

# Who Benefits? Distributional Effects of Fiscal Policy on Household Income

**Preliminary and incomplete version prepared for presentation at the October, 2026 Forum for Macroeconomics and Macroeconomic Policy (FMM) Conference. Please do not quote without permission. \***

Daniel Fernández Romero<sup>†</sup>

Rubén González<sup>‡</sup>

October 19, 2026

## Abstract

This paper examines the distributional effects of fiscal policy shocks. We combine household-level income data with three externally identified fiscal shocks: two deficit-driven fiscal consolidation shocks and one legislated tax shock. We identify dynamic responses using reduced-form and instrumental-variable local projections using an administrative micro dataset tracking Spanish tax records from 1999 to 2016. Fiscal shocks generate sizeable aggregate income effects that operate mainly through labor income. The gains from a fiscal expansion are across the bottom 99% of tax filers on impact and become progressive—an inverted S-shape—once the multi-year implementation of fiscal packages is taken into account, with no evidence of gains at the top. Tax shocks generate a related but less redistributive pattern. Counterfactual simulations translate these heterogeneous responses into income inequality changes. Overall, expansionary fiscal policy lowers income inequality among tax filers.

**Keywords:** *income distribution, fiscal shocks, tax shocks, individual heterogeneity, micro-data.*

**JEL Classification:** *D31, E63, C55.*

---

\*The authors acknowledge financial support received through the project PID2022-136664OA-I00 by The Ministry of Science and Innovation. The authors gratefully acknowledge the Institute for Fiscal Studies (IEF) for granting access to the representative sample of administrative income data from the Household Panel of the Personal Income Tax (IRPF, for its acronym in Spanish). We thank Beatriz de Blas, Aránzazu de Juan, Pilar Poncela, Víctor Sancibrián, Martín Almuzara, Ana Silva, Gabriel Pérez Quirós, and Alessandro Giovannelli for helpful comments. We also thank the attendees of the UAM Economic Theory Seminar, the participants of the XXVII Applied Economics Meeting and the 50th Symposium of the Spanish Economic Association, where this work was presented.

<sup>†</sup>Universidad Autónoma de Madrid (UAM). Departamento de Economía Aplicada. E-mail: [daniel.fernandezr@uam.es](mailto:daniel.fernandezr@uam.es)

<sup>‡</sup>Universidad Complutense de Madrid (UCM). Departamento de Economía Aplicada, Pública y Política. E-mail: [rubenmgo@ucm.es](mailto:rubenmgo@ucm.es)

# 1 Introduction

The design and implementation of fiscal policy play a central role in shaping aggregate income growth and its distribution across households. This makes its analysis central for understanding macroeconomic dynamics and the associated distributional implications. Fiscal policy, through its control of taxation, public spending, and transfer programs, directly influences disposable income, social insurance, and human capital accumulation. How does fiscal policy affect income growth across household groups? Can fiscal policy reduce income inequality? Is it labor income or capital income that contributes most to explaining the change in income growth?

We lack precise answers to such questions because the empirical literature has commonly focused on survey data, aggregated time series, or summary measures such as the Gini coefficient. These studies have overlooked important heterogeneities in the effects. For instance, survey data are often top-coded, and large effects in the tails may offset each other in summary measures of inequality, making them difficult to identify in time series ([Amberg, Jansson, Klein, & Picco, 2022](#)). This limitation is particularly important for fiscal policy. While recent work has made substantial progress in estimating the distributional income effects of macroeconomic shocks using administrative micro data, much less is known about the distributional mechanics of fiscal policy and about whether fiscal expansions and consolidations affect income through the same parts of the distribution.

In this paper, we aim to fill this gap by analyzing how fiscal policy affects income growth across the income distribution. To do so, we combine administrative income data from Spain with externally identified fiscal shocks for the period 1999-2016 obtained from [Gil, Martí, Morris, Pérez, and Ramos \(2019\)](#); [Gupta, Jalles, Mulas-Granados, and Schena \(2018\)](#); [Jordà and Taylor \(2025\)](#); [Leigh, Pescatori, Devries, and Guajardo \(2011\)](#). The identification of fiscal policy shocks comes from “narrative” approach. The former capture deficit-driven fiscal adjustments that combine spending and revenue measures and are measured relative to GDP, whereas the latter isolate legislated tax changes. Furthermore, we run simulations that account for heterogeneity in responses to document how inequality would have evolved in the absence of these policies.

Spain provides a compelling case for analyzing the effects of fiscal policy on income. First, Spain implemented extensive fiscal consolidation after the Great Recession, which affected household income. Second, the unemployment rate in Spain is among the highest in Europe, directly influencing labor income and making the labor-income channel especially relevant. Third, as shown in [Figure F1](#) and [Figure J1](#), income inequality in Spain increased over 2003-2016, a period that overlaps closely with our estimation sample. Fourth, Spain combines a progressive personal income tax, a large role for labor income, and a fiscal consolidation episode that affected both spending and taxation, making it a useful setting to study which parts of the distribution bear the income effects of fiscal shocks.

We estimate reduced-form and instrumental variables local projections ([Jordà, 2005](#); [Jordà & Taylor, 2025](#)), comparing household income in periods with fiscal shocks to periods without such shocks. We use detailed administrative tax records to estimate how

fiscal shocks and tax shocks transmit across the income distribution, and we decompose the responses into labor and capital income. These tax records come from Spain’s Tax Agency and are the most reliable source for measuring income and its various components at an annual frequency, making them well suited to study heterogeneous distributional responses. The unit of observation is the tax filing unit. The main sample is the *Pure Personal Income Tax Panel*, which follows tax filers over 1999-2016.

Moreover, because the fiscal shocks are measured at the same annual frequency as the income data, we avoid the time-aggregation issues that arise when high-frequency shocks must be aggregated to match lower-frequency outcomes (Buda et al., 2025; Ramey, 2011).<sup>1</sup> To the best of our knowledge, this is the first study to use administrative tax records to estimate the heterogeneous income effects of fiscal policy across the full income distribution and to compare fiscal shocks with legislated tax shocks.

We find that fiscal tightening has sizeable income effects. A deficit-driven fiscal consolidation of 1% of GDP results in -4.88 p.p. decline in aggregate after-tax income at the peak accumulated response after two years. The effect is mainly explained by labor income, while capital income responds less strongly. These estimates show that fiscal shocks generate economically large income responses in the administrative data. However, the aggregate response masks substantial heterogeneity across the income distribution, which is why we estimate heterogeneous responses by income percentile.

Next, we estimate the impact on each part of the income distribution. The fiscal shock generates income gains across the bottom 99% of tax filers, with no evidence of gains for the top 1%. Once the multi-year implementation of fiscal packages is taken into account through a cumulative shock measure, the incidence becomes progressive: we observe an inverted S-shaped response in total income across groups, with the strongest gains concentrated among lower-income tax filers, progressively smaller effects along the distribution, and no effect at the top. This pattern is consistent with a strong labor income response and relatively muted capital income effects, as in Gechert and Heimberger (2022); Hope and Limberg (2022). This progressive distributional response of total income to fiscal shocks in administrative tax records is the central empirical result of the paper.

Estimating a deficit-driven shock and a tax shock allows us to assess whether the distributional response of fiscal policy reflects asymmetrical effects. We find that tax shocks generate a related but less redistributive pattern: the gains at the bottom are smaller and the differences across groups narrower. Thus, fiscal policy reduces inequality more effectively than tax changes alone.

For robustness, we provide additional evidence using another Household Tax Panel which benefits from the absence of panel attrition (Appendix I). The estimates confirm our main result: we observe the same progressive response to fiscal shocks. We also examine the aggregate effects of the tax shock using a proxy SVAR in Appendix C. Given the weak first stage in that much more demanding aggregate environment, we treat this exercise as a suggestive diagnostic rather than as supporting evidence. In the distributional exercise,

---

<sup>1</sup>Appendix A discusses in detail how aggregating high-frequency shocks to the annual frequency of the income data can introduce endogeneity.

we report dynamic responses over several horizons and focus on the two-year horizon because this is where the responses reach their peak and are comparable to the related literature.

This paper shows that fiscal policy does more than stimulate income: it also redistributes income across low-, middle-, and high-income households. The main contribution is to bring fiscal policy into the administrative-microdata literature on distributional macroeconomic policy effects. The instrumental-variable exercise in [Section 5.2](#) maps the narrative shocks into realized fiscal variables, so the income responses can be interpreted as multipliers. This puts our estimates on the scale of the fiscal-multiplier literature ([Gechert, 2015](#)). Two interpretations govern this comparison. Our estimates refer to the after-tax income of tax filers rather than to GDP, and they are identified mainly from a crisis period in which fiscal effects are plausibly larger than in normal times.

These findings are consistent with how household balance sheets shape the transmission of fiscal policy. Liquidity-constrained, low-income households consume a large share of any additional income, which amplifies short-run stimulus effects, whereas households with high debt burdens use transfers or tax rebates to deleverage rather than to spend ([Dyann, 2012](#); [Kaplan & Violante, 2014](#); [Mian & Sufi, 2010](#)). Anticipated and preannounced tax changes—such as payroll tax adjustments or stimulus rebates—can generate asymmetric consumption responses: constrained households react more strongly, while unconstrained, higher-income households tend to save or invest the additional resources ([Johnson, Parker, & Souleles, 2006](#); [Parker, 1999](#); [Souleles, 2002](#)). Additionally, governments design social protection spending to support vulnerable groups. In practice, this can create complex distributional effects: it first boosts lower-income households and, through labor-market adjustments or tax financing, can pass indirect gains to higher-income groups ([Isiaka, Mihailov, & Razzu, 2025](#)).

The distributional heterogeneity we document is difficult to detect without administrative micro data. Because households differ significantly in their income composition, labor market attachment, tax exposure, and dependence on transfers, the aggregate impact of fiscal policy masks substantial heterogeneity, making it crucial to account for both micro-level channels and macro-level responses. In light of persistently high and rising inequality, alongside evolving fiscal policy toolkits, it is essential that policy evaluation consider both distributional and growth-related effects.

The rest of the paper is organized as follows: [Section 2](#) summarizes the related literature, focusing on the closest papers on distributional macroeconomic policy effects. [Section 3](#) presents the model that will be estimated using data described in [Section 4](#). Hence in [Section 4](#), we describe the data sources and provide a short description of the externally identified fiscal shocks used in the analysis. In [Section 5](#), we report the aggregate income responses, express them in multiplier form through an instrumental-variable local-projection exercise, and present the heterogeneous responses across the income distribution, the inequality simulations, and the contributions of income components. Finally, [Section 6](#) concludes.

## 2 Related Literature

The effects of fiscal policy on income and inequality are one of the most studied questions in empirical macroeconomics. Recent decades have seen a sustained rise in income inequality across Western democracies, a trend that has garnered attention in the wake of the Great Financial Crisis and the accompanying fiscal policy responses (Blanchet & Martínez-Toledano, 2023; Blanco, Bauluz, & Martínez-Toledano, 2021). Yet, despite the large literature on fiscal multipliers and redistribution, we still know relatively little about how fiscal shocks affect income growth across the distribution using administrative micro data.

Our paper is closest to the recent literature that combines macroeconomic shocks with micro-level data to study distributional responses. For instance Coibion, Gorodnichenko, Kueng, and Silvia (2017) study the effects of monetary policy shocks on income and consumption inequality in the United States using survey data. They show that contractionary monetary policy increases inequality and highlight several channels through which aggregate policy shocks can have heterogeneous household-level effects. Amberg et al. (2022) use Swedish administrative income records and show that the income effects of monetary policy shocks are U-shaped over the income distribution: expansionary shocks increase the incomes of low- and high-income individuals relative to middle-income individuals. They also show that labor income explains the response at the bottom, while capital income explains the response at the top. Holm, Paul, and Tischbirek (2021) use Norwegian administrative data and study monetary transmission along the liquid-asset distribution, while Andersen, Johannesen, Jørgensen, and Peydró (2023) use Danish administrative data and find that softer monetary policy increases income, wealth, and consumption more strongly for high-income households. These papers show the value of administrative data for uncovering distributional responses that are difficult to detect with survey data or summary inequality statistics.

Relative to this literature, our paper shifts the focus from monetary policy to fiscal policy. This distinction is important because fiscal policy directly affects disposable income through taxes, transfers, and public spending, while monetary policy operates indirectly through interest rates, asset prices, credit conditions, and labor-market responses. Therefore, the distributional mechanics of fiscal policy need not mirror the U-shaped responses documented for monetary policy. Our contribution is to estimate how externally identified fiscal shocks transmit across the full income distribution using Spanish administrative tax records, and to decompose the responses into labor and capital income.

Our paper also relates to the empirical literature on the heterogeneous and distributional effects of fiscal policy. Using cross-country evidence, Furceri, Ge, Loungani, and Melina (2022) show that fiscal consolidations increase inequality and poverty, while Heimberger (2020) find persistent increases in disposable-income inequality following narrative consolidation episodes, especially when adjustments are large, prolonged, and based on spending cuts. Dhital, Jiang, and Reese (2023) use U.S. household survey data and a time-varying FAVAR to show that expansionary government spending reduces inequal-

ity because lower-income households gain more than higher-income households. Related household-level studies have focused primarily on heterogeneous consumption responses to government spending and tax shocks (Cloyne & Surico, 2017; Giavazzi & McMahon, 2013), rather than on income responses across the full distribution. Complementing this empirical evidence, Ferriere and Navarro (2025) show that the distribution of the tax burden across households shapes government-spending multipliers, highlighting the importance of tax progressivity for both the aggregate transmission and the distributional incidence of fiscal expansions. These studies establish that fiscal policy can have important distributional consequences, but they generally rely on cross-country aggregates, survey data, or outcomes other than household income and do not follow the same tax-filing units across narrowly defined income groups.

A second closely related strand studies the heterogeneous effects of tax policy. Mertens and Montiel Olea (2018) use narratively identified changes in U.S. marginal tax rates and specific income-group time series to estimate income responses, finding positive elasticities throughout the distribution but larger short-run effects at the top. Their estimand, however, is the elasticity of pretax reported income to marginal tax rate, rather than the incidence of broad fiscal shocks on after-tax income. In the Spanish context, Almunia and Lopez-Rodriguez (2019) use the same underlying administrative tax panel to estimate taxable-income elasticities around personal income tax reforms, whereas our analysis combines the longitudinal panel with aggregate fiscal shocks to estimate dynamic after-tax income responses across fixed initial income groups. Zidar (2019) shows that U.S. tax cuts directed toward lower-income groups generate stronger employment growth than cuts for high-income taxpayers, while Hope and Limberg (2022) and Gechert and Heimberger (2022) find that tax cuts targeted at high-income individuals or corporations generate few broad-based economic gains and tend to increase inequality. Relative to this literature, we study both deficit-driven fiscal adjustments and legislated tax shocks, trace their effects across the entire income distribution using administrative microdata, hold group composition fixed over the event-study horizon, and decompose the responses into labor and capital income. Our results therefore characterize the dynamic effects of aggregate fiscal policy rather than the behavioral elasticity of reported income to tax reforms.

Our paper contributes to this literature in three ways. First, we use a large administrative dataset with individual-level income information from 1999 to 2016. Survey data and inequality indices such as the Gini often hide tail heterogeneity in the distribution (Amberg et al., 2022). In contrast, our approach enables a more granular analysis of policy effects across the entire income distribution. Second, we distinguish between reduced-form fiscal shocks and tax shocks. Reduced-form fiscal shocks capture broad fiscal consolidations and expansions that combine spending and revenue measures, whereas tax shocks isolate legislated tax changes. This distinction allows us to assess whether the distributional response reflects general fiscal adjustment or the tax component alone. Third, we decompose total income responses into labor and capital income, which allows us to identify the income components driving the heterogeneous effects.

Our main finding is that fiscal policy generates a progressive response across the income distribution, which takes the form of an inverted S-shape once the multi-year implementation of fiscal packages is accounted for. Expansionary fiscal shocks benefit income most for low-income tax filers. Middle-income groups show smaller gains, while the top of the distribution sees no effect. The response is mainly driven by labor income, with relatively muted capital income effects. Tax shocks generate a related but less redistributive pattern.

### 3 Model

We study the distributional effects of fiscal policy using event-study panel local projections with aggregate fiscal shocks (Almuzara & Sancibrián, 2024; Jordà, 2005; Jordà & Taylor, 2025). Let  $Y_{i,t+h}^k$  denote income component  $k$  of household  $i$  at horizon  $h$ , and let  $Y_{i,t-1}^T$  denote total after-tax household income in  $t - 1$ . We define the dependent variable as

$$\Delta_h y_{it}^k \equiv \frac{Y_{i,t+h}^k - Y_{i,t-1}^k}{Y_{i,t-1}^T} \times 100,$$

so that each coefficient is interpreted as the change in income component  $k$  between  $t - 1$  and  $t + h$ , expressed as a percentage of pre-shock after-tax income. This normalization follows Amberg et al. (2022) and is useful because all income components are scaled by the same pre-shock income measure. As a result, for any additive decomposition of total after-tax income, the estimated response of total income can be decomposed into the contribution of each component. For example, the contribution of labor income and capital income to the total after-tax income response can be obtained directly by comparing their respective coefficients to the coefficient estimated for total after-tax income.

Our design belongs to a large and well-established literature that estimates micro-level responses to externally identified aggregate shocks (Amberg et al., 2022; Andersen et al., 2023; Cloyne, Ferreira, & Surico, 2020; Coibion et al., 2017; Holm et al., 2021). Almuzara and Sancibrián (2024) formalize the econometrics of panel local projections in this setting. As in that literature, the identifying variation is aggregate time-series variation in the shock: the micro panel does not add independent aggregate shocks, but it sharpens the measurement of the response and makes the distributional and component-level decompositions possible. Relative to the monetary applications, which typically exploit a large number of high-frequency policy surprises, our fiscal shocks are annual. This is the price of avoiding the time-aggregation problem documented in Appendix A, and it makes the treatment of pre-shock dynamics, the concentration of consolidation episodes, and inference with a short time dimension particularly important. We address each of these issues below and, for the episode concentration, in the robustness analysis of Appendix F.

Throughout the event-study analysis, households are assigned to an income group using their first observed position in the income distribution. Let  $\tau_i$  denote the first year in which household  $i$  appears in the sample and let  $g_i^0$  denote the percentile group to

which household  $i$  belongs in year  $\tau_i$ . This classification is kept fixed over time. Fixing group membership is important because fiscal shocks may themselves affect the income measure used to rank households. If households were allowed to move across percentile groups during the event window, part of the estimated response would reflect endogenous changes in the composition of each group rather than income changes for a given group of households. Our percentile-specific estimates compare the evolution of the same observed percentile groups, as defined at household entry into the sample, holding group composition fixed over the horizon of the event study.<sup>2</sup>

Our first specification pools all households and estimates, for each income component  $k$  and horizon  $h \in \{-H_-, \dots, -2, 0, 1, \dots, H_+\}$ ,

$$\Delta_h y_{it}^k = \alpha_i^{k,h} + \phi_{g_i^0}^{k,h} + \gamma_h t + \beta_h^{k,s} Z_t^s + \lambda_h^{k'} X_{it} + u_{it}^{k,h,s}, \quad h \neq -1, \quad (1)$$

where  $Z_t^s$  denotes fiscal shock  $s$  in year  $t$ ,  $\alpha_i^{k,h}$  are household fixed effects,  $\phi_{g_i^0}^{k,h}$  are fixed effects for the household's initial income group, and  $\gamma_h t$  is a linear time trend. The coefficient of interest is  $\beta_h^{k,s}$ , which measures the average response of income component  $k$  at horizon  $h$  to fiscal shock  $s$ . The vector  $X_{it}$  includes two lags of the corresponding income component, and fixed effects for the household's region of residence. The term  $\lambda_h^{k'}$  are the coefficients of the control variables. We write the initial-income-group, and region indicators explicitly in (1) to make the composition of the specification transparent. For tax units whose composition and residence do not change over the sample, these indicators are absorbed by the household fixed effects, and they contribute identifying variation only through the units for which they vary over time. Finally,  $u_{it}^{k,h,s}$  is the error term, capturing unobserved characteristics of the change in income component  $k$  for household  $i$  at horizon  $h$ . All regressions are estimated using household sampling weights. The omitted horizon is  $h = -1$ , so the coefficients for  $h < 0$  describe pre-shock dynamics, while the coefficients for  $h \geq 0$  trace the post-shock response. In the baseline implementation, we set  $H_- = 4$  and  $H_+ = 4$ .

Identification requires that, conditional on household fixed effects, baseline income-group fixed effects, the time trend, and the included controls, the timing and intensity of the narrative fiscal shocks are orthogonal to unobserved counterfactual changes in household income. The coefficients at negative horizons provide a diagnostic check of this con-

---

<sup>2</sup>Appendix G reconciles our estimates with earlier results in the literature, particularly those of Amberg et al. (2022). We first replicate their specification without tax-unit fixed effects and allow households to move across contemporaneously defined income groups, even though such transitions may themselves be induced by the fiscal shock. We then address this potential source of endogenous reclassification by assigning each household to the income group in which it is first observed and holding that assignment fixed throughout the sample. Finally, we augment this specification with tax-unit and year fixed effects, using the top 1 percent as the omitted category. The drawback of this last specification is that the resulting coefficients measure differential responses relative to the top 1 percent rather than the absolute response of each income group. To avoid this rescaling, our baseline analysis estimates a separate regression for each fixed initial income group instead of a saturated model containing interactions for all groups simultaneously. This allows us to retain tax-unit fixed effects while estimating the absolute response of each group, without normalizing the coefficients relative to an omitted income category and while avoiding potentially endogenous transitions across income groups. We interpret these estimates as the absolute group-specific responses to the fiscal shock, net of these potential compositional and specification-related pitfalls.

dition by testing whether income was already evolving differentially before years with larger fiscal shocks.<sup>3</sup>

Our second specification estimates the same model separately for each initial income group. For each  $g \in \mathcal{G}$ , we restrict the sample to households such that  $g_i^0 = g$  and estimate

$$\Delta_h y_{it}^k = \alpha_{i,g}^{k,h} + \gamma_{g,h} t + \beta_{g,h}^{k,s} Z_t^s + \lambda_{g,h}^{k'} X_{it} + u_{it,g}^{k,h,s}, \quad h \neq -1. \quad (2)$$

where  $\alpha_{i,g}^{k,h}$  are household fixed effects within initial income group  $g$ ,  $\gamma_h t$  is a linear time trend, and  $\beta_{g,h}^{k,s}$  is the coefficient of interest. It measures the response of income component  $k$  at horizon  $h$  to fiscal shock  $s$  for households belonging to initial income group  $g$ . The vector  $X_{it}$  includes the same demographic and regional controls as in the pooled specification: fixed effects for the region the tax-filer lived. In these percentile-specific regressions, we include only one lag of the corresponding income component.<sup>4</sup> The term  $\lambda_{g,h}^{k'}$  are the coefficients associated to the control variables. Finally,  $u_{it,g}^{k,h,s}$  is the group-specific error term, capturing unobserved characteristics of the change in income component  $k$  for household  $i$  at horizon  $h$  within initial income group  $g$ . All regressions are estimated using household sampling weights. This specification delivers heterogeneous event-study responses across the observed income-percentile groups, while keeping group composition fixed over the event-study horizon. This avoids mechanically attributing to fiscal shocks changes that arise from endogenous movements of households across percentile groups after the shock.

We estimate both specifications for the main income aggregates studied in the paper, namely total after-tax income, total labor income, and total capital income.

The continuous fiscal shock is our baseline treatment. This choice follows the logic of difference-in-differences designs with treatment intensity, where the magnitude of treatment is exploited directly rather than discretized into treated and untreated observations (Callaway, Goodman-Bacon, & Sant'Anna, 2024). In our setting, the size of the narrative fiscal shock is economically informative, so  $\beta_h^{k,s}$  captures the dynamic response of income component  $k$  to a one-unit change in fiscal consolidation at horizon  $h$ .

As a complementary treatment measure, we replace the contemporaneous fiscal shock with a horizon-specific cumulative fiscal shock,

$$C_{t,h}^s = \sum_{j=0}^h Z_{t+j}^s, \quad (3)$$

which measures the cumulative fiscal consolidation between the event year  $t$  and horizon  $h$  (Klein Martins, 2025; Teulings & Zubanov, 2014). We then re-estimate equations

---

<sup>3</sup>Our design is related to, but distinct from, the LP-DiD framework in Dube, Girardi, Jordà, and Taylor (2025). LP-DiD is formulated for settings with discrete treatment timing and explicit clean-control comparisons. In our application, the fiscal shock is aggregate and all households in a given year are exposed to the same macro-fiscal innovation. Identification therefore comes from time variation in treatment intensity, while the panel dimension allows us to trace dynamic and distributional heterogeneity across households.

<sup>4</sup>Each percentile-specific regression is estimated on a substantially smaller subsample than the aggregated specification. Requiring two lags would further reduce the effective sample because observations must be available for an additional preceding year, leading to a sample loss, particularly in the smaller percentile groups.

(1) and (2) replacing  $Z_t^s$  with  $C_{t,h}^s$ . The motivation is the way fiscal adjustment is implemented: consolidation packages are typically legislated as multi-year plans over several budgets (Alesina, Azzalini, Favero, Giavazzi, & Miano, 2018; Alesina, Barbiero, Favero, Giavazzi, & Paradisi, 2015; Alesina, Favero, & Giavazzi, 2015), so the fiscal impulse relevant for household income between  $t$  and  $t + h$  is the total adjustment implemented over that window rather than the initial innovation alone (Carrière-Swallow, David, & Leigh, 2021; Komatsuzaki, 2024).<sup>5</sup> Because  $C_{t,h}^s$  contains measures after  $t$ , its causal interpretation requires that the future measures included in the sum are pre-determined with respect to household income shocks between  $t$  and  $t + h$ . This assumption is plausible in our setting because the large Spanish consolidation measures of 2010–2013 were implemented over multi-year programs adopted under the European Union’s Excessive Deficit Procedure, so the calendar and size of the reforms were committed in advance rather than decided in response to household incomes. Nevertheless, we treat the contemporaneous and cumulative specifications as complementary: the former traces the response to the initial, plausibly unanticipated innovation, while the latter summarizes the response to the implemented multi-year adjustment.

For inference, all specifications use heteroskedasticity-robust standard errors that combine two-way clustering at the household and year levels with a Newey-West HAC correction. Household clustering accounts for arbitrary serial correlation within tax-filing units, whereas year clustering accommodates contemporaneous correlation across households arising from their exposure to the same fiscal shock and other common macroeconomic disturbances. The fiscal shock series and the common component of the regression disturbance may also display residual serial correlation across years. This concern is particularly relevant in local projections because overlapping forward outcomes generate horizon-dependent serial correlation in the regressions. We apply the Newey-West correction separately at each projection horizon using a horizon-specific bandwidth. The resulting covariance estimator accommodates heteroskedasticity, two non-nested dimensions of within-cluster dependence, and serially correlated common disturbances, while incorporating a Driscoll-Kraay type correction for cross-sectional dependence (Driscoll & Kraay, 1998; Newey & West, 1986).

We apply this inference procedure to all reduced-form and instrumental-variable specifications, including the aggregate, distributional, and robustness analyses. Because the shock varies only at the annual level, the effective sample size for  $\beta_h^{k,s}$  is governed by the time dimension of the panel, which is short: the corrections above are asymptotic in the number of years, and the reported confidence bands should be read with this caveat in mind. This choice is the benchmark in micro evidence on the transmission of aggregate shocks to household outcomes (Cloyne et al., 2020) and is consistent with recent work on

---

<sup>5</sup>Related approaches address fiscal measures occurring within the projection horizon by including subsequent fiscal shocks as controls (Klein Martins, 2025; Teulings & Zubanov, 2014). A related strand estimates cumulative or integral fiscal multipliers by relating cumulative output changes to cumulative fiscal changes over the same horizon (Ramey & Zubairy, 2018). Our specification differs in that the cumulative fiscal adjustment itself is used as the horizon-specific treatment variable as Carrière-Swallow et al. (2021), and Komatsuzaki (2024).

panel local projections with aggregate shocks, which emphasizes the strong cross-sectional dependence induced by common aggregate shocks in micro-macro panels ([Almuzara & Sancibrián, 2024](#)).<sup>6</sup> More generally, the panel structure induces serial dependence within households, while the annual fiscal shock creates common dependence across households observed in the same year. Because these two dimensions are distinct and non-nested, one-way clustering is insufficient. This motivation is consistent with evidence that ignoring serial correlation in longitudinal designs can severely over-reject the null ([Bertrand, Duflo, & Mullainathan, 2004](#)), with the use of multi-way clustering under non-nested dependence ([Cameron, Gelbach, & Miller, 2011](#)), and with the principle that standard errors should reflect the level at which treatment is assigned ([Abadie, Athey, Imbens, & Wooldridge, 2023](#)).

## 4 Data

Our analysis uses the *Pure Panel of Personal Income Tax Filers 1999-2016*, an administrative longitudinal dataset produced by the Institute for Fiscal Studies (Instituto de Estudios Fiscales, IEF) using personal income tax records from the Spanish Tax Agency (Agencia Estatal de Administración Tributaria, AEAT). The panel follows tax filing units over the period 1999-2016 and is representative of personal income tax filers in Spain's common fiscal territory, excluding the Basque Country and Navarre, which operate under their own tax administrations.<sup>7</sup> Although incomes below the filing threshold are exempt and the lowest incomes are therefore underrepresented, many low-income individuals file voluntarily to claim withholding refunds or deductions for maternity, large families, or housing, which raises coverage among low-income households ([González, Sala, & Trivín, 2026](#)).

The underlying annual samples are obtained through stratified random sampling by region, income bracket, and main source of income, with sampling weights provided in the data. The unit of observation is the tax filing unit-year. In the raw files, each return is linked to a main filer, and the administrative structure also records spouse and household identifiers when applicable. We conduct the empirical analysis at the level of the main filer-year, which is the appropriate unit for following tax units over time and for measuring income and tax liabilities consistently within Spain's personal income tax system. The raw pure panel contains 5,311,457 filer-year observations over 1999-2016. The base-year sample in 2003 contains 383,237 tax returns, and the 2016 sample contains 227,390 tax returns, implying an official cumulative retention rate of 59.33 percent in the final year.<sup>8</sup> Details on the sample construction are provided in [Appendix D](#). After these data processing, the estimation sample at horizon two contains 861,884 filer-year observations.

The main outcome is real total after-tax income. We define total pre-tax income as the sum of labor income, business income, real-estate capital income, financial income, and

---

<sup>6</sup>For a survey of empirical applications using alternative clustering schemes, see Appendix F of [Almuzara and Sancibrián \(2024\)](#).

<sup>7</sup>In [Appendix K](#), we map the distribution of income across Spanish municipalities, which shows the regional differences in reported income.

<sup>8</sup>In [Appendix I](#), we show that the results are not driven by attrition in the sample.

realized capital gains. Total after-tax income is obtained by subtracting the final personal income tax liability after allowances, deductions, credits, withholdings, and advance payments. This is the preferred income concept throughout the paper because it corresponds to the disposable income measure directly affected by fiscal policy through both the tax schedule and the income components recorded in the tax return. All monetary variables are converted to real terms using the Spanish GDP deflator, base 2015.

These administrative records offer three advantages for the question studied in this paper. First, they jointly record income components and final tax liability within the same tax return, allowing us to distinguish whether fiscal shocks operate through labor income, business income, or capital income. Second, income components are recorded as administrative amounts rather than survey reports, which is essential for studying the upper tail of the distribution. Third, the longitudinal structure allows us to compare income changes for the same tax units over time while assigning households to pre-determined income groups.

The main sample is restricted to individuals aged 26–65. Fiscal shocks affect household income through several channels, but the short- and medium-run transmission that we estimate is expected to operate primarily through market income: employment, hours, wages, capital income, and personal taxation. These margins are most relevant for working-age individuals. By contrast, older tax units receive a larger share of income through pensions and other institutional transfer mechanisms, whose response to fiscal shocks may reflect benefit-indexation rules rather than the labor-market and tax-base channels central to our analysis. [Appendix H](#) reports robustness checks that include older filers.

We assign tax units to income groups using only pre-treatment income information, because the model in [Section 3](#) estimates group-specific responses across the income distribution. [Appendix D](#) provides the exact percentile construction. In brief, individuals are ranked within year using a moving average of real total pre-tax income measured, and are then assigned to eleven groups based on their first observed position: deciles 1–9, the p90–p99 group, and the top 1 percent.


## 4.1 Stylized facts

Two descriptive facts motivate the distributional analysis. First, Spain’s personal income tax is progressive. [Figure 1](#) shows that average tax rates increase along the income distribution, from the lower deciles to the top 1 percent. This matters for our empirical design because a common fiscal shock can translate into different after-tax income responses depending on the tax burden faced by each part of the distribution. In [Appendix E](#), the [Table E1](#) reports descriptive statistics for the main income variables and demographic characteristics by income percentile group.

Second, the composition of pre-tax income changes substantially across income groups. [Figure 2](#) shows that labor income accounts for most pre-tax income over the bulk of the distribution. Capital income is small for most filers but rises at the top, where financial income and capital gains become more relevant. Transfers are more important at the bot-


tom and limited elsewhere in the tax data. These patterns imply that aggregate income responses can mask heterogeneous transmission margins: fiscal shocks may affect lower and middle groups mainly through labor income, while the top of the distribution is more exposed to capital-income and tax-base channels.

**Figure 1:** Average tax rate by income group



*Notes:* The figure reports the average tax rate by income group over 1999–2016. The average tax rate is defined as final net personal income tax liability after deductions, credits, withholdings, and advance payments divided by total pre-tax income. Income groups are defined using the pre-treatment income classification described in [Appendix D](#).

**Figure 2: Income shares by income group**



*Notes:* The figure reports the average shares of labor income, capital income, and transfers in total pre-tax income by income group over 1999–2016. Labor income includes wage and self-employment income. Capital income includes real-estate capital income, financial income, and realized capital gains.

## 4.2 Fiscal consolidation shocks

The identification of fiscal policy shocks through narrative evidence follows the approach pioneered by [Romer and Romer \(2010\)](#) for legislated tax changes. Rather than treating all movements in fiscal variables as shocks, the narrative approach uses contemporaneous policy records, legislative documents, and official statements to classify fiscal actions according to their stated motivation. This distinction is crucial because fiscal variables are highly endogenous to the business cycle: tax revenues and transfers respond automatically to changes in activity, and discretionary measures may also be adopted for short-run stabilization purposes. The relevant shocks are discretionary fiscal actions whose stated motivation is not to respond to macroeconomic conditions.

We use narrative fiscal shocks from previous macro-fiscal literature. These shocks are designed to isolate discretionary fiscal actions that are plausibly orthogonal to contemporaneous cyclical developments. We distinguish between deficit-driven fiscal consolidation shocks and tax shocks.

The broad fiscal consolidation shock is obtained from the narrative consolidation datasets of [Leigh et al. \(2011\)](#), [Guajardo, Leigh, and Pescatori \(2014\)](#), [Gupta et al. \(2018\)](#), and updated in [Jordà and Taylor \(2025\)](#). We use the most recent version narrative shocks,

obtained directly from [Jordà and Taylor \(2025\)](#).<sup>9</sup> These series identify discretionary fiscal measures motivated by deficit reduction, rather than by short-run stabilization objectives. The shock is measured by the estimated budgetary impact of the consolidation package in the year in which the measure is implemented. A positive value denotes fiscal tightening. Since consolidation packages typically combine expenditure cuts and revenue-increasing measures, this variable should be interpreted as a broad deficit-driven consolidation shock.

The tax shock is obtained from the narrative dataset for Spain constructed by [Gil et al. \(2019\)](#). This series records legislated tax changes, their implementation dates, and their estimated revenue effects. This measure provides a direct way to isolate the revenue-side component of discretionary fiscal policy. However, since our income data are annual, the quarterly tax shocks must be aggregated to the year before being merged with the panel.<sup>10</sup> This aggregation reduces the amount of identifying variation and weakens the precision with which tax-specific responses can be estimated. For this reason, our preferred specification uses the deficit-driven fiscal consolidation shock, which is measured directly at the annual frequency and captures the overall discretionary budgetary effort associated with consolidation episodes. We use the tax shock as a complementary exercise to assess the extent to which the distributional effects of fiscal consolidation reflect revenue-side measures.

All fiscal shocks used in the empirical analysis are expressed at the annual frequency and scaled by GDP. We then merge the annual shock series to the panel by year. In the results section, we report responses to an expansionary fiscal shock. Hence, we multiply the consolidation measure by minus one, so that positive values correspond to fiscal expansions and negative values to fiscal tightening.

This distinction matters for the interpretation of the transmission mechanism. A deficit-driven fiscal consolidation may affect tax-unit income through several margins: reductions in public consumption or investment, changes in public-sector employment and wages, cuts in transfers and social spending, and tax increases. By contrast, a tax shock operates only through the tax system, and its incidence depends on which tax bases, rates, exemptions, or brackets are adjusted. For example, changes in personal income taxation are more likely to affect higher taxable-income groups directly, consistent with the increasing average tax rate across the income distribution documented in [Figure 1](#), whereas indirect tax measures may have a broader incidence through real disposable income.


[Figure 3](#) plots the annual fiscal shock series used in the analysis. Fiscal consolidations are small during the early part of the sample and become much larger during the

<sup>9</sup>Appendix B assesses the robustness of our results using the previous version of the fiscal shock series from [Leigh et al. \(2011\)](#), [Guajardo et al. \(2014\)](#), and [Gupta et al. \(2018\)](#)

<sup>10</sup>This aggregation imposes a non-negligible cost. As discussed in [Appendix A](#), aggregating shocks originally identified at a higher frequency may attenuate their identifying content, obscure the timing of transmission, and weaken the link between the shock and the policy margin faced by households. The concern is less severe here than for high-frequency monetary surprises, because the tax series is based on legislated fiscal actions rather than narrow-window asset-price movements. Nevertheless, annual aggregation necessarily discards within-year information and may reduce first-stage relevance and precision. For this reason, we do not use the quarterly tax shock as our baseline source of identification. Instead, we use it as a complementary benchmark to distinguish the distributional effects of broad deficit-driven fiscal consolidations from those associated with a more specific narrative revenue-side shock.


sovereign-debt crisis, especially between 2010 and 2013. This pattern reflects the Spanish fiscal adjustment after the Great Recession. [Figure 4](#) plots the annual tax-shock series. The tax-shock series is also concentrated in the crisis period, but it captures only legislated revenue measures. Its largest realizations coincide with the period of substantial tax increases, including changes to personal income taxation and the taxation of savings income. This timing is useful for interpreting whether the estimated income responses are driven mainly by labor-income effects, direct tax incidence, or capital-income channels.

**Figure 3: Narrative fiscal consolidation shocks**



*Notes:* the figure plots the annual narrative deficit-driven fiscal consolidation shocks used in the empirical analysis over the period 1999-2019. The series measure the estimated budgetary impact of discretionary fiscal measures implemented with the stated objective of reducing the government deficit, expressed as a percentage of GDP. Positive values denote fiscal tightening, whereas negative values denote fiscal expansion. Panel (a) reports the consolidation series obtained from [Leigh et al. \(2011\)](#) and [Gupta et al. \(2018\)](#), supplemented with the authors' compilation for Spain. Panel (b) reports the alternative series obtained from [Guajardo et al. \(2014\)](#) and [Jordà and Taylor \(2025\)](#). These measures may combine expenditure cuts, tax increases, transfer reductions, and other discretionary budgetary actions. In the income specific estimation, the sign of the consolidation shock is reversed so that positive values correspond to expansionary fiscal shocks.

**Figure 4: Narrative tax consolidation shock**



*Notes:* the figure plots the annual narrative tax shock for Spain constructed by [Gil et al. \(2019\)](#) over the sample period 1986–2016. The shock records legislated tax changes and their estimated revenue effects, expressed as a percentage of GDP. The original series is available at quarterly frequency and is aggregated to the year by summing the quarterly realizations before being merged with the annual administrative income panel. Positive values denote revenue-increasing tax measures, or tax-based fiscal tightening, whereas negative values denote tax cuts. Unlike the broad fiscal consolidation measures, this series captures only legislated revenue-side actions and therefore excludes expenditure cuts and other non-tax components of fiscal adjustment. In the income specific estimation, the sign of the tax shock is reversed so that positive values correspond to expansionary tax shocks.

In the empirical analysis, we take the deficit-driven fiscal consolidation shock as our baseline measure. We then repeat the exercise using the tax shock as a complementary specification. If the deficit-driven consolidation shock generates stronger or more persistent responses than the tax shock, this points to expenditure-side or general-equilibrium labor-market channels. If the tax shock generates a similar distributional pattern, this suggests that revenue-side incidence is an important part of the transmission mechanism.

## 5 Results

### 5.1 Aggregate event-study local projections

We begin by examining the aggregate dynamic effects of fiscal policy on household income. Figures 5–8 report the aggregate income responses obtained from the aggregated specification in [Equation 1](#). The estimates reveal a clear contractionary effect of fiscal tightening on household income. Both deficit-driven consolidations and legislated tax increases reduce total after-tax income, with the effects appearing on impact and generally reaching their largest magnitude over the following two years. The decomposition across income sources shows that these aggregate losses are driven primarily by labor income. Capital income also declines initially, but its contribution is substantially smaller and less persistent.


The comparison across shocks further indicates that broad fiscal consolidations generate larger and more persistent income losses than tax increases alone. This distinction is economically informative because the fiscal-consolidation shock captures adjustments op-

erating through several policy margins, whereas the tax shock isolates legislated changes in taxation. Overall, the aggregate estimates establish that the narrative shocks transmit strongly to the household income recorded in the administrative data and identify labor income as the principal channel underlying that transmission.

Figure 5 reports the baseline responses to the continuous deficit-driven consolidation shock. A fiscal tightening immediately reduces total after-tax income by approximately 3.6 percentage points. The decline remains close to 3.5 percentage points after one year, reaches its largest magnitude at  $h = 2$ , at around 4.9 percentage points, and remains sizeable at  $h = 3$ . The 95% confidence intervals exclude zero from  $h = 0$  to  $h = 3$ . By  $h = 4$ , however, the response to the initial fiscal innovation has returned close to zero and is no longer statistically significant.

The decomposition shows that this profile is predominantly explained by labor income. Labor income falls by approximately 2.3 percentage points on impact, 2.9 percentage points after one year, and 5.6 percentage points at  $h = 2$ . The response remains negative at  $h = 3$ , at around 4.2 percentage points, before disappearing at the end of the horizon. In contrast, capital income falls by approximately 0.9 percentage points on impact and 0.5 percentage points after one year, but subsequently returns close to its pre-shock level. Thus, the medium-run decline in household income closely tracks the labor-income response. This pattern is consistent with fiscal consolidation operating through aggregate demand and labor-market conditions, including employment and wage adjustments, rather than primarily through persistent changes in capital returns.

**Figure 5:** Response of income components to a continuous fiscal shock




Notes: The figure reports event-study panel local-projection estimates of the response of income to the continuous fiscal consolidation shock, following Equation 1. Panel (a) reports real total after-tax income, panel (b) real total labor income, and panel (c) real total capital income. For each horizon  $h$ , the dependent variable is the change in the corresponding income component between  $t - 1$  and  $t + h$ , scaled by total after-tax income in  $t - 1$  and expressed in percentage points. The omitted horizon is  $h = -1$ . The fiscal shock is the narrative deficit-driven consolidation shock, measured as a percentage of GDP; positive values denote fiscal tightening. Estimates are weighted using the household sampling factor and include tax unit fixed effects, initial income-group fixed effects, region fixed effects, a linear time trend, and two lags of the corresponding income component. The dark-gray shaded area reports the 68% confidence interval, while the light-gray shaded area reports the 95% confidence interval. Standard errors are heteroskedasticity robust, two-way clustered at the household and year levels, and incorporate a Driscoll-Kraay correction implemented through a Newey-West correction at each specific horizon bandwidth. Outcome changes are winsorized at the 5th and 95th percentiles.

The continuous tax shock in Figure 6 generates a qualitatively similar but somewhat smaller response. A legislated tax increase reduces total after-tax income by around 2.3 percentage points on impact, 3.0 percentage points at  $h = 1$ , and 3.4 percentage points at  $h = 2$ . The response subsequently weakens, falling to approximately 1.5 percentage points

at  $h = 3$  and becoming statistically indistinguishable from zero at  $h = 4$ . Labor income again accounts for most of the medium-run effect: its response increases in magnitude from approximately  $-1.3$  percentage points on impact to  $-3.2$  percentage points at  $h = 2$ , and remains negative at  $h = 3$ . Capital income instead exhibits a front-loaded response, declining by approximately  $0.6$  percentage points on impact and  $0.5$  percentage points after one year, but showing little response thereafter. The tax results therefore combine an immediate reduction in disposable resources with a delayed labor-income adjustment.


**Figure 6:** Response of income components to a continuous tax shock



*Notes:* The figure reports event-study panel local-projection estimates of the response of income to the continuous tax shock, following Equation 1. Panel (a) reports real total after-tax income, panel (b) real total labor income, and panel (c) real total capital income. For each horizon  $h$ , the dependent variable is the change in the corresponding income component between  $t - 1$  and  $t + h$ , scaled by total after-tax income in  $t - 1$  and expressed in percentage points. The omitted horizon is  $h = -1$ . The tax shock is the narrative legislated tax consolidation shock, measured as a percentage of GDP; positive values denote tax increases. Estimates are weighted using the household sampling factor and include tax unit fixed effects, initial income-group fixed effects, region fixed effects, a linear time trend, and two lags of the corresponding income component. The dark-gray shaded area reports the 68% confidence interval, while the light-gray shaded area reports the 95% confidence interval. Standard errors are heteroskedasticity robust, two-way clustered at the household and year levels, and incorporate a Driscoll-Kraay correction implemented through a Newey-West correction at each specific horizon bandwidth. Outcome changes are winsorized at the 5th and 95th percentiles.

The cumulative-shock specifications show that the disappearance of the continuous-shock coefficients at longer horizons should not be interpreted as a complete reversal of the effects of fiscal consolidation. When the fiscal impulse is measured by the total consolidation implemented through each horizon, Figure 7 shows a persistent decline in total after-tax income. The estimated response is approximately  $-3.6$  percentage points on impact, remains between  $-3.3$  and  $-3.8$  percentage points from  $h = 1$  to  $h = 3$ , and is still around  $-2.1$  percentage points at  $h = 4$ . All post-shock coefficients are statistically different from zero at the 95% level. Labor income displays a similarly persistent response, ranging from approximately  $-2.3$  percentage points on impact to  $-3.0$  percentage points at  $h = 2$ , and remaining at  $-2.4$  percentage points four years after the shock. Capital income declines by around  $0.9$  percentage points on impact, but its contribution gradually diminishes and is effectively zero by  $h = 4$ . The persistent cumulative response is therefore mainly a labor-income phenomenon.


**Figure 7: Response of income components to a cumulative fiscal shock**



*Notes:* The figure reports event-study panel local-projection estimates of the response of income to the cumulative fiscal consolidation shock, following Equation 1 and replacing the contemporaneous shock with the cumulative shock defined in Equation 3. Panel (a) reports real total after-tax income, panel (b) real total labor income, and panel (c) real total capital income. For each horizon  $h$ , the dependent variable is the change in the corresponding income component between  $t - 1$  and  $t + h$ , scaled by total after-tax income in  $t - 1$  and expressed in percentage points. The omitted horizon is  $h = -1$ . The cumulative fiscal shock is the sum of narrative deficit-driven consolidation shocks between year  $t$  and horizon  $h$ , measured as a percentage of GDP; positive values denote fiscal tightening. Estimates are weighted using the household sampling factor and include tax unit fixed effects, initial income-group fixed effects, region fixed effects, a linear time trend, and two lags of the corresponding income component. The dark-gray shaded area reports the 68% confidence interval, while the light-gray shaded area reports the 95% confidence interval. Standard errors are heteroskedasticity robust, two-way clustered at the household and year levels, and incorporate a Driscoll-Kraay correction implemented through a Newey-West correction at each specific horizon bandwidth. Outcome changes are winsorized at the 5th and 95th percentiles.

The cumulative tax results in Figure 8 lead to the same general conclusion. Total after-tax income falls by approximately 2.3 percentage points on impact and remains significantly below its pre-shock level throughout the four-year horizon. The response is around  $-2.5$  percentage points at  $h = 1$ ,  $-2.2$  percentage points at  $h = 2$ , and reaches approximately  $-3.2$  percentage points at  $h = 4$ . Labor income becomes increasingly important over the horizon, with its response rising in magnitude from  $-1.3$  percentage points on impact to approximately  $-2.8$  percentage points after four years. Capital income also remains negative, but its response progressively attenuates from approximately  $-0.6$  to  $-0.2$  percentage points. Hence, once the multi-year implementation of legislated tax measures is incorporated, the income losses do not disappear: the initial capital-income response weakens, while the labor-income effect becomes the main source of persistence.

**Figure 8: Response of income components to a cumulative tax shock**



*Notes:* The figure reports event-study panel local-projection estimates of the response of income to the cumulative tax shock, following Equation 1 and replacing the contemporaneous shock with the cumulative shock defined in Equation 3. Panel (a) reports real total after-tax income, panel (b) real total labor income, and panel (c) real total capital income. For each horizon  $h$ , the dependent variable is the change in the corresponding income component between  $t - 1$  and  $t + h$ , scaled by total after-tax income in  $t - 1$  and expressed in percentage points. The omitted horizon is  $h = -1$ . The cumulative tax shock is the sum of narrative legislated tax shocks between year  $t$  and horizon  $h$ , measured as a percentage of GDP; positive values denote tax increases. Estimates are weighted using the household sampling factor and include tax unit fixed effects, initial income-group fixed effects, region fixed effects, a linear time trend, and two lags of the corresponding income component. The dark-gray shaded area reports the 68% confidence interval, while the light-gray shaded area reports the 95% confidence interval. Standard errors are heteroskedasticity robust, two-way clustered at the household and year levels, and incorporate a Driscoll-Kraay correction implemented through a Newey-West correction at each specific horizon bandwidth. Outcome changes are winsorized at the 5th and 95th percentiles.

The continuous and cumulative specifications consequently provide complementary information. The former trace the response to the initial fiscal innovation and show that the largest income losses generally occur one or two years after the shock. The latter account for the fact that fiscal packages are often implemented over several years and reveal a more persistent effect on household income. The cumulative results are relevant for connecting the micro-level estimates to the fiscal-multiplier literature, in which medium-run responses are commonly normalized by the total policy adjustment implemented over the corresponding horizon.

The event-study evidence shows no signs of anticipation. For total after-tax income and labor income, the pre-shock coefficients exhibit no systematic pattern before the fiscal event, while the responses emerge only after the shock. Some departures from this pattern appear in the capital-income specifications. These estimates are inherently noisier because capital income is highly concentrated at the top of the distribution and is characterized by substantial volatility and a greater incidence of extreme observations. The capital-income pre-trends should therefore be interpreted with caution. Taken together, the absence of systematic pre-trends in the main income aggregates and the complementary IV evidence support the interpretation that tax increases generate relatively immediate income losses, whereas broader fiscal consolidations produce more persistent effects operating primarily through labor income.

## 5.2 IV local-projection income multipliers

Having documented the reduced-form income responses to fiscal shocks, we next express the aggregate effects in multiplier form. The IV local-projection analysis in this section provides an economic interpretation of these effect. The narrative shocks strongly predict the corresponding realized fiscal variables—the government deficit and individual

taxes paid—with average post-shock first-stage  $F$ -statistics of 110.14 and 22.51, respectively. Mapping the shocks into realized fiscal changes provides an economically meaningful scale for distinguishing the direct effects of taxation from those associated with broader deficit adjustments. The resulting multipliers point to a more immediate tax channel and a more persistent deficit channel.

We emphasize at the outset that this exercise is an interpretation and scaling device rather than an independent source of identification: the identifying variation remains the annual narrative shocks, and the instrumental-variable step maps them into realized fiscal variables so that the income responses can be read per unit of observed fiscal adjustment rather than per unit of the narrative measure. The resulting coefficients facilitate the economic interpretation and comparison of deficit- and tax-based fiscal adjustments, and they provide a scale that is closer to the fiscal-multiplier literature, which are expressed per one percent of GDP of the narrative shock.

This section estimates aggregate income multipliers using an instrumental-variable version of the local-projection design. The purpose is to complement the reduced-form estimates reported above with a specification that maps externally identified fiscal shocks into realized fiscal variables. We consider two fiscal variables. First, we estimate a tax multiplier using individual personal income taxes paid in the administrative records. The corresponding instrument is the annual aggregated narrative tax shock from [Gil et al. \(2019\)](#). Second, we estimate a deficit multiplier using the general government deficit for Spain from the OECD. The corresponding instrument is the narrative deficit-driven consolidation shock from the fiscal-consolidation literature ([Guajardo et al., 2014](#); [Gupta et al., 2018](#); [Jordà & Taylor, 2025](#); [Leigh et al., 2011](#)).

The exercise follows the cumulative-multiplier logic in [Jordà and Taylor \(2025\)](#); [Ramey \(2016\)](#), but the outcome is measured at the individual level. Hence, the coefficients should be interpreted as income multipliers rather than GDP multipliers: they measure the response of individual total after-tax income to exogenous movements in fiscal variables. Since total after-tax income is one component of the broader income generated in the economy, this specification connects the fiscal-multiplier literature to the micro-distributional analysis of household income.

For each horizon  $h \geq 0$ , we define the cumulative fiscal instrument as

$$C_{t,h}^s = \sum_{j=0}^h Z_{t+j}^s, \quad (4)$$

where  $Z_t^s$  denotes either the narrative tax shock or the narrative deficit-driven fiscal consolidation shock. The cumulative instrument captures the total fiscal adjustment implemented between the event year and horizon  $h$ . This normalization is useful because fiscal packages are often phased in over several years. A specification based only on the impact shock would therefore understate the relevant fiscal impulse at medium horizons.

Let  $\Delta_h y_{it}^T$  denote the percentage change in real total after-tax income between  $t - 1$  and

$t + h$ , scaled by real total after-tax income in  $t - 1$ :

$$\Delta_h y_{it}^T = 100 \times \frac{Y_{i,t+h}^T - Y_{i,t-1}^T}{Y_{i,t-1}^T}. \quad (5)$$

We estimate, separately for each horizon, the following two-stage local-projection model:

$$F_{it,h}^r = \pi_h^r C_{t,h}^s + \alpha_i^{r,h} + \lambda_h^{r'} X_{it} + \eta_{it}^{r,h}, \quad (6)$$

$$\Delta_h y_{it}^T = \beta_h^r \hat{F}_{it,h}^r + \alpha_i^h + \lambda_h' X_{it} + u_{it}^h. \quad (7)$$

The variable  $F_{it,h}^r$  denotes the realized fiscal variable whose multiplier is estimated. In the tax specification,  $F_{it,h}^r$  is the horizon-specific change in personal income taxes paid by tax unit  $i$ . In the deficit specification,  $F_{it,h}^r$  is the OECD general-government deficit for Spain, merged to the micro panel by year. The vector  $X_{it}$  includes lagged total after-tax income, lagged taxes, region controls, and a linear trend. All specifications include individual fixed effects and use the household sampling weights. The omitted horizon is  $h = -1$ , which is normalized to zero. Negative horizons are estimated as placebo responses.


The coefficient  $\beta_h^r$  is the IV income multiplier at horizon  $h$ . In the tax specification, it measures the response of after-tax income to the component of individual taxes paid that is induced by the narrative tax shock. In the deficit specification, it measures the response of after-tax income to the component of the OECD deficit that is induced by the narrative deficit-driven consolidation shock. Since the instruments vary only at the annual aggregate level, the relevant identifying variation is time-series variation in fiscal shocks. The micro panel improves the measurement of the income response and permits later distributional decompositions, but it does not create additional independent aggregate shocks.

For inference, we report confidence interval calculated with Driscoll-Kraay standard errors, allowing for the cross-sectional dependence generated by common aggregate shocks and the serial correlation induced by overlapping local-projection outcomes. As described in the model section, the bandwidth is allowed to vary with the projection horizon ([Almuzara & Sancibrián, 2024](#); [Jordà & Taylor, 2025](#)). Instrument relevance is assessed using the horizon-specific Kleibergen–Paap rk Wald  $F$ -statistic, and we additionally report its average over the post-shock horizons  $h \geq 0$ .

We read this exercise as a scaling device for two reasons. First, the instruments—and, in the deficit specification, the endogenous variable itself—vary only at the annual level, so the first stage is effectively an annual time-series relationship; the reported first-stage statistics should be interpreted with the short time dimension in mind. Second, the two endogenous variables are of a different nature, and the coefficients inherit that difference. In the deficit specification, both the instrument and the endogenous variable are macroeconomic aggregates, so the coefficient scales the income response per unit of aggregate fiscal adjustment. In the tax specification, the endogenous variable is the personal income taxes paid by each filer, which move mechanically with the filer’s own income through the marginal rate schedule. Instrumenting this individual variable with an aggregate shock

yields a coefficient that should be read as a personal-income-tax multiplier for tax filers rather than as the tax multiplier of the macro literature.


**Figure 9:** Aggregate IV-LP income multiplier: cumulative fiscal shock



*Notes:* The figure reports aggregate IV local-projection estimates for real total after-tax income. The endogenous fiscal variable is the OECD general-government deficit for Spain. The instrument is the cumulative narrative deficit-driven fiscal consolidation shock. The dependent variable is the percentage change in total after-tax income relative to the pre-shock baseline. The omitted horizon is  $h = -1$ . The dark-gray shaded area reports the 68% confidence interval, while the light-gray shaded area reports the 95% confidence interval, based on Driscoll-Kraay standard errors. The average first-stage  $F$ -statistic over post-shock horizons  $h \geq 0$  is 110.14.

Figure 9 shows the IV income multiplier associated with the deficit measure. The first stage is strong: the average post-shock  $F$ -statistic is 110.14, indicating that the cumulative narrative consolidation shock is highly relevant for explaining movements in the OECD deficit measure for Spain. The pre-shock coefficients are small and close to zero, providing no evidence of a systematic decline in household income before fiscal consolidations. After the shock, total after-tax income falls immediately, with an impact response of about  $-0.34$  p.p. at  $h = 0$ . The response remains negative at  $h = 1$ , around  $-0.28$  p.p., and becomes larger thereafter, reaching approximately  $-0.40$  p.p. at  $h = 2$ ,  $-0.89$  p.p. at  $h = 3$ , and  $-1.00$  p.p. at  $h = 4$ . This pattern indicates that deficit-driven consolidations have persistent and increasingly large effects on household income, and that the income losses are not reversed within the four-year window.

**Figure 10:** Aggregate IV-LP income multiplier: cumulative tax shock



*Notes:* The figure reports aggregate IV local-projection estimates for real total after-tax income. The endogenous fiscal variable is personal income taxes paid by individuals in the administrative tax records. The instrument is the cumulative narrative tax shock. The dependent variable is the percentage change in total after-tax income relative to the pre-shock baseline. The omitted horizon is  $h = -1$ . The dark-gray shaded area reports the 68% confidence interval, while the light-gray shaded area reports the 95% confidence interval, based on Driscoll–Kraay standard errors. The average first-stage  $F$ -statistic over post-shock horizons  $h \geq 0$  is 22.51.

Figure 10 reports the corresponding tax multiplier. The average post-shock (F)-statistic is 22.51, showing that the cumulative narrative tax shock is relevant for explaining the personal income taxes paid by individuals in the administrative tax records. The tax shock produces a negative response of total after-tax income from the impact year onward. The response is relatively strong in the short run: on impact and over the first post-shock horizons, the tax multiplier is at least as large as, and in some horizons larger than, the deficit multiplier. This is consistent with a direct tax-liability channel, whereby legislated tax increases immediately reduce disposable income through higher taxes paid. However, the response does not strengthen over the horizon. Instead, it remains broadly stable, from approximately  $-0.53$  p.p. on impact to around  $-0.79$  p.p. at  $h = 4$ . By contrast, the deficit-based response is smaller on impact, around  $-0.34$  p.p., but becomes larger at longer horizons, reaching approximately  $-0.89$  p.p. at  $h = 3$  and  $-1.00$  p.p. at  $h = 4$ .

The comparison between Figure 9 and Figure 10 is informative about transmission channels. Tax shocks operate primarily through the direct tax-liability channel: they change the taxes paid by households and therefore affect disposable income mechanically and through incentive effects. This explains why the tax multiplier is relatively strong in the short run. Deficit-driven consolidations, by contrast, capture broader fiscal packages that combine tax increases, spending cuts, transfer changes, public-sector wage adjustments, and changes in public investment. Their effect builds over time and becomes larger than the tax multiplier at longer horizons. The larger long-run deficit multiplier is therefore consistent with an expenditure channel. Fiscal consolidations based on spending cuts can reduce aggregate demand, employment, and wage income more strongly and more persistently than tax-only measures. This interpretation is in line with the fiscal-multiplier

literature, which generally finds larger and more persistent effects for spending-based fiscal policy than for tax-based policy (Gechert, 2015; Gechert & Rannenberg, 2018). It is also consistent with evidence that fiscal consolidations have long-lasting effects on economic activity and income (Fernández Romero, 2025; Klein Martins, 2025).


Overall, this section shows that the main reduced-form results are not driven only by the direct tax channel. When fiscal shocks are used as instruments for realized fiscal variables, both tax and deficit multipliers imply persistent declines in total after-tax income after fiscal tightening. The deficit multiplier becomes larger at longer horizons, suggesting that expenditure-side and general-equilibrium labor-market channels are central for understanding the aggregate income effects of fiscal consolidations.

### 5.3 Impulse Response Heterogeneity

We next examine how the aggregate income response is distributed across initial income groups. Figures 11–14 report the heterogeneous estimates obtained from Equation 2 at horizon  $h = 2$  for total after-tax income, labor income, and capital income. We focus on the two-year horizon because the aggregate responses documented in the previous subsection generally reach their peak effect at this point. Moreover, the dynamic income-group-specific estimates reported in Appendix F show that  $h = 2$  typically captures the peak response across the distribution. This horizon therefore provides the most informative comparison of the incidence of fiscal policy.

The estimates reveal two distinct distributional patterns. For the continuous shocks, total after-tax income responds relatively uniformly across the bottom 99% of the distribution, before changing sharply for the top 1%, producing a sideways J-shaped profile. By contrast, the cumulative specifications display a clearer inverted S-shaped gradient: income gains are largest among low-income tax filers, decline progressively with initial income, and become negligible or negative at the top. Across specifications, labor income accounts for most of this heterogeneity, whereas capital-income responses are smaller and considerably less precise, particularly at the top of the distribution.

**Figure 11: Heterogeneous income responses to an expansionary continuous fiscal shock**




*Notes:* The figure reports income-group-specific event-study panel local-projection estimates at horizon  $h = 2$ , following Equation 2. Panel (a) reports real total after-tax income, panel (b) real total labor income, and panel (c) real total capital income. For each component, the dependent variable is the change between  $t - 1$  and  $t + 2$ , scaled by total after-tax income in  $t - 1$  and expressed in percentage points. The solid line plots the sign-reversed coefficient, so that positive values correspond to an expansionary fiscal shock equal to 1% of GDP. The underlying shock is the narrative deficit-driven fiscal consolidation shock; positive realizations of the original series denote fiscal tightening. Tax-filing units are assigned to their initial income group and remain in that group throughout the analysis. The first nine groups correspond to income deciles, followed by the 90th–99th percentiles and the top 1%. Estimates use household sampling weights and include household fixed effects, region fixed effects, a linear time trend, and two lags of the corresponding income component. The dark-gray shaded area reports the 68% confidence interval, while the light-gray shaded area reports the 95% confidence interval. Standard errors are heteroskedasticity robust, two-way clustered at the household and year levels, and incorporate a Newey–West correction with the horizon-specific bandwidth. Outcome changes are winsorized at the 5th and 95th percentiles. The sample covers 1999–2016.

Figure 11 shows that the initial fiscal innovation produces sizeable income gains across nearly the entire distribution when its sign is reversed to represent an expansion. The response of total after-tax income remains around five to six percentage points throughout most of the distribution, with only a moderate decline among the upper deciles. The top 1% is the clear exception: its point estimate turns negative and its confidence interval widens substantially. Thus, the contemporaneous shock does not generate a pronounced gradient among the bottom 99%, but it does produce a sharp discontinuity at the upper tail.

Labor income closely reproduces this profile and accounts for most of the positive response below the top percentile. The estimates are positive and economically sizeable from the bottom decile through the 90th–99th percentile group. This indicates that fiscal expansions transmit primarily through earnings and employment-related margins rather than through capital income. By contrast, capital-income responses remain close to zero across most of the distribution and become negative only at the top. The upper-tail estimate is imprecise, however, and should not be interpreted as firm evidence of a negative capital-income effect. Capital income is highly concentrated among top earners and contains volatile items such as financial returns and realized capital gains, making estimates for the top percentile particularly sensitive to extreme observations even in administrative data.

**Figure 12: Heterogeneous income responses to an expansionary continuous tax shock**




*Notes:* The figure reports income-group-specific event-study panel local-projection estimates at horizon  $h = 2$ , following Equation 2. Panel (a) reports real total after-tax income, panel (b) real total labor income, and panel (c) real total capital income. For each component, the dependent variable is the change between  $t - 1$  and  $t + 2$ , scaled by total after-tax income in  $t - 1$  and expressed in percentage points. The solid line plots the sign-reversed coefficient, so that positive values correspond to an expansionary tax shock equal to 1% of GDP. The underlying shock is the narrative legislated tax shock from Gil et al. (2019); positive realizations of the original series denote tax increases. Tax-filing units are assigned to their initial income group and remain in that group throughout the analysis. The first nine groups correspond to income deciles, followed by the 90th–99th percentiles and the top 1%. Estimates use household sampling weights and include household fixed effects, region fixed effects, a linear time trend, and two lags of the corresponding income component. The dark-gray shaded area reports the 68% confidence interval, while the light-gray shaded area reports the 95% confidence interval. Standard errors are heteroskedasticity robust, two-way clustered at the household and year levels, and incorporate a Newey–West correction with the horizon-specific bandwidth. Outcome changes are winsorized at the 5th and 95th percentiles. The sample covers 1999–2016.

The continuous tax shock in Figure 12 generates a related but less uniform pattern. Total after-tax income increases across the bottom 90% of the distribution, although not statistically significant, with the largest point estimates located around the lower-middle and middle income groups. The response weakens in the upper deciles and becomes negative for the top 1%. Hence, the tax shock also produces a sideways J-shaped profile, although the variation within the bottom 99% is somewhat greater than for the broader fiscal shock.


Labor income again explains most of the positive response. The gains increase from the bottom toward the middle of the distribution and subsequently decline, becoming small around the 90th–99th percentile group and negative at the top. Capital income contributes little for most tax filers and becomes increasingly negative toward the upper tail. Taken together, these estimates suggest that legislated tax changes affect household income not only through the mechanical change in tax liabilities but also through indirect labor-market responses. In particular, the absence of disproportionately large gains at the top is difficult to reconcile with a purely mechanical tax-liability channel, under which the groups paying the largest absolute tax liabilities would generally receive the largest gains from a tax reduction.

**Figure 13: Heterogeneous income responses to a cumulative expansionary fiscal shock**



Notes: The figure reports income-group-specific event-study panel local-projection estimates at horizon  $h = 2$ , following Equation 2 and replacing the contemporaneous shock with the cumulative shock defined in Equation 3. Panel (a) reports real total after-tax income, panel (b) real total labor income, and panel (c) real total capital income. For each component, the dependent variable is the change between  $t - 1$  and  $t + 2$ , scaled by total after-tax income in  $t - 1$  and expressed in percentage points. The solid line plots the sign-reversed coefficient, so that positive values correspond to a cumulative expansionary fiscal adjustment equal to 1% of GDP. The underlying shock is the narrative deficit-driven fiscal consolidation shock; positive realizations of the original series denote fiscal tightening. Tax-filing units are assigned to their initial income group and remain in that group throughout the analysis. The first nine groups correspond to income deciles, followed by the 90th–99th percentiles and the top 1%. Estimates use household sampling weights and include household fixed effects, region fixed effects, a linear time trend, and two lags of the corresponding income component. The dark-gray shaded area reports the 68% confidence interval, while the light-gray shaded area reports the 95% confidence interval. Standard errors are heteroskedasticity robust, two-way clustered at the household and year levels, and incorporate a Newey–West correction with the horizon-specific bandwidth. Outcome changes are winsorized at the 5th and 95th percentiles. The sample covers 1999–2016.

**Figure 14: Heterogeneous income responses to a cumulative expansionary tax shock**



Notes: The figure reports income-group-specific event-study panel local-projection estimates at horizon  $h = 2$ , following Equation 2 and replacing the contemporaneous shock with the cumulative shock defined in Equation 3. Panel (a) reports real total after-tax income, panel (b) real total labor income, and panel (c) real total capital income. For each component, the dependent variable is the change between  $t - 1$  and  $t + 2$ , scaled by total after-tax income in  $t - 1$  and expressed in percentage points. The solid line plots the sign-reversed coefficient, so that positive values correspond to a cumulative expansionary tax adjustment equal to 1% of GDP. The underlying shock is the narrative legislated tax shock from Gil et al. (2019); positive realizations of the original series denote tax increases. Tax-filing units are assigned to their initial income group and remain in that group throughout the analysis. The first nine groups correspond to income deciles, followed by the 90th–99th percentiles and the top 1%. Estimates use household sampling weights and include household fixed effects, region fixed effects, a linear time trend, and two lags of the corresponding income component. The dark-gray shaded area reports the 68% confidence interval, while the light-gray shaded area reports the 95% confidence interval. Standard errors are heteroskedasticity robust, two-way clustered at the household and year levels, and incorporate a Newey–West correction with the horizon-specific bandwidth. Outcome changes are winsorized at the 5th and 95th percentiles. The sample covers 1999–2016.

The cumulative specifications sharpen the distributional result. In Figure 13, the response of total after-tax income declines almost monotonically with initial income. The

gain is approximately four percentage points in the lowest groups, falls gradually to around three percentage points in the middle of the distribution, and is smaller among the upper deciles. The estimate then falls sharply for the top 1%. The resulting inverted S-shaped profile implies that the income gains generated by a cumulative fiscal expansion are concentrated among low-income tax filers and become progressively weaker as initial income rises.

Labor income remains the principal positive component. Its response is sizeable throughout the bottom 99% and comparatively stable across groups, indicating that stronger employment, hours, or earnings constitute the main transmission mechanism. The steeper gradient in total after-tax income therefore cannot be attributed exclusively to labor income. It also reflects the direct incidence of taxes and other income components that enter the transition from labor and capital income to disposable income. Capital-income responses are modest and relatively flat across the bottom 99%, before turning sharply negative and highly imprecise in the top percentile.

The cumulative tax shock in [Figure 14](#) produces the same qualitative gradient, but with smaller differences across most of the distribution. Total after-tax income rises by close to three percentage points for the lowest-income tax filers and declines gradually toward approximately one-and-a-half to two percentage points in the upper deciles. As with the fiscal shock, the estimate becomes negative for the top percentile. Labor income is positive throughout the bottom 99%, but the response is smaller than under the broader fiscal expansion. Capital income remains close to zero except at the very top, where the point estimate is negative and surrounded by a wide confidence interval.

The contrast between the continuous and cumulative specifications is economically informative. The contemporaneous innovation captures the response to the fiscal action initiated in a particular year and generates relatively similar gains across most tax filers. Fiscal packages, however, are frequently phased in and combine measures implemented over several years. Once these subsequent measures are incorporated, the incidence becomes distinctly progressive: lower-income tax filers benefit more, while gains for upper-income groups are smaller or absent.

Several mechanisms can rationalize this pattern. Lower-income households are more exposed to cyclical employment risk and tend to work in sectors in which labor demand responds strongly to aggregate expenditure, producing a sharper impact of public spending. This can be amplified due to additional aggregate demand expansion due to their higher marginal propensity to consume. In contrast, higher-income households have more stable labor-market outcomes and receive a larger share of their resources from capital income, which shows no positive response to fiscal expansions in our estimates.

The comparison between deficit-driven and tax shocks is also revealing. The fiscal shock moves several policy tools at once, and its effects reach households mainly through the labor market. The tax shock moves tax bills and incentives only, and it redistributes less. This matches the aggregate and IV evidence. Tax changes affect disposable income immediately while deficit adjustments also move labor demand and transfers, which is

why the bottom of the distribution gains the most.

The results for the top 1% should be interpreted with caution. Confidence intervals widen substantially in this group, particularly for capital income. This reflects both the concentration of capital income at the top and the volatility of dividends, financial returns, business income, and realized capital gains. The conclusion is not the precise sign of the top-percentile coefficient, but rather the absence of evidence that fiscal expansions generate income gains at the top comparable to those observed throughout the lower and middle parts of the distribution. Overall, the inverted S-shaped response in the cumulative specifications is a central distributional result: expansionary fiscal policy raises income most strongly among low-income tax filers, primarily through labor income, while producing progressively smaller effects toward the top.

Two robustness exercises reported in [Appendix H](#) address the concentration of consolidation episodes in the crisis period, which is the natural concern for any design identified from Spanish fiscal variation in this sample. First, [Figure H1](#) re-estimates the distributional profile at  $h = 2$  excluding the financial-crisis years (2007–2010) and, separately, the sovereign-debt-crisis years (2010–2013). Excluding the financial-crisis years leaves the fiscal profile essentially unchanged. Excluding the debt-crisis years removes the largest consolidation episodes and therefore most of the identifying variation, so the estimates become noisier. The point estimates nevertheless preserve—and, if anything, steepen—the progressive gradient, with the largest gains at the bottom of the distribution and a negative point estimate for the top group. The distributional pattern is not generated solely by comparing crisis with non-crisis years. Second, [Figure H7](#) separates the deficit-driven consolidation measures into their spending-based and tax-based components. Both components display the largest income gains at the bottom of the distribution, so the progressive incidence is not an artifact of one side of the budget. The two components diverge mainly at the top, where the tax-based component drives the negative response of the highest-income group, consistent with the direct tax-liability channel discussed above.

## 5.4 Inequality

We translate the heterogeneous responses into implications for income inequality among tax filers. We first compute a suite of widely used inequality measures from the actual after-tax incomes of the tax units in our 2014 sample, a post-crisis year late in the sample period. We then simulate the two-year impact of an expansionary fiscal shock by scaling each unit's income by its income-group-specific estimated response at  $h = 2$ , and we recompute the inequality measures on the resulting counterfactual distribution. We consider three shocks: an expansionary deficit shock of 1% of GDP, the same shock measured with the alternative consolidation series, and an expansionary tax shock of 1% of GDP. [Table 1](#) reports the simulations based on the continuous-shock estimates and [Table 2](#) those based on the cumulative-shock estimates. Two features of the exercise should be kept in mind. First, it is a partial-equilibrium accounting device: it applies the estimated group-specific responses holding group composition, re-ranking, and the financing of the fiscal expan-

sion fixed, and it refers to the population of tax filers described in [Section 4](#). Second, the reported changes carry the estimation uncertainty of the underlying group-specific coefficients, which we do not propagate; the results should therefore be read as point-estimate translations of the distributional responses rather than as formal tests.

Under the continuous specification, the equalizing effect of a fiscal expansion operates mainly through the top of the distribution. The Gini coefficient falls from 0.2772 to 0.2745 after the deficit shock, to 0.2753 with the alternative consolidation series, and to 0.2739 after the tax shock. The top income shares decline in all three cases: the top 1% measure falls from 1.327 to between 1.191 and 1.233, and the top 10% measure from 10.685 to between 10.343 and 10.513. Measures of overall dispersion also compress modestly—the standard deviation of log income and the 90th-to-10th and 90th-to-50th percentile ratios all decline—whereas the 50th-to-10th percentile ratio rises slightly. This configuration is exactly what the sideways J-shaped profile of [Figure 11](#) implies: gains that are roughly uniform across the bottom 99%, combined with a negative response at the very top, reduce top concentration while leaving the relative position of the bottom and the middle nearly unchanged. It is also noteworthy that, under the continuous estimates, the tax shock compresses the summary measures at least as much as the deficit shock, reflecting the sharply negative top-percentile response to tax changes documented above.

The simulations based on the cumulative specifications reflect a complementary picture in which the compression extends through the entire distribution. The Gini falls from 0.2772 to between 0.2749 and 0.2756; the top 1% and top 10% measures decline in all three scenarios; and, in contrast to the continuous case, every percentile ratio falls, including the 50th-to-10th ratio. This is the counterpart of the inverted S-shaped profile of [Figure 13](#): the multi-year consolidations of the sample period fell hardest on lower-income tax filers, so, under an expansion, the bottom of the distribution catches up with both the middle and the top. Under the cumulative estimates the deficit shock is the most equalizing of the three scenarios, consistent with the steeper distributional gradient of the broad consolidation measure relative to the legislated tax changes.

Taken together, the two sets of simulations deliver a consistent message with a meaningful nuance. Expansionary fiscal policy reduces income inequality among tax filers in every scenario we consider. The margin along which it does so, however, depends on the treatment concept: the initial fiscal innovation compresses the distribution mainly from the top, whereas the implemented multi-year adjustment also compresses it from below, which is the sense in which fiscal expansions are progressive. The comparison between the deficit and tax scenarios mirrors the heterogeneity results: legislated tax changes redistribute mainly by lowering top incomes, while broad deficit-driven expansions additionally raise the relative income of the bottom of the distribution.

**Table 1:** Counterfactual fiscal policy implications for income inequality measures

|                                  | ↑ Deficit shock |                    | ↑ Deficit shock<br>(alternative) | ↓ Tax shock         | ↑ Deficit shock | ↑ Deficit shock<br>(alternative) | ↓ Tax shock   |
|----------------------------------|-----------------|--------------------|----------------------------------|---------------------|-----------------|----------------------------------|---------------|
|                                  | Initial value   | Two years after 1% | Two years after 1%               | Two years after −1% | Change (×100)   | Change (×100)                    | Change (×100) |
| Gini coefficient                 | 0.2772          | 0.2745             | 0.2753                           | 0.2739              | -0.2709         | -0.1922                          | -0.3301       |
| Top 1% income share              | 1.327           | 1.191              | 1.229                            | 1.233               | -13.5216        | -9.7864                          | -9.3543       |
| Top 10% income share             | 10.685          | 10.448             | 10.513                           | 10.343              | -23.7248        | -17.2523                         | -34.1654      |
| Standard deviation of log income | 0.280           | 0.278              | 0.279                            | 0.279               | -0.2266         | -0.1545                          | -0.1876       |
| Ratio of 90th to 10th percentile | 3.406           | 3.390              | 3.396                            | 3.397               | -1.5650         | -0.9897                          | -0.8382       |
| Ratio of 90th to 50th percentile | 1.707           | 1.691              | 1.695                            | 1.683               | -1.6216         | -1.1771                          | -2.3558       |
| Ratio of 50th to 10th percentile | 1.995           | 2.005              | 2.003                            | 2.018               | 0.9884          | 0.8019                           | 2.2949        |

*Notes:* The table reports several standard measures of income inequality calculated using actual after-tax incomes in 2014 for all tax-filing units in the sample. The second column reports the initial value of each inequality measure. The third and fourth columns report the corresponding counterfactual values two years after an expansionary deficit shock equal to 1% of GDP, using the baseline and alternative deficit shock series, respectively. The fifth column reports the counterfactual value two years after an expansionary tax shock equal to −1% of GDP. The three rightmost columns report the corresponding changes in each inequality measure, calculated as the difference between the counterfactual and initial values multiplied by 100.

**Table 2:** Counterfactual implications of cumulative fiscal policy shocks for income inequality measures

|                                  | ↑ Deficit shock |                    | ↑ Deficit shock<br>(alternative) | ↓ Tax shock         | ↑ Deficit shock | ↑ Deficit shock<br>(alternative) | ↓ Tax shock   |
|----------------------------------|-----------------|--------------------|----------------------------------|---------------------|-----------------|----------------------------------|---------------|
|                                  | Initial value   | Two years after 1% | Two years after 1%               | Two years after −1% | Change (×100)   | Change (×100)                    | Change (×100) |
| Gini coefficient                 | 0.2772          | 0.2749             | 0.2756                           | 0.2753              | -0.2253         | -0.1587                          | -0.1944       |
| Top 1% income share              | 1.327           | 1.261              | 1.292                            | 1.293               | -6.6188         | -3.5104                          | -3.3429       |
| Top 10% income share             | 10.685          | 10.587             | 10.629                           | 10.616              | -9.8420         | -5.6273                          | -6.9181       |
| Standard deviation of log income | 0.280           | 0.277              | 0.278                            | 0.277               | -0.3256         | -0.2458                          | -0.2969       |
| Ratio of 90th to 10th percentile | 3.406           | 3.375              | 3.383                            | 3.378               | -3.0586         | -2.2968                          | -2.7669       |
| Ratio of 90th to 50th percentile | 1.707           | 1.695              | 1.699                            | 1.696               | -1.1665         | -0.8156                          | -1.0391       |
| Ratio of 50th to 10th percentile | 1.995           | 1.991              | 1.991                            | 1.991               | -0.4311         | -0.3940                          | -0.4088       |

*Notes:* The table reports several standard measures of income inequality calculated using actual after-tax incomes in 2014 for all tax-filing units in the sample. The second column reports the initial value of each inequality measure. The third and fourth columns report the corresponding counterfactual values two years after a cumulative expansionary deficit shock equal to 1% of GDP, using the baseline and alternative deficit shock series, respectively. The fifth column reports the counterfactual value two years after a cumulative expansionary tax shock equal to −1% of GDP. The three rightmost columns report the corresponding changes in each inequality measure, calculated as the difference between the counterfactual and initial values multiplied by 100.

## 5.5 Income-Component Decomposition

[Table 3](#) reports the decomposition of the two-year income response to a cumulative expansionary deficit shock of 1% of GDP across income groups, and [Table 4](#) reports the analogous decomposition for the cumulative expansionary tax shock. In both tables, Panel A expresses the contribution of each income component as a share of the group-specific response of total pre-tax income,  $100 \times \beta_g^{k,2} / \beta_g^{PT,2}$ , so that total labor income and total capital income sum to 100% within each group. Panel B reports the estimated coefficients in percentage points for total pre-tax income, taxes, and total after-tax income, which links the pre-tax decomposition to the after-tax responses discussed in the previous subsections. The top 1% is omitted because its estimated responses are noisy and yield unstable contribution ratios. This decomposition identifies which income sources drive the heterogeneous effects of fiscal policy.

**Table 3: Income-component contributions to the cumulative fiscal-shock response**

|                                                                       | Income group |              |              |              |              |              |              |              |              |              |
|-----------------------------------------------------------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
|                                                                       | 0–10         | 10–20        | 20–30        | 30–40        | 40–50        | 50–60        | 60–70        | 70–80        | 80–90        | 90–99        |
| <i>Panel A: Contribution to the total pre-tax income response (%)</i> |              |              |              |              |              |              |              |              |              |              |
| <b>Total labor income</b>                                             | <b>83.74</b> | <b>85.76</b> | <b>88.65</b> | <b>89.28</b> | <b>88.30</b> | <b>84.76</b> | <b>84.69</b> | <b>82.62</b> | <b>86.23</b> | <b>75.20</b> |
| Wage income                                                           | 83.89        | 86.12        | 88.63        | 89.33        | 88.33        | 84.77        | 84.75        | 82.54        | 85.74        | 73.86        |
| Self-employment income                                                | -0.15        | -0.36        | 0.02         | -0.05        | -0.02        | -0.005       | -0.06        | 0.08         | 0.49         | 1.34         |
| <b>Total capital income</b>                                           | <b>16.26</b> | <b>14.24</b> | <b>11.35</b> | <b>10.72</b> | <b>11.70</b> | <b>15.24</b> | <b>15.31</b> | <b>17.38</b> | <b>13.77</b> | <b>24.80</b> |
| Capital income                                                        | 0.03         | 0.10         | 0.06         | 0.08         | 0.10         | 0.23         | 0.56         | 0.27         | 0.38         | 1.32         |
| Finance income                                                        | 16.36        | 13.89        | 11.05        | 10.11        | 11.25        | 14.61        | 13.97        | 16.32        | 12.88        | 26.04        |
| Capital gains income                                                  | -0.13        | 0.25         | 0.23         | 0.53         | 0.34         | 0.39         | 0.78         | 0.79         | 0.51         | -2.56        |
| <i>Panel B: Sum of estimated coefficients (p.p.)</i>                  |              |              |              |              |              |              |              |              |              |              |
| <b>Total pre-tax income</b>                                           | <b>2.77</b>  | <b>2.83</b>  | <b>3.21</b>  | <b>3.48</b>  | <b>3.35</b>  | <b>2.86</b>  | <b>2.87</b>  | <b>2.86</b>  | <b>3.20</b>  | <b>2.70</b>  |
| Taxes (–)                                                             | -0.46        | -0.37        | -0.04        | 0.20         | 0.30         | 0.22         | 0.37         | 0.49         | 0.73         | 0.39         |
| <b>Total after-tax income</b>                                         | <b>3.23</b>  | <b>3.20</b>  | <b>3.25</b>  | <b>3.28</b>  | <b>3.05</b>  | <b>2.64</b>  | <b>2.50</b>  | <b>2.37</b>  | <b>2.46</b>  | <b>2.31</b>  |

*Notes:* The table reports the decomposition of the two-year income response to a cumulative expansionary deficit shock equal to 1% of GDP. Panel A reports the contribution of each income component to the group-specific response of total pre-tax income, calculated as  $100 \times \beta_g^{k,2} / \beta_g^{PT,2}$ , where  $\beta_g^{PT,2}$  denotes the estimated response of total pre-tax income. Total labor income and total capital income therefore sum to 100% within each income group. Panel B reports the estimated coefficients for total pre-tax income, taxes, and total after-tax income. The tax coefficient is reported using the sign with which taxes enter the construction of after-tax income. The top 1% income group (99–100) is omitted because its estimated responses are highly noisy and yield unstable contribution ratios.

**Table 4: Income-component contributions to the cumulative tax-shock response**

|                                                                       | Income group |              |              |              |              |              |              |              |              |              |
|-----------------------------------------------------------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
|                                                                       | 0–10         | 10–20        | 20–30        | 30–40        | 40–50        | 50–60        | 60–70        | 70–80        | 80–90        | 90–99        |
| <i>Panel A: Contribution to the total pre-tax income response (%)</i> |              |              |              |              |              |              |              |              |              |              |
| <b>Total labor income</b>                                             | <b>78.86</b> | <b>81.44</b> | <b>85.82</b> | <b>86.48</b> | <b>84.98</b> | <b>80.03</b> | <b>81.30</b> | <b>79.68</b> | <b>83.47</b> | <b>74.83</b> |
| Wage income                                                           | 78.99        | 81.30        | 85.73        | 86.49        | 84.99        | 79.98        | 81.33        | 79.34        | 82.37        | 73.21        |
| Self-employment income                                                | -0.13        | 0.14         | 0.09         | -0.01        | -0.01        | 0.05         | -0.03        | 0.34         | 1.10         | 1.62         |
| <b>Total capital income</b>                                           | <b>21.14</b> | <b>18.56</b> | <b>14.18</b> | <b>13.52</b> | <b>15.02</b> | <b>19.97</b> | <b>18.70</b> | <b>20.32</b> | <b>16.53</b> | <b>25.17</b> |
| Capital income                                                        | 0.001        | -0.002       | 0.04         | -0.09        | 0.02         | -0.04        | 0.34         | -0.02        | 0.19         | 1.06         |
| Finance income                                                        | 21.15        | 18.08        | 13.75        | 12.96        | 14.70        | 19.66        | 17.81        | 20.58        | 18.00        | 33.35        |
| Capital gains income                                                  | -0.01        | 0.48         | 0.38         | 0.65         | 0.30         | 0.35         | 0.55         | -0.23        | -1.66        | -9.24        |
| <i>Panel B: Sum of estimated coefficients (p.p.)</i>                  |              |              |              |              |              |              |              |              |              |              |
| <b>Total pre-tax income</b>                                           | <b>1.63</b>  | <b>1.71</b>  | <b>2.01</b>  | <b>2.14</b>  | <b>2.02</b>  | <b>1.66</b>  | <b>1.69</b>  | <b>1.73</b>  | <b>1.79</b>  | <b>1.80</b>  |
| Taxes (–)                                                             | -0.49        | -0.47        | -0.22        | -0.05        | 0.04         | 0.00         | 0.15         | 0.27         | 0.41         | 0.33         |
| <b>Total after-tax income</b>                                         | <b>2.12</b>  | <b>2.18</b>  | <b>2.23</b>  | <b>2.18</b>  | <b>1.97</b>  | <b>1.66</b>  | <b>1.54</b>  | <b>1.47</b>  | <b>1.38</b>  | <b>1.47</b>  |

*Notes:* The table reports the decomposition of the two-year income response to a cumulative expansionary tax shock equal to –1% of GDP. Panel A reports the contribution of each income component to the group-specific response of total pre-tax income, calculated as  $100 \times \beta_g^{k,2} / \beta_g^{PT,2}$ , where  $\beta_g^{PT,2}$  denotes the estimated response of total pre-tax income. Total labor income and total capital income therefore sum to 100% within each income group. Panel B reports the estimated coefficients for total pre-tax income, taxes, and total after-tax income. The tax coefficient is reported using the sign with which taxes enter the construction of after-tax income. The top 1% income group (99–100) is omitted because its estimated responses are highly noisy and yield unstable contribution ratios.

Total labor income explains the bulk of the pre-tax income response throughout the distribution. Under the cumulative deficit shock, its contribution ranges between roughly 83% and 89% from the bottom decile through the 80th–90th percentile group, and remains at 75% for the 90th–99th percentiles. Within labor income, wage income accounts for essentially the entire contribution, while self-employment income is close to zero for most

groups and becomes modestly positive only toward the top (1.3% in the 90–99 group). This confirms, at the level of income sources, the central role of wage employment in the transmission of fiscal shocks to household income.

Capital income accounts for the remainder, between 11% and 25% of the pre-tax response depending on the group, and its composition is informative. Almost all of the capital-income contribution comes from financial income, whose share is largest at the two ends of the distribution (16% in the bottom decile and 26% in the 90–99 group under the deficit shock). Real-estate capital income contributes close to nothing everywhere, and realized capital gains contribute little for most groups and negatively for the 90–99 group—markedly so under the tax shock (−9.2%), consistent with the sensitivity of realizations to changes in the taxation of savings income. The heterogeneous effects of fiscal policy therefore do not operate through a broad capital-income channel: outside financial income, capital components are essentially unresponsive.

Panel B reveals a result that the pre-tax decomposition alone would miss: the progressivity of the after-tax response is partly produced by the tax system itself. The response of total pre-tax income is relatively flat across groups (between 2.7 and 3.5 percentage points under the deficit shock), whereas the response of total after-tax income declines steadily from 3.2 points in the bottom deciles to 2.3 points in the 90–99 group. The wedge is the tax response: taxes fall for the bottom groups following the expansion (−0.46 points in the bottom decile), amplifying their disposable-income gains, and rise toward the top (+0.73 points in the 80–90 group), dampening them. A common fiscal impulse thus translates into a progressive after-tax incidence partly because Spain’s progressive personal income tax converts similar pre-tax gains into dissimilar disposable-income gains, reinforcing the labor-market channel documented above.

The decomposition of the cumulative tax shock in [Table 4](#) tells the same story on a smaller scale. Labor income contributes between 75% and 86% of the pre-tax response, financial income accounts for nearly all of the capital contribution, and the tax response again turns from negative at the bottom (−0.49 points) to positive at the top (+0.33 points), producing an after-tax gradient that declines from 2.1–2.2 points in the bottom deciles to about 1.4–1.5 points in the top groups. Overall, the decompositions show that assessments of the distributional consequences of fiscal policy must account not only for who receives the income gains but also for the composition of those gains and for the mediating role of the progressive tax schedule.

[Figure H8](#) in [Appendix H](#) separates these two sources of heterogeneity. It compares the actual group-specific responses with counterfactual responses that shut down within-component heterogeneity, so that differences across groups reflect income composition alone. Below the top, the two profiles have similar shapes: the composition of income accounts for part of the pattern. At the top 1%, they diverge. Composition alone would predict a sizeable positive labor-income response and a small negative capital-income response, while the actual estimates are much weaker for labor income and, under the tax shock, far more negative for capital income. The behavior of the top reflects how its income

components respond, not only the composition of its income.

## 6 Conclusion

This paper leverages a rich administrative dataset of Spanish tax filers observed from 1999 to 2016 to analyze the heterogeneous effects of fiscal policy on income. Merging detailed income and tax records with three externally identified narrative fiscal shock series—two deficit-driven consolidation measures and one legislated tax-shock measure—we employ reduced-form and instrumental-variable local projections to trace the dynamic impact of fiscal policy on total, labor, and capital income across the income distribution. In doing so, the paper brings fiscal policy into the administrative-microdata literature on the distributional effects of macroeconomic shocks, which has so far focused on monetary policy ([Amberg et al., 2022](#); [Andersen et al., 2023](#); [Coibion et al., 2017](#); [Holm et al., 2021](#)).

At the aggregate level, fiscal consolidations generate sizeable and persistent declines in the income of tax filers, driven by labor income while capital income reacts modestly and only in the short run. Expressed per unit of realized fiscal adjustment through the instrumental-variable exercise, legislated tax changes reduce disposable income relatively quickly, whereas deficit-driven adjustments build more gradually and end up having the larger and more persistent effect. This is consistent with the fiscal-multiplier literature, which finds stronger and long-lasting effects for spending-based than for tax-based fiscal policy ([Gechert, 2015](#); [Gechert & Rannenberg, 2018](#)).

Across the distribution, the incidence of fiscal policy is progressive. The initial fiscal innovation raises income across the bottom 99% of tax filers, with no evidence of gains at the top. Once the multi-year implementation of consolidation packages is taken into account, the response takes an inverted S-shape, with the largest gains among lower-income tax filers and progressively smaller effects along the distribution. Labor income is the dominant transmission channel: its response is sizeable and relatively flat across the bottom 99%. Its weight in the total response, however, falls with income, because higher income filers obtain a larger share of their income from other sources. The tax shock shifts this composition toward capital income, mainly financial income, so deficit-driven adjustments operate relatively more through labor income and legislated tax changes relatively more through capital income. Counterfactual simulations translate these responses into inequality outcomes. Expansionary fiscal policy reduces income inequality among tax filers in every scenario we consider. Tax shocks redistribute mainly by lowering top incomes, while deficit-driven expansions also raise the relative income of the bottom. These patterns are robust to the alternative narrative consolidation series and to the attrition-free household panel.

Three features of our setting limit the generalizability of our results. First, the identifying variation comes mainly from the consolidations of 2010–2013, implemented during a recession by a member of a currency union without independent monetary policy. If fiscal effects are larger in a crisis than in normal times, our magnitudes are best read as

estimates of the incidence of consolidation under stress. Second, Spain’s dual labor market, in which employment adjustment is concentrated among temporary contracts held by lower-income workers, is a country-specific amplifier of the results that we document. Finally, the estimates cannot explain anything about households that do not file tax returns. However, what we expect to generalize is the qualitative mechanism: fiscal consolidation falls more on labor income at the bottom and middle of the distribution than on capital income at the top, and the progressivity of the tax system shapes the after-tax incidence.

The findings carry a policy message about the composition of fiscal adjustment. Because deficit-driven consolidations load their income costs disproportionately onto the labor income of lower-income tax filers, while legislated tax measures fall relatively more on top incomes, the distributional cost of a consolidation of a given size depends on how the package is designed—a margin over which governments have direct control. Our spending-versus-tax decomposition and the comparison between deficit and tax shocks provide direct evidence on this margin.

## References

- Abadie, A., Athey, S., Imbens, G. W., & Wooldridge, J. M. (2023). When should you adjust standard errors for clustering? *The Quarterly Journal of Economics*, 138(1), 1–35.
- Alesina, A., Azzalini, G., Favero, C., Giavazzi, F., & Miano, A. (2018). Is it the “how” or the “when” that matters in fiscal adjustments? *IMF Economic Review*, 66(1), 144–188. doi: 10.1057/s41308-017-0047-z
- Alesina, A., Barbiero, O., Favero, C., Giavazzi, F., & Paradisi, M. (2015). Austerity in 2009–13. *Economic Policy*, 30(83), 383–437.
- Alesina, A., Favero, C., & Giavazzi, F. (2015). The output effect of fiscal consolidation plans. *Journal of International Economics*, 96, S19–S42.
- Almunia, M., & Lopez-Rodriguez, D. (2019). The elasticity of taxable income in Spain: 1999–2014. *SERIEs*, 10(3), 281–320.
- Almuzara, M., & Sancibrián, V. (2024). Micro responses to macro shocks. *FRB of New York Staff Report*(1090).
- Altavilla, C., Brugnolini, L., Gürkaynak, R. S., Motto, R., & Ragusa, G. (2019). Measuring euro area monetary policy. *Journal of Monetary Economics*, 108, 162–179.
- Amberg, N., Jansson, T., Klein, M., & Picco, A. R. (2022). Five facts about the distributional income effects of monetary policy shocks. *American Economic Review: Insights*, 4(3), 289–304.
- Andersen, A. L., Johannesen, N., Jørgensen, M., & Peydró, J.-L. (2023). Monetary policy and inequality. *The Journal of Finance*, 78(5), 2945–2989.
- Bauer, M. D., & Swanson, E. T. (2023). A reassessment of monetary policy surprises and high-frequency identification. *NBER Macroeconomics Annual*, 37(1), 87–155.
- Bertrand, M., Duflo, E., & Mullainathan, S. (2004). How much should we trust differences-in-differences estimates? *The Quarterly Journal of Economics*, 119(1), 249–275.

- Blanchet, T., & Martínez-Toledano, C. (2023). Wealth inequality dynamics in Europe and the United States: Understanding the determinants. *Journal of Monetary Economics*, 133, 25–43.
- Blanco, M. A., Bauluz, L., & Martínez-Toledano, C. (2021). Wealth in Spain 1900–2017 a country of two lands. *The Economic Journal*, 131(633), 129–155.
- Buda, G., Carvalho, V. M., Corsetti, G., Duarte, J. B., Hansen, S., Moura, A. S., ... Alves da Silva, G. (2025). The short lags of monetary policy. *Faculty of Economics, University of Cambridge*.
- Callaway, B., Goodman-Bacon, A., & Sant’Anna, P. H. (2024). Event studies with a continuous treatment. In *American Economic Association: Papers and Proceedings* (Vol. 114, pp. 601–605).
- Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2011). Robust inference with multiway clustering. *Journal of Business & Economic Statistics*, 29(2), 238–249.
- Carrière-Swallow, Y., David, A. C., & Leigh, D. (2021). Macroeconomic Effects of Fiscal Consolidation in Emerging Economies: New Narrative Evidence from Latin America and the Caribbean. *Journal of Money, Credit and Banking*, 53(6), 1313–1335. doi: 10.1111/jmcb.12784
- Cloyne, J., Ferreira, C., & Surico, P. (2020). Monetary policy when households have debt: New evidence on the transmission mechanism. *The Review of Economic Studies*, 87(1), 102–129.
- Cloyne, J., & Surico, P. (2017). Household debt and the dynamic effects of income tax changes. *The Review of Economic Studies*, 84(1), 45–81. doi: 10.1093/restud/rdw021
- Coibion, O., Gorodnichenko, Y., Kueng, L., & Silvia, J. (2017). Innocent Bystanders? Monetary policy and inequality. *Journal of Monetary Economics*, 88, 70–89.
- Dhital, S., Jiang, S., & Reese, J. (2023). Effects of monetary and government spending policy on economic inequality. *Journal of Macroeconomics*, 77, 103547.
- Drechsel, T., Lewis, D., Melcangi, D., & Pilossoph, L. (2025). *The investment channel of monetary policy: Disentangling firm heterogeneity*. (Mimeo, University of Maryland)
- Driscoll, J. C., & Kraay, A. C. (1998). Consistent covariance matrix estimation with spatially dependent panel data. *Review of Economics and Statistics*, 80(4), 549–560.
- Dube, A., Girardi, D., Jordà, Ò., & Taylor, A. M. (2025). A local projections approach to difference-in-differences. *Journal of Applied Econometrics*, 40, 741–758. doi: 10.1002/jae.70000
- Durante, E., Ferrando, A., & Vermeulen, P. (2022). Monetary policy, investment and firm heterogeneity. *European Economic Review*, 148, 104251.
- Dynan, K. (2012). Is a household debt overhang holding back consumption? *Brookings Papers on Economic Activity*, 299–344.
- Fernández Romero, D. (2025). The fiscal multiplier in presence of unconventional monetary policy: Evidence for 17 OECD countries. *Economic Modelling*, 147, 107063.
- Ferriere, A., & Navarro, G. (2025). The heterogeneous effects of government spending: It’s all about taxes. *The Review of Economic Studies*, 92(2), 1061–1125. doi: 10.1093/

- Furceri, D., Ge, J., Loungani, P., & Melina, G. (2022). The distributional effects of government spending shocks in developing economies. *Review of Development Economics*, 26(3), 1574–1599.
- Gechert, S. (2015). What fiscal policy is most effective? A meta-regression analysis. *Oxford Economic Papers*, 67(3), 553–580.
- Gechert, S., & Heimberger, P. (2022). Do corporate tax cuts boost economic growth? *European Economic Review*, 147, 104157.
- Gechert, S., & Rannenberg, A. (2018). Which fiscal multipliers are regime-dependent? A meta-regression analysis. *Journal of Economic Surveys*, 32(4), 1160–1182.
- Gertler, M., & Karadi, P. (2015). Monetary policy surprises, credit costs, and economic activity. *American Economic Journal: Macroeconomics*, 7(1), 44–76.
- Giavazzi, F., & McMahon, M. (2013). The household effects of government spending. In A. Alesina & F. Giavazzi (Eds.), *Fiscal policy after the financial crisis* (pp. 103–141). Chicago: University of Chicago Press.
- Gil, P., Martí, F., Morris, R., Pérez, J. J., & Ramos, R. (2019). The output effects of tax changes: narrative evidence from Spain. *SERIEs*, 10, 1–23.
- González, I., Sala, H., & Trivín, P. (2026). Consumption responses to a major minimum wage increase: evidence from Spain. *Journal of Public Economics*, 260, 105710.
- Guajardo, J., Leigh, D., & Pescatori, A. (2014). Expansionary austerity? International evidence. *Journal of the European Economic Association*, 12(4), 949–968.
- Gupta, S., Jalles, J. T., Mulas-Granados, C., & Schena, M. (2018). Planned fiscal adjustments: Do governments fulfil their commitments? *European Union Politics*, 19(3), 383–407.
- Heimberger, P. (2020). The dynamic effects of fiscal consolidation episodes on income inequality: Evidence for 17 OECD countries over 1978–2013. *Empirica*, 47(1), 53–81. doi: 10.1007/s10663-018-9404-z
- Holm, M. B., Paul, P., & Tischbirek, A. (2021). The transmission of monetary policy under the microscope. *Journal of Political Economy*, 129(10), 2861–2904.
- Hope, D., & Limberg, J. (2022). The economic consequences of major tax cuts for the rich. *Socio-Economic Review*, 20(2), 539–559.
- Isiaka, A., Mihailov, A., & Razzu, G. (2025). Fiscal policy and inequality in middle-and high-income countries: redistributive effects of tax and spending shocks. *Macroeconomic Dynamics*, 29, e12.
- Jarociński, M., & Karadi, P. (2020). Deconstructing monetary policy surprises—the role of information shocks. *American Economic Journal: Macroeconomics*, 12(2), 1–43.
- Johnson, D. S., Parker, J. A., & Souleles, N. S. (2006). Household expenditure and the income tax rebates of 2001. *American Economic Review*, 96(5), 1589–1610.
- Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1), 161–182.
- Jordà, Ò., & Taylor, A. M. (2025). Local projections. *Journal of Economic Literature*, 63(1), 59–110.

- Kaplan, G., & Violante, G. L. (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica*, 82(4), 1199–1239.
- Klein Martins, G. (2025). Long-run effects of austerity: An analysis of size dependence and persistence in fiscal multipliers. *Oxford Bulletin of Economics and Statistics*, 87(2), 330–356.
- Komatsuzaki, T. (2024). Fiscal Consolidation and Income Inequality in Latin America and the Caribbean. *Economía LACEA Journal*, 23(1), 74–88. doi: 10.31389/eco.437
- Leigh, M. D., Pescatori, M. A., Devries, M. P., & Guajardo, M. J. (2011). A new action-based dataset of fiscal consolidation. *International Monetary Fund*.
- Mertens, K., & Montiel Olea, J. L. (2018). Marginal tax rates and income: New time series evidence. *The Quarterly Journal of Economics*, 133(4), 1803–1884.
- Mian, A. R., & Sufi, A. (2010). Household leverage and the recession of 2007 to 2009. *National Bureau of Economic Research*.
- Newey, W. K., & West, K. D. (1986). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *National Bureau of Economic Research Cambridge, Mass., USA*.
- Parker, J. A. (1999). The reaction of household consumption to predictable changes in social security taxes. *American Economic Review*, 89(4), 959–973.
- Pérez-López, C., Villanueva García, J., & Molinero-Muñoz, I. (2020). Panel de declarantes de irpf 1999-2016: metodología, estructura y variables. *Documentos-Instituto de Estudios Fiscales*(8), 1.
- Ramey, V. A. (2011). Identifying government spending shocks: It's all in the timing. *The Quarterly Journal of Economics*, 126(1), 1–50.
- Ramey, V. A. (2016). Macroeconomic shocks and their propagation. *Handbook of macroeconomics*, 2, 71–162.
- Ramey, V. A., & Zubairy, S. (2018). Government spending multipliers in good times and in bad: evidence from US historical data. *Journal of political economy*, 126(2), 850–901.
- Romer, C. D., & Romer, D. H. (2010). The macroeconomic effects of tax changes: estimates based on a new measure of fiscal shocks. *American Economic Review*, 100(3), 763–801.
- Souleles, N. S. (2002). Consumer response to the reagan tax cuts. *Journal of Public Economics*, 85(1), 99–120.
- Teulings, C. N., & Zubanov, N. (2014). Is Economic Recovery a Myth? Robust Estimation of Impulse Responses. *Journal of Applied Econometrics*, 29(3), 497–514. doi: 10.1002/jae.2333
- Zidar, O. (2019). Tax cuts for whom? Heterogeneous effects of income tax changes on growth and employment. *Journal of Political Economy*, 127(3), 1437–1472.

**Declaration of generative AI and AI-assisted technologies in the writing process**

During the preparation of this work the authors used ChatGPT (GPT-5.6 Thinking) and Claude Opus 4.8 in order to proofread the text. After using these tools, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

## Appendix A Monetary policy shock aggregation


A natural complement to our analysis of fiscal shocks would be to estimate the distributional effects of monetary policy using high-frequency identified surprises. This appendix documents why we do not pursue that route. The reason is mechanical rather than conceptual: income and inequality outcomes are measured at annual frequency (tax records, income surveys), so any monetary shock series must be aggregated at least to the year. We show that once the leading high-frequency shock measure is aggregated—to months, quarters, and especially years—it no longer predicts the interest rates that households and firms actually face. An “annualized” monetary policy shock is therefore not a valid measure of monetary policy, and regressions of annual inequality outcomes on such a shock would be uninterpretable.

We use the monetary policy shock of [Jarociński and Karadi \(2020\)](#) as we can see in [Figure A1](#), updated through 2025, which isolates unanticipated policy from central-bank information using the “poor man’s” sign restriction: an announcement-window increase in the one-year OIS rate accompanied by a decline in the EURO STOXX 50 is classified as a contractionary monetary policy shock ( $mp_{pm}$ ), while co-movement in the same direction is attributed to an information shock. The underlying asset-price surprises come from the Euro Area Monetary Policy Event-Study Database of [Altavilla, Brugnolini, Gürkaynak, Motto, and Ragusa \(2019\)](#), measured in narrow intraday windows around ECB announcements. As the market policy rate we take the 12-month EURIBOR<sup>11</sup>, the reference rate for the overwhelming majority of adjustable-rate mortgages in Spain and the rate most directly relevant for household finances. The sample runs from January 1999 to April 2025. We aggregate the shock by summing within each month, quarter, and year, and we take period averages of the EURIBOR, examining both its level and its change.

---

<sup>11</sup>Monthly averages, obtained from Bank of Spain series called be1901.


**Figure A1:** Monetary policy shocks from [Jarociński and Karadi \(2020\)](#)



*Notes:* This figure shows the change monetary policy poor man's sign restriction variable through time. The y-axis is in basis points. The panel (a) plots “mp pm” shocks monthly while panel (b) plots “mp pm” aggregated by year. The surprises are extracted from [Jarociński and Karadi \(2020\)](#) shocks.

At the frequency at which the shock is constructed, its connection with interest rates is tight. In the raw [Jarociński and Karadi \(2020\)](#) data, the correlation between the monetary policy shock and the high-frequency OIS surprise in the announcement window (the first principal component of OIS changes in the monetary event window) is 0.78 ( $p < 0.001$ ,  $N = 271$  announcement months). [Figure A2](#) shows why: whenever the sign restriction attributes the window surprise to monetary policy, the shock is the OIS surprise (the points on the diagonal); the remaining announcements are classified as information shocks and set to zero (the points on the horizontal axis). The failure documented below is therefore not a defect of the shock itself—it emerges only through aggregation.

**Figure A2:** At high frequency the shock tracks the OIS surprise



Notes: Each point is a month with an ECB announcement (1999–2024,  $N = 271$ ). The horizontal axis is the high-frequency OIS surprise in the announcement window (first principal component of OIS changes in the monetary event window, in percentage points); the vertical axis is the [Jarociński and Karadi \(2020\)](#) monetary policy shock ( $mp_{pm}$ ). The line is the OLS fit; the correlation is 0.78. Points on the diagonal are announcements classified as monetary policy shocks; points on the horizontal axis are announcements classified as central-bank information shocks. Source: raw data accompanying [Jarociński and Karadi \(2020\)](#), 2024 update.

[Table A1](#) reports the correlation between the aggregated shock and the EURIBOR at each frequency. The correlations with the EURIBOR level are indistinguishable from zero at every frequency (between 0.01 and 0.07, with  $p$ -values above 0.7). The correlations with the change in the EURIBOR are, if anything, negative:  $-0.09$  at monthly,  $-0.20$  at quarterly (significant at the 5% level), and  $-0.18$  at yearly frequency. Taken at face value, the point estimates say that years with more contractionary monetary surprises are years in which interest rates fall.


**Table A1:** Correlation between the aggregated monetary policy shock and the EURIBOR

| Frequency | EURIBOR level |            | $\Delta$ EURIBOR |            | $N$ |
|-----------|---------------|------------|------------------|------------|-----|
|           | $\rho$        | $p$ -value | $\rho$           | $p$ -value |     |
| Monthly   | 0.014         | 0.802      | $-0.087$         | 0.126      | 316 |
| Quarterly | 0.013         | 0.893      | $-0.200$         | 0.049      | 106 |
| Yearly    | 0.073         | 0.718      | $-0.183$         | 0.387      | 27  |

Notes: Pearson correlations between the [Jarociński and Karadi \(2020\)](#) monetary policy shock ( $mp_{pm}$ , summed within each period) and the 12-month EURIBOR (period average of monthly data, Banco de España series be1901), in levels and in first differences. Sample: 1999m1–2025m4.  $p$ -values from the exact  $t$ -test for zero correlation;  $N$  refers to the level specification (first-differencing costs one observation).


Figure A3 shows the relationship at the frequency that matters for inequality analysis: the year. Each point is a year; the fitted line slopes the wrong way for changes and is flat for levels. The year labels make the problem concrete. The largest cumulative contractionary surprise of the sample (2008) was followed by collapsing rates, while the largest rate increase on record (2023) occurred in a year whose summed surprise component was negative—the 2022–2023 tightening was, by then, almost fully anticipated by markets. Figure A4 and Figure A5 show that the same pattern holds at quarterly and monthly frequency, so the failure is not an artifact of the small number of annual observations.

**Figure A3: Yearly aggregation: monetary policy shock vs. EURIBOR**




Notes: Each point is a year (1999–2025). The horizontal axis is the Jarociński and Karadi (2020) monetary policy shock summed within the year; the vertical axis is the yearly average of the 12-month EURIBOR (Panel A) and its year-on-year change in percentage points (Panel B). The line is the OLS fit. The corresponding correlation coefficients,  $p$ -values, and sample sizes are reported in Table A1.

**Figure A4: Quarterly aggregation: monetary policy shock vs. EURIBOR**



Notes: Each point is a quarter (1999q1–2025q2). The horizontal axis is the Jarociński and Karadi (2020) monetary policy shock summed within the quarter; the vertical axis is the quarterly average of the 12-month EURIBOR (Panel A) and its quarter-on-quarter change in percentage points (Panel B). The line is the OLS fit.

**Figure A5: Monthly frequency: monetary policy shock vs. EURIBOR**



Notes: Each point is a month (1999m1–2025m4). The horizontal axis is the [Jarociński and Karadi \(2020\)](#) monetary policy shock. The vertical axis is the monthly average of the 12-month EURIBOR (Panel A) and its monthly change in percentage points (Panel B). The line is the OLS fit.

Three features of high-frequency identification explain this pattern. First, the shock is a 30-minute asset-price surprise, and forward-looking market rates such as the EURIBOR price it in within the same day. A monthly, quarterly, or yearly change in the EURIBOR is dominated by everything else that moved markets over the period, such as macroeconomic news, risk premia, and the anticipated component of policy. Second, the surprises are not distributed randomly over the cycle: contractionary surprises cumulated when the ECB repeatedly cut rates by less than markets expected while rates were falling. Summing over a month, a quarter, or a year does not average this pattern away but accumulates it. The aggregated series inherits the state-dependent component of monetary policy and is correlated with contemporaneous economic conditions, which questions the exogeneity of the shock. Following [Ramey \(2016\)](#), a shock series must be exogenous with respect to current and lagged endogenous variables and must represent unanticipated movements; the high-frequency surprise satisfies both requirements within the announcement window, but its low-frequency aggregate satisfies neither. This is also why [Buda et al. \(2025\)](#) argue that the natural frequency at which to trace out the effects of high-frequency monetary surprises is daily: time aggregation mixes the surprise with the endogenous policy path and blurs the transmission one is trying to identify. Third, the resulting weak-instrument problem at aggregated frequencies is well documented: [Gertler and Karadi \(2015\)](#) embed the surprises in a proxy SVAR on the VAR innovations precisely because regressing observed rates on aggregated surprises does not deliver a first stage, and [Bauer and Swanson \(2023\)](#) show that high-frequency surprises lose most of their predictive content for the policy stance once they are cumulated over time.

However, it is quite common in the empirical literature to regress macroeconomic time series and firm- or household-level outcomes on monthly, quarterly, or annually aggregated high-frequency surprises. For instance, [Durante, Ferrando, and Vermeulen \(2022\)](#) merge the 12-month moving sum of the [Jarociński and Karadi \(2020\)](#) shock to annual firm balance-sheet data. [Jarociński and Karadi \(2020\)](#) themselves aggregate the surprises to the monthly level for both the euro area and the United States in their macro analysis. [Cloyne](#)

et al. (2020) aggregate to the quarterly level with Compustat firm data, Drechsel, Lewis, Melcangi, and Pilossoph (2025) do the same for the United States with firm-level data, and Amberg et al. (2022) aggregate to the yearly level with Swedish administrative tax records. That practice assumes that the sum of exogenous innovations remains exogenous. The evidence above casts doubt on that assumption. The aggregated series no longer tracks the policy rates, and the correlation can even turn negative. This is consistent with contamination by the systematic component of policy.


To inspect the mechanism, we simulate a stylized economy and show that the same shock series can be correlated with the policy interest rate at daily frequency and uncorrelated (or with a “wrong sign” correlation) when the shock is aggregated. The model has the following elements:

(i) A persistent business cycle  $c_t$  follows a daily AR(1) with autoregressive coefficient 0.999. (ii) The central bank meets every 31 trading days ( $\approx 8$  times per year) and moves rates systematically with the cycle: the anticipated policy move at announcement  $j$  is  $a_j = \beta c_{t_j}$  with  $\beta = 12.5$  basis points, and markets price it in gradually over the inter-meeting period. (iii) The surprise on the announcement date is  $\varepsilon_j = \eta_j - \lambda a_j$  with  $\eta_j \sim N(0, 2^2)$  basis points and  $\lambda = 0.025$ : the central bank systematically moves rates slightly less than markets expect, so contractionary surprises cumulate in easing cycles, replicating the pattern observed in the data. (iv) The daily market-rate change adds non-policy news  $u_t \sim N(0, 1)$ .

We simulate 27 years of daily data (matching our sample length) 500 times and compute, in each replication, the correlation between the shock and the interest-rate change at daily, monthly, quarterly, and yearly frequency, summing the shock within each period as in the empirical literature.

Figure A6 and Table A2 report the results. On announcement days the surprise and the daily rate change are almost the same object: the mean correlation is 0.88. Aggregation makes this relationship null monotonically: 0.12 at monthly,  $-0.03$  at quarterly, and  $-0.24$  at yearly frequency—despite the fact that, by construction, every single surprise moves the interest rate one-for-one on impact. The yearly correlation is negative for the same reason as in the data: at low frequencies the systematic component dominates the variance of rate changes, and the surprises are negatively correlated with it. Note that the simulated decay profile (0.88, 0.12,  $-0.03$ ,  $-0.24$ ) closely mirrors the empirical one, even though the simulation was not fitted to those moments: the simulated announcement-day correlation of 0.88 matches the 0.78 correlation between the shock and the OIS surprise in the raw Jarociński and Karadi (2020) data, and the simulated aggregated correlations track the estimates in Table A1 ( $-0.09$ ,  $-0.20$ ,  $-0.18$  for monthly, quarterly, yearly changes). It is enough that the central bank moves rates 2.5 percent less than markets expect ( $\lambda = 0.025$ ) to reproduce the correlations observed in the data.

**Figure A6:** Simulation: correlation between the monetary surprise and the interest-rate change, by aggregation frequency



*Notes:* Monte Carlo simulation of the daily model described in the text (500 replications of 27 years of daily data). Diamonds are the mean correlation between the monetary surprise (summed within each period) and the change in the market interest rate over the same period; capped bars are the 5th–95th percentile range across replications. “Daily” refers to the correlation between the surprise and the same-day rate change on announcement days.

**Table A2:** Simulation: correlation between the aggregated surprise and the interest-rate change

| Frequency | Mean correlation | 5th pctile | 95th pctile |
|-----------|------------------|------------|-------------|
| Daily     | 0.881            | 0.853      | 0.906       |
| Monthly   | 0.121            | −0.009     | 0.244       |
| Quarterly | −0.025           | −0.231     | 0.161       |
| Yearly    | −0.242           | −0.575     | 0.095       |

*Notes:* Mean and 5th–95th percentile range of the correlation across 500 Monte Carlo replications of the daily model described in the text. Within each replication the surprise is summed within the period and correlated with the interest-rate change over the same period.

The implication for our setting is immediate. With annual inequality outcomes, the only feasible monetary regressor is the annualized shock. The annualized shock contains essentially no information about the interest rates agents actually faced. Its correlation with the rate change even has the wrong sign. As the simulation shows, this is not a small-sample accident but a structural consequence of aggregating a surprise that co-moves with the systematic policy stance. Any estimated “effect of monetary policy on inequality” built on it would be uninterpretable, whether the shock is used directly as a regressor (it is no longer exogenous) or as an instrument (it fails the relevance condition outright). Fiscal shocks, by contrast, are legislated, dated, and directly measurable at annual frequency,


which is why this paper focuses on them.

## Appendix B Alternative fiscal shock

As a robustness exercise, this appendix estimates the aggregate and distributional responses using an alternative narrative fiscal consolidation shock obtained directly from [Leigh et al. \(2011\)](#) and [Gupta et al. \(2018\)](#). We consider both the contemporaneous and cumulative versions of the shock, applying the same specifications, outcome definitions, and income-group classification used in the main analysis.


The results confirm the main findings. Fiscal tightening reduces total after-tax income, primarily through labor income, while the response of capital income is comparatively limited. The distributional estimates also retain the main pattern documented in the paper, with expansionary fiscal policy producing larger income gains toward the lower part of the distribution and smaller responses at higher income levels.

**Figure B1:** Response of income components to a continuous alternative fiscal shock



*Notes:* The figure reports event-study panel local-projection estimates of the response of income to the continuous alternative fiscal consolidation shock (from [Gupta et al. \(2018\)](#); [Leigh et al. \(2011\)](#)), following [Equation 1](#). Panel (a) reports real total after-tax income, panel (b) real total labor income, and panel (c) real total capital income. For each horizon  $h$ , the dependent variable is the change in the corresponding income component between  $t - 1$  and  $t + h$ , scaled by total after-tax income in  $t - 1$  and expressed in percentage points. The omitted horizon is  $h = -1$ . The fiscal shock is the narrative deficit-driven consolidation shock, measured as a percentage of GDP; positive values denote fiscal tightening. Estimates are weighted using the household sampling factor and include household fixed effects, initial income-group fixed effects, region fixed effects, a linear time trend, and two lags of the corresponding income component. The dark-gray shaded area reports the 68% confidence interval, while the light-gray shaded area reports the 95% confidence interval. Standard errors are heteroskedasticity robust, two-way clustered at the household and year levels, and incorporate a Driscoll-Kraay correction implemented through a Newey-West correction at each specific horizon bandwidth. Outcome changes are winsorized at the 5th and 95th percentiles.


**Figure B2: Response of income components to a cumulative alternative fiscal shock**



*Notes:* The figure reports event-study panel local-projection estimates of the response of income to the cumulative alternative fiscal consolidation shock (from [Gupta et al. \(2018\)](#); [Leigh et al. \(2011\)](#)), following [Equation 1](#) and replacing the contemporaneous shock with the cumulative shock defined in [Equation 3](#). Panel (a) reports real total after-tax income, panel (b) real total labor income, and panel (c) real total capital income. For each horizon  $h$ , the dependent variable is the change in the corresponding income component between  $t - 1$  and  $t + h$ , scaled by total after-tax income in  $t - 1$  and expressed in percentage points. The omitted horizon is  $h = -1$ . The cumulative fiscal shock is the sum of narrative deficit-driven consolidation shocks between year  $t$  and horizon  $h$ , measured as a percentage of GDP; positive values denote fiscal tightening. Estimates are weighted using the household sampling factor and include household fixed effects, initial income-group fixed effects, region fixed effects, a linear time trend, and two lags of the corresponding income component. The dark-gray shaded area reports the 68% confidence interval, while the light-gray shaded area reports the 95% confidence interval. Standard errors are heteroskedasticity robust, two-way clustered at the household and year levels, and incorporate a Driscoll-Kraay correction implemented through a Newey-West correction at each specific horizon bandwidth. Outcome changes are winsorized at the 5th and 95th percentiles.


[Figure B1](#) and [Figure B2](#) show that the aggregate results are robust to the alternative fiscal shock. Total after-tax income declines following fiscal tightening, with the response driven mainly by labor income, whereas capital income reacts more weakly.

**Figure B3: Expansionary 1% of GDP alternative fiscal shock effect on income**



*Notes:* the figure reports the estimated income responses at horizon  $h = 2$  to an expansionary alternative fiscal shock equal to 1% of GDP. The alternative narrative deficit-driven consolidation shock is obtained from [Leigh et al. \(2011\)](#) and [Gupta et al. \(2018\)](#). The dark-gray shaded area reports the 68% confidence interval, while the light-gray shaded area reports the 95% confidence interval. The sample covers the period 1999–2016.

**Figure B4:** Cumulative expansionary 1% of GDP alternative fiscal shock effect on income



*Notes:* the figure reports the estimated income responses at horizon  $h = 2$  to a cumulative expansionary alternative fiscal shock equal to 1% of GDP. The alternative narrative deficit-driven consolidation shock is obtained from [Leigh et al. \(2011\)](#) and [Gupta et al. \(2018\)](#). The dark-gray shaded area reports the 68% confidence interval, while the light-gray shaded area reports the 95% confidence interval. The sample covers the period 1999–2016.

The heterogeneous responses in [Figure B3](#) and [Figure B4](#) likewise reproduce the main distributional pattern. Income gains are generally larger among lower-income households and become more moderate toward the upper part of the distribution.

## Appendix C Aggregate Effects


This appendix examines whether the tax consolidation shock generates aggregate dynamics that are consistent with the macro-fiscal literature. The exercise is useful as a diagnostic, but it should not be interpreted as the main source of evidence in the paper. The reason is that the aggregate proxy-SVAR is a much more demanding environment for identification than our micro-level design. Once the analysis moves from administrative household records to a small quarterly macroeconomic system, the rich cross-sectional variation in the data is necessarily collapsed into a few aggregate series. At the same time, the specification introduces macroeconomic variables that are jointly determined with fiscal policy and with each other. In such a setting, conditioning on output, government spending, interest rates, their lags, and the lags of the fiscal instrument absorbs a large part of the variation that the narrative shock is meant to isolate.

This trade-off is important for interpreting the estimates. In the administrative data, the identifying variation is used to trace how the same aggregate fiscal shock is transmitted differently across households, income groups, and income sources. In the aggregate VAR, by contrast, the shock must explain movements in a small set of highly endogenous macroeconomic variables over a limited time-series sample. As a result, the first stage is substantially weaker and estimation uncertainty increases. This is reflected in the first-stage F-statistic, equal to 3.46, which indicates that the proxy-SVAR estimates should be read with caution.

We estimate a quarterly proxy-SVAR following [Gil et al. \(2019\)](#). The system includes three endogenous variables: the log of real GDP, the log of real government spending, and the short-term interest rate. The external instrument is the narrative tax consolidation shock. The specification includes a linear time trend, and three lags of the endogenous variables. The impulse responses are computed up to 50 quarters and normalized to a tax consolidation shock of 1% of GDP. Confidence bands are constructed at the 68% and 90% levels using Newey-West standard errors.

[Figure C1](#) reports the results. The point estimates are qualitatively consistent with the standard fiscal transmission mechanism: a positive tax consolidation shock is followed by a decline in GDP, with the largest effect occurring at medium horizons. Government spending initially increases but subsequently falls, suggesting that the tax consolidation shock is accompanied by broader fiscal adjustment over time. However, the confidence bands are wide, especially at the 90% level, and the weak first stage prevents us from drawing strong conclusions from the aggregate VAR alone. The exercise therefore provides suggestive aggregate evidence, but the core contribution of the paper comes from exploiting the administrative micro data to uncover the heterogeneous transmission of fiscal policy across the income distribution.

**Figure C1: Aggregate responses to a tax consolidation shock**



*Notes:* The figure reports impulse responses from a quarterly proxy-SVAR to a positive tax consolidation shock normalized to 1% of GDP. The panels show the responses of real GDP and real government spending. Both variables are expressed in logarithms and reported in percent deviations from baseline. The VAR also includes the three-month interest rate, expressed in percentage points. The specification includes a linear time trend, and three lags of the endogenous variables. The dark grey and light grey shaded areas denote 68% and 90% confidence intervals, respectively, computed using Newey-West standard errors. The sample is 1986Q1–2016Q4. The external instrument is the tax consolidation shock obtained from [Gil et al. \(2019\)](#). The first-stage F-statistic is 3.46, indicating that the aggregate estimates should be interpreted as suggestive rather than conclusive evidence.

## Appendix D Sample construction and variable definitions

This appendix describes the construction of the estimation sample from the *Pure Panel of Personal Income Tax Filers 1999-2016*. The raw administrative files contain tax-return identifiers, main-filer identifiers, household identifiers, spouse or second-filer identifiers, sampling weights, filing status, demographic information, municipality and region of residence, income-source strata, and the income and tax variables reported in the personal income tax return.

The starting point is the pure panel of tax returns observed between 1999 and 2016. Monetary variables are recorded in cents of euros in the raw files and are converted to euros. Nominal amounts are then expressed in real terms using the annual Spanish GDP deflator, normalized to 2015. We use the sampling weights in all estimations.

Our preferred income concept is real total after-tax income. Total pre-tax income is constructed as the sum of five income components:

$$Y_{it}^{pre} = Y_{it}^{lab} + Y_{it}^{bus} + Y_{it}^{re} + Y_{it}^{fin} + Y_{it}^{cg},$$

where  $Y_{it}^{lab}$  denotes labor income,  $Y_{it}^{bus}$  business income,  $Y_{it}^{re}$  real-estate capital income,  $Y_{it}^{fin}$  financial income, and  $Y_{it}^{cg}$  realized capital gains. Total after-tax income is then defined as:

$$Y_{it}^T = Y_{it}^{pre} - Tax_{it},$$

where  $Tax_{it}$  is final personal income tax liability after the application of allowances, deductions, credits, withholdings, and advance payments. In the empirical analysis, all income components are expressed as changes relative to lagged real total after-tax income, so that component responses add up to the response of total after-tax income.

We retain observations with positive real total after-tax income, positive real total pre-tax income, positive real labor income, and positive real wage income. We remove duplicated tax returns using the administrative inclusion indicator supplied in the panel. We also exclude observations with implausible birth-date information and construct age from the main filer's year of birth. The main sample is restricted to filers aged 26–65. As discussed in the main text, this restriction focuses the analysis on the population most exposed to the labor-market and tax-base channels of fiscal policy. Robustness checks include older filers and therefore allow pension-income and transfer-indexation channels to enter the response.

Because the tax filing unit may change over time, we keep a stable filing-status definition over the response window. Specifically, we exclude tax units that switch between individual and joint filing between the base year and the horizon used to measure the income response. This avoids attributing to fiscal shocks income changes mechanically generated by changes in the composition of the tax unit.

The empirical model estimates heterogeneous responses by income percentile group. To construct these groups, we rank filers within each calendar year using a pre-treatment

measure of income. For each filer  $i$  and year  $t$ , we compute:

$$\bar{Y}_{i,t}^{pre} = \frac{1}{m_{it}} \sum_{s \in \{t-4, t-3, t-2\}} Y_{is}^{pre},$$

where  $m_{it}$  is the number of available observations in the interval  $[t - 4, t - 2]$ . The year  $t - 1$  is not used for ranking because it enters the denominator of the response variable, and post-treatment income is never used to assign income groups. We require the filer to be observed in  $t - 1$ , in  $t + h$ , and in at least one year between  $t - 4$  and  $t - 2$ . Within each year, filers are assigned to eleven groups: deciles 1–9, the p90–p99 group, and the top 1 percent.

This ranking procedure serves two purposes. First, it aligns the sample construction with the model in [Section 3](#), where the coefficients of interest are income-group-specific impulse responses. Second, it reduces mechanical re-ranking caused by transitory income fluctuations, such as one-off capital gains, irregular labor payments, or temporary business-income changes.

For each horizon  $h$ , the main dependent variable is:

$$\Delta_h y_{it}^k = 100 \times \frac{Y_{i,t+h}^k - Y_{i,t-1}^k}{Y_{i,t-1}^T},$$

where  $Y^k$  denotes income component  $k$  and  $Y^T$  denotes total after-tax income. This normalization expresses all component responses as a percentage of pre-shock after-tax income and permits an additive decomposition of the total-income response into labor, business, capital, and tax-related components.

Finally, we limit the influence of extreme income-component movements. We remove observations in which the absolute ratio of a component to total after-tax income is greater than or equal to five. We also winsorize the tails of the growth-rate distributions for the main income variables at the 5th and 95th percentiles. These restrictions are applied before merging the annual fiscal shock series by calendar year.

## Appendix E Descriptive statistics by income group

To characterize the household groups used in the distributional analysis, [Table E1](#) reports weighted descriptive statistics by position in the income distribution. The table presents both the average level of the main income aggregates and the composition of pre-tax income across labor and capital income sources.

**Table E1:** Descriptive statistics by income group

| Income group                                                  | Average real income |                  |              |                | Demographic characteristics |       |
|---------------------------------------------------------------|---------------------|------------------|--------------|----------------|-----------------------------|-------|
|                                                               | Pre-tax income      | After-tax income | Labor income | Capital income | Male share                  | Age   |
| <i>Panel A: Income levels and demographic characteristics</i> |                     |                  |              |                |                             |       |
| 0–10                                                          | 15,374.20           | 14,519.77        | 14,077.03    | 883.08         | 0.66                        | 44.68 |
| 10–20                                                         | 22,567.41           | 20,400.70        | 20,873.87    | 1,082.39       | 0.68                        | 45.62 |
| 20–30                                                         | 27,846.50           | 24,383.85        | 25,927.85    | 1,263.37       | 0.67                        | 46.51 |
| 30–40                                                         | 32,674.93           | 27,863.76        | 30,503.53    | 1,545.94       | 0.65                        | 47.67 |
| 40–50                                                         | 37,102.17           | 30,939.96        | 34,791.47    | 1,746.05       | 0.63                        | 49.37 |
| 50–60                                                         | 41,360.18           | 33,998.19        | 38,503.16    | 2,248.92       | 0.67                        | 49.93 |
| 60–70                                                         | 47,399.65           | 38,150.84        | 43,919.29    | 2,826.08       | 0.72                        | 49.83 |
| 70–80                                                         | 57,149.87           | 44,572.01        | 52,540.71    | 3,987.11       | 0.76                        | 50.13 |
| 80–90                                                         | 76,424.46           | 56,703.80        | 69,443.52    | 6,467.23       | 0.81                        | 51.27 |
| 90–99                                                         | 129,005.40          | 89,447.21        | 109,028.27   | 20,495.06      | 0.87                        | 52.82 |
| 99–100                                                        | 460,095.94          | 296,996.97       | 331,083.22   | 144,666.84     | 0.89                        | 53.55 |

| Income group                                             | Labor income |          |       | Capital income |           |               | Taxes |           |
|----------------------------------------------------------|--------------|----------|-------|----------------|-----------|---------------|-------|-----------|
|                                                          | Wage         | Business | Total | Capital        | Financial | Capital gains | Total | Tax share |
| <i>Panel B: Income composition (% of pre-tax income)</i> |              |          |       |                |           |               |       |           |
| 0–10                                                     | 87.45        | 4.39     | 91.84 | 1.16           | 2.73      | 1.28          | 5.17  | 8.23      |
| 10–20                                                    | 90.58        | 2.30     | 92.88 | 1.15           | 1.96      | 1.05          | 4.16  | 13.17     |
| 20–30                                                    | 91.94        | 1.51     | 93.46 | 1.00           | 1.92      | 1.04          | 3.96  | 16.15     |
| 30–40                                                    | 92.66        | 1.10     | 93.76 | 1.04           | 1.99      | 1.07          | 4.10  | 18.30     |
| 40–50                                                    | 93.16        | 0.94     | 94.10 | 1.04           | 2.09      | 1.06          | 4.20  | 19.92     |
| 50–60                                                    | 92.49        | 1.00     | 93.50 | 1.18           | 2.45      | 1.21          | 4.85  | 20.93     |
| 60–70                                                    | 91.65        | 1.28     | 92.93 | 1.39           | 2.75      | 1.37          | 5.51  | 22.44     |
| 70–80                                                    | 90.26        | 1.92     | 92.18 | 1.61           | 3.16      | 1.70          | 6.47  | 24.58     |
| 80–90                                                    | 88.12        | 2.84     | 90.96 | 1.81           | 4.13      | 2.09          | 8.02  | 27.77     |
| 90–99                                                    | 81.14        | 3.92     | 85.06 | 2.32           | 7.67      | 4.31          | 14.29 | 30.31     |
| 99–100                                                   | 65.56        | 4.82     | 70.38 | 2.40           | 16.47     | 11.42         | 30.29 | 31.17     |

*Notes:* The table reports weighted means calculated from the 1999–2016 administrative panel, using the household sampling factor. Income groups correspond to the groups employed in the distributional analysis. All monetary variables are expressed in real euros. Panel A reports mean income levels and demographic characteristics. The male share is the weighted mean of an indicator equal to one for a male principal filer and zero for a female principal filer. Age refers to the age of the principal filer. In Panel B, income shares are first calculated at the household-year level as the relevant income component divided by total pre-tax income and multiplied by 100, and are subsequently averaged within income groups using the sampling weights. They are therefore means of household-level shares rather than ratios of group-level mean incomes. Total labor income is the sum of wage and business income, while total capital income comprises capital income, financial income, and capital gains. The tax share is calculated relative to pre-tax income, with positive values denoting taxes paid.

Average income increases monotonically across the distribution, with a particularly pronounced difference between the 90–99 group and the top 1%. Mean pre-tax income rises from approximately 15,400 euros in the bottom decile to 129,000 euros in the 90–99 group and 460,000 euros in the top 1%. The gap between pre-tax and after-tax income also widens with income. Accordingly, the average tax share increases from 8.2% in the bottom


decile to slightly above 30% in the two highest-income groups, reflecting the progressivity of the personal income tax system.

Labor income is the dominant source of income throughout most of the distribution. Its average share exceeds 90% across much of the middle of the distribution, but declines to 85.1% in the 90–99 group and 70.4% in the top 1%. This decline is offset by the increasing importance of capital income. The capital-income share rises from around 4–5% in the lower and middle groups to 14.3% in the 90–99 group and 30.3% in the top 1%. Financial income and capital gains account for most of this increase at the top.

The income groups also differ in their demographic composition. Both the male share and the average age of the principal filer generally increase across the distribution. The male share rises from roughly two-thirds in the lower groups to almost 90% in the top 1%, while mean age increases from approximately 45 to 54 years.

## Appendix F Additional figures


Figure F1: Inequality measures in Spain



Source: OECD data. Aggregate macroeconomics measures of inequality in Spain from 2003 to 2021.


Figures F2 through F13 report the dynamic income-group-specific responses that underlie the  $h = 2$  estimates presented in the main text, for total after-tax income, total labor income, and total capital income, under the continuous and cumulative versions of the fiscal and tax shocks. Across components and shocks, the responses generally reach their largest magnitude around  $h = 2$ , which motivates the focus on that horizon in the main text.

**Figure F2:** Dynamic responses of total after-tax income by income group: continuous fiscal shock




*Notes:* the figure reports income-group-specific event-study local-projection estimates of the response of real total after-tax income to the continuous deficit-driven fiscal consolidation shock across estimation horizons, following Equation 2. Estimates use household sampling weights and the specification and inference procedure described in Section 3. The sample covers 1999–2016.

**Figure F3:** Dynamic responses of total after-tax income by income group: cumulative fiscal shock




*Notes:* the figure reports income-group-specific event-study local-projection estimates of the response of real total after-tax income to the cumulative deficit-driven fiscal consolidation shock defined in Equation 3, following Equation 2. Estimates use household sampling weights and the specification and inference procedure described in Section 3. The sample covers 1999–2016.

**Figure F4:** Dynamic responses of total after-tax income by income group: continuous tax shock




*Notes:* the figure reports income-group-specific event-study local-projection estimates of the response of real total after-tax income to the continuous legislated tax shock of [Gil et al. \(2019\)](#) across estimation horizons, following [Equation 2](#). Estimates use household sampling weights and the specification and inference procedure described in [Section 3](#). The sample covers 1999–2016.

**Figure F5:** Dynamic responses of total after-tax income by income group: cumulative tax shock




*Notes:* the figure reports income-group-specific event-study local-projection estimates of the response of real total after-tax income to the cumulative legislated tax shock defined in [Equation 3](#), following [Equation 2](#). Estimates use household sampling weights and the specification and inference procedure described in [Section 3](#). The sample covers 1999–2016.

**Figure F6:** Dynamic responses of total labor income by income group: continuous fiscal shock




*Notes:* the figure reports income-group-specific event-study local-projection estimates of the response of real total labor income to the continuous deficit-driven fiscal consolidation shock across estimation horizons, following Equation 2. Estimates use household sampling weights and the specification and inference procedure described in Section 3. The sample covers 1999–2016.

**Figure F7:** Dynamic responses of total labor income by income group: cumulative fiscal shock




*Notes:* the figure reports income-group-specific event-study local-projection estimates of the response of real total labor income to the cumulative deficit-driven fiscal consolidation shock defined in Equation 3, following Equation 2. Estimates use household sampling weights and the specification and inference procedure described in Section 3. The sample covers 1999–2016.

**Figure F8:** Dynamic responses of total labor income by income group: continuous tax shock




*Notes:* the figure reports income-group-specific event-study local-projection estimates of the response of real total labor income to the continuous legislated tax shock of [Gil et al. \(2019\)](#) across estimation horizons, following [Equation 2](#). Estimates use household sampling weights and the specification and inference procedure described in [Section 3](#). The sample covers 1999–2016.

**Figure F9:** Dynamic responses of total labor income by income group: cumulative tax shock




*Notes:* the figure reports income-group-specific event-study local-projection estimates of the response of real total labor income to the cumulative legislated tax shock defined in [Equation 3](#), following [Equation 2](#). Estimates use household sampling weights and the specification and inference procedure described in [Section 3](#). The sample covers 1999–2016.

**Figure F10:** Dynamic responses of total capital income by income group: continuous fiscal shock




*Notes:* the figure reports income-group-specific event-study local-projection estimates of the response of real total capital income to the continuous deficit-driven fiscal consolidation shock across estimation horizons, following Equation 2. Estimates use household sampling weights and the specification and inference procedure described in Section 3. Capital-income responses are noisier than labor-income responses, particularly at the top of the distribution, and should be read with the caveats discussed in the main text. The sample covers 1999–2016.

**Figure F11:** Dynamic responses of total capital income by income group: cumulative fiscal shock




*Notes:* the figure reports income-group-specific event-study local-projection estimates of the response of real total capital income to the cumulative deficit-driven fiscal consolidation shock defined in Equation 3, following Equation 2. Estimates use household sampling weights and the specification and inference procedure described in Section 3. Capital-income responses are noisier than labor-income responses, particularly at the top of the distribution, and should be read with the caveats discussed in the main text. The sample covers 1999–2016.

**Figure F12:** Dynamic responses of total capital income by income group: continuous tax shock



*Notes:* the figure reports income-group-specific event-study local-projection estimates of the response of real total capital income to the continuous legislated tax shock of [Gil et al. \(2019\)](#) across estimation horizons, following [Equation 2](#). Estimates use household sampling weights and the specification and inference procedure described in [Section 3](#). Capital-income responses are noisier than labor-income responses, particularly at the top of the distribution, and should be read with the caveats discussed in the main text. The sample covers 1999–2016.

**Figure F13:** Dynamic responses of total capital income by income group: cumulative tax shock



*Notes:* the figure reports income-group-specific event-study local-projection estimates of the response of real total capital income to the cumulative legislated tax shock defined in [Equation 3](#), following [Equation 2](#). Estimates use household sampling weights and the specification and inference procedure described in [Section 3](#). Capital-income responses are noisier than labor-income responses, particularly at the top of the distribution, and should be read with the caveats discussed in the main text. The sample covers 1999–2016.

## Appendix G Income-Group Assignment, Compositional Changes, and Reconciliation with [Amberg et al. \(2022\)](#)

### G.1 Income-group mobility and endogenous reclassification

This appendix examines how the definition of income groups affects the estimated distributional responses to fiscal policy. In particular, we compare the specification used by [Amberg et al. \(2022\)](#), in which tax units are classified according to their moving average position in the income distribution, with specifications that hold income-group membership fixed throughout the sample using as well a moving average. This distinction is important because contemporaneous income-group assignment may itself respond to the fiscal shock endogenously.

To assess the empirical relevance of this issue, [Table G1](#) reports transition probabilities across income groups using all observed transitions in the panel. For each tax unit, its income group in a given observation is compared with its income group in the next available observation. The matrix therefore summarizes income-group mobility over the full 1999–2016 panel.

Income-group membership is highly persistent, particularly at the tails of the distribution. Among tax units observed in the bottom income group, 89.12 percent remain in that group in their next observation. The corresponding persistence rate for the top 1 percent is 85.53 percent. Persistence is lower in the middle of the distribution, where approximately 69–78 percent of tax units remain in the same group. Mobility is concentrated primarily between adjacent groups. Nevertheless, approximately 22–31 percent of tax units in the central groups are classified in a different income group in their next observation.

This mobility matters for the interpretation of heterogeneous estimates. When income groups are defined using contemporaneous moving average income, the composition of each group can change over time. Group-specific coefficients may therefore reflect both income responses among tax units that remain in the same group and changes in the set of tax units assigned to that group. The concern is not that identification is driven exclusively by tax units that change groups. Rather, such movements provide the margin through which contemporaneous assignment can introduce endogenous compositional changes. Holding income-group membership fixed isolates income changes for a stable set of tax units and prevents the fiscal shock from affecting the classification used to estimate heterogeneous responses.

**Table G1: Transitions across income groups**

| Income group in the current observed period | Income group in the subsequent periods |              |              |              |              |              |              |              |              |              |              | Total  |
|---------------------------------------------|----------------------------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------|
|                                             | 0–10                                   | 10–20        | 20–30        | 30–40        | 40–50        | 50–60        | 60–70        | 70–80        | 80–90        | 90–99        | 99–100       |        |
| 0–10                                        | <b>89.12</b>                           | 10.17        | 0.49         | 0.13         | 0.04         | 0.03         | 0.01         | 0.01         | 0.01         | 0.00         | 0.00         | 100.00 |
| 10–20                                       | 8.26                                   | <b>77.55</b> | 12.69        | 1.07         | 0.24         | 0.10         | 0.05         | 0.02         | 0.01         | 0.01         | 0.00         | 100.00 |
| 20–30                                       | 0.66                                   | 10.32        | <b>73.32</b> | 13.64        | 1.37         | 0.43         | 0.18         | 0.06         | 0.02         | 0.01         | 0.00         | 100.00 |
| 30–40                                       | 0.15                                   | 1.33         | 11.05        | <b>70.16</b> | 14.64        | 1.92         | 0.55         | 0.15         | 0.03         | 0.01         | 0.00         | 100.00 |
| 40–50                                       | 0.06                                   | 0.31         | 1.57         | 11.99        | <b>68.72</b> | 15.35        | 1.62         | 0.32         | 0.06         | 0.01         | 0.00         | 100.00 |
| 50–60                                       | 0.03                                   | 0.14         | 0.60         | 2.10         | 12.70        | <b>69.05</b> | 14.22        | 1.01         | 0.12         | 0.02         | 0.00         | 100.00 |
| 60–70                                       | 0.03                                   | 0.07         | 0.23         | 0.80         | 1.84         | 11.44        | <b>72.53</b> | 12.48        | 0.53         | 0.07         | 0.00         | 100.00 |
| 70–80                                       | 0.02                                   | 0.03         | 0.08         | 0.19         | 0.40         | 1.29         | 9.91         | <b>77.17</b> | 10.65        | 0.25         | 0.00         | 100.00 |
| 80–90                                       | 0.01                                   | 0.02         | 0.02         | 0.04         | 0.08         | 0.15         | 0.69         | 8.46         | <b>83.60</b> | 6.93         | 0.01         | 100.00 |
| 90–99                                       | 0.00                                   | 0.00         | 0.01         | 0.01         | 0.02         | 0.03         | 0.08         | 0.33         | 5.94         | <b>91.86</b> | 1.71         | 100.00 |
| 99–100                                      | 0.00                                   | 0.00         | 0.00         | 0.00         | 0.00         | 0.00         | 0.02         | 0.01         | 0.04         | 14.40        | <b>85.53</b> | 100.00 |
| Total                                       | 9.59                                   | 9.94         | 10.03        | 10.07        | 10.12        | 10.10        | 10.09        | 10.08        | 10.07        | 8.96         | 0.96         | 100.00 |

*Notes:* The table reports transition probabilities across income groups using all observed transitions in the 1999–2016 panel. For each tax unit, rows indicate the income group in the current observed period and columns indicate the income group in the subsequent available period. The two observations need not correspond to consecutive calendar years. Entries are expressed as percentages and sum to 100 within each row. Income groups are defined using total after-tax income, with the upper decile divided into the 90th–99th percentiles and the top 1 percent. The main diagonal reports the probability of remaining in the same income group between observed periods. The final row reports the unconditional distribution across destination income groups.

## G.2 Contemporaneous assignment: the [Amberg et al. \(2022\)](#) specification

The baseline specification in [Amberg et al. \(2022\)](#) can be represented as


$$\Delta_h y_{it}^k = \sum_{g=1}^{11} G_{it,g} \left( \alpha_g^{k,h} + \beta_g^{k,h} F_t \right) + \varepsilon_{it}^{k,h}, \quad (8)$$

where  $G_{it,g}$  indicates the income group to which tax unit  $i$  belongs in the shock year  $t$ , and  $F_t$  denotes the fiscal shock. The specification therefore estimates an intercept and a fiscal-shock response for each contemporaneous income group. It does not include tax-unit or year fixed effects.

We first replicate this specification using the continuous fiscal shock. The resulting estimates are reported in [Figure G1](#). The specification produces a distributional gradient in total after-tax income, with larger responses at the bottom of the distribution. However, because group membership is determined from contemporaneous income, these estimates do not trace the response of a fixed set of tax units. Instead, the population represented by each coefficient may change between years.

This creates a potentially important interpretation problem. Suppose, for example, that a fiscal consolidation reduces the income of a tax unit initially located in the 20th–30th percentile group and causes it to move into the 10th–20th percentile group. Under contemporaneous assignment, the tax unit contributes to the estimated response of the lower group even though it belonged to the higher group before the policy intervention. The estimated heterogeneity may therefore reflect both differential income responses and policy-induced changes in the composition of income groups. The specification remains informative about the response of individuals observed at each contemporaneous position in the distribution, but it does not identify the response of fixed ex ante income groups.

**Figure G1: Income responses under contemporaneous income-group assignment**



*Notes:* The figure reports group-specific estimates of the response to an expansionary fiscal shock equal to 1 percent of GDP, obtained from the specification in Equation 8. Income groups are determined by each tax unit's contemporaneous position in the total after-tax income distribution in year  $t$ . Panel (a) reports total after-tax income, panel (b) total labor income, and panel (c) total capital income. Outcomes are measured as the change between  $t - 1$  and  $t + 2$ , scaled by total after-tax income in  $t - 1$ , and expressed in percentage points. The reported coefficients reverse the sign of the original consolidation shock, so positive values represent the response to a fiscal expansion. The dark and light shaded areas denote 68 and 95 percent confidence intervals, respectively. Standard errors are two-way clustered at the tax-unit and year levels. Estimates use the continuous narrative deficit-driven fiscal consolidation shock obtained from Guajardo et al. (2014) and Jordà and Taylor (2025). The sample period is 1999–2016.


### G.3 Holding income-group assignment fixed

We next replace the contemporaneous indicator  $G_{it,g}$  with  $G_{i,g}^0$ , which denotes the tax unit's income group when it is first observed in the panel. This classification is held fixed in all subsequent years. The exercise changes only the assignment of tax units to income groups and therefore isolates the relevance of compositional changes in the Amberg et al. (2022) specification.

As shown in Figure G2, fixing income-group membership produces a clearer distributional profile. Total after-tax income and labor income respond more strongly at the bottom of the initial income distribution, and the responses decline toward the upper groups. The comparison indicates that contemporaneous reclassification tends to attenuate the gradient by moving tax units across groups during the estimation period.

Nevertheless, this specification still does not include tax-unit or year fixed effects. The estimates are therefore absolute group-specific responses and may partly capture persistent differences between tax units assigned to different initial income groups. Fixing group membership addresses endogenous composition, but it does not by itself remove time-invariant individual heterogeneity.

**Figure G2: Income responses holding initial income-group assignment fixed**



*Notes:* The figure reports group-specific estimates of the response to an expansionary fiscal shock equal to 1 percent of GDP. Tax units are assigned to income groups according to their first observed position in the total after-tax income distribution, and this assignment is held fixed throughout the sample. Panel (a) reports total after-tax income, panel (b) total labor income, and panel (c) total capital income. Outcomes are measured as the change between  $t - 1$  and  $t + 2$ , scaled by total after-tax income in  $t - 1$ , and expressed in percentage points. The reported coefficients reverse the sign of the original consolidation shock, so positive values represent the response to a fiscal expansion. The dark and light shaded areas denote 68 and 95 percent confidence intervals, respectively. Standard errors are two-way clustered at the tax-unit and year levels. Estimates use the continuous narrative deficit-driven fiscal consolidation shock obtained from [Guajardo et al. \(2014\)](#) and [Jordà and Taylor \(2025\)](#). The sample period is 1999–2016.

#### G.4 Tax-unit and year fixed effects: relative income responses

Finally, we retain the fixed initial income-group assignment and introduce tax-unit and year fixed effects:


$$\Delta_h y_{it}^k = \alpha_i^{k,h} + \lambda_t^{k,h} + \sum_{g=1}^{10} \beta_g^{k,h} (F_t \times G_{i,g}^0) + u_{it}^{k,h}, \quad (9)$$

where the top 1 percent,  $g = 11$ , is the omitted income group. Tax-unit fixed effects absorb time-invariant differences across filing units, while year fixed effects absorb shocks and macroeconomic conditions that affect all tax units equally in a given year.

Because the fiscal shock varies only over time, its common main effect is absorbed by the year fixed effects. The interaction coefficients in [Equation 9](#) therefore identify differential responses relative to the omitted top 1 percent rather than absolute income responses. Accordingly, the response of the top 1 percent is normalized to zero in [Figure G3](#). This change in the estimand explains the substantial change in scale relative to the previous two figures: the coefficients now measure how much more or less each initial income group responds than the top 1 percent.

The relative estimates show a pronounced gradient. The income response is substantially larger among tax units initially located at the bottom of the distribution than among those in the omitted top group. The pattern is mainly accounted for by labor income, while the relative capital-income profile is flatter. Thus, once group composition is held fixed and time-invariant tax-unit heterogeneity and common year shocks are absorbed, the distributional differences become particularly visible. These coefficients should not, however, be interpreted as the absolute effects of the fiscal shock on each income group.

**Figure G3: Differential income responses relative to the top 1 percent**



*Notes:* The figure reports differential responses to an expansionary fiscal shock equal to 1 percent of GDP. Tax units are assigned to income groups according to their first observed position in the total after-tax income distribution, and this assignment is held fixed throughout the sample. The regressions include tax-unit and year fixed effects. The top 1 percent is the omitted group; its response is therefore normalized to zero, and all reported coefficients measure responses relative to that group. Panel (a) reports total after-tax income, panel (b) total labor income, and panel (c) total capital income. Outcomes are measured as the change between  $t - 1$  and  $t + 2$ , scaled by total after-tax income in  $t - 1$ , and expressed in percentage points. The reported coefficients reverse the sign of the original consolidation shock, so positive values indicate a larger response to a fiscal expansion than in the top 1 percent. The dark and light shaded areas denote 68 and 95 percent confidence intervals, respectively. Standard errors are two-way clustered at the tax-unit and year levels. Estimates use the continuous narrative deficit-driven fiscal consolidation shock based on [Guajardo et al. \(2014\)](#) and [Jordà and Taylor \(2025\)](#). The sample period is 1999–2016.

## G.5 Reconciliation with the baseline estimates

The preceding exercises show that differences across specifications reflect three distinct empirical choices: whether income groups are defined contemporaneously or ex ante, whether tax-unit and year fixed effects are included, and whether the coefficients represent absolute or relative responses. These dimensions should not be conflated. In particular, the larger coefficients in [Figure G3](#) do not imply larger absolute fiscal effects, since they are measured relative to the response of the top 1 percent.

Our main specification differs from [Equation 9](#) by estimating a separate regression for each fixed initial income group. This approach allows us to include tax-unit fixed effects while recovering an absolute response for every group, rather than a response relative to an omitted category. Under this more demanding specification, the heterogeneous pattern generated by the one-year continuous fiscal shock becomes substantially weaker once time-invariant tax-unit heterogeneity is removed.


The response re-emerges, however, when fiscal policy is measured by the cumulative fiscal adjustment. In that case, tax units initially located at the bottom of the income distribution exhibit the largest income gains following a fiscal expansion, with the response declining toward the upper income groups. This suggests that the economically relevant distributional pattern is associated with the cumulative fiscal intervention rather than being mechanically generated by annual reclassification across income groups.

Taken together, the results support two conclusions. First, income groups should be defined using a predetermined measure and held fixed throughout the event window. Allowing contemporaneous reassignment can introduce post-treatment sorting: if fiscal

policy changes a tax unit's income rank, the unit may be assigned to a different group precisely because of the policy whose heterogeneous effect is being estimated. Second, the clearest distributional effects arise from cumulative fiscal adjustment. The response of low-income groups is therefore not primarily a compositional artifact of movements across percentiles, but reflects their greater exposure to the cumulative effects of fiscal policy.


## Appendix H Additional Heterogeneity Results

**Figure H1:** Total after tax income results with and without crisis ( $h = 2$ ).




Notes: panels (a) and (b) present the effects of a 1% of GDP expansionary fiscal shock and a 1% of GDP expansionary tax shock on total after-tax income, as estimated using Equation 8. The estimates are shown both including and excluding the financial crisis years (2007–2010) and the debt crisis years (2010–2013) from the estimation sample. The estimation horizon is  $h = 2$ .

**Figure H2:** Total after tax income results with and without population above 65 ( $h = 2$ ).




Notes: panels (a) and (b) present the effects of a 1% of GDP expansionary fiscal shock and a 1% of GDP expansionary tax shock on total after-tax income, as estimated using Equation 8. The estimates are shown both including and excluding the population above 65 years old from the estimation sample. The estimation horizon is  $h = 2$ .

**Figure H3:** Total after tax income results with and without control variables ( $h = 2$ ).




Notes: panels (a) and (b) present the effects of a 1% of GDP expansionary fiscal shock and a 1% of GDP expansionary tax shock on total after-tax income, as estimated using Equation 8. The estimates are reported both excluding and including additional control variables: two lags of the dependent variable (red line), two lags of the shock (green line), and demographic and structural controls for age, gender, state, and time fixed effects (orange line). The estimation horizon is  $h = 2$ .

**Figure H4:** Total after tax income results by gender ( $h = 2$ ).




Notes: panels (a) and (b) present the effects of a 1% of GDP expansionary fiscal shock and a 1% of GDP expansionary tax shock on total after-tax income, as estimated using Equation 8. The estimates are reported by gender of the main breadwinner. The estimation horizon is  $h = 2$ .

**Figure H5: Total after tax income results by age ( $h = 2$ ).**




Notes: panels (a) and (b) present the effects of a 1% of GDP expansionary fiscal shock and a 1% of GDP expansionary tax shock on total after-tax income, as estimated using Equation 8. The estimates are reported by age group of the main breadwinner. The estimation horizon is  $h = 2$ .

**Figure H6: Total after tax income results by state ( $h = 2$ ).**




Notes: panels (a) and (b) present the effects of a 1% of GDP expansionary fiscal shock and a 1% of GDP expansionary tax shock on total after-tax income, as estimated using Equation 8. The estimates are disaggregated by region: South includes Andalucía, Región de Murcia, and Extremadura; East includes Aragón, Illes Balears, Cataluña, and Comunidad Valenciana; Center includes Castilla-La Mancha, Castilla y León, and Comunidad de Madrid; North includes Principado de Asturias, Cantabria, Galicia, and La Rioja; and Others includes Canarias, Ceuta, and Melilla. The estimation horizon is  $h = 2$ .

**Figure H7:** Total after tax income results differentiating spend vs tax ( $h = 2$ ).



*Notes:* this figure presents the effects of a 1% of GDP expansionary fiscal shock on total after-tax income, as estimated using Equation 8, separating the narrative consolidation measures into their spending-based and tax-based components. The estimation horizon is  $h = 2$ .

**Figure H8: Income Composition versus Within- Component Heterogeneity.**



*Notes:* Panels A and B show the actual effects of a  $-1\%$  of GDP fiscal policy shock and a  $-1\%$  of GDP tax policy shock, respectively, for the estimation horizon  $h = 2$ , as estimated using Equation 8. The red lines depict the responses obtained when Equation 8 is estimated with counterfactual dependent variables that shut down within-component heterogeneity: for each main income component (labor income and capital income), the counterfactual estimates are constructed by first computing the counterfactual responses of the respective subcomponents and then aggregating them following Amberg et al. (2022). Comparing the actual and counterfactual profiles indicates how much of the distributional heterogeneity reflects differences in income composition across groups versus heterogeneous responses within income components.

## Appendix I Attrition robustness using the expanded household panel

This appendix examines whether the aggregate and distributional income responses reported in the main analysis are affected by attrition in the pure panel. For this purpose, we use the expanded IRPF panel from which the pure panel is constructed. We aggregate all tax declarations belonging to the same household and construct household-level measures of total after-tax income, labor income, and capital income. The unit of observation throughout this appendix is the household rather than the individual tax-filing unit.

Unlike the pure panel, which follows the initially selected tax filers for as long as they remain observable, the expanded panel is refreshed over time to preserve its cross-sectional representativeness. The IEF first retains tax filers selected in the preceding year who are observed again and then completes the target sample size within each sampling stratum by randomly drawing from new tax filers entering the population. The sampling strata are defined by regions, predominant source of income, and income bracket (Pérez-López, Villanueva García, & Molinero-Muñoz, 2020). Thus, households entering the panel are not individually matched to households leaving the sample. Instead, replenishment takes place within relatively detailed sampling strata. Although the panel therefore continues to exhibit household entry and exit, its design prevents the progressive loss of cross-sectional representativeness that characterizes the pure panel.


Using the expanded household panel, we first re-estimate the aggregate event-study panel local projections reported in the main text. We then re-estimate the heterogeneous responses at horizon  $h = 2$ , assigning households to their initial position in the household-income distribution and keeping this classification fixed over time. All specifications include household fixed effects, a linear time trend, and two lags of the corresponding household-income component. We maintain the baseline normalization, household sampling weights, treatment of outliers, and inference procedure.

Figures I1–I4 report the aggregate responses. The results are very similar to those obtained with the pure panel. Fiscal consolidations and tax increases reduce total after-tax household income, with the largest effects generally occurring during the first two post-shock years. As in the baseline estimates, the response is driven predominantly by labor income, while capital income displays a smaller and less persistent response. The cumulative-shock specifications also reproduce the more persistent decline documented in the main analysis.

Figures I5–I8 report the corresponding responses across initial household-income groups. The expanded-panel estimates reproduce the broad distributional profile found in the pure panel. Expansionary fiscal adjustments generate larger income gains toward the lower part of the distribution and progressively smaller responses toward the top, primarily through labor income. The estimates for tax shocks display a related, although generally less pronounced, distributional gradient. The similarity in the sign, timing, composition, and distributional incidence of the responses indicates that neither the aggregate


results nor the heterogeneous patterns documented in the main text are driven by attrition or by the declining representativeness of the pure panel.

**Figure I1:** Aggregate household-income responses to a continuous fiscal shock: expanded panel




*Notes:* The figure reports event-study panel local-projection estimates of the household-income response to the continuous fiscal consolidation shock using the expanded IRPF household panel. Panel (a) reports real total after-tax household income, panel (b) real total household labor income, and panel (c) real total household capital income. Income components are aggregated across all tax declarations belonging to the same household. For each horizon  $h$ , the dependent variable is the change in the corresponding household-income component between  $t - 1$  and  $t + h$ , scaled by total after-tax household income in  $t - 1$  and expressed in percentage points. The omitted horizon is  $h = -1$ . The fiscal shock is the narrative deficit-driven consolidation shock, measured as a percentage of GDP; positive values denote fiscal tightening. Estimates use household sampling weights and include household fixed effects, a linear time trend, and two lags of the corresponding household-income component. The dark-gray shaded area reports the 68% confidence interval, while the light-gray shaded area reports the 95% confidence interval. Standard errors are heteroskedasticity robust, two-way clustered at the household and year levels, and incorporate a Driscoll–Kraay correction implemented through a Newey–West correction with a horizon-specific bandwidth. Outcome changes are winsorized at the 5th and 95th percentiles. The sample covers 1999–2016.

**Figure I2:** Aggregate household-income responses to a continuous tax shock: expanded panel




*Notes:* The figure reports event-study panel local-projection estimates of the household-income response to the continuous tax shock using the expanded IRPF household panel. Panel (a) reports real total after-tax household income, panel (b) real total household labor income, and panel (c) real total household capital income. Income components are aggregated across all tax declarations belonging to the same household. For each horizon  $h$ , the dependent variable is the change in the corresponding household-income component between  $t - 1$  and  $t + h$ , scaled by total after-tax household income in  $t - 1$  and expressed in percentage points. The omitted horizon is  $h = -1$ . The tax shock is the narrative legislated tax shock from Gil et al. (2019), measured as a percentage of GDP; positive values denote tax increases. Estimates use household sampling weights and include household fixed effects, a linear time trend, and two lags of the corresponding household-income component. The dark-gray shaded area reports the 68% confidence interval, while the light-gray shaded area reports the 95% confidence interval. Standard errors are heteroskedasticity robust, two-way clustered at the household and year levels, and incorporate a Driscoll–Kraay correction implemented through a Newey–West correction with a horizon-specific bandwidth. Outcome changes are winsorized at the 5th and 95th percentiles. The sample covers 1999–2016.

**Figure I3: Aggregate household-income responses to a cumulative fiscal shock: expanded panel**




*Notes:* The figure reports event-study panel local-projection estimates of the household-income response to the cumulative fiscal consolidation shock using the expanded IRPF household panel. The specification follows Equation 1, replacing the contemporaneous shock with the cumulative shock defined in Equation 3. Panel (a) reports real total after-tax household income, panel (b) real total household labor income, and panel (c) real total household capital income. Income components are aggregated across all tax declarations belonging to the same household. For each horizon  $h$ , the dependent variable is the change in the corresponding household-income component between  $t - 1$  and  $t + h$ , scaled by total after-tax household income in  $t - 1$  and expressed in percentage points. The omitted horizon is  $h = -1$ . The cumulative fiscal shock is the sum of the narrative deficit-driven consolidation shocks between  $t$  and  $t + h$ , measured as a percentage of GDP; positive values denote cumulative fiscal tightening. Estimates use household sampling weights and include household fixed effects, a linear time trend, and two lags of the corresponding household-income component. The dark-gray shaded area reports the 68% confidence interval, while the light-gray shaded area reports the 95% confidence interval. Standard errors are heteroskedasticity robust, two-way clustered at the household and year levels, and incorporate a Driscoll–Kraay correction implemented through a Newey–West correction with a horizon-specific bandwidth. Outcome changes are winsorized at the 5th and 95th percentiles. The sample covers 1999–2016.

**Figure I4: Aggregate household-income responses to a cumulative tax shock: expanded panel**




*Notes:* The figure reports event-study panel local-projection estimates of the household-income response to the cumulative tax shock using the expanded IRPF household panel. The specification follows Equation 1, replacing the contemporaneous shock with the cumulative shock defined in Equation 3. Panel (a) reports real total after-tax household income, panel (b) real total household labor income, and panel (c) real total household capital income. Income components are aggregated across all tax declarations belonging to the same household. For each horizon  $h$ , the dependent variable is the change in the corresponding household-income component between  $t - 1$  and  $t + h$ , scaled by total after-tax household income in  $t - 1$  and expressed in percentage points. The omitted horizon is  $h = -1$ . The cumulative tax shock is the sum of the narrative legislated tax shocks between  $t$  and  $t + h$ , measured as a percentage of GDP; positive values denote cumulative tax increases. Estimates use household sampling weights and include household fixed effects, a linear time trend, and two lags of the corresponding household-income component. The dark-gray shaded area reports the 68% confidence interval, while the light-gray shaded area reports the 95% confidence interval. Standard errors are heteroskedasticity robust, two-way clustered at the household and year levels, and incorporate a Driscoll–Kraay correction implemented through a Newey–West correction with a horizon-specific bandwidth. Outcome changes are winsorized at the 5th and 95th percentiles. The sample covers 1999–2016.

**Figure I5:** Household-income responses across initial income groups to an expansionary continuous fiscal shock: expanded panel




*Notes:* The figure reports household-income-group-specific event-study panel local-projection estimates at horizon  $h = 2$  using the expanded IRPF household panel and following Equation 2. Panel (a) reports real total after-tax household income, panel (b) real total household labor income, and panel (c) real total household capital income. Income components are aggregated across all tax declarations belonging to the same household. For each component, the dependent variable is the change between  $t - 1$  and  $t + 2$ , scaled by total after-tax household income in  $t - 1$  and expressed in percentage points. The solid line plots the sign-reversed coefficient, so positive values correspond to an expansionary fiscal shock equal to 1% of GDP. The underlying shock is the narrative deficit-driven fiscal consolidation shock; positive realizations of the original series denote fiscal tightening. Households are assigned to their income group according to their first observed position in the household-income distribution, and this classification is kept fixed throughout the analysis. The first nine groups correspond to the 0–10th through 80–90th percentiles, followed by the 90th–99th percentiles and the top 1%. Estimates use household sampling weights and include household fixed effects, a linear time trend, and two lags of the corresponding household-income component. The dark-gray shaded area reports the 68% confidence interval, while the light-gray shaded area reports the 95% confidence interval. Standard errors are heteroskedasticity robust, two-way clustered at the household and year levels, and incorporate a Driscoll–Kraay correction implemented through a Newey–West correction with the horizon-specific bandwidth. Outcome changes are winsorized at the 5th and 95th percentiles. The sample covers 1999–2016.

**Figure I6:** Household-income responses across initial income groups to an expansionary continuous tax shock: expanded panel




*Notes:* The figure reports household-income-group-specific event-study panel local-projection estimates at horizon  $h = 2$  using the expanded IRPF household panel and following Equation 2. Panel (a) reports real total after-tax household income, panel (b) real total household labor income, and panel (c) real total household capital income. Income components are aggregated across all tax declarations belonging to the same household. For each component, the dependent variable is the change between  $t - 1$  and  $t + 2$ , scaled by total after-tax household income in  $t - 1$  and expressed in percentage points. The solid line plots the sign-reversed coefficient, so positive values correspond to an expansionary tax shock, represented by a tax reduction equal to 1% of GDP. The underlying shock is the narrative legislated tax shock from Gil et al. (2019); positive realizations of the original series denote tax increases. Households are assigned to their income group according to their first observed position in the household-income distribution, and this classification is kept fixed throughout the analysis. The first nine groups correspond to the 0–10th through 80–90th percentiles, followed by the 90th–99th percentiles and the top 1%. Estimates use household sampling weights and include household fixed effects, a linear time trend, and two lags of the corresponding household-income component. The dark-gray shaded area reports the 68% confidence interval, while the light-gray shaded area reports the 95% confidence interval. Standard errors are heteroskedasticity robust, two-way clustered at the household and year levels, and incorporate a Driscoll–Kraay correction implemented through a Newey–West correction with the horizon-specific bandwidth. Outcome changes are winsorized at the 5th and 95th percentiles. The sample covers 1999–2016.

**Figure I7:** Household-income responses across initial income groups to a cumulative expansionary fiscal shock: expanded panel



*Notes:* The figure reports household-income-group-specific event-study panel local-projection estimates at horizon  $h = 2$  using the expanded IRPF household panel. The specification follows Equation 2, replacing the contemporaneous shock with the cumulative shock defined in Equation 3. Panel (a) reports real total after-tax household income, panel (b) real total household labor income, and panel (c) real total household capital income. Income components are aggregated across all tax declarations belonging to the same household. For each component, the dependent variable is the change between  $t - 1$  and  $t + 2$ , scaled by total after-tax household income in  $t - 1$  and expressed in percentage points. The solid line plots the sign-reversed coefficient, so positive values correspond to a cumulative expansionary fiscal adjustment equal to 1% of GDP between  $t$  and  $t + 2$ . The cumulative shock measures the total policy adjustment over this period rather than a cumulated income response. The underlying shock is the narrative deficit-driven fiscal consolidation shock; positive realizations of the original series denote fiscal tightening. Households are assigned to their income group according to their first observed position in the household-income distribution, and this classification is kept fixed throughout the analysis. The first nine groups correspond to the 0–10th through 80–90th percentiles, followed by the 90th–99th percentiles and the top 1%. Estimates use household sampling weights and include household fixed effects, a linear time trend, and two lags of the corresponding household-income component. The dark-gray shaded area reports the 68% confidence interval, while the light-gray shaded area reports the 95% confidence interval. Standard errors are heteroskedasticity robust, two-way clustered at the household and year levels, and incorporate a Driscoll–Kraay correction implemented through a Newey–West correction with the horizon-specific bandwidth. Outcome changes are winsorized at the 5th and 95th percentiles. The sample covers 1999–2016.


**Figure I8:** Household-income responses across initial income groups to a cumulative expansionary tax shock: expanded panel



*Notes:* The figure reports household-income-group-specific event-study panel local-projection estimates at horizon  $h = 2$  using the expanded IRPF household panel. The specification follows Equation 2, replacing the contemporaneous shock with the cumulative shock defined in Equation 3. Panel (a) reports real total after-tax household income, panel (b) real total household labor income, and panel (c) real total household capital income. Income components are aggregated across all tax declarations belonging to the same household. For each component, the dependent variable is the change between  $t - 1$  and  $t + 2$ , scaled by total after-tax household income in  $t - 1$  and expressed in percentage points. The solid line plots the sign-reversed coefficient, so positive values correspond to a cumulative tax reduction equal to 1% of GDP between  $t$  and  $t + 2$ . The cumulative shock measures the total tax adjustment over this period rather than a cumulated income response. The underlying shock is the narrative legislated tax shock from Gil et al. (2019); positive realizations of the original series denote tax increases. Households are assigned to their income group according to their first observed position in the household-income distribution, and this classification is kept fixed throughout the analysis. The first nine groups correspond to the 0–10th through 80–90th percentiles, followed by the 90th–99th percentiles and the top 1%. Estimates use household sampling weights and include household fixed effects, a linear time trend, and two lags of the corresponding household-income component. The dark-gray shaded area reports the 68% confidence interval, while the light-gray shaded area reports the 95% confidence interval. Standard errors are heteroskedasticity robust, two-way clustered at the household and year levels, and incorporate a Driscoll–Kraay correction implemented through a Newey–West correction with the horizon-specific bandwidth. Outcome changes are winsorized at the 5th and 95th percentiles. The sample covers 1999–2016.

## Appendix J Spain growth incidence curve


**Figure J1:** Lakner-Milanovic graph for Spain (“mammoth” curve)



*Notes:* the figure is a growth incidence curve for Spain over 2001-2016. The income concept is the sampling income (*renta de muestreo*) that the Institute of Fiscal Studies records for each tax return, summed across all returns filed by the same household, and expressed in nominal euros. We use micro-data from the Household Tax Panel, which provides a representative sample of households with comprehensive income information tracked over time, provided jointly by the Spanish Tax Agency Administration and the Institute of Fiscal Studies (IEF, by its Spanish acronym). Further details on the methodology are available in [Pérez-López et al. \(2020\)](#). Households are sorted into 21 groups: twenty bins of five percentiles each, with the last two groups corresponding to the 95th–99th percentiles and the top 1%. The ranking variable is the household’s average total pre-tax income over the second, third and fourth preceding years, so that households are ordered by a measure of permanent rather than current income. Households are re-ranked in every year, so the curve compares positions in the distribution rather than the same households over time, and it does not separate income growth from changes in the composition of each group. Because three earlier years of income are required to rank a household, the first two years of the panel drop out and the curve covers 2001-2016. The vertical axis reports the cumulative growth of each group’s mean income between 2001 and 2016, expressed in percent. Group means are weighted by the household sampling weights. The red line is the unweighted average of the 21 group-specific growth rates, 4.88%, and therefore summarises the typical income group. The sample keeps households whose principal declarant is aged 26-65 and drops non-positive incomes, leaving 4,001,471 household-year observations. The Basque Country and Navarre operate their own tax administrations and are absent from the panel.

## Appendix K Income map

Figure K1: Spain income map.



*Notes:* the map shows the geographic distribution of average household income across Spanish municipalities, computed from the administrative tax records used in the analysis and expressed in constant 2015 euros. Household income is the sum of all tax returns filed by the household. For married couples who file separately, which are a quarter of all fiscal households by 2016, this means adding the second spouse's return. Municipal averages are weighted by the household sampling weights and pooled over 1999–2016. Municipalities are grouped into five classes of equal size, following the quantile classification used by the Spanish Statistical Office in its *Atlas de Distribución de Renta de los Hogares*. Two adjustments are made before mapping. First, municipal averages are winsorised at the 2nd and 98th percentiles: the panel samples about 2.9% of households, so in the smallest municipalities a single high-income household would otherwise dominate the average. Second, each municipality is replaced by the household-weighted average of itself and its contiguous neighbours; half of the municipalities are estimated from twelve sampled households or fewer, and without this smoothing the shading would reflect sampling noise rather than income. The smoothing induces spatial correlation by construction, so the map is descriptive and cannot be used to support claims about spatial dependence. Municipalities shaded in grey are not covered. The Basque Country and Navarre operate their own tax administrations and are absent from the panel, which accounts for 523 of them; the remaining 919 are municipalities in which no household was sampled, concentrated in the depopulated interior, above all in Guadalajara, Burgos, Cuenca, Teruel and Soria, with a median population of 90 inhabitants. The map covers 6,683 municipalities. It is not directly comparable with the published Atlas, which draws on the full census of tax returns filed with the national and the regional tax agencies rather than on a sample, and therefore requires no smoothing.

## Appendix L Macroeconomic representativeness of the administrative data

The analysis draws on two administrative sources. The distributional results of [Section 5](#) are estimated on the personal income tax panel, and the composition of the labor-income response is examined with the *Muestra Continua de Vidas Laborales* (MCVL), a four percent random sample of individuals with a Social Security record from which we construct a worker-year panel covering 1999–2016. Neither source is designed to reproduce a macroeconomic aggregate: the first records the income reported in a tax return and the second measures earnings through Social Security contribution bases. This appendix documents

that both of them nevertheless track the aggregate wage dynamics of the Spanish economy, and states precisely what the exercise establishes and what it leaves open.

The comparison is the same in both cases. For each source we construct an annual aggregate of labor income, raised to the population with the relevant weight, and compare it with Wages and Salaries from the Spanish National Accounts compiled by the INE, which measures the compensation of employees excluding employers' social contributions. All series are expressed in constant 2015 euros using the same GDP deflator employed throughout the paper, and all are indexed to 2015 = 100.

The indexation determines what the comparison can establish. Because each series is normalized by its own 2015 value, any proportional difference in level between a micro source and the aggregate is removed by construction, and the indices coincide mechanically in the base year. The exercise is therefore informative about co-movement and not about the absolute level of either aggregate. This distinction matters for the choice of statistic. The correlation between two series in log levels is invariant to scale and location, so it cannot detect a proportional bias: a version of a micro aggregate that understated the national-accounts aggregate by forty percent would deliver exactly the same correlation as the series we use. We therefore report the correlation in log levels for completeness but do not rely on it. The statistics we rely on are the correlation in growth rates, the slope of the relation between the two series relative to the 45-degree line, and the magnitude and stability of the gap between the two indices.


## L.1 The personal income tax panel

We use the expanded household panel described in [Appendix I](#) rather than the pure panel. The reason is a property of the two designs rather than a judgment about data quality. The pure panel follows the tax filers selected in the base year for as long as they remain observable, which is exactly what the event-study design of [Section 3](#) requires, but it means that its cross-section ages over the sample period and is not replenished. The expanded panel is refreshed each year within sampling strata to preserve cross-sectional representativeness, and it is therefore the one that can be compared with a contemporaneous aggregate. The weighted count of tax returns bears this out: it rises from 14 million in 1999 to 20 million in 2016, in line with the growth of the Spanish filing population.

One feature of the tax data governs the construction of the sample. Labor income as recorded in the tax return is not confined to wages: contributory pensions and unemployment benefits are declared as *rendimientos del trabajo* and enter the same box as salaries, whereas the national-accounts aggregate classifies them as social transfers and excludes them. They cannot be netted out, because no box in the Spanish personal income tax separates them from wage income in any year of the sample. The only available instrument is therefore the age of the filer, and we restrict the sample to filers aged 65 or younger. Self-employment income requires no such treatment: it is recorded as income from economic activities, a separate concept, and is excluded from the measure used here. The micro aggregate is the sum of gross labor income (*rendimientos íntegros del trabajo*) across all tax

returns in the restricted sample, raised to the population with the sampling weight.


**Figure L1:** The tax panel wage aggregate and the national-accounts aggregate



*Notes:* The figure compares the wage aggregate implied by the IRPF household panel with Wages and Salaries from the Spanish National Accounts (INE). The panel aggregate is the sum of gross labor income across all tax returns filed by individuals aged 65 or younger, raised to the population with the sampling weight. Both series are expressed in constant 2015 euros using the same GDP deflator and indexed to 2015 = 100. Panel (a) reports the two indices in logarithms. Panel (b) reports their annual growth rates, computed as  $(Y_t - Y_{t-1})/Y_{t-1} \times 100$  and expressed in percent.

Figure L1 plots the two series. They share the level path and the turning points of the cycle: the expansion through 2007, the contraction of 2008–2009, the second trough of 2012–2013 and the recovery thereafter. Figure L2 presents the same information as a scatter against the 45-degree line, and Table L1 reports the corresponding statistics.

**Figure L2:** Co-movement between the tax panel and the national-accounts aggregate



*Notes:* Each point is a year. Panel (a) plots the two indices in logarithms and panel (b) their annual growth rates, with the IRPF panel on the vertical axis and the national-accounts aggregate on the horizontal axis. The black dashed line is the 45-degree line, along which the two sources would coincide exactly; it is not a fitted line. The solid line is the least-squares fit estimated on the full sample and the long-dashed line the fit estimated on 2004–2016 only. The corresponding slopes, standard errors and tests of the hypothesis that the slope equals one are reported in Table L1.

The two sources move together closely. Over the full sample the correlation of annual growth rates is 0.907 and the slope of the panel growth rate on the national-accounts growth rate is 1.09 with a standard error of 0.13, so the hypothesis that the panel moves one-for-one with the aggregate is not rejected ( $p = 0.503$ ). Restricting the comparison to 2004–2016 leaves the conclusion unchanged and, if anything, sharpens it: the correlation of growth rates is 0.923 and the slope is 1.02 with a standard error of 0.13 ( $p = 0.861$ ). The correlation in log levels is 0.982 over the full sample. The gap between the two indices averages 3.18 log points and remains within a band of  $-4.97$  to  $7.08$  log points.

**Table L1:** Representativeness of the tax panel wage aggregate

|                           | 1999–2016 | 2004–2016 |
|---------------------------|-----------|-----------|
| Correlation, log levels   | 0.982     | 0.942     |
| Correlation, growth rates | 0.907     | 0.923     |
| Slope, growth rates       | 1.09      | 1.02      |
| Standard error            | (0.13)    | (0.13)    |
| $p$ -value, slope = 1     | 0.503     | 0.861     |
| Mean gap (log points)     | 3.18      | 4.75      |
| Minimum gap               | $-4.97$   | 0.00      |
| Maximum gap               | 7.08      | 7.08      |


*Notes:* The panel aggregate is the sum of labor income across all tax returns filed by individuals aged 65 or younger in the expanded household panel, raised to the population with the sampling weight. Labor income is measured gross of deductible expenses, which is the concept that corresponds to Wages and Salaries in the national accounts. It excludes self-employment income, which the tax return records as income from economic activities. The gap is the difference between the two indices in logarithms, expressed in log points; it is measured relative to the common 2015 normalization, so its level is not informative about absolute differences between the two sources. Correlations in log levels are reported for completeness only: they are invariant to scale and location and therefore cannot detect a proportional difference in level between the two sources. Unlike the pure panel, the expanded panel is replenished each year to preserve cross-sectional representativeness, so its weighted count of tax returns tracks the Spanish filing population. Labor income as recorded in the tax return includes contributory pensions and unemployment benefits, which the national-accounts aggregate excludes and which no box of the Spanish personal income tax separates from wage income; the sample is therefore restricted to filers aged 65 or younger.

## L.2 The MCVL

From the 2006–2016 waves of the MCVL we construct a worker-year panel covering 1999–2016 that records contribution bases, days worked, contract type, employer sector, public-sector status, and basic demographics. The micro aggregate is the sum of annual contribution bases across workers, weighted by full-time-equivalent days worked. Two features of its construction make it a closer counterpart of the national-accounts concept than the tax data: the bases of the self-employed are excluded, because they are chosen by the worker rather than paid as a wage, and so are the contributions paid on behalf of the unemployed, so neither self-employment income nor unemployment benefits enter the measure. Con-


tributory pensions do not enter it either.

**Figure L3: The MCVL wage aggregate and the national-accounts aggregate**



*Notes:* The figure compares the wage aggregate implied by the MCVL with Wages and Salaries from the Spanish National Accounts (INE). The MCVL aggregate is the sum of annual contribution bases across workers, weighted by full-time-equivalent days worked. Both series are expressed in constant 2015 euros using the same GDP deflator and indexed to 2015 = 100. Panel (a) reports the two indices in logarithms. Panel (b) reports their annual growth rates, computed as  $(Y_t - Y_{t-1})/Y_{t-1} \times 100$  and expressed in percent. The 1999–2002 segment of the MCVL is reconstructed retrospectively from the affiliation histories carried by later waves and is subject to survivorship, as discussed in the text.

**Figure L4: Co-movement between the MCVL and the national-accounts aggregate**



*Notes:* Each point is a year. Panel (a) plots the two indices in logarithms and panel (b) their annual growth rates, with the MCVL on the vertical axis and the national-accounts aggregate on the horizontal axis. Hollow gray markers denote the retrospective years 1999–2002 and solid markers the years 2003–2016. The black dashed line is the 45-degree line, along which the two sources would coincide exactly; it is not a fitted line. The solid line is the least-squares fit estimated on the full sample and the long-dashed line the fit estimated on 2003–2016 only. The corresponding slopes, standard errors, and tests of the hypothesis that the slope equals one are reported in [Table L2](#).

Over 2003–2016 the two sources move together closely. The correlation of annual growth rates is 0.847, and the slope of the MCVL growth rate on the national-accounts growth rate is 0.938 with a standard error of 0.170, so the hypothesis that the MCVL moves

one-for-one with the aggregate is not rejected ( $p = 0.722$ ). The same holds in log levels, where the slope is 0.925 with a standard error of 0.110 ( $p = 0.507$ ). The gap between the two indices averages 0.83 log points over this period and remains within a band of  $-2.72$  to 4.55 log points.

**Table L2:** Representativeness of the MCVL wage aggregate

|                           | 1999–2016 | 2003–2016 |
|---------------------------|-----------|-----------|
| Correlation, log levels   | 0.961     | 0.925     |
| Correlation, growth rates | 0.840     | 0.847     |
| Slope, log levels         | 1.32      | 0.93      |
| Standard error            | (0.10)    | (0.11)    |
| $p$ -value, slope = 1     | 0.004     | 0.507     |
| Slope, growth rates       | 1.02      | 0.94      |
| Standard error            | (0.17)    | (0.17)    |
| $p$ -value, slope = 1     | 0.898     | 0.722     |
| Mean gap (log points)     | $-0.93$   | 0.83      |
| Minimum gap               | $-11.62$  | $-2.72$   |
| Maximum gap               | 4.55      | 4.55      |

*Notes:* The MCVL aggregate is the sum of annual contribution bases across workers, weighted by full-time-equivalent days worked. The gap is the difference between the two indices in logarithms, expressed in log points; it is measured relative to the common 2015 normalization, so its level is not informative about absolute differences between the two sources. Slopes are from least-squares regressions of the MCVL series on the national-accounts series, in log levels and in growth rates respectively, and the reported  $p$ -values test the hypothesis that the slope equals one, that is, that the fitted line coincides with the 45-degree line. Correlations in log levels are reported for completeness only: they are invariant to scale and location and therefore cannot detect a proportional difference in level between the two sources. The 1999–2002 segment is reconstructed retrospectively from later MCVL waves and is subject to survivorship, so the 2003–2016 column is the relevant one for representativeness.

The full-sample column of [Table L2](#) isolates the one period in which the two sources disagree. Between 1999 and 2003 the MCVL aggregate grows by 22.6 log points while the national-accounts aggregate grows by 11.3, so the MCVL closes an initial shortfall of roughly twelve log points over four years. Estimated on the full sample, the slope in log levels is 1.32 with a standard error of 0.096, which rejects equality with the 45-degree line ( $p = 0.004$ ); restricted to 2003–2016 the same slope is 0.93 and does not reject. The divergence is a property of how the early years of the panel are constructed rather than a feature of the Spanish wage bill. The MCVL is drawn from waves beginning in 2006, and the years before the first wave are recovered from the affiliation histories that those waves carry. Workers who had left the Social Security records before the extraction date are therefore absent from the early cross-sections, which understates the aggregate at the beginning of the period and produces a mechanical catch-up as the sample approaches the first wave. The employment counts implied by the sample are consistent with this interpretation: for 1999 they fall short of contemporaneous registered affiliation. Notably, the growth-rate

slope is unaffected by this segment, since it remains statistically indistinguishable from one whether or not the retrospective years are included (1.02 with the full sample, 0.94 without it). For the analysis that uses the MCVL we restrict attention to 2003 onward.

### **L.3 What the exercise establishes**

Both administrative sources move one-for-one with the aggregate wage bill of the Spanish economy: the slope of the micro growth rate on the national-accounts growth rate is 1.02 in the tax panel and 0.94 in the MCVL, and neither is distinguishable from one. The two sources therefore reproduce the direction and the magnitude of the year-to-year movements against which the distributional results of [Section 5](#) should be read.

Three limitations should be kept in mind. First, as noted above, the indexation removes any proportional difference in level between a micro source and the aggregate, so the exercise establishes co-movement rather than level accuracy. Second, restricting the tax panel by age is a proxy for excluding pensioners, not an identification of them: it does not capture early retirement or disability pensions, it removes pensioners who continue to work, and in joint returns the age is that of the primary filer; unemployment benefits received by filers below the age threshold remain in the measure and cannot be separated from wage income. Third, contribution bases in the MCVL are censored at the legal maximum base, so its earnings measure understates earnings in the upper tail while the national-accounts aggregate does not. A constant proportional shortfall of this kind is absorbed by the indexation; what would affect the statistics reported here is variation over time in the censored share, which enters the comparison as movement in the gap, and the stability of the gap over 2003–2016 bounds how much of that variation the data admit. Censoring nevertheless constrains what the MCVL can say about the top of the earnings distribution, and we take it into account when interpreting results for high-earning workers.