

The Distributional Effects of Economic Uncertainty

Florian HUBER

University of Salzburg, Austria

Massimiliano MARCELLINO

Bocconi University, Italy

Tommaso TORNESE¹

Università Cattolica del Sacro Cuore, Italy

Abstract

We study the distributional implications of uncertainty shocks by developing a model that links macroeconomic aggregates to the US distribution of earnings and consumption. We find that: initially, the fraction of low-earning workers decreases, while the share of households reporting low consumption increases; at longer horizons, the fraction of low-income workers increases, but the consumption distribution reverts to its pre-shock shape. While the first phase reduces income inequality and increases consumption inequality, in the second stage income inequality rises, while the effects on consumption inequality dissipate. Finally, we introduce Functional Local Projections and show that they yield similar results.

Keywords: Functional VAR, income inequality, density regressions, Functional LP

JEL Codes: C11, C30, E3, D31.

¹We would like to thank Minsu Chang, Jamie Cross, Jim Hamilton, Jiaming Huang, Andrea Renzetti, Jim Stock, Allan Timmermann and participants at the 2024 NBER Summer Institute and at seminars at the European Central Bank, and Bocconi University for useful comments and suggestions. Marcellino and Tornese thank MIUR – PRIN Bando 2017 – prot. 2017TA7TYC for financial support; Huber gratefully acknowledges financial support from the Austrian Science Fund (FWF, grant no. ZK 35). Please address correspondence to: Tommaso Tornese. Department of Economics and Finance, Università Cattolica del Sacro Cuore. *Email:* tommaso.tornese@unicatt.it.

1 Introduction

An extensive line of research stimulated by Bloom (2009) has studied the effects of uncertainty shocks on economic fluctuations.² It is generally agreed that an unexpected increase in the level of uncertainty about the future state of the economy generates a significant drop in output, employment and asset prices. Notwithstanding the importance of uncertainty shocks as drivers of the business cycle, little attention has been devoted to their distributional implications.

In this paper, we investigate the consequences of this type of shocks on the earnings distribution among employed people and on the overall consumption distribution in the US. Focusing on both consumption and earnings is indeed crucial for understanding the propagation mechanisms of uncertainty, and whether or not inter-temporal transfers can protect consumption from shocks that affect income (see e.g. Attanasio and Pistaferri (2016) for a discussion). The main econometric specification we use is a Functional Structural Vector Autoregression (F-SVAR) model, which represents a generalization in the functional space of the popular SVAR model commonly used in the empirical macroeconomics literature. Despite having been abundantly developed in the statistical literature, models for functional data are not yet popular among econometrics practitioners. Nevertheless, many recent methodological and empirical contributions have shown the high potential such models have for economic and financial analysis. Kowal et al. (2017) and Chang et al. (2024) develop Bayesian methods for inference in functional linear models, applying their techniques to analyze the dynamics of yield curves and income distribution respectively. Chang et al. (2016) and Hu and Park (2016) provide methods to analyze functional time series in a frequentist setting. Other applications of models for functional data in econometrics include Diebold and Li (2006) Tsay (2016), Meeks and Monti (2022), Inoue and Rossi (2021), Chang and Schorfheide (2024) and Bjørnland et al. (2023). Most of these papers, with the exception of Chang and Schorfheide (2024), assume that the function of interest is directly observed, while we recognize that only a sample from the

²See Castelnovo (2019) and Fernández-Villaverde and Guerrón-Quintana (2020) for recent surveys of the literature.

distribution of interest is available but the full probability density remains unknown.

In the recent period, macroeconomics research has dedicated more and more attention to the interplay between macroeconomic phenomena and distributional developments. The aggregate effects of distributional dynamics are studied, among others, by [Heathcote et al. \(2010b\)](#), [Athreya et al. \(2017\)](#) and [Rognlie et al. \(2019\)](#); while the distributional implications of aggregate shocks are the focus, for instance, of [Anderson et al. \(2016\)](#), [De Giorgi and Gambetti \(2017\)](#), [Mumtaz and Theophilopoulou \(2017\)](#), [Ahn et al. \(2018\)](#), [Kaplan and Violante \(2018\)](#), and [Bayer et al. \(2020\)](#). By means of a functional time series model, we contribute to the latter strand of the macroeconomics literature by showing that uncertainty shocks produce relevant and non trivial distributional effects.

From an econometric point of view, we adopt a three-step procedure along the lines of [Petersen and Muller \(2016\)](#) and [Petersen et al. \(2022\)](#). In the first step, we interpolate the cross-sectional data for each time period using kernel estimators (see, e.g., [Silverman \(2018\)](#)) to obtain period-by-period continuous earnings and consumption distributions, in order to be able to monitor the entire densities rather than only the observed data. In the second step, we transform the continuous earnings and consumption distributions to remove the unit-integration and non-negativity constraints as in [Petersen and Muller \(2016\)](#) and approximate the resulting curve with a set of basis functions (see e.g. [Ramsay and Silverman \(2002\)](#)). Specifically, we use functional principal components (FPCs), along the lines of [Kneip and Utikal \(2001\)](#) and [Petersen and Muller \(2016\)](#). In the final step, we jointly model the FPCs and a set of macroeconomic and financial indicators with a (Bayesian) VAR. By doing this, we are able to assess the effects of the uncertainty shocks on the FPCs, and then to reconstruct the Impulse Response Functions (IRFs) of the entire earnings and consumption distributions.

An alternative one-step approach would be to simply insert percentiles or moments of the earnings and consumption disaggregated data directly into the third step BVAR. Yet, in the case of percentiles, crossing could be an issue (see e.g. [Chang et al. \(2024\)](#)) and the BVAR dimension could easily become very large, if the goal is to get a granular view of the effects of uncertainty shocks on the earnings or consumption distribution, as

one would have to include a large number of percentiles into the model. A similar issue occurs in the case of moments, with the addition that higher order moments would be hardly interpretable, and their theoretical counterpart could also not exist in the presence of a heavy tailed distribution.³ Other empirical studies, including [Mumtaz and Theophilopoulou \(2017\)](#), [Theophilopoulou \(2022\)](#), and [Choi and Phi \(2023\)](#), only consider inequality measures, such as Gini coefficients or interquantile ranges, as endogenous variables in a VAR model. Such approach, however, conveys an inevitably partial account of the distributional dynamics, possibly providing misleading insights as discussed by [Heathcote et al. \(2010a\)](#) and [Choi and Phi \(2023\)](#).

Due to the three-step procedure, our methodology requires the use of "generated" regressors. Specifically, in the second step we rely on kernel-based estimated continuous distributions. In order to address concerns related to the generated regressors problem, one could either allow for the presence of measurement error, as e.g. in [Chang et al. \(2024\)](#), or use a kernel estimator that ensures consistency also in the presence of truncated distributions, see e.g. [Petersen and Muller \(2016\)](#). We follow the latter approach, in combination with a quite large cross-sectional dimension (about 10.000 units). Similarly, in the third step, we use the FPC in place of the estimated continuous earnings and consumption distributions. It can be shown, however, that the approximation error disappears as the cross-sectional dimension increases, when using a large enough number of FPCs (see e.g. [Kneip and Utikal \(2001\)](#)). In related contemporaneous work, [Huang \(2023\)](#) studies the asymptotic behaviour of a similar three-steps, but fully frequentist, procedure, and shows consistency of the kernel estimators of the density and the functional principal components, under proper assumptions. [Huang \(2023\)](#) also shows that the method works well in a classical context in MC experiments, and discusses an interesting application about the distributional effects of tax cuts in the UK.

To assess the empirical performance of our functional structural VAR in finite sam-

³One may be also tempted to perform standard Principal Component Analysis (PCA) directly on disaggregated data and use them in a standard VAR as in [Bernanke et al. \(2005\)](#). However, disaggregated data are usually collected in the form of repeated cross-section rather than in panel data form, making the use of standard PCA invalid.

ples, we conduct two sets of simulation experiments. In the first one, we generate data from two different types of F-VARs, while, in the second one, we use data simulated from the log-linearized solution of the version of the [Krusell and Smith \(1998\)](#) heterogeneous agents model considered in [Chang et al. \(2024\)](#). In all cases our approach returns estimated responses remarkably close to the true ones.

In our empirical analysis, we find that the propagation mechanism of uncertainty shocks on aggregated macroeconomic and financial data resulting from our functional VAR is similar to that reported by [Jurado et al. \(2015\)](#). Instead, their distributional effects are novel in the literature and can be broken down in two phases. First, in the short-run phase, as employment and output drop, a larger share of low-income workers are likely laid off, while those that keep their jobs see their wage increase relative to the GDP per capita level. This is reflected in a significant reduction of the mass of employed people with an income-to-GDP per capita ratio smaller than unit, which is presumably due to a larger increase of unemployment among the less specialized workers. At the same horizons, the response of the consumption distribution shows that the proportion of households reporting a low level of consumption increases significantly after the shocks, while the mass in the middle part of the distribution decreases. The response is therefore compatible with a short-run propagation mechanism in which some of the workers with a low degree of specialization, and belonging to the middle part of the consumption distribution, become unemployed and have to cut down their consumption level due to their inability to access consumption smoothing channels. On the other hand, in a longer horizon phase, while the effects on unemployment and output dissipate, the mass of workers with a low relative income increases again, supposedly due to the increase of the employment rate among the unskilled workers, and the stronger decrease of their labor productivity generated by the investment foregone in the first phase. At the same time, the consumption distribution reverts to its pre-shock shape. While the first phase reduces the overall earnings disparity among employed people but increases the degree of inequality in the distribution of consumption people, in the second stage the Gini coefficient of the income distribution rises considerably and that of the consumption distributions returns

to its pre-shock level.

Interestingly, this resembles the two-wave economic contraction in response to uncertainty shocks pointed out by [Carriero et al. \(2023\)](#). Just as at the aggregate level the investments foregone following an uncertainty shock generate lower economic growth in the long run, at a more granular level those missed investments may also reduce the productivity of lower-income employees and their compensation as a result.

Finally, we introduce Functional Local Projections (F-LPs), showing that they represent a viable alternative to the estimation of F-SVARs. To the best of our knowledge, we are the first to suggest the use of F-LPs in the literature. The F-LPs are obtained by regressing the FPCs on uncertainty and a set of controls. Our FPCs summarize the dynamics of the cross-section. Treating them independent of each other thus might have adverse effects on the estimated density response and we propose modeling them jointly in a system of equations. Since, as mentioned above, estimation uncertainty around the FPCs disappears in large cross-sections so that one can treat the estimated FPCs as observable variables, one can expect that FPC-based functional LPs and functional SVARs should lead to similar results, as standard LP and SVARs, see [Plagborg-Møller and Wolf \(2021\)](#). Indeed, we show that for both the simulated data and the actual data, the two methods yield comparable results, which provides additional robustness to the distributional effects of uncertainty that we have described.

The remainder of the paper is organized as follows: Section 2 describes the functional time series model and the econometric methods employed. Section 3 tests the ability of the model to track the propagation of structural shocks to the distribution of interest with simulated data while Section 4 extends the empirical analysis of [Jurado et al. \(2015\)](#) to study the distributional effects of macro uncertainty shocks. Section 5 introduces and demonstrates the merits of Functional Local Projections. Section 6 summarizes and concludes. A set of appendices provide additional results.

2 Functional Vector Autoregressions

In this section, we introduce our econometric framework, which is closely related to [Chang et al. \(2024\)](#). The first sub-section develops a model that is capable of capturing dynamic relations between the aggregate economy and the cross-sectional distribution of earnings or consumption. Sub-section 2.2 then discusses the transformation and approximation techniques used to summarize cross-sectional dynamics with a few unconstrained latent factors. In Sub-section 2.3 we summarize our Bayesian prior setup and sketch the posterior simulation algorithm.

2.1 Combining cross-sectional densities with aggregate time series

We aim to model the joint dynamics of a cross-sectional density, labeled $p_t(\xi)$ for $\xi \in \Xi$, with Ξ denoting the domain on which the probability density function (pdf) is defined, and a set of n_v aggregate economic random variables $y_t = [y_{1,t}, \dots, y_{n_v,t}]'$. Time runs from $t = 1, \dots, T$. The cross-sectional density, in our case, will alternately be the distribution of earnings or consumption in the US. We assume that we do not observe $p_t(\xi)$ directly but we are given a set of identically and independently distributed (iid) draws from $p_t(\xi)$, ξ_{it} for $i = 1, \dots, N$. Let $f_t = g(p_t)$ denote an invertible transformation of the density describing the distribution of interest. Additionally, let $\bar{f} = \frac{1}{T} \sum_{t=1}^T f_t$ be the sample mean of f_t , and $\tilde{f}_t = f_t - \bar{f}$ be the deviation of f_t from this mean.

The F-VAR with p lags then consists of two blocks:

$$y_t = c_y + \sum_{l=1}^p B_{l,yy} y_{t-l} + \sum_{l=1}^p \int B_{l,yf}(\acute{x}) \tilde{f}_{t-l}(\acute{x}) d\acute{x} + u_{y,t}, \quad (1)$$

$$\tilde{f}_t(x) = c_f(x) + \sum_{l=1}^p B_{l,fy}(x) y_{t-l} + \sum_{l=1}^p \int B_{l,ff}(x, \acute{x}) \tilde{f}_{t-l}(\acute{x}) d\acute{x} + u_{f,t}(x), \quad (2)$$

where x and \acute{x} are in the domain of the transformed distribution \tilde{f}_t , and $u_{y,t}$ and $u_{f,t}(x)$

are innovations with zero mean and variance given by:

$$\Omega(x, \acute{x}) = \begin{bmatrix} \Omega_{yy} & \Omega_{yf}(\acute{x}) \\ \Omega_{fy}(x) & \Omega_{ff}(x, \acute{x}) \end{bmatrix}.$$

Notice that the innovations are correlated across the aggregate macro series in y_t but also across the two blocks. This implies that structural economic shocks (such as the uncertainty shock we consider in this paper) can impact y_t but also have a direct effect on $\tilde{f}_t(x)$. It is also worth stressing that dynamic inter-dependencies are captured through the inclusion of lagged y_t 's in (2) but also by including lags of the log-densities in (1).

Since the functions $\tilde{f}_t(\cdot)$ and $u_{f,t}(\cdot)$ are continuous, the dimension of the model is infinite. In order to be able to conduct inference about the parameters of the infinite dimensional model, we resort to the Karhunen-Loève theorem (see e.g. [Bosq \(2000\)](#), Theorem 1.5) and write the function $\tilde{f}_t(\cdot)$ as an infinite order expansion:

$$\tilde{f}_t(x) = \sum_{k=1}^{\infty} \zeta_k(x) \times \alpha_{k,t}, \quad (3)$$

with $\zeta_k(\cdot)$ and $\alpha_{k,t}$ denoting time-invariant orthogonal functional bases and time-varying factors, respectively. Conditional on choosing an appropriate truncation level K , we can approximate (3) through:

$$\tilde{f}_t(x) \approx \sum_{k=1}^K \zeta_k(x) \times \alpha_{k,t} = \boldsymbol{\zeta}(x)' \alpha_t, \quad (4)$$

where $\alpha_t = [\alpha_{1,t}, \dots, \alpha_{K,t}]'$ and $\boldsymbol{\zeta}(x) = [\zeta_1(x), \dots, \zeta_K(x)]'$. Likewise, the other components of equations (1) and (2) can be approximated in a similar way:

$$u_{f,t}(x) \approx \boldsymbol{\zeta}(x)' \tilde{u}_{f,t}, \quad c_f(x) \approx \boldsymbol{\zeta}(x)' \tilde{c}_f, \quad B_{l,yf}(\acute{x}) \approx \tilde{B}_{l,yf} \boldsymbol{\delta}(\acute{x}), \quad B_{l,fy}(x) \approx \boldsymbol{\zeta}(x)' \tilde{B}_{l,fy} \quad (5)$$

and $B_{l,ff}(x, \acute{x}) \approx \boldsymbol{\zeta}(x)' \tilde{B}_{l,ff} \boldsymbol{\delta}(\acute{x})$. Here, $\boldsymbol{\zeta}(x)$ and $\boldsymbol{\delta}(\acute{x})$ are K -dimensional vectors of

time-invariant orthogonal functional bases, $\tilde{u}_{f,t}$ is a vector of time-varying factors, and \tilde{c}_f , \tilde{B}_{yf} , \tilde{B}_{fy} , and \tilde{B}_{ff} are matrices of scalar coefficients.

Using these finite approximations, the F-VAR in (1) and (2) can then be written as a finite dimensional VAR:

$$\begin{bmatrix} y_t \\ \alpha_t \end{bmatrix} = \begin{bmatrix} c_y \\ \tilde{c}_f \end{bmatrix} + \sum_{l=1}^p \begin{bmatrix} B_{l,yy} & \tilde{B}_{l,yf}C_\alpha \\ \tilde{B}_{l,fy} & \tilde{B}_{l,ff}C_\alpha \end{bmatrix} \begin{bmatrix} y_{t-l} \\ \alpha_{t-l} \end{bmatrix} + \begin{bmatrix} u_{y,t} \\ \tilde{u}_{f,t} \end{bmatrix}, \quad (6)$$

where $C_\alpha = \int \delta(\acute{x}) \zeta'(\acute{x}) d\acute{x}$, and $u_t = [u'_{y,t}, \tilde{u}'_{f,t}]'$ has zero mean and variance Ω . Conditional on knowing α_t , the model in (6) is a standard VAR model and inference can be performed applying conventional frequentist or Bayesian techniques. If K is large (i.e., the time variation in cross-sectional densities is complex and multi-dimensional and requires many modes of variation to approximate their behaviour), frequentist estimation of the model might suffer from overfitting issues. Hence, our approach will be Bayesian and we describe it in more detail in Sub-section 2.3 below.

It is worth stressing that the covariance matrix Ω is a full matrix. Hence, to attribute a structural interpretation to the model and to recover the structural shocks, we need to back out the structural representation of the model. This can be achieved by introducing suitable restrictions on the matrix A_0^{-1} , such that $\Omega = A_0^{-1} (A_0^{-1})'$, which defines the impact matrix, and then map back the impulse responses (IRFs) of α_t to the original functional space to which $\tilde{f}_t(\cdot)$ belongs. For instance, the structural form of the VAR is exactly identified if we assume that the system (6) is driven by $n = n_v + K$ iid structural shocks with unit variance, ε_t , and that the responses of the variables in the system to the shocks are such that we can write:

$$A_0 u_t = \varepsilon_t, \quad (7)$$

with A_0 an invertible lower-triangular matrix. In this case, the matrix A_0 can be simply found by inverting the lower-triangular Cholesky factor of the estimate obtained for Ω . This type of identifying restrictions are popular in the empirical macroeconomics litera-

ture and will be used in subsequent sections to identify uncertainty shocks. Noticeably, nothing prevents the use of other, more sophisticated, identification strategies in other applications of this class of models. As in standard SVARs, the impact matrix A_0^{-1} could be identified through a variety of second- or higher-moments assumptions (see e.g. Kilian and Lütkepohl (2017)). We focus on the "Cholesky" identification scheme because it is the one adopted by the important contribution of Jurado et al. (2015) to the literature on uncertainty shocks, which represents the foundation of the model we employ in Section 4 to study distributional responses.

2.2 Modeling cross-sectional dynamics using functional principal components

Up to this point, we have not specified the kernel density estimation method we employ, what the transformation $g(\cdot)$ does to the density p_t , and what kind of bases are contained in $\zeta(\cdot)$ and $\delta(\cdot)$. Hence, in this sub-section we describe how we estimate the continuous earnings and consumption distributions starting from an observed sample, we define the transformation $g(\cdot)$ applied to p_t to enforce the necessary non-negativity and unit-integration constraints, and illustrate the method we use to determine the functional basis that best approximates $\tilde{f}_t(\cdot)$ for any given truncation point K . We should mention that in these choices we differ from Chang and Schorfheide (2024), for the reasons discussed below.

2.2.1 Kernel density estimation

Since we want to model the dynamics of the entire continuous distributions of earnings and consumption, but only observe a sample from it, a density estimator has to be applied. To this end, we use a Gaussian kernel estimator specifying the bandwidth size following Silverman's rule of thumb, which has been shown to be a reliable choice in much of the applied statistics literature (see e.g. Kokoszka et al. (2019)).⁴ Because in our application

⁴In our simulations and application, we have experimented with different choices of kernel functions and bandwidth sizes, which did not produce relevant differences in the results.

(and in most practical cases) the support of the distributions of earnings Ξ and consumption are bounded due to the non-negativity of such variables and the censoring adopted by statistical agencies when conducting surveys, we apply a boundary correction to the kernel density estimator by augmenting data with their reflection near the boundaries. More precisely, we adopt the following density estimator:

$$\hat{p}_t(\xi) = \frac{1}{Nh} \sum_{i=1}^N \left\{ k\left(\frac{\xi - \xi_{it}}{h}\right) + k\left(\frac{\xi - \xi_{it}^L}{h}\right) + k\left(\frac{\xi - \xi_{it}^U}{h}\right) \right\}, \quad (8)$$

where $\xi_{it}^L = 2L - \xi_{it}$ and $\xi_{it}^U = 2U - \xi_{it}$, with L and U being respectively the lower and upper bound of the support Ξ .

2.2.2 Log Quantile Density transformation

An important challenge when modeling the dynamics of a distribution is to ensure that it always satisfies the non-negativity and unit-integration constraints intrinsic in the space of distributions. For example, one must always ensure that the predicted distribution h periods after a shock (i.e. the distribution IRF at horizon h) is non-negative at every point in the domain and integrates to one. One could simply set $g(p_t) = p_t$ and apply the approximation (4) to the deviations of the distribution $\tilde{f}_t = p_t - \frac{1}{T} \sum_{t=1}^T p_t$. As the sample mean $\frac{1}{T} \sum_{t=1}^T p_t$ has integral equal to 1 by construction, the unit integration constraint can be enforced by ensuring that each of the K basis functions integrate to 0 (for example by using FPC as basis, which have zero integral by construction). As discussed by Petersen and Muller (2016), however, the fact that the basis have zero integral also implies that they must have some negative parts, which can generate approximations of the distribution taking inadmissible negative values in some regions of the support. This problem is particularly important in the context of a F-VAR, in which the structural shocks move the coefficients α_t that multiply the functional basis. In such a context, if a large shock is estimated to have a large effect on some of the α_t , the negative part of the associated basis can become so important that it leads the predicted distribution to have negative regions after the shock.

To avoid this issue, most of the existing literature has applied the approximation in equation (4) to the logarithm of the density function, which naturally enforces the non-negativity constraint. This however is not sufficient to enforce the unit-integration constraint, and therefore leads again the functional IRFs to leave the space of probability functions. Some of the literature, then, brings the functional IRFs back to the space of distributions by re-normalizing them *ad hoc* to have unit integral.

Although this strategy may seem appealing, [Petersen and Muller \(2016\)](#) warn about the important potential flaws of such practice. The problem originates from the fact that standard functional data methods (e.g. FPCA) are designed to model the behaviour of random functions in the Hilbert space L^2 , while the space of densities is only a subspace of L^2 . A sensible modeling strategy would therefore need either to use functional data methods that restrict the analysis to the space of densities, or to transform the density in an unconstrained function in the Hilbert space L^2 . [Petersen and Muller \(2016\)](#) propose a coherent and straightforward strategy to do the latter. In particular, they suggest to apply standard functional data analysis to the Log Quantile Density (LQD) function associated with the distribution of interest, in place of working directly with the distribution itself. Even if there exists a one-to-one mapping between the LQD and the density function defined on a given support, the LQD has the remarkable advantage of being an unconstrained function in L^2 , and hence no special care is needed when applying (4) and (6) to its de-meaned version \tilde{f}_t . As the mapping between the LQD and the density function is one-to-one, in fact, a simple inverse transformation can be applied to retrieve the IRFs of the distribution of interest from those inferred for the LQD. The resulting distributional impulse responses will then satisfy, by construction, the required non-negativity and unit-integration constraints.

Given the clear advantage of the approach proposed by [Petersen and Muller \(2016\)](#) and the success of its application in [Kokoszka et al. \(2019\)](#) and [Petersen et al. \(2022\)](#), in what follows we define the function $f_t(\cdot)$ to be the LQD associated to the earnings or consumption distribution; the transformation $g(\cdot)$ is therefore the mapping that defines the LQD function associated with $p_t(\xi)$. The LQD associated with a probability distribu-

tion $p_t(\xi)$ is the logarithm of the first derivative of the quantile function $Q(z) = F^{-1}(z)$ (also known as inverse cumulated density function, icdf):

$$f_t(x) = g\{p_t(\cdot)\} = \log \left\{ \frac{d}{dz} Q(z) \Big|_{z=x} \right\}, \quad (9)$$

where $x \in [0, 1]$. The resulting LQD function is therefore defined on $[0, 1]$ and loses information about the support Ξ . Its great advantage, however, is that the LQD is unrestricted and lies in the L^2 space, and thus can be modeled without difficulty.


As the support Ξ is known, the distribution of interest $p_t(\cdot)$ can be easily derived back from $f_t(x)$ using the inverse transformation. It simply amounts to computing the quantile function $Q(x) = \theta \int_0^x \exp[f_t(z)] dz$, where $\theta = \sup_{\xi \in \Xi} \xi \times \left\{ \int_0^1 \exp[f_t(z)] dz \right\}^{-1}$, and then computing the first derivative of the associated cumulative distribution function $\frac{d}{d\xi} F(\xi) = \frac{d}{d\xi} Q^{-1}(\xi)$.

Figure 1 helps to visualize the LQD transformation and to highlight why it is convenient. The figure refers to a *Gamma(2,1)* distribution censored to have support $\Xi = [0, 5]$. The four panels represent the four steps that the mapping $g(\cdot)$ applies to $p_t(\cdot)$ (panel (a)) to find its LQD (panel (d)). In the first three panels, the transformation runs through three alternative ways of describing the same distribution, namely the *pdf*, the *cdf*, and the quantile function. These three objects are familiar, contain exactly the same amount of information, and are all subject to some constraint, that is, the *pdf* must be non-negative and integrate to 1, the *cdf* must be non-negative, non-decreasing, and take values from 0 to 1, the quantile function must be non-decreasing and take values in Ξ . In practice, once the kernel estimate of the *pdf* is available, the *cdf* (and hence its inverse, the quantile function) can be easily obtained by numerical integration. The only unfamiliar function in the figure is the LQD, which is simply the logarithm of the first derivative of the quantile function in panel (c), which can also be easily computed numerically. However, the LQD does not obey to any constraint and therefore represents an ideal modeling device.


As can be seen from the figure, the LQD loses the information about the support, but, since Ξ is known to the researcher, it is easy to retrieve any of the functions in panels

Figure 1: LQD Transformation


a) PDF




b) CDF



c) Quantile



d) LQD



(a),(b), or (c), once the LQD is available. To recover the $Q(z)$ once the LQD is known, for example, it is sufficient to integrate numerically $\exp[LQD]$ over $[0, z]$, and to re-normalize it so that $Q(1) = \sup_{\xi \in \Xi} \xi = 5$.

2.2.3 Functional Principal Component Analysis

As regards the functional basis used in the approximations of equation 4 and 5, the choice of $\delta(\cdot)$ is of no practical relevance as long as its inner product with $\zeta(\cdot)$, C_α , is finite; while the choice of the basis $\zeta(\cdot)$ used to expand the function of interest, $\tilde{f}_t(\cdot)$, represents a central theme in the literature on functional data analysis (for a recent survey, see Wang et al. (2016)).

Various alternative methods are available. For instance, Chang et al. (2024) con-

sider splines to approximate the cross-sectional distribution. While their approach is flexible and non-parametric, the truncation point K (i.e. number of spline knots) needed to approximate sophisticated distributions can easily become very large, which in turn undermines inference for the associated VAR in equation (6). In this paper, we rely on the computation of Functional Principal Component Analysis (FPCA) described in Tsay (2016). We adopt FPCA due to its ease of use and the fact that it can be flexibly applied to summarize the bulk of the dynamic variation observed in $\tilde{f}_t(\cdot)$ through a limited number of factors α_t .

To set the stage, let Z denote a $T \times N_Z$ matrix with $(j, t)^{th}$ element $z_{jt} = \tilde{f}_t(x_{jt})$. We let N_Z denote the number of grid points on which we evaluate the LQD. Typically we set this to a large value (such as 1000). The matrix Z can be decomposed using a truncated singular value decomposition (SVD):

$$Z = SVD' + E \tag{10}$$

where S is a $T \times K$ matrix containing the first K left eigenvectors, V is a $K \times K$ -dimensional diagonal matrix with the first K largest eigenvalues on its main diagonal and D is a $N_Z \times K$ matrix of right eigenvectors. The term E measures the approximation error committed by just considering the first K eigenvectors. In what follows, we are going to use the first K principal components, computed as $A = SV$ with $\alpha_t = A'_t$ denoting the t^{th} row of A . The associated orthogonal basis are therefore the K columns of D .

The key intuition behind FPCA is that we aim to summarize cross-sectional information using a small number of factors. For example, if the density we wish to approximate is a time-varying Gaussian, the first two principal components of Z will be closely related to the time-varying mean and standard deviation. In this case, two factors will be sufficient to adequately approximate the associated \tilde{f}_t . If the cross-sectional density is non-Gaussian and exhibits more complex variation, the number of factors necessary to approximate \tilde{f}_t becomes larger. In practice, several strategies can be used to determine the truncation point K . As with standard factor models, one can rely on information

criteria as in Bai and Ng (2002), perform cross-validation as suggested by Tsay (2016), or stop the truncation when the share of variation explained in the observed cross-section is above a desired threshold. Alternatively, one can try different K to check how robust the analysis is to different choices. In this paper, we follow the latter strategy and repeat each application with different plausible K s, showing that the differences between the resulting functional IRFs of interest are mostly negligible. We also experimented with very small K and found that, although the model was still able to replicate the crucial features of the responses, increasing the truncation point to $K = 3$ for earnings and $K = 5$ for consumption improved the performance of the model, while increasing it further did not generate any relevant differences.

One advantage of this approach is that we introduce relatively little parametric restrictions. We start out by using a kernel density based estimate of the cross-sectional densities and transform it to become a simple function in L^2 . By doing so we can decide on the grid of points at which the LQD is evaluated. Since the number of grid points determines the size of the cross-section, we can use results such as the ones in Stock and Watson (2016) and show that the principal components are consistent estimators of the true factors. What is more important, however, is that estimation uncertainty surrounding the factors declines appreciably. This implies that we do not have to estimate a full state space system but can simply plug in the PCA-based estimates of the latent factors.

Before discussing the estimation of the SVAR, it is worth stressing that the resulting functional VAR resembles a factor-augmented VAR similar to the one proposed in Bernanke et al. (2005). This implies that the choice of how we capture dynamics in the cross-sectional densities impacts the intrinsic nature of the reduced form innovations $u_{\alpha,t}$, but it does not affect the labeling of the first n_v structural shocks in ε_t , $\varepsilon_{y,t}$, in a Cholesky identification scheme. As long as the orthogonal decomposition is able to reflect the time dependence between the aggregate variables and the cross-sectional distribution of interest, the performance of identification schemes used to disentangle the aggregate structural shocks of interest is not influenced by the chosen functional basis.

2.3 Bayesian inference

Once the scores α_t are obtained via FPCA, the VAR in (7) is estimated with Bayesian methods. We specify a natural conjugate Gaussian-Inverse Wishart prior distribution for the reduced form parameters, see, e.g., Sims and Zha (1998); Koop et al. (2010); Giannone et al. (2015).

To simplify prior implementation, we rewrite the model more compactly as:

$$z_t = \Pi x_t + u_t, \quad (11)$$

where $z_t = [y_t', \alpha_t']'$, $x_t = [1, z_{t-1}', \dots, z_{t-p}']'$ and $\Pi = [\Pi_0, \Pi_1, \dots, \Pi_p]$, with $\Pi_0 = [c_y', \tilde{c}_f']'$ and $\Pi_l = \begin{bmatrix} B_{l,yy} & \tilde{B}_{l,yf}C_\alpha \\ \tilde{B}_{l,fy} & \tilde{B}_{l,ff}C_\alpha \end{bmatrix}$ for $l = 1, \dots, p$. This is a standard multivariate regression model which explains the observed aggregate macro series and the functional PCs using only lagged values of z_t .

The natural conjugate Gaussian-Inverse Wishart prior distribution for the reduced form parameters is factorized as

$$p(\text{vec}(\Pi'), \Omega) = p(\Omega) \times p(\text{vec}(\Pi') \mid \Omega),$$

where $p(\Omega)$ is Inverse Wishart with ν degrees of freedom and scale matrix Φ , and $p(\text{vec}(\Pi') \mid \Omega)$ is Gaussian with mean $\text{vec}(\Psi)$ and variance $\Omega \otimes \Gamma$. In specifying the parameters of these priors, we follow the Minnesota tradition of Doan et al. (1984) and set Ψ to be a $n \times m$ ($m = np + 1$) matrix of zeros, except for the $(i, i + 1)^{th}$ element that is set to 1 if the i -th variable of the system is known to be persistent. The component of the prior variance Γ is a diagonal $m \times m$ matrix with the $(i, i)^{th}$ element equal to 10^3 if $i = 1$, and $\frac{\lambda_1^2}{\sigma_j l^{\lambda_2}}$ otherwise, where j and l are the index and lag of the right hand variable to which the j -th column of Π refers, and σ_j is the error variance of an AR(1) model estimated by OLS for the j -th variable. Finally, the parameters of the Inverse Wishart prior $p(\Omega)$ are set to $\nu = n + 2$ and $\Phi = \text{diag}[\sigma_1, \dots, \sigma_2]$. Throughout the paper, the

hyperparameters λ_1 and λ_2 are set to 0.2 and 2 respectively, values commonly used in the VAR literature.

Given these priors, the posterior distribution is also Gaussian-Inverse Wishart

$$p(\text{vec}(\Pi'), \Omega | Y) = p(\Omega | Y) \times p(\text{vec}(\Pi') | Y, \Omega).$$

Specifically, $p(\text{vec}(\Pi') | Y, \Omega)$ is a Gaussian distribution

$$p(\text{vec}(\Pi') | Y, \Omega) = \mathcal{N}(\bar{\Psi}, \bar{\Omega})$$

with variance and mean given by, respectively:

$$\begin{aligned} \bar{\Omega} &= \Omega \otimes \bar{\Gamma}, \quad \bar{\Gamma} = (\Gamma^{-1} + X'X)^{-1}, \\ \bar{\Psi} &= \bar{\Gamma} (\Gamma^{-1}\Psi + X'Y), \end{aligned}$$

with $X = [x'_1, \dots, x'_T]'$ and $Y = [z'_1, \dots, z'_T]'$ denoting stacked data matrices.

The posterior of the covariance matrix follows an inverse Wishart distribution:

$$p(\Omega | Y) = \mathcal{W}^{-1}(\bar{\nu}, \bar{\Phi}).$$

The posterior degrees of freedom are $\bar{\nu} = \nu + T$ and the scaling matrix is $\bar{\Phi} = \Phi + Y'Y + \Psi'\Gamma^{-1}\Psi - \bar{\Psi}'\bar{\Gamma}^{-1}\bar{\Psi}$.

The posterior just described is convenient because draws from the stationary distribution can be easily obtained by direct Monte Carlo, sampling from $p(\Omega | Y)$ first, and then drawing from $p(\text{vec}(\Pi') | Y, \Omega)$. This completes the description of our functional VAR model. In Section 3, we will illustrate that it works well in terms of recovering the true impulse responses using simulated data.

2.4 Construction of distributional IRFs

Before focusing on the application of the described procedure to artificial and real-world data, it is worth describing in more detail how we obtain posterior draws of the distributional IRFs. Following Chang et al. (2024), we assume that the system is at its steady state when the shock hits. For a given posterior draw of the model parameters, the LQD function before the shock is therefore $f_{ss}(\cdot) = \zeta(\cdot)' \alpha_{ss} + \bar{f}$, where α_{ss} is the unconditional mean implied by the model parameters. The steady state distribution to which it corresponds, $p_{ss}(\cdot)$, can then be easily found by applying the inverse transformation, $p_{ss}(\cdot) = g^{-1}(f_{ss}(\cdot))$, described in Section 2.2.2.

The response of the functional PC, α_t , to a structural shock of interest can be computed, for every posterior draw of the SVAR parameters, by standard methods, assuming that only the shock of interest at $h = 0$ differs from zero. The resulting standard IRFs are therefore defined as $IRF_{\alpha,j,d,h} = E \{ \alpha_{t+h} | I_{t-1}, \varepsilon_{t,j} = d, \varepsilon_{t,-j} = 0 \} - E \{ \alpha_{t+h} | I_{t-1}, \varepsilon_t = 0 \}$. These IRFs can then be used to find the expected value of the LQD function h periods after the shock, which we define as: $f_{ss+h}(\cdot) = \zeta(\cdot)'(\alpha_{ss} + IRF_{\alpha,j,d,h}) + \bar{f}$, and to the corresponding distribution of interest, $p_{ss+h}(\cdot) = g^{-1}(f_{ss+h}(\cdot))$, using the inverse transformation described in Section 2.2.2. The distributional IRFs that we report throughout the paper are then simply computed as the difference between $p_{ss+h}(x)$ and $p_{ss}(x)$, for every x in a thin grid belonging to the support Ξ .

An important feature to take into consideration when constructing distributional IRFs is that, although the SVAR is linear, the map from $p(\cdot)$ to $f(\cdot)$ is not, which implies that both the initial level of the distribution, $p_{ss}(\cdot)$, and the size of the shock, d , are relevant and must be chosen carefully. In our applications, $p_{ss}(\cdot)$ is chosen to be the steady state implied by the SVAR parameters, while the shock size is set in terms of standard deviations.

It is also important to notice that, unlike most of the existing literature, we do not need to re-normalize $p_{ss+h}(x)$ to ensure that it is positive and integrates to 1, since the use of the LQD ensures it by construction. In light of the discussion of Petersen and Muller

(2016), we believe that this represents a crucial advantage of our econometric procedure, since the re-normalization step operated by all the existing contributions to our knowledge can generate spurious shapes in the differences between $p_{ss+h}(x)$ and $p_{ss}(x)$.

3 Artificial Data Experiments


In this section we assess the ability of the F-SVAR, estimated through FPCA and the Bayesian routine just described, to capture the propagation of the aggregate structural shocks $\varepsilon_{y,t}$ to the distribution p_t . We do this by running three experiments with data simulated from known Data Generating Processes (DGPs). In the first experiment, the DGP is a simple F-SVAR of the form in (7), in which the endogenous function is the deviation of the LQD from its sample mean, \tilde{f}_t , which has an exact finite factor structure, as in (4). In the second experiment the endogenous function of the F-SVAR is instead the logarithm of the distribution p_t itself, which is also assumed to have an exact finite factor structure. In both cases, at each period a sample is drawn from the generated distribution and assumed to be observed, together with the realizations of the endogenous variables y_t . In the final simulation experiment, the DGP is taken from Section 5 of [Chang et al. \(2024\)](#), which is based on a log-linearized solution of the [Krusell and Smith \(1998\)](#) HANK model. In that case the endogenous function in the DGP represents the distribution of assets among the employed population. Finally, in Section 3.4 we use the three DGPs to examine how the choice of the transformation $g(p_t(\cdot))$ affects the goodness of fit of the approximation delivered by FPCA.

3.1 F-SVAR: DGP 1

The data used for the first experiment are generated by simulating forward the model in (11),⁵ setting $p = 4$, $n_v = 2$, and $K_{true} = 3$. The realizations of the K_{true} -dimensional vector time series α_t are then transformed in LQD functions, using as basis the FPC computed from the LQDs of 50 realizations of a mixture of Gamma distribution, where

⁵The exact model parameters used to simulate the data are reported in Appendix A.

Figure 2: Modes of variation: DGP 1



Notes: The black line depicts the sample mean of the distributions generated by the DGP. Red (blue) lines show the change implied by an increase (decrease) in $\alpha_{k,t}$ of 2 standard deviations.

the mixing distribution is a $Beta(a, b)$, with a and b uniformly distributed in $[0, 3]$. Doing so, f_t is obtained for $t = 1, \dots, 500$, and then transformed into a distribution on $\Xi = [0, 6]$ through applying the inverse mapping g^{-1} . From the resulting distribution, a sample of size $N = 8000$ is drawn at every period and assumed to be observed by the econometrician.


To get an idea of the distributional dynamics generated by the DGP, in Figure 2 we report the modes of variation implied by the DGP, that is the effect of a two standard deviations increase/decrease of each element of α_t on the endogenous distribution.

Once the data have been simulated, we estimate the distribution \hat{p}_t using the kernel estimator described in Sub-section 2.2.1, and then transform it to the LQD $\hat{f}_t = g(\hat{p}_t)$. We apply FPCA to the deviation from the mean of this functional data, setting $K = 7$,⁶ and conduct inference about the VAR parameters employing the Bayesian methods described in the previous section.

Figure 3 shows the h -step-ahead IRFs of p_t to the first $n_v = 2$ structural shocks (ε_y) implied by the DGP (red), together with the posterior median (solid blue) and the 90% credible bands (dashed blue) obtained for the model-based functional IRFs. We consider $h \in \{0, 4, 12, 24\}$. The two rows refer to the two structural shocks, while each panel in every row focuses on one of the horizons considered. As explained in Section 2.4, the distributional IRFs are computed as the difference between the distribution h periods after the shock and that corresponding to the steady-state. Hence, the horizontal axis reports the support Ξ of the distribution. The figure shows that the methodology described in

⁶Results with different truncation points K are reported in the appendix.

Figure 3: Functional impulse responses arising from the F-SVAR: DGP 1



Notes: Red lines show the true responses of $p_t(\xi)$ to one standard deviation shocks to ε_1 (upper panels) and ε_2 (lower panels). The solid blue lines represent the posterior median response, while dashed blue lines delimit the 90% credible bands. h denotes the horizon at which the response is measured.

Section 2 does a remarkably good job in tracking the propagation of the two shocks to the distribution of interest at all horizons. In particular, the model-based IRFs track the horizon-specific IRFs along the cross-section well. Only in some rare cases, the actual IRFs are outside the credible intervals. It is also interesting to note how the functional IRFs change as the number of principal components K included in the estimated model changes. The responses produced by the F-SVAR with $K = 1, 2, 3, 15$ are reported in the appendix and show that, as the number of principal components increases from 1 to 3, the accuracy of the inferred IRFs improves steadily, while for $K \geq 3$ the responses produced by the model are virtually the same. In particular, the model including only one principal component, although able to capture important features of the IRFs, fails to replicate the impact responses adequately. On the other hand, when at least three principal components are included, the model is able to track most of the distributional variation with credible bands that contain the true response most of the time. As seen from Figure A4, increasing the number of factors included to a moderately large number does not translate into substantially larger credible bands. We take this as evidence that

Bayesian prior shrinkage does a good job in limiting the efficiency loss resulting from estimating a larger model.

To check that the desirable properties we have described so far for a specific realization from the DGP are maintained in repeated samples, we generate 200 additional samples from the same DGP and apply the inference procedure described. For every replication and different values of K , we compute the correlation between the true functional responses to the n_v aggregated shocks ε_y at various horizons and the posterior median functional IRF inferred by the estimated F-SVAR. In particular, we consider functional responses at horizons $h = 0, 1, 2, 3, 4, 8, 12, 24$, and repeat the inference procedure for $K = 1, 2, 3, 5, 7, 15$. In addition, we consider the case in which K is set to be the smallest values for which the variance explained by the FPC is at least 90% (we refer to this method as 'scree plot' method using the terminology of Tsay, 2016). The bar charts in Figure 4 show the average correlation coefficients attained across replications. As in the single replication reported above, for K sufficiently large, the responses inferred by the F-SVAR are able to replicate the bulk of the distributional changes generated by the structural shocks in the true model. Moreover, as discussed earlier, increasing the number of principal components from 1 to 3 generates a substantial gain in the accuracy of the functional IRFs, but setting K above 5, or selecting it through the 'scree plot' method, does not produce any additional benefit.

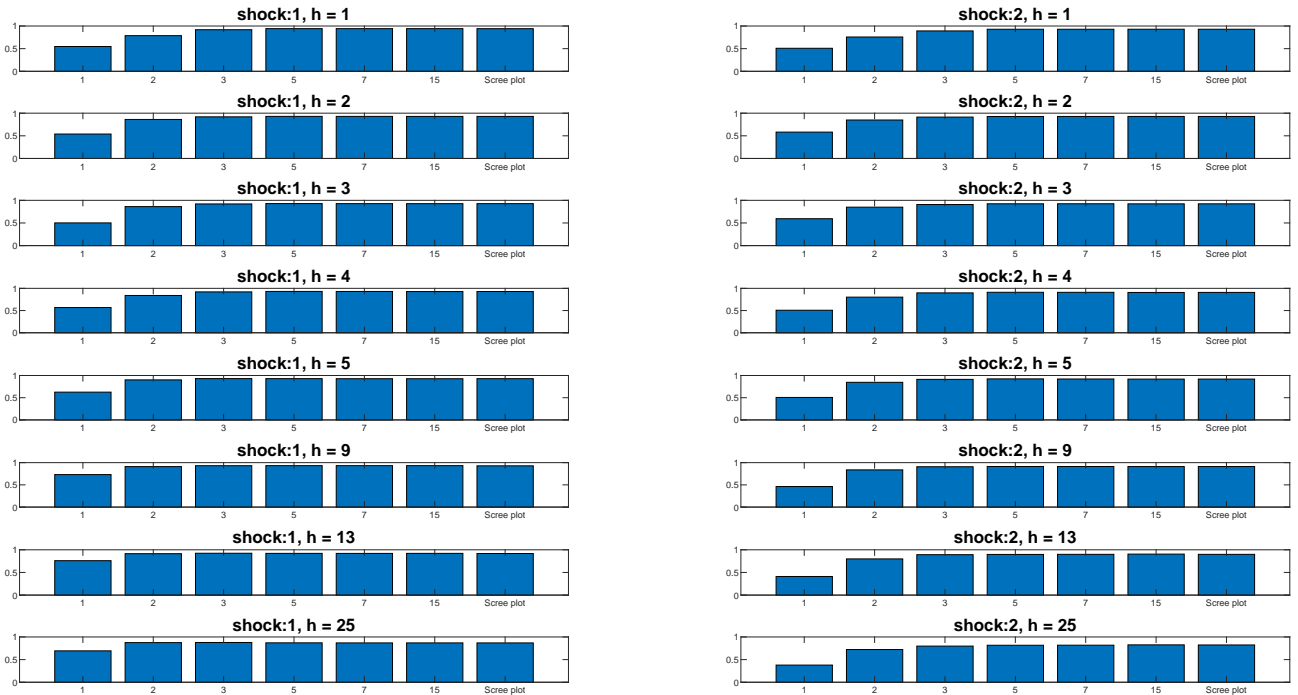
3.2 F-SVAR: DGP 2

The data used in the first experiment were generated from a model that differed only slightly from the model used to conduct inference.⁷ We now want to assess the reliability of our inference procedure in a more challenging setting.

In the second experiment, we increase the degree of mis-specification of the estimated model by assuming that the endogenous function in the DGP is the logarithm of the distribution p_t , rather than the LQD as postulated by the estimated model. Since the

⁷The difference between the DGP and the estimated model is in the functional basis used and in the true number of factors α_t , however the endogenous function in the VAR DGP is the LQD, as postulated by the estimated model.

Figure 4: Average correlation between median and true functional IRF: DGP 1




Notes: The coefficient is computed as the average across Monte Carlo repetitions and refers to the correlation between the point-wise median of the posterior distribution of functional IRFs and the true responses implied by the DGP.

transformation $g(\cdot)$ from p_t to the LQD f_t is highly non-linear, the true dynamics of the f_t 's are also highly non-linear, while our procedure assumes they follow a linear VAR. This second experiment therefore represents a notably more demanding challenge for the estimated model. As in the previous sub-section, we generate 500 observations of the vector time series $[y'_t, \alpha'_t]'$ by simulating the model in (11) forward. We set $p = 4$, $n_v = 2$, and $K_{true} = 3$.⁸ Differently from the first experiment, however, the obtained α_t 's are directly used to generate realizations of $\log p_t$ (with support $\Xi = [0, 6]$) using as basis the first K_{true} FPC of the logarithm of the mixture of Gamma distributions described in the previous sub-section. At every time period in the sample we normalize the obtained p_t to have unit integral and draw a sample of $N = 8000$ observations from it. Figure 5 depicts the modes of variation implied by the DGP.

Figure 6 compares the IRFs of p_t to the ε_y shocks implied by the DGP (red), with

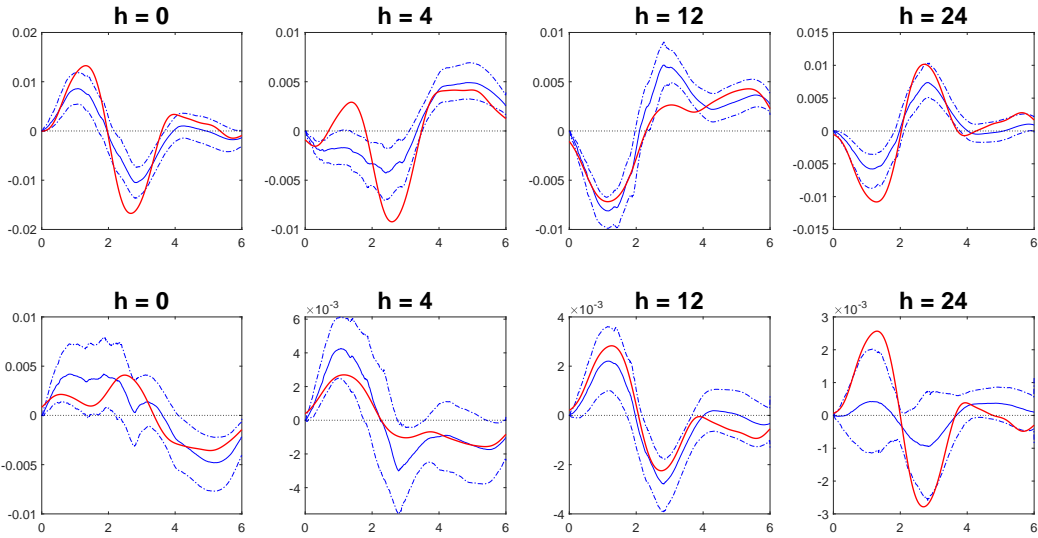
⁸Again, the exact model parameters used to simulate the data are reported in Appendix A.

Figure 5: Modes of variation: DGP 2



Notes: The black line depicts the sample mean of the distributions generated by the DGP. Red (blue) lines show the change implied by an increase (decrease) in $\alpha_{k,t}$ of 2 standard deviations.

Figure 6: Functional impulse responses arising from the F-SVAR: DGP 2



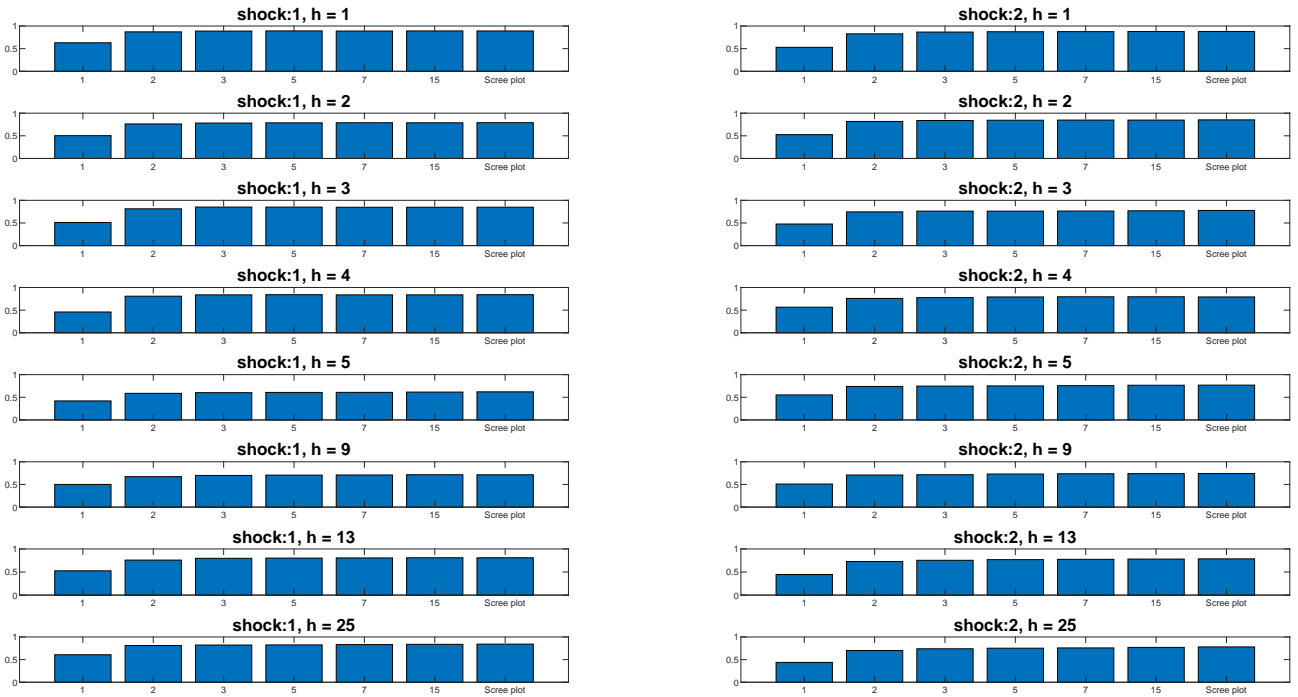
Notes: Red lines show the true responses of $p_t(\xi)$ to one standard deviation ε_1 (upper panels) and ε_2 (lower panels). The solid blue lines represent the posterior median response, while dashed blue lines delimit the 90% credible bands. h denotes the horizon at which the response is measured.

the posterior median (solid blue) and the 90% credible bands (dashed blue) obtained by the estimated model setting $K = 7$.⁹ Despite the higher degree of mis-specification of the estimated model, our procedure is able, also in this case, to replicate satisfactorily the salient features of the responses of the distribution of interest to the first n_v structural shocks, especially at shorter and medium horizons.

Next, as we did in the previous sub-section, we produce 200 additional replications of the experiment to check that the good performance just described is not limited to the

⁹Results with different truncation points K are very similar and are reported in the appendix.

Figure 7: Average correlation between median and true functional IRF: DGP 2



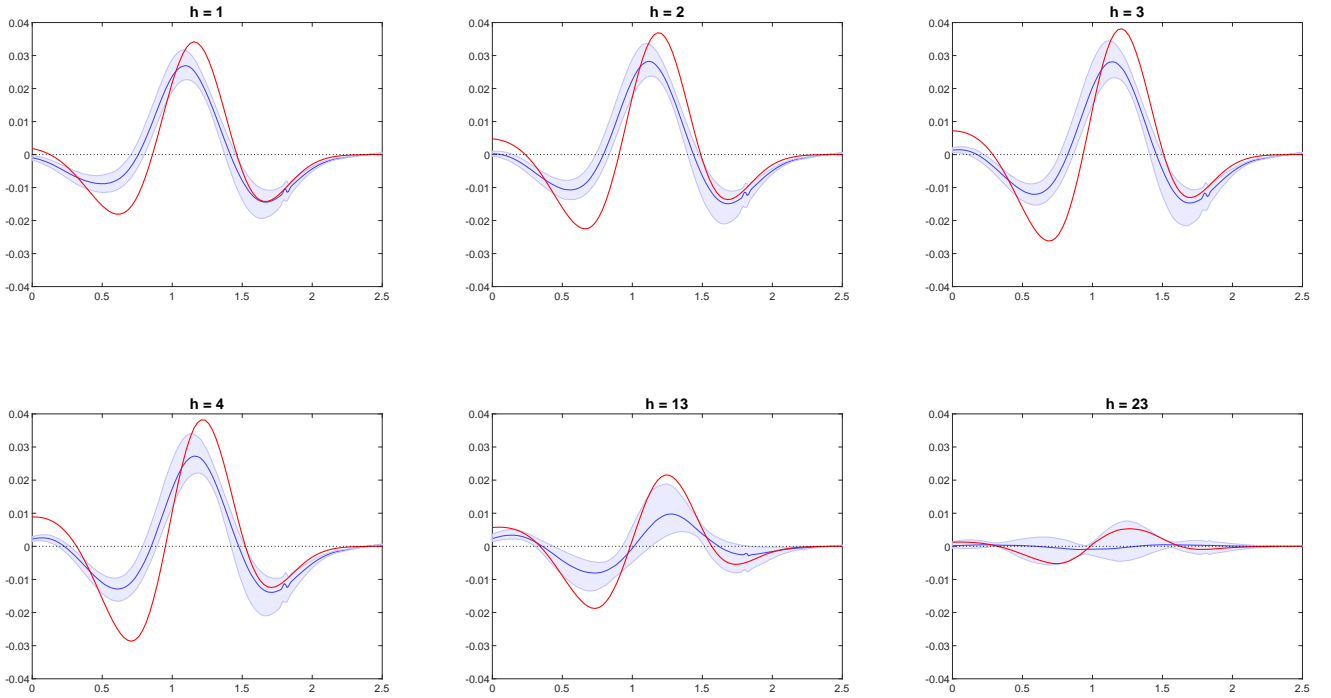
Notes: The coefficient is computed as the average across Monte Carlo repetitions and refers to the correlation between the point-wise median of the posterior distribution of functional IRFs and the true responses implied by the DGP.

particular sample obtained. The average correlation coefficients between the median of the posterior of functional IRFs and the ones implied by the DGPs are again large and are shown in Figure 7. Also in this case the correlation values are generally large, increasing the number of FPCs from 1 to 3 improves the accuracy of the inferred functional responses considerably, but the improvement of setting K to larger values is negligible on average.

3.3 Krusell and Smith (1998) DGP

The third experiment we perform is inspired by [Chang et al. \(2024\)](#) and makes use of the [Krusell and Smith \(1998\)](#) model as DGP. The log-linearized solution of the model proposed by [Krusell and Smith \(1998\)](#) implies a VAR law of motion for the productivity level, the capital stock, the employment level, and the centered moments of the distribution of assets among the employed. [Chang et al. \(2024\)](#) generate 160 artificial observations from this VAR and use them to estimate their proposed model. In this section, we use

Figure 8: Functional impulse responses arising from the F-SVAR: Krusell and Smith (1998) DGP



Notes: Red lines show the true responses of the asset distribution to a one standard deviation productivity shock. The solid blue lines represent the posterior median response, while dashed blue lines delimit the 90% credible bands. h denotes the horizon at which the response is measured. Notice that the horizons for which responses are reported are the same as those in Figure 6 of Chang et al. (2024), the difference in the panels titles only stems from different timing conventions.

the same artificial data kindly made available by the authors.¹⁰ For inference, we set $p = 1$, $K = 7$,¹¹ and specify the natural conjugate prior described in the previous section. Figure 8 compares the functional IRFs implied by the estimated F-SVAR (blue) with the true responses implied by the DGP. While the inferred responses tend in some cases to underestimate the deviation of the shocked distribution from the steady state counterpart, Figure 8 shows that the F-SVAR is still able to reproduce the salient features of the dynamic responses.

¹⁰We obtained the data from Frank Schorfheide’s website: web.sas.upenn.edu/schorf/working-papers/.

¹¹Results with different truncation points K are very similar and are reported in the appendix.

3.4 Alternative Transformations of $p_t(\cdot)$

Although we have shown that the econometric strategy described in Section 2 works well in a variety of settings, and despite we have already discussed the advantages of making use of the LQD transformation, in this section we want to understand how the goodness of fit of the FPC approximation of the distribution is affected when different transformations $g(p_t(\cdot))$ are considered. Specifically we compare the use of the LQD with the cases in which FPCA is performed on the deviation from the mean (i) of density itself, i.e. $g(p_t(\cdot)) = p_t(\cdot)$, and (ii) of the logarithm of the density, i.e. $g(p_t(\cdot)) = \log(p_t(\cdot))$. In order to do that, we implement a cross-validation exercise using samples of distributions generated by the DGPs described in Sections 3.1, 3.2, and 3.3. For each DGP and different truncation points K , we extract FPCs from 80% of the distributions in the sample (randomly selected), and use those to approximate the remaining 20% of the available distributions. More specifically, once we extract the basis $\zeta(x)$ through FPCA from the training sub-sample, we estimate by OLS the α_t coefficients associated with the distributions in the validation sub-sample. The obtained approximations are then used to compute the Mean Integrated Squared Error over the validation sub-sample:

$$MISE = \frac{1}{T} \sum_{t=1}^T \int_{\Xi} (\hat{p}_t(\xi) - p_t(\xi))^2 d\xi \quad (12)$$

where $\hat{p}_t(\xi) = g^{-1}(\bar{f} + \zeta(\xi)' \hat{\alpha}_t)$. We repeat the exercise 100 times and report the average MISE attained by each method in Table 1.

The Table shows that performing FPCA directly on the distribution provides the most efficient approximation, i.e. it achieves the lowest MISE for a given number of FPCs. The nonlinear transformations operated by the logarithm and by the LQD, instead, deteriorate to some extent the goodness of fit of the approximation. Intuitively, FPCA is designed to provide an optimal approximation in the transformed space, but the goodness of fit inevitably decreases once the approximation is projected back to the density space.

The Table, however, shows that the use of the LQD allows a substantially better

approximation of the distribution when compared to the log-transformation. To summarize, although there is some loss in the goodness of fit provided by the implementation of FPCA on the LQD rather than on $p_t(\cdot)$ itself, we think that the advantages discussed in Section 2.2.2, and the superiority in terms of approximation errors when compared to the logarithm transformation, justify our preference for the use of LQD over alternative transformations. In addition, Petersen and Muller (2016) provide examples in which the goodness of the approximation obtained by performing FPCA on the LQD can even exceed that obtained by performing FPCA directly on $p_t(\cdot)$. This mainly happens when large horizontal variation is observed across the distributions in the sample.

Table 1: Cross-validation exercise: MISE

| | | K | | | | |
|------|-------------------|-------|-------|-------|-------|-------|
| | | 1 | 2 | 3 | 4 | 5 |
| DGP1 | $p_t(\cdot)$ | 1 | 0.465 | 0.249 | 0.127 | 0.069 |
| | $\log p_t(\cdot)$ | 1.879 | 1.457 | 0.996 | 0.730 | 0.594 |
| | LQD | 1.085 | 0.644 | 0.530 | 0.370 | 0.307 |
| DGP2 | $p_t(\cdot)$ | 1 | 0.102 | 0.053 | 0.032 | 0.022 |
| | $\log p_t(\cdot)$ | 5.150 | 1.982 | 3.143 | 1.671 | 1.128 |
| | LQD | 2.334 | 1.167 | 0.678 | 0.421 | 0.337 |
| DGP3 | $p_t(\cdot)$ | 1 | 0.598 | 0.486 | 0.395 | 0.324 |
| | $\log p_t(\cdot)$ | 1.908 | 1.857 | 1.467 | 1.339 | 1.260 |
| | LQD | 1.449 | 1.199 | 1.174 | 1.066 | 1.020 |

Ratios relative to the MISE attained by the first approach for $K = 1$.

4 The Distributional Implications of Uncertainty Shocks

Having demonstrated the reliability of the adopted econometric procedure, the aim of this section is to apply these techniques to study the distributional consequences of uncertainty shocks. We start by discussing the dataset and how we identify an uncertainty shock before discussing the reaction of the macroeconomic aggregates to uncertainty shocks. Afterwards, we focus on how the cross-sectional distributions of earnings and consumption reacts to unexpected movements in economic uncertainty.

4.1 Data overview, structural identification and model specification

We analyze the relationship between economic uncertainty, macro aggregates and the distributions of earnings and consumption by building on the linear VAR model put forth in Jurado et al. (2015, JLN). In addition to the aggregate series discussed below, this amounts to including alternately the distribution of earning among employed people or the distribution of consumption as endogenous functional variable. We compile time series of earnings-to-GDP distributions using data constructed by Chang et al. (2024) based on the Current Population Survey (CPS). When we focus on the consumption distribution, we employ the micro data from the Consumption Expenditure Survey (CE) used by Chang and Schorfheide (2024). In this dataset, every observation is divided by the per capita consumption level, so that an observation equal to 1 represents a household consuming the national per capita level.¹² Since in Chang et al. (2024)'s and Chang and Schorfheide (2024)'s datasets the micro data are available at quarterly frequency, we convert the monthly SVAR(12) of JLN into a quarterly F-SVAR(4) model. The sample period for the F-VAR including the earnings distribution runs from 1989:Q1 to 2017:Q3. When we focus on the consumption distribution, the sample period goes from 1990:Q2 to 2016:Q4.

As endogenous aggregate series, we include 11 series in the model. These are: real GDP (\mathbf{rGDP}), real PCE (\mathbf{rPCE}), the GDP deflator (\mathbf{GDPdef}), real wages (\mathbf{rW}), real investments (\mathbf{rINV}), labor productivity (\mathbf{Lprod}), the unemployment rate (\mathbf{unr}), the Federal Funds Rate (\mathbf{ffr}), the S&P500 index ($\mathbf{sp500}$), the M2 growth rate ($\mathbf{M2}$) and the macro-uncertainty measure constructed by JLN ($\mathbf{U}_m \mathbf{h3}$).¹³ All variables are considered in logarithm, except for the interest rate and uncertainty measure.

To these variables we append the K factors α_t obtained through FPCA of the LQD

¹²Both datasets can be downloaded from Frank Schorfheide's website: web.sas.upenn.edu/schorf/working-papers/.

¹³All variables can be downloaded from the FRED MD database, except for the uncertainty measure which is available at www.sydneyludvigson.com/macro-and-financial-uncertainty-indexes.

While real GDP, real PCE, the GDP deflator, real wages, real investments, labor productivity, the unemployment rate are available at quarterly frequency, the Federal Funds Rate, the S&P500 index, the M2 growth rate and the macro-uncertainty measure are aggregated from the monthly series.

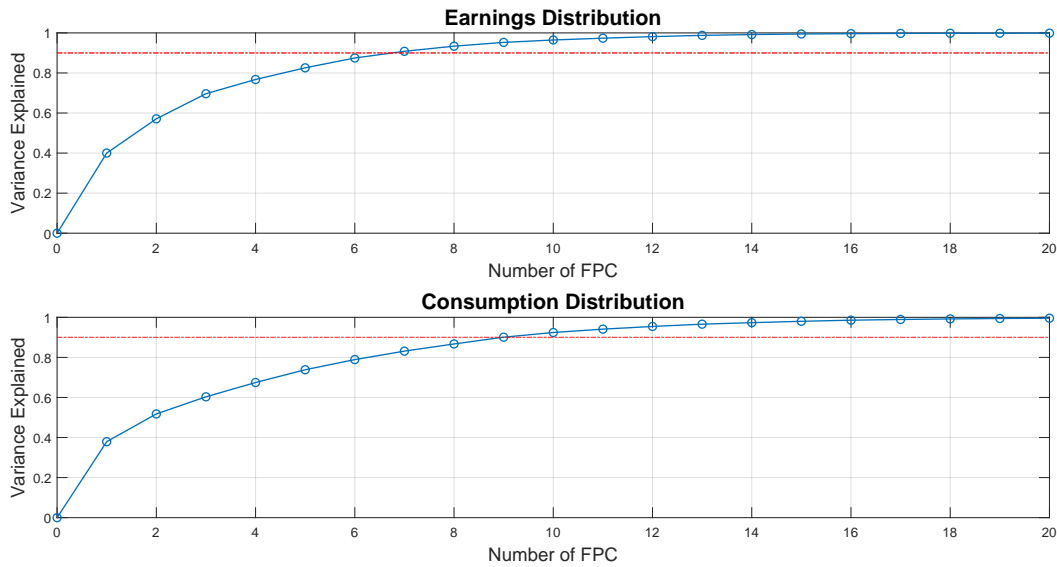
associated to the earnings or the consumption data. The chosen value of K is that for which the cumulative share of in-sample variance explained by the FPCs is closest to 90%. The scree plot in Figure 9 shows that, for earnings-to-GDP data, the first three FPCs are able to explain already 70% of the in-sample variation observed in the demeaned LQD, while the subsequent four PCs explain an additional share that accounts for over 20%. The first seven FPCs are therefore sufficient to explain more than 90% of the total variation observed. For consumption data, on the other hand, the first three FPCs explain 60% of the time-variation, and nine FPCs are enough to explain 90% of the in-sample variance. The relatively small value of FPCs required to reflect the bulk of the variation observed in the function of interest is indeed one of the main advantages of using FPCA instead of alternative functional bases, such as splines, for which a larger number would be needed. Furthermore, in our case, the exact choice of K does not appear to be crucial for the analysis, as long as it is not set to an excessively small number. In fact, we have experimented with different choices of K and noticed that different values did not produce any relevant differences in the results, as long as the value was greater than 2 for earning data and greater than 4 for consumption data.¹⁴

As CPS and CE data are top-coded with censoring values changing several times along the sample period, we limit the support of the distribution to the smallest censoring values applied in the surveys, which is slightly above an earnings-to-DGP per-capita ratio equal to 2 for CPS data, and around 3.2 times the per capita level of consumption for CE data. Our analysis remains therefore silent about the dynamics concerning the higher end of the earnings and consumption distributions, the study of such dynamics is left for future research.

Following JLN, the macro uncertainty shock is identified by ordering the uncertainty measure last among the endogenous variables in a Cholesky identification scheme. The idea behind this strategy is to reflect a structural source of fluctuation in the uncertainty measure that remains after accounting for other contemporaneous developments affecting

¹⁴Figures in Appendix E show functional IRFs obtained by setting $K = 3$ or $K = 15$ for earnings data and $K = 5$ or $K = 15$ for consumption data.

Figure 9: Share of variance explained by Functional PC



Notes: The blue line denotes the cumulated share of in-sample variation explained by the Functional PC. The dashed red line denote the threshold we apply for selecting the truncation point K .

the macroeconomy.

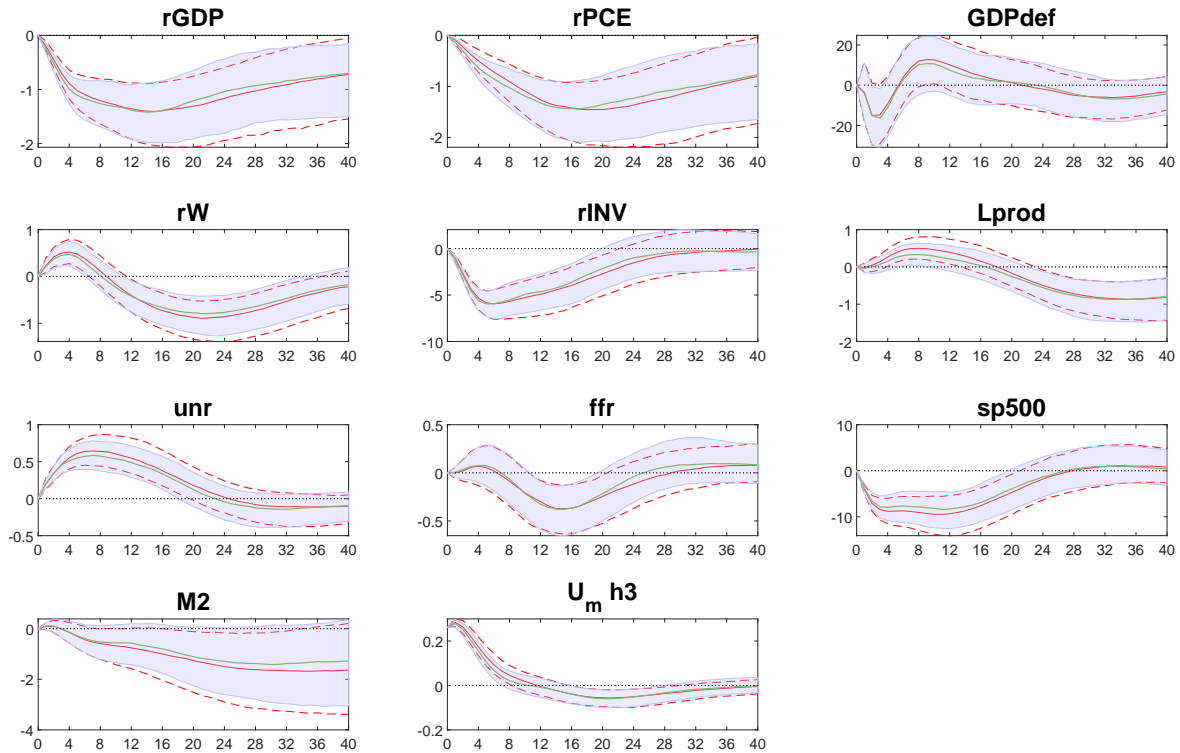
4.2 Dynamic reactions of the aggregate economy to uncertainty shocks

We start our discussion by focusing on the responses of the aggregate macro series first. These are shown in Figure 10. To understand the effects of adding the cross-sectional distribution of earnings or consumption to the VAR, we contrast the reactions of the F-SVAR (in blue) to the ones of a standard SVAR that features aggregate series only (in red).¹⁵ All these responses are to a one standard deviation shock, and we report 68% credible bands in both cases.

At a general level, the figure indicates that the responses are consistent with what has been found in the literature (Bloom (2009), Jurado et al. (2015), Castelnuovo (2019), Carriero et al. (2023)). Real output, consumption, investments and stock prices decline while the unemployment rate increases. As opposed to the earlier literature, we do not find evidence supporting an overshoot in real activity in the medium run. Real Wages

¹⁵In the Figure we report the IRFs generated by the F-VAR including the earnings distribution. Those generated by the F-VAR including the distribution of consumption are virtually identical.

Figure 10: Impulse response of the macroeconomic aggregates to a one standard deviation uncertainty shock



Notes: The solid lines represent the posterior median response, while dashed lines delimit the 68% credible bands. The horizon on the horizontal axis is expressed in quarters. Blue lines are generated by the F-SVAR, red lined are generated by the standard VAR.

and labor productivity display a similar behavior: after a muted short-run reaction, both decline at medium to long response horizons. Comparing the results between our F-SVAR and the SVAR points towards no discernible effects of including information in the form of the cross-section of earnings or consumption to the model.

The fact that our model produces IRFs that are in line with the previous literature gives confidence about the ability of the model to identify the desired shock, and assures that the model, although more densely parameterized than a simple SVAR, is still capable of reproducing the established features of the propagation of uncertainty shocks. It also suggests that the distributional dynamics related to uncertainty shocks, while relevant as we will see at the micro level, do not contain information that is not spanned by the set of variables considered by JLN.

4.3 Dynamic reactions of the cross-sectional distributions

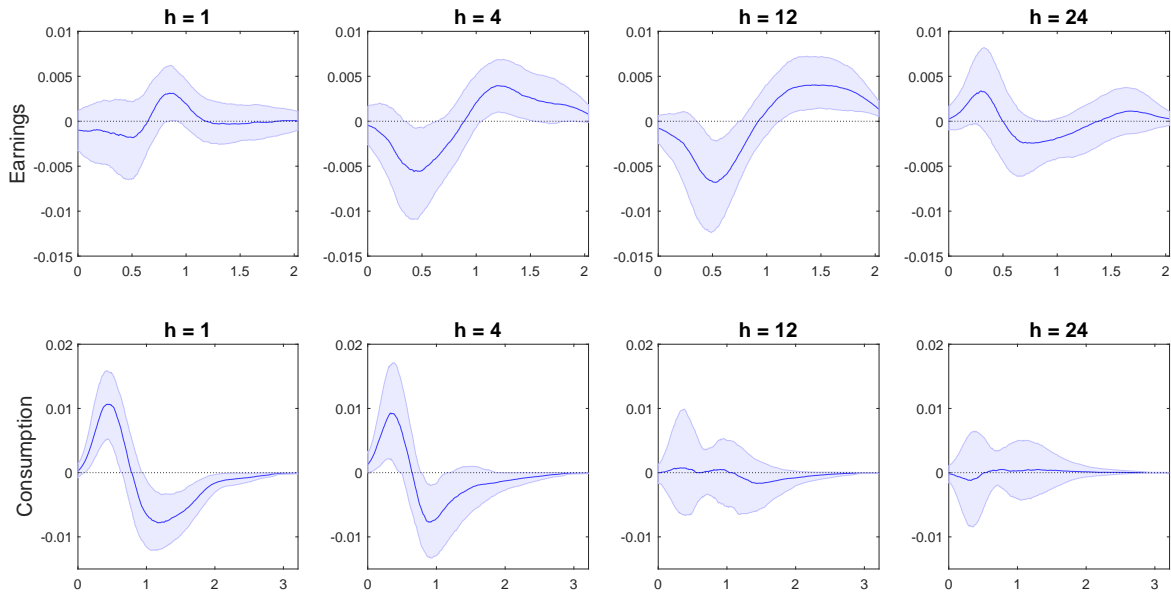
We now turn to how the cross-sectional distribution of earnings-to-GDP reacts to uncertainty shocks.

It is worth emphasizing again at this point that, for each impulse response horizon, we enforce the unit integration of the earnings distribution by construction. As a result, the distribution of earnings for which the changes in response to an uncertainty shock are shown below, only refers to employed people and does not contain information about variations of unemployment generated by the shock. The way we construct the distributional IRFs is therefore different from that in the empirical analysis of [Chang et al. \(2024\)](#), as they re-normalize the distribution at every horizon after the shock to integrate to $(1 - \text{unr})$, implying that the resulting distributional IRFs in their paper also reflect changes in the ratio between employed people and total population. As a consequence of this difference, the functional IRFs we show throughout the paper always integrate to zero, while those shown by [Chang et al. \(2024\)](#) in their empirical section integrate at each horizon to the change of the employment ratio generated by the shock. Despite accounting for the change of the mass of employed people has its merit, it requires artificially re-normalizing the area under the distribution at every horizon. For this reason, we prefer to focus on the distribution of earnings among employed people and to let the distribution integrate to one (and its change to zero) at every horizon, as implied by the use of the LQD transformation.

Figure 11 shows the change in the earnings-to-GDP and in the consumption distribution produced by the identified uncertainty shock at different horizons. In the first three years from the shock (i.e., for $h = 1, 4, 12$), the proportion of employed people receiving an income below the GDP per-capita level (i.e. earnings-to-DGP per-capita ratio equal to unity) decreases significantly, while the share of employed with salary between one and two times GDP per-capita increases. At the same time, the mass in the lower part of the consumption distribution increases significantly, reflecting households cutting down their consumption from an average level to a low level.

When we consider longer-run responses, however, the distributional consequences of uncertainty shocks change markedly. Specifically, the figure for $h = 24$ suggests that the portion of employed people receiving low salaries increases considerably, while the share of people earning salaries at the GDP per-capita level drops. At the same horizon, the median response for the consumption distribution also shows a slight decrease of the mass in the lowest part of the distribution; the uncertainty around this response is however too large to support any conjecture.

Figure 11: Functional responses of the earnings and consumption distributions to a one standard deviation uncertainty shock.



Notes: The first (second) row of the figure shows the difference between the earnings (consumption) distribution after a standard deviation shock and the one prevailing in the steady state. The solid blue lines represent the posterior median response, while dashed blue lines delimit the 68% credible bands. h denotes the horizon at which the response is measured. In the first (second) row, the measure on the horizontal axis is the earnings-to-GDP per-capita (households consumption - to - per-capita consumption) ratio.

A first important consideration when interpreting these results is that the distribution of earnings only refers to employed people, therefore it is very likely that the immediate decrease of the mass in the left part of the distribution does not only reflect people moving from the low-income category to a higher class. It may well be that the decline is also due to a higher proportion of the people laid off formerly belonging to a low-income class. This is strongly supported by the opposite change seen in the distribu-

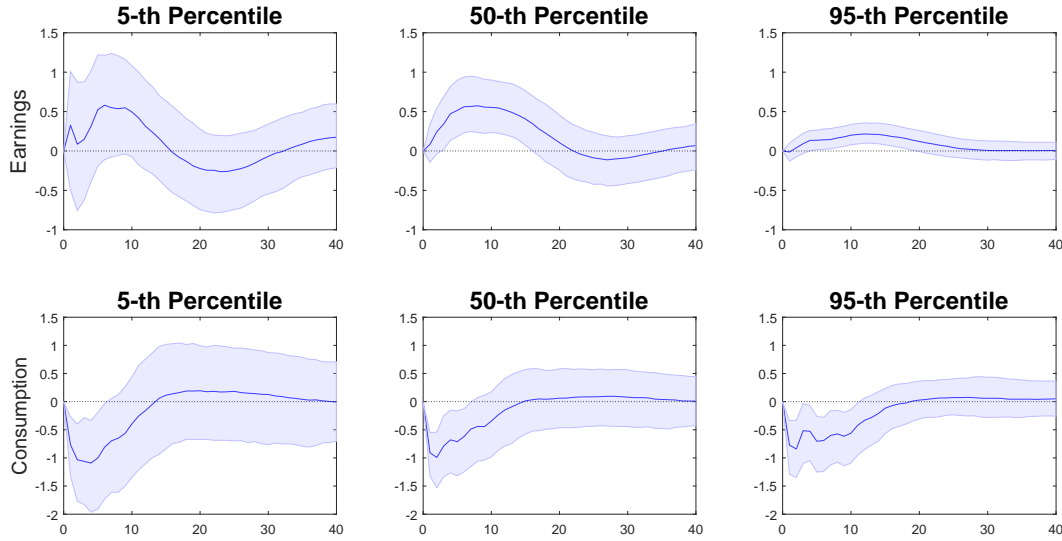
tion of consumption, where a significant share of households move from the central part to the lower tail.

Another important consideration is that the functional IRFs in Figure 11 must be read in conjunction with the IRFs of aggregate variables, in particular with the response of real wages and labour productivity. In fact, Figure 10 shows that the average real wage level actually increases at short horizons after the shock, but then it declines significantly, together with labour productivity, and reaches its trough between five and six years from the shock. This corroborates the existence of two stages in the propagation of uncertainty shocks, also evident from Figure 11.

To gain a better understanding of how uncertainty impacts specific quantiles of the distributions, we now focus on the responses of the 5th, 50th and 95th quantiles of the earnings and consumption distribution to a one standard deviation uncertainty shock. This is shown in Figure 12 and the corresponding IRFs are obtained by picking the relevant quantiles from the posterior of the functional responses of the distributions of interest. From the upper panels, it is clear that, while the entire distribution of earnings shift to the right, the median quantile is the one that displays the most significant percentage increase. This reflects an increase of the share of employed receiving salaries close to the per-capita GDP level, suggesting that the distributional developments triggered by uncertainty tend to concern to a minor extent the people receiving very high income. From the consumption side, the whole distribution shifts to the left, but the largest declines (in percentage) is again observed in the left part of the distribution.

A related but distinct examination that the F-SVAR allows us to consider is the assessment of how the relative weight of each earnings and consumption class changes in response to the identified shock. In Figure 13 we track the share of employed people belonging to four classes, each corresponds to one fourth of the distribution support. As for earnings, the figure suggests that the share of employed people belonging to the bottom income class decreases significantly in a first phase, while the relative weight of the classes whose earnings belong to the right half of the support increases. In a second phase, however, the share of low-income employed increases (although the credible bands

Figure 12: Responses of specific quantiles of the earnings and consumption distribution to a one standard deviation uncertainty shock.



Notes: The first (second) row shows the percentage change of specific percentiles of the of the Earnings-to-GDP ratio (consumption) distribution. The solid blue lines represent the posterior median response, while the dashed area represent the 68% credible bands. The horizon on the horizontal axis is expressed in quarters.

also contain the zero line at larger horizons), at the cost of mid-low income class.

As for consumption, on the other hand, the uncertainty shock affects almost exclusively the relative weights of the low- and mid-low consumption classes. In fact, the figure shows that the fast increase in the proportion of households reporting low consumption levels is associated with an almost identical decrease in the proportion of households reporting consumption level in the second fourth of the support. Meanwhile, the relative weight of the consumption classes belonging to the upper half of the support remain almost unchanged. This makes it even clearer that the decrease of the mass in the left part of the earnings distribution is mainly due to an increase of the pool of unemployed, who are then compelled to cut down consumption, due to their inability to access consumption smoothing channels.

In summary, we conjecture that the propagation of uncertainty shocks happens in two phases. In the short run, the rise in unemployment concerns to a larger extent people at the lower end of the earnings distribution, who therefore cut down their consumption, while the part of low-income workers that is not laid off actually sees its relative

Figure 13: Responses of earnings and consumption classes to a one standard deviation uncertainty shock

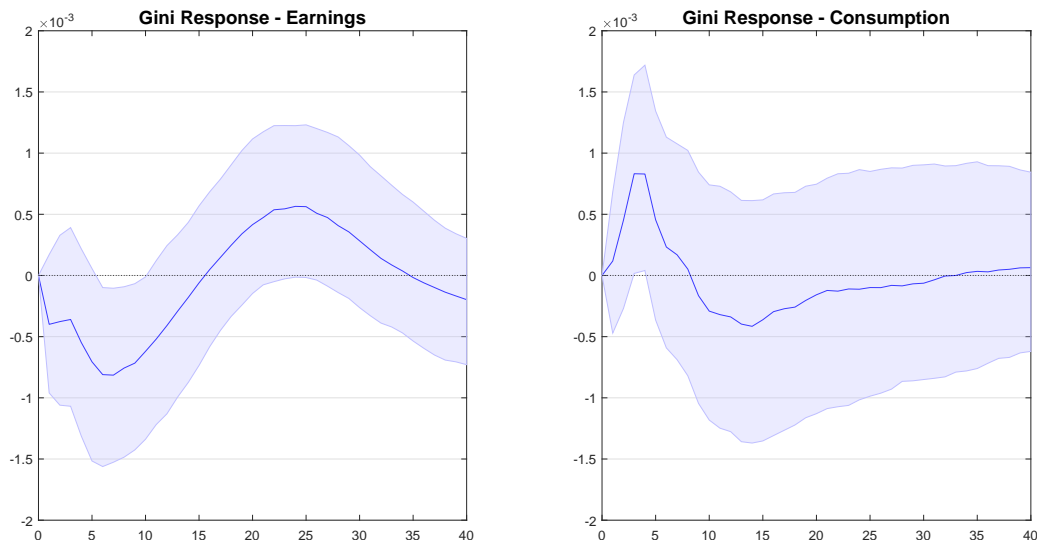


Notes: The first (second) row shows the response of the share of the employed (total) population belonging to specific earnings (consumption) classes. The solid blue lines represent the posterior median response, while the dashed area represent the 68% credible bands. The horizon on the horizontal axis is expressed in quarters.

earnings increased or kept constant. In the subsequent phase, when unemployment is finally reabsorbed, the size of the low-income group rises at the cost of middle-income employed people, who however do not seem to reduce their consumption level. This evidence is in line with the findings of [Choi and Phi \(2023\)](#), who show that uncertainty shocks affect more adversely poor people, while it is in contrast with the results obtained by [Theophilopoulou \(2022\)](#) for the UK, where poor people income and consumption appeared to be sustained more vigorously by social benefits.

It is interesting to see the reflection of the mechanisms just described in the responses of the Gini coefficient, which measures the degree of inequality in the distributions and can be computed from the functional responses produced by the F-SVAR. Figure 14 shows that the short-run developments generated by the uncertainty shock reduces the total level of income inequality among employed people but increases the degree of consumption inequality significantly, while the subsequent phase of the propagation mechanism causes

Figure 14: Responses of the Gini coefficient to a one standard deviation uncertainty shock



Notes: The solid blue lines represent the posterior median response, while dashed blue lines delimit the 68% credible bands. The horizon on the horizontal axis is expressed in quarters.

a considerable rise in the Gini coefficient, which is reabsorbed only eight years after the shock. This pattern is consistent with Fischer et al. (2021) who show that an increase in uncertainty triggers an immediate decline in the Gini coefficient and a subsequent increase after around three years.

It is also worth noticing that the responses depicted in Figure 14 are very similar to those that would be obtained by a standard quarterly SVAR(4) in which the Gini coefficients computed on CPS and CEX data are appended to the vector of macro variables. These responses are reported in Figure A11 and display the same two-phase propagation found by the F-SVAR. Clearly, the similarity of the result points to the robustness of our main findings, but does not imply that a standard SVAR would be sufficient to perform the analysis. The scope of the F-SVAR is much broader, as it gives the possibility to focus on any characteristic of interest of the earnings and consumption distributions, and to enforce the restrictions that are inherent in the space of distributions automatically. On the other hand, any standard SVAR alternative would inevitably limit the analysis to the characteristics included in the vector of endogenous variables, and would not guarantee that responses of these characteristics will be admissible. For example, in a standard

SVAR containing the Gini coefficient, the corresponding responses could take values below zero or larger than one; in a standard SVAR containing quantiles, the implied IRFs could be characterized by quantile crossing or could take values outside of the domain, e.g. negative income or consumption levels.

This discussion demonstrates that simply considering how summary statistics, such as the Gini coefficient, react to uncertainty shocks paints an inevitably partial picture of the overall distributional dynamics, and could run into problems that undermine the reliability of the results. For this reason, we believe the F-SVAR represents a more suitable tool to answer the central empirical question of this paper.

4.4 Investigating the mechanism: Unemployment reactions across educational levels

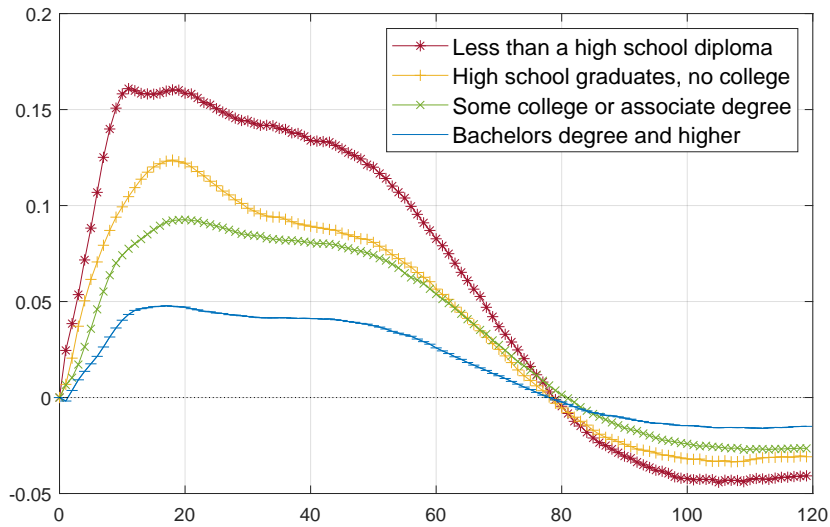
One important conjecture we made in the previous subsection is that the decline in the mass in the left tail of the earnings distribution is driven by changing employment levels across income classes.

To check that this is the case, we study the responses to uncertainty shocks of unemployment rates computed for different categories of the work-force. Unfortunately, we are not aware of any dataset providing employment information for sub-groups of the population based on their previous income level. However, the Bureau of Labour Statistics (BLS) provides unemployment rates computed monthly for sub-groups based on educational attainments, which can serve as a proxy of income classes.¹⁶ In particular, the BLS distinguishes between four educational levels: (i) less than a high school diploma, (ii) high school diploma but no college, (iii) some college but no degree or some associate degree, (iv) Bachelor's degree or higher.

Since the BLS series start in January 1992, using the quarterly dataset would imply too few observations. Hence, to understand how different education-specific unemploy-

¹⁶The average income differential between the categories is indeed substantial and reported periodically by the BLS. Figure A12 shows the evolution of the median weekly earnings level for the four categories in the time period for which it is available.

Figure 15: Dynamic responses of different unemployment categories to a one standard deviation uncertainty shock



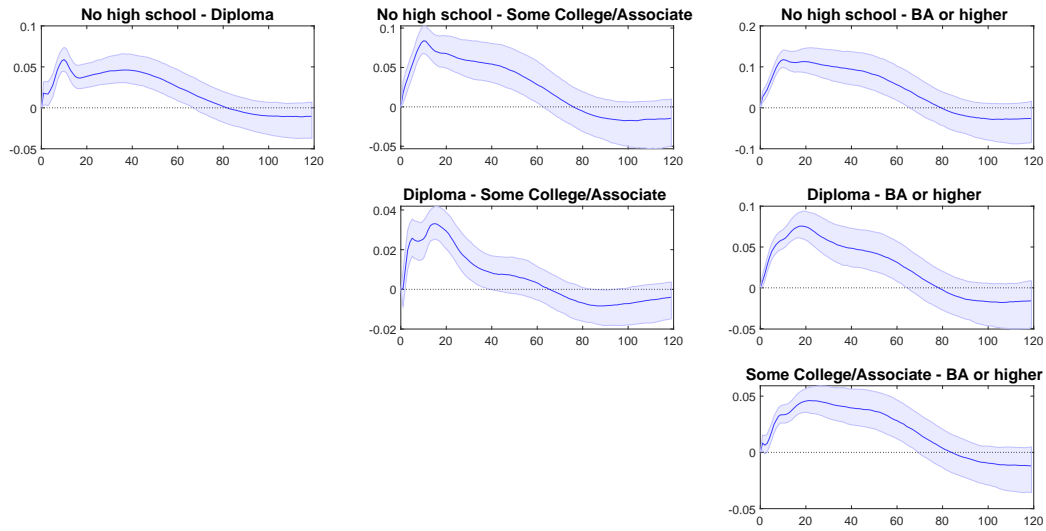
Notes: The four different lines depict the posterior median of different unemployment rates across education levels.

ment rates react to an uncertainty shock we add the (logarithm) of these granular unemployment series in the monthly SVAR of [Jurado et al. \(2015\)](#). The auxiliary VAR model is estimated on a sample going from 1992:M01 (the start of the BLS series) to 2019:M12 (to avoid having to deal with Covid-19-specific outliers in the data).

In [Figure 15](#), we report the (point-wise) posterior median IRFs obtained for the four sub-groups. The responses provide substantial evidence that workers with a lower education attainment, who tend to occupy the left part of the earnings distribution, are those that are lied off more often after an uncertainty shock. The figure tells a remarkably consistent story that employment reactions tend to become weaker with increasing educational attainment.

Considering only posterior medians neglects posterior uncertainty surrounding the estimates of the IRFs and one might ask whether differences across educational categories are statistically different from each other. To shed light on this, we report pairwise differences of the posterior distribution of impulse responses in [Figure 16](#). Since these are the responses of the differences, we can assess whether a specific horizon h IRF between two educational levels differs by simply checking whether the corresponding credible intervals include zero or not.

Figure 16: Differences between IRFs of different categories



Notes: Solid lines show the median difference between IRFs of the unemployment rate computed for different categories of the population. Shaded areas represent 68% credible bands. Categories are: (i) No high school: civilian population 25 years with less than a high school diploma; (ii) Diploma: civilian population 25 years graduated from high school but no college; (iii) Some College/Associate: civilian population 25 years who attended some college but did not graduate, or obtained some associate degree; (iv) BA or higher: civilian population 25 years with a Bachelor’s degree or higher educational attainment.

The pairwise comparisons in the figure reveal that the differences are, in fact, highly significant. Notice that these significant differences persist throughout the impulse response horizon, turning insignificant only after around 80 months. This relatively simple analysis, therefore, supports our conjecture that the short-run decrease of the mass in the most-left part of the earnings distribution highlighted in Figure 11 is, at least in part, due to a greater proportion of the people laid off belonging to a low-income class. If the insights collected from Euro-Area survey responses by Coibion et al. (2024) hold also for US citizens, then the difference in the unemployment risk of different income classes can also explain the results we have obtained for the consumption distribution. In fact, according to their evidence, the households that cut down expenditure more aggressively in the presence of uncertainty are those whose income is more at risk in periods of high uncertainty, which we have shown include less specialized workers.

5 Functional Local Projections

Just as for standard scalar aggregate variables the IRFs can be estimated by Local Projections (LPs, see [Jordà \(2005\)](#)), the response of distributions to the structural shock of interest can be estimated by Functional Local Projections (F-LPs). As a matter of fact, the reason why a F-VAR process was assumed for the joint dynamics of aggregate variables and the endogenous function \tilde{f}_t was to have a tractable model in terms of scalar variables y_t and α_t , which would allow us to infer the responses of the distribution $p_t(\cdot)$ to an uncertainty shock. However, this assumption is not essential, and nothing prevents the use of LPs to perform inference about the IRFs of α_t , which we defined as $IRF_{\alpha,j,d,h}$, and to map them to distributional IRFs as described in [Section 2.4](#). This is the approach we follow in this section to show that all the results we have discussed so far are not determined by the modelling choice described in [2.1](#) or by the priors discussed in [2.3](#).

Although the estimation of F-LPs is straightforward and can simply be done by OLS, two points are worth considering. First, as the steady state distribution, $p_{ss}(\cdot)$, cannot be implied by the model, we take it to be the one associated with the sample average $\bar{f}(\cdot)$, so that $p_{ss}(x) = g^{-1}(\bar{f}(c))$. Second, unlike in standard LP applications, we are not interested in the IRF of one variable at a time. Instead, we are interested in the joint response of the whole vector α_t , $IRF_{\alpha,j,d,h}$, which then determines the distributional IRF as: $p_{ss+h}(\cdot) - p_{ss}(\cdot) = g^{-1}(\zeta(\cdot)'IRF_{\alpha,j,d,h} + \bar{f}) - p_{ss}(\cdot)$. As a result, when conducting frequentist inference, the variance of the estimator of $IRF_{\alpha,j,d,h}$ must account not only for the presence of serial correlation in the residuals arising in standard LP inference, but also for the cross correlation across residuals associated with the different α_{tk} 's. [Newey and West \(1987\)](#) covariance estimator is therefore not enough, [Driscoll and Kraay \(1998\)](#)'s method is required in this context.

Furthermore, since the shock identification strategy used in this paper only relies on timing assumptions, it is easy to incorporate them in the estimation of the LP for α_t , as explained by [Plagborg-Møller and Wolf \(2021\)](#). More specifically, continuing to assume that the uncertainty shock is ordered last in Cholesky-type identification scheme,

$IRF_{\alpha,N,1,h}$ is given by the $\beta_{1,U}^h$ coefficient in the following regression:

$$\alpha_{t+h} = a^h + B_1^h \left[y'_{t,\setminus U}, \alpha'_t \right]' + \beta_{1,U}^h \mathbf{U}_{m,t} + \sum_{l=1}^p B_{l+1}^h \left[y'_{t-l}, \alpha'_{t-l} \right]' + e_{h,t}, \quad (13)$$

where we have defined as $y_{t,\setminus U}$ the vector of all aggregate variables except $\mathbf{U}_{m,t}$.

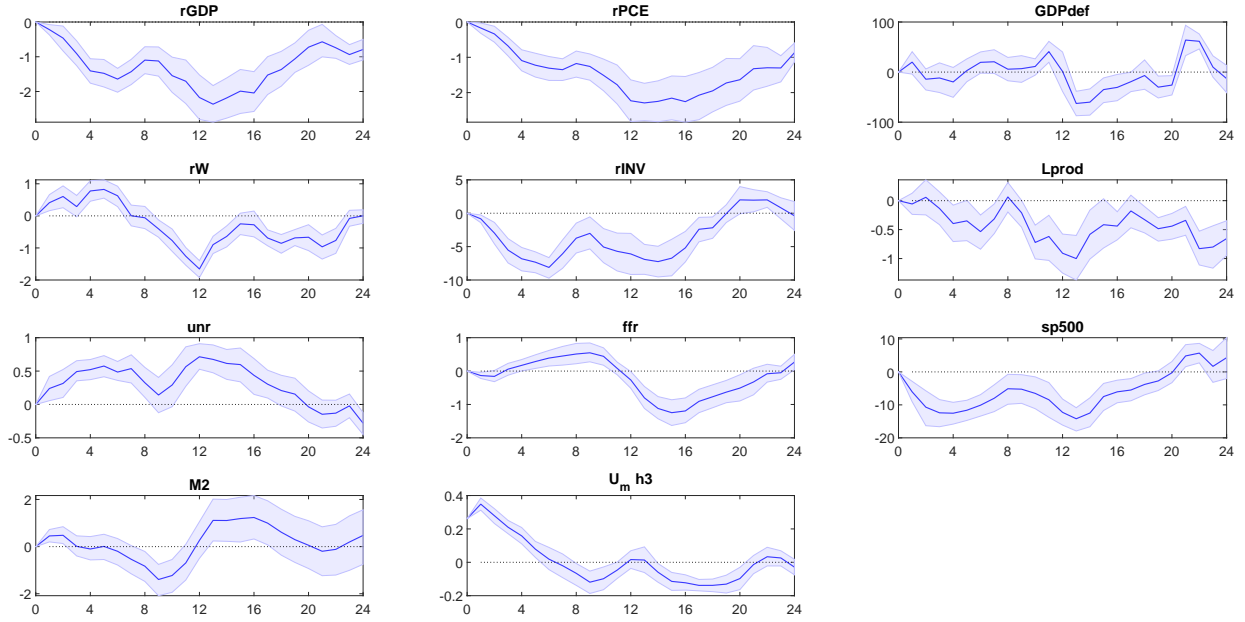
Since in this section we perform inference employing frequentist methods, we cannot make use of priors to shrink irrelevant parameters toward zero.¹⁷ Furthermore, as there are 10 variables in the vector $y_{t,\setminus U}$, the number of parameters to estimate in (13) can become excessively large when we consider many lags, p , and many FPCs, K . For this reason, in this section we set $p = 1$ and $K = 3$ when focusing on the earnings distribution, and $p = 1$ and $K = 5$ when focusing on the consumption distribution.

To demonstrate that F-LPs are a viable alternative to the F-SVAR, we repeated all the three artificial data experiments described in Section 3 performing inferences about $IRF_{\alpha,N,1,h}$ via LP. In the interest of space, the resulting figures are reported in Appendix H. The results show that also F-LPs do a good job in estimating the responses of interest. Moreover, the figures in Appendix H also suggest that the well known variance-bias trade-off that characterizes the comparison between standard LPs and standard SVARs carries over to the comparison between both approaches. In fact, while at short horizons the uncertainty around F-LP-based estimates is very limited, at long horizons the variance of the estimator grows substantially. This does not hold for the F-SVAR in which the credible bands never inflate excessively. Such comparison is definitely worth more research and we leave a thorough assessment for future work.

Additionally, in order to show that the regression in (13) actually allows us to estimate the response to the same shock studied in the previous section, in Figure 17 we report the IRFs of the macro aggregates estimated by LP. Apart from erratic behaviours peculiar to LP, the responses are very similar to those estimated through the F-SVAR and shown in Figure 10.

¹⁷Although performing Bayesian inference on F-LP coefficients is possible following the approach of Ferreira et al. (2023), eliciting priors for such parameters and accounting for the autocorrelation of projection residuals are not standard tasks. For these reasons, we choose to apply standard frequentist methods and to leave the Bayesian treatment of F-LPs for future research.

Figure 17: Impulse response of the macroeconomic aggregates to a one standard deviation uncertainty shock



Notes: The solid blue lines represent the posterior median response, while dashed blue lines delimit the 68% credible bands. The horizon on the horizontal axis is expressed in quarters.

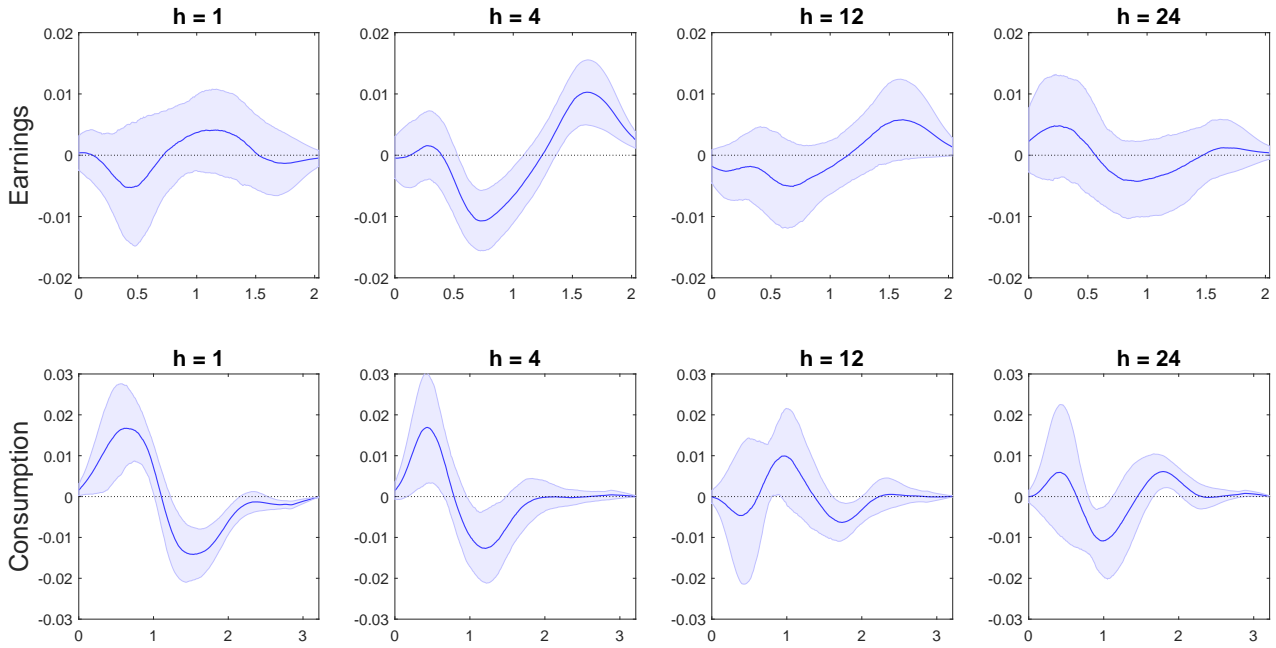
Confident that the LP in (13) is able to retrieve the responses to the uncertainty shock of interest, Figure 18 depicts the distributional IRFs estimated by F-LPs. Although there are some differences with respect to Figure 11, the general features highlighted by the two methods are very similar, especially for horizons up to one year.

6 Conclusion

Since the distributional consequences of uncertainty shocks have remained relatively unexplored by the literature, we introduced a Functional VAR model to study the propagation of such shocks on the earnings and consumption distributions.

After showing that the adopted econometric technique is able to reproduce distributional developments in response of structural shocks in a variety of challenging simulated experiments, we used the Functional VAR methodology to extend the popular analysis of Jurado et al. (2015) and to take into account the dynamics of the earnings and consump-

Figure 18: Distributional LP



Notes: The first (second) row of the figure shows the difference between the earnings (consumption) distribution after a standard deviation shock and the one prevailing in the steady state. The solid blue lines represent the posterior median response, while dashed blue lines delimit the 68% credible bands. h denotes the horizon at which the response is measured. In the first (second) row, the measure on the horizontal axis is the earnings-to-GDP per-capita (households consumption - to - per-capita consumption) ratio.

tion distributions. In particular, the Functional VAR we employ is based on a Functional Principal Component Analysis of the Log Quantile Density associated with the distribution of earnings and consumption, which is in turn estimated starting from available survey data using kernel methods. Furthermore, we also showed that the same type of analysis can be easily carried out estimating Functional LP.

Our results show that the effects of uncertainty shocks unfold in two phases. In a first phase, the economic contraction triggered by the shock is coupled with an increase of unemployment and a drop in stock prices. At this stage, the fraction of employed people earning less than the GDP per capita level decreases, and the share of higher-income workers increases. In a second phase, the unemployment level is gradually reabsorbed and labour productivity decreases. At the same time, the fraction of employed low-income people increases, while the employed medium-income class is reduced in size. As a consequence, the uncertainty shock reduces significantly the total level of earnings

inequality among employed people in the short run, as measured by the Gini coefficient, while in the medium term it causes a considerable rise in it,

A potential explanation for this pattern is that, in the short run, a larger proportion of less specialized low-income workers are laid off due to the economic contraction, while those that are not laid off see their earnings kept constant. In the longer run, however, the pool of unemployed people is reabsorbed, but labour productivity declines due to the previous foregone investments. The lower labor productivity can in turn explain the increase in the share of low-income employed relative to the steady state distribution, at the expense of the middle class.

References

- AHN, S., G. KAPLAN, B. MOLL, T. WINBERRY, AND C. WOLF (2018): “When inequality matters for macro and macro matters for inequality,” *NBER macroeconomics annual*, 32(1), 1–75.
- ANDERSON, E., A. INOUE, AND B. ROSSI (2016): “Heterogeneous consumers and fiscal policy shocks,” *Journal of Money, Credit and Banking*, 48(8), 1877–1888.
- ATHREYA, K., A. OWENS, AND F. SCHWARTZMAN (2017): “Does redistribution increase output? The centrality of labor supply,” *Quantitative Economics*, 8(3), 761–808.
- ATTANASIO, O. P., AND L. PISTAFERRI (2016): “Consumption inequality,” *Journal of Economic Perspectives*, 30(2), 3–28.
- BAI, J., AND S. NG (2002): “Determining the number of factors in approximate factor models,” *Econometrica*, 70(1), 191–221.
- BAYER, C., B. BORN, AND R. LUETTICKE (2020): “Shocks, frictions, and inequality in US business cycles,” *CEPR Discussion Paper No. DP14364*.
- BERNANKE, B. S., J. BOIVIN, AND P. ELIASZ (2005): “Measuring the effects of monetary policy: a factor-augmented vector autoregressive (FAVAR) approach,” *The Quarterly journal of economics*, 120(1), 387–422.
- BJØRNLAND, H. C., Y. CHANG, AND J. CROSS (2023): “Oil and the Stock Market Revisited: A mixed functional VAR approach,” *Available at SSRN*.
- BLOOM, N. (2009): “The impact of uncertainty shocks,” *Econometrica*, 77(3), 623–685.
- BOSQ, D. (2000): *Linear processes in function spaces: theory and applications*, vol. 149. Springer Science & Business Media.
- CARRIERO, A., M. MARCELLINO, AND T. TORNESE (2023): “Macro uncertainty in the long run,” *Economics Letters*, 225, 111067.
- CASTELNUOVO, E. (2019): “Domestic and global uncertainty: A survey and some new results,” *CAMA Working Paper*.
- CHANG, M., X. CHEN, AND F. SCHORFHEIDE (2024): “Heterogeneity and aggregate fluctuations,” *Journal of Political Economy*.
- CHANG, M., AND F. SCHORFHEIDE (2024): “On the Effects of Monetary Policy Shocks on

- Income and Consumption Heterogeneity,” Discussion paper, National Bureau of Economic Research.
- CHANG, Y., C. S. KIM, AND J. Y. PARK (2016): “Nonstationarity in time series of state densities,” *Journal of Econometrics*, 192(1), 152–167.
- CHOI, S., AND J. PHI (2023): “Impact of Uncertainty Shocks on Income and Wealth Inequality,” *CAMA Working Paper 33/2023*.
- COIBION, O., D. GEORGARAKOS, Y. GORODNICHENKO, G. KENNY, AND M. WEBER (2024): “The effect of macroeconomic uncertainty on household spending,” *American Economic Review*, 114(3), 645–677.
- DE GIORGI, G., AND L. GAMBETTI (2017): “Business cycle fluctuations and the distribution of consumption,” *Review of Economic Dynamics*, 23, 19–41.
- DIEBOLD, F. X., AND C. LI (2006): “Forecasting the term structure of government bond yields,” *Journal of econometrics*, 130(2), 337–364.
- DOAN, T., R. LITTERMAN, AND C. SIMS (1984): “Forecasting and conditional projection using realistic prior distributions,” *Econometric Reviews*, 3(1), 1–100.
- DRISCOLL, J. C., AND A. C. KRAAY (1998): “Consistent covariance matrix estimation with spatially dependent panel data,” *Review of economics and statistics*, 80(4), 549–560.
- FERNÁNDEZ-VILLAYERDE, J., AND P. A. GUERRÓN-QUINTANA (2020): “Uncertainty shocks and business cycle research,” *Review of economic dynamics*, 37, S118–S146.
- FERREIRA, L. N., S. MIRANDA-AGRIPPINO, AND G. RICCO (2023): “Bayesian local projections,” *Review of Economics and Statistics*, pp. 1–45.
- FISCHER, M. M., F. HUBER, AND M. PFARRHOFER (2021): “The regional transmission of uncertainty shocks on income inequality in the United States,” *Journal of Economic Behavior & Organization*, 183, 887–900.
- GIANNONE, D., M. LENZA, AND G. E. PRIMICERI (2015): “Prior selection for vector autoregressions,” *Review of Economics and Statistics*, 97(2), 436–451.
- HEATHCOTE, J., F. PERRI, AND G. L. VIOLANTE (2010a): “Unequal we stand: An empirical analysis of economic inequality in the United States, 1967–2006,” *Review of Economic dynamics*, 13(1), 15–51.
- HEATHCOTE, J., K. STORESLETTEN, AND G. L. VIOLANTE (2010b): “The macroeconomic

- implications of rising wage inequality in the United States,” *Journal of political economy*, 118(4), 681–722.
- HU, B., AND J. Y. PARK (2016): “Econometric analysis of functional dynamics in the presence of persistence,” Discussion paper, Mimeo, Indiana University.
- HUANG, J. (2023): “Functional VAR,” *Manuscript*.
- INOUE, A., AND B. ROSSI (2021): “A new approach to measuring economic policy shocks, with an application to conventional and unconventional monetary policy,” *Quantitative Economics*, 12(4), 1085–1138.
- JORDÀ, Ò. (2005): “Estimation and inference of impulse responses by local projections,” *American economic review*, 95(1), 161–182.
- JURADO, K., S. C. LUDVIGSON, AND S. NG (2015): “Measuring uncertainty,” *American Economic Review*, 105(3), 1177–1216.
- KAPLAN, G., AND G. L. VIOLANTE (2018): “Microeconomic heterogeneity and macroeconomic shocks,” *Journal of Economic Perspectives*, 32(3), 167–94.
- KILIAN, L., AND H. LÜTKEPOHL (2017): *Structural vector autoregressive analysis*. Cambridge University Press.
- KNEIP, A., AND K. J. UTIKAL (2001): “Inference for density families using functional principal component analysis,” *Journal of the American Statistical Association*, 96(454), 519–542.
- KOKOSZKA, P., H. MIAO, A. PETERSEN, AND H. L. SHANG (2019): “Forecasting of density functions with an application to cross-sectional and intraday returns,” *International Journal of Forecasting*, 35(4), 1304–1317.
- KOOP, G., D. KOROBILIS, ET AL. (2010): “Bayesian multivariate time series methods for empirical macroeconomics,” *Foundations and Trends® in Econometrics*, 3(4), 267–358.
- KOWAL, D. R., D. S. MATTESON, AND D. RUPPERT (2017): “A Bayesian multivariate functional dynamic linear model,” *Journal of the American Statistical Association*, 112(518), 733–744.
- KRUSELL, P., AND A. A. SMITH, JR (1998): “Income and wealth heterogeneity in the macroeconomy,” *Journal of political Economy*, 106(5), 867–896.
- MEEKS, R., AND F. MONTI (2022): “Heterogeneous beliefs and the Phillips curve,” *Available at SSRN*.

- MUMTAZ, H., AND A. THEOPHILOPOULOU (2017): “The impact of monetary policy on inequality in the UK. An empirical analysis,” *European Economic Review*, 98, 410–423.
- NEWKEY, W. K., AND K. D. WEST (1987): “A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix,” *Econometrica: Journal of the Econometric Society*, pp. 703–708.
- PETERSEN, A., AND H.-G. MULLER (2016): “Functional data analysis for density functions by transformation to a Hilbert space,” *Annals of Statistics*, 44, 183–218.
- PETERSEN, A., C. ZHANG, AND P. KOKOSZKA (2022): “Modeling Probability Density Functions as Data Objects,” *Econometrics and Statistics*, 21, 159–178.
- PLAGBORG-MØLLER, M., AND C. K. WOLF (2021): “Local projections and VARs estimate the same impulse responses,” *Econometrica*, 89(2), 955–980.
- RAMSAY, J. O., AND B. W. SILVERMAN (2002): *Applied functional data analysis: methods and case studies*. Springer.
- ROGNLIE, M., A. AUCLERT, ET AL. (2019): “Inequality and aggregate demand,” Discussion paper, Working Paper 24280, National Bureau of Economic Research, February 2018
- SILVERMAN, B. W. (2018): *Density estimation for statistics and data analysis*. Routledge.
- SIMS, C. A., AND T. ZHA (1998): “Bayesian methods for dynamic multivariate models,” *International Economic Review*, pp. 949–968.
- STOCK, J. H., AND M. W. WATSON (2016): “Dynamic factor models, factor-augmented vector autoregressions, and structural vector autoregressions in macroeconomics,” in *Handbook of macroeconomics*, vol. 2, pp. 415–525. Elsevier.
- THEOPHILOPOULOU, A. (2022): “The impact of macroeconomic uncertainty on inequality: An empirical study for the United Kingdom,” *Journal of Money, Credit and Banking*, 54(4), 859–884.
- TSAY, R. S. (2016): “Some methods for analyzing big dependent data,” *Journal of Business & Economic Statistics*, 34(4), 673–688.
- WANG, J.-L., J.-M. CHIOU, AND H.-G. MÜLLER (2016): “Functional data analysis,” *Annual Review of Statistics and its application*, 3, 257–295.

Online Appendix

A F-SVAR DGP

The reduced form parameters of the VARs used to simulate the data in Section 3.1 and 3.2 are:

$$\Pi_0 = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix};$$

$$\Pi_1 = \begin{bmatrix} 0.85 & -0.15 & 0.15 & 0.15 & -0.25 \\ -0.2 & 0.85 & -0.15 & -0.25 & -0.2 \\ 0.15 & -0.15 & 0.85 & -0.15 & 0 \\ 0.1 & 0.15 & -0.2 & 0.85 & -0.2 \\ -0.25 & 0.15 & 0.15 & 0.15 & 0.85 \end{bmatrix};$$

$$\Pi_2 = \begin{bmatrix} -0.3 & 0.1 & 0.15 & -0.15 & -0.1 \\ -0.1 & -0.3 & 0.1 & 0.15 & 0.15 \\ -0.05 & 0.1 & -0.3 & 0.05 & -0.1 \\ 0.15 & -0.1 & -0.05 & -0.3 & -0.05 \\ 0.15 & -0.15 & 0.1 & -0.1 & -0.3 \end{bmatrix};$$

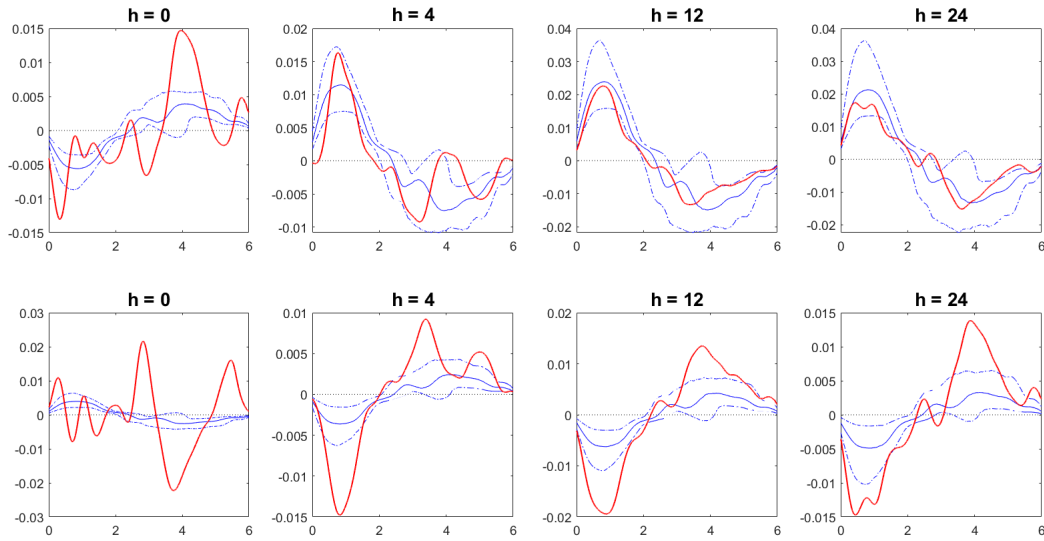
$$\Pi_3 = \begin{bmatrix} 0.15 & 0 & 0 & 0 & 0 \\ 0 & 0.15 & 0 & 0 & 0 \\ 0 & 0 & 0.15 & 0 & 0 \\ 0 & 0 & 0 & 0.15 & 0 \\ 0 & 0 & 0 & 0 & 0.15 \end{bmatrix};$$

$$\Pi_4 = \begin{bmatrix} 0.05 & 0 & 0 & 0 & 0 \\ 0 & 0.05 & 0 & 0 & 0 \\ 0 & 0 & 0.05 & 0 & 0 \\ 0 & 0 & 0 & 0.05 & 0 \\ 0 & 0 & 0 & 0 & 0.05 \end{bmatrix};$$

The *Omega* matrix is formed sampling 100 random matrices R with elements equal to the product of two standard normal draws multiplied by 0.1, and setting $\Omega = R * R'$. the matrix A_0^{-1} is then set to be the lower-triangular Cholesky factor of Ω .

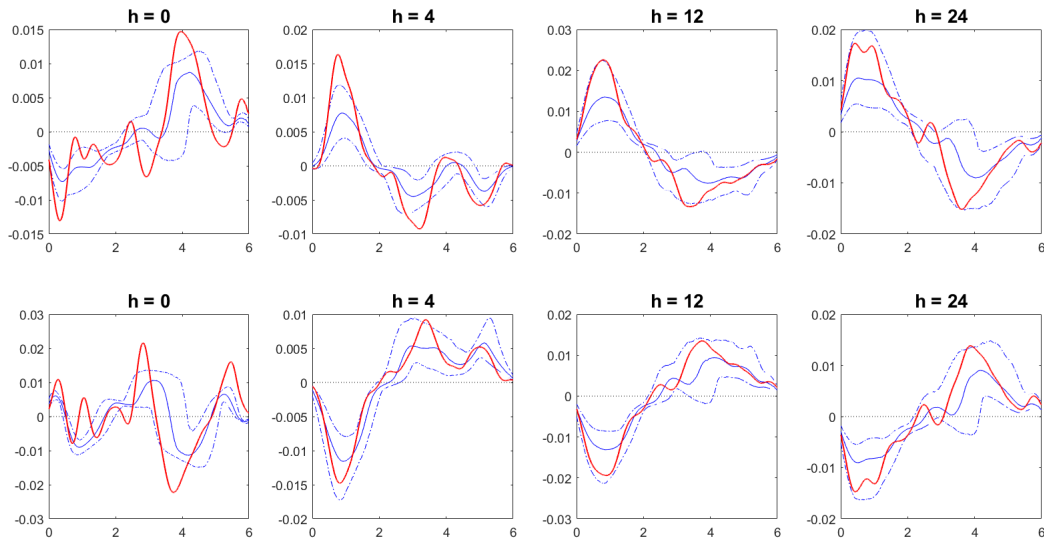
B F-SVAR DGP 1 - Results with Different K

Figure A1: Functional IRFs - F-SVAR DGP 1 - $K = 1$




Notes: Red lines show the true responses of $p_t(\xi)$ to one standard deviation ε_1 and ε_2 . The solid blue lines represent the posterior median response, while dashed blue lines delimit the 90% credible bands. h denotes the horizon at which the response is measured.

Figure A2: Functional IRFs - F-SVAR DGP 1 - $K = 2$




Notes: Red lines show the true responses of $p_t(\xi)$ to one standard deviation ε_1 and ε_2 . The solid blue lines represent the posterior median response, while dashed blue lines delimit the 90% credible bands. h denotes the horizon at which the response is measured.

Figure A3: Functional IRFs - F-SVAR DGP 1 - $K = 3$



Notes: Red lines show the true responses of $p_t(\xi)$ to one standard deviation ε_1 and ε_2 . The solid blue lines represent the posterior median response, while dashed blue lines delimit the 90% credible bands. h denotes the horizon at which the response is measured.

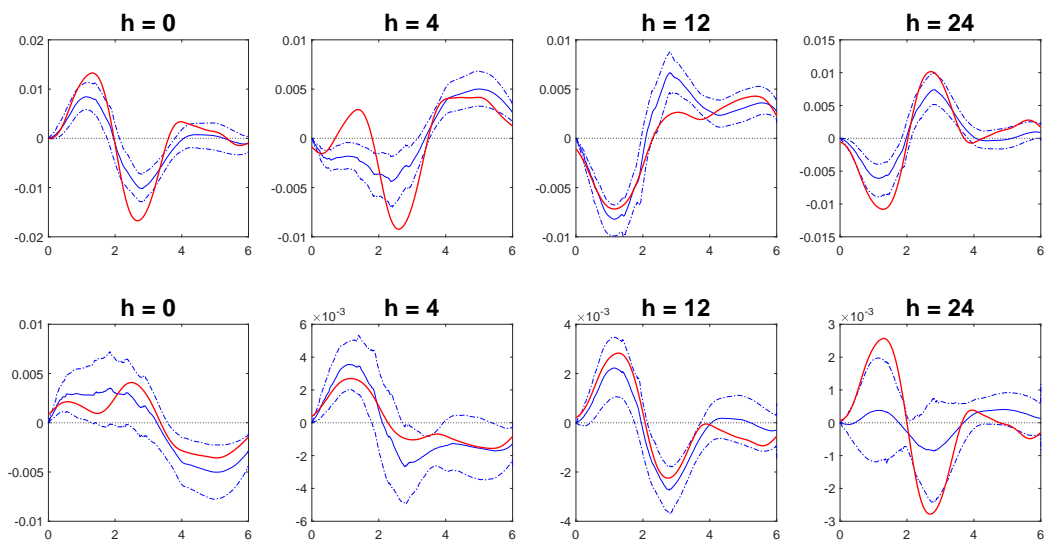
Figure A4: Functional IRFs - F-SVAR DGP 1 - $K = 15$



Notes: Red lines show the true responses of $p_t(\xi)$ to one standard deviation ε_1 and ε_2 . The solid blue lines represent the posterior median response, while dashed blue lines delimit the 90% credible bands. h denotes the horizon at which the response is measured.

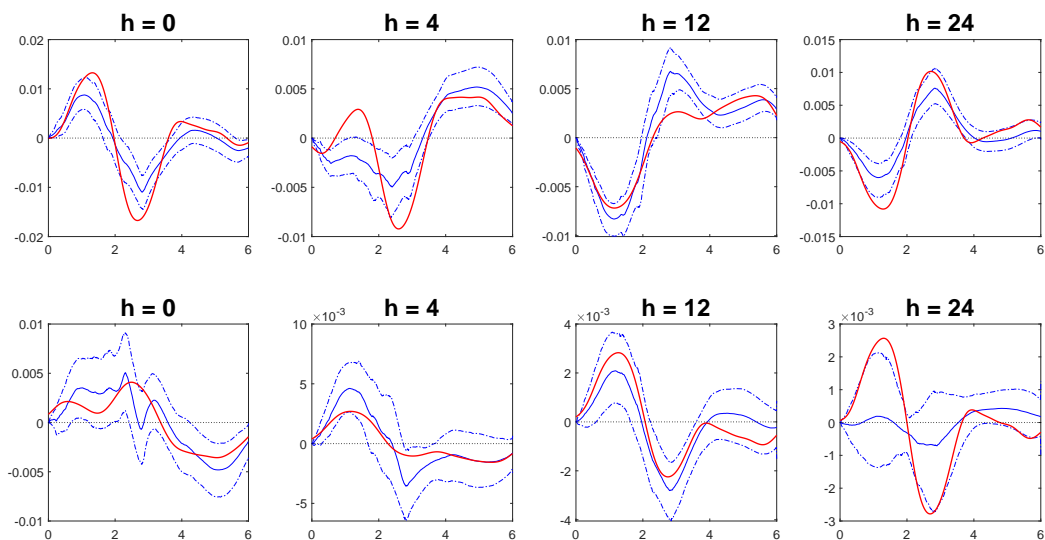
C F-SVAR DGP 2 - Results with Different K

Figure A5: Functional IRFs - F-SVAR DGP 2 - $K = 3$



Notes: Red lines show the true responses of $p_t(\xi)$ to one standard deviation ε_1 and ε_2 . The solid blue lines represent the posterior median response, while dashed blue lines delimit the 90% credible bands. h denotes the horizon at which the response is measured.

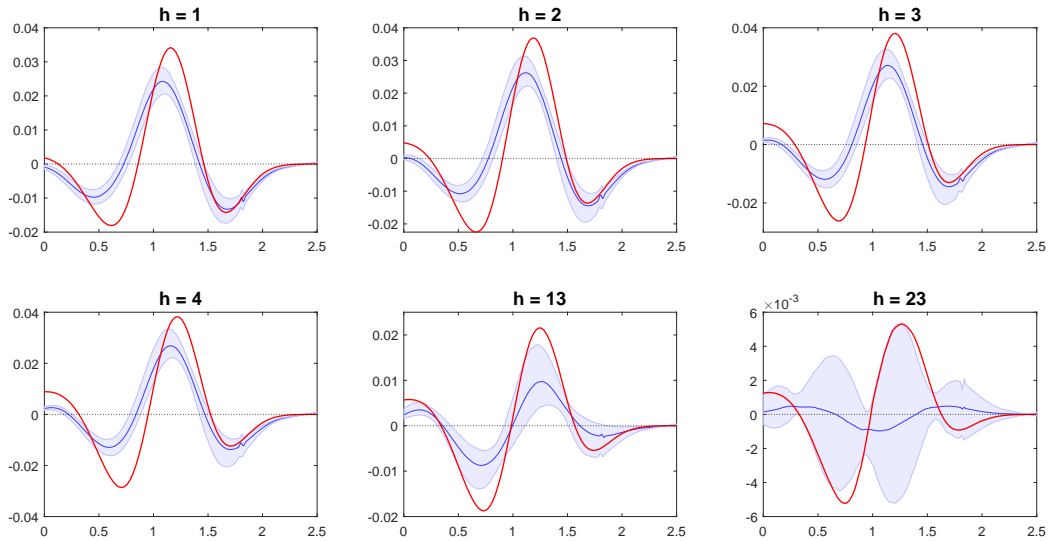
Figure A6: Functional IRFs - F-SVAR DGP 2 - $K = 15$



Notes: Red lines show the true responses of $p_t(\xi)$ to one standard deviation ε_1 and ε_2 . The solid blue lines represent the posterior median response, while dashed blue lines delimit the 90% credible bands. h denotes the horizon at which the response is measured.

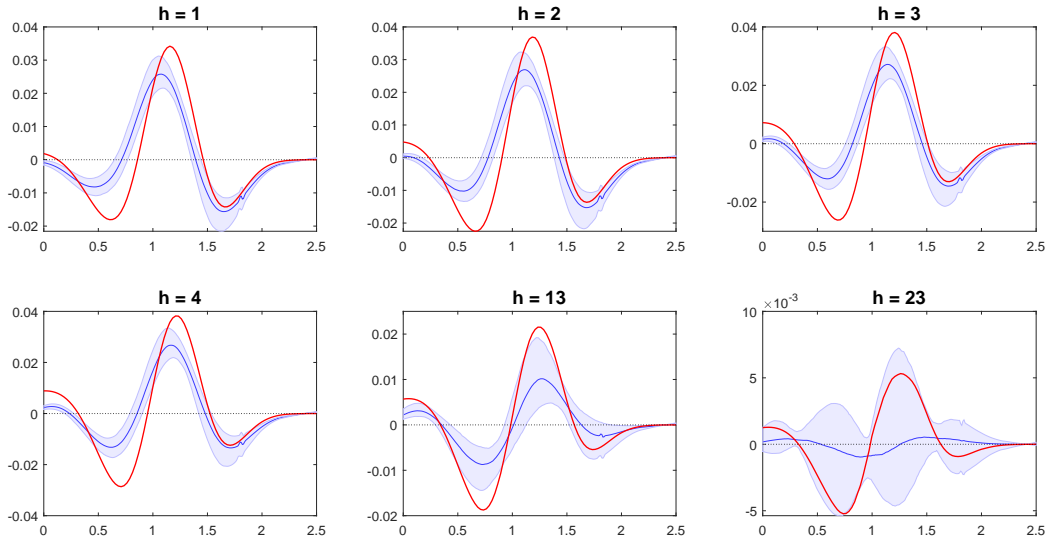
D Krusell and Smith (1998) DGP - Results with Different K

Figure A7: Functional IRFs - Krusell and Smith (1998) DGP - $K = 3$



Notes: Red lines show the true responses of the asset distribution to a one standard deviation productivity shock. The solid blue lines represent the posterior median response, while dashed blue lines delimit the 90% credible bands. h denotes the horizon at which the response is measured. Notice that the horizons for which responses are reported are the same as those in Figure 6 of Chang et al. (2024), the difference in the panels titles only stems from different timing conventions.

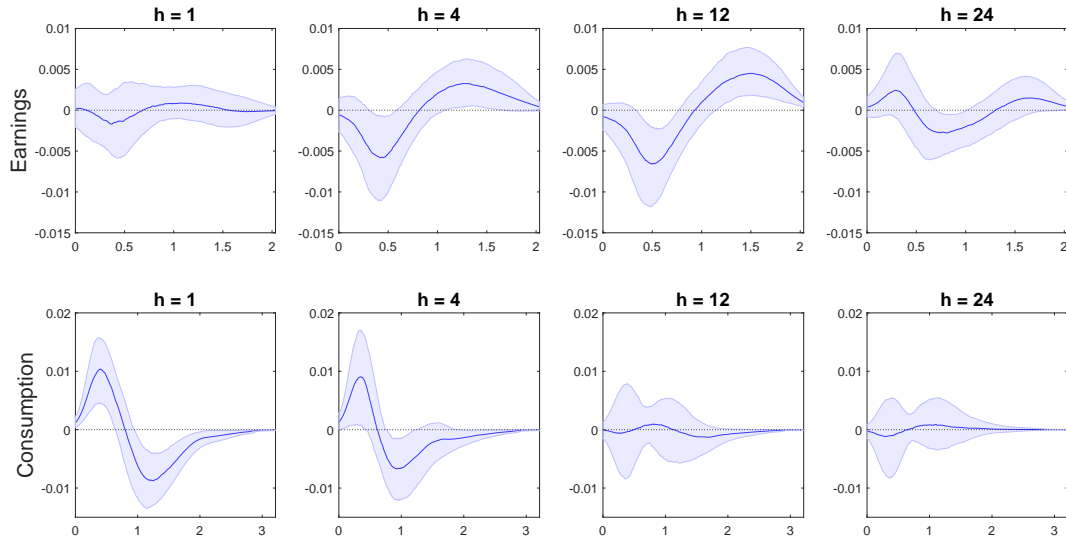
Figure A8: Functional IRFs - Krusell and Smith (1998) DGP - $K = 15$



Notes: Red lines show the true responses of the asset distribution to a one standard deviation productivity shock. The solid blue lines represent the posterior median response, while dashed blue lines delimit the 90% credible bands. h denotes the horizon at which the response is measured. Notice that the horizons for which responses are reported are the same as those in Figure 6 of Chang et al. (2024), the difference in the panels titles only stems from different timing conventions.

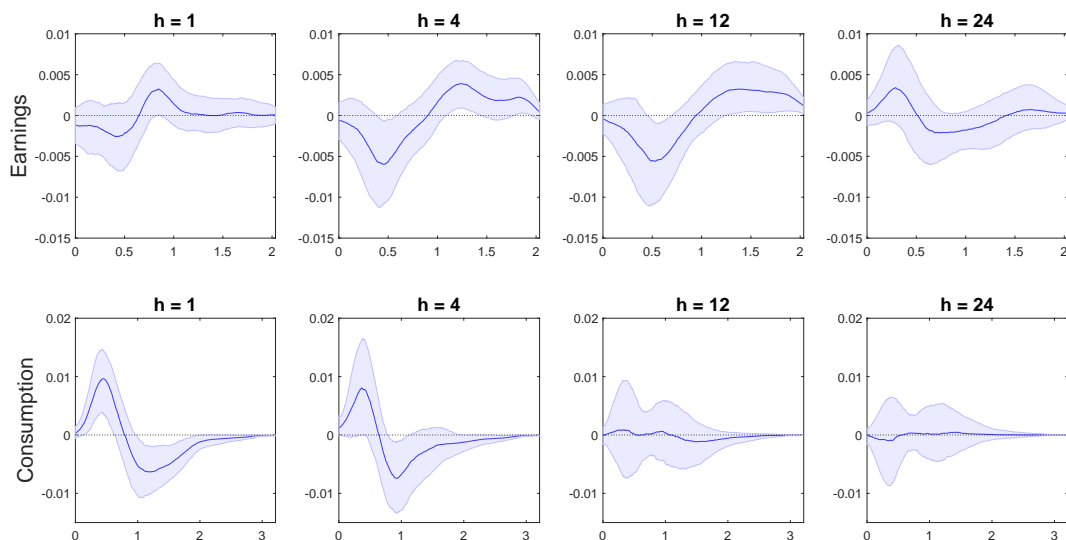
E Uncertainty Shocks - Results with Different K

Figure A9: Functional IRFs - Effects of Uncertainty Shocks on Earnings and Consumption Distribution - $K = 3$ and $K = 5$ respectively



Notes: The solid blue lines represent the posterior median response, while dashed blue lines delimit the 68% credible bands. h denotes the horizon at which the response is measured.

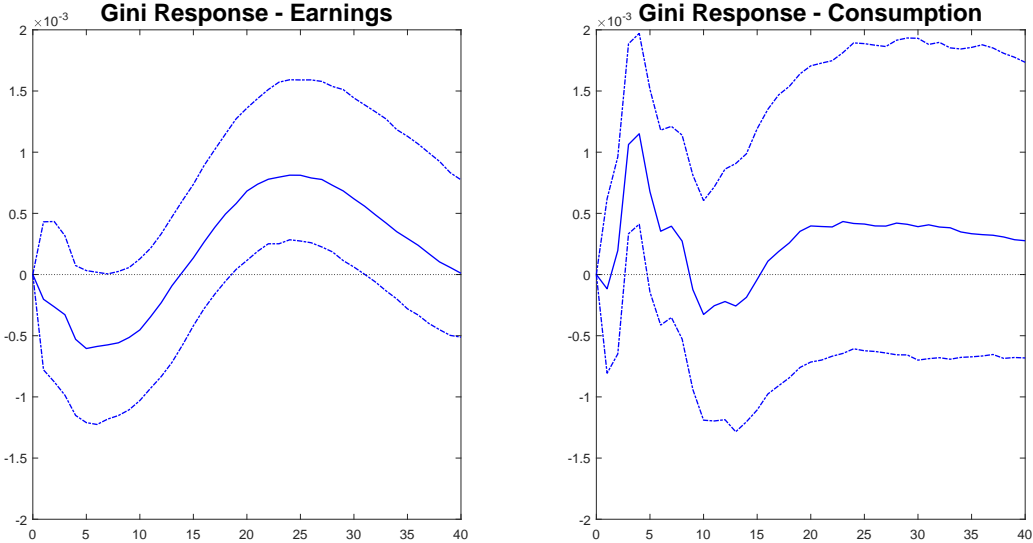
Figure A10: Functional IRFs - Effects of Uncertainty Shocks on Earnings and Consumption Distribution - $K = 15$



Notes: The solid blue lines represent the posterior median response, while dashed blue lines delimit the 68% credible bands. h denotes the horizon at which the response is measured. The measure on the horizontal axis is the earnings-to-DGP per-capita ratio.

F Standard SVAR Results

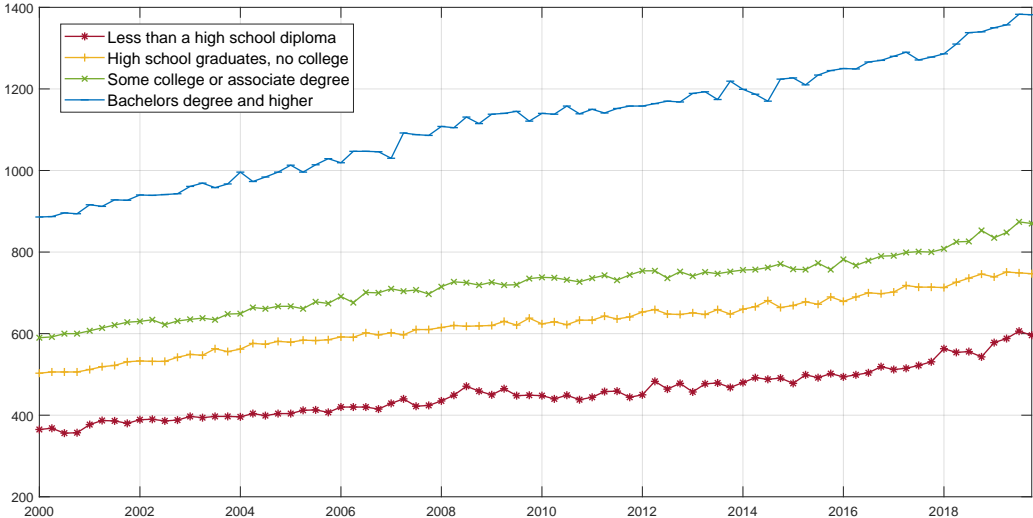
Figure A11



Notes: The solid red lines represent the posterior median response to a one-standard deviation shock, while dashed red lines delimit the 68% credible bands. The horizon is measured in quarters.

G IRFs Unemployment by Educational Attainment

Figure A12: Median weekly earnings

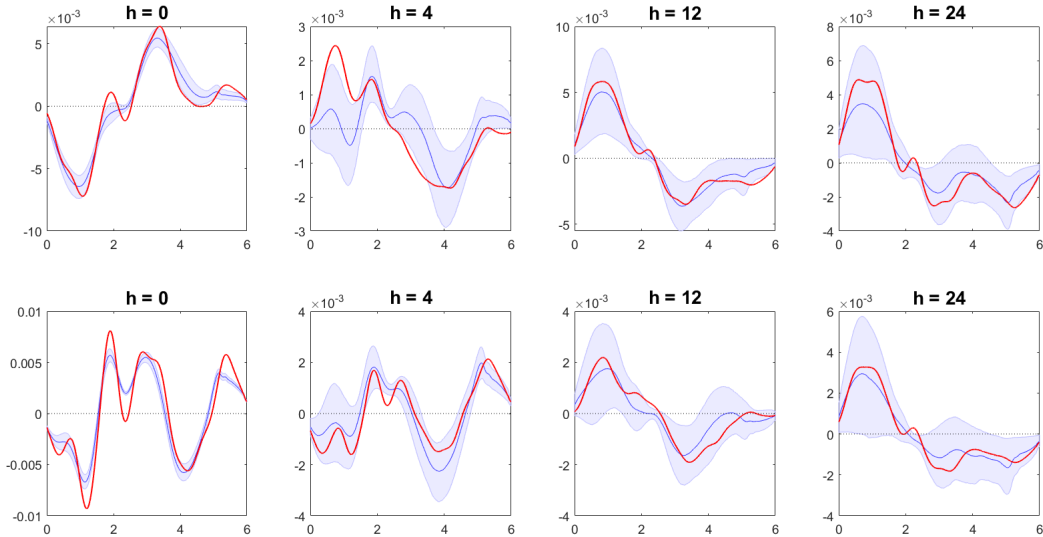


Notes: The lines show the median weekly earnings of the four categories.

H F-LP: Artificial Data Experiments

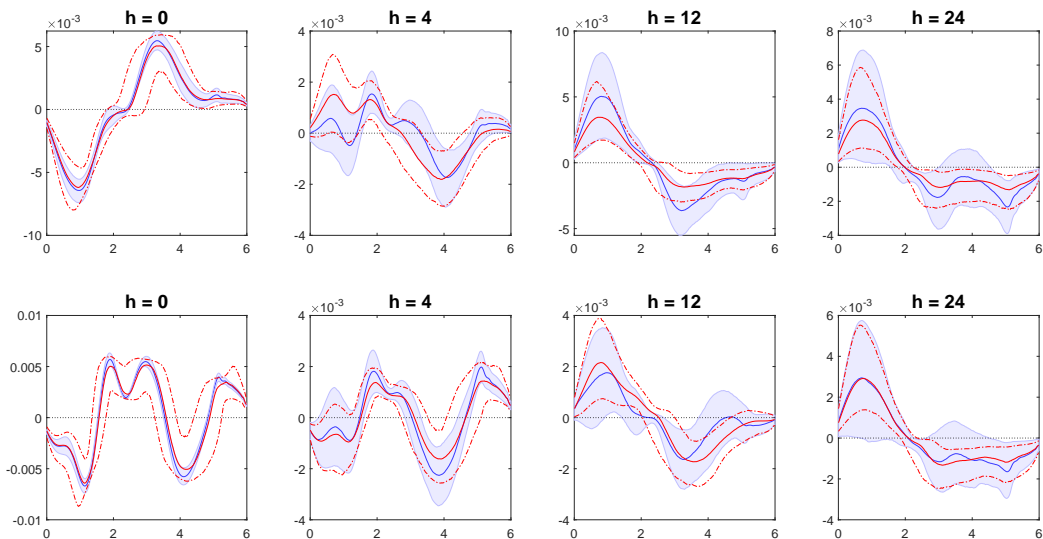
H.1 DGP 1

Figure A13: F-LP: DGP 1



Notes: Red lines show the true responses of $p_t(\xi)$ to one standard deviation shocks to ε_1 (upper panels) and ε_2 (lower panels). The solid blue lines represent the posterior median response, while dashed blue lines delimit the 90% credible bands. h denotes the horizon at which the response is measured.

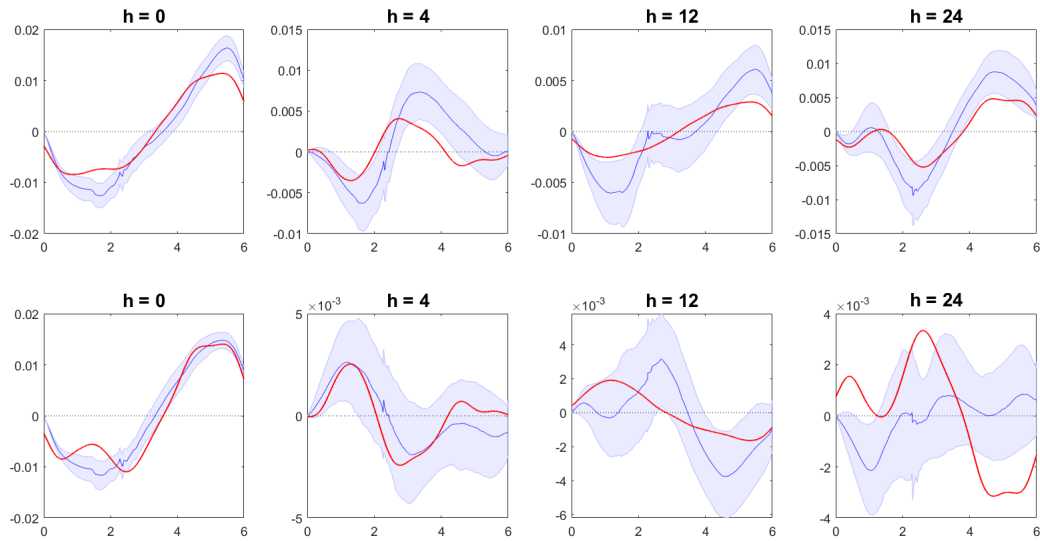
Figure A14: F-LP vs F-SVAR: DGP 1



Notes: The solid red line shows and the dashed red lines show the posterior median and the 90% credible bands of the responses inferred by the F-SVAR. The solid blue lines represent the estimated F-LP, while the light blue area delimits the associated 90% confidence interval.

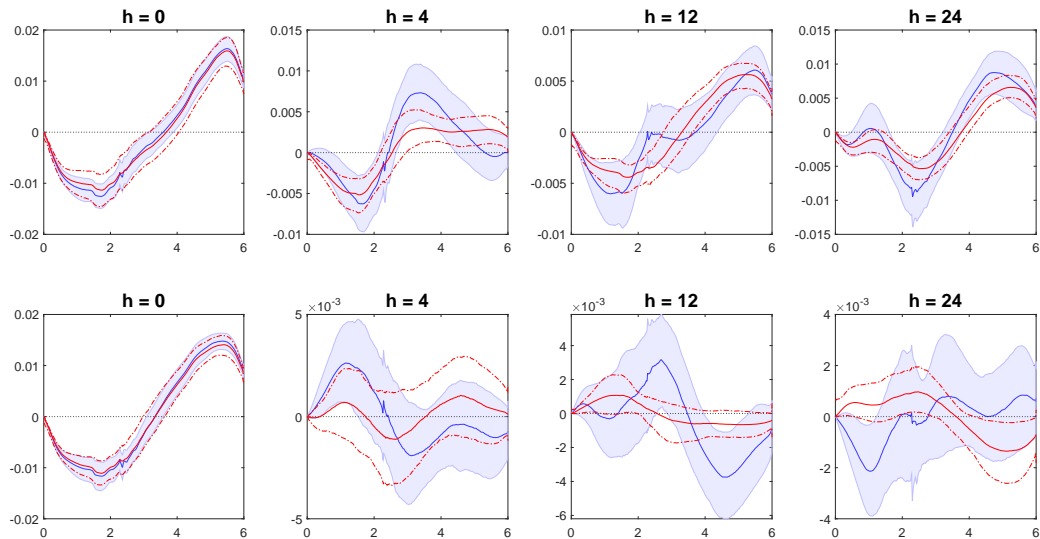
H.2 DGP 2

Figure A15: F-LP: DGP 2



Notes: Red lines show the true responses of $p_t(\xi)$ to one standard deviation ε_1 (upper panels) and ε_2 (lower panels). The solid blue lines represent the posterior median response, while dashed blue lines delimit the 90% credible bands. h denotes the horizon at which the response is measured.

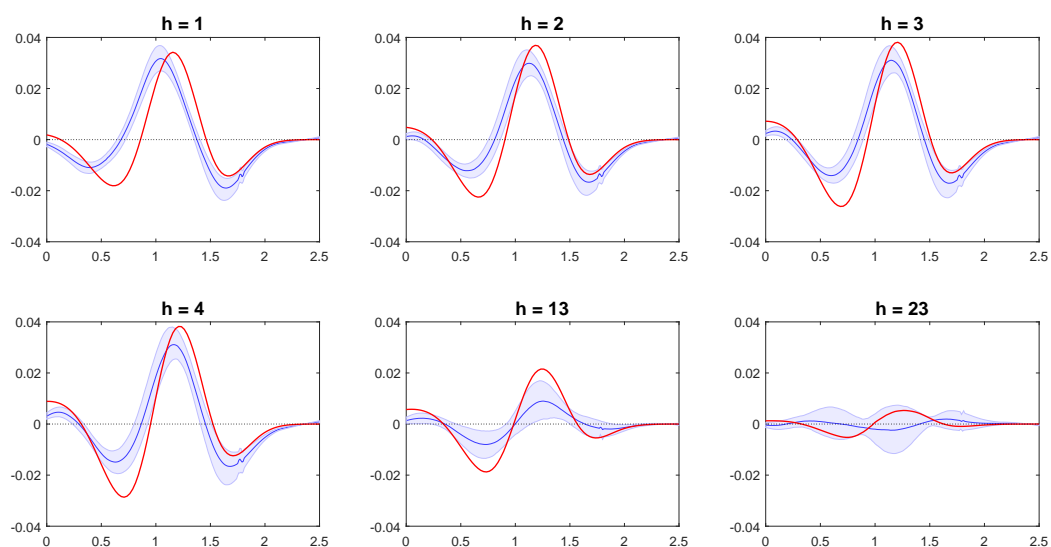
Figure A16: F-LP vs F-SVAR: DGP 2



Notes: The solid red line shows and the dashed red lines show the posterior median and the 90% credible bands of the responses inferred by the F-SVAR. The solid blue lines represent the estimated F-LP, while the light blue area delimits the associated 90% confidence interval.

H.3 Krusell and Smith (1998) DGP

Figure A17: F-LP: Krusell and Smith (1998) DGP



Notes: Red lines show the true responses of the asset distribution to a one standard deviation productivity shock. The solid blue lines represent the posterior median response, while dashed blue lines delimit the 90% credible bands. h denotes the horizon at which the response is measured. Notice that the horizons for which responses are reported are the same as those in Figure 6 of Chang et al. (2024), the difference in the panels titles only stems from different timing conventions.