# Oil, Inflation Expectations, and Household Characteristics: A Nonlinear Heterogeneous Agent VAR Approach

Christiane Baumeister
*University of Notre Dame*
*NBER and CEPR*

Pascal Frank
*University of Notre Dame*

Florian Huber
*University of Salzburg*

Gary Koop
*University of Strathclyde*

ECB-Cleveland Fed Conference on Inflation: Drivers and Dynamics

*September 29-30, 2025*

# Motivation

(1) Growing interest in how macro shocks affect the micro level

- Effects can differ considerably across economic agents
- Existing VAR approaches that use unit-level data:
  - functional VARs (Chang, Chen, and Schorfheide, *JPE* 2024)
  - cross-sectional-unit VAR (Ettmeier, Kim, and Schorfheide, 2025)
  - pseudo VAR (Koop, Mitchell, McIntyre, and Wu, 2025)
  - disaggregated VAR (Baumeister and Hamilton, 2024)

(2) Nonlinear transmission of aggregate shocks

- Sign and size asymmetries
- State dependence

# Contributions

(1) Develop flexible modeling framework that combines these two aspects

⇒ Nonlinear Heterogeneous Agent VAR (HAVAR) Model

- Multivariate time series model
- Nonlinear panel model

(2) Examine how households with different characteristics adjust their monthly inflation expectations in response to oil supply shocks of different sign and magnitude

- ➢ asymmetric effects of oil price shocks on
  - economic activity:
    Hamilton (1996, 2003), Davis and Haltiwanger (2001), Balke et al. (2002), Choi et al. (2018), Knotek and Zaman (2021), Cheikh et al. (2023), Caravello and Martinez-Bruera (2024), Miescu et al. (2024), Forni et al. (2025)
  - inflation:
    Hooker (2000), Hwang and Zhu (2024), Bobeica, Holton, Huber, and Martínez Hernández (2025)
- ➢ sensitivity of inflation expectations to energy price shocks
  Coibion and Gorodnichenko (2015), Wong (2015), Aastveit et al. (2022), Baumeister (2023)

# Outline

- Modeling framework
- Survey data
- Empirical results

# A Nonlinear HAVAR Model

- Capture interactions between
  - a set of $M$ macro aggregates $\boldsymbol{y_t}$
  - a dynamic panel $\boldsymbol{x_t}$ with $N$ units (households, firms, etc.)
- Our model consists of two inter-related blocks:
  - Block 1: Structural macro dynamics

$$\boldsymbol{A y_t} = \boldsymbol{B_1 y_{t-1}} + \cdots + \boldsymbol{B_p y_{t-p}} + \sum_{s=1}^{S} \boldsymbol{\beta_s f_{y,s}(x_{t-1})} + \boldsymbol{u_t},\ \boldsymbol{u_t} \sim \boldsymbol{N(0, D)}$$

  - Block 2: Micro dynamics

$$x_{it} = c_i + \rho_i x_{it-1} + f_x(\underbrace{\boldsymbol{y_t}, \dots, \boldsymbol{y_{t-p}}, \boldsymbol{z_i}}_{\boldsymbol{w_t}}) + \varepsilon_{it}, \qquad \varepsilon_{it} \sim N(0, \sigma_i^2)$$

$\Rightarrow$ two key components:

(1) $f_x$ denotes an unknown function of the macro series and $K$ unit-specific characteristics $\boldsymbol{z_i}$

(2) $f_{y,s}$ denotes pre-selected functions (e.g., mean, variance, skew) of the cross-sectional distribution to capture feedback into macro outcomes $\longrightarrow$ in the empirical application we set this to selected quantiles $Q_s(\boldsymbol{x_{t-1}})$ for s $\in$ $S = \{0.05, 0.10, \dots, 0.90, 0.95\}$

# Choosing $f_x$

- Impose parametric assumptions

  - Linear model

    The simplest specification would be a model that is linear in $\boldsymbol{w_t}$:

    $$f_x(\boldsymbol{w_t}, \boldsymbol{z_i}) = \psi(\boldsymbol{z_i})' \boldsymbol{w_t}$$

    where $\psi(\boldsymbol{z_i})$ is a vector of linear parameters that might depend on $\boldsymbol{z_i}$

    $\Rightarrow$ Implication: Shocks to $\boldsymbol{y_t}$ trigger linear reactions of $x_{it+h}$ for $h = 0, \dots, H$

  - Non-linear model

    Requires to take a stance on the functional form of $f_x$

    $\Rightarrow$ Implications: * hard to justify specific form of non-linearity at the micro level
    * risks model mis-specification

    While it would be possible to entertain different forms of non-linearities and discriminate between them using criteria of model adequacy, this is computationally cumbersome.

# Choosing $f_x$

- Model non-linearities non-parametrically

  $\Rightarrow$ treat $f_x$ as unknown, place a prior on it, and estimate it

  - This approach can be interpreted as a prior over the space of unknown functions.
  - Nonparametric machine-learning methods that can be used to approximate $f_x$:
    - Gaussian processes (Hauzenberger et al., JBES 2025)
    - Neural networks (Farrell et al., Ecma 2021)
    - Bayesian Additive Regression Trees (Chipman, George, and McCulloch, AoAS 2010)

**BUT** unrestricted modeling of $f_x$ is computationally and statistically challenging (in particular, if $N$ and $T$ are large)

$\Rightarrow$ we offer a novel solution based on the finance literature

# Introducing Additional Restrictions

- Assume a factor structure in $\boldsymbol{x_t}$ with factor loadings being a function of $\boldsymbol{z_i}$

  (Kelly, Pruit, and Su, JFE 2019; Gu, Kelly, and Su, JoE 2021)

- Yields a semi-parametric specification

$$f_x(\boldsymbol{w_t}, \boldsymbol{z_i}) = \underbrace{\mathbf{z}_i'\gamma}_{\gamma_i}\, q_t\,, \qquad q_t \sim N\left(\mu(\boldsymbol{w_t}), \tau_q^2\right), \qquad \tau_q^2 \sim G^{-1}(c, d)$$

where

- $\gamma$ is a $K \times 1$ vector of loadings associated with latent factor $q_t$
- $\mu$ is a (unknown) function that takes $\boldsymbol{w_t}$ as input
- $\tau_q^2$ is a variance parameter

# Inferring $\mu(\boldsymbol{w_t})$

- The state equation of $q_t$ can be written as:

$$q_t = \mu(\boldsymbol{w_t}) + \eta_t\,, \quad \eta_t \sim N(0, \tau_q^2)$$

- We use Bayesian Additive Regression Trees (BART) to infer $\mu(\boldsymbol{w_t})$

  → good empirical performance: Huber et al. (JoE 2021), Clark et al. (IER 2022), Clark et al. (AoAS 2023) Clark et al. (JBES 2024), Baumeister et al. (2025)

- BART approximates using a sum-of-trees model

$$\mu(\boldsymbol{w_t}) \approx \sum_{s=1}^{S} \mathfrak{t}(\boldsymbol{w_t} \mid T_s, m_s)$$

  where $\mathfrak{t}$ is a tree function parametrized by tree structures $T_s$ and the terminal node parameters $m_s$

  → Overfitting is alleviated using (regularization) priors for the tree structures and the terminal node parameters based on the standard setup of Chipman et al. (AoAS 2010)

# Summary of Key Model Features

- Different sources of asymmetries
  1. unit-specific characteristics in $\boldsymbol{z_i}$ which induce cross-sectional heterogeneity
  2. sign of a structural shock
  3. size of a structural shock

     trigger non-linear effects through $\mu(\boldsymbol{w_t})$
  4. initial conditions of the model

- Remarks:
  - ➢ Adding the quantiles to the VAR equation implies that asymmetries w.r.t. shocks in the micro panel could spill back to the macro VAR
  - ➢ Similar to a factor-augmented VAR where the quantiles of $\boldsymbol{x_{t-1}}$ ensure that the distributional shifts can have an effect on $\boldsymbol{y_t}$
  - ➢ Relatively simple and fast Markov Chain Monte Carlo sampler to produce draws from the joint posterior
  - ➢ Can produce two types of non-linear IRFs: unit-level and distributional IRFs

# Macro Block and Structural Identification

- Any standard identification scheme can be used to recover the structural VAR in the macro block

Our application for the Euro area:

- $\boldsymbol{y_t}$ includes $M = 5$ variables
  - $s_t^{oil}$ : oil supply shock from Baumeister and Hamilton (AER 2019)
  - $p_t^{oil}$ : real price of Brent crude oil (deflated with Euro area HICP) in logs
  - $\pi_t^{ener}$ : energy component of HICP in logs
  - $cci_t$ : consumer confidence index in logs
  - $i_t$ : 3-month Euribor as a proxy for the policy rate
- Sample period: Jan 2004 – Dec 2024

  → downweight covid observations (Hamilton, 2025)
- Lag length: $p = 12$

# Structural Model

$$s_t^{oil} = \qquad \boldsymbol{b}_1'\boldsymbol{y}_{t-1} + \sum_{s\in S} \beta_{1s} Q_s(\boldsymbol{x}_{t-1}) + u_{1t}$$

$$p_t^{oil} = \alpha_{ps} s_t^{oil} + \boldsymbol{b}_2'\boldsymbol{y}_{t-1} + \sum_{s\in S} \beta_{2s} Q_s(\boldsymbol{x}_{t-1}) + u_{2t}$$

$$\pi_t^{ener} = \alpha_{\pi s} s_t^{oil} + \alpha_{\pi p} p_t^{oil} + \boldsymbol{b}_3'\boldsymbol{y}_{t-1} + \sum_{s\in S} \beta_{3s} Q_s(\boldsymbol{x}_{t-1}) + u_{3t}$$

$$cci_t = \alpha_{cs} s_t^{oil} + \alpha_{cp} p_t^{oil} + \alpha_{c\pi} \pi_t^{ener} + \boldsymbol{b}_4'\boldsymbol{y}_{t-1} + \sum_{s\in S} \beta_{4s} Q_s(\boldsymbol{x}_{t-1}) + u_{4t}$$

$$i_t = \alpha_{is} s_t^{oil} + \alpha_{ip} p_t^{oil} + \alpha_{i\pi} \pi_t^{ener} + \alpha_{ic} cci_t + \boldsymbol{b}_5{}'\boldsymbol{y}_{t-1} + \sum_{s\in S} \beta_{5s} Q_s(\boldsymbol{x}_{t-1}) + u_{5t}$$

- Identification: combine recursive structure with sign priors on vector of structural parameters $\boldsymbol{\alpha}$

(Baumeister and Hamilton, Ecma 2015; Plagborg-Møller and Wolf, Ecma 2021; Chan, QE 2022)

# Matrix of Contemporaneous Relations

- Signs on elements in $\boldsymbol{A}$ (negative of $\boldsymbol{\alpha}$)

$$\boldsymbol{A} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ + & 1 & 0 & 0 & 0 \\ + & - & 1 & 0 & 0 \\ - & + & + & 1 & 0 \\ + & - & - & - & 1 \end{pmatrix}$$

- Implemented with truncated Gaussians (see Botev, JRSSB 2017)

# Survey Data on Quantitative Inflation Expectations

- The microdata is from the *Harmonized Business and Consumer Survey* provided by the European Commission.
- The survey is based on a repeated cross section.
- Survey question
  - ➢ Participants are first asked about their qualitative inflation expectations.
  - ➢ If they expect a change, they are asked:
    *(Q61) "By how many per cent do you expect consumer prices to go up/down in the next 12 months? (Please provide a single figure estimate.)*
    *Consumer prices will increase by ……,…%/ decrease by ……,…%."*
- The dataset is monthly and starts in January 2004.
- On average, around 21,000 participants are surveyed each month in all Euro area member countries using a harmonized questionnaire and common timetable.

# A Pseudo Panel for Quantitative Inflation Expectations

- The survey also collects demographic data, including gender, age, income, occupation, and education.
- We construct a **pseudo panel** along the following dimensions (Deaton JoE 1985, Verbeek 2008):
  - income quartiles (4 groups)
  - gender (2 groups)
  - age (4 groups)
  - across 10 countries

  yields 305 pseudo individuals (after eliminating 'odd' combinations)
- For each pseudo individual, we compute the average of their subjective inflation expectations.
- Due to data gaps in some countries, the analysis focuses on **ten Euro area member countries**: Germany, France, Italy, Spain, Belgium, Austria, Finland, Estonia, Luxembourg, and Slovenia.

# Illustration: Median inflation expectations by socioeconomic groups

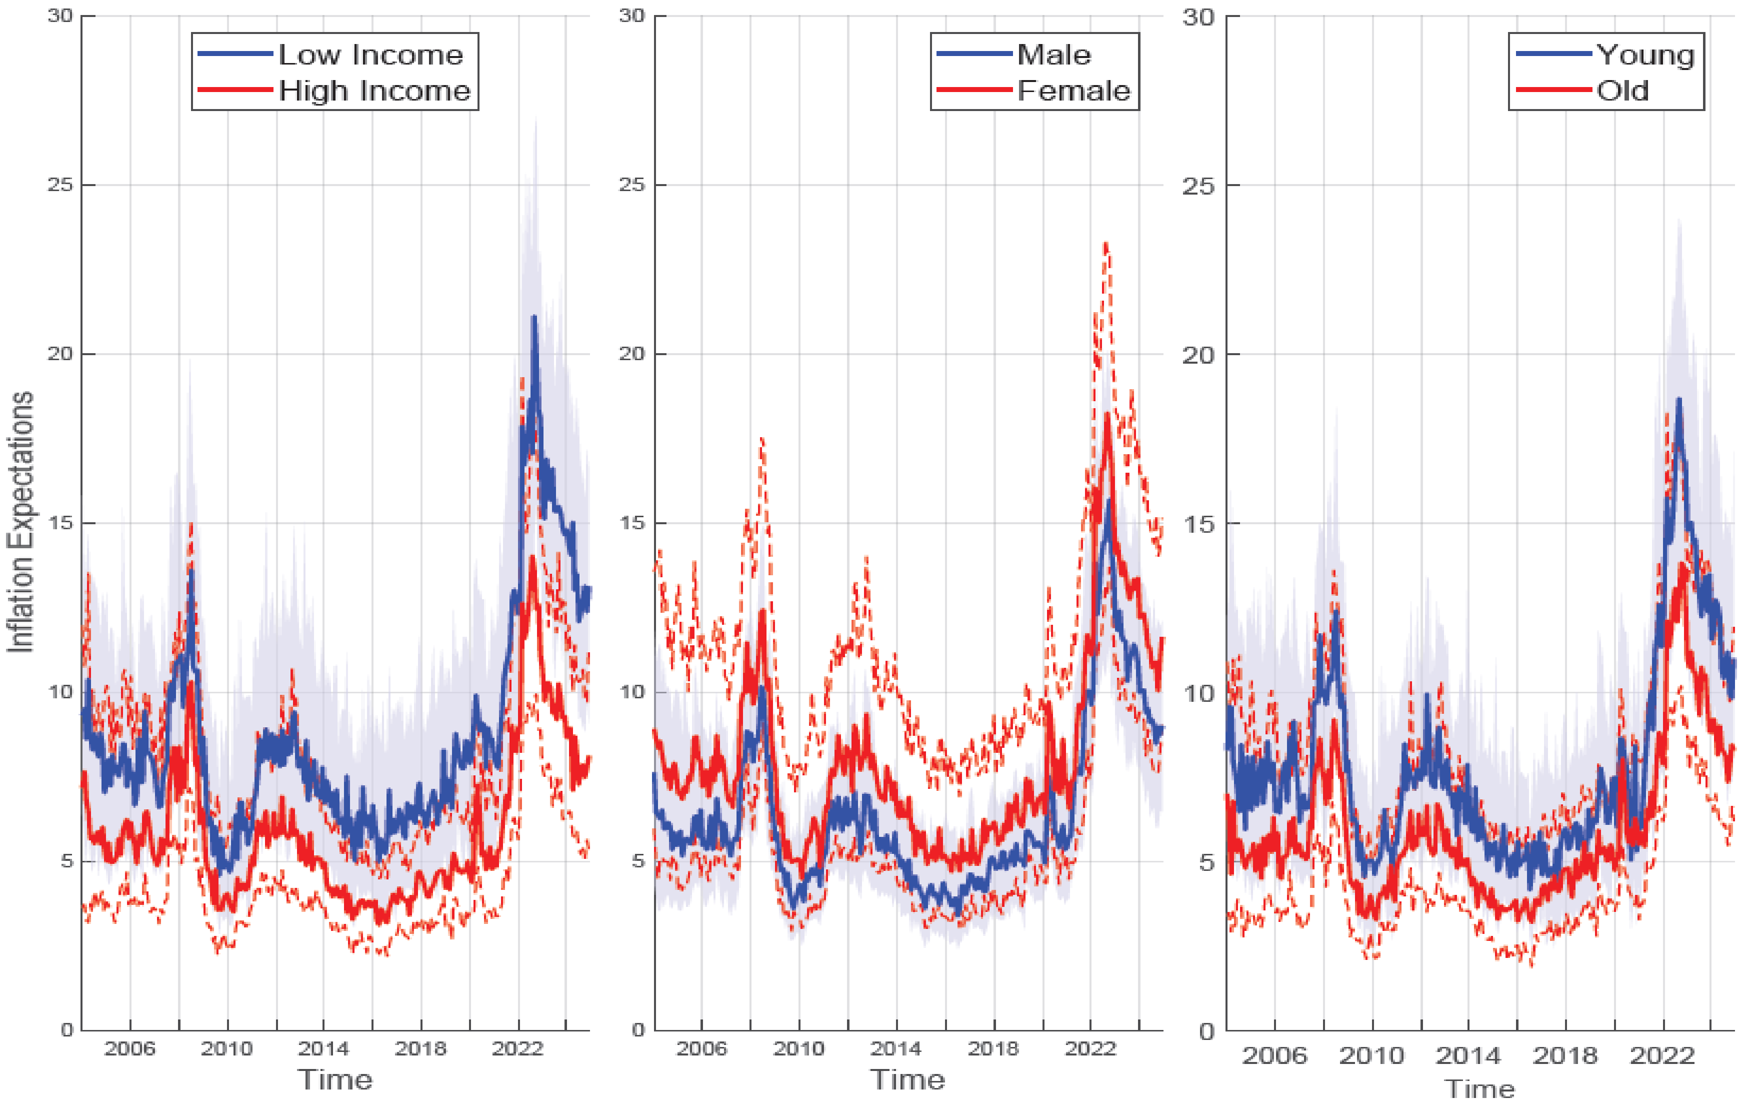


Rich heterogeneity, especially when comparing high- and low-income groups

# Evolution of the Distribution of Inflation Expectations

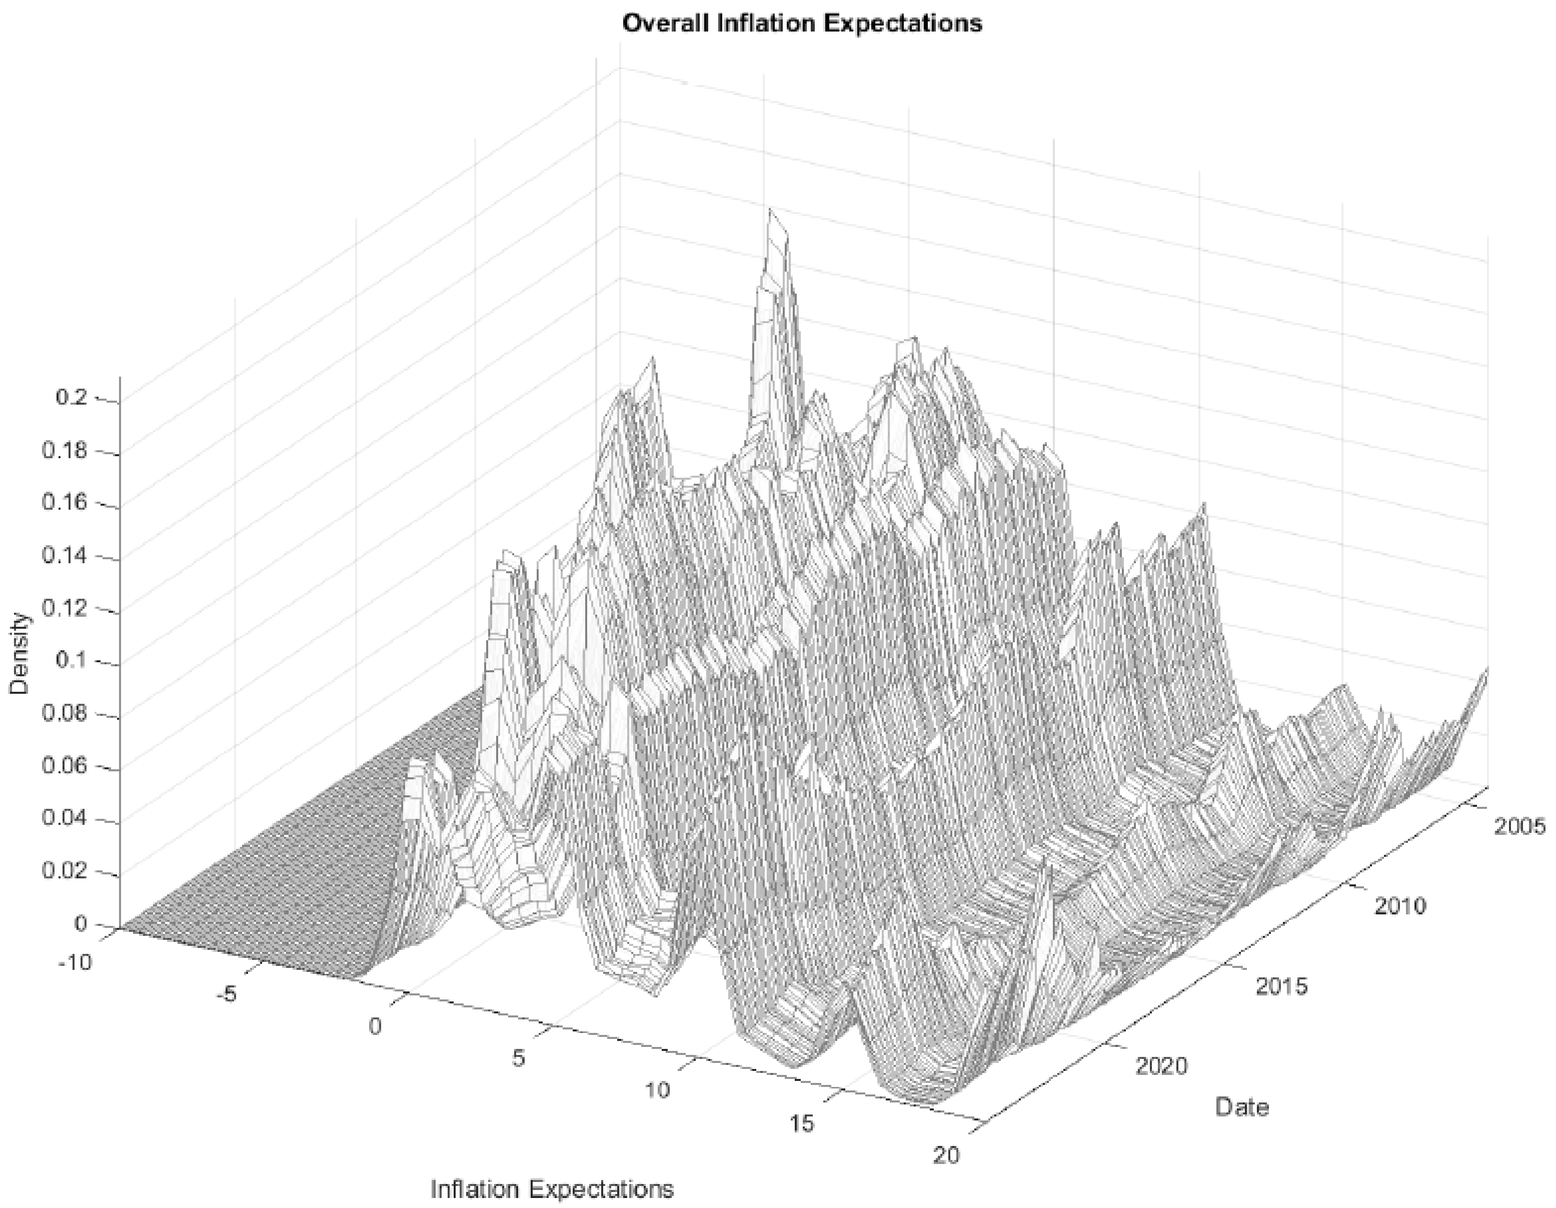

# Dynamic Effects of a 2 SD Oil Supply Shock on Macro Variables

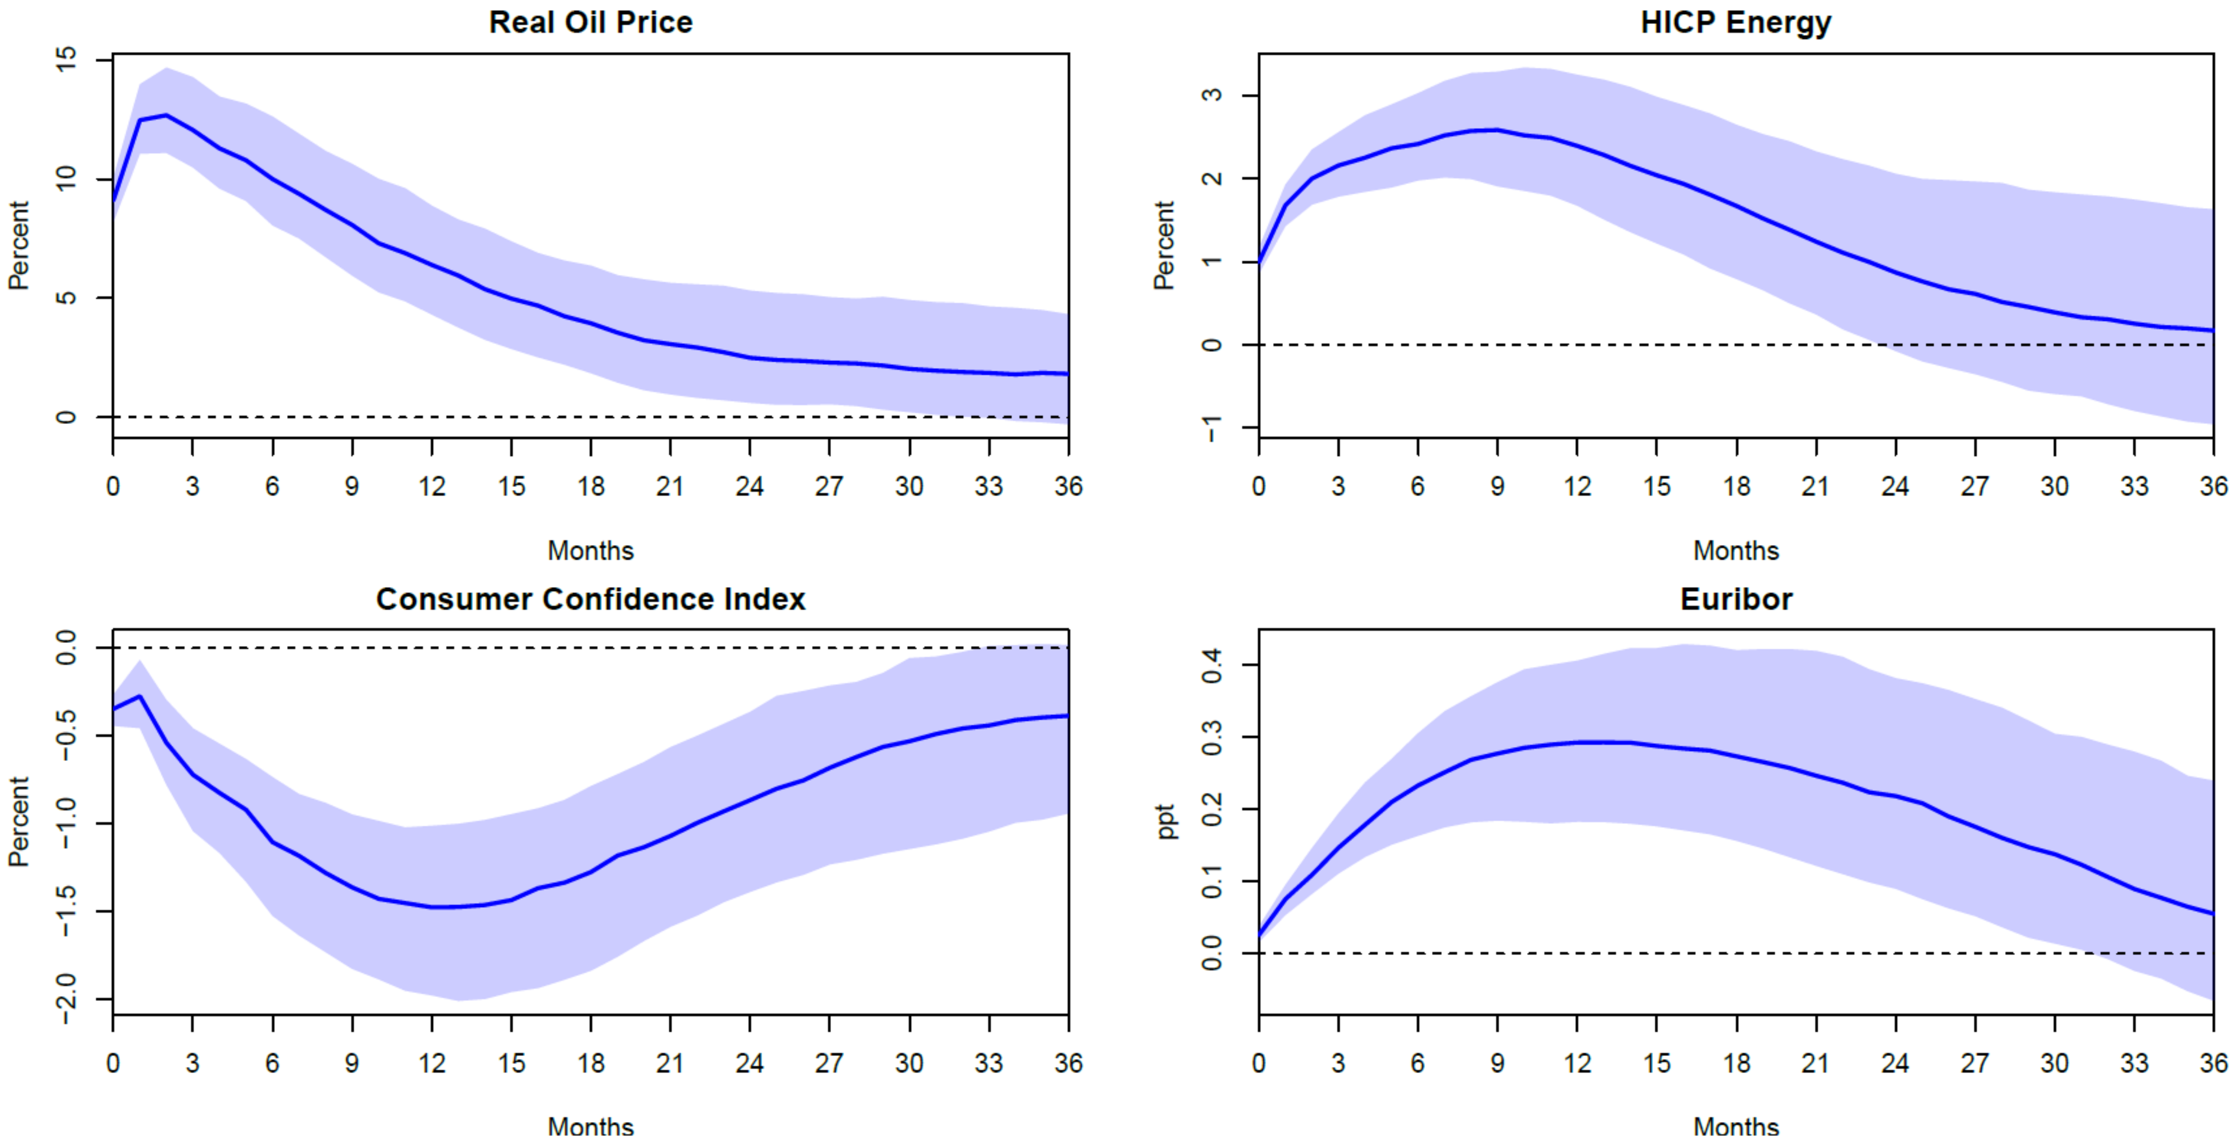

# Responses of Aggregated Inflation Expectations to Oil Supply Shocks of Different Size and Sign

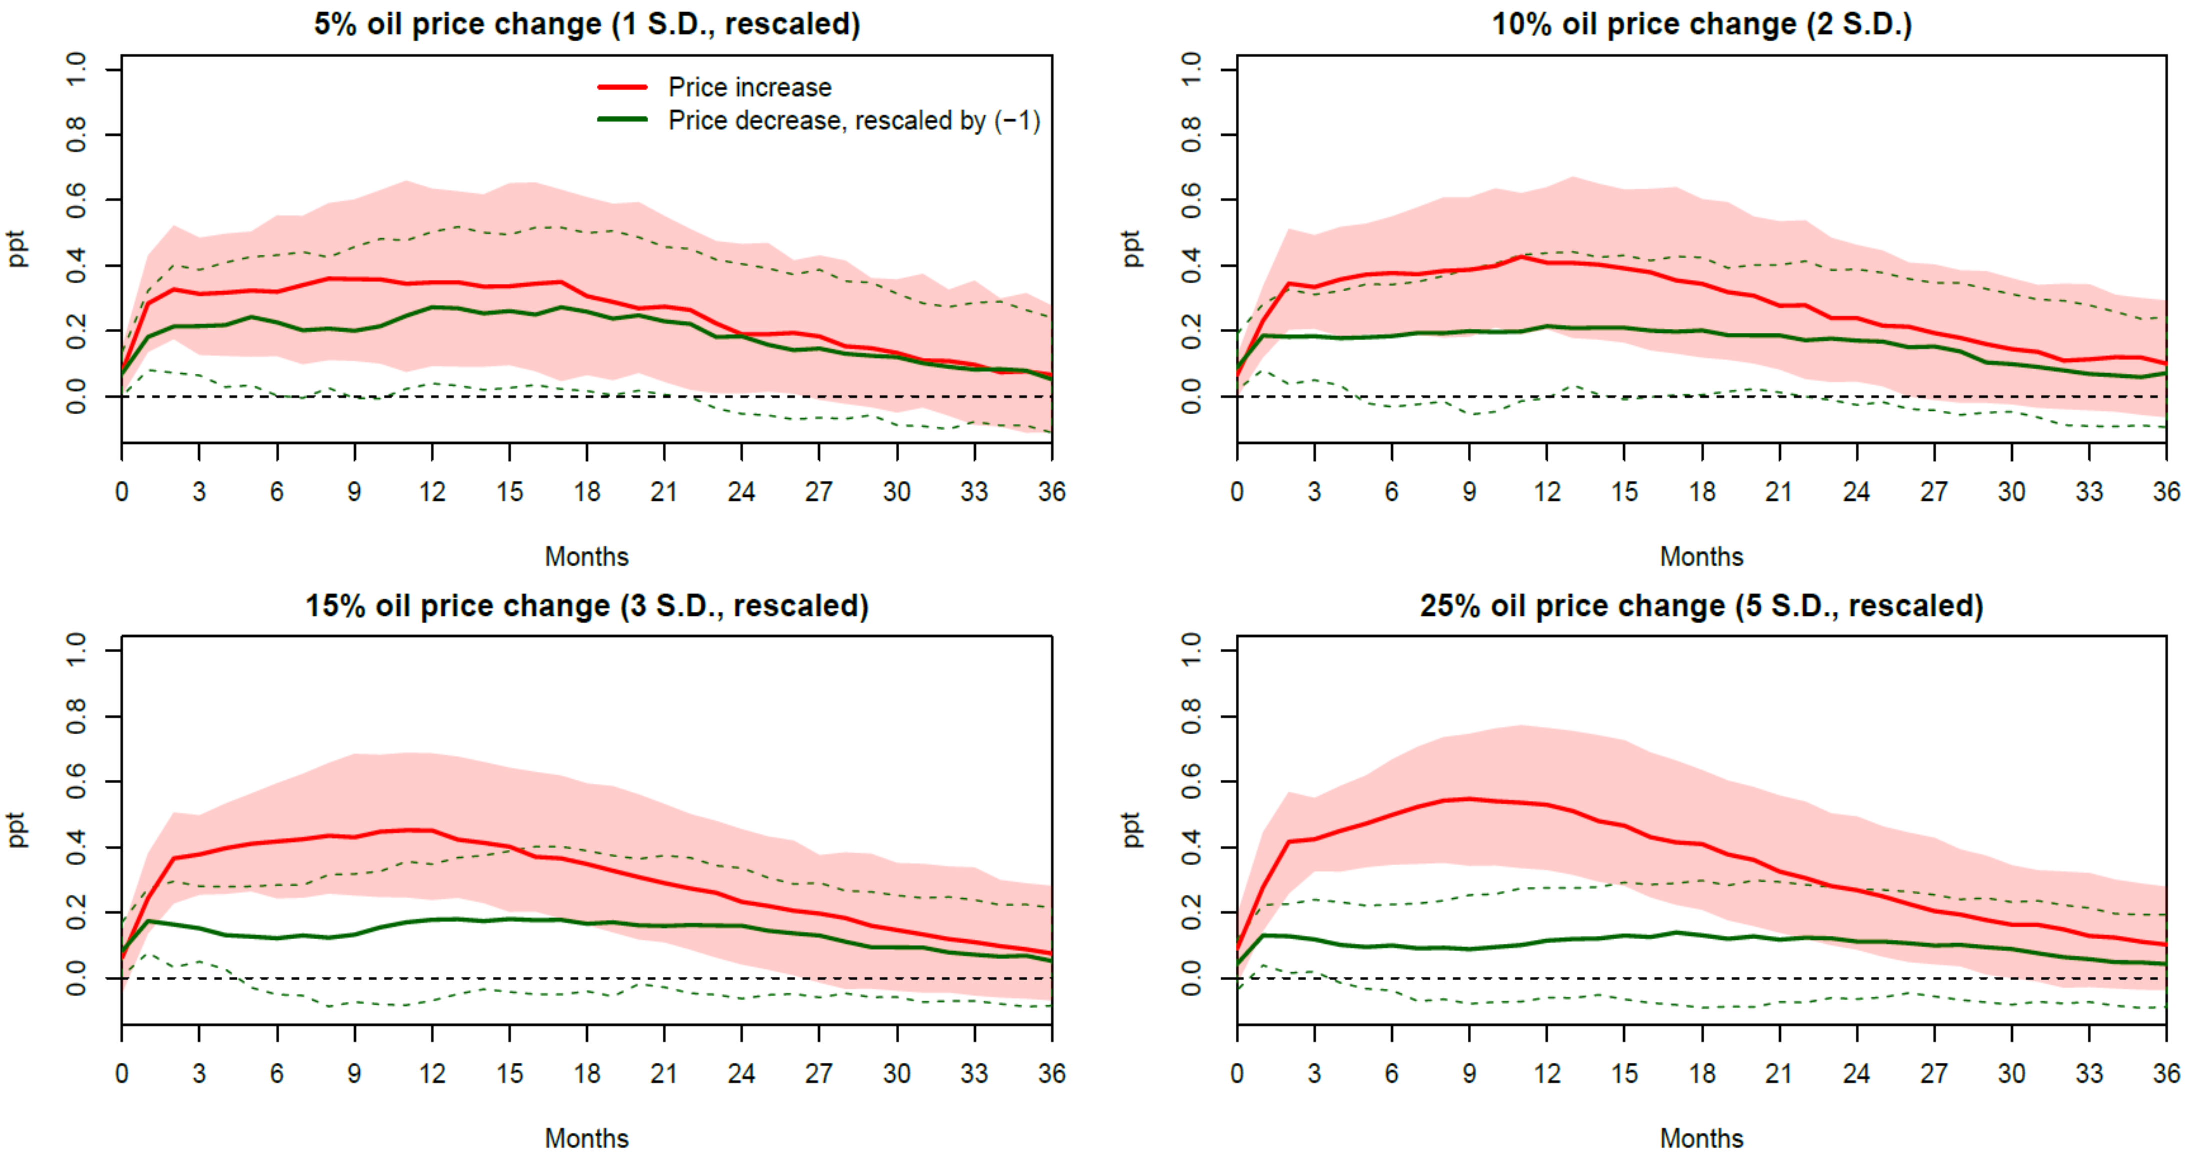

# Dynamic Response of Aggregated Inflation Expectations

- Asymmetry:

  Oil price increases trigger stronger and more persistent increase in inflation expectations compared to price decreases of similar magnitude

  → households react more strongly to cost-push pressures than deflationary signals

- Nonlinearity:

  The gap widens with larger shocks

  → asymmetry becomes more pronounced for large oil supply shocks

# Digging Deeper: Unit-Level Asymmetries and Heterogeneities

Key questions:

- Is there heterogeneity in the reaction of inflation expectations at the disaggregated level?
- Do these responses differ with respect to the size and the sign of the oil supply shocks?

→ look at how sub-groups with different socioeconomic characteristics and across countries react

# Heterogeneity and Asymmetries: Unit-Level Peak Responses by Country and Income

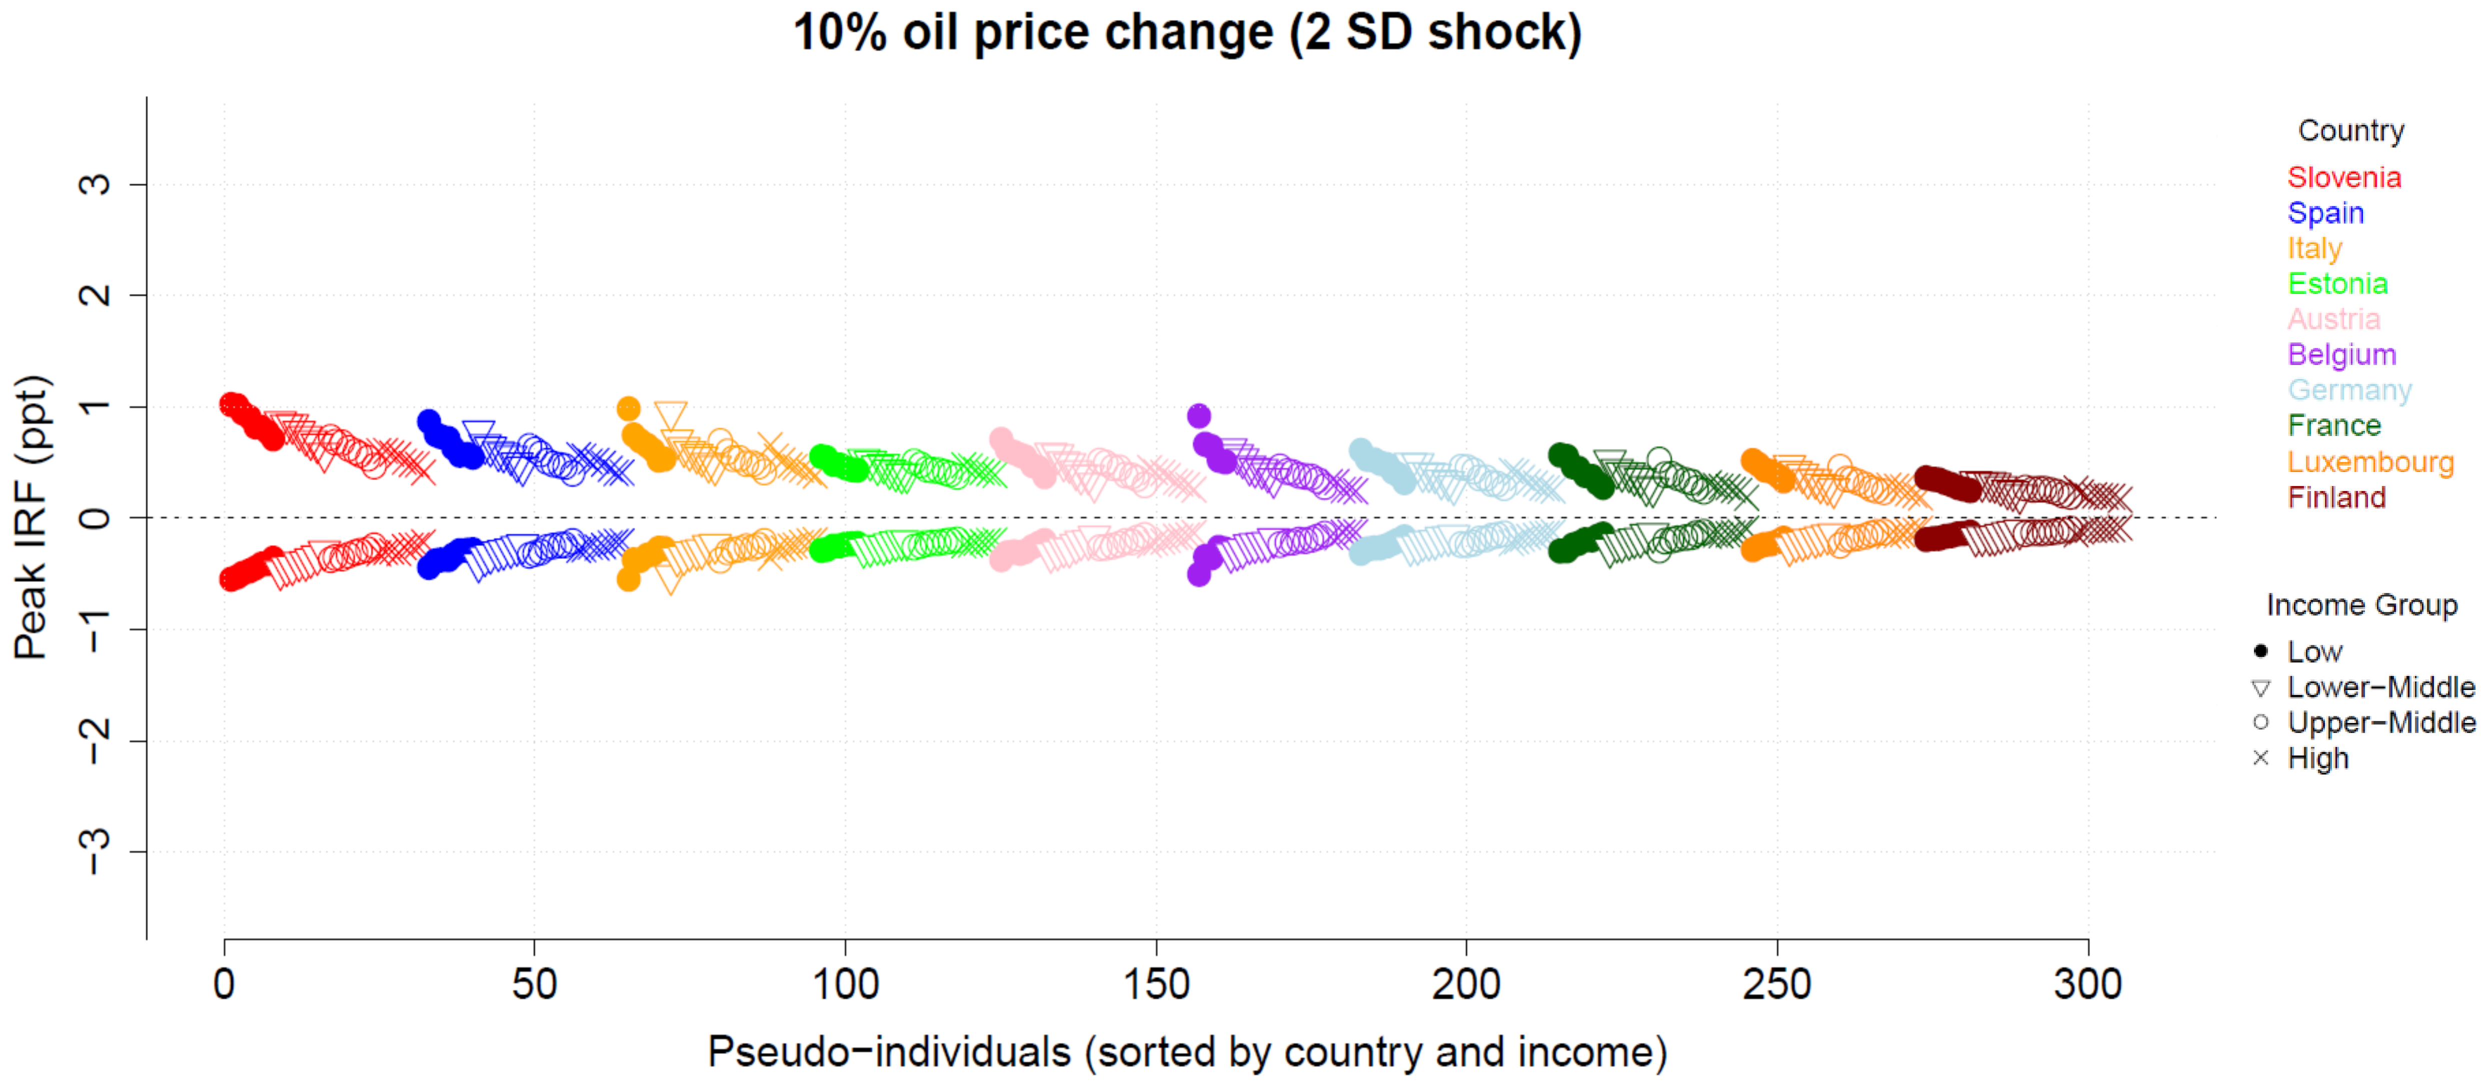

# Heterogeneity and Asymmetries: Unit-Level Peak Responses by Country and Income

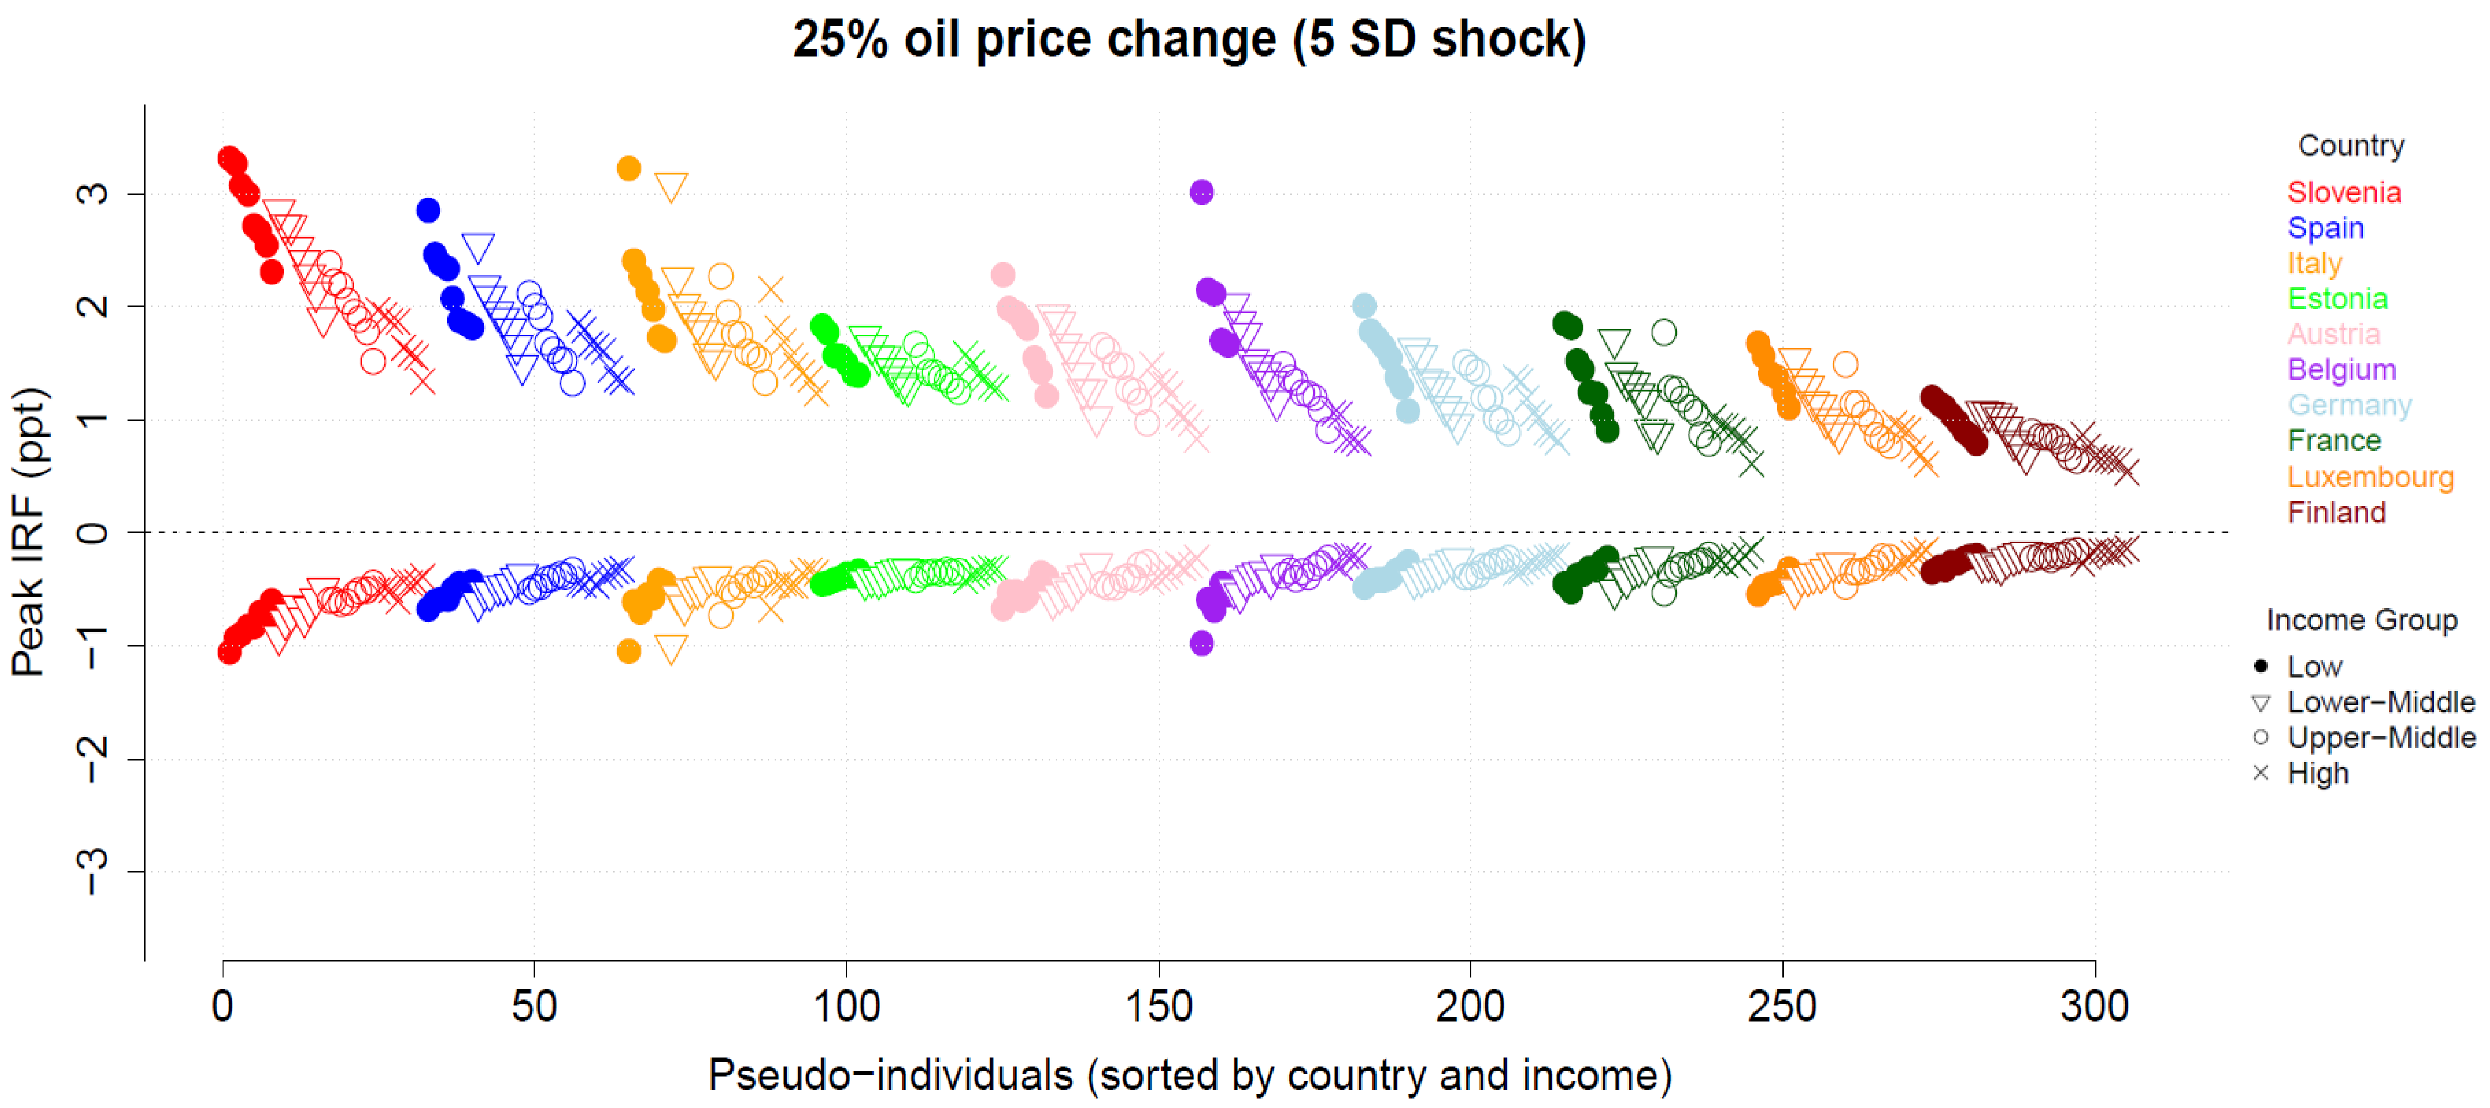

# Asymmetries and Heterogeneity: What Explains this Pattern?

## (1) Asymmetries

- While for smaller shocks the responses to negative and positive shocks are basically mirror images, there are strong asymmetries for large shocks across countries and income groups.

- Large price-increasing shocks are costly for households → marginal utility of updating information is higher in the case of a price increase

⟹ asymmetry can be traced back to rational (in)attention

# Asymmetries and Heterogeneity: What Explains this Pattern?

## (2) Heterogeneity

- There are pronounced differences both across countries and by income group within countries.

- "*Static*" explanation:

  Energy consumption shares measure exposure to shocks.

  - Cross-country differences:

    The higher the relative energy consumption shares of households, the *larger* the adjustment of inflation expectations should be.

# Cross-Country Differences: Energy Weight in Consumption Basket

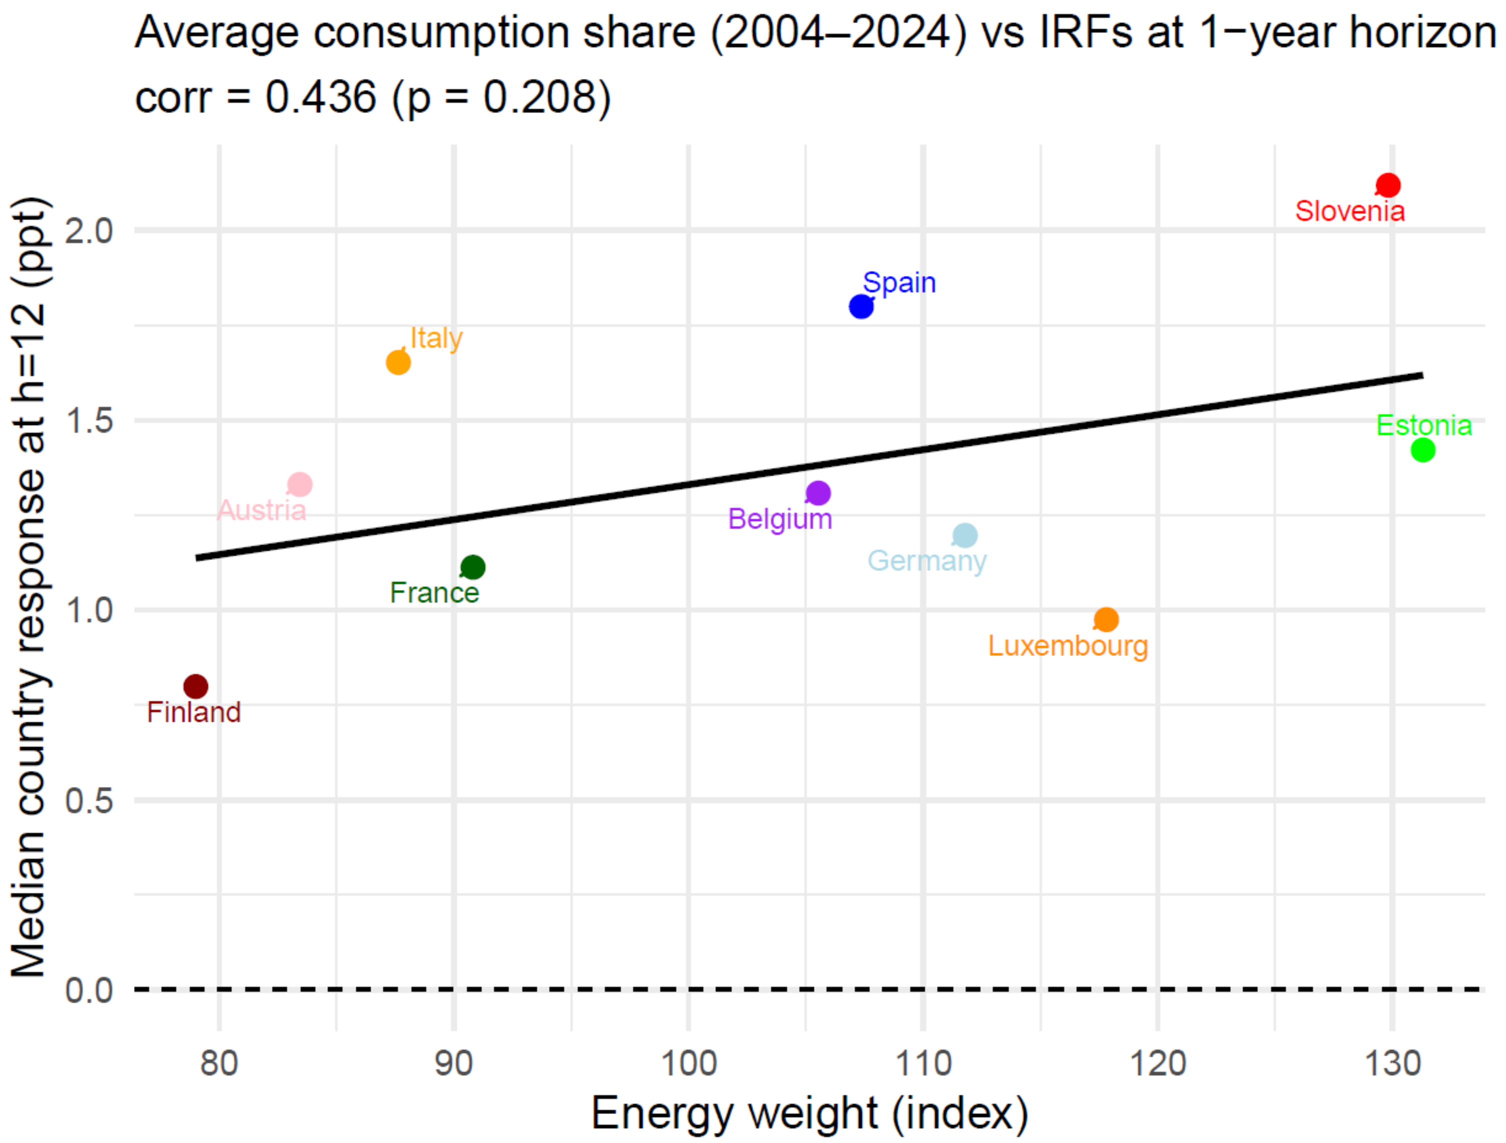

# Asymmetries and Heterogeneity: What Explains this Pattern?

## (2) Heterogeneity

- There are pronounced differences both across countries and by income group within countries.
- "*Static*" explanation:
  Energy consumption shares measure exposure to shocks.
  - Cross-country differences:
    The higher the relative energy consumption shares of households, the *larger* the adjustment of inflation expectations should be.
  - Within-country differences by income:
    Low-income households spend a larger share of their budget on energy, which should make them *more* sensitive to oil price shocks.

# Within-Country Differences: Energy Consumption Shares by Income Group

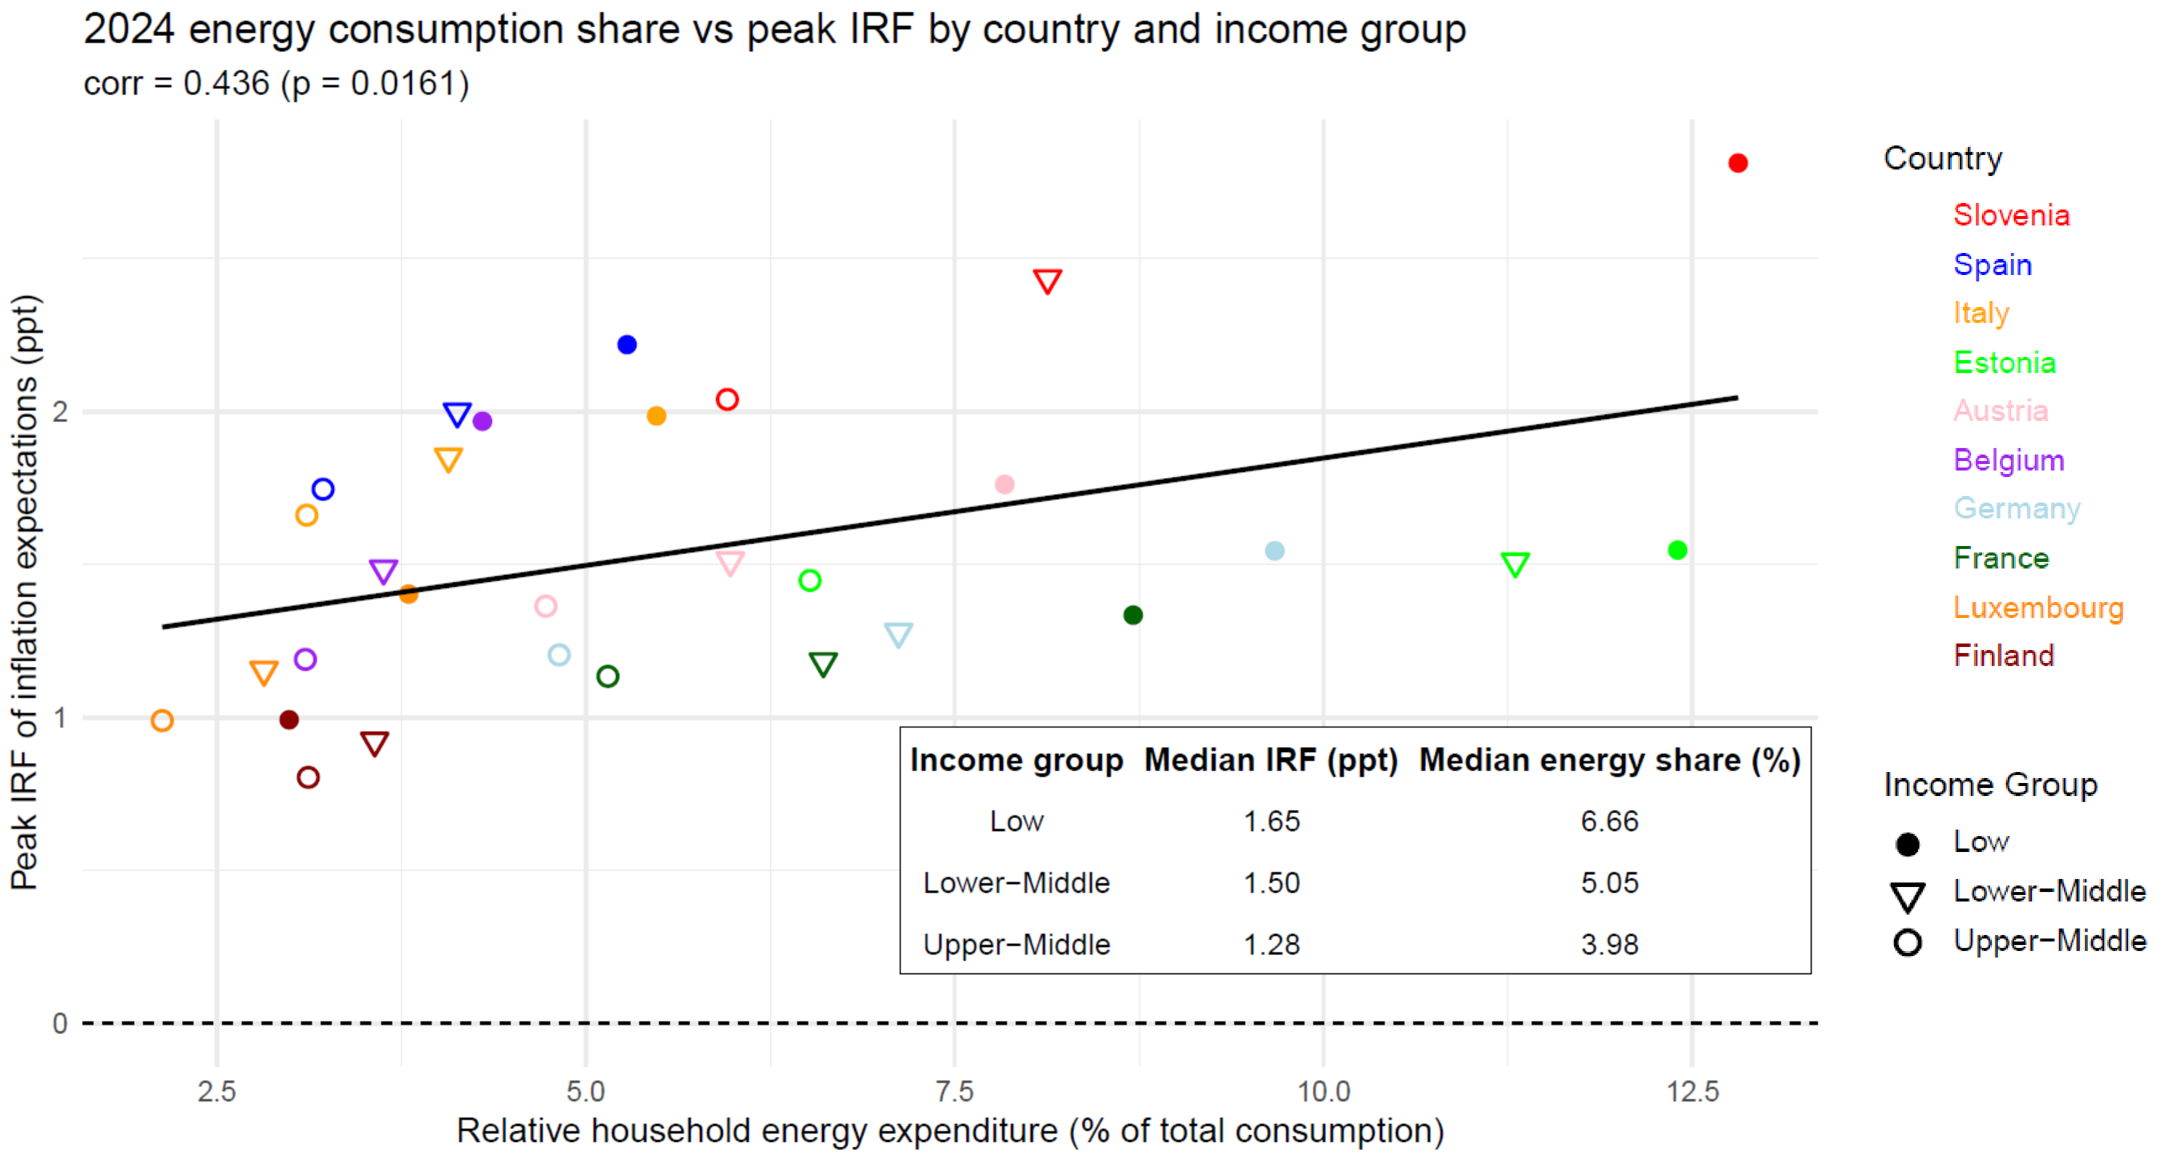

# Asymmetries and Heterogeneity: What Explains this Pattern?

## (2) Heterogeneity

- There are pronounced differences both across countries and by income group within countries.

- "*Dynamic*" explanation:

  An oil supply shock triggers a contractionary monetary policy response.

  - Homeowners with mortgages pay closer attention to monetary policy and thus anticipate that tighter policy will dampen inflation (Ahn et al., JME 2024)
    → reaction to oil supply shocks should be *less* pronounced
  - Hand-to-mouth households are more adversely affected by monetary policy, generating stagflationary correlation beliefs (Kamdar and Ray, 2025)
    → reaction to oil supply shocks should be *stronger*

# Monetary Policy Channels and Cross-Country Household Inflation Expectation Responses

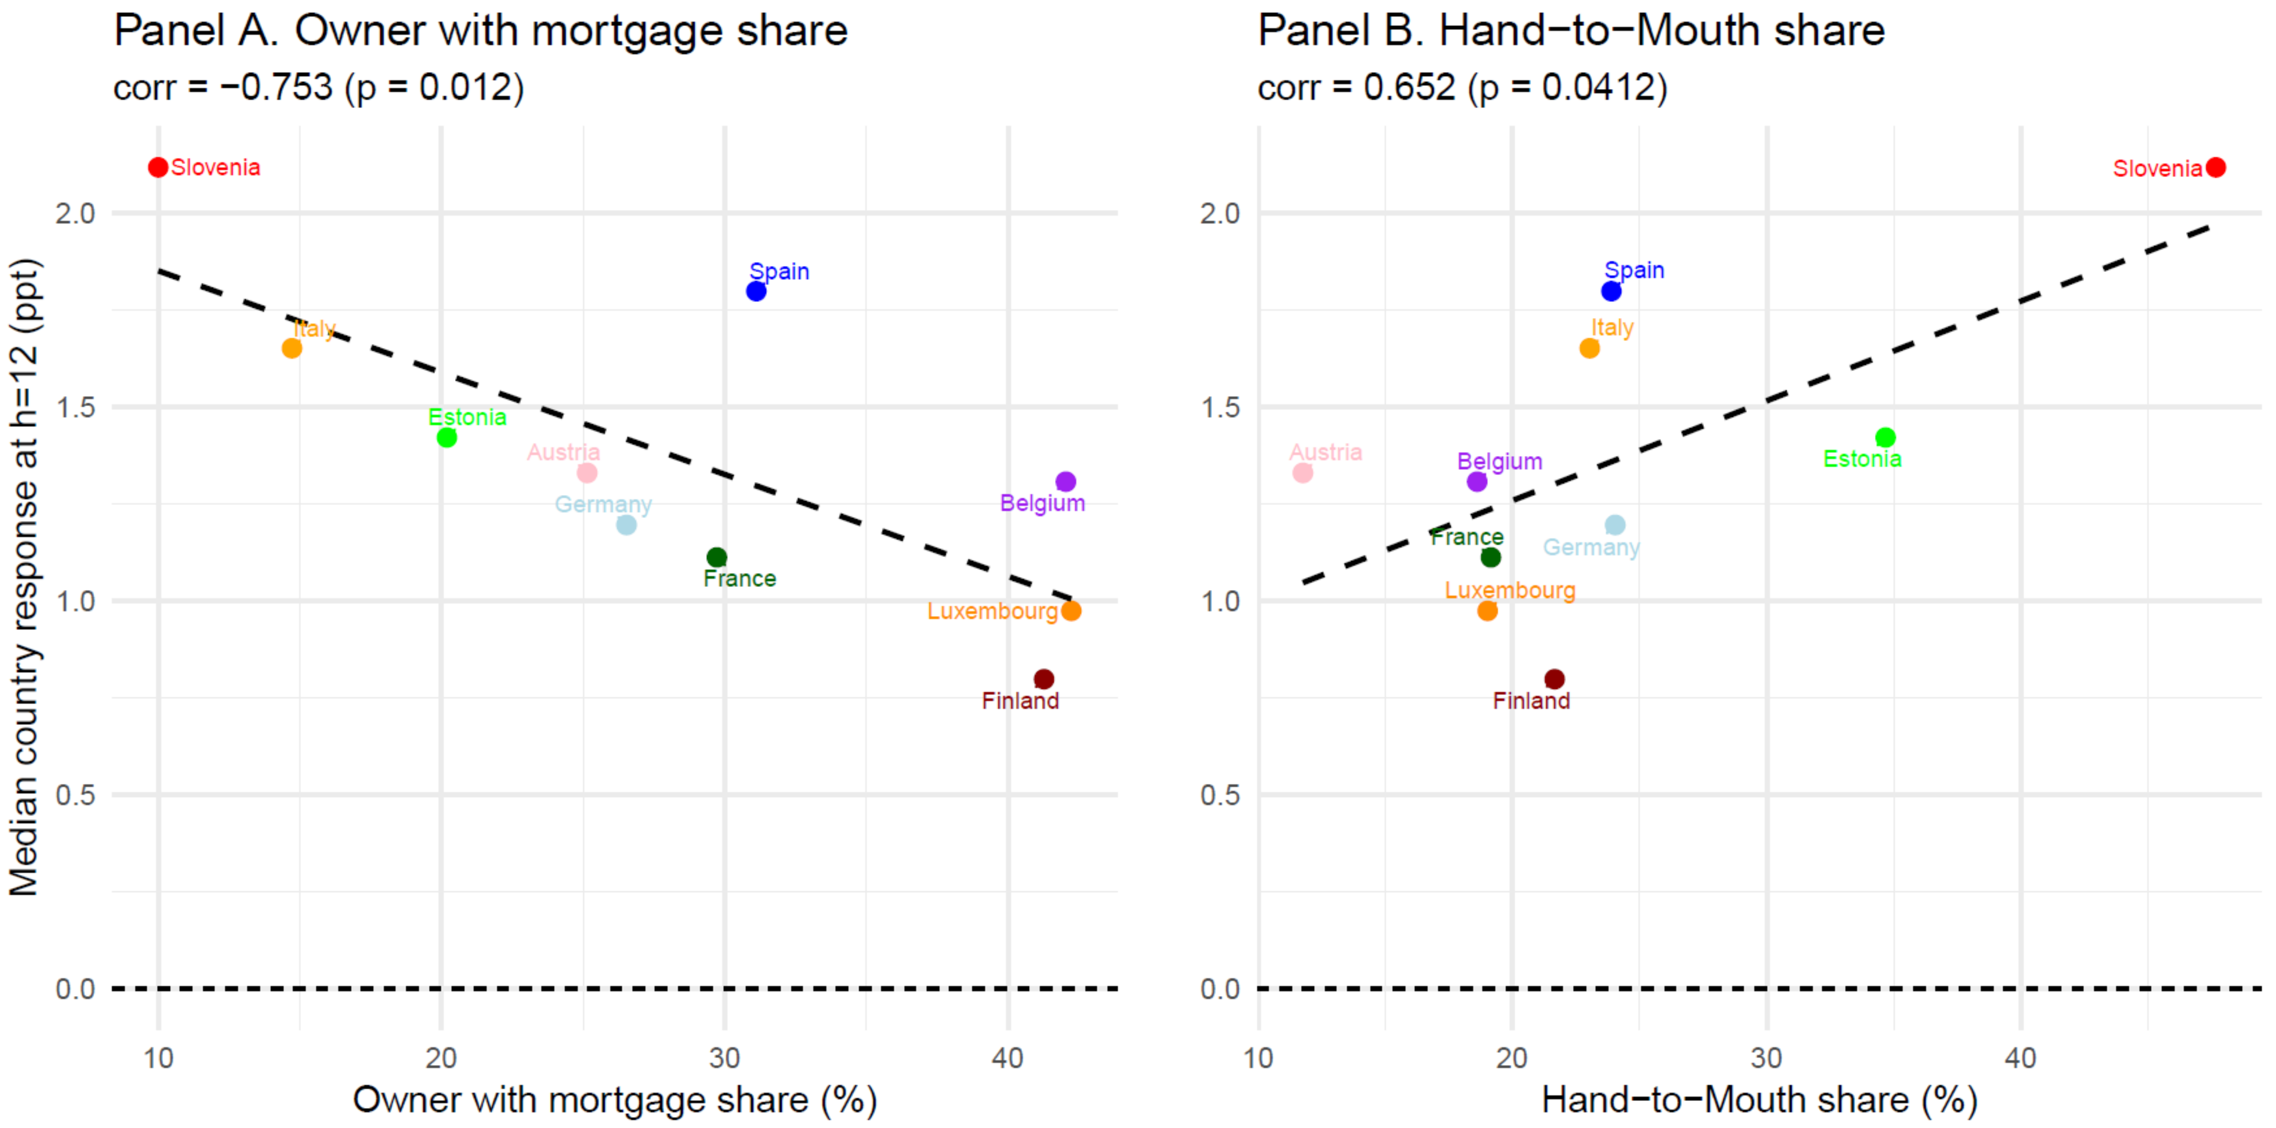

# Constructing Density Responses

Steps:

- Reconstruct the *actual* distribution by drawing $N_{jt}$ observations from a Gaussian distribution with group-level means and empirical standard deviation accounting for relative weights of pseudo individuals

  $\rightarrow$ mixture of Normals
- Use the posterior draws to obtain predictive distributions via kernel smoothing
- Do this twice – once with and once without the shock

# Density IRFs after 1 Year for Small and Large Shocks

- Median forecast of the cross-sectional distributions with and without shock

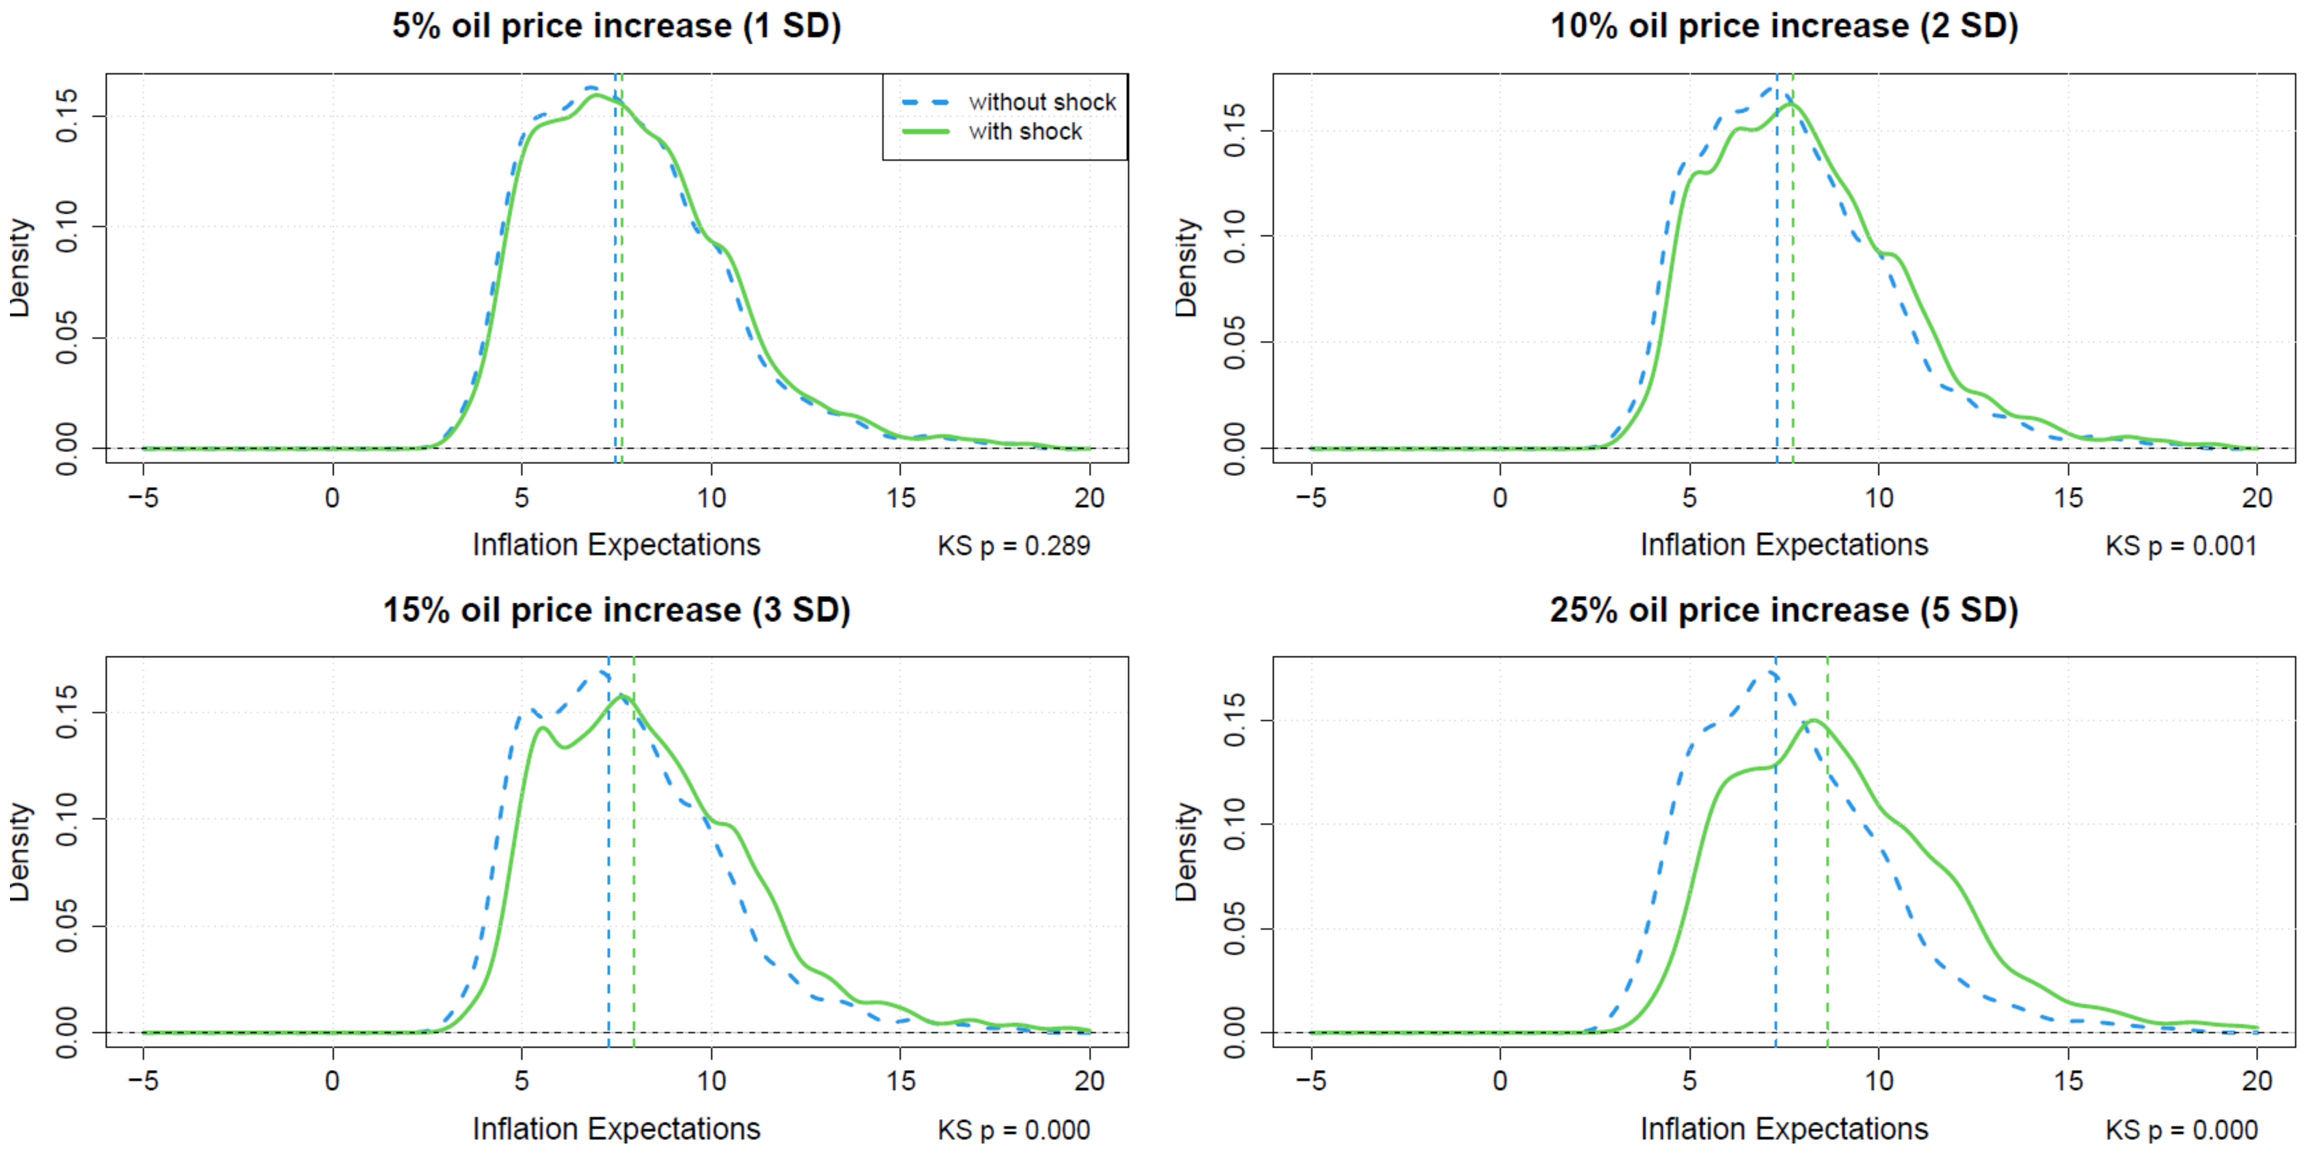

# Differences in Distributional Responses

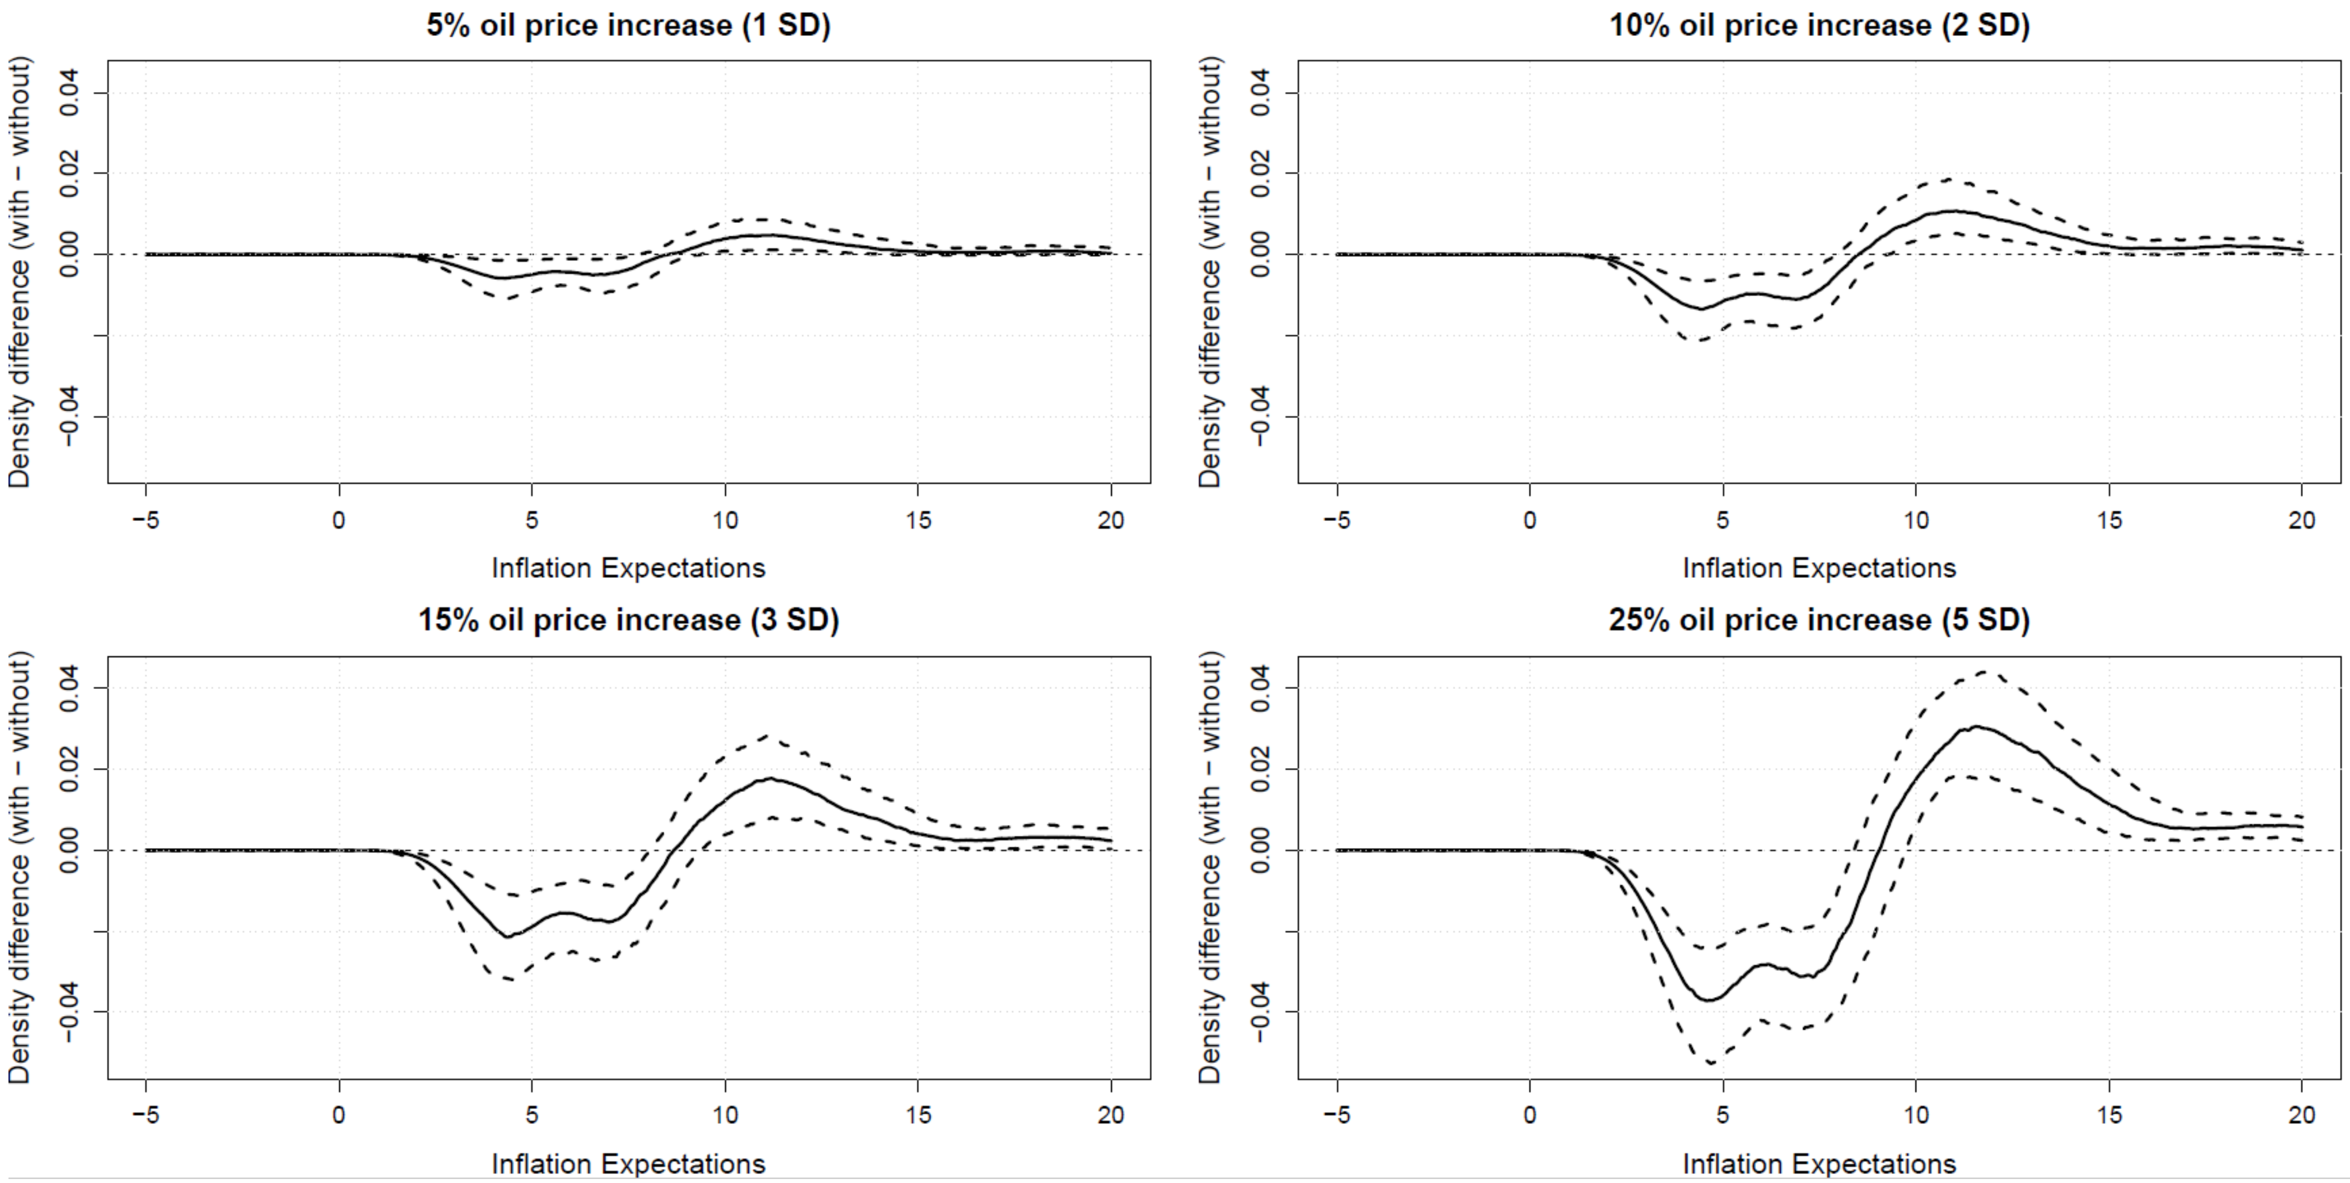

# Distributional Shifts in Inflation Expectations

- Estimate the share of the population whose inflation expectations are shifted above or below a certain threshold
- Some notation:
  - Let $p(\boldsymbol{x}_{t+h}|u_t^{oil\ supply} = c, \mathrm{I}_t)$ denote the h-step-ahead predictive distribution given that the oil supply shock in time $t$ was $c$

    $\rightarrow$ posterior mean of this conditional predictive density is $\overline{\boldsymbol{x}}_{t+h}^{c}$
  - Likewise, $p(\boldsymbol{x}_{t+h}|\,\mathrm{I}_t)$ is the unconditional predictive distribution with posterior mean $\overline{\boldsymbol{x}}_{t+h}^{uc}$
- Compute downside/upside pseudo probabilities:

$$F_{\gtrless,t+h}^{s}(\pi^e) = \frac{1}{N}\sum_{n=1}^{N} \mathbb{I}(\bar{x}_{n,t+h}^{(s)} \lesseqgtr \pi^e) \quad \text{for } s = \{c, uc\}$$

$\rightarrow$ difference $\Delta(\pi^e)_{\lesseqgtr,t+h} = F_{\gtrless,t+h}^{c}(\pi^e) - F_{\gtrless,t+h}^{uc}(\pi^e)$ tells us how inflation expectations in terms of the overall shares in the population shift

# Distributional Shifts in Inflation Expectations

| $h$ | $\Delta(2)_{<,t+h}$ | $\Delta(5)_{>,t+h}$ | $\Delta(8)_{>,t+h}$ |
|---|---|---|---|
| Oil price increase (5 SD) | | | |
| 0 | 0.00 | 2.68 | 2.93 |
| 3 | 0.00 | 9.47 | 15.76 |
| 6 | 0.00 | 10.56 | 19.45 |
| 12 | 0.00 | 9.97 | 21.71 |
| Oil price decrease (5 SD) | | | |
| 0 | 0.00 | -0.92 | -1.93 |
| 3 | 0.00 | -3.60 | -4.44 |
| 6 | 0.00 | -2.93 | -3.52 |
| 12 | 0.00 | -2.77 | -3.77 |

**Table:** Differences in pseudo probabilities, $\Delta = F^{c}_{t+h} - F^{uc}_{t+h}$.

# Density IRFs after 1 Year for Small and Large Shocks: Comparing Young and Old

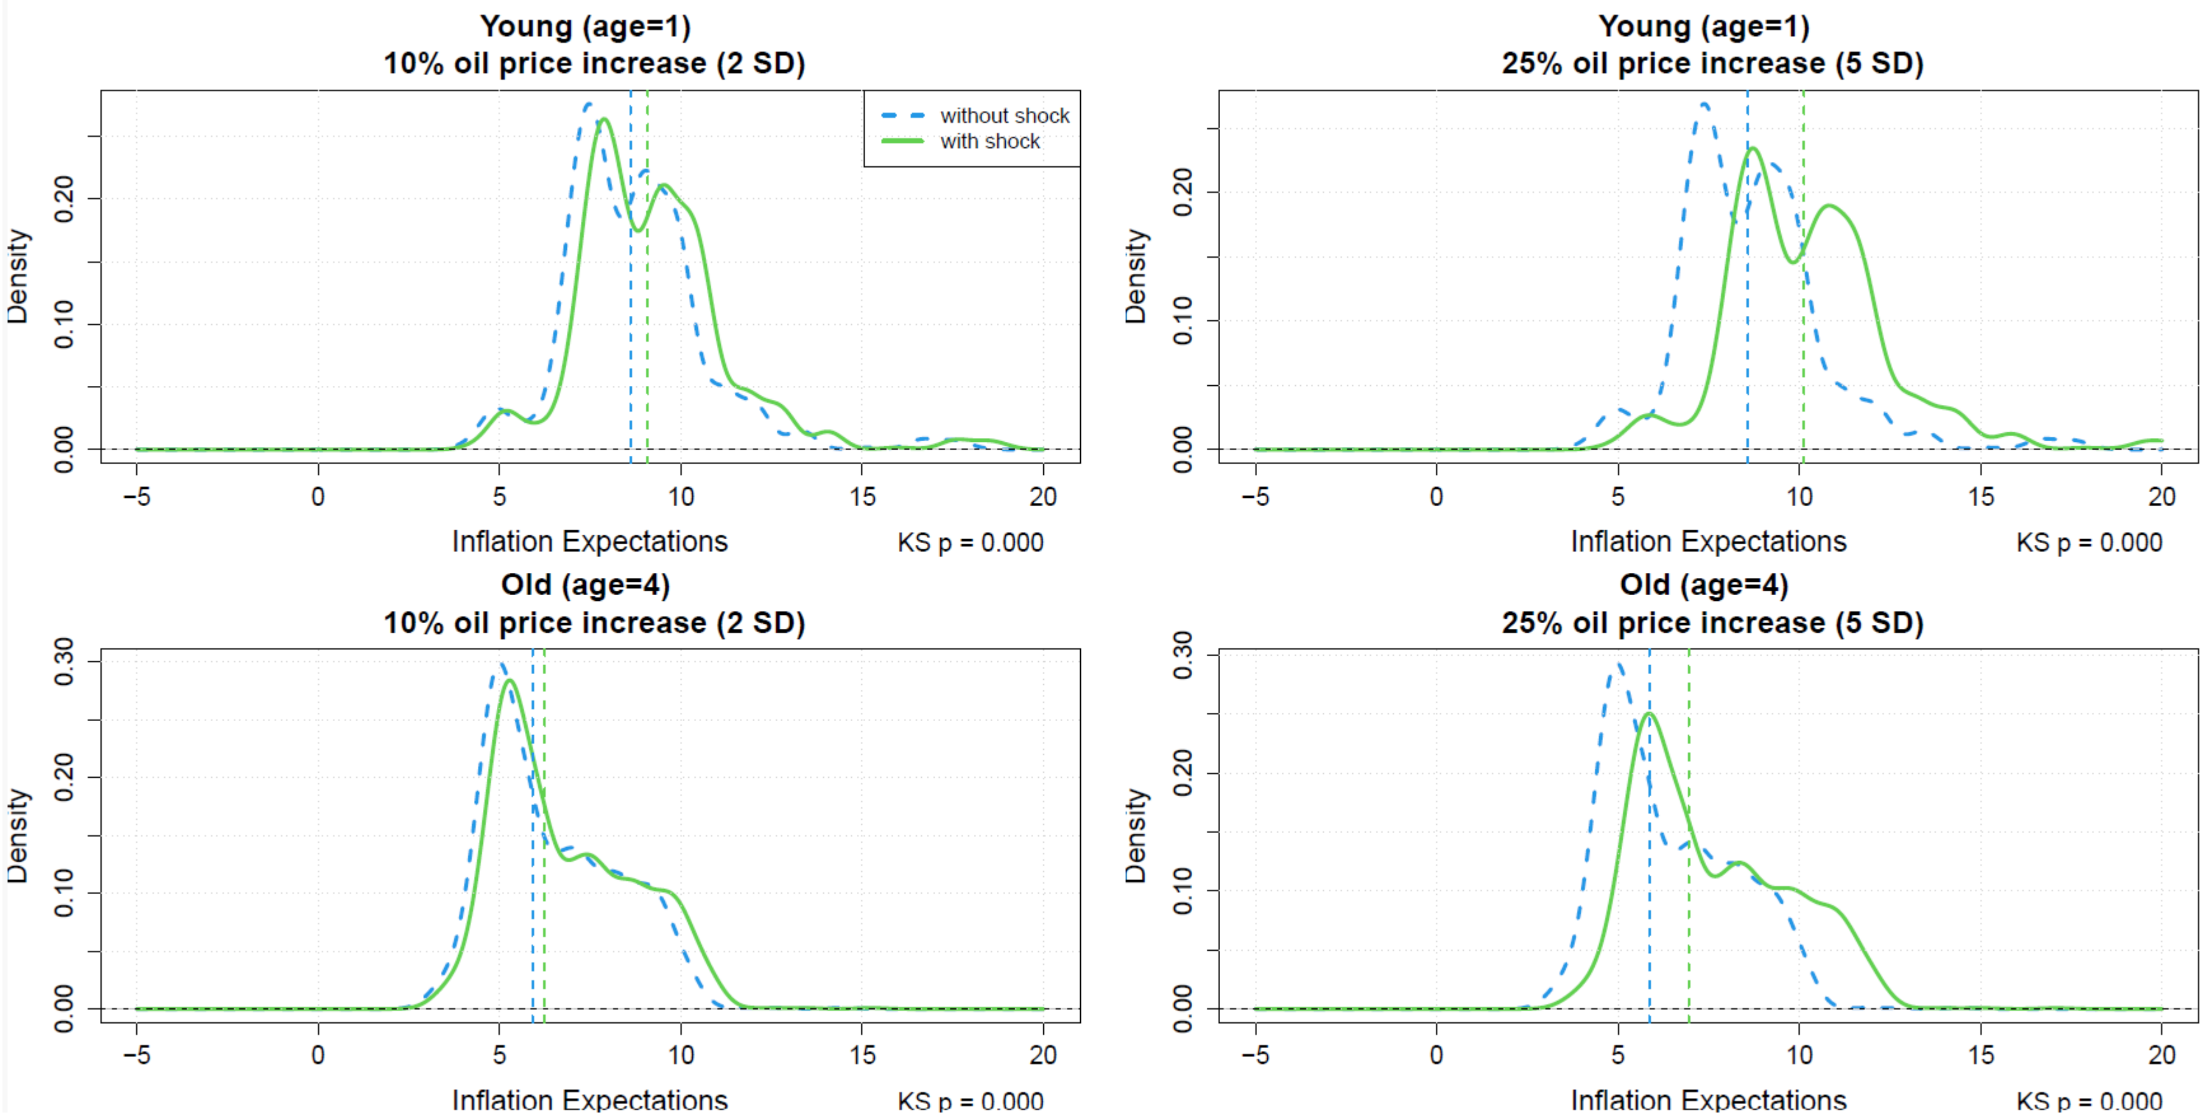

# Differences in Distributional Responses: Comparing Young and Old

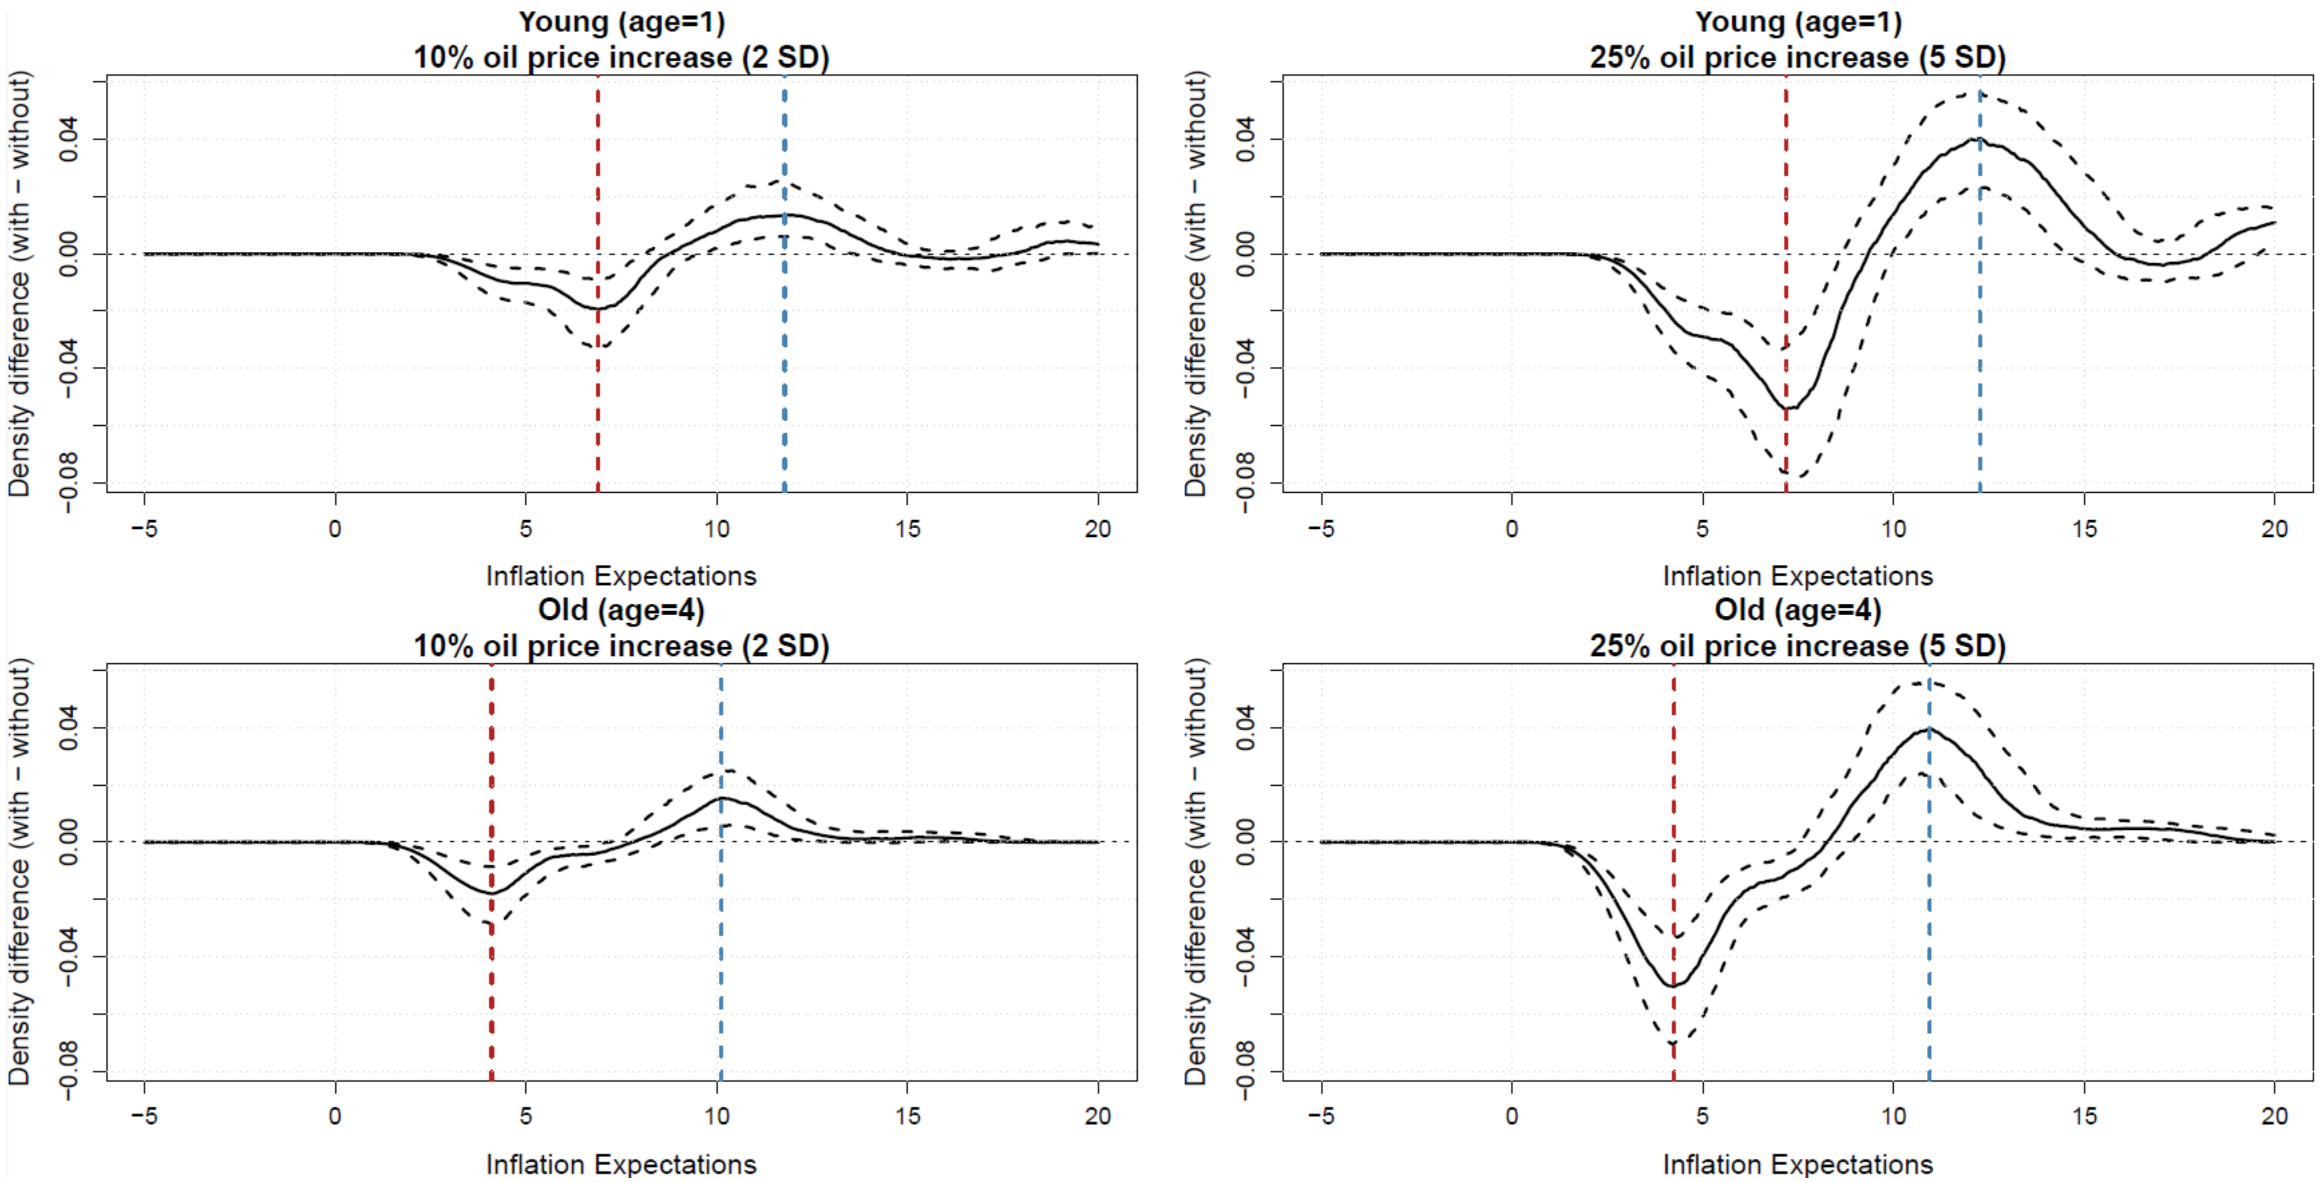

# Evidence on Age Differences

- Older individuals (65+):

  Probability mass shifts from inflation expectations below 5% (without a large shock) to above 10% once the shock occurs.

- Younger individuals (16-29):

  Probability mass shifts from already above 5% further upward to above 10% after the shock.

- Potential channel:

  The stronger adjustment among older individuals is consistent with selective memory recall of past experiences, triggering a more pronounced belief update (Gennaioli, Leva, Schoenle, and Shleifer, 2025).

# Conclusion

- Methodological takeaways:
  - ➢ We develop a scalable micro-macro modeling framework that links a standard macro SVAR to a nonlinear panel model
  - ➢ The model captures non-linear macro-to-micro propagation of shocks, while remaining agnostic about the form of non-linearities
    - We use machine learning to infer them (here: BART)
    - To achieve parsimony and economic interpretability, we use a factor model where the loadings that are determined by household characteristics
  - ➢ Model is tractable and computation is fast thanks to efficient MCMC sampler
- New empirical findings:
  - ➢ Oil supply shocks trigger asymmetric and heterogeneous reactions in household inflation expectations across Euro area countries
    - Unexpected oil price increases due to supply disruptions trigger stronger increases in inflation expectations and this sign asymmetry becomes more pronounced if shocks are large.
    - Lower income households, older individuals, and countries with a larger energy expenditure share display a more pronounced adjustment.

# APPENDIX

# A Note on the Weighting Scheme

- For aggregation across IRFs and to construct distributional responses, we use survey and country weights.
  - ➢ Survey weights: The Commission provides weights that take account of gender, age, occupation, size of household, region, size of town, and education level in the survey.
  - ➢ Country weights: When we consider aggregate inflation expectations responses or the distributional responses, we treat it as a "Euro-area wide" distribution; thus, to aggregate the 10 countries, we use the share of the respective country in the overall private consumption expenditures. Shares are pretty constant over time, so we use the average from 2004 to 2024 to aggregate the results.