ONLINE APPENDIX

THE DECLINE OF THE LABOR SHARE: NEW EMPIRICAL EVIDENCE

DRAGO BERGHOLT[§], FRANCESCO FURLANETTO[†] AND NICOLÒ MAFFEI FACCIOLI[‡]

October, 2020

This appendix includes the following material: Appendix A offers additional details on the theoretical framework used in the main text. Appendix B explains the econometric approach and how we operationalize the theory robust sign restrictions. Appendix C documents and discusses a battery of robustness exercises. Appendix D provides additional figures.

A ADDITIONAL DETAILS ON THE THEORETICAL MODEL

A.1 THE INITIAL STEADY STATE USED FOR SIMULATIONS

The steady state in the baseline, theoretical model follows recursively given an initialization of s_l, s_k, s_d, A_l, Υ and L:

$$
\begin{aligned}
r &= \beta^{-1} - 1 \\
P_I &= \Upsilon^{-1} \\
r^k &= P_I \left(r + \delta \right) \\
\alpha_l &= \frac{s_l}{1 - s_d} \\
\alpha_k &= \frac{s_k}{1 - s_d} \\
\mathcal{M}_p &= \frac{1}{1 - s_d} \\
A_k &= r^k s_k^{\frac{1}{\eta - 1}} \left(\frac{\mathcal{M}_p}{\alpha_k} \right)^{\frac{\eta}{\eta - 1}} \\
Y &= \left(\frac{\alpha_l}{\mathcal{M}_p s_l} \right)^{\frac{\eta}{\eta - 1}} A_l L \\
W &= s_l \frac{Y}{L} \\
K &= \frac{s_k Y}{r^k}
\end{aligned}
$$

[§]Norges Bank. P.O. Box 1179 Sentrum, 0107 Oslo, Norway. E-mail: drago.bergholt@norges-bank.no.

[†]Norges Bank and BI Norwegian Business School. P.O. Box 1179 Sentrum, 0107 Oslo, Norway. Corresponding author. E-mail: francesco.furlanetto@norges-bank.no.

[‡]Universitat Autònoma de Barcelona, Barcelona GSE and IGIER, Università Bocconi. E-mail: nicolo.maffei@barcelonagse.eu. The author gratefully acknowledges financial support from the "La Caixa-Severo Ochoa International Doctoral Fellowship".

$$
\begin{aligned}
I &= \delta K \\
X &= P_I I \\
C &= Y - X \\
\mathcal{D} &= s_d Y \\
\epsilon_p &= \frac{\mathcal{M}_p}{\mathcal{M}_p - 1} \\
\mathcal{M}_w &= \mathcal{M}_p \\
\epsilon_w &= \frac{\mathcal{M}_w}{\mathcal{M}_w - 1} \\
\Psi &= \frac{W}{\mathcal{M}_w L^{\varphi} C} \\
\Lambda &= C^{-\sigma} \exp\left(-\Psi \frac{(1-\sigma) L^{1+\varphi}}{1+\varphi}\right)
\end{aligned}
$$

The steady state of the New Keynesian model (see Appendix A.4.2) is identical, except that we also have to solve for nominal variables: given a choice of gross inflation Π (we set $\Pi = 1$), we have $i_p = \frac{\Pi}{\beta} - 1$ and $\Pi_w = \Pi$.

A.2 LONG-RUN INCOME SHARES

Given the definition of a long-run equilibrium in the main text, we set out to derive expressions of long-run income shares. We start with the profit income share. It follows from the definition of profits and optimal price-setting behavior:

$$
\bar{s}_{d,t} = \frac{\bar{\mathcal{D}}_t}{\bar{Y}_t} = 1 - \frac{1}{\bar{\mathcal{M}}_{p,t}}
$$

Thus, the long-run profit income share depends only on firms' markup, which is assumed exogenous in the baseline model. In order to derive the long-run capital income share we note that

$$
\bar{r}^k_t = \bar{\Upsilon}_t^{-1} \left[\beta^{-1} - (1-\delta)\right]
$$

The expression for firms' optimal capital demand can then be used to arrive at the following long-run capital share:

$$
\bar{s}_{k,t} = \frac{\bar{r}^k_t \bar{K}_{t-1}}{\bar{Y}_t} = \left(\frac{\bar{\alpha}_{k,t}}{\bar{\mathcal{M}}_{p,t}}\right)^{\eta} \left(\frac{\beta^{-1} - (1-\delta)}{\bar{\Upsilon}_t \bar{A}_{k,t}}\right)^{1-\eta}
$$

This expression shows that automation (firms' markup) raises (lowers) the capital income share. The effects of investment-specific or capital-biased technologies depend qualitatively on whether or not η is higher than one. Labor-augmenting technology and labor markups have no long-run effects on the capital share. Finally, the labor income share is found by substituting the two expressions derived above into the identity $\bar{s}_{l,t} + \bar{s}_{k,t} + \bar{s}_{d,t} = 1$:

$$
\bar{s}_{l,t} = \frac{\bar{W}_t \bar{L}_t}{\bar{Y}_t} = \frac{1}{\bar{\mathcal{M}}_{p,t}} \left[1 - \bar{\alpha}^{\eta}_{k,t} \left(\frac{\beta^{-1} - (1-\delta)}{\bar{\Upsilon}_t \bar{A}_{k,t}} \bar{\mathcal{M}}_{p,t}\right)^{1-\eta}\right]
$$

This is the equation used in the main text.

A.3 DIRECT PRODUCTIVITY EFFECTS OF AUTOMATION

Next, we set out to characterize a situation in which an automation shock—a permanent rise in $\alpha_{k,t}$—impacts total factor productivity (TFP). Such a direct productivity effect is incorporated in some micro-founded models of task-based automation, see for example Acemoglu and Restrepo (2018) and Martinez (2019). One concern in this respect is that the direct productivity effect may lead to a violation of the sign restrictions we use to identify the automation shock. In particular, if automation translates into higher TFP, then the real wages may rise in the long run, thus, potentially invalidating our sign restrictions. However, this is only a concern if the real wage increase takes place within the horizon at which we impose our sign restrictions.

To fix ideas, we restrict attention to the long-run equilibrium described above but consider a situation in which the production function is equal to

$$\bar{Y}_t = \bar{A}_t \left[\bar{\alpha}_{l,t} \left(\bar{A}_{l,t} \bar{L}_t \right)^{\frac{\eta-1}{\eta}} + \bar{\alpha}_{k,t} \left(\bar{A}_{k,t} \bar{K}_{t-1} \right)^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}} .$$

This is the same production function as the one used in the main text, except that now we have an additional stochastic supply shifter, $\bar{A}_t$. This variable is interpreted as TFP, and it may or may not be a function of automation. The general formulation used by Acemoglu and Restrepo (2018) implies that $\bar{A}_t$ depends on $\bar{\alpha}_{k,t}$, with $\frac{\partial \bar{A}_t}{\partial \bar{\alpha}_{k,t}} \geq 0$. We do not specify the pass-through from $\bar{\alpha}_{k,t}$ to $\bar{A}_t$ here, but loosely refer to it as the direct productivity effect of automation. Firms' optimal factor demand schedules, taking TFP changes into account, follow below:

$$\bar{r}_t^k \bar{\mathcal{M}}_{p,t} = \bar{\alpha}_{k,t} \left(\bar{A}_t \bar{A}_{k,t} \right)^{\frac{\eta-1}{\eta}} \left(\frac{\bar{Y}_t}{\bar{K}_{t-1}} \right)^{\frac{1}{\eta}}$$

$$\bar{W}_t \bar{\mathcal{M}}_{p,t} = \bar{\alpha}_{l,t} \left(\bar{A}_t \bar{A}_{l,t} \right)^{\frac{\eta-1}{\eta}} \left(\frac{\bar{Y}_t}{\bar{L}_t} \right)^{\frac{1}{\eta}}$$

Using the same procedure as in subsection A.2, we can show that

$$\bar{s}_{l,t} = \frac{\bar{W}_t \bar{L}_t}{\bar{Y}_t} = \frac{1}{\bar{\mathcal{M}}_{p,t}} \left[1 - \bar{\alpha}_{k,t}^{\eta} \left(\frac{\beta^{-1} - (1-\delta)}{\bar{\Upsilon}_t \bar{A}_t \bar{A}_{k,t}} \bar{\mathcal{M}}_{p,t} \right)^{1-\eta} \right] .$$

Thus, in general, a direct productivity effect mutes (amplifies) the long-run labor share decline in response to automation if $\eta < 1$ ($\eta > 1$), given that $\frac{\partial \bar{A}_t}{\partial \bar{\alpha}_{k,t}} \geq 0$.[1] We are not in a position to draw further conclusions at this level of generality, but our normalization strategy, following Cantore and Levine (2012) and others, facilitates further insights at the expense of an additional assumption. The normalization of $\bar{A}_{k,t}$, in particular—tailored to make the initial steady state independent of η—implies that $\frac{\beta^{-1}-(1-\delta)}{\bar{\Upsilon}_t \bar{A}_t \bar{A}_{k,t}} \bar{\mathcal{M}}_{p,t} = \bar{\alpha}_{k,t}$. This is evident from the long-run labor share expression above. With this normalization, the long-run labor share collapses to

$$\bar{s}_{l,t} = \frac{\bar{W}_t \bar{L}_t}{\bar{Y}_t} = \frac{1 - \bar{\alpha}_{k,t}}{\bar{\mathcal{M}}_{p,t}},$$

[1] One could in principle also imagine that automation affects the levels of labor or capital augmenting technologies $\bar{A}_{l,t}$ and $\bar{A}_{k,t}$. We do not consider such cases here, but note that pass-through to $\bar{A}_{l,t}$ has zero effect on the long-run labor share, while the implications of pass-through to $\bar{A}_{k,t}$ (or $\bar{\Upsilon}_t$) is determined by whether or not η is greater than unity.

so that we can parameterize the initial steady state of the economy without any reference to the value of η.

What about real wages? Combing expressions, we can show that the general long-run solution for the real wage is

$$\bar{W}_t = \bar{s}_{l,t}\frac{\bar{Y}_t}{\bar{L}_t} = \frac{\bar{A}_t\bar{A}_{l,t}}{\bar{\mathcal{M}}_{p,t}}\left(\left(\frac{1}{\bar{\alpha}_{l,t}}\right)^{\eta}\left[1-\bar{\alpha}_{k,t}^{\eta}\left(\frac{\beta^{-1}-(1-\delta)}{\Upsilon_t\bar{A}_t\bar{A}_{k,t}}\bar{\mathcal{M}}_{p,t}\right)^{1-\eta}\right]\right)^{\frac{1}{1-\eta}}$$

Taking the derivative with respect to $\bar{\alpha}_{k,t}$ and imposing the normalization, we arrive at

$$\frac{\partial\bar{W}_t}{\partial\bar{\alpha}_{k,t}} = \frac{\bar{A}_{l,t}}{\bar{\mathcal{M}}_{p,t}}\frac{\partial\bar{A}_t}{\partial\bar{\alpha}_{k,t}}.$$

Thus, in the very long run, the real wage will rise as long as $\frac{\partial\bar{A}_t}{\partial\bar{\alpha}_{k,t}} > 0$. In our baseline model, in contrast, $\frac{\partial\bar{A}_t}{\partial\bar{\alpha}_{k,t}} = 0$ and automation does not affect the real wage permanently. However, even with productivity effects present, the real wage may still decline in the short and medium run, as it takes time to adjust capital. Whether or not the real wage declines initially, depends on how sensitive TFP is to changes in automation.[2] In the main text, we effectively restrict attention to a scenario in which automation only has modest effects on productivity, consistent with the findings in US data by Acemoglu and Restrepo (2020).

A.4 ALTERNATIVE THEORETICAL ASSUMPTIONS

This section documents how robust our sign restrictions are to (i) mis-measurement of profit income, and (ii) the inclusion of various real and nominal frictions.

A.4.1 MEASUREMENT OF PROFITS

A potential issue with the analysis in Section 3.1 in the main text concerns our measurement of profit income. The profit variable displayed in Figure 2 and Figure 3 is model-consistent and interpreted as sales net of factor payments. However, empirical measurements of profits might be distorted by the inclusion of some unobserved, intangible capital income (Karabarbounis and Neiman, 2019). Therefore, as a robustness check we now take the extreme view that *all capital income* is counted as profits in data, and simply refer to profit revenues $D_{k,t}$ as non-labor income:

$$D_{k,t} = D_t + r_t^k K_{t-1} = Y_t - W_t L_t$$

Given this new measure, we re-evaluate the model's implied sign restrictions. Figure A.1 compares impulse responses of pure profits, D_t, with those of non-labor income $D_{k,t}$. The medium- to long-run signs of either variable are largely identical for all shocks. Our only disclaimer in this regard is that, conditional on investment-specific technology shocks, about 6% of the models imply a decline in $D_{k,t}$ at horizons relevant for our sign restrictions.

[2] A micro-founded discussion along these lines is found in Acemoglu and Restrepo (2018), see their Figure 6.

Figure A.1: Monte Carlo results: pure profits vs. non-labor income in the baseline model

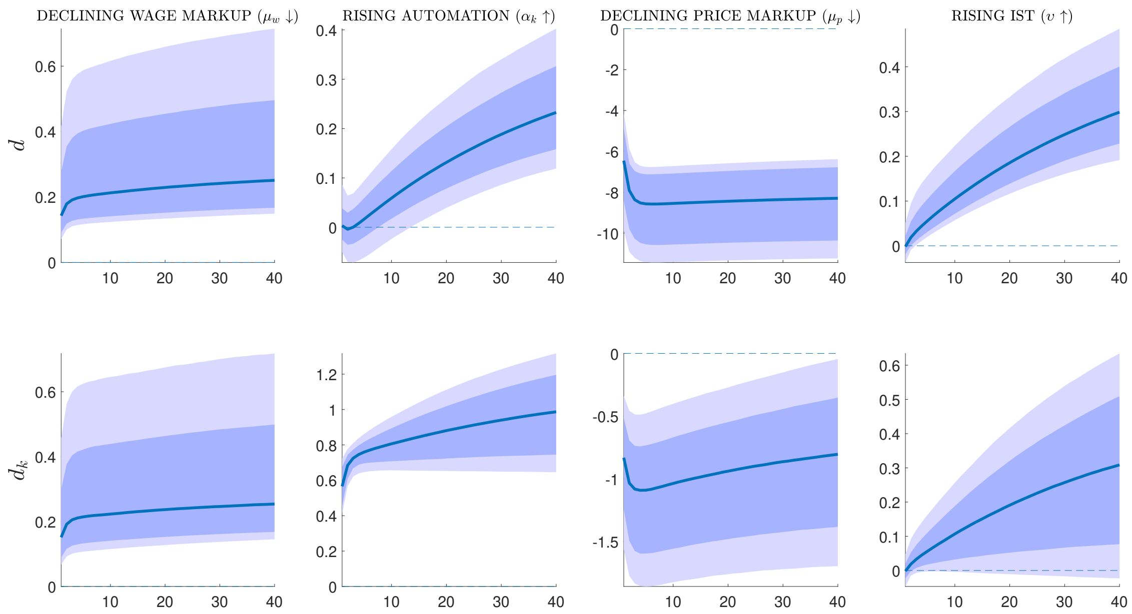

Note: Median (solid line), 90%, and 68% credible bands based on 10000 draws. Pure profits (d) and non-labor income (d_k) are both expressed in percentage deviations from initial values.

A.4.2 Real and Nominal Frictions

The theoretical model presented in the main text abstracts from a number of commonly used real and nominal frictions. One potential concern, therefore, is that our sign restrictions might be violated at certain frequencies if these frictions are included. This section incorporates a few "bells and whistles" into the baseline, theoretical model. We add (i) habit formation in consumption, (ii) adjustment costs in investments, (iii) variable capital utilization, (iv) nominal price stickiness, and (v) nominal wage stickiness. We also allow for partial indexation to past inflation in price and wage setting. Finally, we specify (vi) a Taylor-type rule for monetary policy. While the two models share identical long-run properties, the extended version implies different dynamics in the short to medium run. A brief summary of the additions to our baseline model follows:

External habit formation: The period utility is changed to

$$\mathcal{U}_t = \frac{\left(C_t - hC_{t-1}\right)^{1-\sigma}}{1-\sigma} \exp\left(-\Psi \frac{(1-\sigma)\, L_t^{1+\varphi}}{1+\varphi}\right).$$

Investment adjustment costs: We assume a convex investment adjustment cost, so that

$$K_t = (1-\delta)\, K_{t-1} + \left[1 - \frac{\chi}{2}\left(\frac{I_t}{I_{t-1}} - 1\right)^2\right] I_t.$$

Variable capital utilization: Wholesale firms rent effective capital services $\bar{K}_t = U_t K_{t-1}$, where U_t is the utilization rate of capital. Higher utilization comes at a cost $AC_{u,t}$ paid by

Figure A.2: Monte Carlo results: New Keynesian model with additional bells and whistles

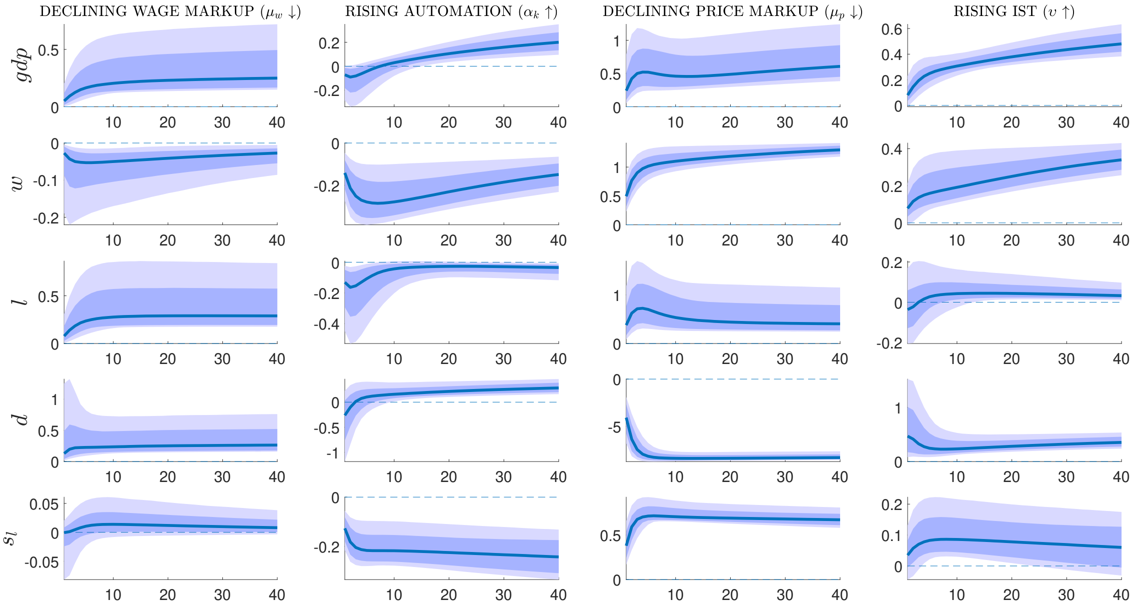

Note: Median (solid line), 90%, and 68% credible bands based on 10000 draws. Income shares are expressed in percentage point deviations from initial values. Remaining variables are expressed in percentage deviations.

households who own the capital, where

$$AC_{u,t} = \xi_u' (U_t - 1) + \frac{\xi_u \xi_u'}{2} (U_t - 1)^2 .$$

Nominal price stickiness: We incorporate price stickiness á la Rotemberg (1982). Nominal price adjustments are costly for wholesale firms. We also allow for partial indexation to past inflation and specify the cost function as

$$AC_{p,t} = \frac{\xi_p}{2} \left(\frac{\Pi_{jp,t}}{\Pi_{p,t-1}^{\gamma_p} \Pi_p^{1-\gamma_p}} - 1 \right)^2 Y_t .$$

Nominal wage stickiness: Wage stickiness á la Rotemberg (1982) is the final extension. Nominal wage adjustments come at a cost paid by households:

$$AC_{w,t} = \frac{\xi_w}{2} \left(\frac{\Pi_{nw,t}}{\Pi_{p,t-1}^{\gamma_w} \Pi_p^{1-\gamma_w}} - 1 \right)^2 L_t .$$

Monetary policy: Nominal rigidities imply the need to specify a nominal anchor. To this end we assume a Taylor type rule for the policy rate $i_{p,t}$:

$$1 + i_{p,t} = (1 + i_{p,t-1})^{\rho_i} \left[(1 + i_p) \left(\frac{\Pi_{p,t}}{\Pi_p} \right)^{\rho_\pi} \left(\frac{GDP_t}{GDP_{t-1}} \right)^{\rho_y} \right]^{1-\rho_i}$$

Table A.1: Additional parameter bounds in the model with real and nominal frictions

		M	*LB*	*UB*
h	Consumption habits	0.45	0	0.9
χ	Investment adjustment cost elasticity	5	0	10
ξ_u	Capital utilization cost elasticity	1.525	0.05	3
θ_w	Calvo parameter, wages	0.4	0	0.8
θ_p	Calvo parameter, prices	0.4	0	0.8
γ_w	Indexation, wages	0.375	0	0.75
γ_p	Indexation, prices	0.375	0	0.75
ρ_i	Interest inertia, Taylor rule	0.45	0	0.9
ρ_π	Inflation weight, Taylor rule	2	1	3
ρ_y	Output growth weight, Taylor rule	0.5	0	1

Note: Bounds for the uniform distributions. Notation: M $\rightarrow$ median; LB $\rightarrow$ lower bound; UB $\rightarrow$ upper bound. The parameters θ_p and θ_w represent the probabilities of being stuck with old prices and wages in the Calvo model. They do not appear in our model because we use Rotemberg pricing. However, we exploit the first order equivalence between Calvo and Rotemberg pricing in order to back out ξ_p and ξ_w, given θ_p and θ_w. The parameters σ_p, σ_w, σ_v, and σ_{α_k} are normalized so that impulse responses are computed conditional on a long-run change in $\mathcal{M}_{p,t}$, $\mathcal{M}_{w,t}$, Υ_t, and $\alpha_{k,t}$ of 1 percent.

The Fisher equation $(1 + i_{p,t}) = (1 + r_t)\,\Pi_{t+1}$ links nominal to real outcomes. We also note that wage adjustment costs enter $s_{l,t}$, utilization adjustment costs enter $s_{k,t}$, while price adjustment costs enter $s_{d,t}$. However, these shares still sum to one, and the long run properties of the model are unaffected. Finally, we note that the New Keynesian model captures the neoclassical setup as a special case ($h = \chi = \xi_p = \xi_w = 0$ and $\xi_u \rightarrow \infty$).

Bounds for the additional parameters are summarized in Table A.1 (all parameters appearing in the baseline model are drawn from the same distributions as those used in the main text). Figure A.2 documents the distributions of theoretical impulse responses when we include the real and nominal frictions just described. Importantly, the impulse responses are qualitatively similar across models even after a few periods, and the signs are identical from quarter 16 and onwards. We conclude, therefore, that the sign restrictions used in the main text are robust to the inclusion of real and nominal frictions.

B Bayesian Estimation of the VAR model

Consider the reduced form VAR model presented in Section 3.2:

$$Y_t = C + \sum_{j=1}^{p} A_j Y_{t-j} + u_t$$

The process above can be stacked in a more compact form as follows:

$$\mathbf{Y} = \mathbf{X}B + \mathbf{U}$$

where:
1) $\mathbf{Y} = (Y_{p+1}, ..., Y_T)'$ is a $(T-p)$ x n matrix, with $Y_t = (Y_{1,t}, ..., Y_{n,t})'$.
2) $\mathbf{X} = (\mathbf{1}, \mathbf{Y}_{-1}, ..., \mathbf{Y}_{-p})$ is a $(T-p)$ x $(np+1)$ matrix, where $\mathbf{1}$ is a $(T-p)$ x 1 matrix of ones and $\mathbf{Y}_{-k} = (Y_{p+1-k}, ..., Y_{T-k})'$ is a $(T-p)$ x n matrix.
3) $\mathbf{U} = (u_{p+1}, ..., u_T)'$ is a $(T-p)$ x n matrix.
4) $B = (C, A_1, ..., A_p)'$ is a $(np+1)$ x n matrix of coefficients.
Vectorizing the equation above, we obtain:

$$\mathbf{y} = (I_n \otimes \mathbf{X})\beta + \mathbf{u}$$

where $\mathbf{y} = vec(\mathbf{Y})$, $\beta = vec(B)$, $\mathbf{u} = vec(\mathbf{U})$ and $\mathbf{u} \sim N(0, \Sigma \otimes I_{T-p})$.
Given the assumption of normality of the reduced-form errors, $u_t \sim N(0, \Sigma)$, we can express the likelihood of the sample, conditional on the parameters of the model and the set of regressors $\mathbf{X}$, as follows:

$$L(\mathbf{y}|\mathbf{X}, \beta, \Sigma) \propto |\Sigma \otimes I_{T-p}|^{-\frac{T-p}{2}} exp\left\{\frac{1}{2}(\mathbf{y} - I_n \otimes \mathbf{X}\beta)'(\Sigma \otimes I_{T-p})^{-1}(\mathbf{y} - I_n \otimes \mathbf{X}\beta)\right\}$$

Denote $\hat{\beta} = vec(\hat{B})$, where $\hat{B} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y}$ is the OLS estimate, and let $S = (\mathbf{Y} - \mathbf{X}\hat{B})'(\mathbf{Y} - \mathbf{X}\hat{B})$ be the sum of squared errors. Then we can rewrite the likelihood as follows:

$$\begin{aligned} L(\mathbf{y}|\mathbf{X}, \beta, \Sigma) \propto & |\Sigma \otimes I_{T-p}|^{-\frac{T-p}{2}} exp\left\{\frac{1}{2}(\beta - \hat{\beta})'(\Sigma^{-1} \otimes \mathbf{X}'\mathbf{X})(\beta - \hat{\beta})\right\} \\ & exp\left\{-\frac{1}{2}tr(\Sigma^{-1}S)\right\} \end{aligned}$$

By choosing a non-informative (flat) prior for B and Σ that is proportional to $|\Sigma|^{-\frac{n+1}{2}}$, namely:

$$p(B|\Sigma) \propto 1$$
$$p(\Sigma) \propto |\Sigma|^{-\frac{n+1}{2}}$$

We can compute the posterior of the parameters given the data at hand using Bayes rule, as follows:

$$\begin{aligned} P(B, \Sigma|\mathbf{y}, \mathbf{X}) & \propto L(\mathbf{y}|\mathbf{X}, \beta, \Sigma)p(B|\Sigma)p(\Sigma) \\ & = |\Sigma|^{-\frac{T-p+n+1}{2}} exp\left\{\frac{1}{2}(\beta - \hat{\beta})'(\Sigma^{-1} \otimes \mathbf{X}'\mathbf{X})(\beta - \hat{\beta})\right\} exp\left\{-\frac{1}{2}tr(\Sigma^{-1}S)\right\} \end{aligned}$$

This posterior distribution is the product of a normal distribution for β conditional on Σ and an inverted Wishart distribution for Σ. Thus, we draw β conditional on Σ from:

$$\beta|\Sigma, \mathbf{y}, \mathbf{X} \sim N(\hat{\beta}, \Sigma \otimes (\mathbf{X}'\mathbf{X})^{-1})$$

and Σ from:

$$\Sigma|\mathbf{y}, \mathbf{X} \sim IW(S, v)$$

through Gibbs sampling, where $v = T - p - np - 1$.

The identification procedure described in **??** is performed using the algorithm of Rubio-Ramírez, Waggoner, and Zha (2010), which consists of the following steps, for each given draw from the posterior of the reduced-form parameters:

1. Draw a nxn matrix W from $N(0_n, I_n)$ and perform a QR decomposition of W, with the diagonal of R normalized to be positive and $QQ' = I_n$.

2. Let S be the lower triangular Cholesky decomposition of Σ and define $A = SQ'$. Compute the candidate impulse responses as $IRF_j = C_j A$, where C_j are the reduced form impulse responses from the Wold representation, for $j = 0, ..., J$. If the set of impulse responses satisfies all the sign restrictions, store them. If not, discard them and go back to the first step.

We repeat steps 1 and 2 until M impulse responses that satisfy the sign restrictions are obtained. The resulting set A, together with the reduced-form estimates, characterizes the set of structural VAR models that satisfy the sign restrictions.

C Robustness tests and additional results

In this section we report an extensive battery of robustness tests: first, we perform a Monte Carlo study in order to assess whether our empirical approach is able to identify the main drivers of the labor income share. Second, we check for robustness with respect to the horizon in which sign restrictions are imposed, as well as lag specifications in the SVAR, samples and priors. Third, we exploit additional sign and zero restrictions implied by the theoretical framework. Finally, we inspect the implications for labor market variables such as routine- and non-routine employment.

C.1 A Monte Carlo Experiment

First we assess whether our empirical approach is suitable for identifying the main drivers of the labor income share in data. As shown by Paustian (2007), SVAR models can have a hard time recovering the structural shocks of interest if these shocks have a sufficiently small variance. One may be concerned that our automation shock for example, which explains a rather small share of the macroeconomic variables in the system, is subject to this drawback. To this end, we set up a two-step procedure in order to evaluate the baseline SVAR model's ability to recover the relevant shocks for the labor share. The procedure is effectively a controlled experiment: first, we use a version of the theoretical model with fixed parameters in order to generate an artificial dataset which includes time series of real GDP, real wages, hours, and real profits (i.e. the same variables that are used as observables in the main text). Second, we estimate the baseline SVAR model on those artificial data, and investigate whether the SVAR model is able to recover the true impulse responses and variance decompositions. Most parameter values are set equal to the median of their uniform distributions in the main text (the median values correspond well with typical values used in the literature), with one exception. The labor-capital elasticity η is set to 0.5, so that the true impulse responses illustrate factor complementarity. This is a common value, but lower than that estimated by Karabarbounis and Neiman (2014). In order to obtain a case of comparison, we scale the shock variances (the innovations are drawn from normal distributions) so that each shock's relative importance for the labor share, *40 quarters ahead*, resembles what we get with the SVAR estimated on real data.[3] Note, however, that the variance decompositions of other macroeconomic variables in the model generated data could be very different from the SVAR's variance decompositions shown in Figure 6.

First we assess the implications of *identification uncertainty* and *misspecification* embedded in the SVAR model.[4] To this end, we generate one large sample of 10,000 observations and estimate the SVAR on those data using the same econometric procedure and baseline restrictions described in the main text. We keep 1,000 draws that satisfy the sign restrictions. Figure D.1 Panel (a) compares the estimated impulse responses of the labor share from the SVAR model with the true model counterparts (i.e. those from the theoretical model, in red), as well as the estimated and true variance decompositions. The SVAR model does remarkably well in matching the responses of the labor share to the

[3]This scaling implies the following calibration: $\sigma_p^2 = 0.0056$, $\sigma_w^2 = 0.1406$, $\sigma_\upsilon^2 = 0.1139$, and $\sigma_{\alpha_k}^2 = 0.0506$.

[4]The SVAR is misspecified as we include only a subset of all the state variables in the SVAR model.

important shocks. Not only is it able to recover the correct signs of the responses after horizon 16, but in most cases the median labor share responses closely resemble those of the true model even quantitatively. There is one exception: The SVAR overstates the labor share decline after a wage markup shock. As a result, it also overestimates the importance of wage markups in the variance decomposition. We suspect that this discrepancy would be limited, had we added the long-run restriction that the wage-markup shocks cannot affect the labor share permanently.

Next, we add estimation or *small-sample uncertainty* to the picture, as in Canova and Paustian (2011). To this end we generate 200 samples of the same length as our baseline sample with real data for the US economy (144 observations). For each generated sample, we estimate the SVAR and keep 1,000 draws that satisfy the sign restrictions. This leaves us with 200,000 posterior draws. Results are reported in Figure D.1 Panel (b). Again, the SVAR model is able to match pretty well the true model, both qualitatively (in terms of signs) and quantitatively. If anything, it underestimates the importance of automation shocks when the sample is small, while we still tend to overestimate the importance of the wage markup. But in total, our identification procedure seems successful in recovering the main drivers of the labor share.

C.2 Alternative horizons, lag specifications, samples and priors

The baseline SVAR model presented in the main text is estimated using 4 lags, imposing the sign restrictions of Table 2 at a horizon 16 quarters ahead, using the variables in differences with a flat prior and on a quarterly sample that spans 1983Q1-2018Q3. Here we check the robustness of our results to changes in all of these specifications. For the sake of exposition, we present only the variance decompositions of the labor share corresponding to the different sensitivity checks, but the complete set of results is available upon request.

The results are shown in Figure D.2. The first two rows present the variance decompositions of the labor share using, respectively, different horizons and lag specifications. Changing the horizon at which the sign restrictions are imposed does not seem to affect the results presented in the previous section. The same is true if we use a different lag specification, although the role of price markups in explaining labor share fluctuations becomes slightly higher at long-run horizons when we include more feedback in the system. In the first two panels of the third row, we first expand the sample to go back to 1948Q1 and then restrict it from 1990Q1 onwards. Interestingly, price markups seem to have significantly less explanatory power in the first decades after the second World War. This evidence supports the view that firms' market power started to rise in the beginning of the 1980s and then accelerated in the 1990s and 2000s. The third panel, instead, presents the variance decomposition of the labor share using annual data in the estimation of the baseline model. The three panels of the fourth row refer to three different exercises. In the previous two, we use two different prior specifications: we estimate the VAR in levels using the dummy observation prior proposed by Sims and Zha (1998) and the priors for the long run (PLR) of Giannone, Lenza, and Primiceri (2019), which resemble our baseline specification in differences when infinitely tight. Differently from our baseline empirical framework, in these cases shocks do not necessarily have permanent effects. In the latter, we consider the median-target impulse responses proposed by Fry and Pagan

(2011). Overall, the results are in line with our baseline, although price markups seem slightly less relevant when we use the VAR in levels with the sum of coefficients prior or PLR.

C.3 Measurement of profits vs. other non-labor income

In our baseline SVAR for the nonfarm business sector, and for our extensions to the business and manufacturing sectors, we used data on output, wages, hours and deflators for the particular sector of interest, but profits for the whole economy (the after tax measure with IVA and CCadj from the BEA). The same measure of profits at the sectoral level is available for the nonfinancial corporate sector only. This is because this sector is the only sector for which the BLS has income-side data, and thus the only sector for which the BLS can break out non-labor payments into non-labor costs and profits. Non-labor payments are defined as "the excess of current-dollar output in an economic sector over corresponding labor compensation, and include non-labor costs, corporate profits and the profit-type income of proprietors. Non-labor costs include consumption of fixed capital, taxes on production and imports less subsidies, net interest and miscellaneous payments, and business current transfer payments". Thus, non-labor payments include both profits and capital income. In our theoretical model, this measure would correspond to the difference between output Y_t and labor costs $W_t L_t$. In what follows, we perform the following exercise: instead of including a measure of profits for the whole economy (or for the non-financial corporate sector), we use non-labor payments but impose the same sign restrictions as in the baseline setup. The sign restrictions are likely to hold even in this case, as shown in subsubsection A.4.1. Results are shown in Figure D.3 and Figure D.4. While the variance decomposition changes for some of the other variables in the system, the labor share results are robust to the use of non-labor payments instead of profits.

C.4 Additional medium- and long-run identification restrictions

In our baseline specification of Table 2, we have shown a minimum set of identifying restrictions that are sufficient to set apart the four shocks under consideration. The theoretical model, however, provides additional overidentifying restrictions that could potentially be exploited, both to check for robustness and as a mean to sharpen the inference. In this section, therefore, we investigate how the results are affected by various medium-run sign restrictions as well as long-run zero restrictions. All the restrictions we consider are consistent with the theoretical framework presented in the main text.

Figure D.5 shows the results when we add additional *medium-run sign restrictions* to those in the baseline. The first row adds the restriction that hours increase in response to a negative price markup shock. In the second row, we impose the restriction that hours increase in response to a positive investment-specific technology shock. The third row presents results when we combine restrictions: hours increase in response to both price markup and investment-specific technology shocks, profits increase in response to wage markup shocks and automation. All additional restrictions are imposed 16 quarters ahead. Our baseline results are largely confirmed. If anything, the evidence in favor of capital-labor complementarity becomes stronger once we restrict hours to increase in response to investment technology.

Next, we investigate whether the baseline results are robust to selected *long-run zero restrictions* as well. Some of these have a well-grounded theoretical justification and may help us to further sharpen the identification of shocks. Here we add the following long-run zero restrictions to the baseline identification scheme summarized in Table 2: (i) only the price markup shock can have a permanent effect on the profit share $\mathcal{D}_t/Y_t$, and (ii) the real wage W_t, and labor productivity Y_t/L_t, cannot be permanently affected by wage markup shocks. The first restriction is consistent with standard models of monopolistic competition, but does not necessarily hold if we consider more strategic interaction in goods markets.[5] The second restriction is based on the neo-classical assumption that real wages and labor productivity are fully determined in the long run by technical change and other supply-side factors. Our theoretical framework in Section 2 illustrates this assumption: factor prices and inputs—hence the labor income share—are pinned down by the production function as well as equations 9 and 10.

Figure D.6 summarize the results when we add the zero restrictions just described to the baseline identification scheme. The restrictions are implemented following Arias, Rubio-Ramírez, and Waggoner (2018). Long run restrictions on wages and labor productivity are shown in the first row, long-run restrictions on the profit share in the second row, and finally the combination of both in the last row. Qualitatively, our main results are robust to these additional restrictions. Quantitatively, however, the inclusion of long-run restrictions facilitate a tighter identification of the dynamics. The labor share responses to investment-specific technology shocks, for example, are more precisely estimated with our zero restrictions on the profit share. Zero restrictions on real wages and labor productivity, in contrast, imply a greater role for wage markup shocks at short horizons. But automation and price markup shocks remain the main drivers of the labor income share across all specifications.

C.5 A further investigation of labor market variables

C.5.1 Wage bargaining, labor participation and employment

It is well known that wage mark-up shocks and labor supply shocks are observationally equivalent in the standard neoclassical model (and thus also in our baseline model). However, the two shocks propagate differently in models with search and matching frictions and endogenous participation, as shown by Foroni, Furlanetto, and Lepetit (2018). In fact, a decrease in the bargaining power of workers leads to a negative co-movement between output and participation while an exogenous increase in labor supply leads to a positive co-movement between output and participation for a broad range of parameterizations. In addition, the former shock leads to a negative co-movement between output and unemployment while the latter leads to a positive co-movement. Therefore, using data on participation and/or unemployment makes it possible to disentangle the two shocks in our SVAR, as done in Foroni et al. (2018). To this end we extend the baseline SVAR model by identifying a labor supply shock using data on the participation rate according to the restrictions presented in Table C.1.[6] Results are summarized in Figure D.7 and

[5] Since we assume monopolistic competition among atomistic firms, the profit share becomes a function solely of the price markup, which is exogenous.

[6] Results are very similar when using data on the unemployment rate and are available upon request.

Table C.1: Sign Restrictions

	Wage Bargaining	Automation $\alpha_k \uparrow$	Firms' $\mathcal{M}_p \downarrow$	IST $\Upsilon \uparrow$	Labor Supply
GDP	+	+	+	+	+
Wages	-	-	+	+	-
Hours	+	-	/	/	+
Profits	/	/	-	+	/
Participation	-	/	/	/	+

Note: The restrictions are imposed at quarter 16.

Figure D.8. Although both labor market shocks turn out to play a minor role for the labor share dynamics, they are important drivers of hours and participation (and to some extent also GDP and wages).

C.5.2 Routine and non-routine employment

Finally, we disentangle the separate effects on routine- and non-routine labor, respectively. From a theoretical point of view, it seems reasonable that routine employment, which is more exposed to automation of production tasks, may react stronger to the automation shock. In order to check this hypothesis, we use data on routine and non-routine per capita employment, obtained from Zhang (2019), which in turn updates the data constructed by Jaimovich and Siu (2020). Results are shown in Figure D.9. Interestingly, while wage markup shocks are important drivers of both routine and non-routine employment, automation has a much stronger negative effect on routine employment at short horizons. Also price markups are more important for routine than non-routine employment. When compared with the baseline variance decomposition of hours in the main text, non-routine employment has a very similar behavior to hours. Routine employment, in contrast, features a much larger role on impact for automation and, over the entire horizon, for price markups.

D Additional figures

Figure D.1: Monte Carlo results with artificial data generated from the theoretical model

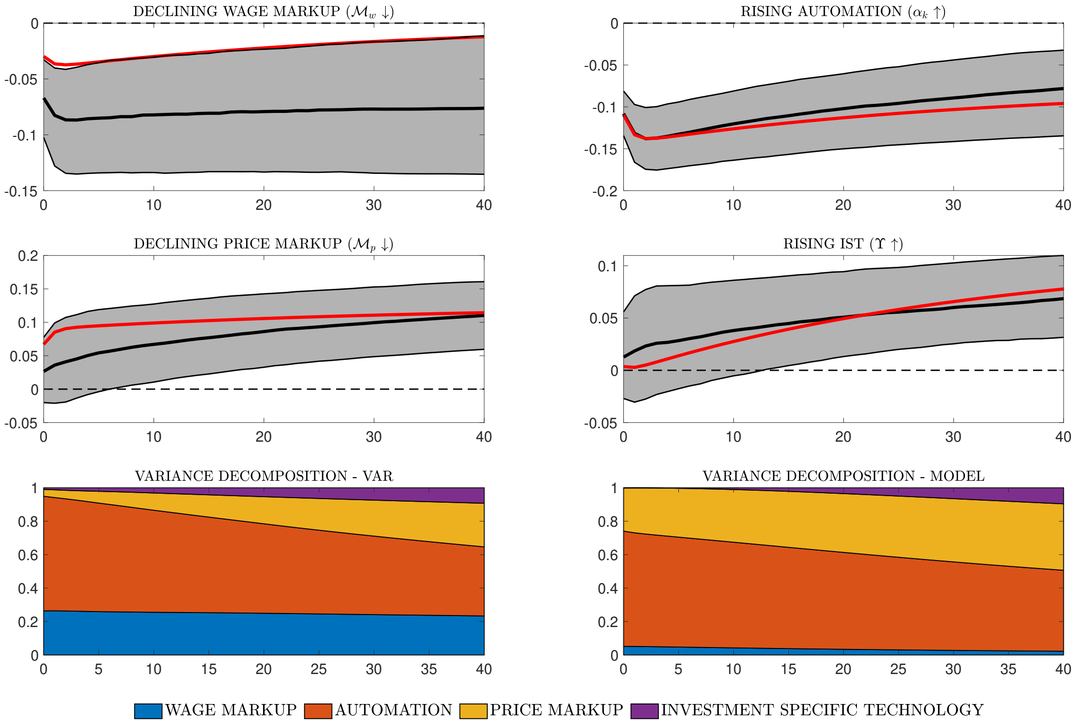

(a) One large sample of 10,000 observations

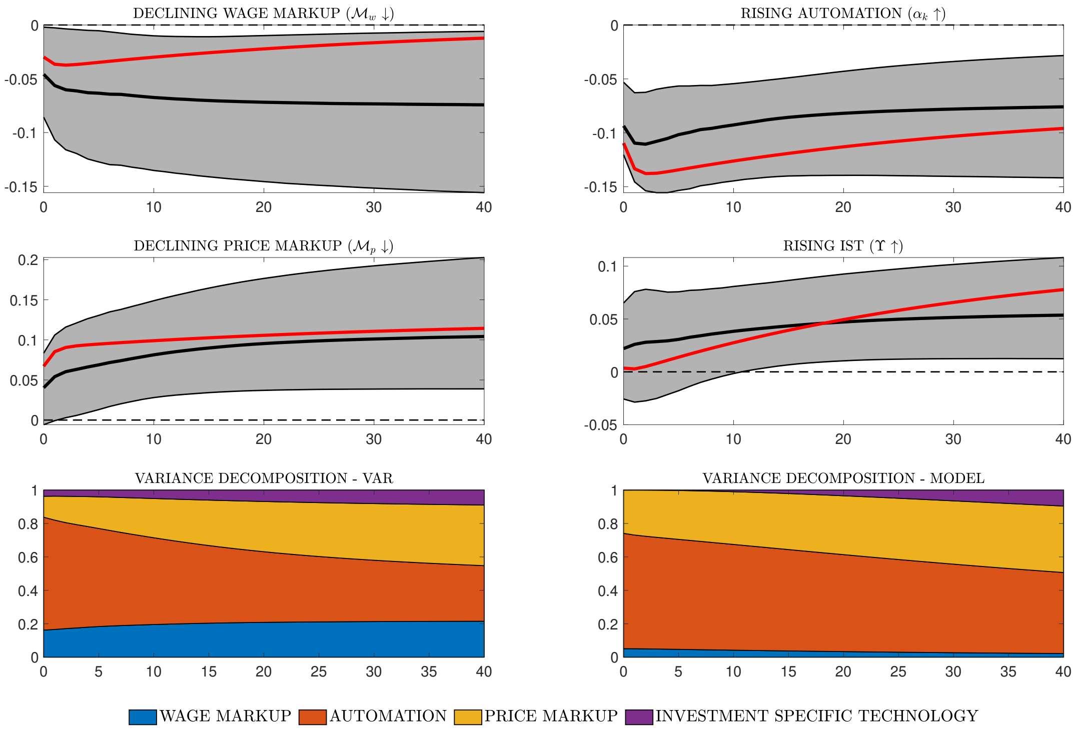

(b) 200 small samples, each of 144 observations

Figure D.2: Sensitivity checks

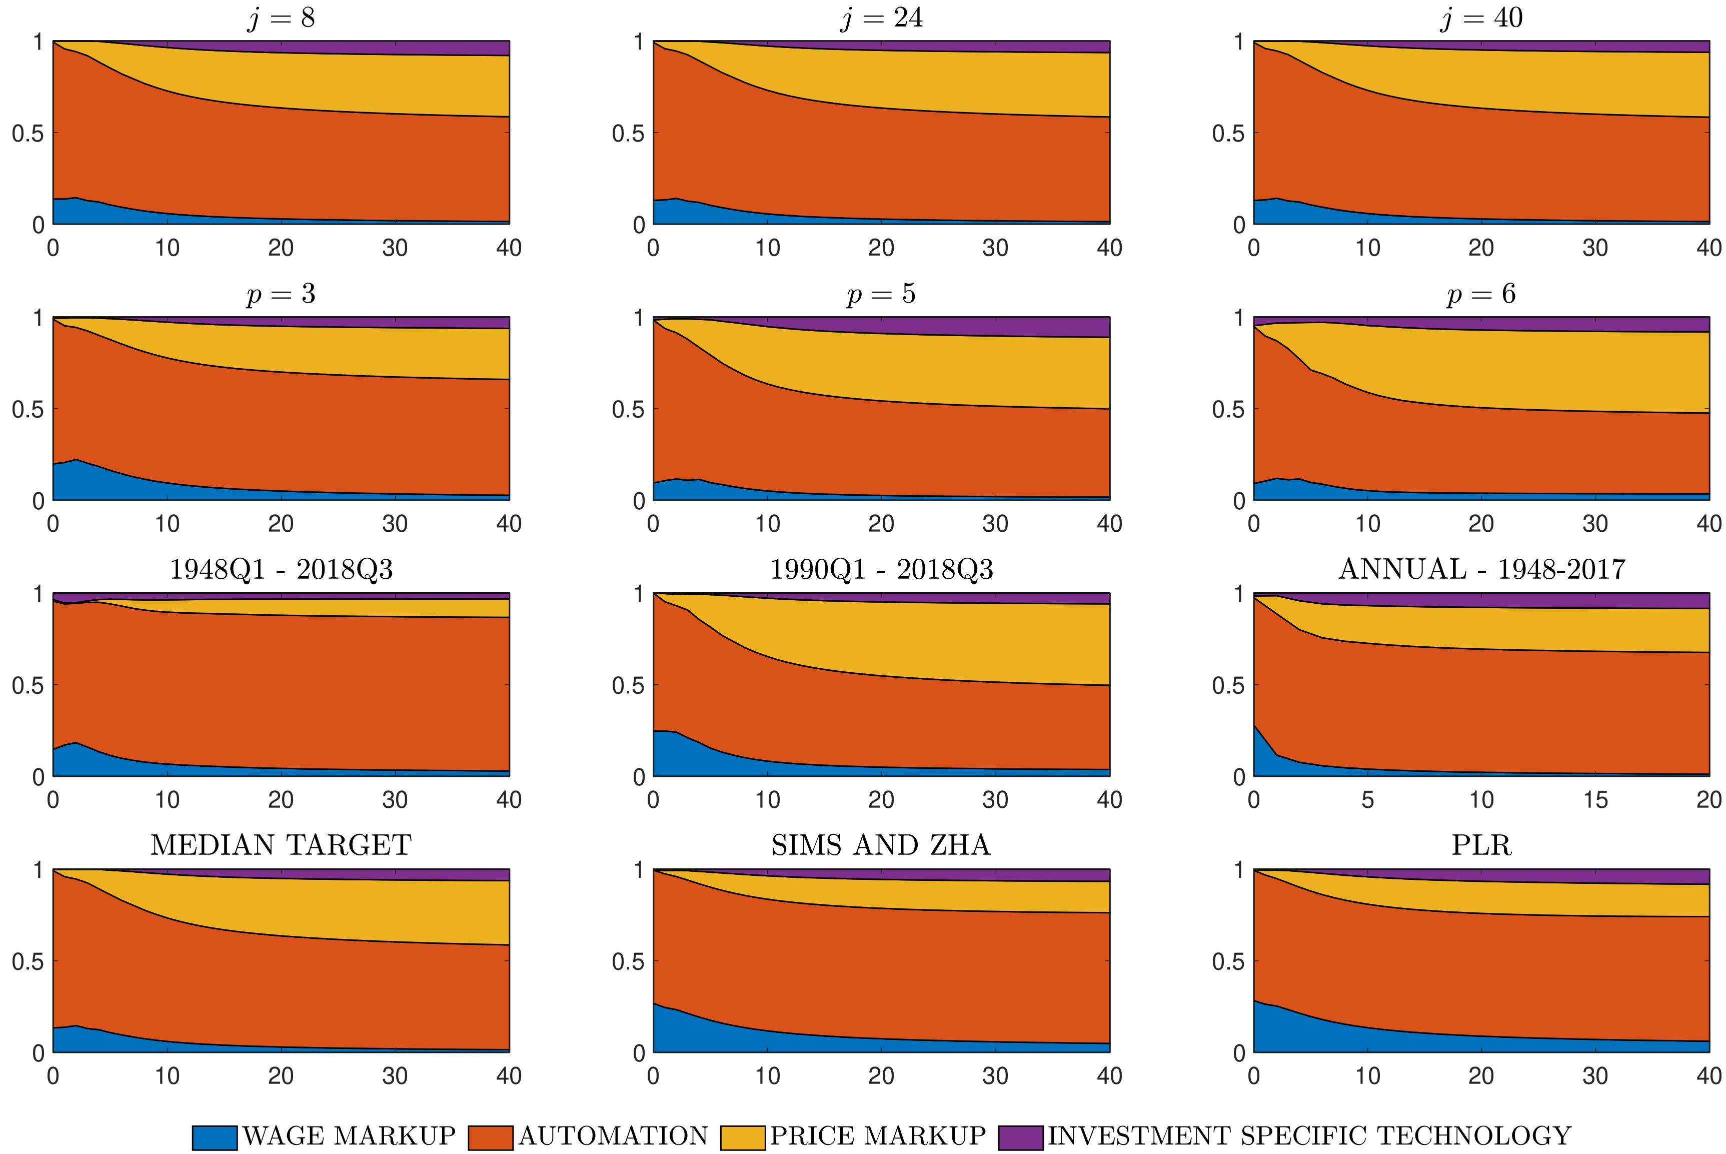

Note: The colored areas represent the point-wise median contributions of each identified shock to the forecast error variance of the labor share (in levels) at horizons $j = 0, 1, \ldots, 40$ using the baseline identifying restrictions.

Figure D.3: Impulse responses with non-labor payments

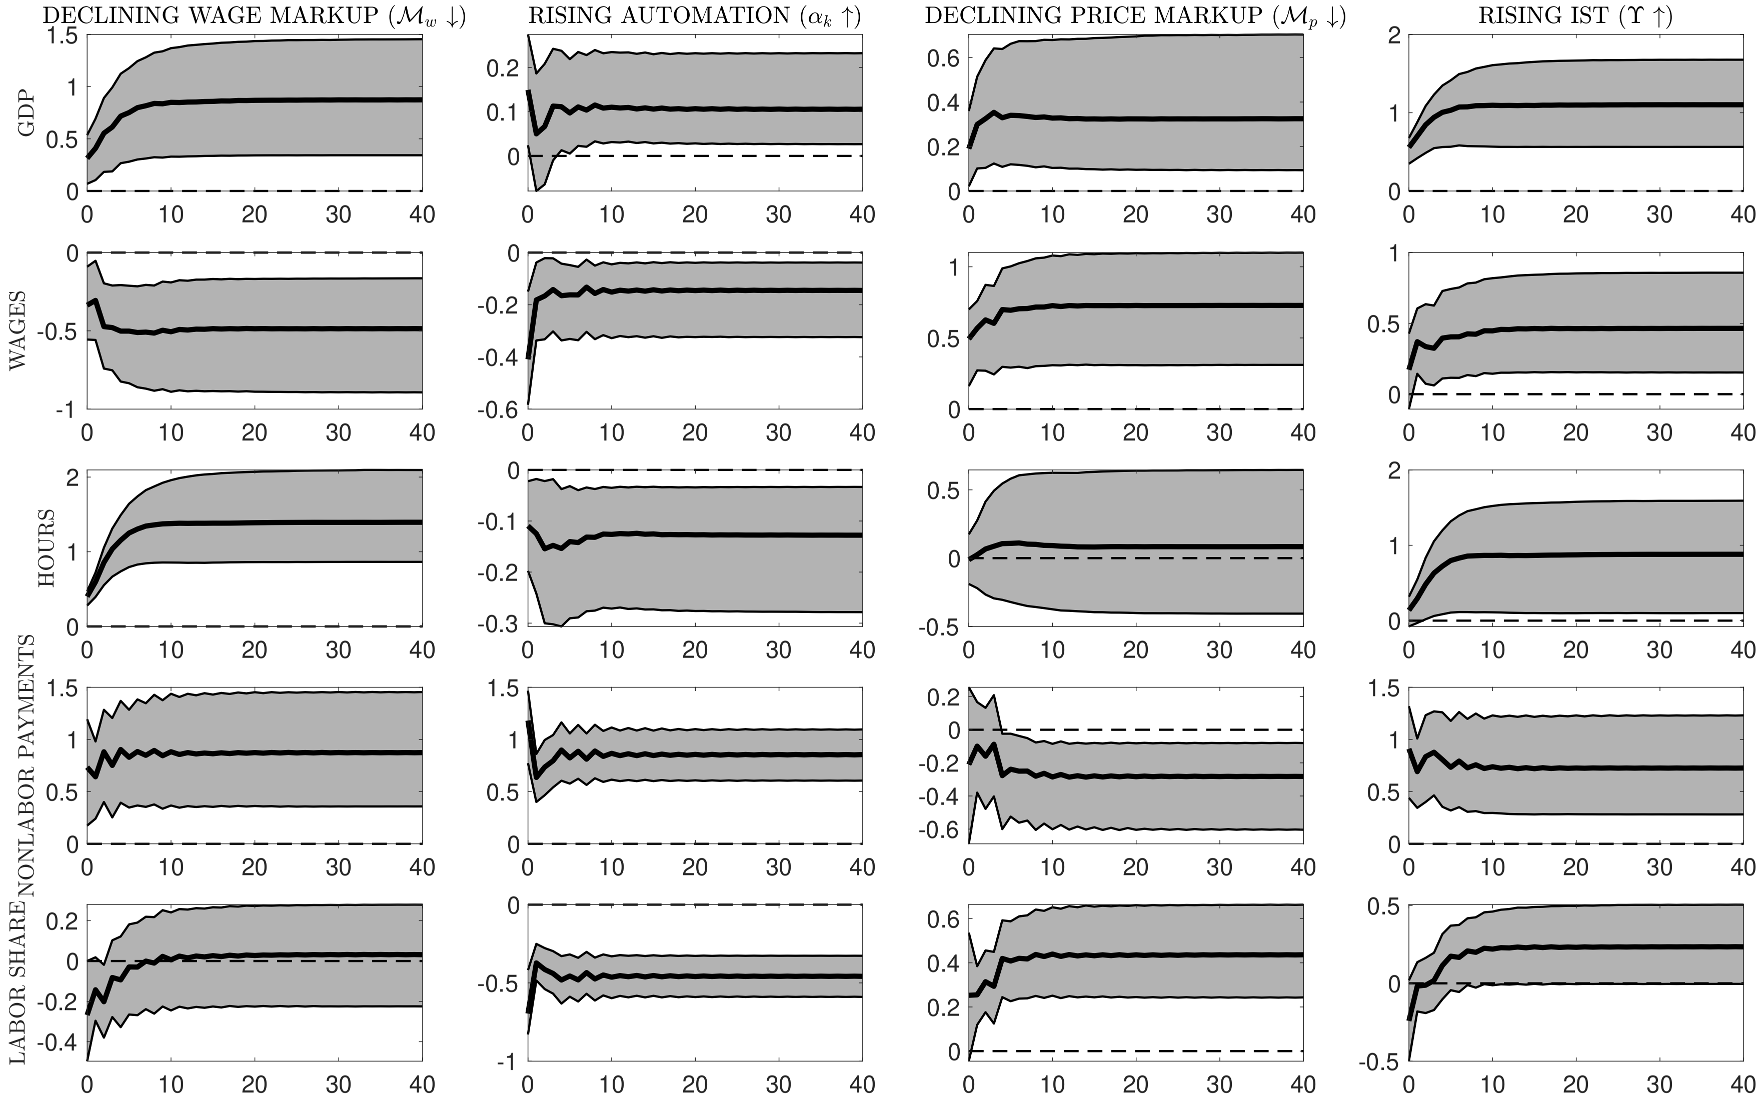

Note: Posterior distributions of cumulated impulse responses to an estimated shock of one standard deviation using the baseline identifying restrictions. Median (solid line) and 68% probability density intervals (shaded area) based on 10,000 draws. The median and the percentiles are defined at each point in time. The colored areas represent the point-wise median contributions of each identified shock to the forecast error variance of the labor share (in levels) at horizons $j = 0, 1, \ldots, 40$ using the baseline identifying restrictions.

Figure D.4: Variance decompositions with non-labor payments

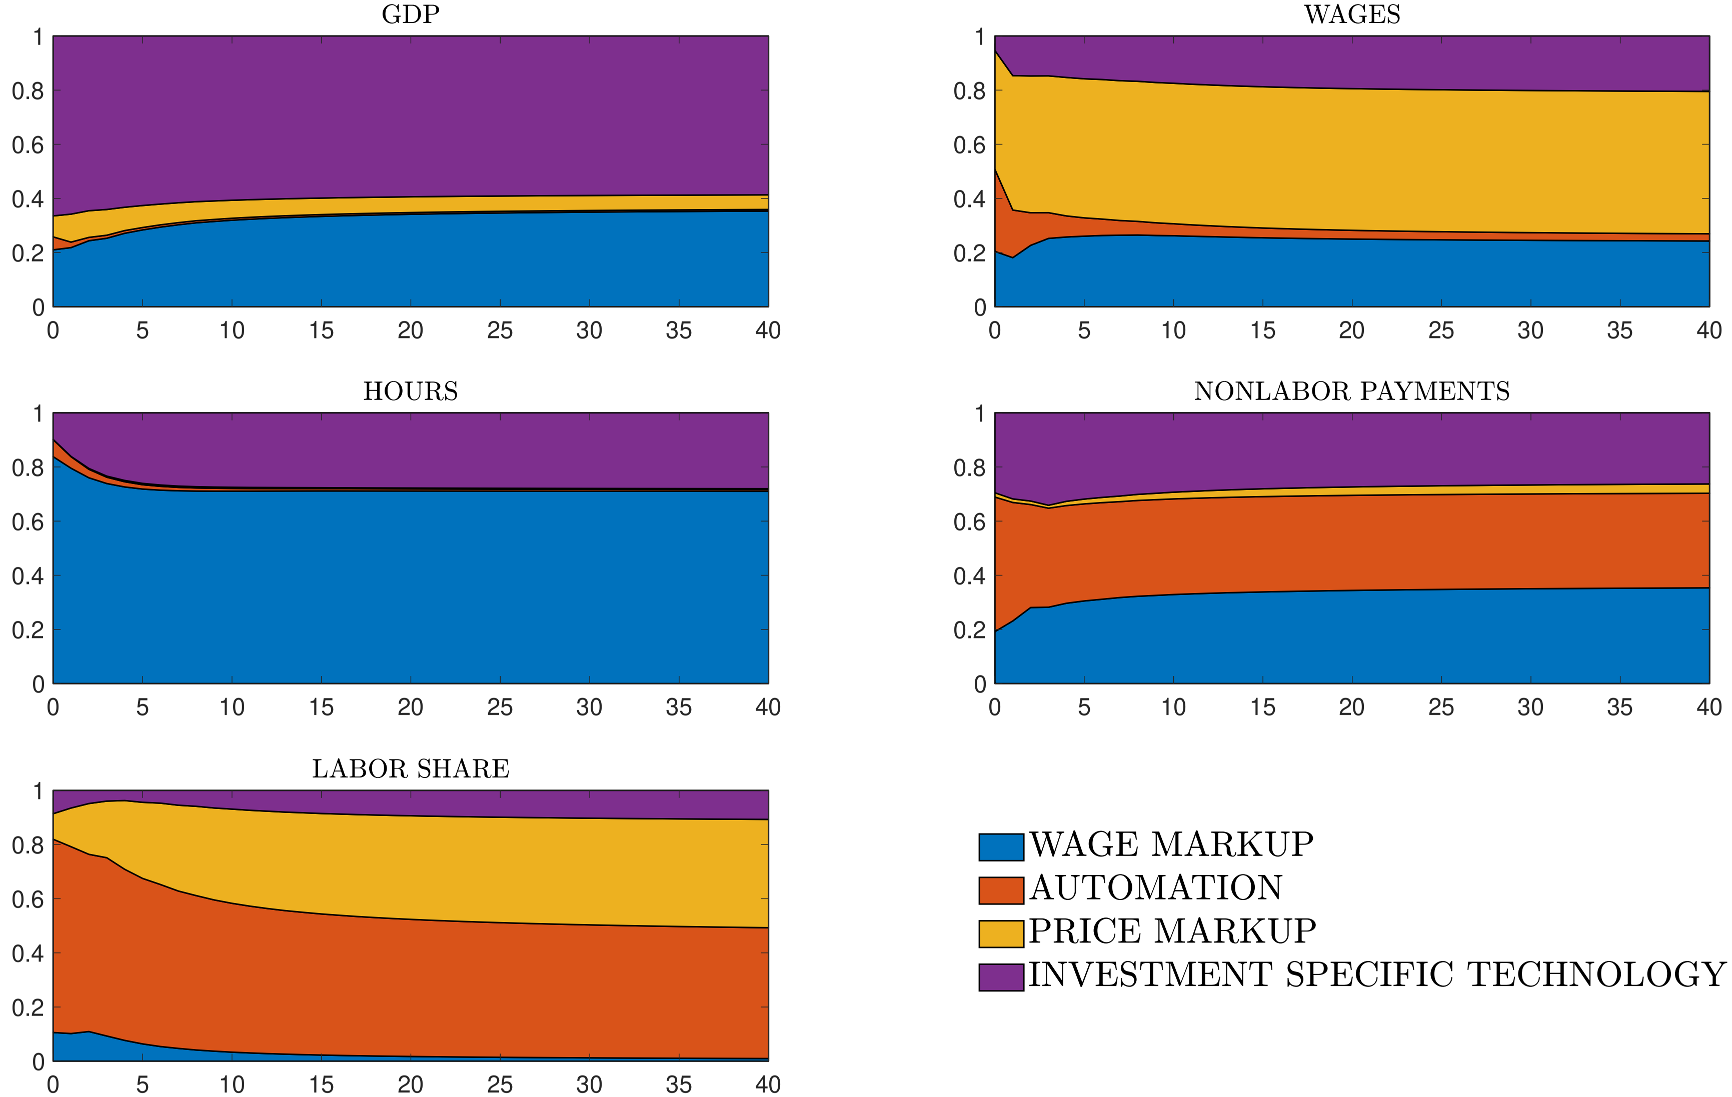

Note: Posterior distributions of cumulated impulse responses to an estimated shock of one standard deviation using the baseline identifying restrictions. Median (solid line) and 68% probability density intervals (shaded area) based on 10,000 draws. The median and the percentiles are defined at each point in time. The colored areas represent the point-wise median contributions of each identified shock to the forecast error variance of the labor share (in levels) at horizons $j = 0, 1, \ldots, 40$ using the baseline identifying restrictions.

Figure D.5: Additional medium-run sign restrictions

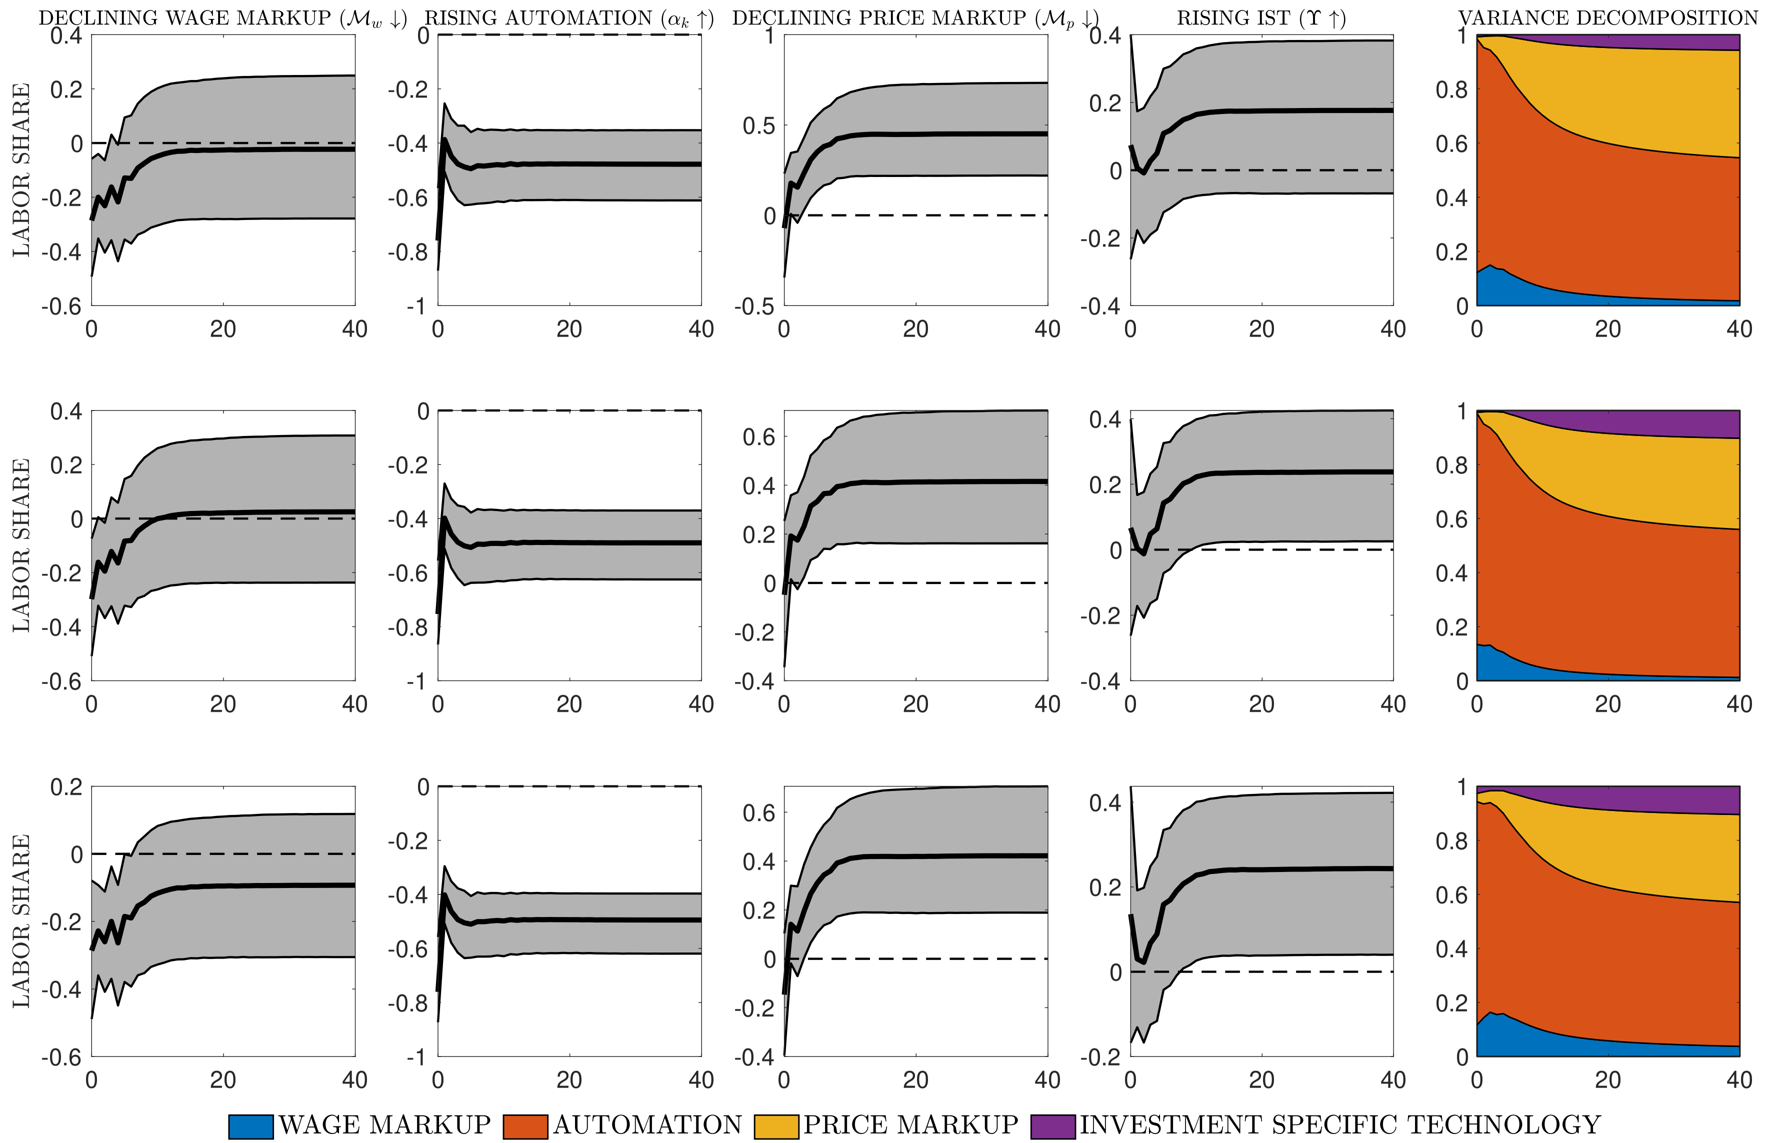

Note: Posterior distributions of cumulated impulse responses to an estimated shock of one standard deviation using the baseline identifying restrictions. Median (solid line) and 68% probability density intervals (shaded area) based on 10,000 draws. The median and the percentiles are defined at each point in time. The colored areas represent the point-wise median contributions of each identified shock to the forecast error variance of the labor share (in levels) at horizons $j = 0, 1, \ldots, 40$ using the alternative identifying restrictions.

Figure D.6: Additional long-run zero restrictions

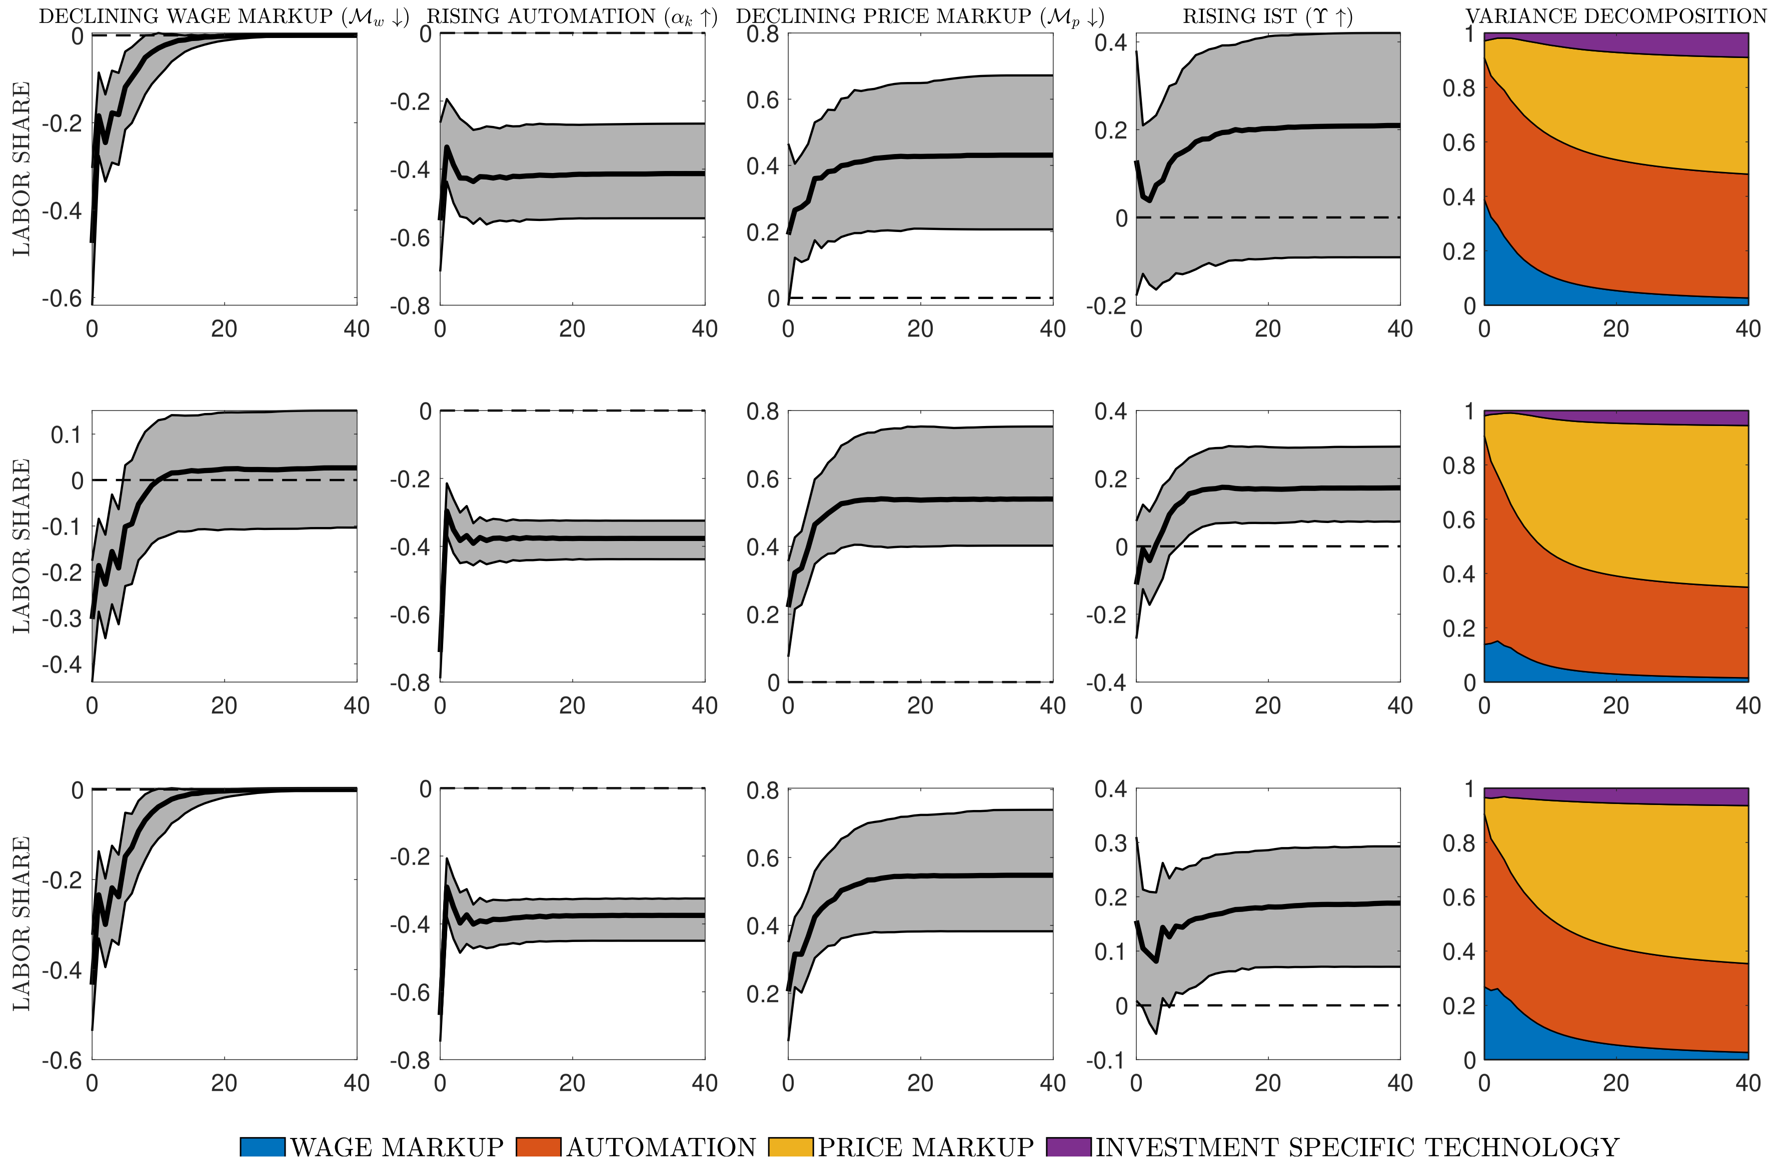

Note: Posterior distributions of cumulated impulse responses to an estimated shock of one standard deviation using the baseline identifying restrictions. Median (solid line) and 68% probability density intervals (shaded area) based on 1,000 draws. The median and the percentiles are defined at each point in time. The colored areas represent the point-wise median contributions of each identified shock to the forecast error variance of each variable (in levels) at horizons $j = 0, 1, \ldots, 40$ using the alternative identifying restrictions.

Figure D.7: Impulse responses with participation

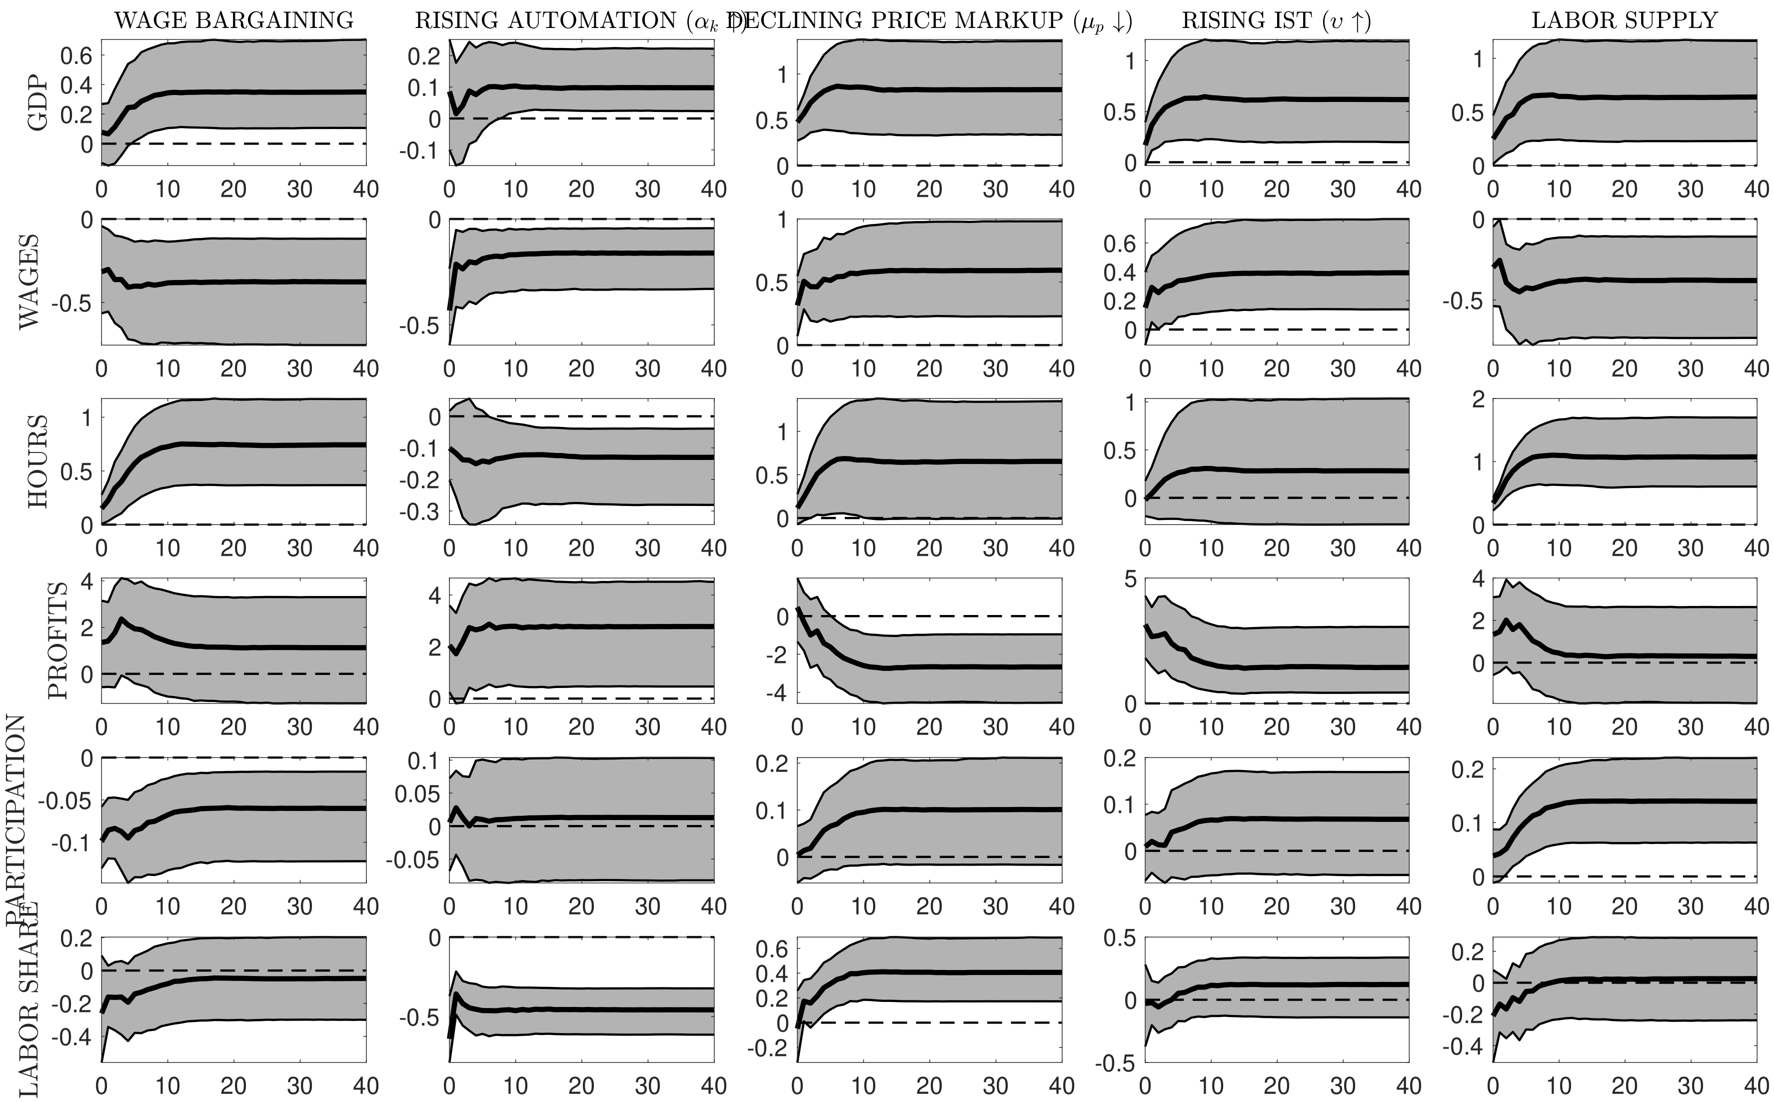

Note: Posterior distributions of cumulated impulse responses to an estimated shock of one standard deviation using the baseline identifying restrictions. Median (solid line) and 68% probability density intervals (shaded area) based on 10,000 draws. The median and the percentiles are defined at each point in time. The colored areas represent the point-wise median contributions of each identified shock to the forecast error variance of the labor share (in levels) at horizons $j = 0, 1, \ldots, 40$ using the baseline identifying restrictions.

Figure D.8: Variance decompositions with participation

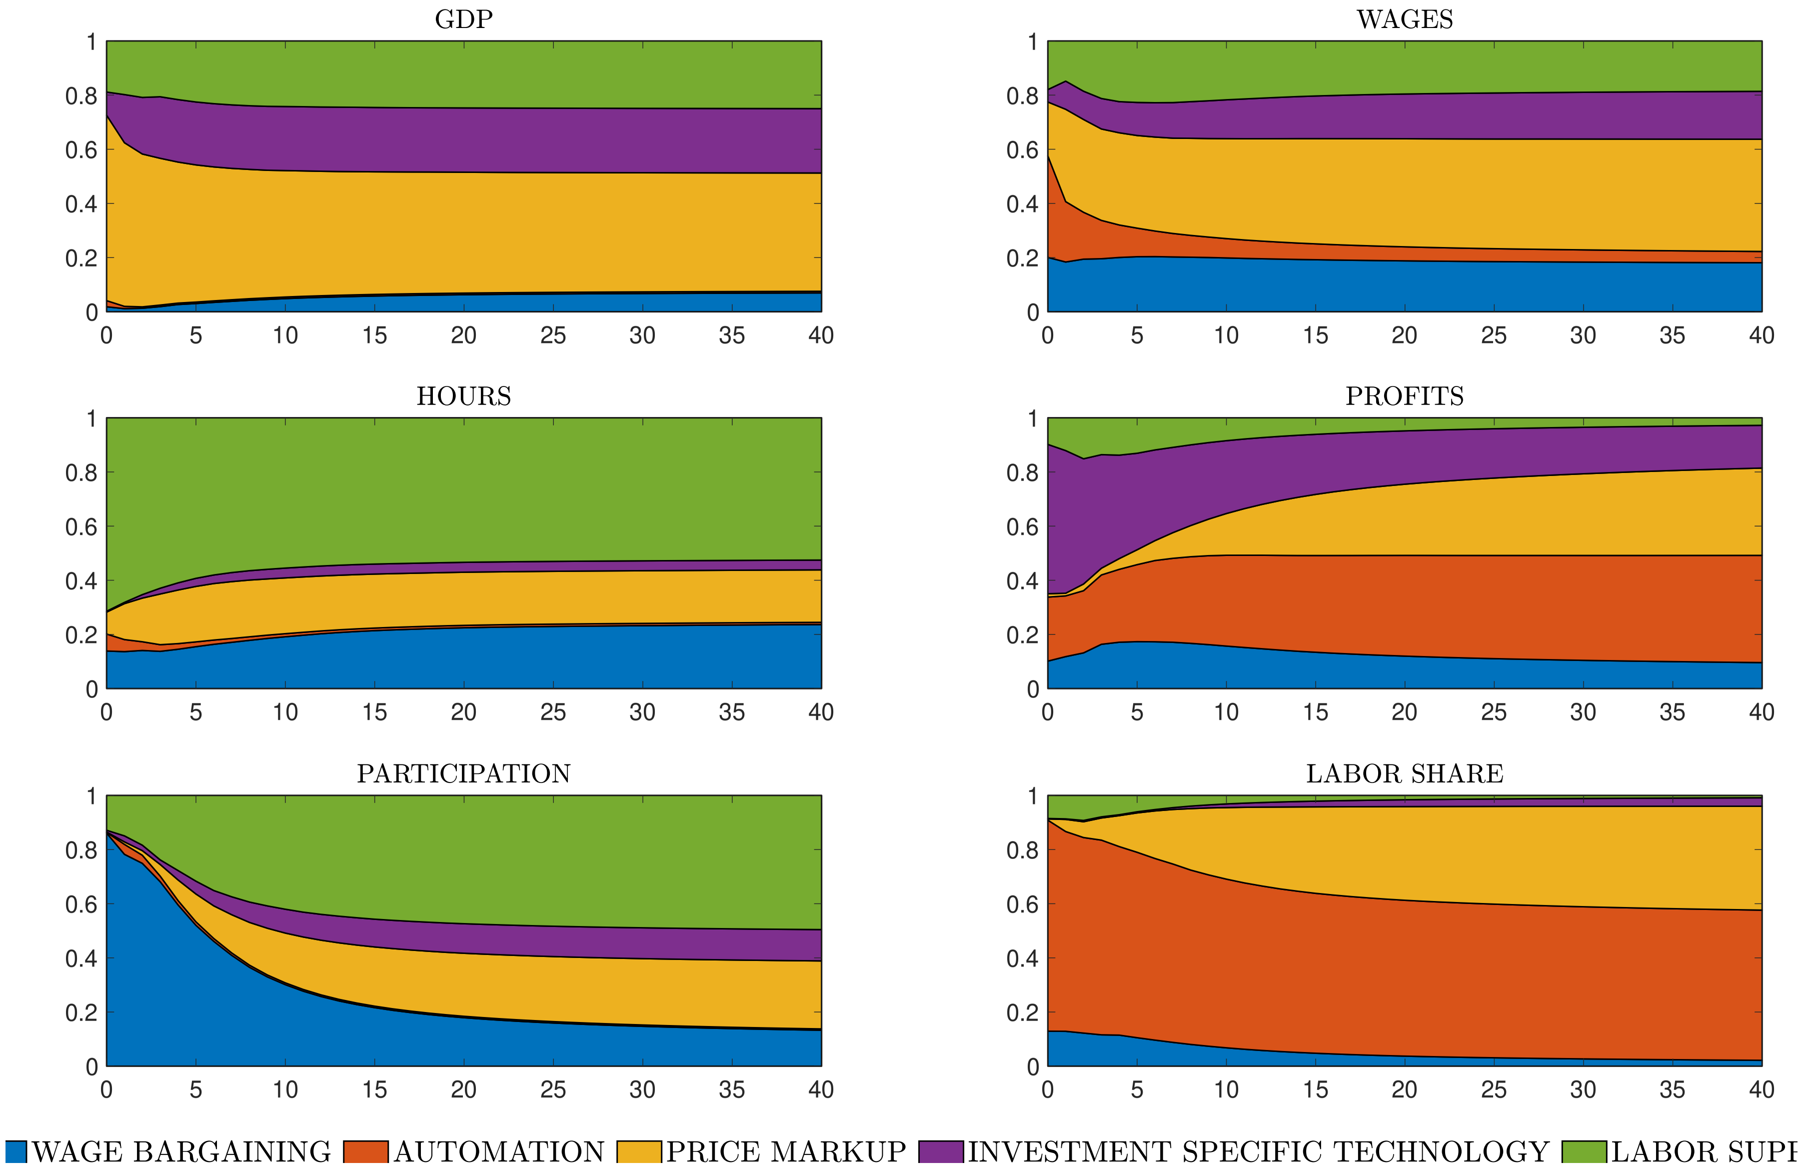

Note: Posterior distributions of cumulated impulse responses to an estimated shock of one standard deviation using the baseline identifying restrictions. Median (solid line) and 68% probability density intervals (shaded area) based on 10,000 draws. The median and the percentiles are defined at each point in time. The colored areas represent the point-wise median contributions of each identified shock to the forecast error variance of the labor share (in levels) at horizons $j = 0, 1, \ldots, 40$ using the baseline identifying restrictions.

Figure D.9: Routine- versus non-routine employment

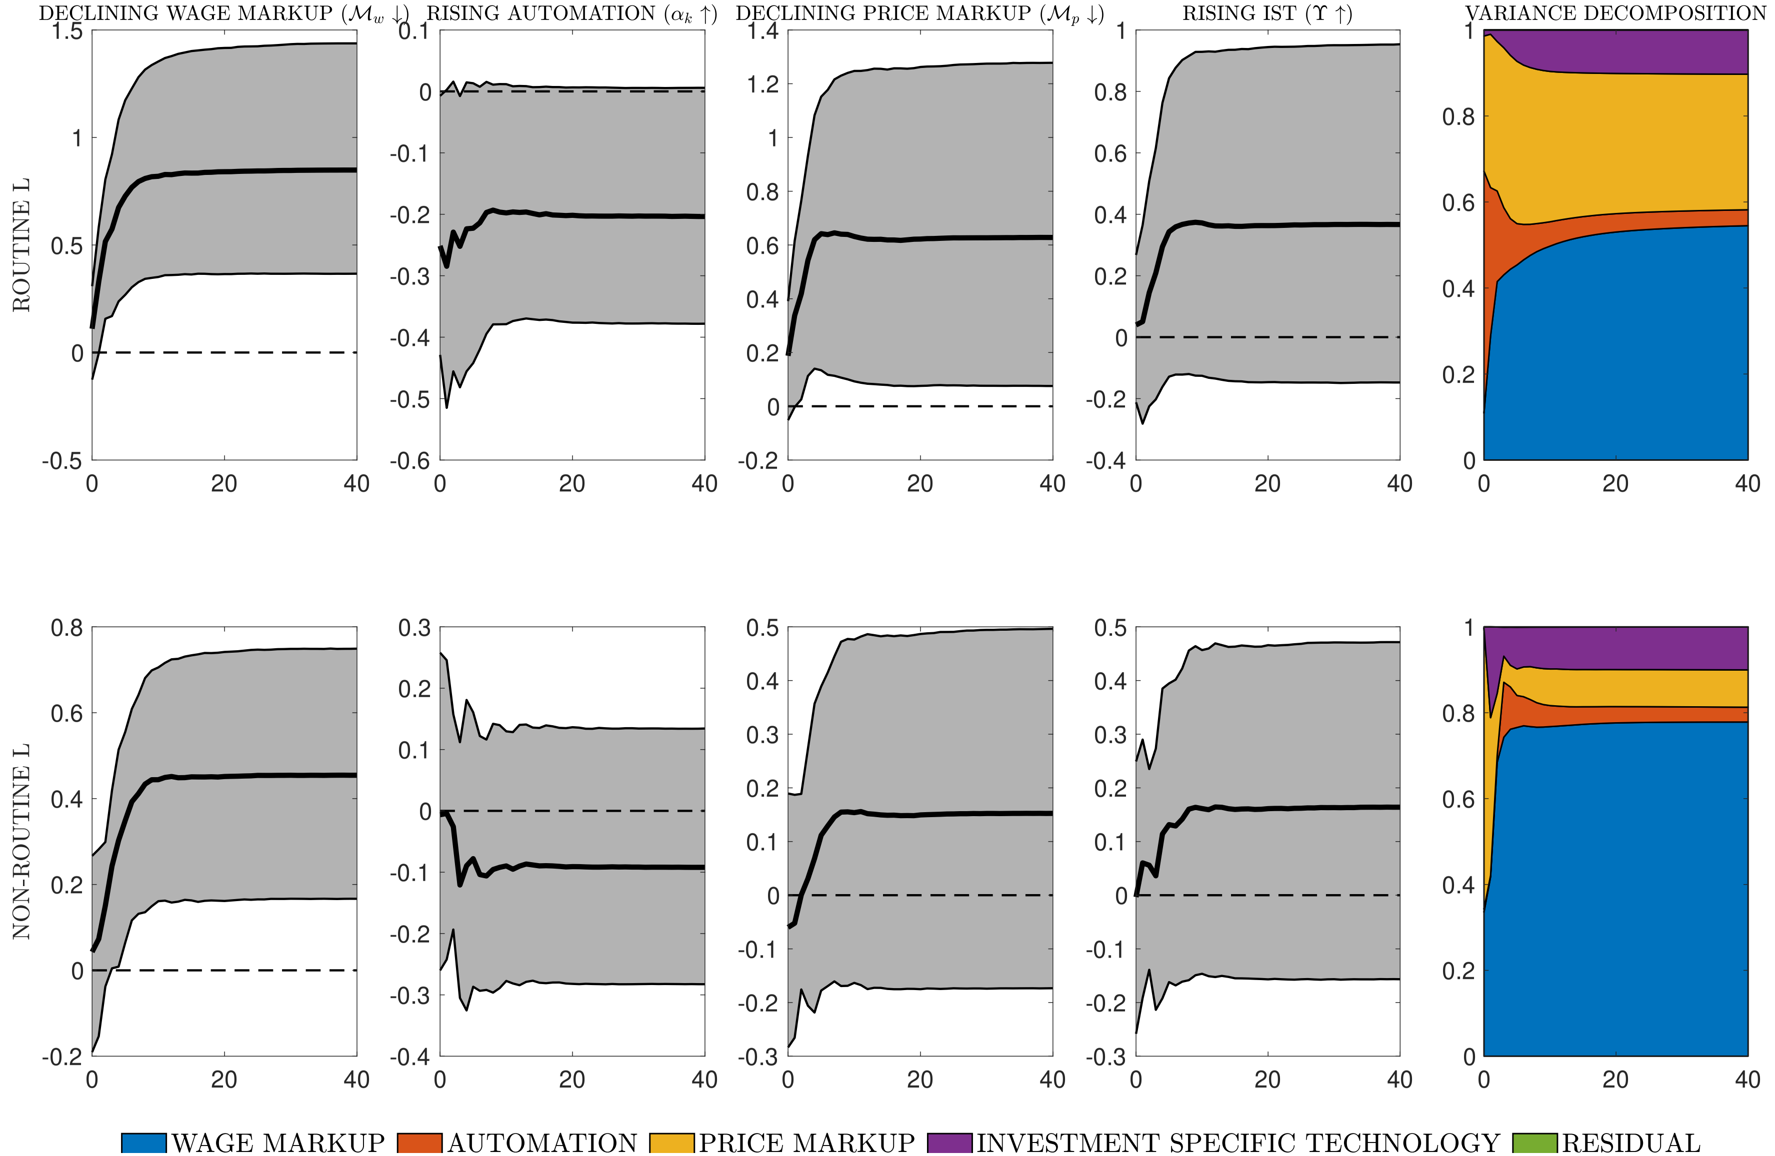

Note: Posterior distributions of cumulated impulse responses to an estimated shock of one standard deviation using the baseline identifying restrictions. Median (solid line) and 68% probability density intervals (shaded area) based on 10,000 draws. The median and the percentiles are defined at each point in time. The colored areas represent the point-wise median contributions of each identified shock to the forecast error variance of each variable (in levels) at horizons $j = 0, 1, \ldots, 40$ using the alternative identifying restrictions.

REFERENCES

Acemoglu, D. and P. Restrepo (2018). The Race Between Machine and Man: Implications of Technology for Growth, Factor Shares and Employment. *American Economic Review 108*(6), 1488–1542.

Acemoglu, D. and P. Restrepo (2020). Robots and Jobs: Evidence from US Labor Markets. *Journal of Political Economy 128*(6), 2188–2244.

Arias, J. E., J. F. Rubio-Ramírez, and D. F. Waggoner (2018). Inference based on structural vector autoregressions identified with sign and zero restrictions: Theory and applications. *Econometrica 86*(2), 685–720.

Canova, F. and M. Paustian (2011). Business cycle measurement with some theory. *Journal of Monetary Economics 58*(4), 345–361.

Cantore, C. and P. Levine (2012). Getting normalization right: Dealing with dimensional constants in macroeconomics. *Journal of Economic Dynamics and Control 36*(12), 1931–1949.

Foroni, C., F. Furlanetto, and A. Lepetit (2018). Labor Supply Factors and Economic Fluctuations. *International Economic Review 59*(3), 1491–1510.

Fry, R. and A. Pagan (2011). Sign Restrictions in Structural Vector Autoregressions: A Critical Review. *Journal of Economic Literature 49*(4), 938–960.

Giannone, D., M. Lenza, and G. E. Primiceri (2019). Priors for the Long Run. *Journal of the American Statistical Association 114*(526), 565–580.

Jaimovich, N. and H. E. Siu (2020). Job Polarization and Jobless Recoveries. *The Review of Economics and Statistics 102*(1), 129–147.

Karabarbounis, L. and B. Neiman (2014). The Global Decline of the Labor Share. *The Quarterly Journal of Economics 129*(1), 61–103.

Karabarbounis, L. and B. Neiman (2019). Accounting for Factorless Income. *NBER Macroeconomics Annual 33*(1), 167–228.

Martinez, J. (2019). Automation, Growth and Factor Shares. Manuscript.

Paustian, M. (2007). Assessing Sign Restrictions. *The B.E. Journal of Macroeconomics 7*(1), 1–33.

Rubio-Ramírez, J. F., D. F. Waggoner, and T. Zha (2010). Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference. *Review of Economic Studies 77*(2), 665–696.

Sims, C. A. and T. Zha (1998). Bayesian Methods for Dynamic Multivariate Models. *International Economic Review 39*(4), 949–968.

Zhang, W. (2019). Deciphering the causes for the post-1990 slow output recoveries. *Economics Letters 176*(C), 28–34.