

The Distributional Effects of Oil Supply News shocks*

Theo Drossidis¹, Haroon Mumtaz², and Angeliki Theophilopoulou³

²Queen Mary University of London

¹Brunel University

³Brunel University

This version: May 23, 2024.

Abstract

This paper uses high frequency data on the distribution of US income to investigate the heterogeneous effects of oil supply news shocks. Using a FAVAR with an external instrument, We show that these shocks have large negative effects on the left and right tail of the distribution. For low income individuals, the effect is driven by a decline in wages and proprietor's income, while a fall in corporate profits and interest income drives the effect for affluent individuals.

JEL Classification: C32, E32, Q54

Keywords: Oil shock; income inequality; FAVAR; External instrument identification.

* We thank an anonymous referee, the editor Eric Young and participants at the 2024 RCEA conference for useful comments.

1 Introduction

The recent energy crisis has again focused attention on oil prices. A large empirical literature has established the importance of oil shocks for economic fluctuations. In a recent important contribution, [Känzig \(2021\)](#) uses a narrative approach to identify oil supply news shocks, i.e. unexpected fluctuations in current and future oil supply, and shows that these disturbances have a sizeable effect on US industrial production and CPI inflation. By applying this identification approach, [Känzig \(2021\)](#) builds on a large literature that reaches similar conclusions for oil market disturbances (see for e.g. [Hamilton \(2003\)](#), [Baumeister and Kilian \(2016\)](#), [Caldara *et al.* \(2019\)](#)).

One common feature of this literature is the focus on aggregate macroeconomic outcomes. In this paper, we exploit high-frequency data on the distribution of income and its components for the US to investigate the distributional effects of oil supply news shocks. We use a factor augmented VAR (FAVAR) to jointly model the oil market, macroeconomic variables and income in deciles of distribution. The oil supply shock is identified using the external instrument approach of [Känzig \(2021\)](#). The analysis leads to three key findings:

1. While an adverse oil supply shock has the largest effect at the left tail of the income distribution, the income of affluent individuals also declines relative to the median.
2. For individuals at the left tail, the decline in income is driven by a sharp fall in wages and proprietor's income.
3. At the right tail, income declines as the shock pushes down components of capital income such as interest income and corporate profits.

Our paper is related to [Berisha *et al.* \(2021\)](#) who examine the impact of oil production and dependency on the annual Gini coefficient for US states in a reduced-form setting. Our analysis is an extension of [Berisha *et al.* \(2021\)](#), as we identify an oil supply shock taking into account the effect of news and examine how the distribution and components of income are affected rather than focusing on one measure of inequality. Our paper is also related to [Del Canto *et al.* \(2023\)](#) who examine the effect of oil supply shocks on households with differing levels of educational attainment.

The paper is organised as follows: The data and empirical model is described in Section 2 while Section 3 describes the main results.

2 Empirical model and data

To estimate the impact of oil supply news on income for different groups of the population, we use a factor-augmented VAR (FAVAR) model. The model is defined by the VAR:

$$Y_t = c + \sum_{j=1}^P \beta_j Y_{t-j} + u_t \quad (1)$$

where $Y_t = \begin{pmatrix} z_t \\ \hat{F}_t \end{pmatrix}$, where z_t denotes a set of variables pertaining to the oil market: the real price of oil, world oil production and world oil inventories. \hat{F}_t represents factors that summarize information in a panel of macroeconomic and financial series and the individual-level data on income and its components, described below. The factors are estimated using the non-stationary factor model of [Barigozzi *et al.* \(2021\)](#). Denote X_t as the $(M \times 1)$ data matrix that contains the panel of macroeconomic and financial series that summarize information about the economy, and also includes income data at the dis-aggregated level. The observation equation of the FAVAR is defined as:

$$X_t = c + b\tau + \Lambda F_t + \xi_t \quad (2)$$

where c is an intercept, τ denotes a time-trend, F_t are the R non-stationary factors, Λ is a $M \times R$ matrix of factor loadings, and ξ_t are idiosyncratic components that are allowed to be $I(1)$ or $I(0)$. Note that the idiosyncratic components corresponding to the disaggregated income data can be interpreted as shocks that are specific to those groups and also capture possible measurement errors. The shocks to equation (1) represent macroeconomic or common shocks. It is the response to these common shocks that is relevant to our investigation. This ability to estimate the effect of macroeconomic shocks while taking into account idiosyncratic errors via equation (2) is a key advantage of the FAVAR over a VAR, where these two sources of fluctuations are harder to separate (see [De Giorgi and Gambetti \(2017\)](#) and [Cantore *et al.* \(2023\)](#)). Moreover, by incorporating a large data set, the FAVAR reduces the problem of information deficiency (see e.g. [Forni and Gambetti \(2014\)](#)) and shock deformation (see e.g. [Canova and Ferroni \(2022\)](#)).

2.1 Identification of the Oil supply news shock

To identify the oil supply news shock, we use an external instrument approach (see e.g. [Stock and Watson \(2008\)](#) and [Mertens and Ravn \(2013\)](#)). The residuals u_t are related to

structural shocks ε_t via:

$$u_t = A_0 \varepsilon_t \quad (3)$$

where $\text{cov}(u_t) = \Sigma = A_0 A_0'$. We denote the shock of interest as ε_{1t} and the remaining disturbances as ε_{-t} . Identification of ε_{1t} is based on the instrument m_t that satisfies the relevance and exogeneity conditions: $\text{cov}(m_t, \varepsilon_{1t}) = \alpha \neq 0$ and $\text{cov}(m_t, \varepsilon_{-t}) = 0$. As discussed in the technical appendix, these conditions can be combined with the covariance restrictions to obtain an estimate of the relevant column of the contemporaneous impact matrix A_0 . In our benchmark model, we employ the instrument constructed by [Känzig \(2021\)](#) which is based on the variation in oil futures prices around OPEC announcements. [Känzig \(2021\)](#) provides evidence to suggest that the instrument is relevant and exogenous.

2.2 Data and Estimation

As noted above, X includes both aggregate and individual-level data. The aggregate data is taken from the [Fred-MD](#) database. This consists of 134 variables covering industrial production, employment, consumer prices, asset prices, interest rates, exchange rates and spreads¹.

The data on individual level income is obtained from the [Real Time Inequality](#) database constructed by [Blanchet et al. \(2022\)](#). [Blanchet et al. \(2022\)](#) construct *monthly* distributions of income, wealth and their components by statistically matching the annual distributional national accounts of [Piketty et al. \(2017\)](#) with the current population survey and the survey of consumer finances in order to incorporate demographic information. They then construct monthly variables by re-scaling each component of income and using information on the distribution of wages from monthly and quarterly survey and administrative data. We use *factor income* as our benchmark income measure. Factor income is the sum of labour and capital income.²

We define 10 groups based on the deciles of factor income: P_1, P_2, \dots, P_{10} . P_1 includes individuals that fall below the tenth percentile of factor income, P_2 denotes individuals above the tenth percentile but below the twentieth percentile and so on. We construct average factor income, capital and labour income in each of these groups. In addition, we calculate the average of the main components of capital and labour income in each group. All of these income variables are deflated by the national income deflator and included in

¹A full list of these variables is available on [FRED-MD website](#).

²As in [Blanchet et al. \(2022\)](#) labour income is defined as the sum of wages and 0.7 times proprietors income. Capital income is the sum of 0.3 times proprietors income, corporate profits, interest income, rental income net of corporate taxes and non-mortgage interest payments.

X. The sample ranges from 1976M1 to 2017M12.³

The number of factors in the FAVAR model is chosen via the information criteria of [Bai and Ng \(2002\)](#). This procedure suggests the presence of 15 factors. The lag length is set at 12.⁴ The parameters of the VAR model in (1) are estimated using a Bayesian approach. We use a Markov chain Monte-Carlo algorithm to approximate the posterior distributions.⁵ We employ 11,000 iterations, retaining every 10th draw after a burn-in period of 1000.

3 Empirical results

Before turning to the effect of the oil supply news shock on the distribution of income, we show the response of selected aggregate variables to this shock in [Figure 1](#). These results broadly support the conclusions reached by [Känzig \(2021\)](#). A 10% increase in the oil price leads to an increase in oil inventories and the median response of oil production is negative at medium horizons, albeit with large error bands. The shock depresses both global and US industrial production and leads to increase in the US unemployment rate and CPI. The shock has a limited effect on short-term interest rates but affects financial conditions adversely, with the BAA spread increasing and the stock market index declining.

3.1 Impact on the distribution of income

[Figure 2](#) shows our main result. The left panels of the figure show the median response of total income, labour income and capital income, averaged in each group defined by deciles of total income. The right panels present the response of these variables in each decile group, along with 90% error bands at the 2 year horizon. The top row of the figure shows that the oil supply shock has the largest effect on income of individuals on the left tail—for the first decile, total income declines by 1% at the 2 year horizon. The impact is smaller towards the center of the distribution with the income of individuals in groups P_6 and P_7 falling by less than 0.5 %. However, for the top 10 %, the effect of the shock appears to be relatively larger. The second row of the figure shows that the impact on the left tail is driven by the large negative reaction of labour income. In contrast, capital income, that constitutes a larger proportion of income at the right tail, barely reacts significantly below

³As discussed in [Känzig \(2021\)](#), the instrument is only available from 1984M4 and the estimation of the A_0 matrix uses this sample.

⁴Our main results are robust to the number of factors and lags.

⁵The prior and posterior distributions for the VAR parameters are standard and described in the appendix.

the median at the 2 year horizon. For high income individuals, capital income declines substantially driving the larger reaction of total income observed for this group.

In a related paper, [Del Canto *et al.* \(2023\)](#) find that the oil supply news shock has a smaller effect on labour income for households with high educational attainment (Bachelors degree or higher) relative to households where the head has only obtained some college education. However, they report a relatively muted impact of the shock on households with even lower educational attainment. As discussed in their paper, educational attainment is likely to be correlated with the permanent component of income, while the distributions considered in our paper pertain to total income.⁶

Figure 3 shows the reaction of some of the main components of labour and capital income to the oil shock and suggests two key conclusions.⁷ First, the shock leads to a decline in labour income at the left tail as both wages and proprietor’s income declines. Second, capital income is adversely affected at the right tail – the shock is associated with a fall in interest income for groups P_7 to P_{10} and corporate profits for the top decile, possibly as a result of the rise in corporate spreads and fall in interest rates.

3.2 Robustness

We carry out a number of robustness checks that are presented in detail in the technical appendix. A summary is as follows:

1. Identification: As discussed in [Känzig \(2021\)](#), the oil supply news shock can also be identified under weaker assumptions: i.e. allowing for the possibility that other disturbances occur at the same time as the news shock. One method to accomplish this is identification via heteroscedasticity, which only requires the assumption that the variance of oil supply news shocks increases around OPEC announcements while the variance of other shocks remains unchanged. We show in the technical appendix that we obtain very similar results to the benchmark when the oil shock is identified using this approach. In particular, the shock has the largest effect on income at the left tail. We also use the time series of the oil supply shock from the VAR model of [Baumeister and Hamilton \(2019\)](#) as an alternative instrument. As shown in the appendix, the impulse responses obtained from this approach are broadly similar to the benchmark case.

⁶[Del Canto *et al.* \(2023\)](#) obtain labour income from the Current Population Survey. One crucial advantage of the [Blanchet *et al.* \(2022\)](#) database over the CPS is the fact that it incorporates information from the Survey of Consumer finances and may provide more accurate estimates of income at the right tail of the distribution.

⁷Note that interest income is defined as income from currency bonds and deposits.

2. Specification and model: The results are also preserved for FAVAR models with alternative lag lengths and number of factors. As a further check, we estimate the VAR model of [Känzig \(2021\)](#) adding the deciles of income measures one by one to the original set of endogenous variables. As shown in the appendix the results from this model support the benchmark conclusions regarding the distributional effects of the shock.

4 Conclusions

This paper shows that adverse oil supply news shocks have a heterogeneous effect on the US income distribution. While the impact of the shock is largest at the left tail of the distribution, more affluent individuals are also significantly affected by the shock. An examination of the components of income suggests that these results are driven by a decline in the labour income of the former group and capital income of the latter.

References

- Bai, Jushan and Serena Ng, 2002, Determining the Number of Factors in Approximate Factor Models, *Econometrica* **70**(1), 191–221.
- Barigozzi, Matteo, Marco Lippi and Matteo Luciani, 2021, Large-dimensional Dynamic Factor Models: Estimation of Impulse-Response Functions with I(1) cointegrated factors, *Journal of Econometrics* **221**(2), 455–482.
- Baumeister, Christiane and James D. Hamilton, 2019, Structural Interpretation of Vector Autoregressions with Incomplete Identification: Revisiting the Role of Oil Supply and Demand Shocks, *American Economic Review* **109**(5), 1873–1910.
- Baumeister, Christiane and Lutz Kilian, 2016, Forty years of oil price fluctuations: Why the price of oil may still surprise us, *Journal of Economic Perspectives* **30**(1), 139–60.
- Berisha, Edmond, Carolyn Chisadza, Matthew Clance and Rangan Gupta, 2021, Income inequality and oil resources: Panel evidence from the United States, *Energy Policy* **159**, 112603.
- Blanchet, Thomas, Emmanuel Saez and Gabriel Zucman, 2022, Real-Time Inequality, *NBER Working Papers 30229*, National Bureau of Economic Research, Inc.
- Caldara, Dario, Michele Cavallo and Matteo Iacoviello, 2019, Oil price elasticities and oil price fluctuations, *Journal of Monetary Economics* **103**, 1–20.
- Canova, Fabio and Filippo Ferroni, 2022, Mind the gap! Stylized dynamic facts and structural models, *American Economic Journal: Macroeconomics* **forthcoming**.
- Cantore, Cristiano, Filippo Ferroni, Haroon Mumtaz and Angeliki Theophilopoulou, 2023, A tail of labor supply and a tale of monetary policy, *Discussion Papers 2308*, Centre for Macroeconomics (CFM).
- De Giorgi, Giacomo and Luca Gambetti, 2017, Business cycle fluctuations and the distribution of consumption, *Review of Economic Dynamics* **23**, 19 – 41.
- Del Canto, Felipe N., John R. Grigsby, Eric Qian and Conor Walsh, 2023, Are Inflationary Shocks Regressive? A Feasible Set Approach, *NBER Working Papers 31124*, National Bureau of Economic Research, Inc.
- Forni, Mario and Luca Gambetti, 2014, Sufficient information in structural VARs, *Journal of Monetary Economics* **66**(C), 124–136.

- Hamilton, James D., 2003, What is an oil shock?, *Journal of Econometrics* **113**(2), 363–398.
- Känzig, Diego R., 2021, The macroeconomic effects of oil supply news: Evidence from OPEC announcements, *American Economic Review* **111**(4), 1092–1125.
- Mertens, Karel and Morten O. Ravn, 2013, The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States, *American Economic Review* **103**(4), 1212–47.
- Piketty, Thomas, Emmanuel Saez and Gabriel Zucman, 2017, Distributional National Accounts: Methods and Estimates for the United States*, *The Quarterly Journal of Economics* **133**(2), 553–609.
- Stock, James H. and Mark W. Watson, 2008, What’s New in Econometrics - Time Series, *Lecture 7*, National Bureau of Economic Research, Inc.

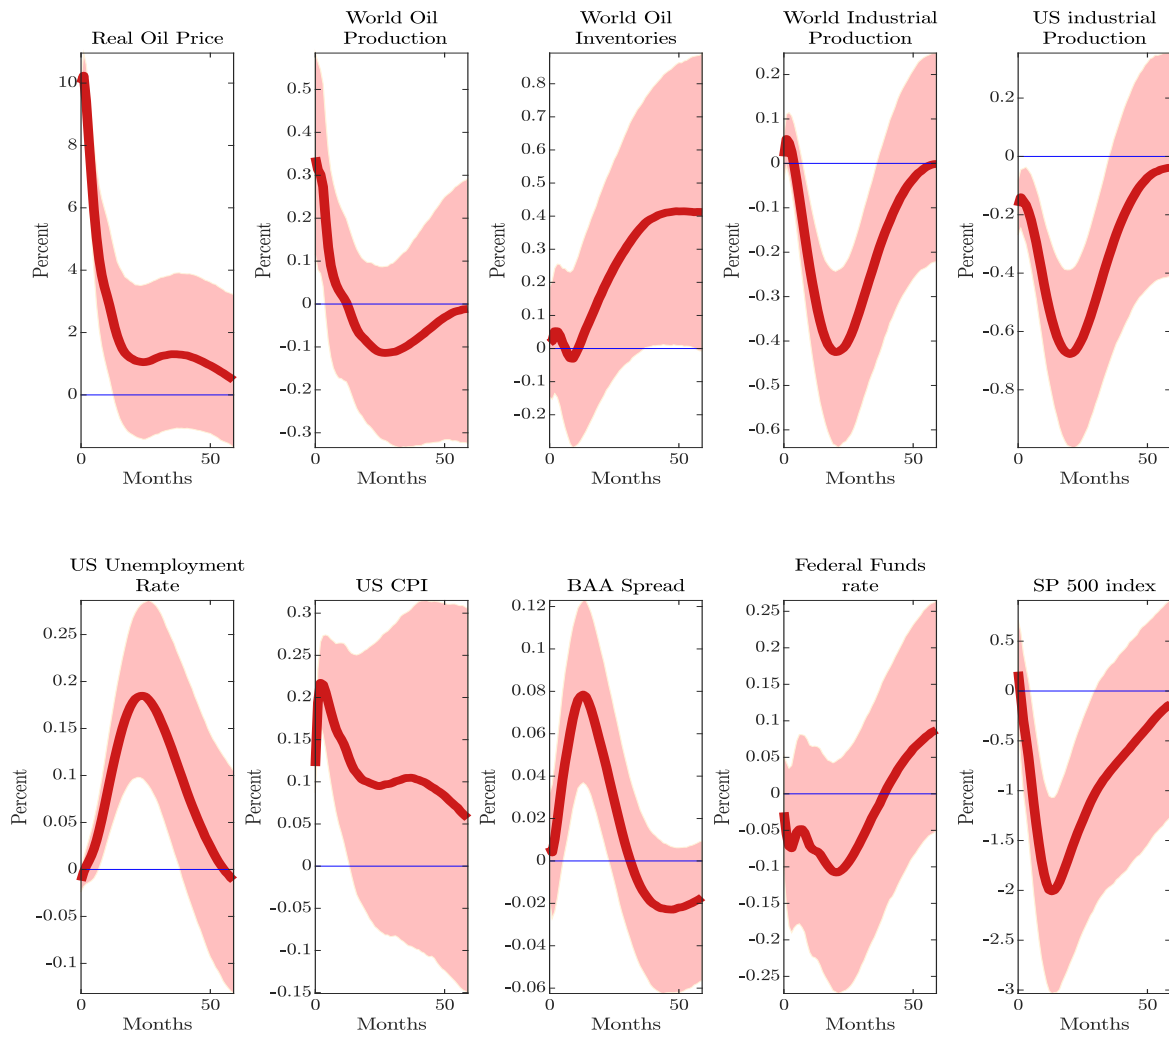


Figure 1: Impulse response functions of selected variables to an oil supply news shock. The shock is normalised to increase the oil price by 10%. The solid lines are the medians while the shaded area represents the 90% error band

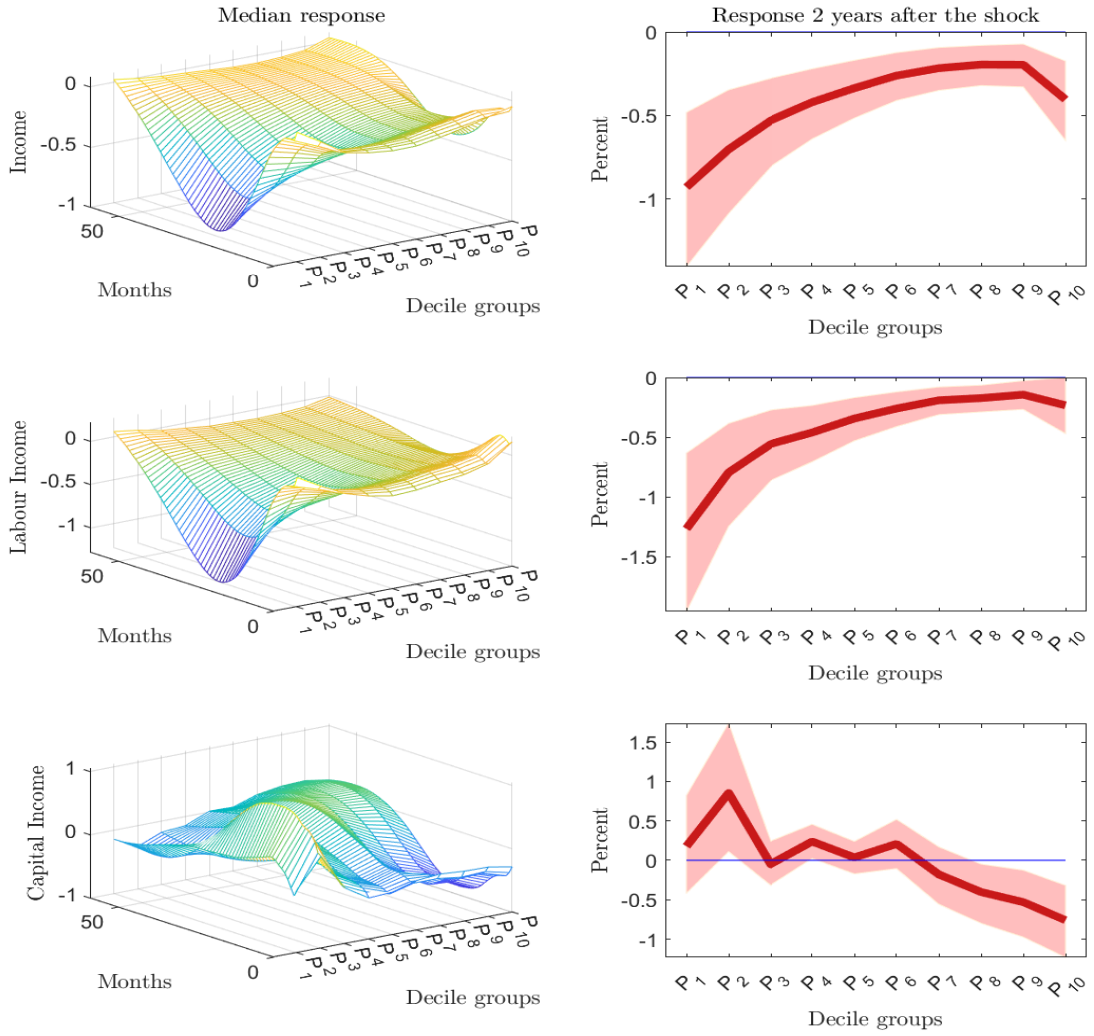


Figure 2: Impulse response functions of total income, labour income and capital income. The shock is normalised to increase the oil price by 10%. The first column shows the median response. The right panel shows the response at the 2 year horizon. The solid lines are the medians while the shaded area represents the 90% error band. P_1, P_2, \dots, P_{10} denotes the decile groups.

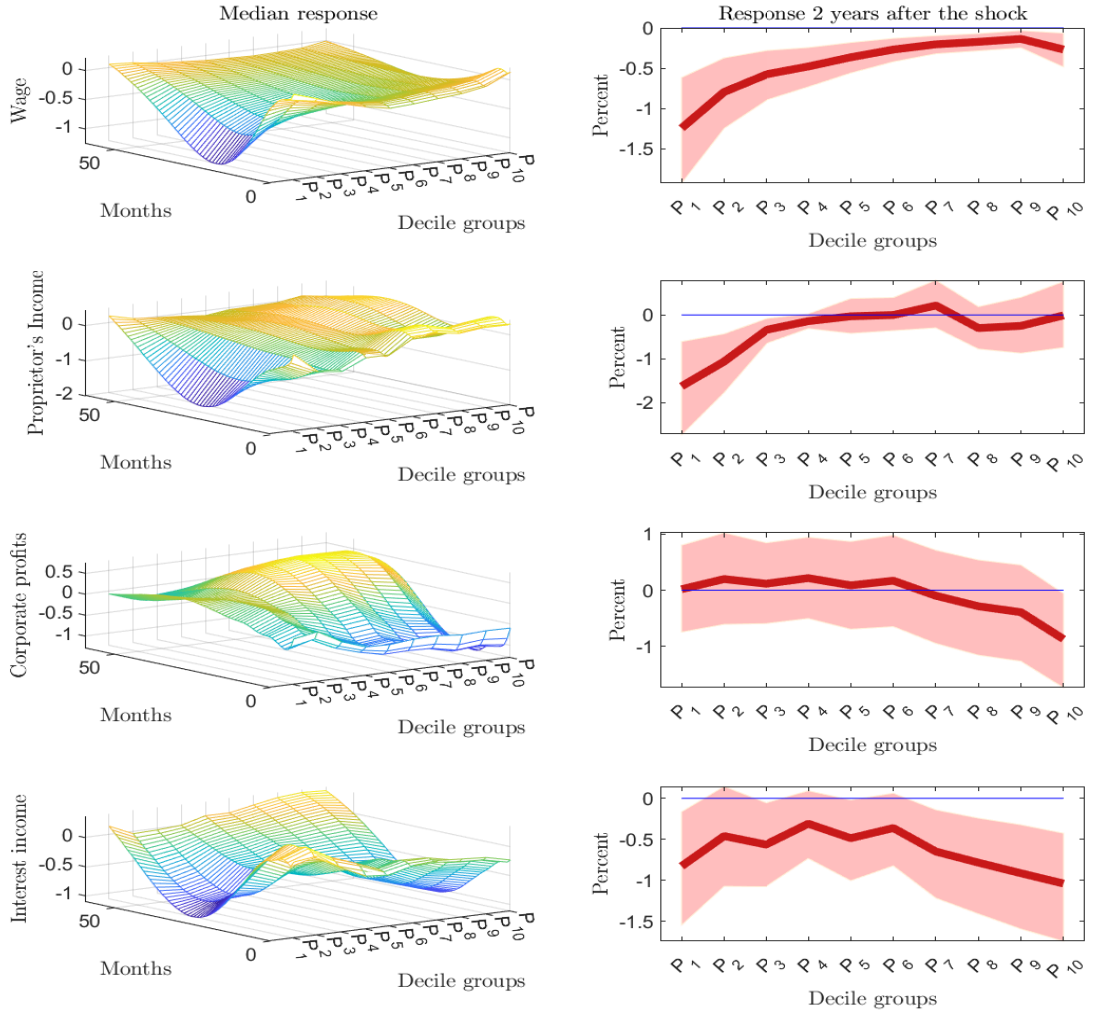


Figure 3: Impulse response functions of the components of labour and capital income The shock is normalised to increase the oil price by 10%. The first column shows the median response. The right panel shows the response at the 2 year horizon. The solid lines are the medians while the shaded area represents the 90% error band. P_1, P_2, \dots, P_{10} denotes the decile groups.