SVERIGES RIKSBANK
WORKING PAPER SERIES

465

The Role of Firm Heterogeneity for the Transmission of Aggregate Shocks

Michele Lenza, Giuseppe Pagano Giorgianni, Lorenza Rossi and Ettore Savoia

April 2026

WORKING PAPERS ARE OBTAINABLE FROM

www.riksbank.se/en/research
Sveriges Riksbank • SE-103 37 Stockholm
Fax international: +46 8 21 05 31
Telephone international: +46 8 787 00 00

The Role of Firm Heterogeneity for the Transmission of Aggregate Shocks

Michele Lenza* Giuseppe Pagano Giorgianni†

Lorenza Rossi‡ Ettore Savoia§

Sveriges Riksbank Working Paper Series No.465¶
April 29, 2026

Abstract

We study whether firm-level heterogeneity helps explain U.S. macroeconomic fluctuations in response to aggregate shocks. Using quarterly Compustat and CRSP data from 1986 to 2025, we construct two revenue-based statistics inspired by the Melitz (2003) model: the average firm and the marginal near-default firm. These statistics summarize key features of the firm distribution. We augment a Bayesian VAR with these measures and compare its performance to a standard aggregate VAR and to a functional VAR that incorporates the full cross-sectional distribution of firm revenues.

We find that firm-level heterogeneity contains information not captured by aggregate variables. Including the two statistics allows the VAR to closely replicate the impulse responses obtained using the functional VAR and improves out-of-sample forecast accuracy. These findings are robust to a replication using UK data.

Keywords: Firm heterogeneity; Entry and exit; Business cycles; Bayesian VAR; Functional VAR; Selection and default; Sufficient statistics; Aggregate shocks.
JEL Classification: E32; E37; D22; C32; G33.

---

*European Central Bank and CEPR
†Sapienza University of Rome
‡Lancaster University, FREcon, Bailrigg LA14YW, Lancaster (UK). E-mail: l.rossi@lancaster.ac.uk
§Sveriges Riksbank and CeMoF
¶The opinions expressed in this article are the sole responsibility of the authors and should not be interpreted as reflecting the views of Sveriges Riksbank, European Central Bank and the Eurosystem.

1 Introduction

Macroeconomic models increasingly study how heterogeneity across households and firms shapes aggregate fluctuations. The Global Financial Crisis and the Covid-19 recession made clear that aggregate downturns often coincide with pronounced cross-sectional reallocation, and that some shocks operate through changes in the distribution of agents rather than through a representative unit. Yet heterogeneity is not a free lunch. Modeling distributions increases state dimensionality, complicates analysis, and reduces transparency. This raises a practical question: if the goal is to understand the transmission of aggregate shocks, how much heterogeneity is needed, and of what kind? This is the question we address in this paper.

A growing literature has addressed this question for households, with mixed results. Some studies find limited aggregate implications of household heterogeneity (Debortoli and Galí, 2022; Chang and Schorfheide, 2022; Chang et al., 2024; Bayer et al., 2024; Huber et al., 2024; Lenza and Tristani, 2025), while others emphasize amplification through inequality and borrowing constraints (Bilbiie et al., 2023; Krueger et al., 2016). Taken together, this evidence suggests that, even when micro data are informative for many purposes, household distributional dynamics often do not materially alter aggregate responses to business-cycle shocks in VAR-type frameworks.

The case for focusing on firms is distinct. Firms are not only heterogeneous but also central to the canonical transmission channels of aggregate shocks. They make investment and employment decisions, face credit constraints, and determine entry and exit. As aggregate conditions change, differences in productivity, size, and financial constraints lead firms to adjust unevenly across the cross-section. These heterogeneous responses can shape the propagation of shocks and influence aggregate fluctuations.

The macroeconomic literature on heterogeneous firms has expanded substantially over the past decades and can be broadly grouped into two traditions: Hopenhayn (1992)-type industry dynamics models and Melitz (2003)-type models.[1] Both frameworks share common

[1] The Hopenhayn (1992) partial equilibrium model characterizes an industry with endogenous entry and exit and heterogeneous firms whose size and survival depend on idiosyncratic productivity draws. General equilibrium extensions in this tradition, including Clementi and Palazzo (2016), Khan and Thomas (2008), and Khan and Thomas (2013), show how financial frictions, capital irreversibility, and firm turnover shape amplification and persistence. In the Melitz (2003) tradition, firm heterogeneity operates through selection and the extensive margin. Applications such as Ghironi and Melits (2005) and Ghironi and Melitz (2007) integrate this structure into New Keynesian open-economy models. More recent contributions, including Rossi (2019), Ascari et al. (2023), and Fasani et al. (2023), embed endogenous firm dynamics into models of technology, supply, and uncertainty shocks. A related strand emphasizes heterogeneous firms in the presence of financial frictions, for example Gertler and Gilchrist, 1994a, Gertler and Gilchrist, 1994b,

foundations, including firm heterogeneity, entry and exit, and selection effects. The Melitz (2003) structure offers an important analytical advantage. Aggregation does not require tracking the full distribution of firms but can instead be summarized by two sufficient statistics, namely the average firm and the marginal firm close to default. This makes the framework particularly attractive for macroeconomic applications, where tracking the full distribution would otherwise be computationally burdensome. While this parsimonious representation may not capture how aggregate shocks reshape the entire firm distribution, a small number of statistics may still be sufficient to characterize aggregate responses to such shocks. This is especially relevant in VAR models, where including many cross-sectional variables leads to parameter proliferation and weaker identification.

This tension between distributional richness and empirical tractability raises a central question: what is the minimal representation of firm heterogeneity needed to capture the transmission of aggregate shocks? It remains unclear how much, and which dimensions, of firm heterogeneity must be incorporated into empirical macroeconomic models, particularly VARs, to explain the response to aggregate shocks.

Recently, Lenza and Savoia (2024) estimate a functional VAR (FVAR) that combines euro area macro aggregates with the full cross-sectional distribution of firm revenues and show that firm-level distributions contain information not spanned by aggregate variables. This finding is particularly striking because it contrasts with the evidence for households and points to an omitted-information problem in standard aggregate VARs.

Marcellino et al. (2024), develop a functional VAR that accounts for firm-level heterogeneity along multiple dimensions to study the propagation of total factor productivity (TFP) shocks and their effects on aggregate fluctuations and on the cross-sectional allocation of capital and labor.

In this paper, we go beyond documenting distributional responses or assessing the importance of firm heterogeneity. We provide new empirical evidence on the role of firm heterogeneity in the transmission of aggregate shocks and, importantly, develop a theoretical framework based on Melitz (2003)-type models that guides the identification of sufficient statistics of the firm distribution. These statistics summarize the relevant cross-sectional information for aggregate dynamics without requiring the full distribution to be explicitly modeled.

We combine U.S. quarterly macroeconomic data with firm-level information from Compustat and CRSP over a long sample (1986Q1-2025Q3).

Gilchrist et al., 2014, Arellano et al., 2019, Cooley and Quadrini, 2006, Ottonello and Winberry, 2020, and Cloyne et al., 2020.

Our starting point is empirical: we ask whether the distribution of firm outcomes contains information that changes the estimated transmission of business-cycle shocks. Methodologically, we build on Chang et al. (2024), treating the firm distribution as an evolving object embedded in a VAR system. The advantage of this approach is that it permits general transmission mechanisms at the macro level (the VAR component) and highly flexible dynamics in the distribution (the functional component). But this strength also highlights a central issue for economic modeling: if the data say that "the distribution matters", which aspect of the distribution should a modeler incorporate in a tractable macro model?

To discipline the modeling implications of this evidence, we turn to a parsimonious micro-founded framework: a Melitz (2003)-type model with endogenous entry and exit. As discussed above, a key advantage of this framework is its analytical tractability.

In this environment, the cross-sectional distribution of active firms can be summarized by two economically meaningful statistics: the cutoff productivity, which identifies the marginal firm close to exit, and the average productivity of surviving firms, which identifies the average firm. Under the Pareto benchmark used here, average productivity is proportional to the cutoff productivity, so the two objects should not be interpreted as independent state variables. We nevertheless retain both because they capture two economically distinct margins: selection at the cutoff and average conditions among surviving firms.

Our paper takes the sufficient-statistics implication of the Melitz (2003) framework to the data. Directly measuring firm-level productivity would require strong assumptions on production functions and introduce substantial measurement challenges. Instead, we work with firm revenues. Within the Melitz (2003) framework, we show that the revenues of the average and marginal firms are monotonic functions of average productivity and cutoff productivity, respectively, and therefore provide valid empirical counterparts to these theoretical objects. Guided by this mapping, we measure the average firm using mean revenues and proxy the marginal firm using revenues of firms close to exit, identified through distance-to-default measures in the spirit of Ottonello and Winberry (2020). We then include these revenue-based statistics in an otherwise standard aggregate Bayesian VAR to assess whether they capture the relevant cross-sectional information. We find that they do: a small number of revenue moments are sufficient to summarize the role of firm heterogeneity in the transmission of aggregate shocks.

Our empirical design compares three models. The first is a conventional aggregate

BVAR estimated on U.S. macroeconomic variables (real GDP, real consumption, real investment, unemployment, CPI inflation, a short-term nominal interest rate, and the excess bond premium). The second is a high-dimensional functional VAR that augments the macro system with the cross-sectional density of firm revenues, approximated using a spline-basis representation. The third is a parsimonious augmented VAR that adds only our two sufficient statistics to the aggregate BVAR, retaining tractability while being grounded in theory. Across these models, we study (i) whether firm distributions represent omitted information for macro dynamics, (ii) whether a small number of statistics can summarize this information, and (iii) whether incorporating firm heterogeneity improves empirical performance.

The results deliver a clear message. First, firm heterogeneity matters for the estimated transmission of aggregate business-cycle shocks in the United States. When we augment the BVAR with firm revenues' distribution, generalized impulse responses to a GDP-identified business-cycle shock differ quantitatively from those obtained in an aggregate-only BVAR, indicating that the cross-section contains information relevant for propagation. Second, we find that we do not need the whole distribution to capture these effects. A VAR that includes only the average and marginal firm statistics closely tracks the responses implied by the functional VAR, suggesting that these two objects summarize the distributional information that is relevant for aggregate dynamics. Third, the parsimonious representation delivers comparable, and in several cases improved, predictive performance. In an out-of-sample recursive forecasting exercise, augmenting the BVAR with the two statistics improves short-horizon forecast accuracy for several key macro variables, and performs competitively relative to models that use the full distribution. Our results are robust to excluding major crisis episodes. Re-estimating the models without the Covid-19 period and, separately, without the 2008–2009 Global Financial Crisis leaves the main conclusions unchanged: the VAR augmented with the sufficient statistics continues to replicate the functional VAR responses, whereas the aggregate-only VAR does not.

Taken together, our findings imply that heterogeneity among U.S. firms plays a meaningful role in shaping aggregate fluctuations in response to aggregate shocks, and that this role can be captured by a low-dimensional, economically interpretable representation inspired by Melitz (2003). In the Appendix, we show that these findings are robust to a replication using UK data, further corroborating the usefulness of the Melitz framework as a theoretical guide for identifying empirically relevant features of the firm distribution.

We do not interpret the Melitz (2003) model as a literal structural representation of the economy. Instead, we use it as a benchmark to identify cross-sectional objects that a broad class of heterogeneous-firm models and empirical evidence on firm dynamics suggest are relevant for aggregate transmission. In particular, we focus on statistics capturing both the average firm and the marginal firms that drive entry and exit decisions. In particular, we focus on statistics capturing both the average firm and the marginal firms that drive entry and exit decisions.

Our objective is not to validate this specific framework or to claim that it outperforms richer models with more detailed forms of heterogeneity. Rather, we build on empirical regularities that are common across a wide class of models with heterogeneous firms. Firm dynamics shape aggregate outcomes, entry and exit are business-cycle relevant, and exit is closely related to firms' default risk.

From this perspective, the statistics we construct should be interpreted as practical sufficient statistics for firm heterogeneity. By embedding them in a VAR framework, we allow the data to assess their relevance for macroeconomic dynamics and the transmission of aggregate shocks, without imposing the structural restrictions of any particular model.

A natural question is why firm-level distributions appear to contain incremental information for aggregate dynamics, whereas heterogeneity across households seems to play a more limited role. One possible explanation is that aggregate shocks affect firms asymmetrically. Firms close to financial distress may experience changes in survival risk and respond more strongly in terms of investment, hiring, and production. In models with endogenous selection, this corresponds to firms near the exit margin. More generally, shifts in the lower tail of the firm distribution can influence aggregate outcomes through selection and state-dependent adjustments, and these effects may not be fully captured by aggregate variables alone.

By contrast, household income and consumption may be partly stabilized by government automatic stabilizers, such as taxes and transfer programs. Although we do not directly test this mechanism for households, our findings are consistent with the view that firms near distress play a key role in shaping the informational content of firm-level distributions for aggregate dynamics.

The remainder of the paper proceeds as follows. Section 2 develops the Melitz-based argument for sufficiency and describes the construction of the average and marginal firm statistics. Section 3 describes the data, construction of firm-level distributions and default-based measures, outlines the functional VAR framework and our Bayesian estimation strategy, and evaluates the ability of the VAR augmented with the sufficient statistics to match the functional VAR and its forecasting performance. Section 3.3 presents the results and discusses forecasting performance of the model and robustness checks that are reported in the Appendix. Section 4 concludes.

2 A Melitz-Type Framework and Sufficient Statistics

To guide the empirical analysis, we consider a stripped-down real version of a Melitz (2003) economy with endogenous entry and exit. The framework is used solely to identify which features of the firm distribution are theoretically relevant for aggregate transmission.

Firms are heterogeneous in idiosyncratic productivity $z_{i,t}$, drawn from a Pareto distribution with support $[z_{\min}, \infty)$ and shape parameter ξ. Production uses only labor,

$$y_{i,t} = z_{i,t} l_{i,t}. \tag{1}$$

Intermediate varieties are aggregated through a CES technology with elasticity of substitution $\theta > 1$.

Revenues. Firm i faces CES demand: $y_{i,t} = \left(\frac{p_{i,t}}{P_t}\right)^{-\theta} Y_t$, where P_t is the aggregate price index and Y_t aggregate output. The marginal cost of firm i is $w_t/z_{i,t}$. Profit maximization implies a constant markup over marginal cost, $p_{i,t} = \frac{\theta}{\theta-1}\frac{w_t}{z_{i,t}}$. Substituting into demand yields equilibrium revenues

$$r_{i,t} = p_{i,t} y_{i,t} = \Omega_t z_{i,t}^{\theta-1}, \tag{2}$$

where, $\Omega_t \equiv P_t^{\theta} Y_t \left(\frac{\theta}{\theta-1} w_t\right)^{1-\theta}$. Revenues are therefore strictly increasing in productivity.

Exit and the Cutoff. Firms face a fixed operating cost f_t. Profits are given by

$$\pi_{i,t} = r_{i,t} - w_t l_{i,t} - f_t. \tag{3}$$

Since revenues and scale are increasing in productivity, profits are increasing in $z_{i,t}$. There exists an endogenous cutoff productivity $\bar{z}_t$ such that firms operate if and only if $z_{i,t} \geq \bar{z}_t$, defined by

$$\pi_t(\bar{z}_t) = 0. \tag{4}$$

Entry. Prospective entrants pay a sunk entry cost EC_t and draw productivity from the Pareto distribution. Upon drawing productivity, they operate only if $z_{i,t} \geq \bar{z}_t$.

Free entry implies

$$EC_t = v_t(\tilde{z}_t), \tag{5}$$

where $v_t(\bar{z}_t)$ denotes the discounted value of profits conditional on operating, i.e. for firms with productivity above the cutoff.[2] Entry therefore, depends on the expected profitability of surviving firms, while exit is governed by the cutoff productivity.

Firm Dynamics. Let N_t denote the mass of operating firms. Under the Pareto distribution, the exit probability implied by the cutoff is

$$\eta_t = 1 - \left(\frac{z_{\min}}{\bar{z}_t}\right)^{\xi}. \tag{6}$$

Under the Pareto distribution, the average productivity of surviving firms is proportional to the cutoff, $\tilde{z}_t = \left[\frac{\xi}{\xi-(\theta-1)}\right]^{\frac{1}{\theta-1}} \bar{z}_t$. Since both incumbents and entrants draw productivity before producing, the mass of firms evolves according to: $N_t = (1-\eta_t)\left(N_{t-1} + N^E_{t-1}\right)$.

Average and Marginal Firms. The marginal firm is the firm with productivity $\bar{z}_t$. Its revenue is

$$r_t^{marg} = r(\bar{z}_t) = \Omega_t \bar{z}_t^{\theta-1}. \tag{7}$$

The revenue of the average firm is instead,

$$r_t^{avg} = r(\tilde{z}_t) = \Omega_t \tilde{z}_t^{\theta-1}. \tag{8}$$

as shown in Melitz (2003) this coincides with the average revenues.

Aggregation. Aggregate output is obtained by integrating over operating firms,

$$Y_t = \left[\int_{\bar{z}_t}^{\infty} y_{i,t}^{\frac{\theta-1}{\theta}} g(z_{i,t}) dz_i\right]^{\frac{\theta}{\theta-1}}. \tag{9}$$

Under the Pareto assumption, aggregate output depends on the mass of firms and on the cross-sectional distribution of productivity among operating firms. Similarly, aggregate revenues satisfy

$$R_t = \int_{\bar{z}_t}^{\infty} r_{i,t} g(z_{i,t}) dz_i = N_t r_t^{avg}. \tag{10}$$

[2] As shown in the Online Appendix $v_t(\bar{z}_t)$ is already weighted by the probability of survival.

Thus, conditional on the mass of firms N_t, aggregate output and aggregate revenues depend only on the average firm. Selection operates through the cutoff $\bar{z}_t$, which determines both the exit probability and, through the Pareto structure, average productivity $\tilde{z}_t$.

Under the Pareto structure, the truncated distribution of operating firms is characterized by the cutoff productivity $\bar{z}_t$, while average productivity $\tilde{z}_t$ is a deterministic function of that cutoff. We nevertheless retain both the marginal and the average firm because they capture two economically meaningful margins: the selection margin and the average conditions among surviving firms. Their revenue counterparts, (r_t^{marg}, r_t^{avg}), are therefore useful observable proxies in the empirical analysis.

2.1 RBC Extension and State-Space Representation

The Melitz-type structure described above can be embedded in a simplified real business cycle (RBC) framework without capital accumulation, augmented with heterogeneous firms, endogenous entry, and exit. The objective of this extension is not to provide a fully structural quantitative model, but to demonstrate that the mechanisms highlighted above admit a compact state-space representation. In particular, the model can be reduced to a low-dimensional dynamic system in which firm heterogeneity affects aggregate fluctuations through entry, exit, and selection—precisely the margins we examine empirically in the VAR. The full RBC model and the derivation of the reduced state-space representation are provided in a separate Technical Appendix.

In the RBC version, aggregate productivity evolves according to

$$\widehat{a}_{t+1} = \rho_a \widehat{a}_t + \varepsilon^a_{t+1}, \tag{11}$$

and firm dynamics are governed by the zero-profit condition for the marginal firm, the free-entry condition, and the law of motion for the mass of firms.

Log-linearizing the equilibrium conditions around the steady state yields a system that can be reduced to three variables: aggregate output $\widehat{y}_t$, exit probability $\widehat{\eta}_t$, and aggregate productivity $\widehat{a}_t$. Defining the state variable

$$s_t \equiv \widehat{a}_t,$$

and the vector of forward-looking variables

$$x_t \equiv \begin{bmatrix} \widehat{y}_t \\ \widehat{\eta}_t \end{bmatrix},$$

the equilibrium admits a minimal state-space representation of the form

$$E_t x_{t+1} = A_{xx} x_t + A_{xs} s_t, \tag{12}$$

with

$$s_{t+1} = \rho_a s_t + \varepsilon^a_{t+1}. \tag{13}$$

All remaining equilibrium variables, including the cutoff productivity $\bar{z}_t$, average productivity $\tilde{z}_t$, average and marginal revenues, entry, the mass of firms, consumption, and labor, can be recovered as linear functions of (x_t, s_t). In particular, average and marginal revenues are linear functions of the state vector and therefore provide a low-dimensional summary of firm heterogeneity in the dynamic equilibrium. Because these revenue-based statistics span the same information about selection and aggregate conditions as (x_t, s_t) in the log-linearized system, they can be used as empirical sufficient statistics for firm heterogeneity in the VAR.

Mapping to the Data and Link to the VAR. The theoretical framework highlights two cross-sectional objects that are central for aggregate transmission: the marginal firm, defined by the cutoff productivity $\bar{z}_t$, and the average firm, defined by the average productivity of surviving firms $\tilde{z}_t$. In the RBC extension, both objects are linear functions of the low-dimensional state vector that drives aggregate dynamics.

Productivity is not directly observable in the data. However, under monopolistic competition, firm revenues are monotonic functions of productivity. In particular, revenues of the marginal and average firms are direct functions of $\bar{z}_t$ and $\tilde{z}_t$, respectively. This monotonic mapping allows us to construct empirical counterparts to the theoretical sufficient statistics using revenue data, without imposing additional assumptions on production functions or requiring firm-level productivity estimation.

In the empirical analysis, we estimate a VAR and identify an innovation to GDP using a Cholesky decomposition. This GDP shock should be interpreted as a reduced-form forecast error, which in general reflects a combination of structural disturbances.[3]

[3] The theoretical model does not impose a one-to-one mapping between this reduced-form shock and the primitive productivity innovation. Rather, it clarifies how aggregate fluctuations, whatever their source,

Since average and marginal revenues are linear functions of the state vector in the log-linearized equilibrium, embedding these statistics in the VAR provides a direct empirical test of whether a low-dimensional representation of firm heterogeneity is sufficient to capture the transmission of aggregate GDP shocks.

3 Empirical Analysis

3.1 Data

Our empirical analysis relies on U.S. quarterly data spanning the period 1986Q1–2025Q3. The sample covers multiple business cycles, including the Great Recession and the COVID-19 crisis, allowing us to study firm dynamics across episodes of substantial macroeconomic stress. Macroeconomic aggregates are drawn from the FRED database. The baseline vector includes real GDP, real consumption, real investment, the unemployment rate, GDP deflator, the Federal Fund rate, and the Excess Bond Premium (EBP), constructed following Favara et al. (2016). These variables jointly capture the key dimensions of aggregate fluctuations: real activity, labor market conditions, price dynamics, the stance of monetary policy, and financial conditions. The inclusion of the EBP is particularly important, as it provides a forward-looking measure of credit market stress and risk premia that may interact with firm-level default risk and selection dynamics.

Firm-level data are constructed by merging quarterly balance-sheet data from Compustat with stock market data from CRSP. Total firms' revenues from Compustat serve as our primary firm-level measure. They are used both to introduce heterogeneity into the functional VAR and to construct the two sufficient statistics that summarize the cross-sectional distribution.[4]

To capture the conditions of firms near the exit margin, we compute distance-to-default (DD) and expected default probability (EDF) measures. Distance to Default and Expected Default Frequency are computed according to the Merton (1974) model, following the methodology of by Bharath and Shumway (2008). The construction uses long-term debt and debt in current liabilities from Compustat, together with equity returns, stock prices (or bid–ask averages), and shares outstanding from CRSP. Distance-to-default provides a forward-looking measure of balance-sheet vulnerability, while EDF translates this distance

propagate through firm entry, exit, and selection.

[4]In monopolistic competition environments such as the Melitz framework, firm revenues are monotonic functions of firm productivity. This property allows revenues to serve as observable proxies for the theoretical objects that characterize the cross-sectional distribution.

into a probability of default. These measures allow us to identify firms that are close to default and therefore near the endogenous exit threshold emphasized by Melitz-type models with selection. All the details about the construction of the firm level data on distance to default and its expected probability to default counterpart are reported in Appendix A.1.

Distance-to-Default and Expected Default Probability. Figure 1 plots the aggregate distance-to-default and expected default probability over time constructed using the micro data.

Figure 1: Distance-to-Default and Expected Default Probability (1986Q1–2025Q3)

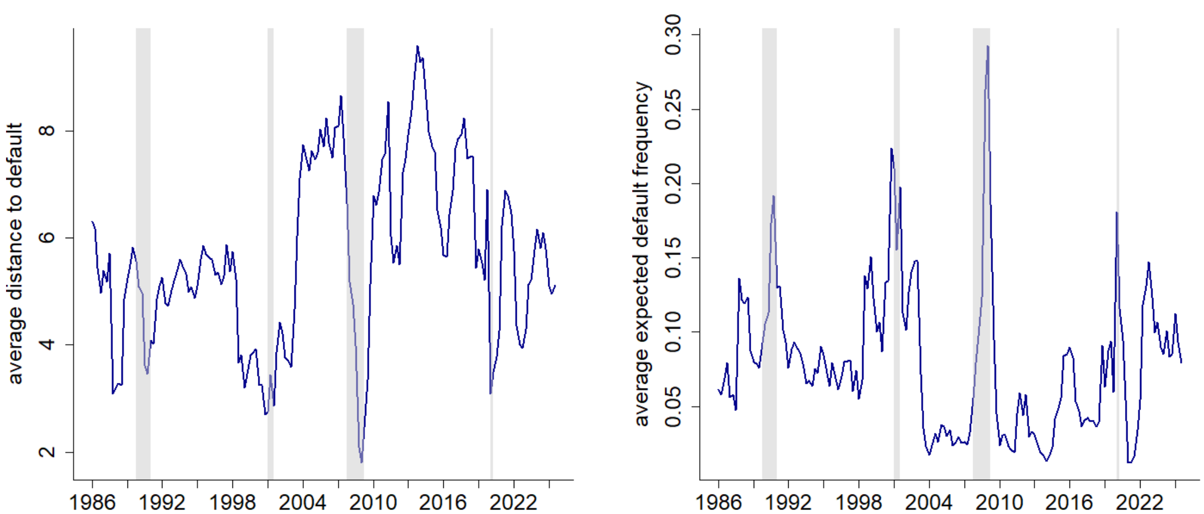

Notes: The plots report the average distance to default (right-hand-side) and expected default frequency (right-hand-side). Shaded areas represent NBER recessions. Source: Compustat, CRSP. Sample: 1986Q1–2025Q3.

Both measures display strong cyclical variation. Distance-to-default declines sharply during recessions, while expected default probability spikes during periods of financial stress, notably around the Global Financial Crisis and the COVID-19 recession. This behavior supports the interpretation of these measures as capturing conditions of firms near the exit margin.

Average and Marginal Revenues. By combining revenue information with default-based measures, we construct empirical counterparts to the average and marginal firms that play a central role in the theoretical framework.

The average revenue is constructed as the cross-sectional mean of firm-level revenues in the sample and is normalized by aggregate GVA.

We proxy the model's marginal firm, defined by the exit cutoff, using firms that are empirically close to default. We start from distance-to-default measures constructed from CRSP equity prices and Compustat balance-sheet items and map them into expected default probabilities (EDF). This step is useful because EDF is directly interpretable as an exit-risk object and is therefore conceptually closer to the model's exit probability η_t than the raw distance-to-default.

We define the marginal-firm statistic as the average revenue among firms with high EDF. Our baseline threshold is EDF $\geq 12\%$, which identifies financially distressed firms, i.e. firms with an expected one-year default probability of at least 12%. This threshold is consistent with the average annual default probability, η_{t+1}, commonly used in calibrations of the Melitz (2003)-type models for the U.S. economy (see for example Bilbiie et al., 2012, Fasani et al., 2023, among many others).. In the Appendix A.4, we show that our results are robust to alternative threshold values, considering both a lower threshold of 9% and a higher threshold of 15%. Firms at or above this threshold are classified as marginal.[5]

Remarkably, our results are robust when using the mean instead of the median for both the marginal and the average revenue measures.

Figure 2 reports the evolution of average firm revenues and revenues of firms close to default, both expressed as a percentage of aggregate GVA.

Average revenues track aggregate activity closely, while revenues of firms close to default display substantially stronger cyclical sensitivity. In downturns, marginal-firm revenues decline disproportionately, consistent with intensified selection dynamics. This pattern is precisely what the Melitz framework predicts: aggregate fluctuations operate partly through movements in the cutoff productivity and changes in the survival margin.

These two series constitute our empirical counterparts to the theoretical average and marginal firms.

[5]Figure 10 in the Appendix reports the share of firms in our sample classified as marginal in each quarter.

Figure 2: Average and Marginal Revenues

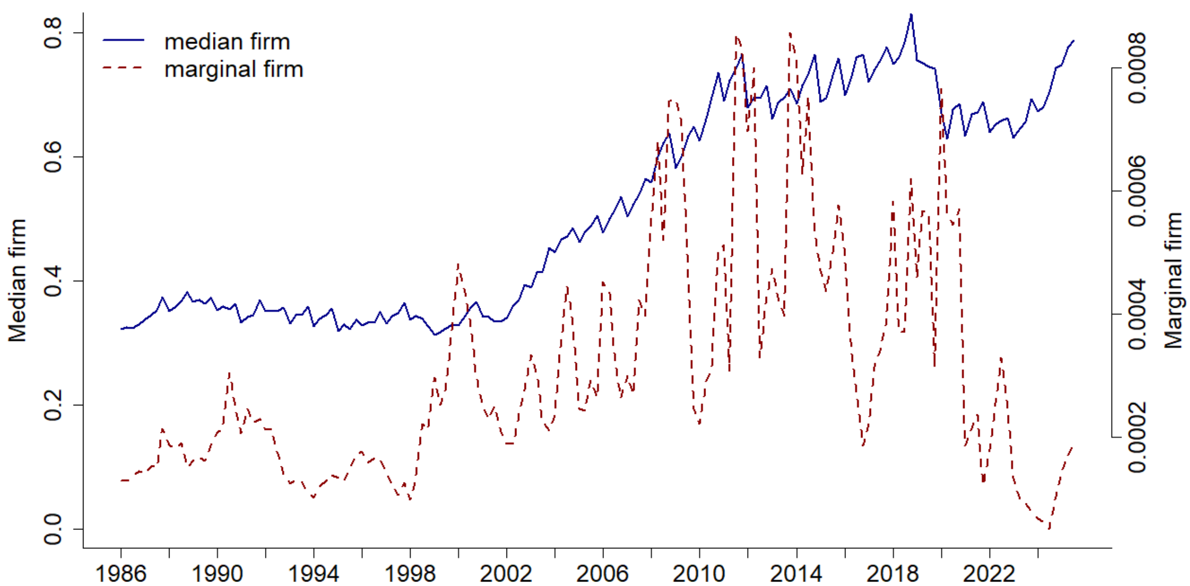

Notes: The figure reports the median of revenues for the entire distribution (blue solid line) and for the marginal firm, identified as the firm for which the expected default frequency is equal or above 12% (red dashed line). Source: Compustat, U.S. Bureau of Economic Analysis. Sample: 1986Q1–2025Q3.

3.2 Empirical Models

Our empirical strategy is designed to assess whether the cross-sectional distribution of firms contains information that is not spanned by standard macroeconomic aggregates and, more importantly, whether this information can be summarized by a small number of economically motivated statistics. To this end, we proceed in three steps, moving from a purely aggregate representation to increasingly richer descriptions of firm heterogeneity.

We begin with a standard Bayesian vector autoregression (BVAR) that includes only macroeconomic variables.

Let Y_t denote the vector of transformed macroeconomic variables used in the VAR, including real GDP, real consumption, real investment, the unemployment rate, CPI inflation, the short-term nominal interest rate, and the Excess Bond Premium. The baseline specification is estimated on these transformed series with four lags.

$$Y_t = A(L)Y_{t-1} + u_t, \tag{14}$$

where $A(L)$ is a matrix lag polynomial and u_t is a vector of reduced-form innovations. We adopt Minnesota priors within a hierarchical Bayesian framework following Giannone et al. (2015) and Lenza and Primiceri (2022), which allows the hyperparameters to be determined by the data. This benchmark model provides a flexible yet disciplined representation of aggregate dynamics and serves as a reference point for evaluating the incremental contribution of firm-level information.

We then consider a richer framework that incorporates the full cross-sectional distribution of firm revenues. Following Chang et al. (2024), we specify a functional VAR in which macroeconomic variables interact with the entire log-density of firm revenues, denoted by $l_t(x)$. The system can be written as

$$\begin{cases} Y_t = B(L)Y_{t-1} + \int B_{yl}(L,x)l_{t-1}(x)dx + v_{y,t}, \\ l_t(x) = B_{ly}(L,x)Y_{t-1} + \int B_{ll}(L,x,s)l_{t-1}(s)\,ds + v_{l,t}(x) \end{cases} \tag{15}$$

where the evolution of macro aggregates depends on the past distribution of firm revenues, and the distribution itself responds to past macroeconomic conditions. Because this representation is infinite-dimensional, we approximate the density using a finite set of basis functions. Specifically, we write

$$l_t(x) = \sum_{i=1}^{K} \alpha_{i,t}\zeta_i(x), \tag{16}$$

where K indicates the dimension of the basis expansion, $\{\zeta_i(x)\}_{i=1}^K$ is a set of basis functions and $\alpha_{i,t}$ are time-varying coefficients. This transformation reduces the model to a finite-dimensional VAR in (Y_t, α_t). The coefficients $\alpha_{i,t}$ summarize the dynamics of the firm revenue distribution, and impulse responses and forecasts of the density can be recovered from their evolution. We consider a set of different models with K ranging from 4 to 12. Our baseline model corresponds to $K = 12$. In figure 3 we show that the basis expansion provides a good approximation of the distribution of the data over time.

Figure 3: Cross-sectional distribution of firm revenues

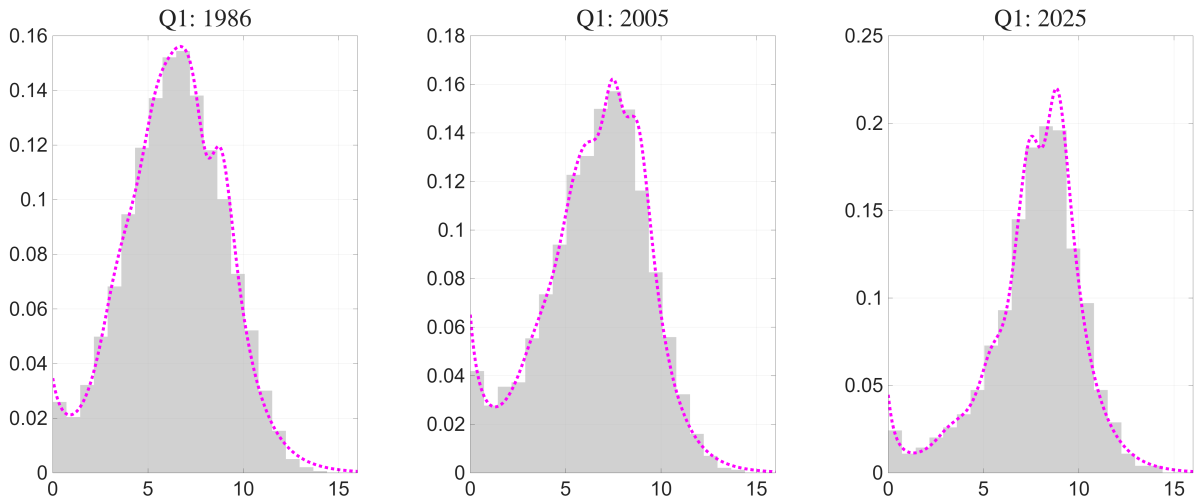

Notes: The grey bars show the histogram of the empirical distribution of firm revenues. The solid line shows the log-spline density estimate for $K = 12$ at three points in time.

This functional VAR provides a benchmark for evaluating the informational content of the entire cross-section.

Our key contribution is to assess whether the informational content of the full distribution can instead be captured by a low-dimensional representation motivated by the Melitz framework. Rather than including the entire density, we augment the macro BVAR with two observable statistics: average firm revenues and average revenues of firms close to default. The first series captures movements in the central tendency of the distribution and serves as a proxy for the average firm. The second series isolates firms near the exit margin—identified using distance-to-default measures—and serves as a proxy for the marginal firm.

This augmented specification, which we refer to as the "Melitz-BVAR," allows us to test whether these two economically grounded statistics replicate the information contained in the full distribution. If the Melitz-BVAR performs comparably to the functional VAR in terms of impulse responses and predictive content, this would imply that the cross-sectional information relevant for aggregate transmission can be summarized by a small number of sufficient statistics. The empirical comparison therefore provides a direct test of whether firm heterogeneity matters for aggregate dynamics and whether its macroeconomic relevance can be captured without modeling the full distribution explicitly.

3.3 Impulse Responses to a Business-Cycle Shock

This section conducts an impulse response analysis of the main macroeconomic variables following an identified aggregate shock, to assess whether the distribution of firm revenues contains information not spanned by aggregate variables, thus constituting omitted information in the sense of Lenza and Savoia (2024). Specifically, we identify a business-cycle shock in both the BVAR and the FVAR by imposing a recursive ordering in which real GDP is placed first. This shock should not be interpreted as structural; rather, it represents a reduced-form innovation capturing the main drivers of GDP fluctuations, in the spirit of Angeletos et al. (2020). Figure 14 reports the impulse response functions for the two models, scaled so that GDP exhibits the same impact effect across specifications. The FVAR dynamics (blue areas) are more attenuated than those implied by the aggregate-only specification (red areas), supporting the view that the firm-level distribution contains information relevant for the transmission of aggregate shocks.

We then repeat the same exercise for an aggregate model augmented with our set of sufficient statistics. In Figure 15, we compare the impulse responses to the same shock obtained from this statistic-augmented specification with those generated by the heterogeneous model with $K = 12$. The estimated impulse response functions display a high degree of overlap. Crucially, this suggests that the two sufficient statistics contain enough information from the firm-level distribution to capture most of the dynamics identified by the FVAR, consistent with the intuition of Melitz (2003)-type models.

Forecasting Performance Table 1 in the Appendix shows that the FVAR outperforms the VAR in terms of out-of-sample forecasting performance. The Melitz-VAR provides forecasting performance broadly similar to the BVAR but remains slightly weaker than the preferred FVAR specification. We also compare the model's performance using alternative statistics to summarize the cross-sectional distribution of firm revenues, including the variance, the interquantile range (95th–5th percentile), and the 95th percentile, which captures large firms. In all cases the Melitz-VAR, based on the marginal-firm threshold outperforms these alternatives in terms of forecasting accuracy. Finally, 2 compares the performance of the three models by repeating the forecasting exercise in rolling windows, with a window width of 60 quarters, obtaining comparable results.

Figure 4: Impulse response functions to a business cycle shock

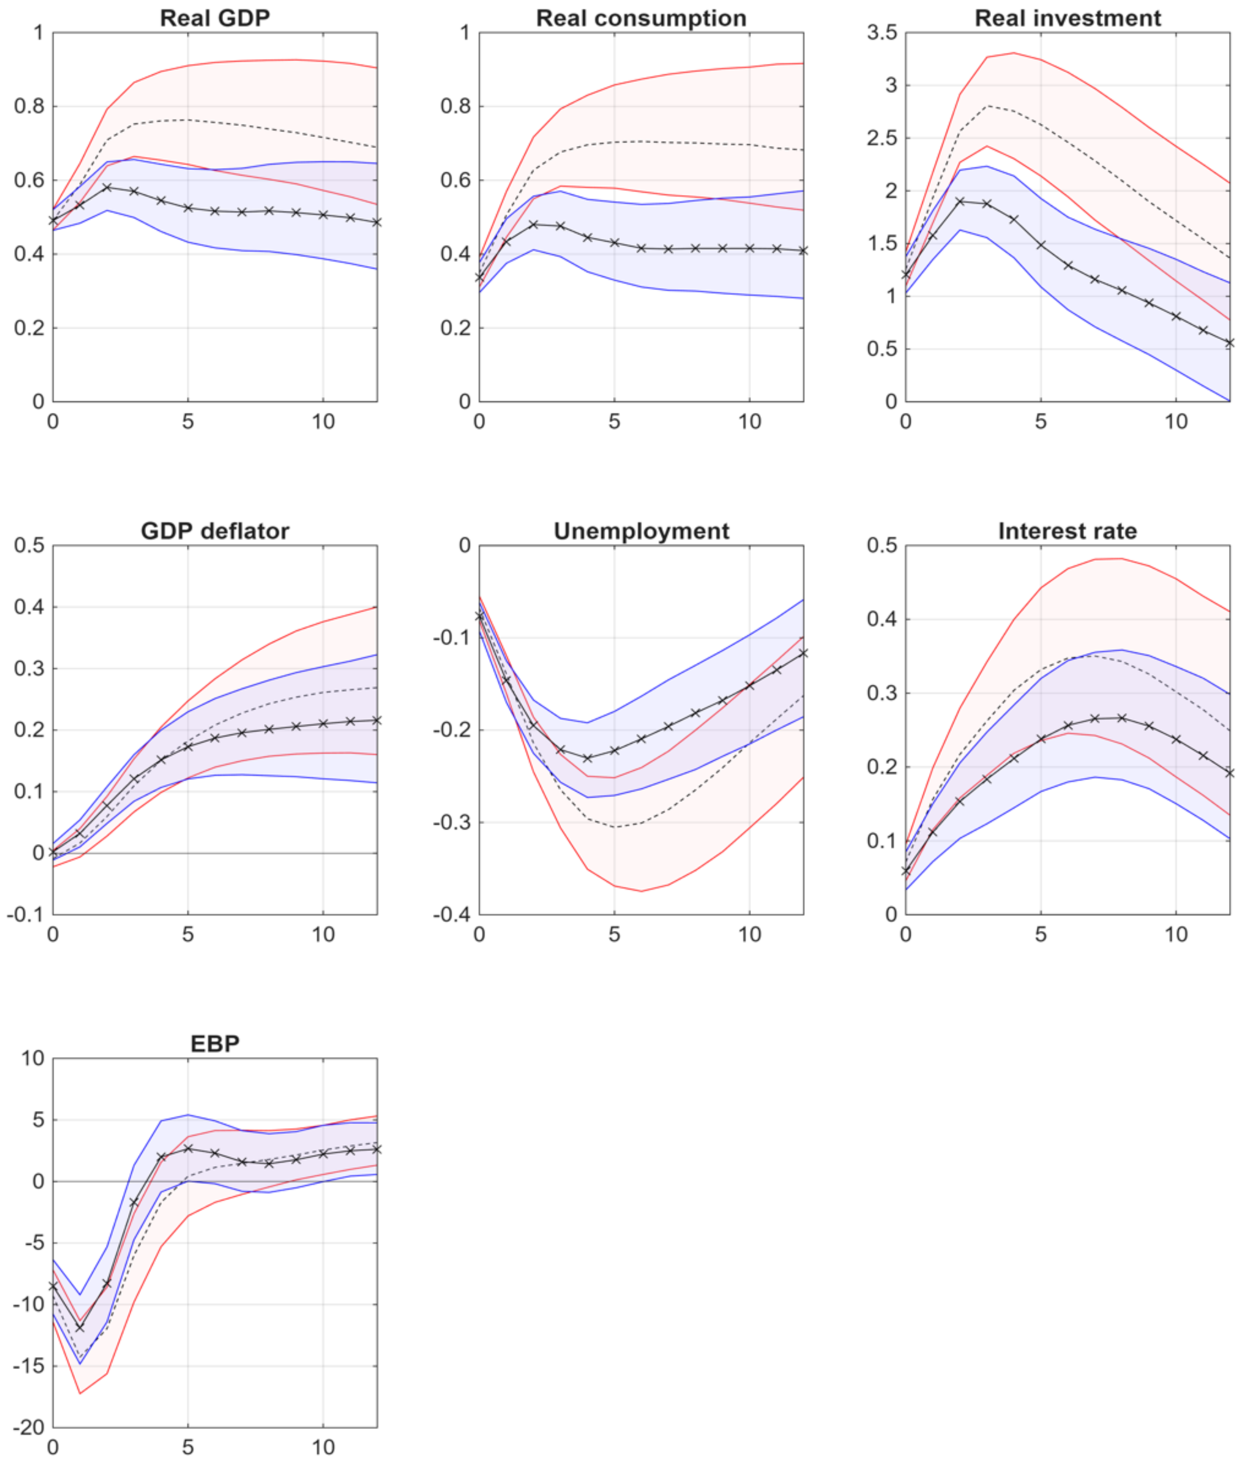

Notes: The figure reports the impulse response functions of real gdp, real consumption, real investment, GDP deflator, unemployment rate, interest rate and EBP to a business cycle shock. **Blue Area**: FVAR, K=12; **Red Area**: aggregate BVAR; Dashed line: BVAR median; Solid line with crosses: FVAR median. Source: U.S. Bureau of Economic Analysis, U.S. Bureau of Labor Statistics, Board of Governors of the Federal Reserve System, Favara et al. (2016). Sample: 1986Q1–2025Q3.

Figure 5: Impulse response functions to a business cycle shock

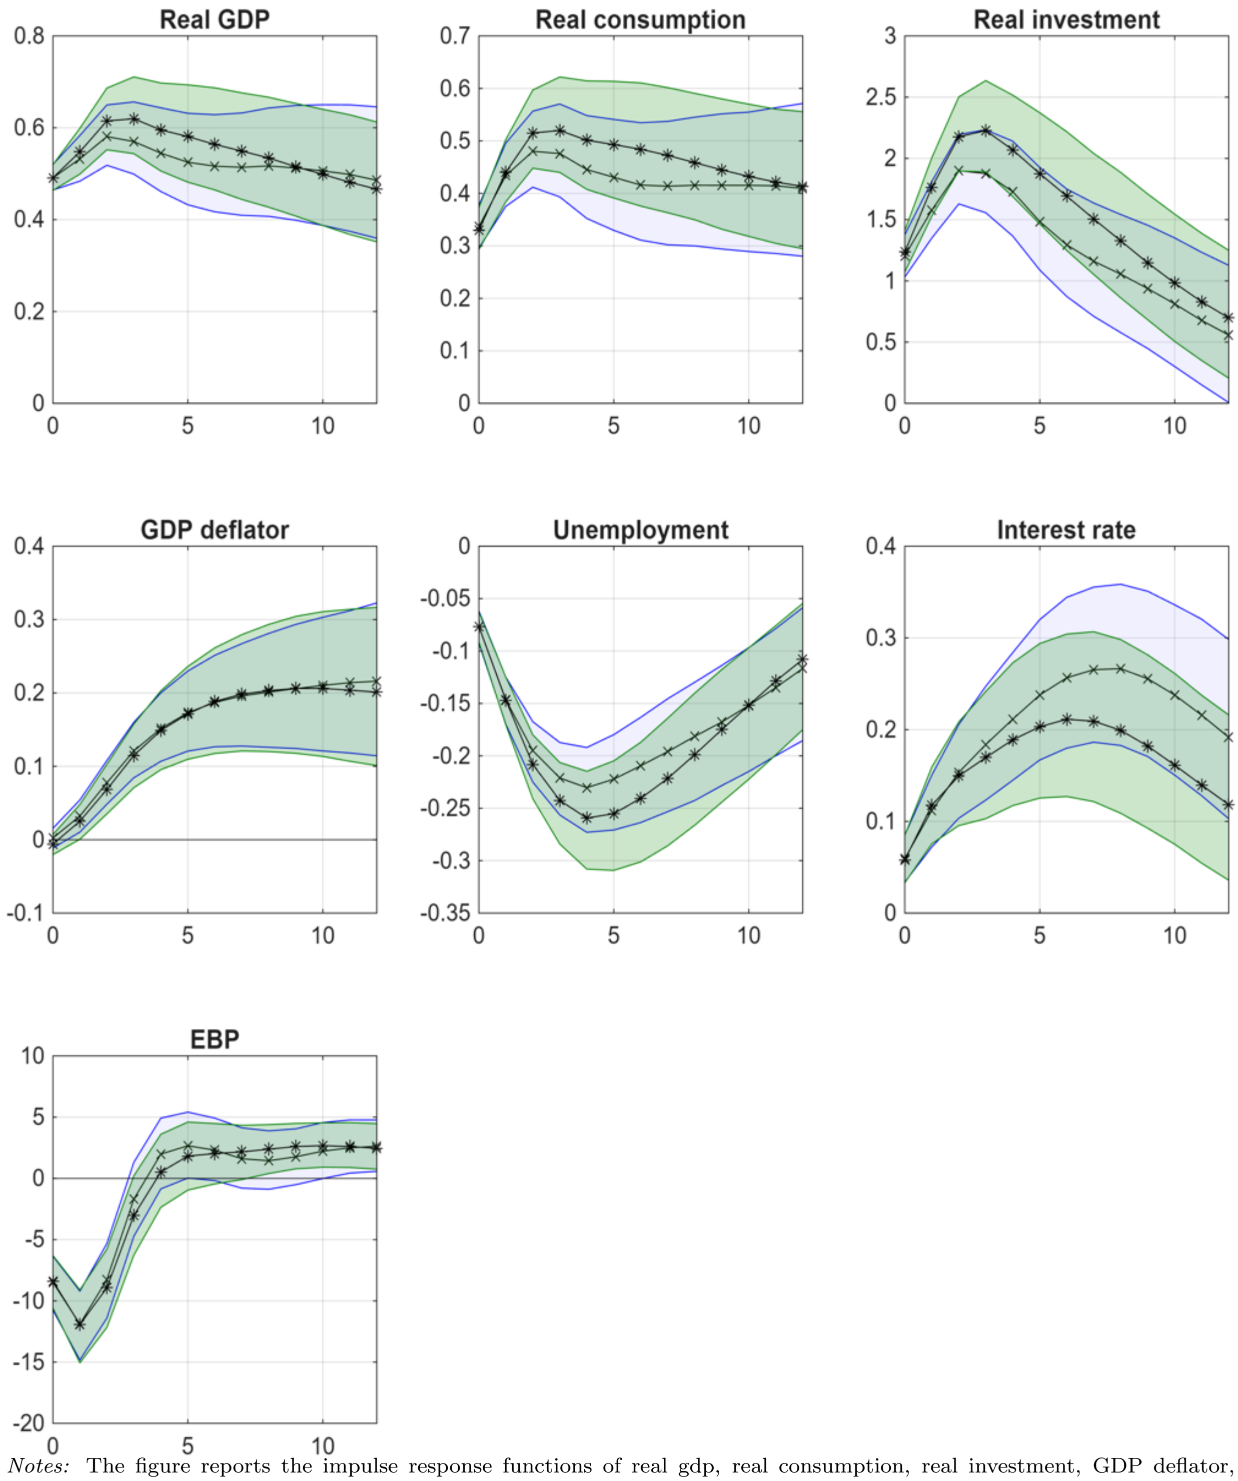

Notes: The figure reports the impulse response functions of real gdp, real consumption, real investment, GDP deflator, unemployment rate, interest rate and EBP to a business cycle shock. **Green Area**: Melitz VAR; **Blue Area**: FVAR, k=12; Solid line with stars: Melitz VAR median; Solid line with crosses: FVAR median. Source: U.S. Bureau of Economic Analysis, U.S. Bureau of Labor Statistics, Board of Governors of the Federal Reserve System, Favara et al. (2016). Sample: 1986Q1–2025Q3.

3.4 Robustness.

We conduct a broad set of robustness exercises to assess whether our baseline results depend on particular episodes, threshold choices, or the institutional environment of the U.S. sample. The goal is to verify that the informational content of the cross-sectional distribution of firm revenues is not specific to a small number of extreme aggregate events or to the exact definition of financially distressed firms, and that it also extends beyond the U.S. setting. To this end, we complement the U.S. analysis with a parallel exercise for the United Kingdom, repeating the same empirical analysis using UK firm-level data constructed from Compustat Global and Compustat Securities. For the UK case we focus on the 2001–2025 period because coverage and data completeness for UK listed firms are substantially more balanced from 2000 onward, in line with recent work using Compustat for UK firms over the 2000s (see for example Indraccolo, 2025), while the broader literature on Compustat Global also stresses thinner coverage in earlier years.[6]

To ensure that our findings are not driven by extreme aggregate events, we first conduct two additional robustness exercises on the U.S. sample. We re-estimate all specifications excluding the COVID-19 period (see Figures 6 and 7 in Appendix A.3). We then re-estimate the models on the baseline sample excluding the years of the 2008–2009 Great Financial Crisis (see Figures 8 and 9 in Appendix A.3). These episodes represent two of the largest aggregate shocks in the sample and could disproportionately affect both the cross-sectional distribution of firms and aggregate dynamics.

In both exercises, the results remain qualitatively unchanged. The impulse response functions generated by the Melitz-augmented VAR continue to closely overlap with those of the FVAR across horizons, whereas the standard aggregate VAR fails to replicate the richer dynamics identified by the heterogeneous specification. This evidence indicates that our results are not driven by these major crisis episodes, but instead reflect systematic information contained in the cross-sectional distribution of firm revenues.

We also assess robustness to the definition of the marginal firm. Appendix A.4 shows that the results are robust to alternative thresholds based on the firm expected default probability measure.

Finally, we replicate the full baseline analysis using UK data. The UK evidence from the IRFs comparing the Melitz-Bvar with the FVAR and the BVAR reported from figures 14 and 15 in the Appendix B, as well as the forecast analysis presented in Tables 3 and 4, confirms the main message of the paper: incorporating information from the cross-

[6] See for example Strauss and Yang (2019).

sectional distribution of firms substantially improves the ability of the Melitz-augmented VAR to reproduce the dynamics captured by the heterogeneous specification, while the standard aggregate VAR remains less informative. Hence, our results are not specific to the U.S. sample, but remain robust in an alternative country setting.

4 Conclusion

This paper asks how much firm heterogeneity is needed to capture the transmission of aggregate shocks. Using U.S. quarterly data from 1986Q1 to 2025Q3,
We show that the cross-sectional distribution of firm revenues contains information not spanned by aggregate variables. Augmenting a Bayesian VAR with the full firm-level distribution alters the estimated transmission of business-cycle shocks, indicating that firm heterogeneity represents omitted information in aggregate-only models.

Our central finding is that the full distribution is not necessary to capture these effects. Guided by a Melitz-type framework with endogenous selection, we identify two revenue-based sufficient statistics—the average firm and the marginal near-default firm—that summarize the relevant cross-sectional information. A parsimonious VAR augmented with these statistics closely replicates the impulse responses of a high-dimensional functional VAR and delivers comparable, and in some cases improved, short-horizon forecasting performance.

These results suggest that firm heterogeneity matters for aggregate dynamics, particularly through firms close to default, whose behavior reflects selection and default margins. At the same time, the relevant information can be captured using a low-dimensional and economically interpretable representation. Robustness exercises excluding the Covid-19 period and the Global Financial Crisis confirm that these findings are not driven by extreme events.

Taken together, our findings show that the informational content of firm heterogeneity for aggregate transmission can be captured without modeling the entire cross-section. A small number of economically motivated statistics are sufficient to summarize the relevant distributional information, and the Melitz-type framework provides a transparent guide for identifying these objects.

References

Angeletos, George-Marios, Fabrice Collard, and Harris Dellas (2020) "Business-Cycle Anatomy," *American Economic Review*, 110 (10), 3030–3070, 10.1257/aer.20181174.

Arellano, Cristina, Yan Bai, and Patrick J Kehoe (2019) "Financial frictions and fluctuations in volatility," *Journal of Political Economy*, 127 (5), 2049–2103.

Ascari, Guido, Stefano Fasani, Jakob Grazzini, and Lorenza Rossi (2023) "Endogenous uncertainty and the macroeconomic impact of shocks to inflation expectations," *Journal of Monetary Economics*, 140, S48–S63.

Bayer, Christian, Benjamin Born, and Ralph Luetticke (2024) "Shocks, frictions, and inequality in US business cycles," *American Economic Review*, 114 (5), 1211–1247.

Bharath, Sreedhar T. and Tyler Shumway (2008) "Forecasting Default with the Merton Distance to Default Model," *The Review of Financial Studies*, 21 (3), 1339–1369.

Bilbiie, Florin O, Fabio Ghironi, and Marc J Melitz (2012) "Endogenous entry, product variety, and business cycles," *Journal of Political Economy*, 120 (2), 304–345.

Bilbiie, Florin O, Giorgio Primiceri, and Andrea Tambalotti (2023) "Inequality and business cycles,"Technical report, National Bureau of Economic Research.

Chang, Minsu, Xiaohong Chen, and Frank Schorfheide (2024) "Heterogeneity and aggregate fluctuations," *Journal of Political Economy*, 132 (12), 4021–4067.

Chang, Minsu and Frank Schorfheide (2022) "On the Effects of Monetary Policy Shocks on Earnings and Consumption Heterogeneity."

Clementi, Gian Luca and Berardino Palazzo (2016) "Entry, exit, firm dynamics, and aggregate fluctuations," *American Economic Journal: Macroeconomics*, 8 (3), 1–41.

Cloyne, James, Clodomiro Ferreira, and Paolo Surico (2020) "Monetary policy when households have debt: new evidence on the transmission mechanism," *The Review of Economic Studies*, 87 (1), 102–129.

Cooley, Thomas and Vincenzo Quadrini (2006) "Monetary policy and the financial decisions of firms," *Economic Theory*, 27 (1), 243–270.

Debortoli, Davide and Jordi Galí (2022) "Idiosyncratic Income Risk and Aggregate Fluctuations," *National Bureau of Economic Research, Working Paper Series* (29704).

Fasani, Stefano, Haroon Mumtaz, and Lorenza Rossi (2023) "Monetary policy uncertainty and firm dynamics," *Review of Economic Dynamics*, 47, 278–296.

Favara, Giovanni, Simon Gilchrist, Kurt F. Lewis, and Egon Zakrajšek (2016) "Updating the Recession Risk and the Excess Bond Premium," FEDS Notes 2016-10-06, Board of Governors of the Federal Reserve System, 10.17016/2380-7172.1836.

Gertler, Mark and Simon Gilchrist (1994a) "Monetary policy, business cycles, and the behavior of small manufacturing firms," *The quarterly journal of economics*, 109 (2), 309–340.

——— (1994b) "Monetary Policy, Business Cycles, and the Behavior of Small Manufacturing Firms," *The Quarterly Journal of Economics*, 109 (2), 309–340.

Ghironi, Fabio and Marc J Melits (2005) "International trade and macroeconomic dynamics with heterogeneous firms," *The Quarterly Journal of Economics*, 120 (3), 865–915.

Ghironi, Fabio and Marc J Melitz (2007) "Trade flow dynamics with heterogeneous firms," *American Economic Review*, 97 (2), 356–361.

Giannone, Domenico, Michele Lenza, and Giorgio E Primiceri (2015) "Prior selection for vector autoregressions," *Review of Economics and Statistics*, 97 (2), 436–451.

Gilchrist, S, J. W. Sim, and E. Zakrajšek (2014) "Uncertainty, financial frictions, and investment dynamics," *Working Paper*.

Hopenhayn, Hugo A (1992) "Entry, exit, and firm dynamics in long run equilibrium," *Econometrica: Journal of the Econometric Society*, 1127–1150.

Huber, Florian, Massimiliano Marcellino, and Tommaso Tornese (2024) "The Distributional Effects of Economic Uncertainty," *arXiv preprint arXiv:2411.12655*.

Indraccolo, Leonardo (2025) *Bridging the Gap: Understanding the UK-US Productivity Decoupling*: International Monetary Fund.

Khan, Aubhik and Julia K. Thomas (2008) "Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics," *Econometrica*, 76 (2), 395–436.

Khan, Aubhik and Julia K Thomas (2013) "Credit shocks and aggregate fluctuations in an economy with production heterogeneity," *Journal of Political Economy*, 121 (6), 1055–1107.

Krueger, Dirk, Kurt Mitman, and Fabrizio Perri (2016) "Macroeconomics and household heterogeneity," in *Handbook of macroeconomics*, 2, 843–921: Elsevier.

Lenza, Michele and Giorgio E Primiceri (2022) "How to estimate a vector autoregression after March 2020," *Journal of Applied Econometrics*, 37 (4), 688–699.

Lenza, Michele and Ettore Savoia (2024) "Do we need firm data to understand macroeconomic dynamics?"Technical report, Sveriges Riksbank Working Paper Series.

Lenza, Michele and Oreste Tristani (2025) "Heterogeneity in Macroeconomics."

Marcellino, Massimiliano, Andrea Renzetti, and Tommaso Tornese (2024) "Firm Heterogeneity and Macroeconomic Fluctuations: a Functional VAR model," *arXiv preprint arXiv:2411.05695*.

Melitz, Marc J (2003) "The impact of trade on intra-industry reallocations and aggregate industry productivity," *econometrica*, 71 (6), 1695–1725.

Merton, Robert C (1974) "On the pricing of corporate debt: The risk structure of interest rates," *The Journal of finance*, 29 (2), 449–470.

Ottonello, Pablo and Thomas Winberry (2020) "Financial Heterogeneity and the Investment Channel of Monetary Policy," *Econometrica*, 88 (6), 2473–2502.

Rossi, Lorenza (2019) "The overshooting of firms' destruction, banks and productivity shocks," *European Economic Review*, 113, 136–155.

Strauss, Ilan and Jangho Yang (2019) "The Global Investment Slowdown: Corporate Secular Stagnation and the Draining of the Cash Flow Swamp," *Institute for New Economic Thinking. INET Oxford Working Paper No*, 16, 1–69.

A Appendix

A.1 Details on Distance-to-Default Computation

In order to identify the marginal firms, we adopt the Merton (1974) model to compute each firm distance to default and expected default frequency, following the procedure proposed by Bharath and Shumway (2008). First, we merge balance sheet data for US listed firms from Compustat with CRSP stock market data, by using the GVKEY as an unique identifier. Then, we use this joint information to compute the Distance to Default. We construct the market value of equity as:

$$E_{it} = |P_{it}| \times \mathrm{N}_{it}, \tag{17}$$

where P_{it} denotes the end-of-quarter stock price and N_{it} the number of shares outstanding. The face value of debt is defined as

$$F_{it} = \mathrm{DLC}_{it} + 0.5 \times \mathrm{DLTT}_{it}, \tag{18}$$

where DLC_{it} is debt in current liabilities and DLTT_{it} is long-term debt. The total firm value is initialized as

$$A_{it}^{(0)} = E_{it} + F_{it}. \tag{19}$$

Under the Merton (1974) framework, equity is modeled as a call option on firm assets with strike price equal to the face value of debt. Let A_{it} denote the market value of assets and $\sigma_{A,it}$ its volatility. The Black–Scholes relationship implies

$$E_{it} = A_{it}\Phi(d_{1,it}) - F_{it}e^{-r_{f,t}T}\Phi(d_{2,it}), \tag{20}$$

where

$$d_{1,it} = \frac{\ln(A_{it}/F_{it}) + \left(r_{f,t} + \frac{1}{2}\sigma_{A,it}^2\right)T}{\sigma_{A,it}\sqrt{T}}, \tag{21}$$

$$d_{2,it} = d_{1,it} - \sigma_{A,it}\sqrt{T}, \tag{22}$$

$r_{f,t}$ is the risk-free rate, T is the time horizon (set to one year), and $\Phi(\cdot)$ denotes the cumulative distribution function of the standard normal distribution.

Because A_{it} and $\sigma_{A,it}$ are not directly observable, we estimate them iteratively. Asset volatility is initialized using the annualized standard deviation of equity returns. We then recover A_{it} from the option-pricing equation and update $\sigma_{A,it}$ exploiting the relationship

$$\sigma_{E,it} = \sigma_{A,it} \frac{A_{it}}{E_{it}} \Phi(d_{1,it}), \tag{23}$$

where $\sigma_{E,it}$ is the observed equity volatility. The procedure is repeated until convergence.

Given the estimated asset value and volatility, the Distance to Default is computed as

$$DD_{it} = \frac{\ln(A_{it}/F_{it}) + \left(\mu_{it} - \frac{1}{2}\sigma^2_{A,it}\right) T}{\sigma_{A,it}\sqrt{T}}, \tag{24}$$

where μ_{it} is the expected asset return.

In the implementation, we estimate it as the annualized mean of daily log asset returns, computed over the previous 250 trading days. Finally, the Expected Default Frequency is defined as

$$EDF_{it} = \Phi(-DD_{it}). \tag{25}$$

Firms with an EDF_{it} at or above a sufficiently high threshold are classified as marginal firms, as they operate closer to the default boundary and are therefore more financially vulnerable. We set the threshold equal to 12% which is the same value calibrated in the Melitz-type models (see for example Bilbiie et al., 2012, Fasani et al., 2023, among many others).

A.2 Model validation and selection

Out of Sample Forecast The differences in the estimated impulse response functions across models may reflect the inclusion of measurement noise in our FVAR specification or differences in model complexity. In addition, since the functional model includes between 10 and 19 variables, the resulting degree of shrinkage—relative to the smaller aggregate specification—may also help explain the discrepancies in the impulse responses. In order to address these concerns, we conduct a simple forecasting exercise. We split the sample into a training set, spanning from 1986Q1 to 1999Q4, and a test set covering the period from 2001Q1 to 2025Q3. We then generate one-step-ahead out-of-sample forecasts using an expanding window scheme for each model over the test period. For each functional VAR specification, we compute the relative mean squared error with respect to the aggregate VAR. Finally, we repeat the same exercise for the Melitz VAR. This procedure allows us to assess whether the functional model provides a better approximation of the true data-generating process and to select the preferred functional specification based on its out-of-sample forecasting performance relative to the aggregate model. In addition, it enables us to evaluate the performance of the Melitz VAR relative to the specification that incorporates the full distribution.

The results are reported in Table 1. The FVAR models generally perform similarly to, or better than, the BVAR. Among them, the specification with $K = 12$ delivers the best performance in terms of relative mean squared error.

In comparison, the Melitz VAR delivers forecasting performance broadly similar to the BVAR but remains slightly weaker than the preferred FVAR specification.

We also compare the model's performance using alternative statistics to summarize the cross-sectional distribution of firm revenues, including the variance, the interquantile range (95th–5th percentile), and the 95th percentile, which captures large firms. In all cases the Melitz-VAR, based on the marginal-firm threshold outperforms these alternatives in terms of forecasting accuracy.

In Table 2 we compare the performance of the three model by repeating the forecasting exercise in rolling window, with a window width of 60 quarters, obtaining comparable results.[7]

	Period	**GDP**	**Cons.**	**Inv.**	**Infl.**	**Unemp.**	**Int.Rate**	**EBP**	**Average**
K=4	2001–2025	0.885	0.896	0.831	0.940	0.863	0.978	1.083	0.925
K=6	2001–2025	0.908	0.935	0.815	1.005	0.891	1.109	1.145	0.973
K=8	2001–2025	0.900	0.953	0.858	0.973	0.967	1.068	1.086	0.972
K=10	2001–2025	0.921	1.030	0.807	0.975	0.842	1.067	1.079	0.960
K=12	2001–2025	0.935	1.004	0.816	0.982	0.834	1.062	1.071	0.957
Melitz	2001–2025	0.928	0.958	0.935	0.972	0.927	1.009	1.076	0.972
Melitz*	2001–2025	0.981	0.985	0.985	1.004	0.996	1.004	1.034	0.998
SD	2001–2025	0.965	0.975	0.971	1.007	0.991	0.992	1.024	0.989
IQR	2001–2025	0.986	0.987	1.014	1.033	0.984	1.068	1.024	1.014
Q95	2001–2025	0.966	0.965	0.991	1.058	0.979	1.041	1.040	1.006

Table 1: **Expanding Windows Forecasts.** Relative MSE. Values smaller than one indicate that the FVAR model is more accurate than the VAR. Melitz is the VAR with average revenues and the average revenues of firms close to default. Melitz* is the VAR with average revenues only.

[7] In this latter case the best performance in terms of rmse is provided by the $K = 10$ specification.

	Period	**GDP**	**Cons.**	**Inv.**	**Infl.**	**Unemp.**	**Int.Rate**	**EBP**	**Average**
K=4	2001–2025	0.867	0.837	0.962	1.022	0.701	1.169	1.244	0.972
K=6	2001–2025	0.832	0.765	1.023	1.121	0.642	1.278	1.335	1.000
K=8	2001–2025	0.827	0.754	1.134	0.949	0.525	1.256	1.188	0.948
K=10	2001–2025	0.777	0.747	1.084	0.898	0.475	1.254	1.240	0.925
K=12	2001–2025	0.824	0.763	1.124	0.906	0.517	1.308	1.167	0.944
Melitz	2001–2025	0.898	0.933	0.976	0.964	0.834	1.050	1.144	0.971
Melitz*	2001–2025	0.957	0.976	0.940	1.011	0.934	1.101	1.044	0.995

Table 2: **Rolling Windows Forecasts.** Relative MSE. Values smaller than one indicate that the FVAR model is more accurate than the VAR. Melitz denotes the VAR augmented with sufficient statistics (average revenues and average revenues of firms close to default).

A.3 Additional Robustness Checks

A.3.1 Excluding the Covid Period

Figure 6: Impulse response functions to a business cycle shock

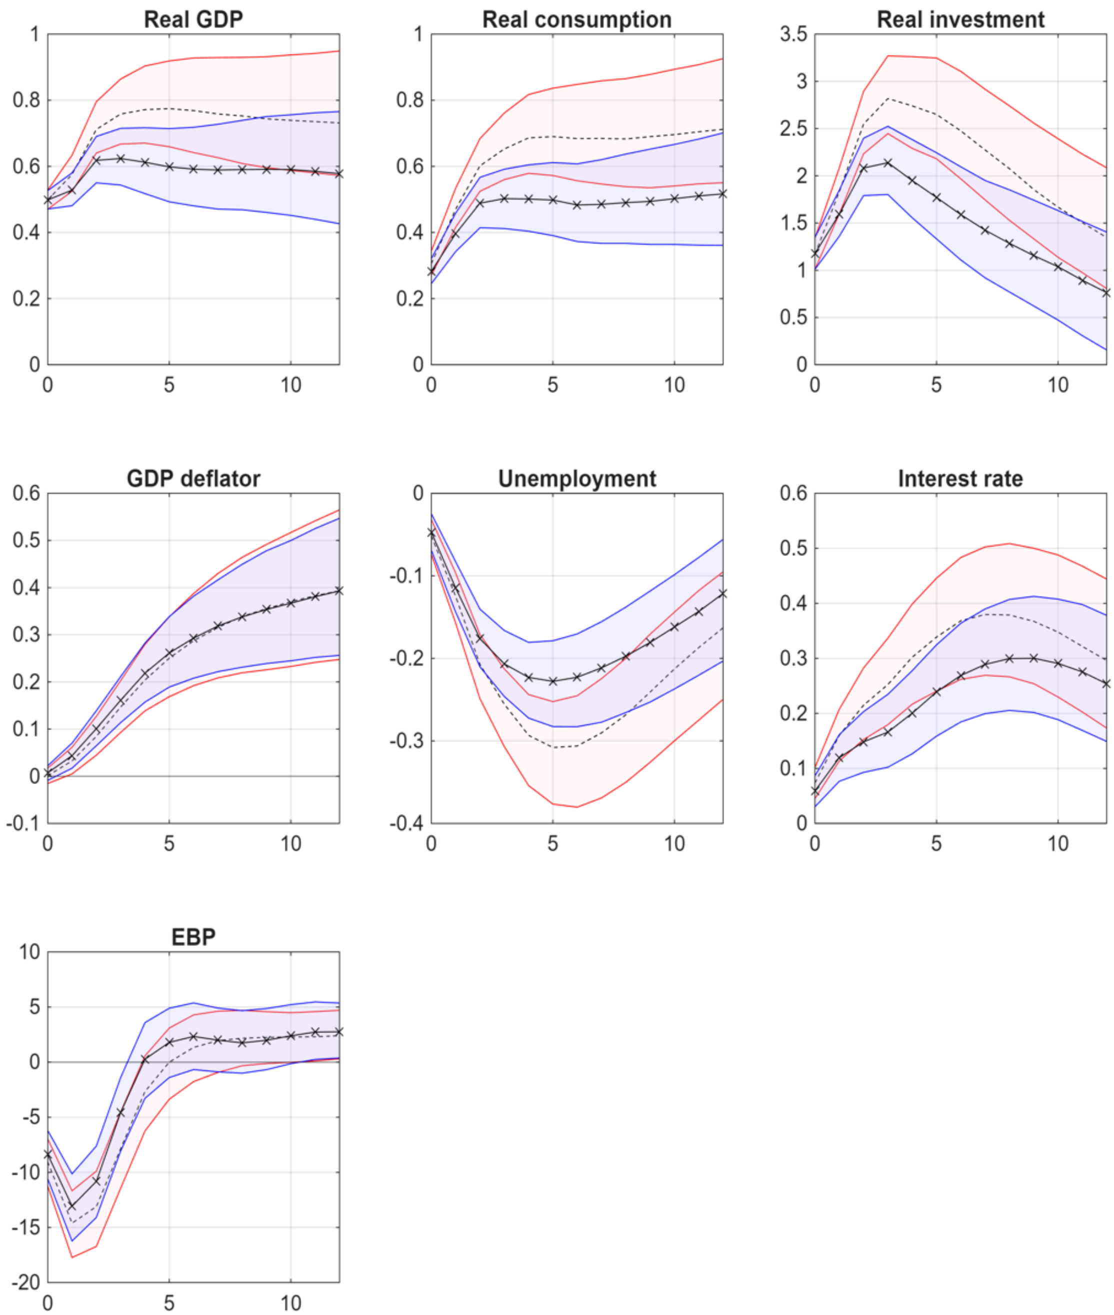

Notes: The figure reports the impulse response functions of real gdp, real consumption, real investment, GDP deflator, unemployment rate, interest rate and EBP to a business cycle shock, with the exclusion of year 2020 from the sample. **Blue Area**: FVAR, k=12; **Red Area**: aggregate BVAR; Dashed line: BVAR median; Solid line with crosses: FVAR median. Source: U.S. Bureau of Economic Analysis, U.S. Bureau of Labor Statistics, Board of Governors of the Federal Reserve System, Favara et al. (2016). Sample: 1986Q1–2025Q3.

Figure 7: Impulse response functions to a business cycle shock

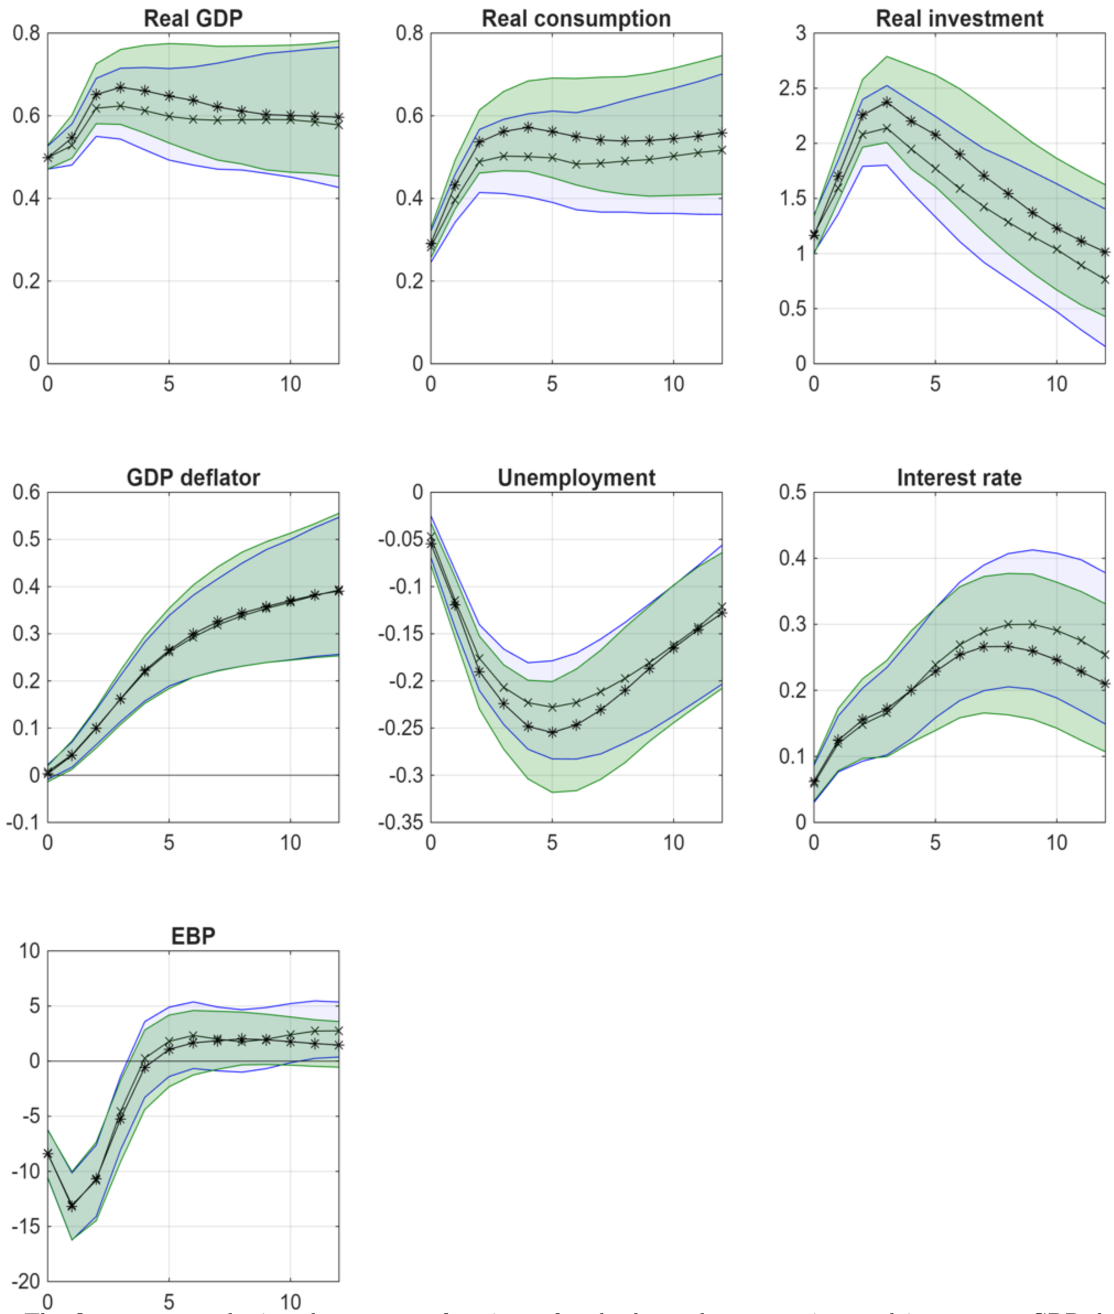

Notes: The figure reports the impulse response functions of real gdp, real consumption, real investment, GDP deflator, unemployment rate, interest rate and EBP to a business cycle shock, with the exclusion of year 2020 from the sample. **Green Area**: Melitz VAR; **Blue Area**: FVAR, k=12; Solid line with stars: Melitz VAR median; Solid line with crosses: FVAR median. Source: U.S. Bureau of Economic Analysis, U.S. Bureau of Labor Statistics, Board of Governors of the Federal Reserve System, Favara et al. (2016). Sample: 1986Q1–2025Q3.

A.3.2 Excluding the Global Financial Crisis

Figure 8: Impulse response functions to a business cycle shock

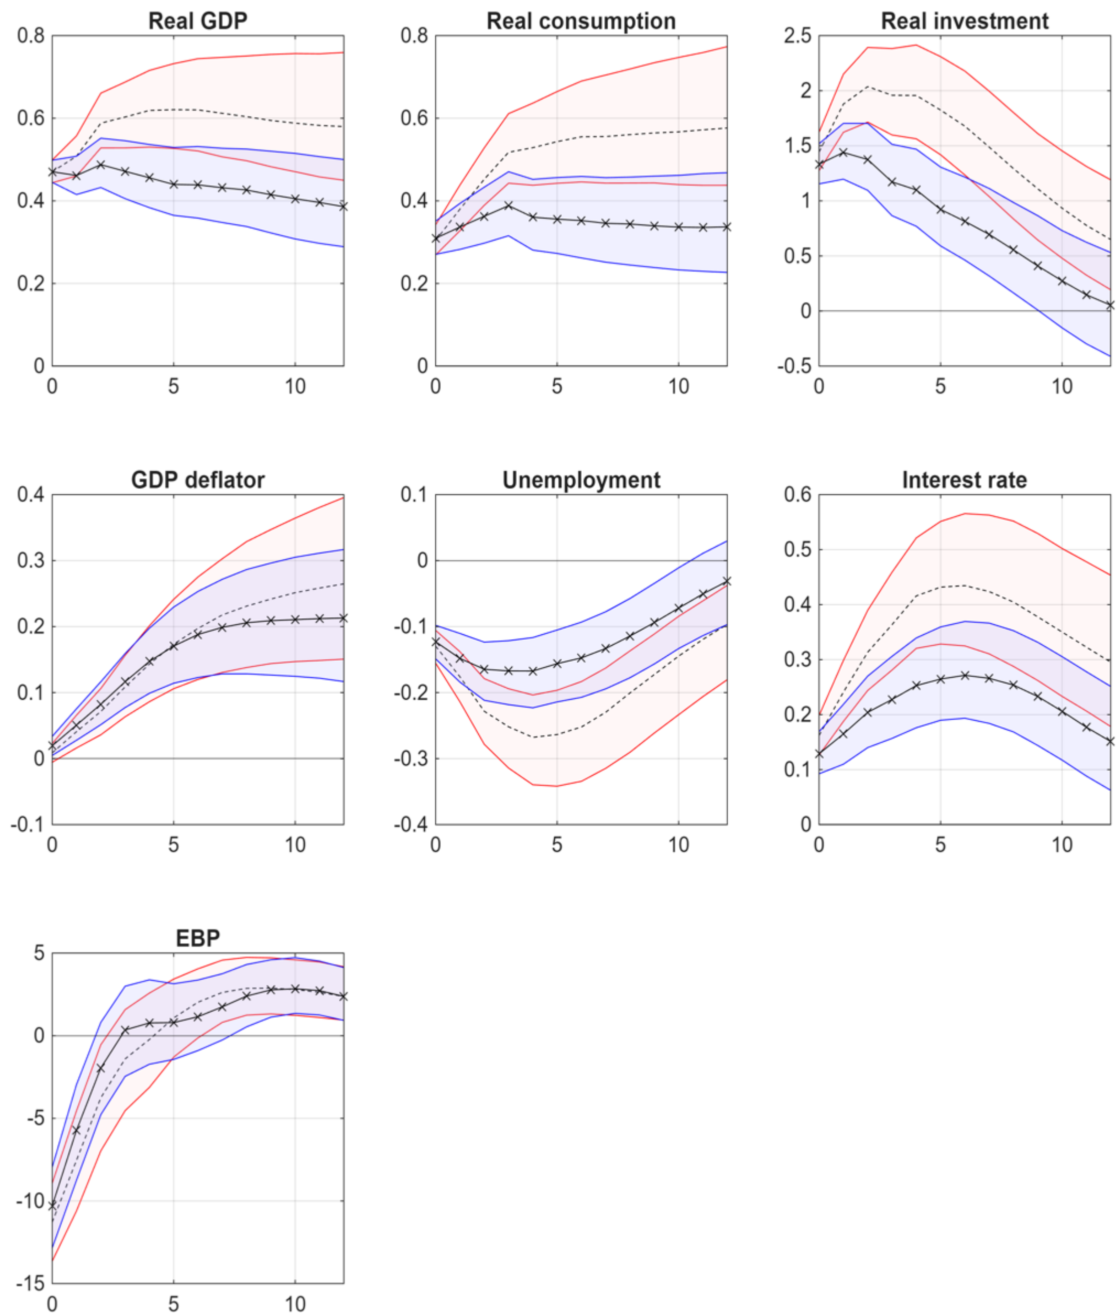

Notes: The figure reports the impulse response functions of real gdp, real consumption, real investment, GDP deflator, unemployment rate, interest rate and EBP to a business cycle shock, with the exclusion of year 2008 from the sample. **Blue Area**: FVAR, K=12; **Red Area**: aggregate BVAR; Dashed line: BVAR median; Solid line with crosses: FVAR median. Source: U.S. Bureau of Economic Analysis, U.S. Bureau of Labor Statistics, Board of Governors of the Federal Reserve System, Favara et al. (2016). Sample: 1986Q1–2025Q3.

Figure 9: Impulse response functions to a business cycle shock

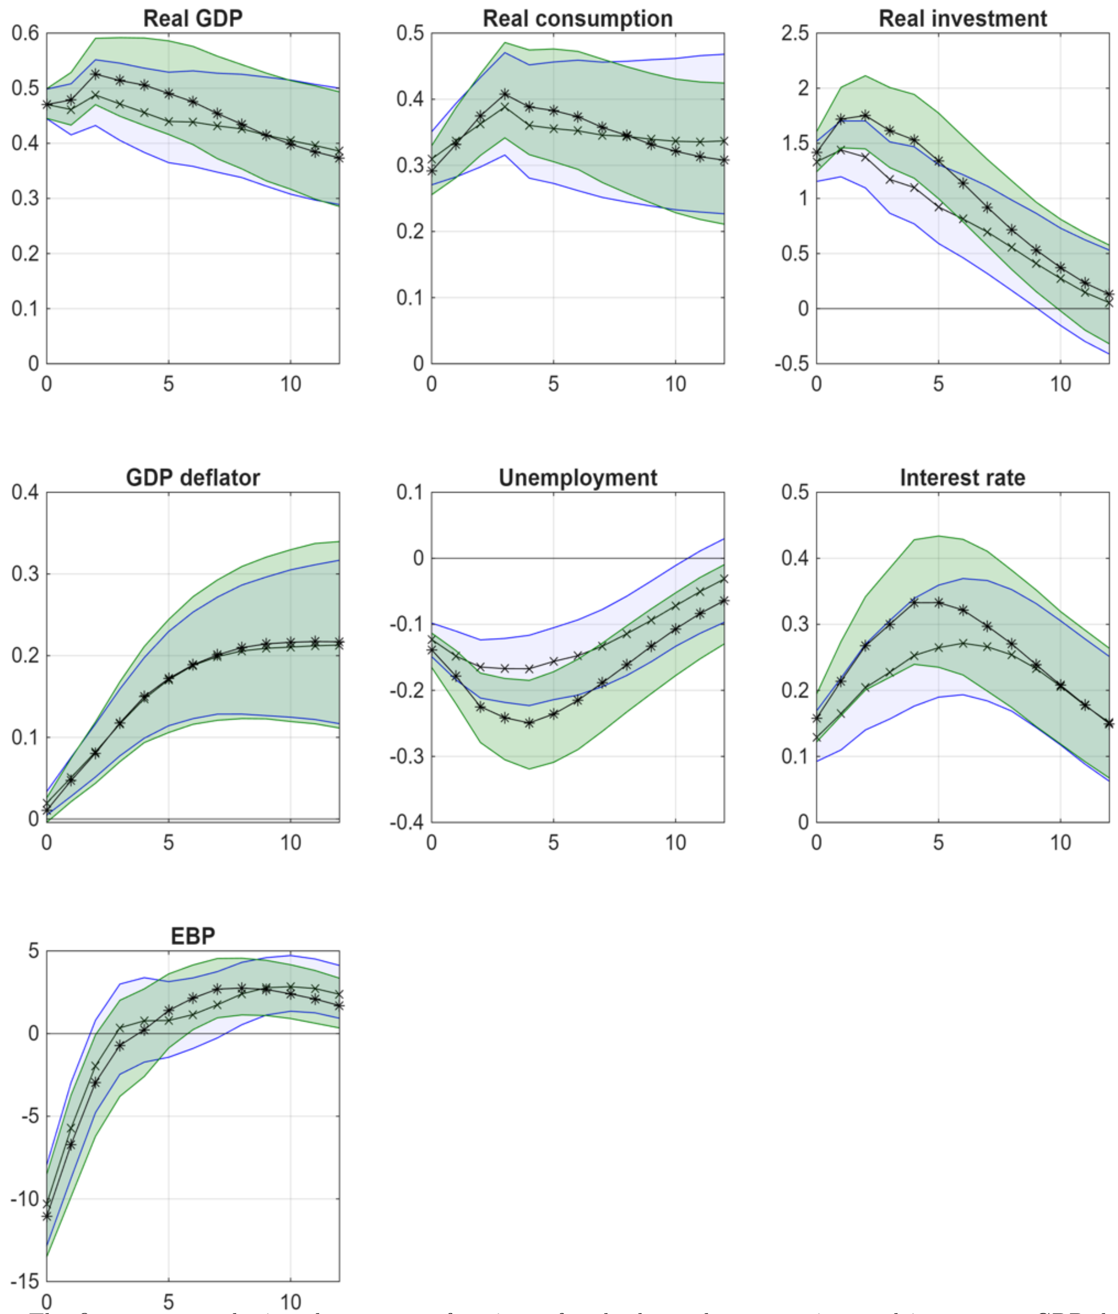

Notes: The figure reports the impulse response functions of real gdp, real consumption, real investment, GDP deflator, unemployment rate, interest rate and EBP to a business cycle shock, with the exclusion of year 2008 from the sample. **Green Area**: Melitz VAR; **Blue Area**: FVAR, K=12; Solid line with stars: Melitz VAR median; Solid line with crosses: FVAR median. Source: U.S. Bureau of Economic Analysis, U.S. Bureau of Labor Statistics, Board of Governors of the Federal Reserve System, Favara et al. (2016). Sample: 1986Q1–2025Q3.

A.4 Alternative Thresholds for the Marginal Revenues

In Figure 10, we report the share of firms in our sample classified as marginal in each quarter. We compare our baseline threshold of 12 percent, which represents our baseline value that aligns with the Melitz-type models calibration, with alternative thresholds of 9 and 15 percent. Under our baseline threshold (EDF $\geq$ 12%), the marginal firms constitute, on average, 12% of the sample annually. About 13% of firms are classified as marginal on an annual basis before 2000, while this share declines to roughly 9% after. This pattern is consistent with the documented decline in U.S. business dynamism over the past two decades.

Figure 10: Share of Firms Classified as Marginal

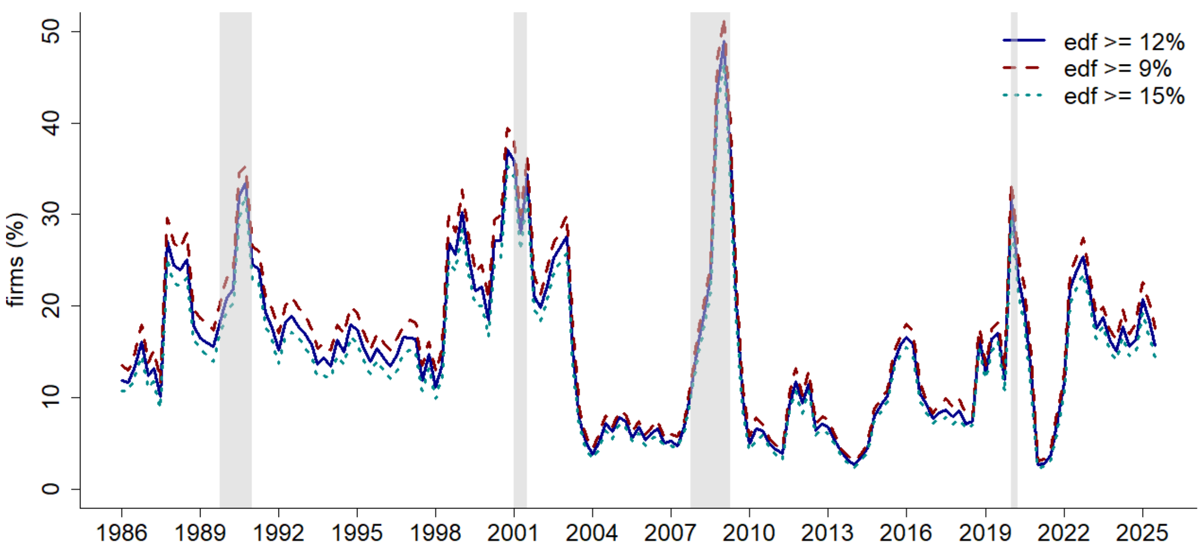

Notes: The figure reports the quarterly share of firms classified as marginal under EDF thresholds of 9%, 12%, and 15%. Source: Compustat, U.S. Bureau of Economic Analysis. Sample: 1986Q1–2025Q3.

In figure 12 and 11 we show that choosing an alternative threshold does not dramatically alter our result, and the impulse response functions deviate only marginally from our baseline. Same results apply for the out of sample forecasting properties.

Figure 11: Impulse response functions to a business cycle shock

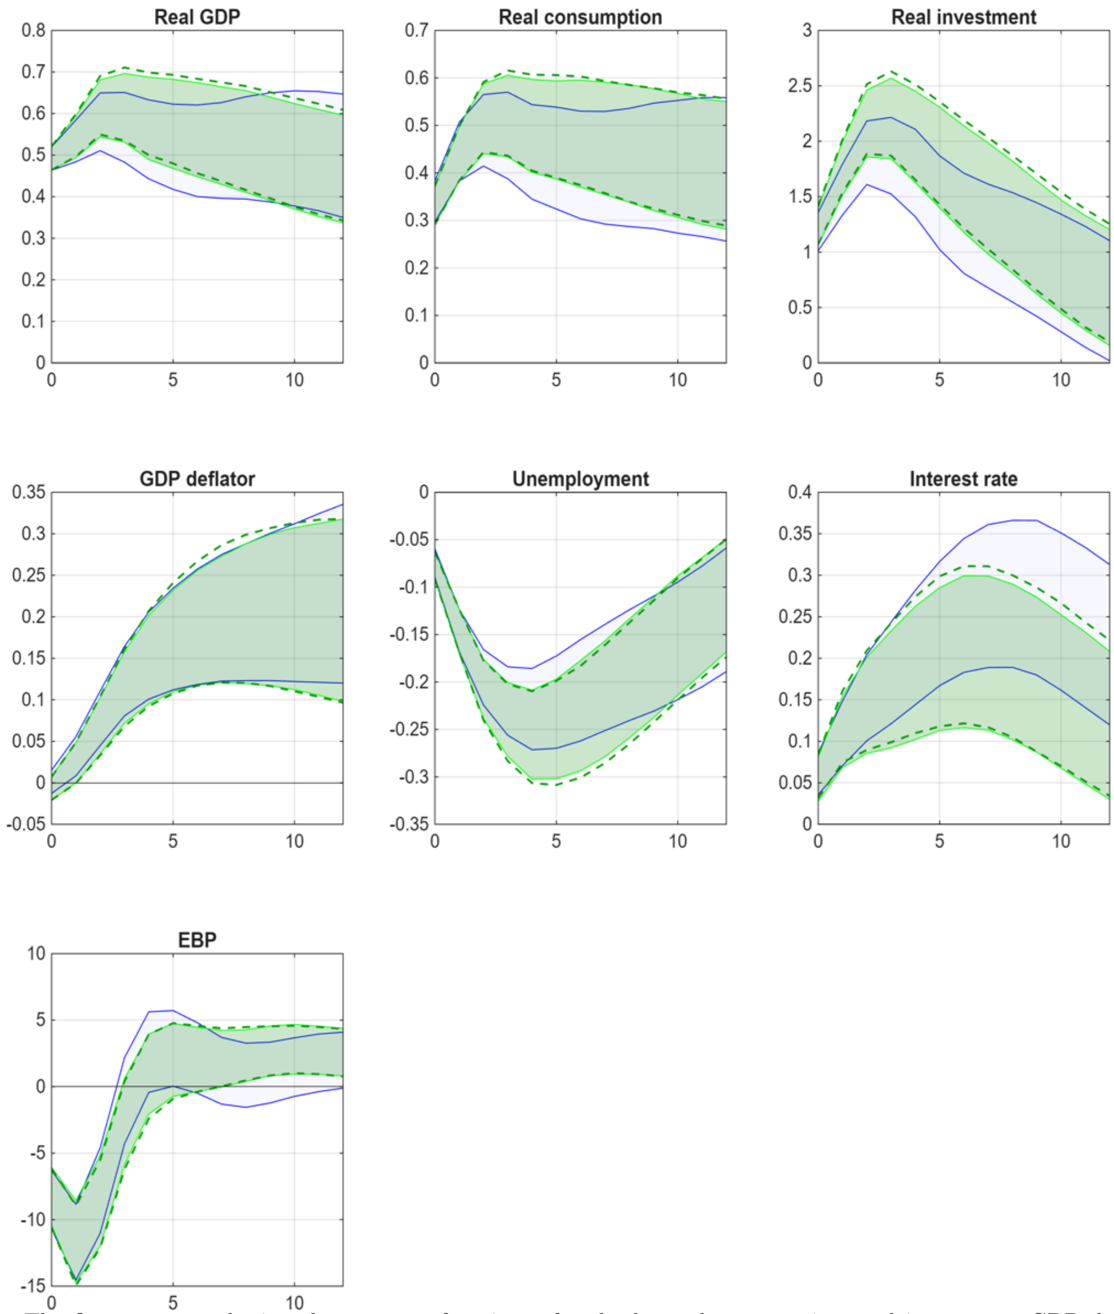

Notes: The figure reports the impulse response functions of real gdp, real consumption, real investment, GDP deflator, unemployment rate, interest rate and EBP to a business cycle shock. **Green dashed line**: Melitz VAR, with baseline 12% threshold; **Green Area**: Melitz VAR, with 9% threshold; **Blue Area**: FVAR, K=12; Source: U.S. Bureau of Economic Analysis, U.S. Bureau of Labor Statistics, Board of Governors of the Federal Reserve System, Favara et al. (2016). Sample: 1986Q1–2025Q3.

Figure 12: Impulse response functions to a business cycle shock

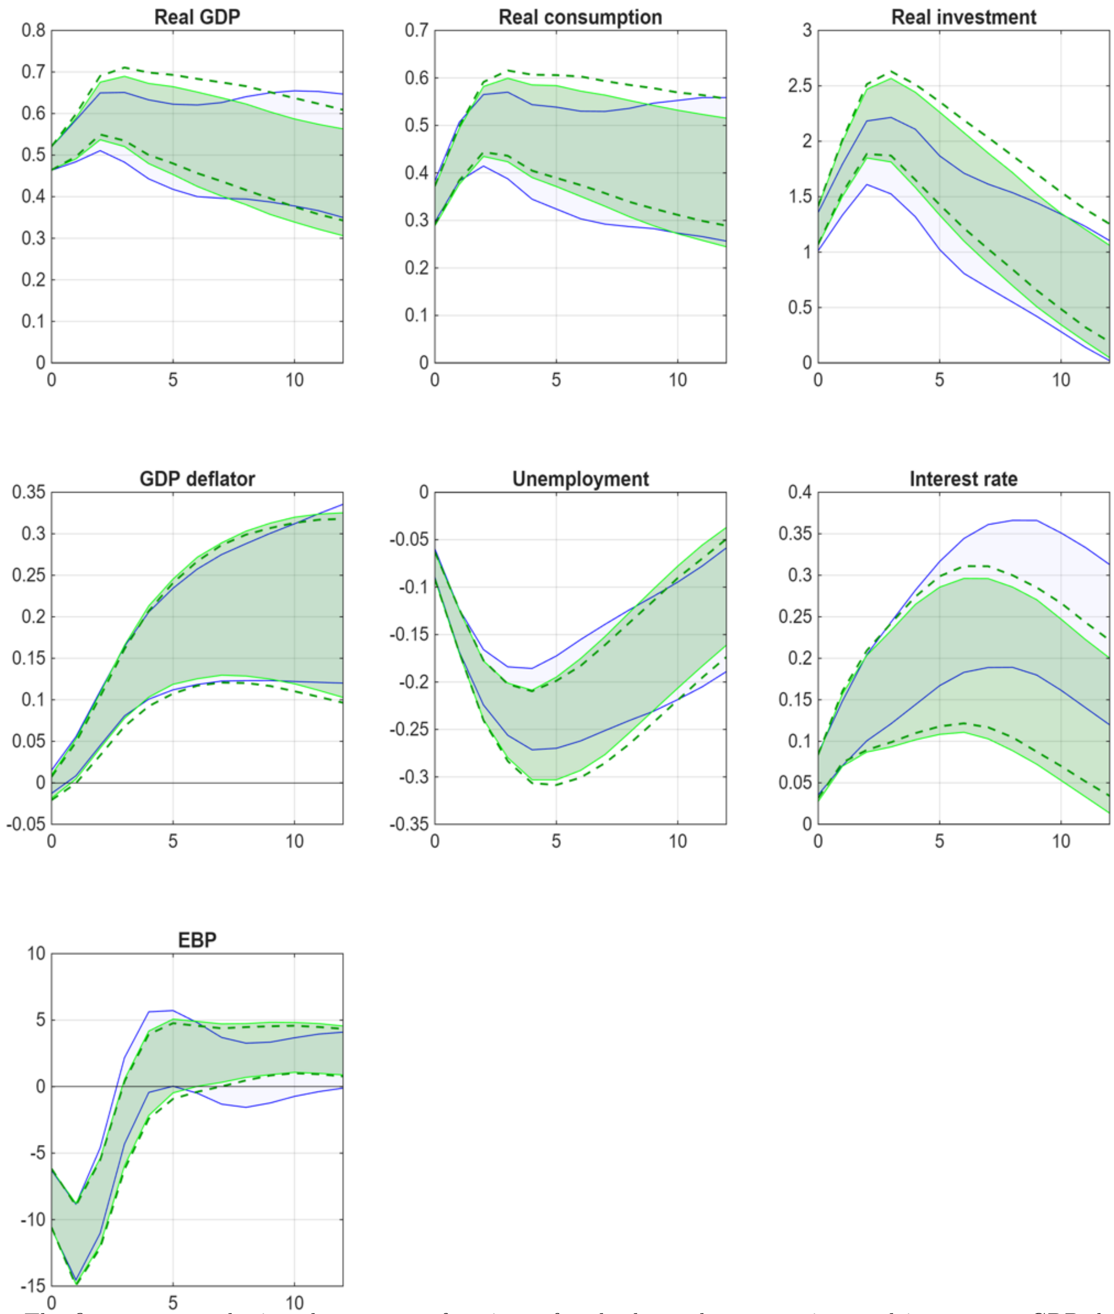

Notes: The figure reports the impulse response functions of real gdp, real consumption, real investment, GDP deflator, unemployment rate, interest rate and EBP to a business cycle shock. **Green dashed line**: Melitz VAR, with baseline 15% threshold; **Green Area**: Melitz VAR, with 24% threshold; **Blue Area**: FVAR, K=12; Source: U.S. Bureau of Economic Analysis, U.S. Bureau of Labor Statistics, Board of Governors of the Federal Reserve System, Favara et al. (2016). Sample: 1986Q1–2025Q3.

B Results UK

Figure 13: Distance-to-Default and Expected Default Probability (2001Q1–2025Q2)

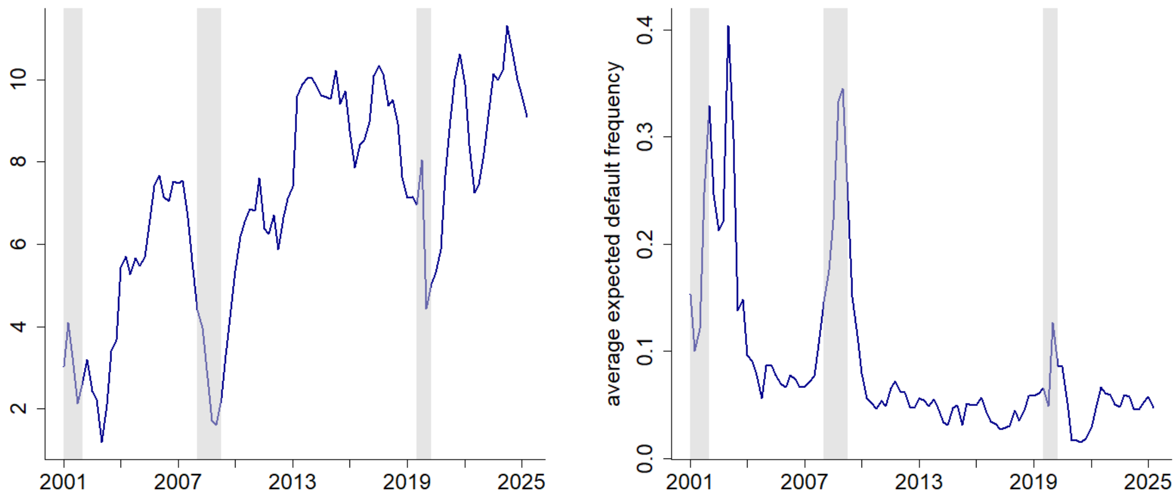

Notes: The plots report the average distance to default (right-hand-side) and expected default frequency (right-hand-side). Shaded areas represent OECD recessions. Source: Compustat, CRSP. Sample: 2001Q1–2025Q2.

Figure 14: Impulse response functions to a business cycle shock

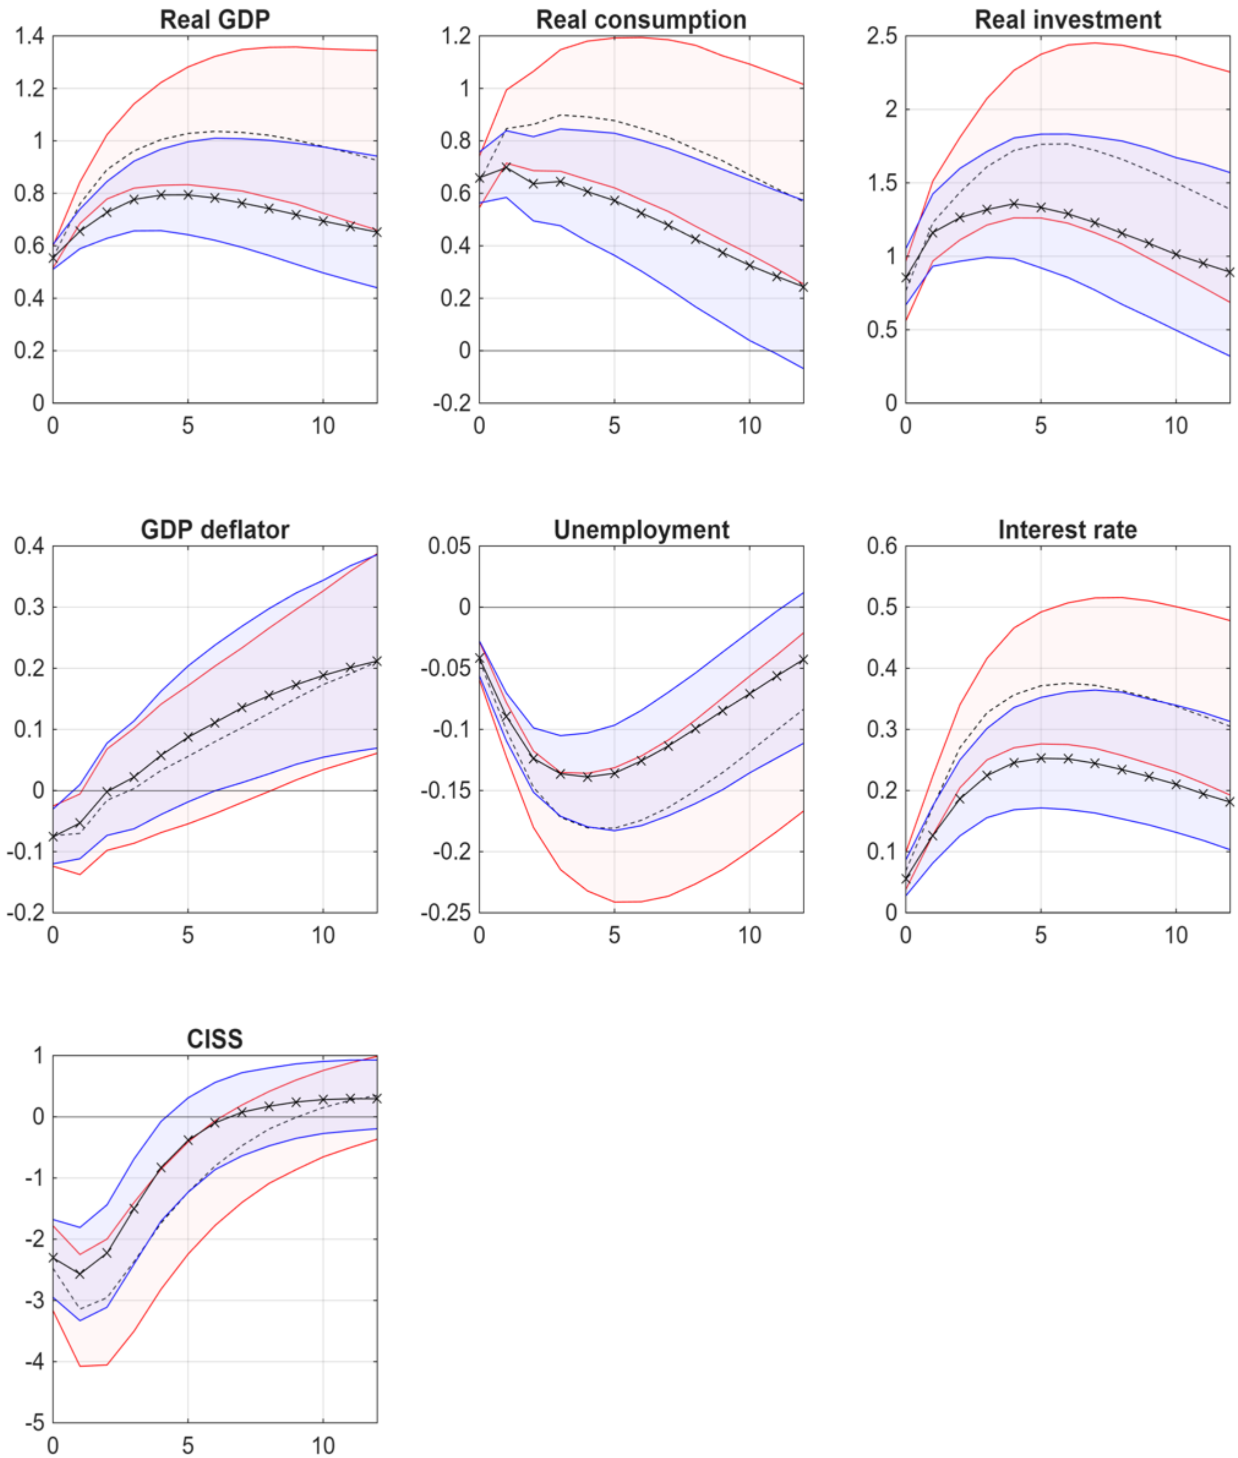

Notes: The figure reports the impulse response functions of real gdp, real consumption, real investment, GDP deflator, unemployment rate, interest rate and CISS to a business cycle shock. **Blue Area**: FVAR, K=10; **Red Area**: aggregate BVAR; Dashed line: BVAR median; Solid line with crosses: FVAR median. Source: Office for National Statistics, ECB. Sample: 2001Q1–2025Q2.

Figure 15: Impulse response functions to a business cycle shock

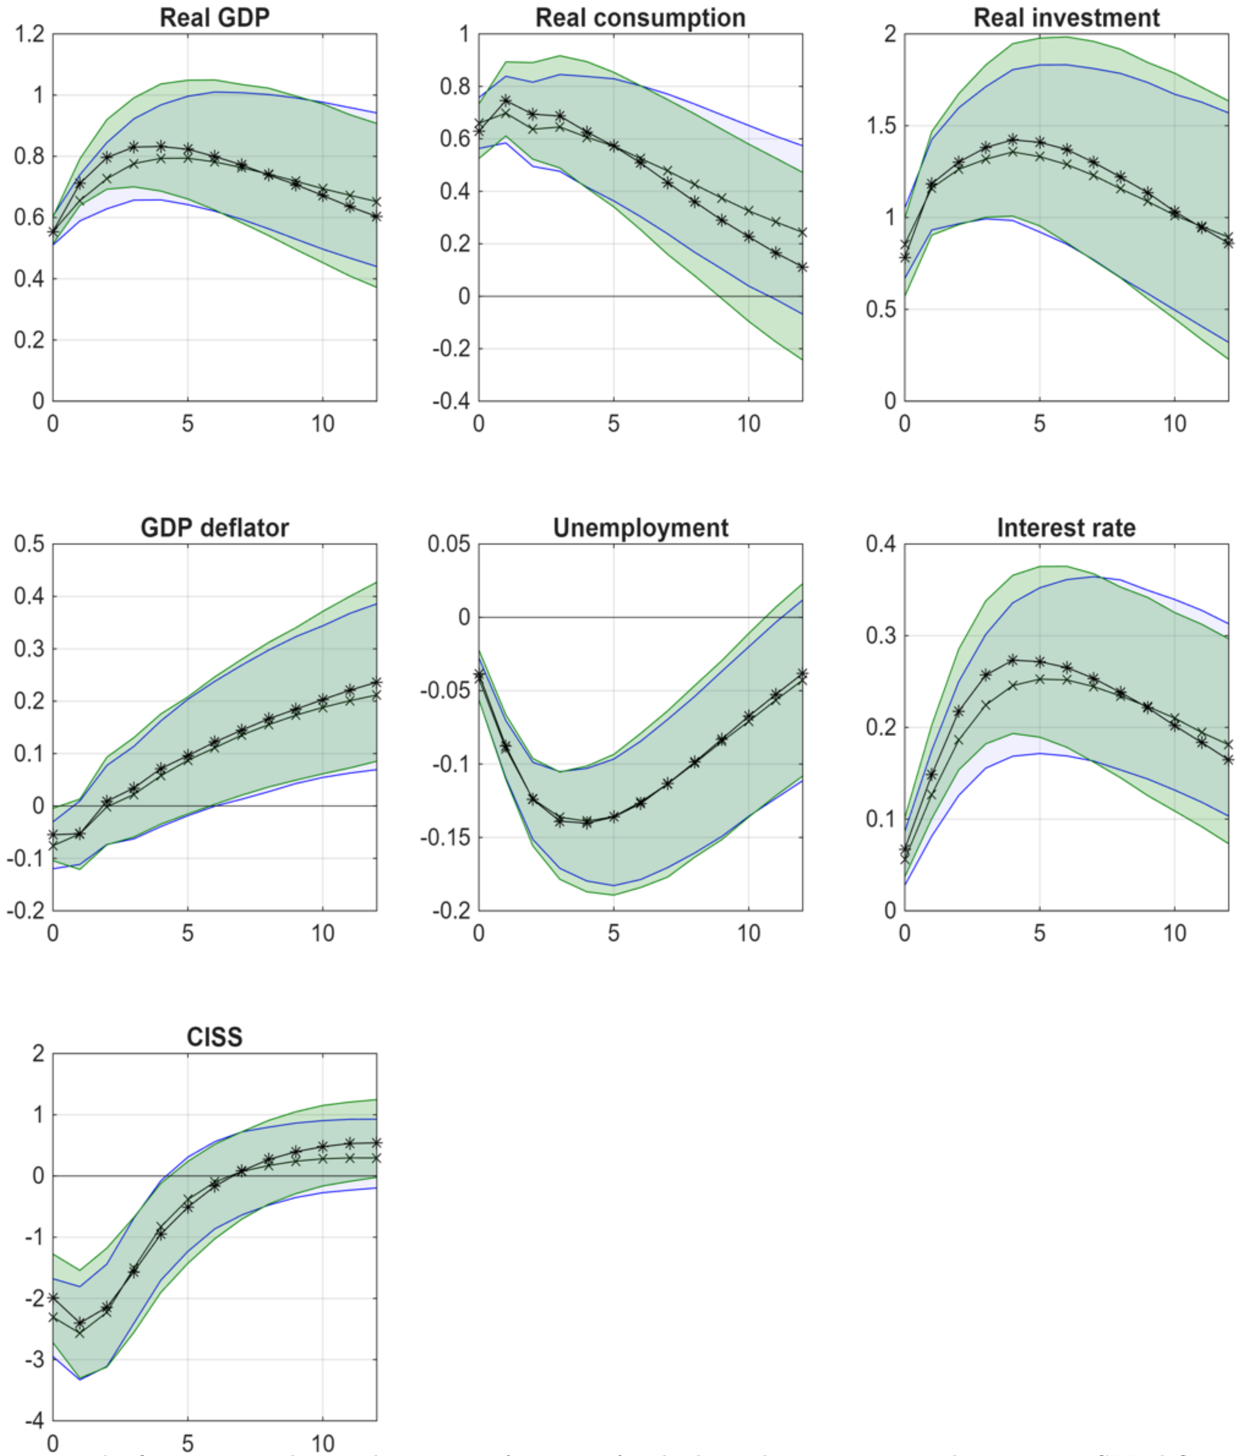

Notes: The figure reports the impulse response functions of real gdp, real consumption, real investment, GDP deflator, unemployment rate, interest rate and CISS to a business cycle shock. **Green Area**: Melitz VAR; **Blue Area**: FVAR, K=10; Solid line with stars: Melitz VAR median; Solid line with crosses: FVAR median. Source: Office for National Statistics, ECB. . Sample: 2001Q1–2025Q2.

	Period	GDP	Cons.	Inv.	Infl.	Unemp.	Int.Rate	CISS	Average
K=4	2016–2025	0.735	0.738	0.721	0.751	0.668	0.804	0.749	0.738
K=6	2016–2025	0.603	0.603	0.571	0.597	0.623	0.767	0.628	0.627
K=8	2016–2025	0.591	0.588	0.547	0.589	0.803	0.773	0.558	0.636
Melitz	2016–2025	0.670	0.666	0.643	0.674	0.651	0.737	0.603	0.663
Melitz*	2016–2025	0.731	0.733	0.728	0.739	0.899	0.895	0.721	0.778

Table 3: **Expanding Windows Forecasts.** Relative MSE. Values smaller than one indicate that the FVAR model is more accurate than the VAR. Melitz is the VAR with average revenues and the average revenues of firms close to default. Melitz* is the VAR with average revenues only.

	Period	GDP	Cons.	Inv.	Infl.	Unemp.	Int.Rate	CISS	Average
K=4	2016–2025	0.592	0.592	0.551	0.591	0.724	0.633	0.568	0.607
K=6	2016–2025	0.544	0.549	0.511	0.565	0.691	0.640	0.588	0.584
K=8	2016–2025	0.539	0.536	0.480	0.558	0.795	0.528	0.624	0.580
Melitz	2016–2025	0.674	0.674	0.636	0.689	0.670	0.659	0.606	0.658
Melitz*	2016–2025	0.721	0.721	0.718	0.731	0.982	0.753	0.745	0.767

Table 4: **Rolling Windows Forecasts.** Relative MSE. Values smaller than one indicate that the FVAR model is more accurate than the VAR. Melitz denotes the VAR augmented with sufficient statistics (average revenues and average revenues of firms close to default).

Recent Working Papers:

For a complete list of Working Papers published by Sveriges Riksbank, see www.riksbank.se

The Macroeconomic Effects of Trade Tariffs: Revisiting the Lerner Symmetry Result *by Jesper Lindé and Andrea Pescatori*	2019:363
Biased Forecasts to Affect Voting Decisions? The Brexit Case *by Davide Cipullo and André Reslow*	2019:364
The Interaction Between Fiscal and Monetary Policies: Evidence from Sweden *by Sebastian Ankargren and Hovick Shahnazarian*	2019:365
Designing a Simple Loss Function for Central Banks: Does a Dual Mandate Make Sense? *by Davide Debortoli, Jinill Kim and Jesper Lindé*	2019:366
Gains from Wage Flexibility and the Zero Lower Bound *by Roberto M. Billi and Jordi Galí*	2019:367
Fixed Wage Contracts and Monetary Non-Neutrality *by Maria Björklund, Mikael Carlsson and Oskar Nordström Skans*	2019:368
The Consequences of Uncertainty: Climate Sensitivity and Economic Sensitivity to the Climate *by John Hassler, Per Krusell and Conny Olovsson*	2019:369
Does Inflation Targeting Reduce the Dispersion of Price Setters' Inflation Expectations? *by Charlotte Paulie*	2019:370
Subsampling Sequential Monte Carlo for Static Bayesian Models *by David Gunawan, Khue-Dung Dang, Matias Quiroz, Robert Kohn and Minh-Ngoc Tran*	2019:371
Hamiltonian Monte Carlo with Energy Conserving Subsampling *by Khue-Dung Dang, Matias Quiroz, Robert Kohn, Minh-Ngoc Tran and Mattias Villani*	2019:372
Institutional Investors and Corporate Investment *by Cristina Cella*	2019:373
The Impact of Local Taxes and Public Services on Property Values *by Anna Grodecka and Isaiah Hull*	2019:374
Directed technical change as a response to natural-resource scarcity *by John Hassler, Per Krusell and Conny Olovsson*	2019:375
A Tale of Two Countries: Cash Demand in Canada and Sweden *by Walter Engert, Ben Fung and Björn Segendorf*	2019:376
Tax and spending shocks in the open economy: are the deficits twins? *by Mathias Klein and Ludger Linnemann*	2019:377
Mind the gap! Stylized dynamic facts and structural models *by Fabio Canova and Filippo Ferroni*	2019:378
Financial Buffers, Unemployment Duration and Replacement Labor Income *by Mats Levander*	2019:379
Inefficient Use of Competitors' Forecasts? *by André Reslow*	2019:380
How Much Information Do Monetary Policy Committees Disclose? Evidence from the FOMC's Minutes and Transcripts *by Mikael Apel, Marianna Blix Grimaldi and Isaiah Hull*	2019:381
Risk endogeneity at the lender/investor-of-last-resort *by Diego Caballero, André Lucas, Bernd Schwaab and Xin Zhang*	2019:382
Heterogeneity in Households' Expectations of Housing Prices – Evidence from Micro Data *by Erik Hjalmarsson and Pär Österholm*	2019:383
Big Broad Banks: How Does Cross-Selling A Affect Lending? *by Yingjie Qi*	2020:384
Unemployment Fluctuations and Nominal GDP Targeting *by Roberto Billi*	2020:385
FAQ: How do I extract the output gap? *by Fabio Canova*	2020:386

Drivers of consumer prices and exchange rates in small open economies *by Vesna Corbo and Paola Di Casola*	2020:387
TFP news, stock market booms and the business cycle: Revisiting the evidence with VEC models *by Paola Di Casola and Spyridon Sichlimiris*	2020:388
The costs of macroprudential deleveraging in a liquidity trap *by Jiaqian Chen, Daria Finocchiaro, Jesper Lindé and Karl Walentin*	2020:389
The Role of Money in Monetary Policy at the Lower Bound *by Roberto M. Billi, Ulf Söderström and Carl E. Walsh*	2020:390
MAJA: A two-region DSGE model for Sweden and its main trading partners *by Vesna Corbo and Ingvar Strid*	2020:391
The interaction between macroprudential and monetary policies: The cases of Norway and Sweden *by Jin Cao, Valeriya Dinger, Anna Grodecka-Messi, Ragnar Juelsrud and Xin Zhang*	2020:392
Withering Cash: Is Sweden ahead of the curve or just special? *by Hanna Armelius, Carl Andreas Claussen and André Reslow*	2020:393
Labor shortages and wage growth *by Erik Frohm*	2020:394
Macro Uncertainty and Unemployment Risk *by Joonseok Oh and Anna Rogantini Picco*	2020:395
Monetary Policy Surprises, Central Bank Information Shocks, and Economic Activity in a Small Open Economy *by Stefan Laséen*	2020:396
Econometric issues with Laubach and Williams' estimates of the natural rate of interest *by Daniel Buncic*	2020:397
Quantum Technology for Economists *by Isaiah Hull, Or Sattath, Eleni Diamanti and Göran Wendin*	2020:398
Modeling extreme events: time-varying extreme tail shape *by Bernd Schwaab, Xin Zhang and André Lucas*	2020:399
The Effects of Government Spending in the Eurozone *by Ricardo Duque Gabriel, Mathias Klein and Ana Sofia Pessoa*	2020:400
Narrative Fragmentation and the Business Cycle *by Christoph Bertsch, Isaiah Hull and Xin Zhang*	2021:401
The Liquidity of the Government Bond Market – What Impact Does Quantitative Easing Have? Evidence from Sweden *by Marianna Blix Grimaldi, Alberto Crosta and Dong Zhang*	2021:402
Five Facts about the Distributional Income Effects of Monetary Policy *by Niklas Amberg, Thomas Jansson, Mathias Klein and Anna Rogantini Picco*	2021:403
When domestic and foreign QE overlap: evidence from Sweden *by Paola Di Casola and Pär Stockhammar*	2021:404
Dynamic Macroeconomic Implications of Immigration *by Conny Olovsson, Karl Walentin, and Andreas Westermark*	2021:405
Revisiting the Properties of Money *by Isaiah Hull and Or Sattath*	2021:406
The cost of disinflation in a small open economy vis-à-vis a closed economy *by Oleksandr Faryna, Magnus Jonsson and Nadiia Shapovalenko*	2021:407
On the Performance of Cryptocurrency Funds *by Daniele Bianchi and Mykola Babiak*	2021:408
The low-carbon transition, climate commitments and firm credit risk *by Sante Carbone, Margherita Giuzio, Sujit Kapadia, Johannes Sebastian Krämer, Ken Nyholm and Katia Vozian*	2022:409
Seemingly Irresponsible but Welfare Improving Fiscal Policy at the Lower Bound *by Roberto M. Billi and Carl E. Walsh*	2022:410
Pension Reform and Wealth Inequality: Evidence from Denmark *by Torben M. Andersen, Joydeep Bhattacharya, Anna Grodecka-Messi and Katja Mann*	2022:411

Inflation Targeting or Fiscal Activism? *by Roberto M. Billi*	2022:412
Trading volume and liquidity provision in cryptocurrency markets *by Daniele Bianchi, Mykola Babiak and Alexander Dickerson*	2022:413
DISPERSION OVER THE BUSINESS CYCLE: PASSTHROUGH, PRODUCTIVITY, AND DEMAND *by Mikael Carlsson, Alex Clymo and Knut-Eric Joslin*	2022:414
Electoral Cycles in Macroeconomic Forecasts *by Davide Cipullo and André Reslow*	2022:415
The Curious Incidence of Monetary Policy Across the Income Distribution *by Tobias Broer, John Kramer and Kurt Mitman*	2022:416
Central Bank Mandates and Monetary Policy Stances: through the Lens of Federal Reserve Speeches *by Christoph Bertsch, Isaiah Hull, Robin L. Lumsdaine, and Xin Zhang*	2022:417
The Political Costs of Austerity *by Ricardo Duque Gabriel, Mathias Klein and Ana Sofia Pessoa*	2022:418
Central bank asset purchases: Insights from quantitative easing auctions of government bonds *by Stefan Laséen*	2023:419
Greenflation? *by Conny Olovsson and David Vestin*	2023:420
Effects of foreign and domestic central bank government bond purchases in a small open economy DSGE model: Evidence from Sweden before and during the coronavirus pandemic *by Yildiz Akkaya, Carl-Johan Belfrage, Paola Di Casola and Ingvar Strid*	2023:421
Dynamic Credit Constraints: Theory and Evidence from Credit Lines* *by Niklas Amberg, Tor Jacobson, Vincenzo Quadrini and Anna Rogantini Picco*	2023:422
Stablecoins: Adoption and Fragility *by Christoph Bertsch*	2023:423
CBDC: Lesson from a Historical Experience *by Anna Grodecka-Messi and Xin Zhang*	2023:424
Do Credit Lines Provide Reliable Liquidity Insurance? Evidence from Commercial-Paper Backup Lines *by Niklas Amberg*	2023:425
Price Pass-Through Along the Supply Chain: Evidence from PPI and CPI Microdata *by Edvin Ahlander, Mikael Carlsson and Mathias Klein*	2023:426
Cash for Transactions or Store-of-Value? A comparative study on Sweden and peer countries *by Carl Andreas Claussen, Björn Segendorf and Franz Seitz*	2023:427
Fed QE and bank lending behaviour: a heterogeneity analysis of asset purchases *by Marianna Blix Grimaldi and Supriya Kapoor*	2023:428
Monetary policy in Sweden after the end of Bretton Woods *by Emma Bylund, Jens Iversen and Anders Vredin*	2023:429
Banking Without Branches *by Niklas Amberg and Bo Becker*	2024:430
Climate impact assessment of retail payment services *by Niklas Arvidsson, Fumi Harahap, Frauke Urban and Anissa Nurdiawati*	2024:431
Four Facts about International Central Bank Communication b*y Christoph Bertsch, Isaiah Hull, Robin L. Lumsdaine, and Xin Zhang*	2024:432
Optimal Monetary Policy with r* < 0 *by Roberto Billi, Jordi Galí, and Anton Nakov*	2024:433
Quantitative Easing, Bond Risk Premia and the Exchange Rate in a Small Open Economy *by Jens H. E. Christensen and Xin Zhang*	2024:434
Supply-Chain Finance: An Empirical Evaluation of Supplier Outcomes *by Niklas Amberg, Tor Jacobson and Yingjie Qi*	2024:435
Optimal Contracts and Inflation Targets Revisited *by Torsten Persson and Guido Tabellini*	2024:436
Potential Climate Impact of Retail CBDC Models *by Niklas Arvidsson, Fumi Harahap, Frauke Urban and Anissa Nurdiawati*	2024:437

Sveriges Riksbank
Visiting address: Brunkebergs torg 11
Mail address: se-103 37 Stockholm

Website: www.riksbank.se
Telephone: +46 8 787 00 00, Fax: +46 8 21 05 31
E-mail: registratorn@riksbank.se