturing	CES3000000008
Services	CES0800000008
Trade, Transportation, and Utilities	CES4000000008
Wholesale Trade	CES4142000008
Retail Trade	CES4200000008
Transportation and Warehousing	CES4300000008
Utilities	CES4422000008
Information	CES5000000008
Financial Activities	CES5500000008
Professional and Business Services	CES6000000008
Education and Health Services	CES6500000008
Leisure and Hospitality	CES7000000008
Other Services	CES8000000008

Costs release (Nonfarm Business Sector: Hours of All Persons) and is also divided by the population.

Lastly, our measure of real stock prices comes from Robert Shiller, and was download from on June 1, 2019.

We augment these variables with 19 measures of the aggregate and sectoral wages. Our preferred wage measure is comes from the BEA National Accounts, series code A132RC, and consists of wage and salary compensation for private industries. To arrive at an hourly wage, we divide this hours series used above and the GDP deflator.

The additional elements of wage panel includes (i) median weekly earnings divided by the GDP deflator from the BLS's Current Population Survey (ii) the new hire real wage series produced by Basu and House (2016) and downloaded from and (iii) sixteen additional hourly wage series originating from the super-sector classification level of the BLS's Current Establishment Survey. These series are listed in Table 5. We download each form the FRED database in nominal terms divid them by the GDP deflator to arrive at real hours wages.

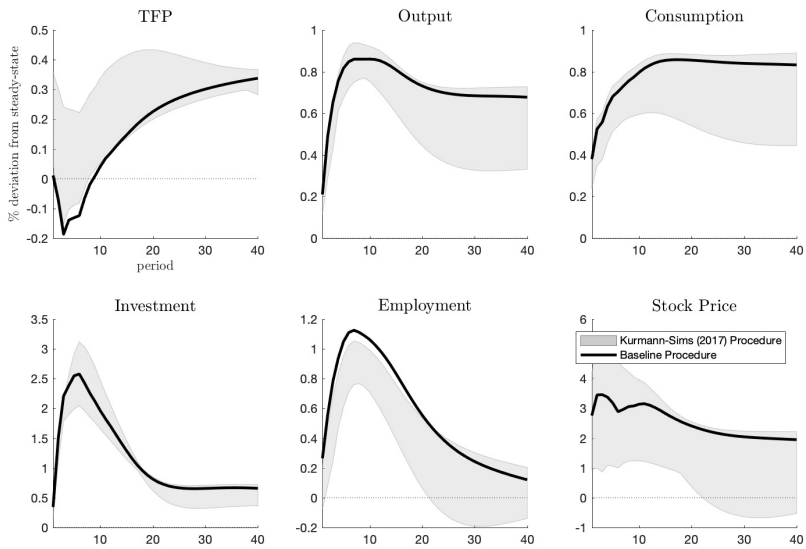

Figure 9: Comparison with the Kurmann and Sims (2017) procedure for target horizons of 20 through 160 quarters.

D Robustness

D.1 Empirical Exercise

Our empirical impulse responses are robust to (i) changing the number of lags in the VAR (ii) running in VECM imposing one, two, or more trends in the data and (iii) expanding the set of observables in Y_t to include additional variables, such as alternative labor market indicators.

As noted in the main text, our agnostic identification procedure yield impulse responses that very similar to those implied by the News identification procedure of Kurmann and Sims (2017). Those authors identify a news shock as the shock that explains the largest portion of the forecast error of TFP as some distant horizon, taking a horizon of 80 quarters as their baseline. Figure 9 presents the impulse responses from our estimation along with range of point estimates generated by apply their procedure for horizons between 20 and 160 quarters. The figure shows that the impulse responses implied by this alternative procedure are quite similar to our own, except with respect to the stock price (which is marginally significant in any case.)

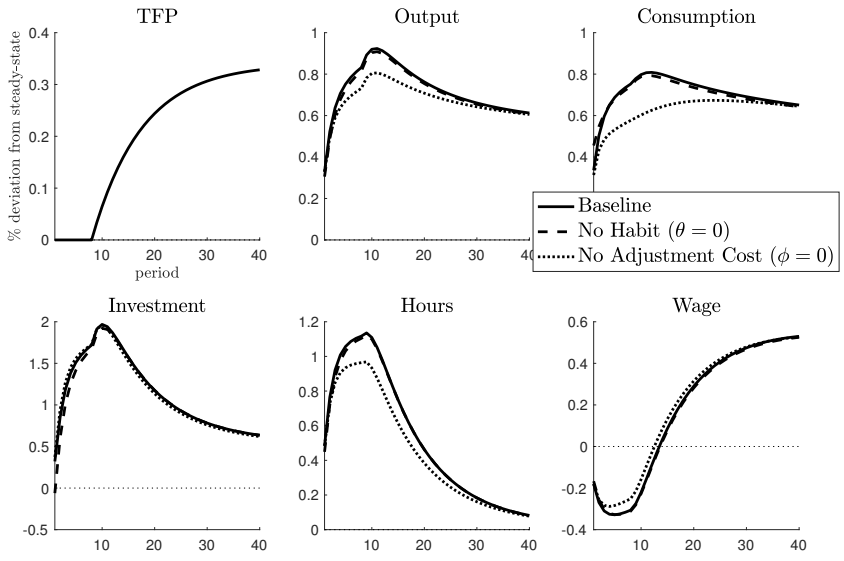

Figure 10: Robustness

D.2 Model

As discussed in the text, our parsimonious model provides an excellent fit for the data. Beyond the flow-based model of the wage, our model incorporates two standard features from the DSGE estimation literature that help it to match the data: an external habit in consumption (θ) and capital adjustment costs (ϕ). Figure 10 reports impulse responses from our estimated flow-wage model when we separately shut down these features.

Shutting down habit formation (dashed line) generates a counterfactually large impact response in consumption and a counterfactual negative impact response from investment, as expected. Shutting down capital adjustment costs, on the other hand, is associated with a marginally smaller decline in the wage during the anticipation period, and a corresponding reduction in the magnitude in the response of hours, output and consumption.