

TI 2026-083/VI
Tinbergen Institute Discussion Paper

Austerity and Productivity: Evidence from Euro Area Regions

*Yao Chen*¹

*Ricardo Duque Gabriel*²

*Mathias Klein*³

*Ana Sofia Pessoa*⁴

¹ Erasmus University Rotterdam

² Federal Reserve Board

³ Sveriges Riksbank

⁴ International Monetary Fund

Austerity and Productivity: Evidence from Euro Area Regions ^{*}

Yao Chen[†] Ricardo Duque Gabriel[‡] Mathias Klein[§] Ana Sofia Pessoa[¶]

September 2026

Abstract

We find that fiscal consolidations raise productivity significantly and persistently. In our data, a consolidation that lowers regional public spending by 1% on impact raises TFP by about 1% on impact and about 2% after two years, even as economic activity contracts. We document this using two innovations. First, we build a novel regional utilization-adjusted TFP series for Europe, covering 139 NUTS 2 regions in ten euro area countries over 1999–2019, which corrects the Solow residual for cyclical variation in capital utilization and labor effort. Second, we isolate plausibly exogenous spending cuts with a Bartik-style instrument that interacts pre-accession regional spending shares with narratively identified national consolidation episodes. Productivity gains are larger than those observed in comparable downturns without consolidation and arise almost entirely within sectors and coincide with net exit of establishments, rather than reflecting reallocation across sectors.

JEL classification: E32, E62, H53.

Keywords: Fiscal consolidation, total factor productivity, euro zone.

^{*}We thank Ricardo Amaro, Davide Auteri, Tamon Asonuma, Francesco Bianchi, Olivier Coibion, Diego Comin, Harald Fadinger, Thiemo Fetzter, Matthias Meier, Karel Mertens, Nitya Pandalai Nayar, Alp Simsek, and seminar participants at the Banque de France, Banco de Portugal, the Annual Meeting of the VfS Standing Field Committee in International Economics, the Ghent Workshop on Empirical Macroeconomics, and the Annual CEPR MEF Symposium for valuable comments and suggestions. The opinions expressed in this article are the sole responsibility of the authors and should not be interpreted as reflecting the views of the Federal Reserve Board, the Federal Reserve System, the Sveriges Riksbank, and the IMF, its Executive Board, or IMF management. Yao Chen gratefully acknowledges financial support from the NWO grant 406.XS.24.02.051.

[†]Erasmus School of Economics, Erasmus University Rotterdam; CEPR; Tinbergen Institute. Email: y.chen@ese.eur.nl.

[‡]Federal Reserve Board. Email: ricardo.f.duquegabriel@frb.gov.

[§]Sveriges Riksbank, Research Division. Email: mathias.klein@riksbank.se.

[¶]International Monetary Fund. Email: apessoa@imf.org.

1 Introduction

Over the last two decades, many European economies experienced a significant slowdown in productivity growth (OECD 2015; IMF 2023), coinciding with a period of significant fiscal spending cuts. Growth performance was particularly dismal in Southern European countries, which implemented the largest consolidation programs. Among policymakers, this coincidence has fueled the idea that contractionary fiscal policy *decreases* productivity growth (Cournède et al. 2013; European Fiscal Board 2019). This policy narrative carries weight in ongoing discussions regarding the design of fiscal policy as public debt has risen sharply due to the COVID-19 pandemic and a renewed period of geopolitical tension. Against this background, our paper investigates the link between fiscal consolidations and productivity using granular data from euro area regions. In contrast to the policy narrative, we find robust empirical evidence that fiscal austerity significantly *increases* productivity.

Establishing this result requires two ingredients: isolating exogenous variation in government spending and measuring productivity. For the first, we exploit variation across euro area regions, following recent U.S. studies that use regional variation to estimate the effects of fiscal policy (Chodorow-Reich et al. 2012; Nakamura and Steinsson 2014; Bernardini et al. 2020). Our main data source is ARDECO, a database of European regional accounts, and our unit of observation is a NUTS 2 region, roughly comparable to a U.S. state. The regional approach is particularly attractive in our setting because all regions in our sample share the ECB’s policy stance, so time fixed effects absorb monetary policy, a confounder in national-level estimates.

To further isolate exogenous changes in regional government spending, we follow Gabriel et al. (2026) and instrument regional spending with a Bartik-type measure (Bartik 1991) that interacts predetermined regional sensitivities to changes in national government expenditures with the narrative national consolidation measures of Adler et al. (2024). The narrative series captures only changes in the national spending motivated by budget deficit reduction rather than by current or prospective economic conditions, so it should be uncorrelated with contemporaneous developments in economic activity. When interacted with predetermined regional sensitivities, these national consolidation shocks translate into heterogeneous spending cuts across regions.

The second ingredient — utilization-adjusted productivity series — does not exist at the regional level. Therefore, we construct a novel measure of the utilization-adjusted total factor productivity (TFP) for NUTS 2 regions, following Imbs (1999) and correcting the Solow residual for the cyclical variation in capital utilization and labor effort. The series covers 139 regions in ten euro area countries over 1999–2019. Aggregated at the country level, it aligns with existing utilization-adjusted measures (Basu et al. 2006; Jordà et al. 2024; Comin et al. 2025), while showing much more cross-sectional dispersion than national aggregates. This regional dispersion allows us to

estimate the productivity response more precisely.

We estimate the causal effect of austerity on productivity using instrumental variable local projections. The Bartik measure is a strong instrument for regional spending reduction, with a first-stage F-statistic well above conventional thresholds. We find that a consolidation that lowers regional public spending by one percent on impact raises utilization-adjusted TFP by about one percent on impact and about two percent after two years, relative to less exposed regions. The productivity response contrasts with a broad-based contraction in economic activity: regional GDP, employment, investment, and consumption all decline.

The finding survives a battery of robustness checks, including alternative productivity measures, different ways of constructing the instrument, additional controls, and changes in sample composition. We also probe identification directly: the estimates are unchanged when the shock is purged of predictable components, and placebo instruments built from shocks unrelated to fiscal policy do not predict regional spending.

A natural concern is that *any* recession raises measured productivity. However, we find a significantly larger productivity increase during recessions that coincide with fiscal consolidations. Productivity growth is thus systematically higher in recessions that coincide with consolidations, pointing toward specific transmission mechanisms underlying fiscal consolidations.

Two channels can generate the total productivity response: gains within sectors and reallocation across sectors. Separating tradable from non-tradable sectors, we find that the aggregate response is almost entirely a within-sector phenomenon, with non-tradables accounting for about three quarters of it. Across sectors, austerity does shift resources away from non-tradable activities, which are more exposed to government demand and less productive, but the shift is too small to matter for aggregate TFP. The within-sector gains have several possible sources: efficiency improvement, restructuring, or a cleansing effect that shifts activity toward more productive firms (Caballero and Hammour 1994; Osotimehin and Pappadà 2017). The regional-level data do not allow us to measure firm-level efficiency or restructuring directly. Firm dynamics, however, are consistent with cleansing: austerity is followed by a significant decline in the number of local units (establishments), with the sectoral pattern on impact matching that of the TFP response.

Related literature. Our paper is related to various strands of literature. We are part of the recent and fast-growing literature that uses sub-national data to estimate the impact of fiscal policy (Nakamura and Steinsson 2014; Bernardini et al. 2020; Chodorow-Reich et al. 2012; Auerbach et al. 2020; Gabriel et al. 2026). While the majority of these studies focus on output and employment outcomes, we provide new evidence on the effects of fiscal policy on productivity.

In terms of identification, we build on the narrative approach of Guajardo et al. (2014), Alesina et al. (2015), and Jordà and Taylor (2016), who use narratively identified fiscal consolidation episodes

to estimate the aggregate impact of austerity shocks. We combine these narrative consolidation episodes at the aggregate level with regional sensitivities to national spending changes to identify fiscal policy shocks at the regional level. This generates substantially more cross-sectional variation than country-level estimates, reduces estimation uncertainty, and sidesteps confounding factors such as endogenous monetary policy responses that can affect aggregate estimates.

A small but growing literature studies the effects of fiscal policy on productivity at the regional level. Gabriel et al. (2023) find that fiscal expansions raise labor productivity and the Solow residual in the Euro Area, and Canova and Pappa (2025) document short-run labor productivity gains from ERDF structural fund disbursements across European regions. We contribute to this literature by using newly constructed utilization-adjusted TFP series at the regional level. Labor productivity conflates efficiency gains with changes in capital intensity, while the Solow residual is mechanically procyclical through its utilization component (see Basu et al. 2006; Fernald and Wang 2016; Comin et al. 2025). As a result, standard productivity measures may be expected to show losses under austerity even absent genuine efficiency changes – a concern that motivates our utilization adjustment.¹

More broadly, our findings complement a growing literature documenting long-run productivity gains from government spending increases. Recent work shows that public expenditures can raise TFP through innovation-related channels, including public R&D spillovers (Antolin-Diaz and Surico 2025) and learning-by-necessity effects associated with defense spending (Ilzetzki 2024). In contrast, we show that spending contractions can also raise measured productivity when the demand withdrawn falls mainly on non-tradable activities, which depend heavily on public demand and tend to exhibit lower R&D intensity and productivity levels (Galindo-Rueda and Verger 2016). The gains arise mostly within these sectors, consistent with cleansing, rather than through a reallocation of resources across sectors. They are not costless, however: R&D expenditure shows little response at short horizons but is about 20 percent lower after four years, and patent applications decline as well. Consolidations thus raise measured productivity quickly, consistent with selection, while reducing innovation inputs and outputs with a lag; whether this eventually weighs on productivity beyond our horizon is a question our data cannot answer.

Outline. The remainder of the paper is organized as follows. Section 2 describes the data we use and the construction of the TFP measure. Section 3 presents the methodology. Section 4 shows our empirical results. Finally, Section 5 concludes.

¹At the national level, the closest study is Bardaka et al. (2021), who find that consolidations reduce TFP in a panel of OECD countries. The differences in terms of identification (narrative Bartik vs. panel VAR), aggregation (regional vs. country), and measurement (utilization-adjusted TFP vs. non-adjusted TFP) plausibly account for the opposite sign.

2 Data

In our analysis, we draw on a broad set of annual data covering the period from 1999 to 2019 for 139 regions in ten European countries: Austria, Belgium, Finland, France, Germany, Ireland, Italy, Netherlands, Portugal, and Spain. In the following, we describe the main variables used in our analysis. Table A.1 in the Appendix provides more information on the regional structure, and A.2 provides additional information regarding data definitions and sources.

2.1 Economic data

To measure regional economic condition, we rely on data from the Annual Regional Database of the European Commission’s Directorate General for Regional and Urban Policy (ARDECO), which is a highly disaggregated dataset across sectoral and regional dimensions. The database contains several long time-series indicators for European regions at different statistical scales and expands the Cambridge Econometrics Dataset used by much of the literature on European regional dynamics.

The database provides regional measures for output (gross domestic product (GDP) and gross value added (GVA)), investment, earnings, hours worked, and employment for different economic sectors like industry, construction, financial, non-financial, and non-market services. The dataset is an annual unbalanced panel starting in 1980 for the European Union (EU) and some European Free Trade Association (EFTA) and candidate countries.² By construction, ARDECO’s regional data is consistent with the commonly used national accounts data (Lequiller and Blades 2014). In particular, as explained in Appendix A, the regional ARDECO time series are constructed in such a way that the country aggregates equal the corresponding time series in the National Accounts reported by the annual macro-economic database of the European Commission’s Directorate General for Economic and Financial Affairs (AMECO).

The data are divided into NUTS (Nomenclature of Territorial Units for Statistics) regions. NUTS is a geocode standard for referencing the subdivisions of countries for statistical purposes. The hierarchy of three NUTS levels (NUTS 1, 2, 3) is established by Eurostat in agreement with each member state, and for most countries, the respective NUTS level corresponds to a specific administrative division within the country. ARDECO provides most data series at these regional disaggregation levels except for NUTS 3, for which it only reports measures on population, employment, GDP, and GVA.³

²The data used for our analysis was collected in June 2024.

³Our final dataset excludes the French overseas regions of Guadeloupe (FRY1), Martinique (FRY2), Guyane (FRY3), La Réunion (FRY4), and Mayotte (FRY5), as well as the Spanish autonomous cities of Ceuta (ES63) and Melilla (ES64), the Finnish autonomous region of Åland (FI20), Corsica (FRM0), the Spanish island regions of Illes Balears (ES53) and Canarias (ES70), and the Portuguese autonomous regions of Açores (PT20) and Madeira (PT30). In addition, we

Official data on final consumption expenditure of the general government (henceforth, government spending) is not available at the European regional level. Hereinafter, in the spirit of Brueckner et al. (2023) and closely following Gabriel et al. (2023) and Gabriel et al. (2026), we use the sum of GVA and intermediate consumption of the non-market sector as a proxy for government spending. GVA of the non-market sector is computed as the sum of compensation to employees (including social contributions), consumption of fixed capital (which measures the decline in the value of fixed assets owned as a result of normal wear, tear, and obsolescence), and taxes less subsidies on production. Because GVA of the non-market sector does not include intermediate consumption, which is, however, one of the main components of government spending, we use input-output (IO) tables from the PBL EUREGIO database (Thissen et al. 2018) to calculate regional intermediate consumption shares of the non-market sector, which we then add to the GVA of the non-market sector.

Our regional measure (GVA plus intermediate consumption of the non-market sector) is a valid proxy for government spending for several reasons. First, as previously mentioned, ARDECO’s regional data is consistent with the national accounts data by construction. By definition, there exists a close link between government spending and the GVA of the non-market sector. Consequently, almost the entire variation in the GVA of the non-market sector refers to activities by the general government. Second, government spending and our proxy measure show very similar statistical properties. Both measures are very tightly linked—regressions at the national and regional level suggest a strong and significant relationship between both measures, with estimated coefficients close to 1. We will thus refer to our regional proxy for government spending as government spending throughout the paper. More details on the justification of our proxy choice are given in Appendix A.

Another novel contribution of our analysis is the use of regional private-sector consumption data, which has been largely unavailable in existing European regional databases. While ARDECO provides comprehensive coverage of output, employment, and investment at the NUTS 2 level, regional private consumption expenditure is notably absent. We obtain consumption data from Oxford Economics, which constructs estimates of consumer spending at the NUTS 2 level for all European regions.⁴

exclude the 2015 level of GVA for the Southern Region of Ireland (IE05), which removes TFP growth for 2015–2016 for that region. Irish GDP grew by 26 percent in 2015 as multinational firms relocated intellectual-property and leasing assets to Ireland, raising measured value added without any corresponding change in real activity; the effect is concentrated in the Southern Region, where the unadjusted Solow residual would otherwise jump by 53 percent in a single year. Results are robust to excluding Ireland altogether (Appendix Table E.1).

⁴We are grateful to Ricardo Amaro and Anthony Light from Oxford Economics for sharing these data. The consumption series is measured on a ‘residence basis,’ capturing the spending of people living in each region rather than spending that occurs within the region. This measurement concept is more appropriate for our purposes as it better reflects actual household consumption behavior.

2.2 Regional utilization-adjusted TFP

2.2.1 Construction

Throughout, we measure productivity for the private sector, defined as the economy net of non-market services. National accounts measure non-market output at input cost, so a Solow residual computed over the total economy carries no information about technology in that sector. All series entering our TFP construction are therefore private-sector series, obtained by aggregating across the market sectors—agriculture, industry, construction, distributive and information services, and financial and business services (NACE A–N)—and excluding public administration, education, health, and other non-market services (NACE O–U).

We construct a measure of utilization-adjusted TFP for each NUTS 2 region in two steps. First, we derive the private-sector Solow residual, following the growth-accounting methodology of the Penn World Tables (PWT). This measure captures both productivity improvements and fluctuations in unobserved labor and capital utilization. Second, we remove the utilization component using a model-based adjustment adapted from Imbs (1999). This produces a measure of productivity that is insulated from cyclical factor-intensity movements.⁵

Production framework. For each region, private-sector output (private-sector GVA) is produced according to

$$Y_t^p = A_t (u_t K_t^p)^{1-\alpha_t} (e_t H_t^p)^{\alpha_t} = \tilde{A}_t (K_t^p)^{1-\alpha_t} (H_t^p)^{\alpha_t}, \quad (1)$$

where K_t^p is the private-sector capital stock, predetermined in the last period, H_t^p denotes private-sector hours worked, u_t the rate of capital utilization, e_t labor effort, and α_t the output elasticity of labor. The term A_t denotes utilization-adjusted TFP, while

$$\tilde{A}_t = A_t u_t^{1-\alpha_t} e_t^{\alpha_t},$$

is the corresponding (unadjusted) Solow residual.

Step 1: Private-sector Solow residual. Abstracting from utilization, log-differentiating (1) yields

$$\Delta \ln \tilde{A}_t = \Delta \ln Y_t^p - (1 - \alpha_t) \Delta \ln K_t^p - \alpha_t \Delta \ln H_t^p. \quad (2)$$

Following PWT and EU-KLEMS, α_t is proxied by the smoothed private-sector labor income share (labor compensation to GVA, two-year moving average). ARDECO reports private-sector

⁵Several studies have constructed regional TFP measures (e.g. Beugelsdijk et al. 2018; Siller et al. 2021) relying on the Solow residual. For our purposes, it is essential to incorporate adjustments for time-varying utilization, as a large literature documents that fluctuations in labor effort and capital utilization account for a substantial share of the cyclicity in the Solow residual (see Burnside et al. 1995; Basu and Kimball 1997).

output and hours, but only a total-economy capital stock; we estimate the private-sector stock by perpetual inventory:

$$K_t^p = K_{t-1}^p + I_t^p - D_t^p, \quad (3)$$

where I_t^p and D_t^p are private-sector gross fixed capital formation and consumption of fixed capital, both from ARDECO and summed over the market sectors. The initial level of private-sector capital stock in 1995 is estimated by scaling each region's total capital stock by the national share of private-sector investment in that year: $K_{i,0}^p = K_{i,0} \times \frac{I_0^p}{I_0}$.⁶ The resulting Solow residual ($\Delta \ln \tilde{A}_t$) reflects both productivity changes and fluctuations in factor utilization.

Step 2: Adjustment for labor effort and capital utilization. To isolate productivity from cyclical fluctuations in factor utilization, we adjust $\tilde{A}_t$ for variation in labor effort e_t and capital utilization u_t . The adjustment follows from the optimality conditions of a model in which households derive utility from consumption, hours, and effort, while firms optimally choose utilization and effort as inputs alongside hours and the capital stock. Higher capital utilization entails additional depreciation costs, while greater labor effort requires compensating wage premia. The resulting optimality conditions governing these adjustments are summarized below, with the full model and derivations provided in Appendix B.

Effort adjustment. Optimizing households jointly choose hours and effort by equating the marginal disutility of each activity to its corresponding marginal return in labor income. Firms, in turn, choose labor services on both the extensive (hours) and intensive (effort) margins such that the marginal product of effective labor equals the marginal cost of supplying it. As shown in Appendix B, combining the household first-order conditions for hours and effort with the firm's optimality condition yields a mapping between hours worked and effort supplied:

$$e_t = (H_t^p)^{\frac{1+v}{1+v_e}}, \quad (4)$$

where the elasticities (v, v_e) reflect the curvature of the disutilities from hours and effort. This mechanism aligns with the interpretation in Basu et al. (2006), who emphasize that fluctuations in hours serve as an observable proxy for broader movements in labor utilization and work intensity. Following Jordà et al. (2024), we set the curvature parameters of disutility from hours and effort to unity, $v = v_e = 1$.⁷

⁶Gardiner et al. (2020) estimates of private-sector capital stock are available only until 2016. When using them to construct a new TFP measure as a robustness check, results remain unchanged.

⁷In practice, we implement (4) using hours relative to their HP-filtered trend, $e_t = (H_t^p / \bar{H}_t^p)^{(1+v)/(1+v_e)}$, so that $e_t = 1$ when hours are at trend. The constant in the production function absorbs this normalization, leaving the growth rates in (6) unaffected. The exponent, by contrast, depends on the two labor supply elasticities; Appendix Figure D.1(a) shows that our baseline results are robust to alternative values of v and v_e .

Capital-utilization adjustment. The real cost of utilization per unit of capital rises with the intensity of use u_t , reflecting the additional depreciation and maintenance associated with more intensive operation. Since these costs apply to the entire capital stock, total nominal utilization costs are proportional to K_t^p and the replacement-cost price of investment goods Q_t . Firms choose u_t by equating the marginal cost of higher utilization to the marginal product of utilized capital. Because output enters both this condition and the optimality condition for hours, combining the two links the optimal utilization rate directly to the wage bill relative to capital costs (see Appendix B):

$$\frac{u_t}{\bar{u}} = \left(\frac{W_t H_t^p}{K_t^p Q_t} \middle/ \frac{\bar{W} \bar{H}^p}{\bar{K}^p \bar{Q}} \right)^{1/\beta_k}, \quad (5)$$

where $W_t H_t^p$ denotes nominal private-sector labor compensation, Q_t the investment-goods price index, and $\beta_k > 0$ the curvature of the utilization cost function. A higher β_k implies that utilization costs rise more steeply with intensity, dampening the firm's utilization response. The intuition is straightforward: since the capital stock is predetermined, firms facing stronger demand optimally raise utilization in proportion to labor payments, so u_t comoves with the labor-to-capital cost ratio. Trend values $(\bar{W} \bar{H}^p / \bar{K}^p \bar{Q})$ are obtained by HP-filtering the corresponding ratio, and $\bar{u}$ denotes the steady-state utilization rate. We set $\beta_k = 4$ following Anzoategui et al. (2019).⁸

Utilization-adjusted TFP. Combining (2)–(5), growth in utilization-adjusted TFP is given by

$$\Delta \ln A_t = \Delta \ln \tilde{A}_t - (1 - \alpha_t) \Delta \ln u_t - \alpha_t \Delta \ln e_t, \quad (6)$$

where the utilization terms are constructed using (4) and (5).

Alternative elasticity estimates. Using income shares to proxy α_t presumes perfect competition. Under imperfect competition, cost shares should be used instead (Comin et al. 2025). We retain the income-share baseline because regional cost-of-capital data are not directly observable, so constructing cost-share-based elasticities requires additional assumptions that may introduce their own measurement error (see Appendix D.2 for details). As a robustness check, we estimate the cost of capital following Comin et al. (2025) and recompute TFP using the resulting cost-share-based elasticities; our main findings are robust to using cost-share-based elasticities instead (Section 4.2).


2.2.2 Properties of the regional TFP series

Summary statistics and spatial patterns. Across all NUTS 2 regions, the average annual utilization-adjusted TFP growth rate over 1999–2019 is 0.21 percent, with a cross-regional standard deviation of 0.73; at the region-year level, annual growth rates have a mean of 0.21 percent and a

⁸Our baseline results are robust to alternative values of β_k ; see Appendix Figure D.1(b).

standard deviation of 2.55. Figure 1 displays the average annual TFP growth rate for each region in our sample. Substantial heterogeneity emerges across the euro area: regions in Finland, Eastern and Southern Germany, and Eastern Austria exhibit persistently strong productivity growth over the sample period, whereas productivity declines are concentrated in Southern Italy and parts of Southern Spain.⁹

Figure 1: Average Annual TFP Growth Rates



Notes: The figure displays the average annual utilization-adjusted TFP growth by NUTS 2 region over 1999–2019. Darker shades indicate higher productivity growth. The sample covers 139 regions across ten euro area countries: Austria, Belgium, Finland, France, Germany, Ireland, Italy, the Netherlands, Portugal, and Spain. Regions in gray are outside the sample. See Table A.1 for details on sample coverage.


For comparison, labor productivity displays substantially higher average growth and similar dispersion (Appendix Figure C.4). Across regions, mean annual labor-productivity growth over 1999–2019 is 1.06 percent with a cross-regional standard deviation of 0.79; at the region-year level, annual growth rates have a mean of 1.06 percent and a standard deviation of 2.94. The two maps are positively correlated across regions (0.58), so regions with faster TFP growth also tend

⁹The regions with the highest average annual TFP growth are IE06 (Eastern and Midland, +3.00 percent), DED5 (Leipzig, +1.71 percent), and DE24 (Oberfranken, +1.62 percent). At the other end of the distribution, the lowest rates are observed in IE04 (Northern and Western, -1.74 percent), ITF5 (Basilicata, -1.23 percent), and ITF2 (Molise, -1.21 percent).

to have faster labor-productivity growth, without the two coinciding. Labor productivity can obscure underlying differences in productivity performance across regions because it conflates productivity dynamics with changes in capital stock and factor utilization. Notwithstanding, as shown below, our main empirical results are robust to using labor productivity instead of TFP.

Regional versus national dispersion. Figure 2 compares the distribution of annual TFP growth using regional versus country-level aggregates. Incorporating regional information markedly increases the observed dispersion in productivity dynamics. Ireland exhibits the highest average growth and the largest dispersion in both the regional and national series. The much wider dispersion in regional TFP growth provides substantially more identifying variation than national aggregates, thereby improving the precision of our estimated productivity responses to fiscal consolidations.

Figure 2: Distribution of Country- and Regional-Level Annual TFP Growth




Notes: Distribution of annual utilization-adjusted TFP growth for region-level and country-level series, 1999–2019.

Cyclicality. TFP is, on average, procyclical: the correlation between annual TFP growth and annual output growth is 0.51, significant at the 1 percent level. Figure 3 illustrates the positive relationship at the regional level.¹⁰

¹⁰The strength of cyclicality varies across countries. For example, TFP growth is negatively correlated with output growth in Spain (−0.23, significant), only weakly correlated in Portugal (+0.11, insignificant), and moderately correlated in the Netherlands (+0.26) and Ireland (+0.30). In all other countries, correlations range between 0.56 and 0.78.

Figure 3: Utilization-Adjusted TFP Growth and Output Growth



Notes: Scatter plot of annual utilization-adjusted TFP and GDP growth at the regional level, 1999–2019.

2.2.3 Comparison with aggregate national TFP measures

To validate our series, we benchmark it against widely used national TFP measures. We aggregate regional Solow residual and utilization-adjusted TFP growth rates into national series using Törnqvist indices, weighting each region by the two-year moving average of its share in national GVA.

Non-utilization-adjusted TFP. Our aggregated national Solow residuals correlate at 0.8–0.9 with the non-adjusted TFP measures from the Penn World Table and EU-KLEMS (Table 1, Panel A). This is noteworthy because our series is based exclusively on private-sector activity, whereas both benchmarks are economy-wide, indicating that private-sector dynamics account for most of the variation in aggregate TFP.

Utilization-adjusted TFP. When we aggregate our utilization-adjusted TFP series to the national level, the resulting measure remains positively and significantly correlated with the national utilization-adjusted TFP estimates of Comin et al. (2025), Basu et al. (2006), and Jordà et al. (2024) (Table 1, Panel B). Correlations vary with how each approach handles utilization. Basu et al. (2006) and Comin et al. (2025) purge utilization through an instrumental-variables (IV) procedure based on utilization proxies, and Comin et al. (2025) additionally estimate the output elasticity from cost rather than revenue shares. The IV approach relies on an exclusion restriction requiring TFP to be exogenous to the instruments used to purge utilization effects; this assumption is difficult to

maintain in our setting, where the objective is to study whether TFP responds endogenously to demand fluctuations. By contrast, the utilization adjustment of Imbs (1999), which we adopt as our baseline, is derived from firms’ and households’ optimality conditions and does not require a regression-based identification strategy. Our series is closest to Jordà et al. (2024), who apply the same adjustment. Despite the methodological differences, our finding remains robust when we recompute adjusted TFP using cost-share based elasticities (Appendix D.2) or when applying the instrumental variables procedure used in Basu et al. (2006) and Comin et al. (2025) (Appendix D.1).¹¹

Table 1: Correlation of Aggregated Regional TFP with National Benchmarks

Benchmark series	Correlation	N	Countries	Period
<i>Panel A: Solow residual (no utilization adjustment)</i>				
Penn World Table	0.796***	199	10	2000–2019
EU–KLEMS	0.891***	181	10	2000–2019
<i>Panel B: Utilization-adjusted TFP</i>				
Comin et al. (2025)	0.395***	100	5 (DE, ES, FR, IT, NL)	2000–2019
Basu et al. (2006) adjustment	0.455***	75	4 (DE, ES, FR, IT)	2000–2018
Jordà et al. (2024)	0.545***	128	8 (BE, DE, ES, FI, FR, IT, NL, PT)	2000–2015

Notes: Pairwise correlations between the growth rate of our national series—obtained by aggregating the regional series with Tornqvist weights—and national TFP growth from external sources, computed on the country-years for which both series are available. Panel A correlates the unadjusted private-sector Solow residual with the non-utilization-adjusted TFP series of the Penn World Table and EU–KLEMS. Panel B correlates the utilization-adjusted TFP measure with the utilization-adjusted series of Comin et al. (2025) (Europrod database), the Basu et al. (2006)-adjusted series from the Comin et al. (2025) data package, and Jordà et al. (2024) (authors’ data package). All series are demeaned within country. Ireland in 2015 is excluded (see footnote 3). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

2.3 Narrative austerity episodes

Our data for narrative fiscal consolidations comes from Adler et al. (2024) and spans from 1978 to 2020. Building on Devries et al. (2011), Alesina et al. (2015), and Alesina et al. (2020), we use this variable to address the potential endogeneity of shifts in fiscal variables using the “narrative” approach in the spirit of Romer and Romer (2010) and carefully dividing variables into spending- and tax-based consolidations.

The measure is constructed by examining contemporaneous OECD policy documents that outline the economic situation, fiscal consolidation strategy, and major consolidation measures for each of the OECD member countries. The country notes in each report lay out the government’s

¹¹We also aggregate our regional utilization indices using Tornqvist weights and compare them with national manufacturing capacity-utilization indicators from the European Commission’s Business and Consumer Survey (BCS). Despite differences in sectoral coverage, the correlation is 0.5 and significant at the 1 percent level, further supporting the informational content of our utilization measure.

rationale for pursuing fiscal adjustment, making it possible to isolate consolidations motivated by deficit reduction aiming to correct its long-run trend, and those driven by other motives unrelated to the state of the business cycle. Consolidations are measured in terms of their impact on total revenue and expenditure (relative to a baseline without policy intervention) and scaled by the output level prior to the intervention announcement.

The main advantage of identifying fiscal consolidations via the narrative measure, compared to changes in the cyclically-adjusted primary balance (CAPB) as suggested by Alesina and Ardagna (2010), is that they are exogenous to current economic developments, while changes in the CAPB are correlated with the business cycle. Guajardo et al. (2014) show that there is a significant positive correlation between GDP forecast revisions and changes in the CAPB, whereas the null hypothesis of no correlation between forecast revisions and the narrative measure cannot be rejected.

Alesina et al. (2020) classify as spending-based consolidations all measures related to government spending and investment, including expenditure on goods and services, salaries, managing costs of state-provided services (such as education and healthcare), and government gross fixed capital formation expenditures. Regarding tax-based consolidations, they account for all direct and indirect tax changes. Adler et al. (2024) combine the Devries et al. (2011) and Alesina et al. (2020) datasets, after verifying their consistency, and extend them through 2020.

Throughout the paper, the narratively identified austerity episodes at time t in country I ($\tilde{g}_{I,t}$) measure only spending-based consolidations, excluding episodes driven by significant changes in the tax system. The regional government spending proxy used in the analysis does not include tax revenues. It mainly encompasses the public wage bill and, to a lesser extent, the consumption of fixed capital and intermediate consumption. Therefore, excluding consolidation episodes driven by significant changes in the tax system allows for a stronger and clearer relationship between the narrative national austerity episodes and the regional government spending measures.¹² Following the definition in Devries et al. (2011), Adler et al. (2024) construct $\tilde{g}_{I,t}$ as the sum of unanticipated shifts in government spending at time t – ($g_{I,t}^u$) – and changes in spending that are implemented at time t but had been announced in previous periods – ($g_{I,t-1,t}^a$).

For our sample, we observe 62 consolidation episodes and 2 reversals of consolidations, which corresponds to 30% of all country-year observations. The consolidation episodes are expressed relative to lagged GDP. The mean (median) consolidation amounts to 0.88% (0.65%) of GDP. The largest intervention of 3.76% of GDP occurred in Portugal in 2012 during the Sovereign Debt Crisis. Figure A.1 presents the individual austerity episodes for each country.

Consolidation episodes are not confined to the periphery: the European core (Austria, Belgium,

¹²For the identification strategy described in section 3, focusing on spending-based fiscal adjustments implies maximizing the link between the exogenous shift variable and the outcome variable of the first-stage regression.

Finland, France, Germany, and the Netherlands) accounts for 33 episodes and the periphery (Ireland, Italy, Portugal, and Spain) for 29, so both groups contribute substantial identifying variation. Episodes are also spread over time: consolidations take place in the early and mid-2000s, during the sovereign debt crisis, and in the years that followed, so that identification does not rest on a single crisis period.

Unconditional relationship between TFP growth and austerity-driven spending changes.

Figure 4 provides some first insights into the relationship between government spending changes and TFP growth. It shows for the euro area regions in our sample and for one of the main consolidation years in the sample, namely 2012, the relation between austerity-driven government spending growth and TFP growth. There is a significant negative correlation, that is, reductions in public spending are associated with increases in productivity growth. While Figure 4 is informative about the unconditional correlation between TFP growth and fiscal consolidations, it does not provide a causal interpretation. In the rest of the paper, we conduct a thorough econometric analysis to investigate whether austerity causes an increase in TFP.

Figure 4: TFP Annual Growth and Austerity-driven Spending Changes in 2012



Notes: The chart plots on the y-axis one-year changes in TFP and on the x-axis one-year changes in austerity-driven government spending per capita. For this, government spending is first regressed on the austerity shocks and region and time fixed-effects. The correlation is -0.21 and statistically significant at 5% confidence level.

3 Empirical Strategy

We estimate the dynamic effects of austerity on regional economic outcomes using an instrumental variable local projections approach closely following the econometric specification by Gabriel et al. (2026). For each horizon $h \in \{0, 1, 2, 3, 4\}$, we estimate:

$$\frac{Y_{i,t+h} - Y_{i,t-1}}{Y_{i,t-1}} = \beta_h \frac{G_{i,t} - G_{i,t-1}}{G_{i,t-1}} + \gamma_h(L)X_{i,t-k} + \alpha_{i,h} + \delta_{t,h} + \varepsilon_{i,t+h} \quad \forall h \in \{0, 1, 2, 3, 4\}, \quad (7)$$

where $\frac{Y_{i,t+h} - Y_{i,t-1}}{Y_{i,t-1}}$ is the cumulative growth rate of TFP or other outcomes of interest, including per capita GVA, employment, consumption, and investment; $(L)X_{i,t-k}$ is a vector of controls; $\alpha_{i,h}$ are region fixed-effects to control for region-specific (unobserved) characteristics; and $\delta_{t,h}$ are time fixed-effects that account for common aggregate trends including monetary policy. The control vector includes two lags of the dependent variable and two lags of regional per capita government spending growth. The coefficient β_h directly yields the cumulative response of the outcome to a one percent decline in per capita government spending $G_{i,t}$. Standard errors are clustered at the region level throughout.

Estimating equation (7) by OLS is subject to two endogeneity concerns. First, government spending and TFP may respond to common demand or supply shocks, so any estimated relationship could reflect structural disturbances unrelated to fiscal consolidations. Second, higher productivity may itself reduce public spending needs, generating reverse causality.

To address both concerns, we instrument the change in regional government spending with a Bartik-type instrument (Bartik 1991):

$$\frac{\overline{G_i}}{\overline{G_I}} \times \tilde{g}_{I,t}, \quad (8)$$

where $\tilde{g}_{I,t}$ is the narrative national consolidation measure described in Section 2.3, and $\overline{G_i}$ and $\overline{G_I}$ are average *per capita* government spending in region i and country I over the five years before the accession of country I to the euro area, respectively.

The ratio $s_i \equiv \frac{\overline{G_i}}{\overline{G_I}}$ measures each region's predetermined sensitivity to national spending changes. Regions with $s_i > 1$ receive a disproportionate share of per capita national spending and are, therefore, expected to experience larger spending cuts when the national government implements austerity measures. The instrument scales the national consolidation shock by this regional sensitivity, so that regions more exposed to national spending face larger cuts when their country consolidates. Figure C.3 in the Appendix shows a heat map of s_i , which has a mean close to one and substantial cross-sectional variation ranging from 0.70 to 2.21.¹³

In the taxonomy of Borusyak et al. (2022, 2025), our design is closest to the path that relies on the exogeneity of the shocks.¹⁴ Because s_i averages one within each country, the identifying variation

¹³The lowest shares are in Andalucía (Spain, 0.701), Niederösterreich (Austria, 0.702), and Norte (Portugal, 0.741); the highest in Région de Bruxelles-Capitale (Belgium, 2.210), Lazio (Italy, 1.655), and Valle d'Aosta (Italy, 1.601).

¹⁴Borusyak et al. (2022, 2025) distinguish two identification paths for shift-share instruments. Path 1 treats the shares as fixed and relies on the exogeneity of the shocks: with a sufficiently large number of independent shocks,

stems predominantly from differences in the timing and size of national consolidations across countries and years, which s_i scales according to each region's exposure. The instrument therefore draws on many country-year shifts, including 62 consolidation episodes, rather than on a single common shock. The narrative construction of $\tilde{g}_{I,t}$ is designed to exclude consolidations motivated by current or prospective economic conditions, the defining property of the narrative approach, and year fixed effects absorb shocks common to all euro area regions, including monetary policy. The key identifying assumption is thus that national consolidation episodes are unrelated to other country-specific shocks that affect regional productivity differentially. The shares, in turn, need only be predetermined and uncorrelated with differential regional exposure to such concurrent national shocks beyond the spending channel, not as-good-as-randomly assigned across regions. Both conditions are defensible. The shares are computed from the five years before euro accession and are nearly time-invariant: the average within-region standard deviation over 1994–2019 is 0.032.¹⁵ Hence, they do not respond to post-accession fiscal or economic developments. Moreover, because the data are disaggregated to the NUTS 2 level, it is implausible that individual regions influence national fiscal decisions, ruling out the main reverse-causality threat to share exogeneity (Gabriel et al. 2023). Section 4.2 reports checks aimed at country-level confounders: controls for sovereign spreads, public debt, IMF programs, and the global financial crisis, and estimates that exclude one country at a time.

An alternative and complementary interpretation follows Goldsmith-Pinkham et al. (2020), who show that the shift-share estimator can be written as a Rotemberg-weighted average of region-specific local average treatment effects, where the weights are proportional to each region's contribution to the identifying covariance between the instrument and the treatment. Under this interpretation, a causal reading of $\hat{\beta}_h$ requires only that the identifying variation be concentrated in episodes where $\tilde{g}_{I,t}$ is plausibly exogenous, which are precisely the narrative consolidation episodes we exploit. The Rotemberg diagnostics reported in Section 4.2 confirm that regional weights are uncorrelated with s_i ($\rho = 0.062$), so high-exposure regions do not mechanically drive identification, and that dropping the observations with the highest and most negative weights (where the instrument and treatment move in opposite directions) leaves the baseline results qualitatively unchanged.

Figure 5 shows the first-stage regression, extended to a five-year horizon:


$$\frac{G_{i,t+h} - G_{i,t-1}}{G_{i,t-1}} = \beta_h \frac{\overline{G_i}}{\overline{G_I}} \times \tilde{g}_{I,t} + \gamma_h(L) X_{i,t-k} + \alpha_{i,h} + \delta_{t,h} + \varepsilon_{i,t+h} \quad \forall h \in \{0, 1, 2, 3, 4\}, \quad (9)$$

idiosyncratic components average out and the instrument isolates exogenous variation. Path 2 treats the shocks as fixed and relies on the exogeneity of the shares, requiring that exposure shares are as-good-as-randomly assigned across units.

¹⁵Our baseline finding is also robust to replacing the pre-euro-accession average s_i with the lagged time-varying share $s_{i,t-1}$ shown (Appendix Table E.2).

where the vector of controls includes lags of government spending and TFP growth. Regional government spending falls significantly and persistently following the austerity shock. The response is normalized so that spending declines by 1% on impact; three years after the shock, the reduction reaches its trough at around 2.6 percent, and spending remains significantly below its pre-shock level at the end of the forecast horizon. The Kleibergen–Paap rk Wald F-statistic of 26.25 is above the Olea and Pflueger (2013) threshold of 23 for a 5% critical value, confirming instrument relevance.

Figure 5: First Stage: Government Spending Response



Notes: Figure 5 plots the percent change in per capita government spending in response to a one percent fiscal shock. Bands are 90% (dark) and 95% (light) confidence intervals.

4 Results


4.1 Baseline results

Figure 6 shows our main baseline results. Panel (a) shows the output response, in terms of private-sector GVA and panel (b) reports the response of our newly constructed utilization-adjusted TFP measure. A fiscal consolidation induces a significant decline in output, very much in line with previous literature (Guajardo et al. 2014; Gabriel et al. 2026). After a significant decline of roughly 1% already in the impact period, private-sector GVA reaches its trough three years after the shock with a fall of almost 3%. Other important economic indicators respond consistently with this pattern: investment contracts sharply, the private-sector capital stock declines persistently, private-sector hours worked and wages both fall, pointing to a broad-based deterioration in economic activity and labor market conditions (see Appendix C.1 for a detailed discussion).

Notably, productivity moves in the opposite direction of output following the fiscal contraction. On impact, TFP rises by about 1%, and the response peaks at about 2% after two years. Four years after the shock, productivity remains elevated at around 1.6%, with relatively tight confidence

bands throughout. Thus, our baseline results suggest that while fiscal consolidations lower general economic activity, they significantly and persistently increase productivity in the short and medium term. The utilization adjustment is essential for this result: the unadjusted Solow residual responds far less (Appendix Figure C.2).

Figure 6: Response of Output and TFP to Austerity Shock



Notes: Impulse responses of TFP and other growth components to a one percent austerity-driven decline in government spending, estimated as in equation (7). Bands are 90% (dark) and 95% (light) confidence intervals.

We complement our baseline regression with a forecast error variance decomposition (FEVD) exercise to quantify the contribution of austerity shocks to the variation in economic outcomes (Appendix C.4), finding that austerity shocks account for up to 4% of the variation in productivity and up to 13% of the variation in GDP.

After establishing the robustness of our baseline (Section 4.2), we show that the productivity gains are larger than in ordinary recessions and therefore not simply a cyclical phenomenon (Section 4.3). Sections 4.4 to 4.6 then show that firm dynamics point to cleansing within sectors.

4.2 Robustness checks

We subject the baseline TFP result to a range of specification changes. Table 2 reports the estimates, with the baseline repeated at the top for reference. We first assess how robust the baseline is to the construction of productivity, of the exposure share, and of the sample, and then turn to the validity of our identification.

Robustness A potential concern is that the measured *productivity* gains reflect the construction of our utilization-adjusted TFP measure rather than genuine productivity improvements. We consider three alternatives. Row A1 replaces the Imbs adjustment with the approach of Basu et al. (2006) (BFK), which corrects the Solow residual for utilization by proxying it with changes

in hours per worker; Appendix D.1 provides details. Row A2 relaxes perfect competition by replacing revenue-share-based input elasticities with cost-share-based ones following Comin et al. (2025), while retaining the Imbs utilization adjustment; Appendix D.2 provides details. Both alternatives yield positive, persistent, and significant TFP responses. The cost-share estimates are close to the baseline. The BFK estimates are somewhat attenuated, consistent with the exclusion restriction underlying that approach being difficult to sustain in our setting.¹⁶ Row A3 dispenses with the production function entirely, replacing private-sector TFP with private-sector labor productivity (private-sector GVA divided by private-sector hours worked), which requires no assumptions on market structure, on factor elasticities, or on the utilization adjustment. The response is positive and significant at most horizons, though less persistent than the TFP response, as labor productivity conflates efficiency gains with cyclical movements in the capital stock and its utilization.¹⁷

The result is likewise robust to alternative constructions of the *exposure share* s_i . Replacing the baseline per-capita government consumption share with a per-capita GDP share (Row B1), a total spending share (Row B2), the regional government-to-output ratio (Row B3), all yield positive and statistically significant TFP responses of comparable magnitude.¹⁸

We next ask whether particular *observations* or *countries* account for the estimate. Following Goldsmith-Pinkham et al. (2020), we decompose the IV estimator into Rotemberg weights. The weights are concentrated in a relatively small number of observations, and negative weights are non-negligible, reflecting heterogeneity in how the instrument predicts spending changes. The regional weights are, however, essentially uncorrelated with the exposure share s_i (6.2 percent), so high-exposure regions do not mechanically drive identification. Rows C1–C2 drop observations in the bottom and top 1% of the weight distribution, confirming that neither tail, including observations carrying negative weights, accounts for the result. Appendix Table E.2 reports further checks excluding the top 10% and top 5% of regions ranked by s_i ; results are qualitatively unchanged.¹⁹ We also rerun our baseline specification, excluding one country at a time, and extending the sample to the island regions omitted from the baseline. The results remain

¹⁶The BFK procedure instruments the utilization proxy with oil price and monetary policy surprises (Jarociński and Karadi 2020). If these instruments affect TFP directly, as evidence suggests for monetary policy (see Jordà et al. 2024; Ma and Zimmermann 2023), the exclusion restriction is violated.

¹⁷We consider further alternative TFP series. First, Appendix Table E.2 replaces the perpetual-inventory capital stock we constructed with the regional estimates of Gardiner et al. (2020); the results remain robust. Second, Appendix C.3 replicates the exercise using national-level TFP. The response is positive but imprecise, reflecting the loss of cross-sectional variation. Lastly, our baseline result is robust to alternative parametrizations of the Imbs utilization adjustment (Appendix D.3).

¹⁸The result is also robust to using lagged time-varying share $s_{i,t-1}$, which allows the regional elasticity to national spending to vary over time (Appendix Table E.2).

¹⁹Appendix Table E.2 also confirms robustness to simultaneously excluding Ireland and Portugal, the two program countries with the most severe fiscal adjustments, and to controlling for active IMF arrangements.

materially unchanged (Appendix Table E.1 and Table E.2), including when we use Driscoll and Kraay (1998) standard errors that account for cross-sectional dependence.

Table 2: Robustness Checks

	$h = 0$	$h = 1$	$h = 2$	$h = 3$	$h = 4$
Baseline	1.03*** (0.24)	1.93*** (0.45)	2.10*** (0.50)	1.63*** (0.42)	1.61*** (0.51)
Panel A: Productivity measures					
(A1) BFK-adjusted TFP	0.65*** (0.25)	1.07** (0.43)	0.86* (0.49)	0.94** (0.44)	1.11*** (0.43)
(A2) Cost-share TFP	1.31** (0.59)	2.52*** (0.95)	3.15*** (1.19)	2.17** (1.00)	1.20 (1.02)
(A3) Labor productivity	0.87*** (0.25)	1.75*** (0.41)	1.64*** (0.46)	0.84** (0.38)	0.65 (0.51)
Panel B: Exposure share measures					
(B1) GDP per capita share	1.02*** (0.23)	1.91*** (0.42)	2.11*** (0.50)	1.71*** (0.43)	1.71*** (0.50)
(B2) Total gov. spending share (levels)	0.82*** (0.15)	1.46*** (0.37)	1.48*** (0.42)	1.24*** (0.28)	1.34** (0.57)
(B3) Gov. spending as share of regional output	1.08*** (0.28)	2.06*** (0.51)	2.24*** (0.53)	1.63*** (0.44)	1.58*** (0.56)
Panel C: Rotemberg weights					
(C1) Drop bottom 1% Rotemberg weights	0.85*** (0.18)	1.58*** (0.34)	1.65*** (0.38)	1.27*** (0.32)	1.31*** (0.38)
(C2) Drop top 1% Rotemberg weights	1.21** (0.54)	2.90*** (1.10)	3.88** (1.54)	3.03** (1.32)	2.81** (1.36)
Panel D: Identification					
(D1) Control for tradable share	1.03*** (0.24)	1.92*** (0.47)	2.08*** (0.55)	1.63*** (0.53)	1.62** (0.64)
(D2) Additional controls	0.70*** (0.26)	1.56*** (0.49)	1.50*** (0.55)	1.00* (0.55)	1.17 (0.82)

Notes: The dependent variable is utilization-adjusted TFP growth except in rows A1–A3. *Panel A:* (A1) replaces the Imbs utilization adjustment with the IV-based correction of Basu et al. (2006); (A2) uses cost-share-based input elasticities following Comin et al. (2025) while retaining the Imbs adjustment; (A3) uses private-sector labor productivity (GVA per hour) as the dependent variable. *Panel B:* (B1)–(B3) reconstruct the exposure share s_i using the regional GDP per capita share (B1), the total government spending share in levels (B2), and government spending as a share of regional output (B3). *Panel C:* Rotemberg weights follow Goldsmith-Pinkham et al. (2020); (C1)–(C2) drop the region-year observations in the bottom and top 1% of the weight distribution. *Panel D:* (D1) adds the interaction of the shock with the lagged tradable GVA share (demeaned), instrumented by the interaction with $z_{i,t}$; the coefficient shown is the effect at the mean tradable share. (D2) adds lags of GDP per capita, capital stock, and hours worked growth. Standard errors clustered at the region level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Identification A causal interpretation requires that the exposure shares s_i not be correlated with regional characteristics that independently drive the TFP response. The most prominent candidate is sectoral composition: s_i is negatively correlated with the tradable sector GVA share ($\rho = -0.46$), and non-tradable sectors respond more strongly to austerity shocks (Section 4.4).²⁰ Row D1 addresses this directly by augmenting the baseline with the interaction $\tilde{s}_i \times \Delta G_{i,t}$, instrumented by $\tilde{s}_i \times z_{i,t}$, where $\tilde{s}_i$ is the lagged tradable GVA share, demeaned so that the reported coefficient is the effect at the mean share. The coefficient on government spending is virtually unchanged and significant at all horizons, confirming that sectoral composition does not confound the baseline.

To guard against the possibility that pre-existing regional trends drive the result, Row D2 adds lags of the growth rates of GDP per capita, the capital stock, and hours worked as controls and the estimates are consistent with the baseline. Row D2 adds lags of the growth rates of GDP per capita, the capital stock, and hours worked to guard against pre-existing regional trends; the estimates match the baseline. Appendix Table E.2 reports further checks: a GFC dummy interacted with the instrument, addressing the concern that high-exposure share regions were hit harder in 2008–09 and followed a different recovery path; the national debt-to-GDP ratio and sovereign spreads, capturing financing conditions that may drive both consolidation and productivity; and IMF-program participation, capturing the structural reforms bundled with those programs. The results are broadly unchanged throughout.

The validity of the national *austerity shock* is addressed in Appendix E.3. To account for the predictability critique of Jordà and Taylor (2016), we replace the shock with the residual from a first-step regression on lagged macroeconomic controls. The results closely match the baseline. We then run two falsification exercises. First, we replace the austerity shock with oil supply news shocks and with monetary policy surprises. Both yield near-zero first-stage F-statistics, confirming that these common shocks have no predictive power for regional government spending, and the associated TFP estimates are correspondingly large and imprecise. Second, we randomly permute the narrative consolidation episodes across country–year observations, preserving the distribution of shocks while breaking their true assignment. As shown in Appendix Table E.3, Panel C, the permuted instrument has no predictive power for government spending: the Kleibergen–Paap F-statistic falls to about 0.1, against 26 in the baseline, and the TFP response is insignificant at all horizons. The baseline results are therefore driven by the actual timing and country assignment of the consolidation episodes.

²⁰We regress s_i on within-sample time averages of regional characteristics. Beyond the tradable share, we also test for correlations with average changes in employment and in the unemployment rate, and with the public sector share. None is significantly correlated with s_i , suggesting that these regional characteristics do not confound the identifying variation in our shift-share instrument.

4.3 Austerity-recessions and non-austerity recessions

Our empirical results reveal that fiscal consolidations raise TFP while lowering economic activity. Whether downturns in general raise productivity is, however, an open question. The “cleansing” view, which dates back to Schumpeter (1942) and was formalized by Caballero and Hammour (1994), holds that recessions disproportionately eliminate low-productivity firms and thereby raise the average efficiency of surviving producers (see also Osotimehin and Pappadà 2017). The theoretical literature also identifies countervailing forces: Barlevy (2002) shows that recessions can have a “sullyng” effect, as impaired on-the-job search prevents reallocation toward better matches and can offset cleansing gains. The productivity gains we estimate could therefore reflect a generic recession effect rather than anything specific to austerity. In the following, we investigate whether austerity-driven recessions produce different productivity outcomes than other downturns, and find that TFP growth is systematically higher in recessions that coincide with fiscal consolidations than in other downturns.

We extend our baseline specification (7) and estimate by OLS, for each horizon $h = 0, \dots, 4$, the following regression, in which the narrative instrument is used only to classify observations into states:


$$\begin{aligned}
 z_{i,t+h} = & I_{i,t}^{ra} \left[\zeta_h^{ra} + \beta_h^{ra} \frac{G_{i,t} - G_{i,t-1}}{G_{i,t-1}} + \gamma_h^{ra}(L) X_{i,t-1} \right] \\
 & + I_{i,t}^r \left[\zeta_h^r + \beta_h^r \frac{G_{i,t} - G_{i,t-1}}{G_{i,t-1}} + \gamma_h^r(L) X_{i,t-1} \right] \\
 & + (1 - I_{i,t}^{ra} - I_{i,t}^r) \left[\zeta_h^e + \beta_h^e \frac{G_{i,t} - G_{i,t-1}}{G_{i,t-1}} + \gamma_h^e(L) X_{i,t-1} \right] + \alpha_{i,h} + \delta_{t,h} + u_{i,t+h}.
 \end{aligned} \tag{10}$$

$I_{i,t}^{ra}$ is a dummy variable that takes a value of one in year t when region i contemporaneously experiences negative per capita GVA growth and implements fiscal consolidation measures (i.e., when the Bartik instrument is larger than zero). On the other hand, $I_{i,t}^r$ is a dummy variable that takes a value of one when the regional per capita GVA growth rate is negative, and we identify no fiscal consolidation (when our Bartik instrument is equal to zero). Thus, $I_{i,t}^{ra}$ captures recessions that coincide with austerity (“austerity-recessions”), and $I_{i,t}^r$ measures economic downturns that are not directly related to fiscal consolidations but can be described as a combination of different negative shocks that lead to lower economic activity (“non-austerity recessions”). We also include a dummy for all remaining episodes when there is positive economic growth (economic expansions), $1 - I_{i,t}^{ra} - I_{i,t}^r$, to use the entire variation of the sample. The coefficients ζ_h^{ra} and ζ_h^r capture the average impact of austerity recessions and non-austerity recessions, respectively, on productivity

growth. In addition, β_h^{ra} and β_h^r indicate the marginal effect of lowering regional government spending by 1% in austerity recessions and non-austerity recessions, respectively. The control vector $X_{i,t-1}$ contains two lags of TFP growth, of government spending growth, and of per capita GVA growth, each interacted with the state indicators. If ζ_h^{ra} is larger (smaller) than ζ_h^r , this would imply that economic downturns driven by fiscal consolidations lead to a larger (smaller) increase in productivity than other downturns. The same logic also applies to the marginal effect coefficients β_h^{ra} and β_h^r .

Figure 7 presents the estimation results, where the left panel shows the difference between β_h^{ra} and β_h^r and the right panel shows the difference between ζ_h^{ra} and ζ_h^r . The difference in the average recession effect ($\zeta_h^{ra} - \zeta_h^r$) is positive and statistically significant at all horizons, peaking at about 2 percentage points two years after the shock. TFP growth is on average higher in austerity recessions than in other recessions. Furthermore, the difference in the marginal coefficients is also positive at all horizons and statistically significant at most of them. Thus, TFP growth is systematically higher in recessions that coincide with fiscal consolidations than in other downturns, both on average and per unit of spending cut, indicating that our main results do not simply reflect a downturn effect common to recessions of any origin.

Figure 7: Difference in Responses between Austerity-recessions and Non-austerity Recessions



Notes: Panels 7a and 7b on the first row show the difference of the marginal and average effects between austerity recessions and normal recessions on the TFP growth estimated through Equation 10. Bands are 90% (dark) and 95% (light) confidence intervals.

4.4 Within-sector productivity vs. reallocation across sectors

That austerity recessions differ from ordinary downturns points to something specific in how fiscal consolidations withdraw demand. Two channels can account for the TFP increase. First, productivity may increase *within* sectors, for instance through operational and managerial

efficiency gains, restructuring, or a cleansing mechanism that reallocates activity toward more productive firms. Second, austerity may induce *reallocation* across sectors, whereby resources shift away from sectors that are highly exposed to government demand—directly, through public procurement, and indirectly, through the local spending of public-sector wage income—and toward sectors that are less dependent on public spending.

This reallocation mechanism is particularly relevant in the presence of heterogeneous exposure to the non-market sector. Using the EUREGIO regional input–output (IO) tables introduced in Section 2, we measure for each sector the share of private-sector gross output absorbed by the non-market sector.²¹ The results shown in Figure 8 reveal that non-tradables are substantially more exposed than tradables, indicating a much stronger dependence on government demand.²² Because exposure is measured region by region, this is not a feature of a few countries: the gap is present in nine of the ten countries, and roughly two thirds of the cross-regional variation in exposure is within rather than between countries (Figure F.1 and Tables F.1 and F.2). In particular, whereas the non-tradables sector supplies on average about 4% of its output to the non-market sector, the respective number for the tradables sector is about 2%; the corresponding shares of output sold to households are about 30 and 16 percent (Figure 8, panel b). Because non-tradables also exhibit persistently lower average annual TFP growth, -0.7 percent versus 1.6 percent for tradables, a contraction in government demand is likely to reduce the relative weight of non-tradables and shift activity toward higher-productivity-growth sectors. As a result, aggregate private-sector TFP may increase mechanically even in the absence of within-sector productivity improvements.


In what follows, we quantify the relative importance of within-sector productivity gains and between-sector reallocation using a standard within–between decomposition of the impulse response of aggregate TFP.

Decomposition framework. We focus on decomposing the TFP response for horizons $h = 0$ and $h = 1$, for which the first-order approximation underlying the within–between decomposition is most reliable. To illustrate the procedure, we present the decomposition for $h = 0$; the year-one case ($h = 1$) follows analogously.

²¹The IO tables are not part of ARDECO. EUREGIO (Thissen et al. 2018) reports, for each NUTS 2 region and year from 2000 to 2010, the sales of 14 producing sectors (NACE Rev. 1.1) to the industries and final-demand categories of every region. ARDECO instead provides industry-level data at the 10-sector NACE Rev. 2 aggregation, which we use for the sectoral TFP measures. The two classifications align at the level we use: EUREGIO’s agriculture, mining and energy, and five manufacturing sectors correspond to ARDECO sectors A and B–E; construction, distribution, hotels and restaurants, transport and communication, financial intermediation, and real estate and business services correspond to F, G–J, and K–N; and non-market services correspond to O–U. Throughout, we classify sectors A–E (Agriculture, forestry, and fishing; Mining and quarrying; Manufacturing; Electricity, gas, steam, and air-conditioning supply; Water supply, sewerage, waste management, and remediation activities) as tradable, and the remaining market sectors (NACE codes F–N) as non-tradable.

²²Using estimates for final government consumption by sector from the IO tables also shows higher exposure of the non-tradables sector.

Figure 8: Exposure to the Non-Market Sector and to Household Consumption



Notes: Panel (a): share of each sector's gross output sold as intermediate inputs to the non-market services sector (public administration, education, health, and other non-market services) of the same country. Panel (b): share of each sector's gross output sold to household final consumption in the same country. Source: EUREGIO regional input-output tables, 2000–2010; unweighted averages across the sample regions of the ten countries. The tradables sector comprises agriculture, mining, and industry, while the non-tradables sector comprises services, finance, and construction.

Let $s \in \{\text{tradable}, \text{non-tradable}\}$ index sectors. Aggregate private-sector TFP growth satisfies (Hulten 1978)

$$\frac{A_{i,t} - A_{i,t-1}}{A_{i,t-1}} \approx \sum_s D_{i,t}^s \frac{A_{i,t}^s - A_{i,t-1}^s}{A_{i,t-1}^s}, \quad (11)$$

where $A_{i,t}^s$ denotes sectoral TFP and $D_{i,t}^s$ is the sector's Domar weight. In our baseline estimation, the coefficient β_0 measures the impact response of aggregate TFP growth to a government spending shock. Sector-specific local projections (using the same specification) yield sectoral impact responses β_0^s , and a local projection for the weights $D_{i,t}^s$ delivers δ_0^s , the impact response of sectoral Domar weights. Taking a first-order approximation of the aggregation identity with respect to the shock and evaluating it at sample averages of Domar weights and TFP-growth rates yields


$$\beta_0 \approx \underbrace{\sum_s \bar{D}^s \beta_0^s}_{\text{within-sector response}} + \underbrace{\sum_s \frac{\bar{A}_{i,t}^s - \bar{A}_{i,t-1}^s}{\bar{A}_{i,t-1}^s} \delta_0^s}_{\text{between-sector (reallocation) response}}. \quad (12)$$

The first term captures how the shock affects aggregate TFP by changing sectoral TFP growth, holding sectoral weights fixed at their sample averages $\bar{D}^s$. The second term captures how the induced reallocation of Domar weights across tradable and non-tradable activities contributes to the aggregate impact response.²³

²³Formally, define aggregate and sectoral TFP growth over $[t-1, t+h]$ as $\psi_{i,t+h}(\varepsilon_t) \equiv (A_{i,t+h} - A_{i,t-1})/A_{i,t-1}$

Within-sector productivity responses. Figure 9 reports sectoral impulse responses of TFP following an austerity shock. Both sectors' TFP rises after the shock. The tradable sector responds gradually, reaching roughly 2 percent four years after the shock. The non-tradable response is front-loaded: about 1 percent on impact and peaking near 2 percent after one year.²⁴

Figure 9: TFP Responses in Tradable and Non-Tradable Sectors



Notes: Bands are 90% (dark) and 95% (light) confidence intervals.

Sectoral shares and reallocation. Sectoral weights are constructed using value-added shares: $D_{i,t}^s = \text{ValueAdded}_{i,t}^s / \text{Total Value Added}_{i,t}$.²⁵ Figure 10 shows that the non-tradable value-added share declines following an austerity shock, consistent with a shift in activity away from government-exposed sectors.


Quantifying within and between contributions. Table 3 summarizes the decomposition expressed as fractions of the overall TFP response. The rows of “within-sector TFP” show the contribution of each sector relative to the total effect, while the “Total within-sector contribution” sums these shares between the tradable and non-tradable sectors. The “Reallocation (Shares)” rows analogously report the proportion of the total TFP response attributable to changes in

and $\psi_{i,t+h}^s(\varepsilon_t) \equiv (A_{i,t+h}^s - A_{i,t-1}^s) / A_{i,t-1}^s$, where we make explicit the dependence on the spending shock ε_t . The aggregation result implies $\psi_{i,t+h}^s(\varepsilon_t) \approx \sum_s D_{i,t}^s(\varepsilon_t) \psi_{i,t+h}^s(\varepsilon_t)$. Differentiating both sides with respect to ε_t gives $\frac{\partial \psi_{i,t+h}^s}{\partial \varepsilon_t} \approx \sum_s \frac{\partial D_{i,t}^s}{\partial \varepsilon_t} \psi_{i,t+h}^s + \sum_s D_{i,t}^s \frac{\partial \psi_{i,t+h}^s}{\partial \varepsilon_t}$. Evaluating this expression at the sample means $(\bar{D}^s, \bar{\psi}_h^s)$, links the aggregate impulse response at horizon h to the sectoral impulse responses and the response of the Domar weights: $\beta_h \approx \sum_s \bar{\psi}_h^s \delta_h^s + \sum_s \bar{D}^s \beta_h^s$.

²⁴Appendix F presents TFP and labor-productivity impulse responses at a finer sectoral level. The productivity increase is larger and earlier in non-tradable sectors, particularly construction and trade, transport and information services, while within tradables the TFP response of industry builds up gradually.

²⁵Ideally, weights would be based on gross output, $D_{i,t}^{s,GO} = \text{GrossOutput}_{i,t}^s / \text{Total Value Added}_{i,t}$, because gross output internalizes intermediate-input linkages (Hulten 1978; Baqaee and Farhi 2019). However, sector-level gross-output data are not available for NUTS 2 regions over the full sample. Using value-added shares may, therefore, understate the reallocation due to input-output linkages.

Figure 10: Value-Added Share Response of Non-Tradable Sectors



Notes: Impulse response of the non-tradable sector's share in private-sector value added. Bands are 90% (dark) and 95% (light) confidence intervals.

sectoral weights, with the sum across sectors yielding the “Total reallocation contribution”. At both horizons, the within-sector component accounts for essentially the entire TFP effect (99 percent), about three quarters of which comes from non-tradables. The reallocation component is negligible (about 1 percent): the shift in value-added shares in Figure 10, of two to four tenths of a percentage point, is correctly signed but far too small to move aggregate TFP through composition. The two terms together account for about 85 percent of the estimated aggregate response; the remainder reflects the first-order approximation and the smaller sample for which sectoral TFP is available.

Table 3: Decomposition of the Aggregate TFP Response


Horizon		Impact $h = 1$	
Within-sector TFP	Tradables	0.23	0.15
	Non-tradables	0.77	0.85
	<i>Total within-sector contribution</i>	<i>0.99</i>	<i>0.99</i>
Reallocation (Shares)	Tradables	0.00	0.00
	Non-tradables	0.00	0.00
	<i>Total reallocation contribution</i>	<i>0.01</i>	<i>0.01</i>

Notes: Values are derived from a first-order approximation as in (12) and expressed as a fraction of the total TFP effect. The within-sector component sums sectoral TFP responses weighted by average sectoral shares. The reallocation component captures changes in aggregate TFP due to shifts in sectoral value-added shares, holding sectoral TFP constant.

4.5 Austerity shocks and firm exit

The decomposition raises the question of why productivity rises within sectors. Cleansing is the natural candidate (Section 4.3). We examine firm dynamics to assess whether it is operative here. We measure net firm exit as the cumulative percent decline in the number of local units, using Eurostat’s regional Structural Business Statistics, which report local units by NACE section from 2008 onwards. The estimation window is therefore 2008–2019. The statistics do not cover agriculture or financial services, so tradables are NACE sections B–E and non-tradables F–J and L–N, the paper’s sector split without these two sections. We use the region-years in which all sections of both groups are reported, so that the two sectors are compared on the same 274 region-years; most German regions are not covered.²⁶ Using the same local-projection specification as in the baseline analysis with this outcome, we estimate impulse responses of net firm exit separately for tradable and non-tradable sectors. Figure 11 displays the results.

Figure 11: Number of Local Units by Sectors



Notes: Eurostat SBS local units, 2008–2019; tradables = NACE B–E, non-tradables = F–J and L–N; both panels are estimated on the same 274 region-years (see text). Impulse responses to a one percent austerity-driven decline in government spending, estimated as in equation 7. Bands are 90% (dark) and 95% (light) confidence intervals.

Austerity shocks lead to a net decline in local units in both sectors. The magnitudes are economically meaningful: after three years, the number of local units declines by roughly 6 percent in the non-tradable sectors and by about 2 percent in the tradable sectors. These differences in firm dynamics parallel the sectoral TFP responses documented above. The non-tradable sector exhibits both a sharper decline in the number of local units and a larger increase in TFP, with both responses sizable on impact. This joint pattern is consistent with a cleansing mechanism in which the shock accelerates the exit of less productive units and raises average sectoral productivity,

²⁶Appendix Figure F.7 shows that the results are similar when each sector uses all region-years available for it (330 for tradables and 313 for non-tradables).

particularly in sectors more exposed to the austerity shock.

We emphasize, however, that our data do not permit a clean separation between cleansing and other within-sector channels (such as within-firm efficiency gains, labor reallocation toward more productive uses, and technology adoption). Establishment-level productivity measures would be required for this purpose. Nonetheless, the combination of net exit and rising sectoral productivity suggests that cleansing dynamics plausibly play a role in driving the within-sector component of the regional TFP response.

4.6 The composition of government spending

Our results complement the existing literature finding that government spending *increases* raise TFP when public spending shifts toward R&D (Antolin-Diaz and Surico 2025) or through learning-by-necessity (Ilzetzki 2024). We find that a spending *contraction* can also raise TFP, in a setting where the demand withdrawn falls mainly on non-tradable sectors, which tend to have lower R&D intensity (Galindo-Rueda and Verger 2016). The direction of the fiscal impulse is not sufficient to sign the productivity effect; its composition and the horizon matter as well.

Our estimates of research activity are consistent with this reading. Following a spending contraction, R&D spending shows little response at short horizons and declines only later: by the fourth year, total R&D expenditure is about 20 percent lower and patent applications 10 to 16 percent lower, with the decline precisely estimated for public R&D and too imprecisely for private R&D to draw conclusions (Appendix Figures G.1 and G.2).²⁷ Therefore, the innovation channel presented by Antolin-Diaz and Surico (2025) for defense spending expansions does not account for the productivity gains we estimate; these gains arise within the four-year window, before research activity moves, consistent with selection. The two sets of results are complementary rather than contradictory: spending expansions raise productivity through innovation with long lags, whereas consolidations raise measured productivity quickly, consistent with cleansing, while reducing innovation inputs and outputs with a lag. Whether the latter eventually weighs on productivity beyond our horizon is a question our data cannot answer.

5 Conclusion

This paper provides causal evidence that fiscal austerity raises productivity at the regional level in the euro area, despite a contraction in economic activity. Decompositions indicate that the aggregate effect is almost entirely due to within-sector improvements; austerity does shift resources

²⁷Part of the public R&D response is mechanical: the compensation of R&D personnel in government and higher education enters non-market GVA and therefore our spending proxy.

toward tradable sectors with higher productivity growth, but this reallocation contributes little to aggregate TFP. The within-sector gains are consistent with a cleansing mechanism. Productivity growth is higher in recessions that coincide with austerity than in other downturns.

The results speak against the view that consolidation is inherently detrimental to productivity. Gains can arise when the demand withdrawn falls on sectors that are relatively unproductive and less R&D-intensive. They are not costless, however: within four years, consolidations also reduce R&D expenditure and patent applications, so the medium-run gains from selection may come at the expense of the long-run innovation base. The productivity consequences of fiscal policy therefore depend on the composition of the spending that changes and on the horizon considered.

References

- Adler, G., C. Allen, G. Ganelli, and D. Leigh (2024). An updated action-based dataset of fiscal consolidation. *IMF working papers WP/24/210*.
- Alesina, A. and S. Ardagna (2010). Large changes in fiscal policy: taxes versus spending. *Tax policy and the economy* 24(1), 35–68.
- Alesina, A., C. Favero, and F. Giavazzi (2015). The output effect of fiscal consolidation plans. *Journal of International Economics* 96, S19–S42.
- Alesina, A., C. Favero, and F. Giavazzi (2020). *Austerity: When it Works and when it Doesn't*. Princeton University Press.
- Antolin-Diaz, J. and P. Surico (2025). The long-run effects of government spending. *American Economic Review* 115(7), 2376–2413.
- Anzoategui, D., D. Comin, M. Gertler, and J. Martinez (2019). Endogenous technology adoption and r&d as sources of business cycle persistence. *American Economic Journal: Macroeconomics* 11(3), 67–110.
- Auerbach, A. J., Y. Gorodnichenko, and D. Murphy (2020). Macroeconomic frameworks: Reconciling evidence and model predictions from demand shocks. *NBER Working Paper* 26365.
- Baqae, D. R. and E. Farhi (2019). The macroeconomic impact of microeconomic shocks: Beyond hulten's theorem. *Econometrica* 87(4), 1155–1203.
- Bardaka, I., I. Bournakis, and G. Kaplanoglou (2021). Total factor productivity (tfp) and fiscal consolidation: How harmful is austerity? *Economic Modelling* 94, 908–922.
- Barlevy, G. (2002). The sullyng effect of recessions. *Review of Economic Studies* 69(1), 65–96.
- Bartik, T. J. (1991). Who benefits from state and local economic development policies? *WE Upjohn Institute for Employment Research*.
- Basu, S., J. G. Fernald, and M. S. Kimball (2006). Are technology improvements contractionary? *American Economic Review* 96(5), 1418–1448.
- Basu, S. and M. S. Kimball (1997). Cyclical productivity with unobserved input variation.
- Bernardini, M., S. D. Schryder, and G. Peersman (2020). Heterogeneous government spending multipliers in the era surrounding the great recession. *Review of Economics and Statistics*, 1–45.
- Beugelsdijk, S., M. J. Klasing, and P. Millionis (2018). Regional economic development in europe: the role of total factor productivity. *Regional Studies* 52(4), 461–476.
- Born, B., , G. Müller, J. Pfeifer, and S. Wellmann (2020). Different no more: Country spreads in advanced and emerging economies. *CEPR Discussion Papers* 14392.
- Borusyak, K., P. Hull, and X. Jaravel (2022). Quasi-experimental shift-share research designs. *Review of Economic Studies* 89(1), 181–213.
- Borusyak, K., P. Hull, and X. Jaravel (2025). A practical guide to shift-share instruments. *Journal of Economic Perspectives* 39(1), 181–204.
- Brueckner, M., E. Pappa, and A. Valentinyi (2023). Geographic cross-sectional fiscal spending multipliers and the role of local autonomy: Evidence from european regions. *Journal of Money,*

- Credit and Banking* 55(6), 1357–1396.
- Burnside, C., M. Eichenbaum, and S. Rebelo (1995). Capital utilization and returns to scale. *NBER macroeconomics annual* 10, 67–110.
- Caballero, R. J. and M. L. Hammour (1994). The cleansing effect of recessions. *American Economic Review* 84(2), 135–139. Papers and Proceedings.
- Canova, F. and E. Pappa (2025). The macroeconomic effects of eu regional structural funds. *Journal of the European Economic Association* 23(1), 327–360.
- Chen, X., S. Elsherbiny, R. D. Gabriel, Ò. Jordà, C. H. Kim, M. Schularick, and A. M. Taylor (2022). Documentation on the jst database update 2016–2020, 6th release, july 2022.
- Chodorow-Reich, G., L. Feiveson, Z. Liscow, and W. G. Woolston (2012). Does state fiscal relief during recessions increase employment? evidence from the american recovery and reinvestment act. *American Economic Journal: Economic Policy* 4, 118–45.
- Coibion, O., Y. Gorodnichenko, L. Kueng, and J. Silvia (2017). Innocent bystanders? monetary policy and inequality. *Journal of Monetary Economics* 88, 70–89.
- Comin, D., J. Quintana, T. Schmitz, and A. Trigari (2025). Revisiting productivity dynamics in europe: A new measure of utilization-adjusted tfp growth. *Journal of the European Economic Association* 23(4), 1598–1633.
- Cournède, B., A. Goujard, Álvaro Pina, and A. de Serres (2013). Choosing fiscal consolidation instruments compatible with growth and equity. OECD Economic Policy Papers 7, OECD Publishing.
- Devries, P., J. Guajardo, D. Leigh, and A. Pescatori (2011). A new action-based dataset of fiscal consolidation.
- Driscoll, J. and A. Kraay (1998). Consistent covariance matrix estimation with spatially dependent panel data. *The Review of Economics and Statistics* 80, 549–560.
- European Fiscal Board (2019, Aug). Assessment of eu fiscal rules with a focus on the six and two-pack legislation. Reports 2019, European Fiscal Board.
- Fernald, J. G. and J. C. Wang (2016). Why has the cyclical of productivity changed? what does it mean? *Annual Review of Economics* 8(1), 465–496.
- Gabriel, R. D., M. Klein, and A. S. Pessoa (2023). The effects of government spending in the eurozone. *Journal of the European Economic Association* 21(4), 1397–1427.
- Gabriel, R. D., M. Klein, and A. S. Pessoa (2026, 01). The political costs of austerity. *The Review of Economics and Statistics* 108(1), 145–161.
- Galindo-Rueda, F. and F. Verger (2016). OECD taxonomy of economic activities based on R&D intensity. OECD Science, Technology and Industry Working Papers 2016/04, OECD Publishing, Paris.
- Gardiner, B., B. Fingleton, and R. Martin (2020). Regional disparities in labour productivity and the role of capital stock. *National Institute Economic Review* 253, R29–R43.
- Goldsmith-Pinkham, P., I. Sorkin, and H. Swift (2020, August). Bartik instruments: What, when, why, and how. *American Economic Review* 110(8), 2586–2624.

- Gorodnichenko, Y. and B. Lee (2020). Forecast error variance decompositions with local projections. *Journal of Business & Economic Statistics* 38(4), 921–933.
- Guajardo, J., D. Leigh, and A. Pescatori (2014). Expansionary austerity? international evidence. *Journal of the European Economic Association* 12(4), 949–968.
- Hall, R. E. and D. W. Jorgenson (1967). Tax policy and investment behavior. *The American Economic Review* 57(3), 391–414.
- Hulten, C. R. (1978). Growth accounting with intermediate inputs. *The Review of Economic Studies* 45(3), 511–518.
- Huo, Z., A. A. Levchenko, and N. Pandalai-Nayar (2023). Utilization-adjusted TFP across countries: Measurement and implications for international comovement. *Journal of International Economics*, 103753.
- Ilzetzki, E. (2024). Learning by necessity: Government demand, capacity constraints, and productivity growth. *American economic review* 114(8), 2436–2471.
- Imbs, J. M. (1999). Technology, growth and the business cycle. *Journal of Monetary Economics* 44(1), 65–80.
- IMF (2023). *Navigating Global Divergences*. World Economic Outlook.
- Jarociński, M. and P. Karadi (2020). Deconstructing monetary policy surprises—the role of information shocks. *American Economic Journal: Macroeconomics* 12(2), 1–43.
- Jordà, Ò., S. R. Singh, and A. M. Taylor (2024). The long-run effects of monetary policy. *Review of Economics and Statistics*, 1–49.
- Jordà, Ò. and A. M. Taylor (2016). The time for austerity: estimating the average treatment effect of fiscal policy. *The Economic Journal* 126(590), 219–255.
- Lequiller, F. and D. W. Blades (2014). *Understanding national accounts, second edition*. OECD Paris.
- Ma, Y. and K. Zimmermann (2023). Monetary policy and innovation. Technical report, National Bureau of Economic Research.
- Nakamura, E. and J. Steinsson (2014). Fiscal stimulus in a monetary union: Evidence from us regions. *American Economic Review* 104(3), 753–92.
- OECD (2015). *The Future of Productivity*. Paris: Organization for Economic Co-operation and Development.
- Olea, J. L. M. and C. Pflueger (2013). A robust test for weak instruments. *Journal of Business & Economic Statistics* 31(3), 358–369.
- Osoimehin, S. and F. Pappadà (2017). Credit frictions and the cleansing effect of recessions. *The Economic Journal* 127(602), 1153–1187.
- Romer, C. D. and D. H. Romer (2010). The macroeconomic effects of tax changes: estimates based on a new measure of fiscal shocks. *American Economic Review* 100(3), 763–801.
- Schumpeter, J. A. (1942). *Capitalism, Socialism and Democracy*. New York: Harper & Brothers.
- Siller, M., T. Schatzer, J. Walde, and G. Tappeiner (2021). What drives total factor productivity growth? an examination of spillover effects. *Regional Studies* 55(6), 1129–1139.

Thissen, M., M. Lankhuizen, F. van Oort, B. Los, and D. Diodato (2018). Euregio: The construction of a global IO database with regional detail for europe for 2000–2010.

APPENDIX

Appendix A Data Appendix

ARDECO - Regional European Data ARDECO is the Annual Regional Database of the European Commission’s Directorate General for Regional and Urban Policy and is maintained and updated by the Joint Research Centre. It is a highly disaggregated dataset across both sectoral and sub-regional dimensions. The database contains a set of long time-series indicators for EU regions at various statistical scales (NUTS 0, 1, 2, and 3 level) using the NUTS 2021 regional classification. The dataset includes data on demography, labor markets, capital formation and domestic product by six sectors. The six sectors are (1) agriculture, forestry and fishing; (2) industry excluding construction; (3) construction; (4) wholesale, retail, transport, accommodation, and food services; (5) information and communication, financial and business services; and (6) non-market services.

ARDECO data is an annual unbalanced panel covering the period 1980–2023 for the European Union (EU) and some European Free Trade Association (EFTA) and candidate countries; our sample uses data from 1999–2019. Its main data source is Eurostat (the Statistical Office of the European Commission), supplemented, where necessary, by other appropriate national and international sources. ARDECO is constructed in such a way that the country aggregates of its various time series equal the corresponding time series in the AMECO dataset referring to the National Accounts. Starting in 2002, euro area countries have published national series in EUR. National currency data for all years prior to the switch of the country to EUR have been converted using the irrevocably fixed EUR conversion rate. We use the series with real variables expressed in 2015 constant prices in EUR.

Regional government spending measure We now explain in detail why GVA (plus intermediate consumption of the non-market sector) is a valid proxy for government spending following the argument in Gabriel et al. (2023).

First, as previously mentioned, ARDECO’s regional data is consistent with the national accounts data by construction. By definition, there exists a close link between government spending and the GVA of the non-market sector, however, they differ in two dimensions: actors and composition. Regarding the first, even though the non-market sector includes other institutional units, the general government is the main actor responsible for changes in non-market GVA.

In particular, the non-market sector consists of six sub-sectors: “Public administration and defense”, “Education”, “Human health and social work”, “Arts, entertainment and recreation”, “Other service activities,” and “Activities of household and extra-territorial organizations and bodies.” The first sub-sector, “Public administration and defense,” refers to activities by the general

government, but not all government bodies are automatically classified under this sub-sector. For example, a secondary school administered by the central or local government is classified as “Education,” and a public hospital is allocated to “Human health and social work.” Thus, the two sub-sectors “Education” and “Human health and social work” are also closely linked to the general government in the national accounts, while the last three sub-sectors are linked only loosely.

Concerning the second dimension, we now describe the compositional differences between non-market GVA and government spending. In the national accounts, government spending is defined as follows:

$$\begin{array}{l}
 \hline
 \text{Final consumption expenditure of the general government} \\
 = \text{Gross value added of the general government} \\
 + \text{Intermediate inputs of the general government} \\
 + \text{Social transfers in kind purchased market production} \\
 - \text{Market output and output for own final use} \\
 - \text{Payments for non-market output} \\
 \hline
 \end{array}$$

GVA of the general government is the major component of government spending and fully accounted in the GVA of the non-market sector. Country level data show that GVA of the general government accounts for approximately 71.1% of the final consumption expenditure of the general government.¹ Thus, our proxy measures the single-most dominant source of government expenditures. However, the main difference between government spending and the GVA of the general government are the intermediate inputs and the social transfers in kind. When again looking at country level data, we find that GVA and intermediate consumption account for about 99% of government spending.

To include intermediate consumption in our government spending measure, we use input-output tables from the PBL EUREGIO database. This is the first time-series (annual, 2000–2010) of global IO tables with regional detail for the entire large trading bloc of the European Union. This database allows for a regional analysis at the NUTS 2 level consistent with our baseline method. The tables merge data from World Input-Output Database (the 2013 release) with regional economic accounts and inter-regional trade estimates developed by PBL Netherlands Environmental Assessment Agency and complemented with survey-based regional input-output data for a limited number of countries. All data used are survey data, and only non-behavioral assumptions have been made to estimate the EUREGIO dataset. These two general rules of data construction allow empirical analyses focused on impacts of changes in behavior without endogenously having this

¹According to data collected from [Eurostat](#) for the sample between 1995 and 2016.

behavior embedded already by construction. More detailed information can be found in Thissen et al. (2018).

The PBL EUREGIO database provides estimates for intermediate consumption of the non-market sector at the NUTS 2 level from 2000–2010. Thus, we adjust regional GVA of the non-market sector by a region-specific time-invariant scaling factor to include intermediate consumption in our government spending measure to obtain our proxy for regional government spending.

Table A.1: NUTS 2 Regions in Final Sample

Country	NUTS 2 Regions
Austria	9
Belgium	11
Finland	4
France	21
Germany	38
Ireland	3
Italy	21
Netherlands	12
Portugal	5
Spain	15
Total	139

Notes: Our final dataset excludes the French overseas regions of Guadeloupe (FRY1), Martinique (FRY2), Guyane (FRY3), La Réunion (FRY4), and Mayotte (FRY5), as well as the Spanish autonomous cities of Ceuta (ES63) and Melilla (ES64), the Finnish autonomous region of Åland (FI20), Corsica (FRM0), the Spanish island regions of Illes Balears (ES53) and Canarias (ES70), and the Portuguese autonomous regions of Açores (PT20) and Madeira (PT30).

Sample and Variables Description

Table A.2: Variables Description

Variable Name	Computation	Definition [Source]
Capital Stock (Private-sector)	$K_{t-1}^p + I_t^p - D_t^p$	Private-sector capital stock estimated via perpetual inventory method [ARDECO]
Capital Utilization	See Equation (5)	Rate of capital utilization derived from firm optimality conditions [ARDECO]
Debt-to-GDP		Gross general government debt as a share of GDP, national [OECD]
Depreciation		Consumption of fixed capital in private sector [ARDECO]
Domar Weights	Sectoral GVA / Private-sector GVA	Sectoral value-added shares used in TFP decomposition [ARDECO]
Employment		Total employment [ARDECO]
Exposure to Non-market Sector	Sales to non-market industry / gross output	Share of a sector's gross output sold to the non-market sector of the same country, by region and sector, 2000–2010 [PBL EUREGIO]
Exposure to Households	Sales to households / gross output	Share of a sector's gross output sold to household final consumption in the same country [PBL EUREGIO]
GDP _{pc}	GDP / population	Regional gross domestic product per capita [ARDECO]
Gov. Spending _{pc}	(Non-market GVA + intermediate consumption) / population	Regional government spending proxy per capita [ARDECO, PBL EUREGIO]
Hours Worked (Private-sector)		Total hours worked in private sector [ARDECO]
Hourly Wage	Private-sector compensation / private-sector hours worked	Regional average compensation per hour in the private sector [ARDECO]
IMF Program		Indicator for an active IMF arrangement in the country [IMF MONA]
Intermediate Consumption		Non-market sector intermediate consumption shares [PBL EUREGIO]
Investment _{pc}	Private-sector GFCF / population	Private-sector gross fixed capital formation per capita [ARDECO]
Investment Price Index		Price index for investment goods [ARDECO]
Labor Effort	See Equation (4)	Labor effort index derived from household optimality conditions [ARDECO]


Continued on next page

Table A.2 – continued from previous page

Variable Name	Computation	Definition [Source]
Labor Productivity	Private-sector (GVA / Hours)	Value added per hour worked [ARDECO]
Labor Share	Private-sector compensation / Private-sector GVA	private-sector labor income share [ARDECO]
Local Units		Number of local units (establishments) by NACE section, 2008 onwards [Eurostat SBS]
Narrative Austerity		Spending-based fiscal consolidation episodes [Adler et al. 2024]
Patent Applications		PCT patent applications per million inhabitants, by inventor region and priority year [OECD REGPAT]
Population		Total regional population [ARDECO]
Private-sector Consumption		Consumer spending of people living in the region [Oxford Economics]
Private-sector GVA		Gross value added of the private sector [ARDECO]
R&D Expenditure		Intramural R&D expenditure, total, business and government sectors, deflated to 2015 euros [Eurostat]
Solow Residual	See Equation (2)	Private-sector TFP growth without the utilization and effort adjustment [ARDECO]
Sovereign Spread		10-year government bond yield relative to Germany [OECD]
TFP (utilization-adjusted)	See Equation (6)	Total factor productivity adjusted for capital utilization and labor effort [ARDECO]
Unemployment Rate		Regional unemployment rate, used for the GFC-intensity control [Eurostat]

Notes: See Section 2 for detailed variable construction. We thank Ricardo Amaro and Anthony Light from Oxford Economics for sharing the data on private-sector consumer spending with us.

Figure A.1: Austerity by Country



Notes: Austerity episodes are identified according to Adler et al. (2024) and present the size of fiscal consolidations based on the previous year's GDP in percent.

Appendix B Utilization and Effort Adjustment: Model and Derivation

We adjust the TFP series for capital utilization and labor effort using a model-based procedure similar to that developed in Imbs (1999). This appendix presents the model and the derivation of the adjustment formulas used in the main text.

B.1 Model environment

Consider a representative household and firm. Time is discrete and infinite. Households consume c_t , supply hours H_t^p at effort level e_t , and invest in one-period bonds B_t , paying the nominal interest rate R_t . Preferences are given by GHH utility:

$$E_t \sum_{k=0}^{\infty} \beta^k u(c_{t+k}, H_{t+k}^p, e_{t+k}), \quad (\text{B.1})$$

$$u(c_t, H_t^p, e_t) = \frac{(c_t - \psi(H_t^p)^{1+v} - \psi_e e_t^{1+v_e})^{1-\sigma} - 1}{1-\sigma}, \quad (\text{B.2})$$

with parameters $\beta \in (0, 1)$, $v > 0$, $v_e > 0$, $\psi, \psi_e > 0$, and $\sigma > 0$. Consumption goods are purchased at price P_t .

The household's budget constraint is

$$c_t + \frac{B_t}{P_t} = \frac{W_t(e_t)}{P_t} H_t^p + \frac{B_{t-1} R_{t-1}}{P_t} + \frac{D_t}{P_t}, \quad (\text{B.3})$$

where $W_t(e_t)$ is the nominal wage inclusive of effort premia and D_t denotes firms' dividend payments.

B.2 Firms

Private-sector output is produced according to

$$Y_t^p = (e_t H_t^p)^{\alpha_t} (u_t K_t^p)^{1-\alpha_t}, \quad (\text{B.4})$$

where K_t^p is tangible private-sector capital and $u_t \in (0, 1)$ is the utilization rate. The elasticity α_t corresponds to the labor share in output.

Capital evolves according to

$$K_{t+1}^p = (1 - \delta_k) K_t^p + \Psi\left(\frac{K_{t+1}^p}{K_t^p}\right) K_t^p, \quad (\text{B.5})$$

where $\Psi(\cdot)$ is the adjustment-cost function, satisfying $\Psi(\tilde{g}) = \delta_k$, $\Psi'(\tilde{g}) = 1$, and $\Psi''(\tilde{g}) = \epsilon_k > 0$. Capital utilization generates maintenance costs $a(u_t)K_t^p$, where $a(u_t) = \psi_k u_t^{\beta_k}$ with $\psi_k, \beta_k > 0$.

Firms maximize expected discounted profits:

$$\max_{H_t^p, e_t, u_t, K_{t+1}^p} E_t \sum_{k=0}^{\infty} \tilde{\beta}_{t,t+k} \left[P_{t+k} Y_{t+k}^p - W_{t+k}(e_{t+k}) H_{t+k}^p - Q_{t+k} \Psi\left(\frac{K_{t+k+1}^p}{K_{t+k}^p}\right) K_{t+k}^p \right. \quad (\text{B.6})$$

$$\left. - Q_{t+k} a(u_{t+k}) K_{t+k}^p \right], \quad (\text{B.7})$$

where Q_t is the investment-goods price.

B.3 Effort adjustment

The household's first-order conditions imply:

$$\frac{W_t(e_t)}{P_t} = \psi(1+v)(H_t^p)^v, \quad (\text{B.8})$$

$$\frac{W'_t(e_t) H_t^p}{P_t} = \psi_e(1+v_e) e_t^{v_e}. \quad (\text{B.9})$$

Profit maximization implies:

$$\alpha_t P_t \frac{Y_t^p}{H_t^p} = W_t(e_t), \quad (\text{B.10})$$

$$\alpha_t P_t \frac{Y_t^p}{e_t} = W'_t(e_t) H_t^p. \quad (\text{B.11})$$

Combining these gives:

$$\frac{(H_t^p)^{1+v}}{e_t^{1+v_e}} = \frac{\psi_e}{\psi} \frac{1+v_e}{1+v}. \quad (\text{B.12})$$

Normalizing constants to one and setting $v = v_e = 1$ (following Jordà et al. 2024), we obtain the effort index:

$$e_t = (H_t^p)^{\frac{1+v}{1+v_e}} = H_t^p. \quad (\text{B.13})$$

B.4 Capital-utilization adjustment

From the firm's FOCs we obtain:

$$\alpha_t P_t \frac{Y_t^p}{u_t} = a'(u_t) K_t^p Q_t = \psi_k \beta_k u_t^{\beta_k-1} K_t^p Q_t. \quad (\text{B.14})$$

Combining with the labor FOC and rearranging yields:

$$u_t^{\beta_k} = \left(\frac{W_t(e_t)H_t^p}{K_t^p Q_t} \middle/ \frac{\overline{WH^p}}{\overline{K^p Q}} \right), \quad (\text{B.15})$$

and thus the proportional index:

$$\frac{u_t}{\bar{u}} = \left(\frac{W_t H_t^p}{K_t^p Q_t} \middle/ \frac{\overline{WH^p}}{\overline{K^p Q}} \right)^{1/\beta_k}. \quad (\text{B.16})$$

Trend values $\overline{WH^p}/\overline{K^p Q}$ are obtained using the HP filter. We set $\beta_k = 4$ following Anzoategui et al. (2019).

B.5 Utilization- and effort-adjusted TFP

Given the production function (B.4), utilization-adjusted TFP is defined as:

$$A_t = \frac{Y_t^p}{(u_t K_t^p)^{1-\alpha_t} (e_t H_t^p)^{\alpha_t}}. \quad (\text{B.17})$$

With time-varying α_t , following PWT we compute growth rates as:

$$\Delta \ln A_t = \Delta \ln Y_t^p - (1 - \alpha_t)(\Delta \ln K_t^p + \Delta \ln u_t) - \alpha_t(\Delta \ln H_t^p + \Delta \ln e_t). \quad (\text{B.18})$$

Appendix C Additional Results

C.1 Broader macroeconomic effects of austerity

Figure C.1 presents the responses of key macroeconomic variables to the austerity shock. The first row shows aggregate economic activity and employment: GDP declines significantly, reaching a trough of almost 2% one year after the shock, before slowly reverting to its pre-shock level. Total employment also falls significantly in the wake of the fiscal consolidation, though the effect dissipates over the medium term. Private-sector employment falls by more, about 3% after two years, and remains below its pre-shock level at the end of the horizon.

Figure C.1: Response of other macro-indicators to Austerity Shock




Notes: Impulse responses to a one percent austerity-driven decline in government spending, estimated as in equation 7. Bands are 90% (dark) and 95% (light) confidence intervals.

The second row turns to expenditure components. Private-sector investment, public investment, and private-sector consumption all contract significantly following the shock, with investment adjusting faster than consumption: private investment troughs after one year at about 8% and partly recovers, whereas consumption keeps falling to about 4% after three years. Public investment also declines, so consolidations cut the investment component of public spending as well. The third row documents the capital stock and the labor market: the private-sector capital stock declines gradually, and both hourly wages and hours worked fall significantly, indicating that austerity compresses labor income through both the extensive and intensive margins.

Taken together, the fiscal consolidation triggers a broad-based contraction in real economic activity, factor inputs, and labor income. Figure C.2 shows that the utilization adjustment is essential to read the productivity response: the unadjusted private-sector Solow residual responds by less than half a percent and insignificantly over the first three years, because the sharp fall in factor utilization masks the underlying efficiency gains that the adjusted measure recovers.


Figure C.2: Utilization-adjusted TFP vs. unadjusted Solow residual



Notes: Impulse responses of the baseline utilization-adjusted TFP measure (solid line, 90% dark and 95% light bands) and of the unadjusted private-sector Solow residual (dashed line, 90% band) to a one percent austerity-driven decline in government spending, estimated as in equation 7.


C.2 Additional regional statistics

Figure C.3: Sample regions and the share s_i



Notes: The figure depicts the map of European NUTS 2 regions with the share s_i used in Bartik instrument construction.

Figure C.4: Average Annual Labor Productivity Growth




Notes: Average annual growth of private-sector labor productivity (private-sector GVA per hour worked) between 1999-2019.

C.3 Country-level analysis

We repeat the exercise at the country level with an entirely different dataset. Government spending is total government expenditure from the 6th release of the Macrohistory database (Chen et al. 2022), deflated with the consumer price index and divided by population, and its change is instrumented with the national narrative consolidation shock of Adler et al. (2024), i.e. the shift $\tilde{g}_{I,t}$ of our regional instrument without the exposure share. The sample covers the nine countries of our regional sample included in the Macrohistory database (all except Austria) over 1999–2019. We consider three measures of national TFP: the Penn World Tables index (*rtfpna*, available through 2017), the utilization-adjusted TFP growth series of Huo et al. (2023), which ends in 2007 and therefore leaves few observations, and our regional utilization-adjusted measure aggregated to the national level with Törnqvist weights. With nine countries, the responses are noisier, but they are positive throughout, and our aggregated measure responds significantly.

Figure C.5: National-level TFP responses



Notes: Country-level local projections for nine countries (all in our sample except Austria), 1999–2019: cumulative TFP growth on the change in real government expenditure per capita from the Macrohistory database, instrumented with the national narrative consolidation shock, with one lag of the controls, country and year fixed effects, standard errors clustered by country, and horizons up to three years. Given the small country panel, bands are 68% (dark) and 90% (light) confidence intervals, wider than in the regional figures.

C.4 FEVD exercise

Whereas our main analysis focuses on characterizing whether austerity shocks affect productivity, below we also assess the quantitative importance of this relationship. In doing so, we conduct a forecast error variance decomposition (FEVD) exercise. The local projection framework allows computing the contribution of the austerity shocks to the forecast error variance of our variables of interest. First, we consider the share of the variance in TFP that can be accounted for by austerity shocks from 1999 until 2019. The fraction of the variance in productivity at different horizons accounted for by austerity shocks can be recovered directly from the estimates

of Equation (7). This measure therefore provides a metric of the extent to which austerity shocks are quantitatively important in driving productivity dynamics.

We closely follow Born et al. (2020), who extend the approach by Coibion et al. (2017); Gorodnichenko and Lee (2020) to a panel setting. In particular, we compute the variance share of the regional consolidation shock at horizon h as the R^2 of the following regression:

$$\hat{u}_{i,t+h} = \lambda_0 \hat{\epsilon}_{i,t+h} + \dots + \lambda_h \hat{\epsilon}_{i,t} + v_{i,t+h}. \quad (\text{C.1})$$

where $\hat{u}_{i,t+h}$ is the h -step-ahead forecast error of the outcome in (7), i.e. the residual from projecting cumulative growth between $t - 1$ and $t + h$ on the information available at $t - 1$ (the lagged controls and the region and time fixed effects), and $\hat{\epsilon}_{i,t}$ is the regional austerity shock, $s_i \times \tilde{g}_{I,t}$, net of region and time fixed effects.

Table C.1: Forecast error variance decomposition

Horizon	TFP	GDP
1	1.8 %	3.3 %
2	3.4 %	6.5 %
3	4.3 %	9.6 %
4	4.0 %	12.6 %

Notes: Forecast error variance decomposition of productivity based on local projections (C.1).

Appendix D Alternative Productivity Measures

D.1 BFK utilization adjustment

As an alternative to our baseline TFP measure, which is adjusted for utilization following Imbs (1999), we implement the approach of Basu et al. (2006) (BFK). This method corrects the Solow residual for variations in capacity utilization and effort by estimating an instrumental-variable regression of the Solow residual on changes in a utilization proxy, typically hours worked or a survey-based measure.² Common instruments in this framework include monetary policy and oil price shocks.

We apply the BFK adjustment to our regional Solow residual by estimating an IV regression with regional fixed effects, where changes in band-pass-filtered hours per worker (following Basu et al. (2006)) serve as the proxy for utilization. As instruments, we use the lagged change in the excess bond premium, a measure of financial shocks, and lagged high-frequency monetary policy surprises (Jarociński and Karadi 2020). The resulting BFK-adjusted TFP growth rate is highly correlated with our baseline Imbs-adjusted TFP (correlation = 0.72, significant at the 1% level). When using the BFK-adjusted regional TFP, we continue to find that productivity increases following an austerity shock (Table 2), though the magnitude of the response is somewhat muted relative to the baseline.

The BFK approach relies on the exclusion restriction that TFP is exogenous to the chosen instruments. This assumption may be difficult to sustain in our setting, as TFP itself responds endogenously to demand shocks.³ By contrast, the utilization adjustment proposed by Imbs (1999) does not rely on a regression framework and therefore avoids such identification assumptions, which is why we maintain the Imbs-adjusted TFP as our baseline measure.

D.2 Cost-share-based TFP

As an alternative to our baseline measure, we compute TFP using *cost shares* instead of revenue shares to estimate input elasticities, thereby relaxing the assumption of perfect competition and allowing for firm markups. In particular, the output elasticity of labor is given by

$$\frac{\text{Total private-sector labor compensation}_t}{\text{Total private-sector labor compensation}_t + \text{Private-sector capital costs}_t}. \quad (\text{D.2})$$

While data on the labor compensation is available at ARDECO, capital costs have to be estimated.

²See, among others, Comin et al. (2025) and Huo et al. (2023).

³Moreover, recent evidence suggests that monetary policy shocks can affect both productivity (Jordà et al. 2024) and innovation activity (Ma and Zimmermann 2023).

The private-sector capital costs equal the product of the rental rate of capital, R_t , and the nominal capital stock, $K_t Q_t$. The rental rate of capital is computed following Hall and Jorgenson (1967):

$$R_t = 1 + r_t - (1 - \delta^k) \frac{Q_t}{Q_{t-1}}, \quad (\text{D.3})$$

where r_t is the interest rate and δ^k the depreciation rate. For each country, we estimate

$$r_t = \text{GovBondYield}_t + \text{BaaSpread}_t. \quad (\text{D.4})$$

This corresponds to a simplified version of the equation used in Comin et al. (2025). In their baseline, Comin et al. (2025) estimate r_t as the sum of a risk-free rate and a weighted average of the risk premia on debt and equity. Since equity risk premia are available for only five countries in our sample, we adopt this simplified specification that omits the equity premium. For those five countries, the correlation between our r_t and the estimate using equity premium is high (0.87).

For the government bond yield, we use the 10-year benchmark rate reported by the OECD. The Baa spread is defined as the yield differential between euro-area BBB-rated and AAA-rated corporate bonds, as obtained from Datastream.⁴ The price of capital goods is measured using national-level investment goods price indices from the ARDECO database. The depreciation rate is constructed at the NUTS2 level and corresponds to the implicit rate inferred from the evolution of the regional capital stock. Except for the estimation of the output elasticity, all remaining steps in the computation of TFP follow the baseline specification.


Table 2 demonstrates that our baseline findings are robust to computing TFP using cost shares rather than revenue shares.

D.3 Sensitivity to utilization adjustment parameters

Our utilization-adjusted TFP measure depends on the calibration of the effort elasticities (ν , ν_e) and the capital-utilization curvature (β_k). Figure D.1 shows that the baseline TFP response is robust to plausible alternative values of these parameters. The point estimates under all calibrations lie within or very close to the baseline confidence bands.

⁴The exact series from Datastream we use are (1) SPEIA3E for the euro area AAA corporate bond yield, and (2) SPEIB3E for the euro area BBB corporate bond yield.

Figure D.1: Sensitivity of TFP Response to Utilization-Adjustment Parameters



Notes: The figure reports the TFP impulse response under alternative calibrations of the utilization-adjustment parameters. Panel (a) varies the effort elasticity v while keeping $v_e = 1$: the baseline sets $v = v_e = 1$ (solid green line), with alternatives $v = 1/2$ (navy dashed) and $v = 2$ (red dashed). Panel (b) varies the capital-utilization curvature β_k : the baseline sets $\beta_k = 4$ following Anzoategui et al. (2019) (solid green), with alternatives $\beta_k = 2$ (navy dashed) and $\beta_k = 6$ (red dashed). Shaded areas denote 90% and 95% confidence bands for the baseline specification.

Appendix E Robustness and Identification

E.1 Sample robustness

Table E.1 re-estimates the baseline specification excluding one country at a time. The TFP response remains positive and significant at every horizon in all ten subsamples; the peak response at $h = 2$ ranges from 1.8 percent without the Netherlands to 3.2 percent without Germany, whose 38 regions carry the largest weight in the sample.

Table E.1: Robustness Check: Excluding Countries

	$h = 0$	$h = 1$	$h = 2$	$h = 3$	$h = 4$
Baseline	1.03*** (0.24)	1.93*** (0.45)	2.10*** (0.50)	1.63*** (0.42)	1.61*** (0.51)
# Obs	1,941	1,941	1,941	1,941	1,941
Austria	1.07*** (0.24)	2.02*** (0.46)	2.16*** (0.52)	1.66*** (0.43)	1.60*** (0.52)
# Obs	1,815	1,815	1,815	1,815	1,815
Belgium	0.93*** (0.19)	1.82*** (0.39)	1.92*** (0.44)	1.56*** (0.37)	1.65*** (0.46)
# Obs	1,787	1,787	1,787	1,787	1,787
Finland	1.05*** (0.24)	1.93*** (0.45)	2.06*** (0.51)	1.63*** (0.43)	1.63*** (0.53)
# Obs	1,885	1,885	1,885	1,885	1,885
France	0.86*** (0.21)	1.71*** (0.41)	2.01*** (0.46)	1.49*** (0.36)	1.38*** (0.42)
# Obs	1,647	1,647	1,647	1,647	1,647
Germany	1.74*** (0.51)	2.92*** (0.87)	3.21*** (1.01)	2.56*** (0.88)	2.35*** (0.99)
# Obs	1,409	1,409	1,409	1,409	1,409
Ireland	0.94*** (0.29)	1.69*** (0.44)	1.93*** (0.57)	1.62*** (0.61)	1.66*** (0.76)
# Obs	1,904	1,904	1,904	1,904	1,904
Italy	0.96*** (0.21)	1.87*** (0.40)	1.99*** (0.45)	1.64*** (0.38)	1.68*** (0.47)
# Obs	1,647	1,647	1,647	1,647	1,647
Netherlands	0.92*** (0.22)	1.70*** (0.42)	1.83*** (0.47)	1.29*** (0.35)	1.17*** (0.38)
# Obs	1,773	1,773	1,773	1,773	1,773
Portugal	1.07*** (0.28)	2.14*** (0.52)	2.44*** (0.56)	1.83*** (0.42)	1.90*** (0.59)
# Obs	1,871	1,871	1,871	1,871	1,871
Spain	1.16*** (0.34)	1.94*** (0.64)	1.91*** (0.69)	1.35*** (0.54)	1.41*** (0.63)
# Obs	1,731	1,731	1,731	1,731	1,731

Notes: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

E.2 Additional robustness check

Table E.2 collects a broader set of robustness checks. The first two rows repeat the baseline with region-clustered and with Driscoll and Kraay (1998) standard errors. The baseline results are roughly insensitive to sample composition (extending coverage to the island regions omitted from the baseline); to using alternative capital data to construct the TFP series; to using a time varying exposure share s_i to construct the instrument; and dropping the regions with the highest exposure shares. A further set of checks addresses the concern that consolidations and productivity may both respond to the GFC, and finally, we add some control for fiscal space. The main result also survives when dropping Ireland and Portugal from the sample—the two program countries with the most severe adjustments—and controlling for the presence of an IMF program. The last row of the table broadens the spending measure to total government spending, adding public investment (proxied by non-market-sector gross fixed capital formation) to government consumption, with the exposure share in the instrument recomputed accordingly. The TFP response remains positive and significant at all horizons, at about half the baseline magnitude.

Table E.2: Additional Robustness Checks

	$h = 0$	$h = 1$	$h = 2$	$h = 3$	$h = 4$
Baseline	1.03*** (0.24)	1.93*** (0.45)	2.10*** (0.50)	1.63*** (0.42)	1.61*** (0.51)
# Obs	1,941	1,941	1,941	1,941	1,941
Baseline, Driscoll-Kraay SE	1.03** (0.40)	1.93*** (0.66)	2.10*** (0.52)	1.63*** (0.28)	1.61*** (0.46)
# Obs	1,941	1,941	1,941	1,941	1,941
W/ islands	1.04*** (0.21)	1.89*** (0.39)	2.01*** (0.44)	1.57*** (0.37)	1.52*** (0.42)
# Obs	2,025	2,025	2,025	2,025	2,025
Gardiner's capital stock	0.92*** (0.29)	1.61*** (0.44)	1.76*** (0.56)	1.70*** (0.59)	1.75** (0.70)
# Obs	1,496	1,496	1,496	1,496	1,496
Time varying exposure share	0.98*** (0.22)	1.84*** (0.42)	1.99*** (0.45)	1.52*** (0.37)	1.47*** (0.46)
# Obs	1,941	1,941	1,941	1,941	1,941
Drop top 10% s_i	1.20*** (0.30)	2.19*** (0.54)	2.41*** (0.58)	1.95*** (0.45)	2.07*** (0.57)
# Obs	1,745	1,745	1,745	1,745	1,745

Continued on next page

Table E.2 – continued from previous page

	$h = 0$	$h = 1$	$h = 2$	$h = 3$	$h = 4$
Drop top 5% s_i	1.16*** (0.28)	2.20*** (0.50)	2.45*** (0.54)	1.97*** (0.43)	2.11*** (0.54)
# Obs	1,843	1,843	1,843	1,843	1,843
Control GFC dummy * Shock	0.94*** (0.22)	1.61*** (0.36)	1.69*** (0.37)	1.31*** (0.40)	1.49*** (0.52)
# Obs	1,941	1,941	1,941	1,941	1,941
Control national debt-to-GDP	1.01*** (0.24)	1.90*** (0.46)	2.08*** (0.52)	1.63*** (0.42)	1.62*** (0.48)
# Obs	1,941	1,941	1,941	1,941	1,941
Control sovereign spread	1.16*** (0.39)	2.60*** (0.81)	2.76*** (0.88)	1.89*** (0.64)	2.01*** (0.71)
# Obs	1,941	1,941	1,941	1,941	1,941
Drop Ireland & Portugal	0.96** (0.41)	2.02*** (0.65)	2.59*** (0.86)	2.06** (0.84)	2.34** (1.06)
# Obs	1,834	1,834	1,834	1,834	1,834
Control for IMF program	0.99*** (0.25)	1.95*** (0.53)	2.19*** (0.65)	1.69*** (0.51)	1.67*** (0.47)
# Obs	1,941	1,941	1,941	1,941	1,941
Total gov. spending (consumption + investment)	0.52*** (0.12)	0.98*** (0.23)	1.07*** (0.25)	0.82*** (0.21)	0.81*** (0.26)
# Obs	1,941	1,941	1,941	1,941	1,941

Notes: Dependent variable is utilization-adjusted TFP growth throughout. The first row repeats the baseline; the second reports the same estimates with Driscoll and Kraay (1998) standard errors (bandwidth $h + 1$), robust to cross-sectional dependence. Extended sample adds island regions excluded from the baseline (Åland, Corsica, Madeira, Azores, Balearic Islands, Canary Islands, and French overseas territories). Gardiner et al. (2020) capital stock replaces the perpetual-inventory series, available through 2016. Time-varying share replaces the pre-euro-accession average s_i with the lagged share $s_{i,t-1}$. The following rows drop observations with large exposure share s_i , as indicated. Then, we control for the GFC by adding an interaction between a GFC dummy and the instrument. We also control for the GFC intensity using changes in the unemployment rate during this period, normalized by the country's mean and standard deviation. We control for the fiscal space by adding lagged national debt-to-GDP ratios and sovereign spreads to the baseline regression. The next two rows exclude Ireland and Portugal, the two program countries with the most severe fiscal adjustments, and control for a lagged indicator of an active IMF arrangement (MONA database). The last row replaces the spending measure and the instrument with total government spending, i.e. government consumption plus public investment, the latter proxied by non-market-sector gross fixed capital formation; the exposure share s_i in the instrument is recomputed for total spending. Standard errors clustered at the region level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

E.3 Shock validity

Jordà and Taylor (2016) argue that fiscal consolidation plans are predictable based on observable macroeconomic

conditions, implying that even narrative-identified fiscal shocks may not be fully exogenous. To address this concern, we residualize the national consolidation shock. Specifically, at the country-year level, we regress austerity shocks on two lags of GDP growth, government consumption (as a share of GDP), and the debt-to-GDP ratio. We then construct an alternative instrument by interacting the residual from this regression with the predetermined exposure share, s_i . The results are reported in Panel A of Table E.3. Using the residualized shock leaves the structural estimates essentially unchanged, although the first-stage F-statistics become less strong and fall below the usual rule-of-thumb values.

While Panel A addresses concerns that fiscal consolidations may be systematically related to macroeconomic conditions, a separate concern is that the shift-share instrument captures regions' exposure to broader aggregate shocks rather than to fiscal consolidations per se. This would violate the condition on the shares stated in Section 3: that s_i is uncorrelated with regions' differential exposure to other shocks. To assess this, Panel B of Table E.3 replaces the fiscal shock with two aggregate shocks that are unlikely to operate through the fiscal channel captured by s_i : (i) oil supply news shocks and (ii) monetary policy surprises. If s_i captures exposure to fiscal consolidations specifically, these shocks should not predict regional government spending when scaled by s_i . They do not: the first-stage F-statistics are low, and the associated TFP estimates are imprecise and statistically insignificant. These results are consistent with the baseline estimates reflecting a fiscal channel rather than exposure to generic aggregate shocks.

As a further check on the validity of our instrument, we conduct a placebo exercise. We randomly permute the IMF consolidation episodes across country-year observations, preserving the total number and distribution of shocks. If our results were driven by spurious correlations rather than genuine fiscal consolidation, we would expect the permuted instrument to generate similar responses. Panel C of Table E.3 confirms that the placebo instrument has no predictive power for government spending (the first-stage F-statistic falls from 26 to about 0.1) and produces no significant TFP response.


Table E.3: Shock Validity

	$h = 0$	$h = 1$	$h = 2$	$h = 3$	$h = 4$
Panel A: Predictability check (residualized instrument)					
Residual	1.44*** (0.53)	2.67*** (1.00)	2.91** (1.18)	1.88** (0.87)	1.65* (0.89)
Panel B: Falsification tests					
Oil shock $\times$ exposure share	3.68 (17.33)	3.31 (16.75)	-1.52 (8.47)	10.60 (51.91)	2.97 (17.41)
MP shock $\times$ exposure share	0.32 (0.76)	1.03 (1.58)	0.57 (1.77)	1.20 (2.15)	0.86 (1.90)
Panel C: Placebo test (randomly permuted shift)					
Permuted shift $\times$ exposure share	-1.14 (4.58)	-0.49 (4.37)	-1.81 (7.55)	-3.01 (12.78)	-0.88 (7.70)
# Obs	1,944	1,943	1,942	1,941	1,941


Notes: The dependent variable is utilization-adjusted TFP growth (Imbs-PWT measure) as in the baseline specification. Panel A reports the results obtained when replacing the national consolidation shock in the instrument with its residual component, purged of predictable variation related to economic conditions. Specifically, in a first step, the country-year austerity shock is regressed on lagged GDP growth, lagged government consumption (as a share of GDP), and the lagged debt-to-GDP ratio. The alternative instrument is then constructed by interacting the residual from this regression with the predetermined exposure share, s_i . Panel B presents two falsification exercises in which the baseline consolidation shock is replaced by oil supply news shocks (Baumeister-Hamilton principal-component measure) and monetary policy shocks (Jarociński and Karadi (2020)), respectively. The annual oil shock is constructed as the sum of monthly principal-component shocks to oil futures prices around OPEC announcement dates. Panel C reports the TFP response when the instrument is built from a placebo shift: the narratively identified consolidation episodes are randomly permuted across country-year observations (fixed seed), preserving their number and distribution, and interacted with the exposure share s_i ; the Kleibergen-Paap first-stage F statistic of the placebo instrument is about 0.1 (26.3 in the baseline). Standard errors clustered at the region level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix F Sector-Level Additional Results

Figure F.1: Exposure of the Tradables and Non-tradables to the Non-market Sector and to Household Consumption



(a) Non-market sector



(b) Household consumption

Notes: The figure shows the distribution across regions of the exposure of tradables and non-tradables over time. Exposure is the share of a sector's gross output sold, within the same country, to the non-market services sector (panel a) or to household final consumption (panel b), computed from the EUREGIO regional input-output tables. Boxes show the interquartile range across the sample regions, whiskers the adjacent values; outside values are omitted.

Table F.1: Summary Statistics of the Exposure of Tradables and Non-tradables to the Non-market Sector by Country

Country	Variable	Obs	Mean	Std. Dev.	Min	Max
All	Tradables	1,617	2.18	1.13	0.00	7.49
All	Nontradables	1,617	4.19	1.41	0.00	11.04
AT	Tradables	99	2.08	0.27	1.55	2.62
AT	Nontradables	99	4.37	0.49	3.33	5.30
BE	Tradables	121	1.02	0.35	0.43	1.95
BE	Nontradables	121	3.51	0.60	2.54	5.33
DE	Tradables	451	1.79	0.53	0.12	3.34
DE	Nontradables	451	4.62	0.79	0.21	6.77
ES	Tradables	209	3.56	1.38	0.00	7.49
ES	Nontradables	209	3.56	1.39	0.04	6.38
FI	Tradables	55	2.53	0.99	0.37	4.57
FI	Nontradables	55	5.07	2.70	2.39	11.04
FR	Tradables	242	2.75	0.72	0.63	4.05
FR	Nontradables	242	4.59	0.70	2.73	6.77
IE	Tradables	22	1.52	0.46	0.84	2.27
IE	Nontradables	22	2.59	0.69	1.52	3.85
IT	Tradables	231	1.40	1.25	0.00	4.58
IT	Nontradables	231	2.99	2.07	0.00	6.87
NL	Tradables	132	2.32	0.55	1.14	3.72
NL	Nontradables	132	5.29	0.80	3.16	7.56
PT	Tradables	55	3.21	0.61	2.04	4.36
PT	Nontradables	55	4.81	0.51	3.99	6.14


Notes: The table describes the exposure to the non-market sector by country. Exposure is the share of a sector's gross output sold to the non-market services sector of the same country, for tradables and non-tradables, computed from the EUREGIO regional input-output tables, 2000–2010. Observations are region-years.

Table F.2: Summary Statistics of the Exposure of Tradables and Non-tradables to Household Consumption by Country

Country	Variable	Obs	Mean	Std. Dev.	Min	Max
All	Tradables	1,617	16.49	6.65	0.23	40.61
All	Nontradables	1,617	29.99	6.95	1.66	57.54
AT	Tradables	99	11.72	2.56	6.19	17.65
AT	Nontradables	99	30.57	1.71	25.49	33.17
BE	Tradables	121	7.28	1.26	4.86	10.32
BE	Nontradables	121	22.04	1.24	19.24	25.14
DE	Tradables	451	17.21	3.22	1.46	23.86
DE	Nontradables	451	30.78	3.70	1.66	36.50
ES	Tradables	209	15.22	7.00	3.40	34.37
ES	Nontradables	209	28.64	8.49	14.01	52.38
FI	Tradables	55	7.71	4.12	0.23	15.58
FI	Nontradables	55	29.04	4.66	18.96	38.29
FR	Tradables	242	21.96	2.37	13.83	27.50
FR	Nontradables	242	30.05	1.19	27.34	33.24
IE	Tradables	22	7.54	1.50	5.43	10.89
IE	Nontradables	22	19.42	1.80	16.16	23.07
IT	Tradables	231	22.88	6.74	8.53	40.61
IT	Nontradables	231	37.58	10.65	22.33	57.54
NL	Tradables	132	10.15	1.98	4.93	13.66
NL	Nontradables	132	24.72	1.79	18.05	27.34
PT	Tradables	55	20.87	2.33	16.83	24.72
PT	Nontradables	55	30.68	3.02	25.89	36.60


Notes: The table describes the exposure to household consumption by country. Exposure is the share of a sector's gross output sold to household final consumption in the same country, for tradables and non-tradables, computed from the EUREGIO regional input-output tables, 2000–2010. Observations are region-years.

Figure F.2: TFP Response by Sector




Notes: Impulse responses to a one percent austerity-driven decline in government spending, estimated as in equation 7. Bands are 90% (dark) and 95% (light) confidence intervals.

Figure F.3: Labor Productivity by Hour for Tradables and Non-tradables Sectors




Notes: Impulse responses to a one percent austerity-driven decline in government spending, estimated as in equation 7. Bands are 90% (dark) and 95% (light) confidence intervals.

Figure F.4: Labor Productivity by Hour




Notes: Impulse responses to a one percent austerity-driven decline in government spending, estimated as in equation 7. Bands are 90% (dark) and 95% (light) confidence intervals.

Figure F.5: GVA by Tradables and Non-tradables Sector




Notes: Impulse responses to a one percent austerity-driven decline in government spending, estimated as in equation 7. Bands are 90% (dark) and 95% (light) confidence intervals.

Figure F.6: GVA by Sector



Notes: Impulse responses to a one percent austerity-driven decline in government spending, estimated as in equation 7. Bands are 90% (dark) and 95% (light) confidence intervals.

Figure F.7: Number of Local Units by Sector: Each Panel on Its Own Sample




Notes: Same series and specification as Figure 11 (Eurostat SBS local units, 2008–2019; tradables = NACE B–E, non-tradables = F–J and L–N, sums defined only when all sections are reported), but each panel uses all region-years available for that sector (330 for tradables, 313 for non-tradables) instead of the common sample of 274 used in the main text. Impulse responses to a one percent austerity-driven decline in government spending, estimated as in equation 7. Bands are 90% (dark) and 95% (light) confidence intervals.

Appendix G Innovation Responses: R&D Spending and Patent Applications


This appendix reports the responses of regional innovation inputs and outputs to the austerity shock discussed in Section 4. R&D expenditure is intramural R&D expenditure by NUTS 2 region from Eurostat (table `rd_e_gerdreg`), for all sectors of performance (overall), the business enterprise sector (private), and the government sector (public), deflated with the national investment deflator from ARDECO (2015 euros). Patent applications are PCT applications per million inhabitants from the OECD Regional Patent Statistics database (fractional count, by inventor region and priority year, 1995–2022); they are available at the NUTS 2 level for seven of our countries (AT, ES, FI, IE, IT, NL, PT) and only at the NUTS 1 level for Belgium, France, and Germany. For each series, the dependent variable is the cumulative percent change between $t - 1$ and $t + h$, and the specification is otherwise identical to equation 7.

Figure G.1: Response of R&D Expenditure to Austerity



Notes: Impulse responses of R&D expenditure (cumulative percent change) to a one percent austerity-driven decline in government spending, estimated as in equation 7. Source: Eurostat regional R&D statistics, deflated to 2015 euros. Bands are 90% (dark) and 95% (light) confidence intervals.

Figure G.2: Response of Patent Applications to Austerity



Notes: Impulse responses of PCT patent applications per million inhabitants (cumulative percent change) to a one percent austerity-driven decline in government spending, estimated as in equation 7. Source: OECD Regional Patent Statistics. Panel (a) restricts the sample to the seven countries with NUTS 2 data. Panel (b) extends the sample to all ten countries by assigning the NUTS 1 patent value to each NUTS 2 region within. Bands are 90% (dark) and 95% (light) confidence intervals.