Consumption Responses to a Major Minimum Wage Increase: Evidence from Spain*

Ignacio González[†] Hector Sala[‡] Pedro Trivín[§]

June 29, 2026

Abstract

This paper examines the effects of minimum wage increases on household consumption, focusing on Spain's 2019 reform, which raised the wage floor by an unprecedented 22.3 percent in a low-inflation environment. Using high-frequency confidential card-based transaction data, we exploit geographic variation in exposure to the reform. We find that municipalities more exposed to the reform increased consumption by 4.5% relative to less-exposed municipalities, with the largest gains concentrated in categories such as electronics, leisure, and spending at restaurants and hotels. We complement this analysis with household-level evidence from the Spanish Household Budget Survey. Consistent with the transaction data, affected households increased their consumption substantially—by 10.2 percent—and exhibited a marginal propensity to consume of 0.8, with spending gains also concentrated in discretionary categories. Combining evidence from both sources, our findings are suggestive of sizeable local demand spillovers. Firm-level sectoral evidence further suggests that the costs of the reform were primarily absorbed through the exit of less productive firms rather than through broad employment or price adjustments.

JEL Codes: D12, J31, R28.

Keywords: Minimum Wage, Consumption, Transaction Data, Discretionary Spending, Marginal Propensity to Consume.

*We sincerely appreciate feedback from Arnab Bhattacharjee, Manudeep Bhuller, David Card, Antoni-Italo de Moragas, Edoardo di Porto, Ferrán Elias, Jordi Galí, Amos Golan, Marcel Jansen, Francis Kramarz, Attila Lindner, Núria Rodriguez-Planas, Adam Scavette, and Javier Vazquez-Grenno, as well as seminar participants at the University of Oviedo, the 16th Labor Economics Conference of the Spanish Labor Economics Association, the Society of Government Economists 2026 Annual Conference, the Workshop on Frontiers of Macro-Labour Economics in Edinburgh, and the XIV Workshop in Public Policy Design in Girona. We thank Banc Sabadell for providing data and technical support, and the Data Lab at Bank of Spain, the Spanish National Statistics Institute, the Social Security Administration, and the Spanish Institute for Fiscal Studies for assistance in accessing data. We are grateful to the Spanish Institute for Fiscal Studies for its financial support. Héctor Sala gratefully acknowledges financial support from the Spanish Ministry of Science and Innovation through grant PID2022-136482OB-I00. Pedro Trivin received funding from the University of Eastern Piedmont through the Bando SSHINE 2025 grant.

[†]Department of Economics, American University. ignaciog@american.edu

[‡]Department of Applied Economics, Universitat Autònoma de Barcelona and IZA Network at LISER. hector.sala@uab.cat

[§]DISEI, University of Eastern Piedmont. pedro.trivin@uniupo.it

1 Introduction

The minimum wage has long been central to economic policy debates, with much of the discussion focused on its effects on employment, labor costs, and income distribution. While these labor market impacts have been extensively studied, especially since the seminal work of Card and Krueger (1994), other economic implications of minimum wage policy remain less understood. In particular, relatively little emphasis has been placed on household consumption, despite its fundamental role in shaping both individual welfare and aggregate economic activity.

The relationship between minimum wages and household consumption is relevant for several reasons. If higher minimum wages raise the earnings of low-wage workers with high marginal propensities to consume, they can stimulate aggregate demand and thereby affect economic activity, particularly in areas with a high concentration of these workers. Beyond aggregate effects, responses at the household level may themselves be informative. The extent to which a minimum wage increase provides financial relief is reflected both in how much of the additional income is spent on current consumption and in how it is allocated across different goods. Differences in the level and composition of spending can therefore shed light on the welfare gains generated by the policy. Given the substantial body of research showing that consumption provides a more informative measure of household well-being than income (Deaton, 1992; Attanasio and Davis, 1996), evaluating minimum wage policies through these lenses offers a broader perspective on their economic and welfare implications than an approach focused exclusively on labor market outcomes.

This paper studies how minimum wage increases affect household consumption, exploiting Spain's 2019 reform as a natural experiment. The reform raised the statutory minimum wage by 22.3 percent, one of the largest in recent European history (European Central Bank, 2022). Unlike gradual, inflation-indexed adjustments, the 2019 increase was explicitly aimed at strengthening the purchasing power of low-wage workers and was implemented in a low-inflation environment, with the intention that most of the increase would translate into meaningful real wage gains. From an empirical perspective, the increase constitutes a large and abrupt intervention, allowing us to explore potential nonlinearities in the effects of the wage change on consumption behavior that would be difficult to detect in the case of more modest adjustments. Spain also offers a relevant empirical setting given the characteristics of its labor market. The country has long stood out among advanced economies for its persistently high unemployment rate—of 15.26 percent in 2018—and pronounced regional and demographic disparities in employment outcomes. These conditions heighten uncertainty about the net consumption effect, since potential employment losses could partially offset income gains for affected households.

Our empirical strategy follows a two-pronged approach. First, we use high-frequency, confidential transaction data from point-of-sale (POS) terminals and credit and debit card payments, provided by one of the largest commercial banks in Spain and a market leader in POS coverage.[1] This dataset, available at the municipal level, allows us to exploit geographic variation in exposure to the minimum wage reform to identify its impact on local consumption. We construct measures of local exposure to the reform using administrative income tax records, which capture the pre-reform distribution of wages and enable us to identify municipalities more intensely affected by the minimum wage increase. The frequency and broad coverage of spending categories in our transaction data allow us to precisely measure short-run consumption responses while mitigating concerns that commonly affect survey-based measures of household spending, such as recall error and non-response (Meyer et al., 2015), and selective attrition, which is concentrated among lower socioeconomic status individuals (Fitzgerald et al., 1998), a group that overlaps substantially with minimum wage workers.[2] Because we measure transactions, our data capture local general equilibrium effects—such as demand spillovers across households and firms—that are absent from more traditional analyses based on survey data.

Our baseline specification indicates that the minimum wage increase raised local consumption by 4.5 percent. The magnitude of the effect is economically meaningful and remains significant across a wide range of alternative specifications and exposure measures. Furthermore, our results are robust to potential geographic spillovers: they continue to hold when outcomes are aggregated to larger geographic units and when we explicitly test for consumption spillovers using spatial econometric models, with estimated effects ranging from 3.6 to 5.3 percent across specifications.

We further address concerns that tax-based exposure measures may suffer from measurement error by constructing an alternative indicator of exposure using Social Security administrative data. Although this indicator is only available at a higher level of geographic aggregation than our tax-based counterpart, it includes information on wages and hours worked. Using this alternative approach yields similar results. We also use Social Security data to document the first stage and show that wages at the bottom of the distribution rose more in areas more exposed to the minimum wage increase, consistent with income gains from the minimum wage driving the consumption response.[3]

When we turn to the composition of the consumption response, we find substantial

[1] Throughout the paper, we use the terms transaction data, credit card, and card interchangeably.

[2] Note that, while low-income households are frequently presumed to rely more heavily on cash, evidence from the 2019 ECB SPACE survey suggests that card usage remains broadly consistent across the income distribution. Low-income households represent a substantial share of credit card spending and use cards at rates only marginally lower than those of higher-income groups (see Appendix B for details). This alleviates the concern that low-income households' spending might be poorly captured by transaction data.

[3] We rely on Social Security administrative data to document the first stage because wages provide a more direct measure of the policy-induced wage response than annual earnings observed in tax records.

heterogeneity across spending categories. The largest increases are concentrated in categories such as electronics (20.2 percent), leisure activities (11.7 percent), and restaurants and hotels (8.7 percent), suggesting that the adjustment largely reflects a broadening of the consumption basket following a substantial income gain. As we show in Appendix E, this shift in spending patterns is consistent with a standard model of nonhomothetic preferences.

Second, we complement the transaction data analysis with a more traditional approach using household-level survey data from the Spanish Household Budget Survey (EPF), which allows us to examine consumption patterns among households directly affected by the policy. Although subject to the survey-measurement concerns discussed above, this data provides a valuable complement for two reasons. First, it captures total household consumption, including cash expenditures excluded from our baseline analysis. Second, it isolates individual consumption responses at the household level, abstracting from the local general equilibrium effects that are embedded in the municipal-level transaction estimates. This approach therefore captures a different margin of adjustment, and its estimates are not directly comparable in magnitude to those derived from transaction data. However, taken together, the two approaches provide mutually reinforcing evidence of the reform's effects. Our preferred specification using household-level survey data shows that households affected by the reform increased consumption by 10.2 percent relative to unaffected households. Using exposure to the minimum wage reform as an instrument for income changes, we estimate a marginal propensity to consume (MPC) of 0.80, consistent with previous estimates in the literature for low-income households (Jappelli and Pistaferri, 2014). As in the transaction-level results, spending gains are again concentrated in discretionary consumption categories.

By connecting the household-level consumption responses to the aggregate spending effects observed at the municipality level, we obtain a back-of-the-envelope estimate of the size of local spending spillovers associated with the minimum wage increase. Under some assumptions, our estimates imply a multiplier between 2.1 and 2.5, suggesting that local demand spillovers substantially amplified the initial spending response of affected households. Although this calculation should be interpreted as illustrative rather than as a precise estimate, the implied multiplier exceeds the range typically reported for output multipliers following government spending shocks, which often lies between 1.5 and 2 (Chodorow-Reich, 2019).

Finally, we examine firm-side adjustments to better understand how the reform's costs were absorbed. We find no evidence of significant price pass-through, either across municipalities or sectors, indicating that the observed consumption increases reflect real rather than nominal gains, at least over the time window covered by our data. We also find no significant increase in unemployment; if anything, employment rose modestly in more exposed areas, ruling out the possibility that job losses offset income gains and dampened the consumption response. However, sectoral analysis provides some evidence that more exposed sectors experienced contractions in the number of firms and workers, while total sectoral value added and profits remained stable. This pattern is consistent with the exit of lower-productivity firms, a mechanism similar to that documented by Rao and Risch (2025) for the United States.

The literature on minimum wages and employment is extensive, including contributions such as Card (1992); Card and Krueger (1994); Dube et al. (2010); Meer and West (2016); Clemens and Wither (2019); Cengiz et al. (2019); Dustmann et al. (2021); and Azar et al. (2024). There is also a growing body of work on the effects of minimum wages on income and income distribution (e.g., Dickens et al., 1999; Autor et al., 2016; Dube, 2019; Harasztosi and Lindner, 2019; Engbom and Moser, 2022). However, the literature examining the relationship between minimum wages and consumption remains relatively recent and limited.

Studies of the US (Aaronson et al., 2012; Cooper et al., 2020; Alonso, 2022) and China (Dautović et al., 2026) have found that higher minimum wages can boost household spending, though the composition of consumption expenditures varies across countries. Aaronson et al. (2012), using household survey data covering the late 1980s to the mid-2000s, finds that minimum wage increases led to higher consumption of durable goods—particularly vehicles—among a small group of households with access to collateralized credit. Similarly, Cooper et al. (2020), using city-level data from 1999 to 2017 and exploiting variation in local minimum wages, find that minimum wage hikes led to increased spending on food, especially food away from home. Alonso (2022), using county-level grocery retail sales from 2006 to 2014, documents a positive effect of minimum wage increases on consumption of nondurables.

Relative to these papers, our contribution is threefold. First, we exploit sharper policy variation and more granular consumption data than has typically been available in this literature. Second, we study an economy with broad social insurance coverage and a relatively high minimum wage relative to the median, providing evidence from a distinct institutional setting and thereby strengthening the external validity of existing findings. Third, by combining local transaction data with household-level survey evidence, we can approximate the magnitude of the local demand spillovers generated by the reform.

In terms of results, we find substantial increases in spending on leisure and electronics, which were not emphasized in earlier studies. We observe increased spending in restaurants, consistent with Cooper et al.'s (2020) findings on food away from home, though we find no significant effects on food consumed at home, unlike Alonso (2022).

Our results diverge more sharply from those of Dautović et al. (2026), who uses a representative panel of urban Chinese households and finds that the income from minimum wage increases is spent mostly on health care and education, likely reflecting the ubiquity of intrafamily and intergenerational obligations in this context. While such obligations are also prevalent in Spain, such an effect may not be present there because of widespread access to publicly funded health care and education among low-income households.

A related literature has examined the effects of minimum wage hikes on prices. Studies using scanner data from the US, such as Leung (2021) and Renkin et al. (2022), find a positive relationship between minimum wage increases and grocery prices. Ashenfelter and Jurajda (2022) reach a similar conclusion using data from McDonald's restaurants. For Germany, Link (2024) finds that firms responded to minimum wage hikes by increasing prices more often, a response tied to their decision to preserve employment. Although price effects are not the central focus of our analysis, we find little evidence of meaningful price increases in our setting. One possible explanation is that our analysis relies on broader geographic or sectoral variation, which may mask more pronounced price adjustments at the firm level, particularly among smaller firms where the minimum wage hike represents a larger cost shock. Alternatively, Spanish firms may have absorbed higher labor costs through adjustments on other margins (Clemens, 2021), or price responses could materialize with a lag, beyond the short-term horizon of our study.

The remainder of the paper is structured as follows. Section 2 describes the institutional background of the 2019 minimum wage increase in Spain. Section 3 introduces the data sources used throughout the paper and their intended use. Section 4 outlines the empirical strategy based on municipal-level transaction data and presents the main results. Section 5 complements this analysis with household-level evidence from the EPF. In Section 6 we compare both approaches and discuss aggregate local effects from our transaction data. In Section 7, we present evidence on how prices, employment, and various firm-side variables respond to the increase in the minimum wage. Section 8 concludes.

2 Institutional background

Spain's 2019 minimum wage (MW) increase represented a major policy shift, reflecting broader efforts to address real wage stagnation that had persisted since the Great Recession. In the past, MW adjustments in Spain were typically moderate and incremental, aimed primarily at keeping pace with inflation. These changes were usually negotiated through tripartite discussions involving employers' associations, trade unions, and the government. The 2019 hike, however, stemmed from a political agreement rather than a consensus-based negotiation. The proposal emerged in mid-2018, during efforts to secure parliamentary support for the 2019 budget of the newly formed center-left government. In this context, the left-wing coalition Unidos Podemos made its support conditional on several social policy measures, including a substantial MW increase. The agreement was announced on October 11, 2018, approved by the Council of Ministers in December, and took effect on January 1, 2019.

Implemented in a context of relatively low inflation (1.2 percent in 2018 and 0.7 percent in 2019), the reform introduced an unprecedented 22.3 percent increase in the minimum wage—the largest single-year rise in four decades. The annual MW rose from €10,302.60 in 2018 to €12,600.00 in 2019, corresponding to a monthly increase from €858.55 to €1,050.00 (Figure 1). For comparability with similar reforms in other countries, this corresponds to an hourly minimum wage of €5.37 in 2018 and €6.56 in 2019, assuming 1,920 working hours per year (8 hours per day, 4 weeks per month).

Figure 1: Minimum wage in Spain

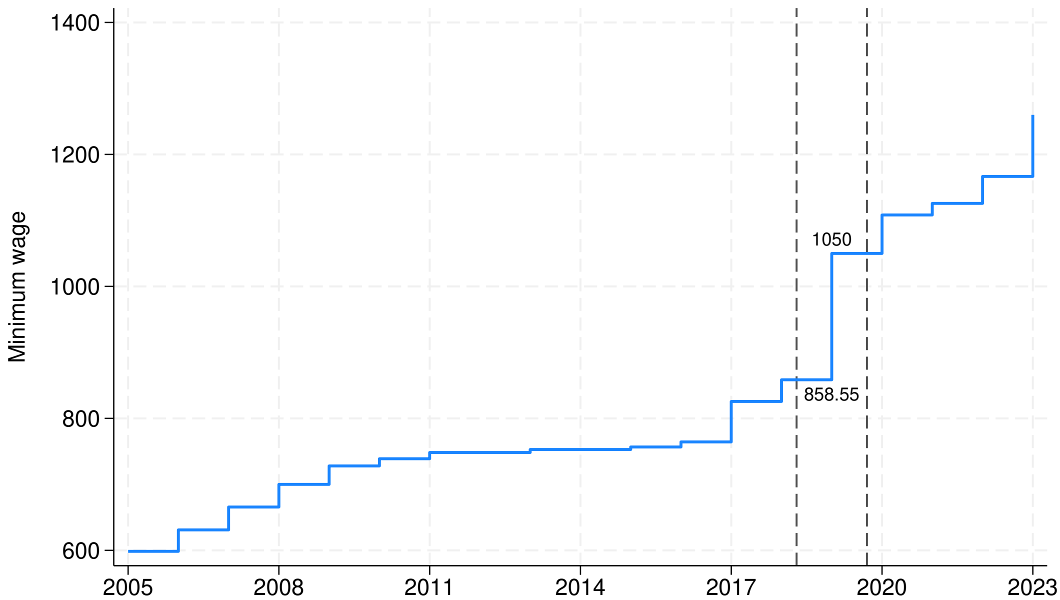

Notes: Information is reported on a 12-month basis. Annual salaries in Spain may also involve 14 payments, in the case of including two extra payments in summer and at Christmas; these may be disbursed at those times or prorated across the year. Accordingly, media outlets sometimes report the monthly minimum wage as rising from €735.90 in 2018 to €900.00 in 2019, based on the 14-payment structure. *Source:* Spanish Ministry of Labor and Social Economy.

The MW is set at the national level and applies uniformly to all workers, regardless of age, sector, or region, with proportional adjustments for part-time and temporary contracts.[4] As a result, the increase affected a broad cross-section of the labor market. According to the Wage Structure Survey from the Spanish National Statistics Institute

[4]Despite Spain's prominent dual labor market, the reform did not significantly alter the aggregate composition of contract types. Labor Force Survey data show the share of permanent contracts among dependent employees was stable at 73.2 percent in 2018 and 73.7 percent in 2019, with fixed-term contracts at 26.8 percent and 26.3 percent, respectively.

(INE), 14.10 percent of workers earned between 0 and 1 times the MW in 2018. By 2019, this share had increased to 18.18 percent, reflecting the extensive reach of the reform and its economy-wide impact (INE, 2024). Gorjon et al. (2024) show that workers affected by the reform were disproportionately women, young workers, and individuals employed in small establishments and in certain low-wage service sectors such as hospitality, commerce, administrative support, and household and recreational services.

3 Data sources

We combine multiple administrative, survey, financial, and transaction datasets to study the effects of the 2019 minimum wage increase from three complementary perspectives: regional exposure, household consumption behavior, and firm-side adjustment. This section describes each data source, its coverage, and its role in the empirical analysis.

3.1 Data for the geographic analysis

The regional analysis combines three datasets that provide information on local exposure to the reform, consumption outcomes, and wage responses. The role of each dataset within the regional analysis is described in Section 4.

Transaction data. We use confidential, anonymized municipality-level transaction records based on credit card and POS operations provided by Banc Sabadell, one of Spain's largest banks and the market leader in POS devices, with an estimated 20 percent share of the national market.[5] The dataset contains weekly transaction data spanning January 2018 to December 2019, disaggregated into 48 consumption categories.

Benchmarking against National Accounts suggests that the dataset covers approximately 3.7 percent of total private consumption in 2018 and 4.0 percent in 2019, with coverage remaining stable across regions and over time. Moreover, the regional distribution of Banc Sabadell POS transactions closely mirrors the geographic distribution of total POS spending in Spain. Additional evidence on the representativeness of the dataset is provided in Appendix B.

Income tax declarations. To measure geographic exposure to the reform, we use the 2018 Personal Income Tax Sample from the Spanish *Instituto de Estudios Fiscales* (IEF), which draws directly from the Spanish Tax Agency (AEAT). The dataset contains detailed information on annual labor income, demographic characteristics, and municipality of residence for 3,011,866 personal income tax declarations, representing approximately 14 percent of all declarations filed nationally. We use these records to characterize the

[5] According to the Asociación Española de Banca, Banc Sabadell is the fourth-largest Spanish bank by assets, after Santander, BBVA, and CaixaBank (AEA, 2023).

distribution of labor income across municipalities and to construct the exposure measures described in Section 4.

Social Security administrative data. We complement the analysis with the *Muestra Continua de Vidas Laborales* (MCVL), an administrative dataset derived from Spanish Social Security records that provides detailed longitudinal information on individuals' employment histories. The MCVL consists of a random sample covering approximately 4 percent of all Social Security-affiliated individuals each year, corresponding to roughly 1.2 million individuals per wave. The dataset includes complete employment and unemployment spell histories, Social Security contributions, part-time coefficients, and geographic area of residence.

The MCVL serves several purposes throughout the paper. Within the regional analysis, it is used to construct an alternative exposure measure and to document the first-stage effects of the reform on the wage distribution (Sections 4.3.2 and 4.5). In the firm-side analysis, it provides sector-level exposure measures, employment outcomes, and control variables for the price and sector-margin analyses (Section 7). Depending on the exercise, the MCVL is used at different levels of aggregation and temporal frequencies, drawing on data from 2018–2019. The specific aggregation and time windows employed in each exercise are described in the corresponding empirical sections.

3.2 Data for the household analysis

For the household-level analysis, we use the EPF, an annual survey conducted by INE that samples approximately 21,000 households per wave. Unlike the transaction data, the EPF captures total household expenditures, including spending conducted outside electronic payment systems. The survey records detailed expenditure information following the ECOICOP classification system and provides information on household income, employment characteristics, demographics, and wealth proxies. Sampling weights allow the construction of nationally representative estimates.

The EPF incorporates a rotating panel design in which a subset of households is reinterviewed in consecutive waves. We exploit this feature to follow households before and after the reform using waves covering 2017–2019. Details of sample construction and estimation procedures are provided in Section 5.

3.3 Data for the firm-side analysis

The firm-side analysis examines how firms adjusted to higher labor costs using data on prices, unemployment, and firm financial accounts. Their use within the firm-side analysis is discussed in Section 7.

Business Margins Observatory. To study price responses, we use quarterly sectoral price indices from the *Observatorio de Márgenes Empresariales* (Business Margins Observatory), covering the period from the first quarter of 2018 to the fourth quarter of 2019. The Observatory is a joint initiative of the Ministry of Economy, Trade and Business, the Bank of Spain, and the AEAT established to monitor corporate profit margins. It provides quarterly price indices for 72 sectors constructed using methodologies consistent with those applied in the AEAT's Sales, Employment, and Wages reports, incorporating industrial price indices for manufacturing, service price indices for service-sector activities, and consumer price indices for retail and transport. Sector classifications follow NACE codes at varying levels of aggregation.

Registered unemployment data. To study local labor market conditions, we use monthly registered unemployment records from the Spanish Public Employment Service (SEPE), spanning January 2018 to December 2019 at the municipality level. We derive municipal unemployment rates by dividing the number of registered unemployed individuals by the working-age population using annual population estimates from INE.

Firm-level financial data. To examine broader firm-side adjustments—including value added, profits, firm counts, employment, productivity, markups, and markdowns—we use the Bank of Spain's *Central de Balances Integrada* (CBI), which contains annual financial information for non-financial companies obtained from accounts filed in the *Registro Mercantil*. For the 2015–2019 period covered in our analysis, the dataset includes on average approximately one million firms per year, which we aggregate at the sector-year level in our analysis.

4 Regional approach

4.1 Variable construction and estimation sample

The regional analysis exploits cross-municipal variation in the share of workers affected by the reform to identify its causal effect on local consumption.[6] We describe the construction of the exposure measure, the local consumption outcome, and the resulting estimation sample.

Exposure measure. Using the IEF income tax sample, we define exposure to the reform in a given municipality as the share of working-age individuals whose annual earnings in 2018 fell between €8,000 and €12,600—the earnings range directly affected

[6] Municipalities are the lowest tier of government in Spain and are responsible for local public services such as urban planning, infrastructure, social services, and basic administration. Spain's subnational government structure consists of 17 Autonomous Communities, 50 provinces, and more than 8,000 municipalities, with municipalities operating at the local level beneath the regional and provincial tiers.

by the 2019 minimum wage increase.[7]

This measure can be computed for 5,520 of Spain's 8,124 municipalities.[8] Exposure ranges from 0 to 60 percent, indicating substantial geographic heterogeneity. Figure 2 presents the geographic distribution of exposure across municipalities. As expected, large municipalities—particularly those in the metropolitan areas of Madrid, Barcelona, and Valencia—tend to exhibit lower exposure, reflecting their higher average wage levels relative to the new minimum.

Figure 2: Exposure to the 2019 MW increase

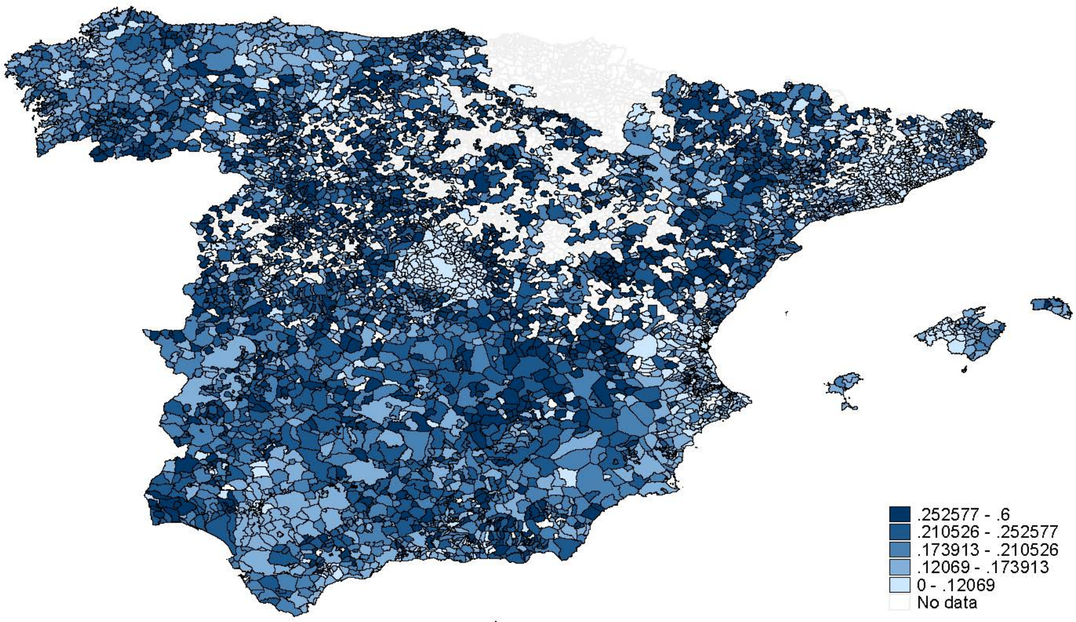

Notes: Share of tax declarations from individuals aged 16–65 with an annual income in the range €8,000–€12,600. Data are available for 5,520 municipalities. *Source:* IEF and AEAT.

A potential limitation of the tax data could be that low-earning workers are generally not required to file income tax returns if annual earnings fall below €22,000 (€12,643 with multiple payers), implying that some affected workers may be missing from the sample. However, many low-income individuals file voluntarily to claim refunds or deductions related to maternity, large families, or housing benefits, increasing coverage among the target population. In addition, this limitation would only affect our exposure measure if under-coverage varies systematically across municipalities. Table A.1 in Appendix A shows that province-level under-coverage rates are uncorrelated with the composition of

[7] The lower bound is set below the 2018 minimum wage to account for workers who may have been affected despite not working full-time or throughout the full year. We later show robustness to alternative exposure definitions.

[8] We exclude municipalities with fewer than 10 tax declarations to reduce measurement noise. The Basque Country and Navarre are not included because personal income taxes in these regions are administered by regional tax agencies rather than the AEAT. Most remaining gaps correspond to sparsely populated municipalities.

the low-wage workforce, suggesting that differential under-coverage is unlikely to systematically bias our exposure measure.

Consumption outcome. We measure local consumption using Banc Sabadell transaction records aggregated from weekly to monthly frequency. To minimize the influence of tourism and better capture domestic spending responses, we exclude transactions made with foreign cards. We combine point-of-sale transactions made with Spanish cards with customer-side transaction data from Banc Sabadell-issued cards, allowing us to capture expenditures processed through Banc Sabadell terminals as well as purchases made by Banc Sabadell customers.[9]

Consumption values are expressed in real 2021 euros using province-, month-, and product-specific CPI adjustments based on the ECOICOP classification.[10] Our baseline outcome is monthly real consumption normalized by municipality's population. To study heterogeneity in spending responses, we aggregate the bank's 48 transaction categories into 11 broader consumption groups.[11]

Econometric sample. The estimation sample is determined by both data availability and measurement considerations. First, we restrict the sample to municipalities with a balanced panel of monthly consumption, thereby avoiding composition effects.[12] Moreover, the data provided by Banc Sabadell is extracted directly by their technology department, without pre-treatment or outlier adjustments. To reduce noise, we exclude municipalities in the bottom 5 percent of the per capita consumption distribution and the top 5 percent with the largest consumption variations.

Second, because our identification strategy relies on geographic variation in exposure, we impose additional restrictions intended to limit contamination from cross-municipality consumption flows. For example, national tourism presents a challenge, as touristic locations—often dominated by low-wage, service-sector jobs—may appear in the upper range of our exposure distribution. If these areas attract national tourists, our estimates could be biased upward by capturing increased consumption from low-income workers temporarily relocating from control to treatment municipalities. To address this, we conservatively exclude potentially touristic municipalities, defined as those where average summer per

[9]To avoid double-counting, transactions made with Banc Sabadell cards are excluded from the POS dataset when combining the two sources.

[10]Although Banc Sabadell categories do not directly follow ECOICOP definitions, we harmonize transaction categories to their closest classification groups.

[11]Appendix A, Table A.2, reports spending composition across groups and shows stable expenditure shares between 2018 and 2019. Monthly consumption per capita is, on average, €38.2 in 2018 and €41 in 2019.

[12]In the Banc Sabadell data, missing/zero values do not display a clear economic pattern and are disproportionately concentrated in small municipalities. In these cases, limited transaction coverage raises concerns about the reliability of measured consumption. Focusing on a balanced panel of municipalities with consistently positive observed spending helps mitigate these concerns.

capita consumption is at least 50 percent higher than during the rest of the year.[13]

After applying these restrictions, the final estimation sample contains 1,999 municipalities, covering approximately 78 percent of the Spanish population (around 36 million individuals). The reduction from the 5,520 municipalities with exposure data reflects exclusions that disproportionately affect small, sparsely populated localities. Figure A.1 in Appendix A shows the exposure distribution in the final estimation sample, and Table A.3 documents that less exposed municipalities tend to be larger, wealthier, and have lower unemployment rates.

4.2 Empirical strategy

The nationwide scope of the policy and its simultaneous implementation across all regions rule out the use of a natural control group. Instead, our identification strategy exploits geographic variation in exposure to the MW hike, measured by the share of individuals affected by the increase within each geographic unit. Following the approach of Card (1992), we compare outcomes across areas with differing levels of exposure to the reform.[14]

Recent econometric studies have shown that standard difference-in-differences (DiD) designs may yield biased estimates in settings with continuous treatment and no stayers (de Chaisemartin et al., 2024; Callaway et al., 2024). To address this concern, we adopt the aggregation approach proposed by Callaway et al. (2024), which demonstrates that partial aggregation across treatment intensities can yield interpretable causal parameters.

Specifically, we classify our geographic units—municipalities—into two groups based on their exposure to the MW increase: those above and those below the median exposure level. This binary classification ensures balanced group sizes and enables a comparison between municipalities with high and low exposure, under two standard assumptions: (i) municipalities with a larger share of individuals affected by the reform experience a stronger treatment effect; and (ii) in the absence of the reform, both groups would have followed parallel trends in the outcome variable.[15]

Using monthly data from January 2018 to December 2019, we estimate a DiD model to assess the causal effect of the MW increase on consumption. The baseline specification

[13]We later show that our results are robust to alternative thresholds for this exclusion criterion. In Section 4.3.1, we explicitly employ spatial econometric models to assess and bound the potential influence of geographic spillovers on our estimates.

[14]Numerous studies exploit geographic variation in exposure to minimum wage changes, including Stewart (2002), Caliendo et al. (2018), Dustmann et al. (2021), and Jiménez (2023). This strategy is also widely used to assess the effects of other policies and economic shocks—for example, Black et al. (2005), Autor et al. (2013), and Yagan (2019), among many others.

[15]For empirical applications of this identification strategy, see, for example, Bartik et al. (2019) and Caliendo and Wittbrodt (2022).

is:

$$\ln(c_{mt}) = \alpha_m + \beta_1(\text{Treat}_m \times \text{Post}_t) + \phi_{rt} + \epsilon_{mt}, \tag{1}$$

where $\ln(\text{c}_{mt})$ is the natural logarithm of real per capita consumption in municipality m at time t; α_m denotes municipality fixed effects, which control for time-invariant heterogeneity across municipalities; and Treat_m is a binary indicator equal to 1 if municipality m has above-median exposure to the MW increase. Post_t equals 1 for months in the post-reform period (2019) and 0 for the pre-reform period (2018). The term ϕ_{rt} represents region-by-time fixed effects, which account for time-varying local economic conditions, and ϵ_{mt} is the error term, clustered at the municipality level. The coefficient β_1 captures the average treatment effect on the treated (ATT) of the MW increase on consumption.

A key identifying assumption of the DiD approach is that treated and untreated municipalities would have followed parallel trends in the absence of the reform. To validate this assumption, we estimate a dynamic DiD model that examines outcome differences quarter by quarter, using the pre-treatment period as the baseline.[16] The dynamic specification is:

$$\ln(c_{mt}) = \alpha_m + \sum_{q=-4}^{4} \beta_q \cdot (\text{Treat}_m \times \text{Quarter}_q) + \phi_{rt} + \epsilon_{mt}, \tag{2}$$

where β_q represents the relative impact of the reform in quarter q, with $q = 1$ corresponding to the reform's implementation quarter. Coefficients for $q < 0$ allow us to test for pre-existing trends, while those for $q > 0$ capture the dynamic effects of the reform over time.[17]

4.3 Effects on aggregate local consumption

4.3.1 Baseline results

We begin by estimating the effect of the reform on aggregate local consumption. Figure 3 presents the quarterly coefficients from the event study based on equation (2).[18] The first key finding is the absence of statistically significant coefficients during the pre-intervention period. This suggests that consumption trends in high- and low-exposure municipalities

[16] In this analysis, we keep using monthly data but present coefficients by quarter for better graphical representation.

[17] A potential identification concern is that changes in transaction coverage could mechanically generate differential consumption growth across municipalities. In particular, our estimates may be biased if Banc Sabadell expanded disproportionately between 2018 and 2019 in municipalities with greater exposure to the minimum wage increase. Appendix B addresses this concern by showing that province-level changes in transaction coverage over the sample period are uncorrelated with our exposure measure. This evidence suggests that differential expansion is unlikely to be a significant source of bias.

[18] Figure A.2 in Appendix A presents the corresponding event study using the continuous treatment variable instead of the binary indicator.

did not diverge prior to the implementation of the minimum wage reform, supporting the parallel trends assumption.

Figure 3: Event study: Consumption

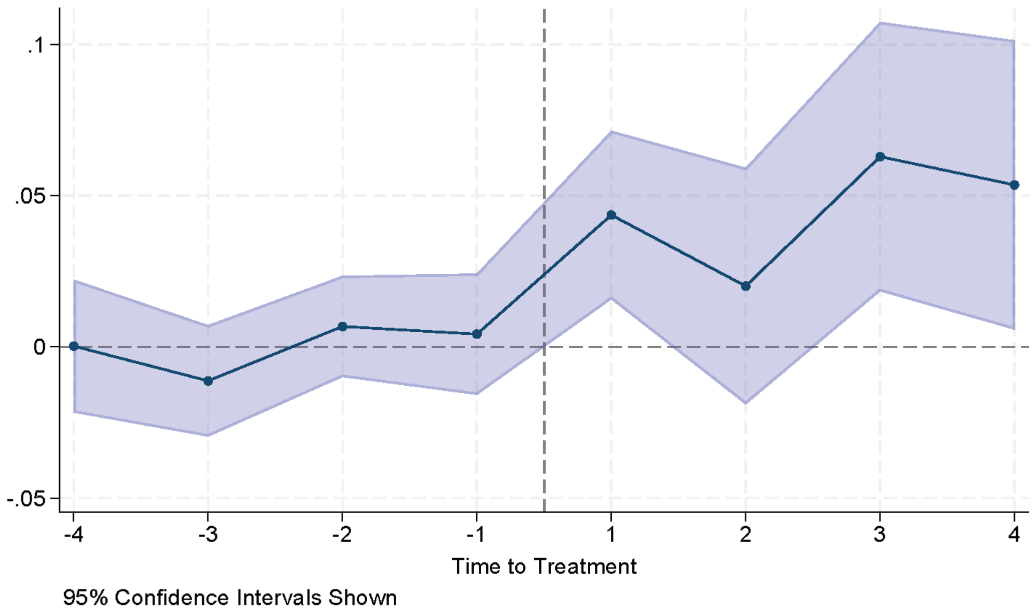

Notes: Dynamic DiD coefficients from equation (2). The horizontal axis shows quarters relative to the reform ($q = 1$ is the first post-reform quarter). Shaded areas are 95% confidence intervals, with standard errors clustered at the municipality level. Regression includes municipality and region-by-time fixed effects. The sample is restricted to municipalities with a balanced panel of 24 monthly observations. 1,999 municipalities; 47,976 observations. *Source:* IEF, AEAT, INE, and Banc Sabadell.

Following the reform, we document a notable increase in consumption in treated municipalities. In the first quarter post-implementation, consumption rises by nearly 5 percent, indicating a significant and immediate response to the minimum wage increase. This effect persists over the subsequent quarters, although the coefficient for the second quarter is somewhat smaller in magnitude and estimated with less precision. Table C.1 in Appendix A reports the corresponding static DiD estimates based on equation (1), showing that municipalities more exposed to the reform increased consumption by approximately 4.5 percent relative to less exposed municipalities.

Robustness: Alternative specifications and geographical spillovers. The baseline estimate is robust to a broad set of alternative specifications, exposure measures, and sample definitions. The estimated effect remains close to the baseline when adding predetermined controls, using a continuous exposure measure, redefining exposure, reclassifying treatment status, excluding peak tourism periods, varying the tourism-exclusion threshold, and broadening the sample to include smaller municipalities and observations with zero consumption. Estimated effects generally remain within the 2–5 percent range and become larger when comparing municipalities at the extremes of the exposure distribution, consistent with stronger treatment intensity generating larger responses. Appendix C reports these exercises in detail (Table C.1).

A distinct concern is spatial interdependence: consumption responses in one municipality may spill over into neighboring areas, potentially violating the Stable Unit Treatment Value Assumption (SUTVA). We address this concern by restricting the sample to larger municipalities, aggregating outcomes to large urban areas, and estimating spatial econometric models that explicitly account for geographic dependence and neighboring exposure. Across all specifications, the coefficient on neighboring exposure is small and statistically insignificant, while the direct effect of municipal exposure remains positive and highly significant. These findings suggest that the observed consumption response primarily reflects local effects rather than cross-municipal propagation. Appendix C reports the full set of spatial exercises, including alternative weighting matrices and spatially corrected standard errors (Tables C.2 and C.3).

4.3.2 Exposure measure from Social Security administrative data

Using tax declarations to measure exposure to the 2018 reform allows us to leverage the granularity of our consumption data. However, this approach is subject to potential measurement error, as we lack details on work hours, months employed, and whether declared income came from wages or unemployment benefits. While the use of a binary exposure indicator alleviates these concerns, we further assess the robustness of our findings by constructing an alternative exposure measure based on Social Security administrative data.

We draw on the MCVL described in Section 3.1, whose longitudinal information on individuals' complete employment trajectories allows us to construct high-quality measures of pre-reform exposure at the local level. Using this data, we compute workers' 2018 hourly wages directly from contributions and hours, a more precise measure of the policy's bite than the annual earnings observed in tax records.[19]

Relative to tax data, the primary limitation of the MCVL is its lower level of geographic detail. The dataset reports municipality of residence only for individuals living in localities with more than 40,000 inhabitants; for smaller areas, residence is identified at the province level. To maximize geographic coverage, we retain the 171 municipalities exceeding this population threshold and aggregate workers residing in smaller municipalities to the province level (50 provinces). This aggregation not only expands the set of usable geographic units but also provides an additional robustness check against potential

[19] As is standard in administrative data, constributions are top and bottom-coded. We recover uncensored hourly wages following Bonhomme and Hospido (2017). Appendix D details the full construction of the exposure measure.

spillovers across small neighboring municipalities.

Finally, following previous literature on minimum wage (Card and Krueger 1994 and Dustmann et al. 2021), we exploit the richness of the MCVL to compute a continuous exposure measure for each geographic area, which we denote GAP. This measure captures both the share of workers directly affected by the reform and the intensity of the implied wage shock. It is defined as follows:

$$\mathrm{GAP}_g = \frac{\sum_{i \in g} h_i \max\{0,\ \mathrm{MW} - w_i\}}{\sum_{i \in g} h_i\, w_i}, \tag{3}$$

where h_i denotes the total number of hours worked by worker i in 2018 residing in geographic unit g, MW is the new 2019 minimum wage, and w_i is the worker's hourly wage in 2018. The numerator is the additional pay required to lift every worker in g up to the new minimum; dividing by the unit's wage bill and multiplying by 100 yields the percentage increase in the aggregate wage bill needed to comply with the reform.[20]

Empirically, we estimate a version of equation (1) in which we again split the treatment variable above and below the median exposure value.[21] Given the smaller number of geographic areas, we do not include region-by-time fixed effects in this specification. Instead, our preferred regression incorporates baseline controls—(log) population, share of women, and share of the working-age population in 2018—interacted with time fixed effects.

The results are presented in Table 1. Column (1) reports estimates from a two-way fixed effects model without controls; column (2) adds our baseline controls interacted with time; and column (3) uses the continuous treatment measure. All coefficients are positive and highly significant. Our preferred specification indicates that treated areas experienced a 4.8 percent increase in consumption following the reform, consistent with our baseline findings. Columns (4)–(6) replicate the analysis using provinces as the geographic unit. The results remain positive and comparable in magnitude, with our preferred estimate showing a 4.2 percent increase in consumption due to the minimum wage reform. Overall, these findings strengthen our confidence in the validity of the empirical strategy and alleviate concerns related to the exposure measure or potential spatial spillovers, as the results here are also robust to aggregation at larger geographic levels.

[20] Table A.4 in Appendix A presents the 10 geographic areas that were most and least affected by the reform.

[21] Table A.5 in Appendix A compares pre-reform characteristics of municipalities with low and high exposure to the MW reform using this MCVL-based criterion, showing patterns very similar to those obtained using municipality-level tax data. The higher level of geographic aggregation required for the MCVL analysis naturally reduces the influence of missing values, measurement error, and idiosyncratic outliers, which tend to be more prevalent in small or highly touristic areas. This greater statistical stability allows us to use the full consumption dataset provided by the bank in this exercise.

Table 1: Consumption Responses Using Social Security-Based Measures of Minimum Wage Exposure

	MCVL groups			Province groups		
	No Controls	Controls	Cont.	No Controls	Controls	Cont.
	(1)	(2)	(3)	(4)	(5)	(6)
Treat x Post 2019	0.034	0.048	0.108	0.030	0.042	0.124
	(0.007)***	(0.006)***	(0.021)***	(0.011)***	(0.010)***	(0.025)***
Controls	No	Yes	Yes	No	Yes	Yes
Mean outcome	3.61	3.61	3.61	3.40	3.40	3.40
Observations	5304	5304	5304	1200	1200	1200
Groups	221	221	221	50	50	50

Notes: Standard errors clustered at the geographic unit level in parentheses. The dependent variable is the logarithm of real per capita consumption. Columns (1)–(3) exploit variation across 221 MCVL geographic units (municipalities with more than 40,000 inhabitants and provinces for smaller municipalities). Columns (4)–(6) aggregate all units to the 50 Spanish provinces. Treatment is defined as above-median exposure to the 2019 MW reform, based on the hourly wage gap measure computed from the MCVL (see equation (3)). Columns (1) and (4) include only geographic-area and time fixed effects. Columns (2) and (5) add baseline controls (i.e., the 2018 values of the logarithm of population, share of working-age population, and the share of women) interacted with time fixed effects. Columns (3) and (6) replace the binary treatment indicator with the continuous exposure. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. *Source:* MCVL, INE, and Banc Sabadell.

4.4 The composition of the consumption response

Having established the overall positive impact of the minimum wage reform on aggregate consumption, we now examine whether this effect varies across different spending categories. Given the limited evidence of geographic spillovers documented above, we return to our preferred municipality-level specification, which allows us to exploit finer geographic variation in consumption responses. Understanding how households reallocate their expenditures in response to an income increase offers valuable insights into consumption priorities and the mechanisms driving the aggregate consumption response.

We analyze the impact of the 2019 MW increase across eleven broad consumption categories. Note that the more disaggregated the data, the greater the likelihood that some municipalities will have missing values for specific categories in a given month. To maintain the sample as comparable as possible to that used in our baseline analysis, we consider the average per capita consumption by municipality for each category over the years 2018 and 2019.[22]

Figure 4 presents the DiD results, revealing significant heterogeneity across consumption groups. The overall increase in consumption observed earlier is primarily driven by higher expenditures on electronics (20.2 percent), leisure (11.7 percent), and restaurants and hotels (8.7 percent). This suggests that the additional household income generated by the minimum wage increase was largely directed toward more income-elastic, discretionary goods, reflecting a broadening of the consumption basket.

In contrast, we observe a decline in spending on food and clothing (6.7 percent). This reduction, alongside the increase in expenditures on restaurants and hotels, suggests

[22] Table A.6 in Appendix A provides details on the number of municipalities included for each category.

Figure 4: DiD: Consumption components

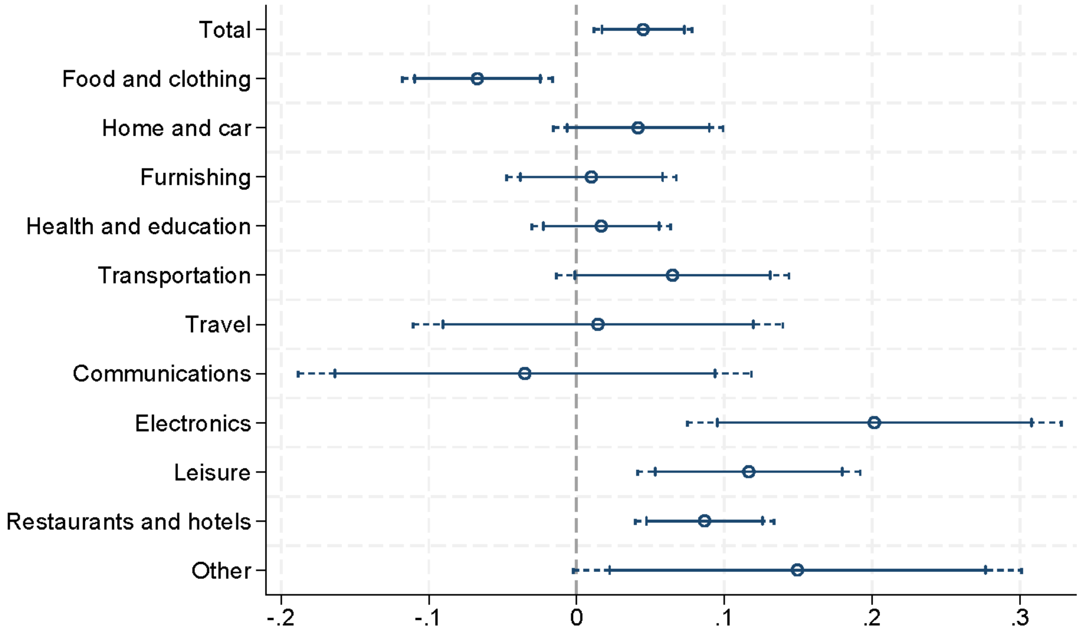

Notes: Each dot reports the DiD coefficient from equation (1), using the yearly average logarithm of real per capita consumption in the indicated category as the dependent variable. Horizontal lines show 90% (solid) and 95% (dashed) confidence intervals, with standard errors clustered at the municipality level. Regressions include municipality and region-by-time fixed effects. *Source:* IEF, AEAT, INE, and Banc Sabadell.

a potential substitution effect, with households shifting from home-prepared meals to dining out. Alternatively, the decline may reflect budget reallocation, as newly affordable discretionary items—such as electronics—prompt households to reduce spending on goods that, while typically treated as staples, are not entirely inelastic, such as certain more substitutable food or clothing items. We explore this in greater detail using household-level data.[23]

Finally, no significant effects are detected for the remaining categories. The stability in transportation and communications spending may reflect the lower income elasticity of these categories, as such expenses are typically tied to commuting and work-related requirements rather than discretionary choices. Likewise, the lack of response in health and education expenditures may be explained by Spain's public healthcare and education systems, which provide comprehensive coverage and thereby insulate these categories from short-term income fluctuations. Similarly, the absence of significant effects in housing-related categories—such as home and car expenses or furnishings—may be due to indivisibilities: these items often involve large, one-time purchases that require prior savings or access to credit, making them less responsive to short-term increases in income, particu-

[23] Additionally, we observe a sizable but less precisely estimated increase in expenditures classified as "other" (15.0 percent).

larly for credit-constrained, low-income households. In any case, although not statistically significant at conventional levels, both home and car and transportation expenditures display a positive response that is near significance, suggesting a potential, albeit modest, response to the reform.[24]

4.5 First-stage

The reduced-form estimates presented above show that the minimum wage reform raised aggregate local consumption. The natural mechanism behind this effect is that workers at the bottom of the wage distribution experienced meaningful income gains as a result of the reform. In this section, we verify this channel directly using Social Security administrative data from the MCVL, which allow us to examine wage changes across the full distribution within geographic areas.

Relying solely on average income measures, such as those publicly available at the municipal level, may obscure the effects of a policy that primarily raises wages for individuals at the bottom of the distribution. In particular, we compute hourly wages at different points of the distribution and test whether wages in the lower part of the distribution increased more in municipalities that were more exposed to the reform. Following the discussion in Section 4.3.2, we estimate a first-stage event-study specification that adapts Equation (2) to the MCVL setting:

$$\ln(Wh^i_{gt}) = \alpha_g + \delta_t + \sum_{q=-4}^{4} \beta_q \cdot (\text{Treat}_g \times \text{Quarter}_q) + X_g + \epsilon_{gt}, \tag{4}$$

where $\ln(Wh^i_{gt})$ is the natural logarithm of the real hourly wage at different percentiles i of the wage distribution at year t; α_g and δ_t denote, respectively, fixed effects at the MCVL geographic area and time fixed effects; Treat_g is a binary indicator equal to 1 if geographic area g has above-median exposure to the MW increase, X_g represents our baseline controls interacted with time fixed effects, and ϵ_{gt} is the error term, clustered at the geographic unit level. The coefficient β_q captures the ATT of the MW increase on income.

Figure 5 presents the event-study results. Two findings stand out. First, there is no evidence of differential pre-trends across areas with high and low exposure, supporting the parallel trends assumption in this setting. Second, following the reform, wages at the bottom of the distribution rose significantly more in higher-exposure areas. The hourly wage at the 5th percentile increased by 4.8 percent in treated relative to control areas,

[24] In Appendix E, we show that a simple framework with nonhomothetic preferences can rationalize our findings by illustrating how an income gain can generate larger proportional increases in spending on discretionary goods.

with this effect reaching approximately 2.5 percent at the 10th percentile. Effects decay rapidly up the distribution and are negligible above the 30th percentile, precisely as one would expect from a policy targeting the wage floor.[25]

Figure 5: Event study: First-stage

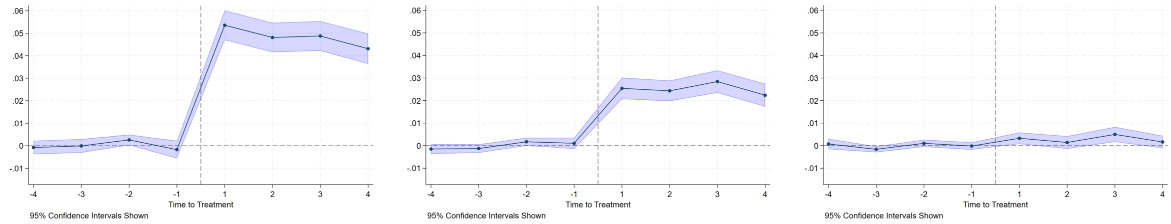

(a) Pct 5 (b) Pct 10 (c) Pct 30

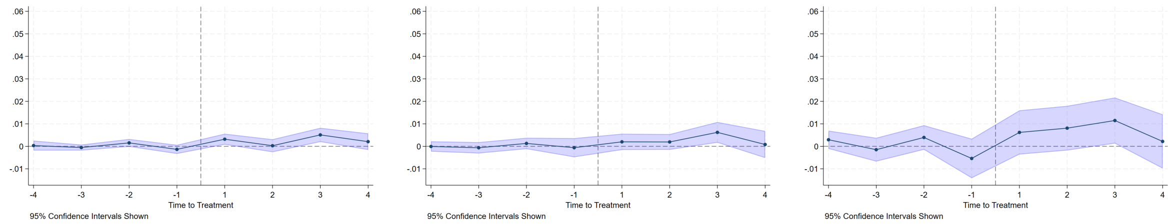

(d) Pct 50 (e) Pct 70 (f) Pct 90

Notes: Dynamic DiD coefficients from equation (4). The dependent variable in each panel is the log of the real hourly wage at the indicated percentile of the wage distribution within each MCVL geographic unit. The horizontal axis shows quarters relative to the reform ($q = 1$ is the first post-reform quarter). Shaded areas are 95% confidence intervals, with standard errors clustered at the geographic-unit level. Regressions include geographic-unit and time fixed effects, along with baseline controls interacted with time fixed effects. 221 geographic-units; 5,304 observations. *Source:* MCVL, INE, and Banc Sabadell.

Table A.8 in Appendix A further reports effects on the total wage bill accruing to workers in each segment of the distribution, providing a robustness check on the hourly wage results and a measure more directly relevant to the consumption channel. While hourly wages capture the price of labor, the wage bill reflects the aggregate purchasing power transferred to each group, accounting for the number of workers and hours worked. What drives household spending is not the wage rate alone, but the total income flowing to affected workers. Consistent with this, Panel A shows that the wage bill increased by 8.6 percent for workers in the bottom 0–5 percentile range and by 4.6 percent for those in the 5–10 percentile range, with effects declining monotonically up the distribution. These

[25] Table A.7 in Appendix A shows the results of the static difference-in-differences where we also show the effect on the average hourly wage. We find a much smaller effect at the upper end of the wage distribution, although it remains positive and statistically significant. This pattern may reflect regional trends related to the business cycle. To address this concern, we replicate the analysis using yearly data for the period 2015–2019 and explicitly account for potential differences in pre-trends across geographic areas similar to Dustmann et al. (2021). The results for the lower part of the distribution remain very similar, while effects above the 30th percentile are no longer statistically different from zero. The results are available upon request. Importantly, this issue should not affect our baseline consumption estimates, since those regressions include region-by-time fixed effects and therefore compare municipalities within the same broader geographic area.

wage bill gains represent a meaningful income injection concentrated among high-MPC households, providing direct evidence linking income gains to the aggregate consumption response documented above.[26]

5 Household approach

5.1 Data and empirical strategy

Our previous analyses rely on transaction data from POS terminals and credit card payments. These data capture only the card component of spending, so any consumption settled in cash is not directly observed. If households adjust cash and card spending differently in response to the reform, or if the cash-card mix varies across spending categories, some effects may not be fully reflected. To address this, we complement the geographic analysis with household-level evidence on total consumption and an alternative identification strategy.

Specifically, we use the EPF (see Section 3.2) exploiting its rotating panel structure, in which a subset of households is interviewed in two consecutive waves.[27] Because the minimum-wage reform directly affects wage workers, we restrict the sample to households whose primary earner was in salaried employment in 2018, before the reform took effect. We focus on the 2018–2019 period to compare household outcomes immediately before and after the policy change and estimate the following specification:[28]

$$\Delta y_{ht} = \beta_0 + \beta_1 \text{Treat}_h + \gamma \mathbf{X}_{ht} + \epsilon_{ht}, \tag{5}$$

where y_{ht} denotes household consumption expenditure or household income, measured alternatively in levels and logs. Monetary variables are expressed in 2021 euros and adjusted

[26]Panel B of Table A.8 decomposes the wage bill effect using the accounting identity $\ln(\text{wage bill}) = \ln(N) + \ln(h/N) + \ln(Wh)$, where N denotes employment, h/N hours worked per employee, and Wh the average hourly wage within each segment. Overall employment increased by a modest 0.8 percent, an effect discussed in detail in Section 7.2. Since percentile groups are recalculated after accounting for employment changes, the employment component is similar across groups by construction; the informative margins are therefore hours per worker and the average hourly wage. At the bottom 0–5 percentile range, the 8.6 percent wage bill effect reflects a 3.0 percent increase in hours worked per employee and a 4.7 percent increase in the average hourly wage. For the 5–10 percentile range, the 4.6 percent wage bill effect is primarily driven by the increase in the average hourly wage (3.5 percent), with hours per worker contributing a small and statistically insignificant share; the segment-level wage estimate of 3.5 percent lies between the point estimates at the 5th and 10th percentiles (4.8 and 2.5 percent, respectively), as expected. Taken together, the wage bill results corroborate the hourly wage estimates; both show a monotonically decreasing effect of the minimum wage reform across the distribution, confirming that the reform generated income gains concentrated at the bottom.

[27]An earlier draft of this paper used the EPF as a repeated cross-section.

[28]See Campos and Reggio (2015), who use a similar specification with the EPF to study the effects of unemployment on consumption.

using the modified OECD equivalence scale for household comparability. Treat_h is a binary indicator equal to 1 if household h belongs to the treatment group and 0 otherwise. A household is considered treated if its primary earner is a wage earner with monthly earnings below €1,000 in 2018.[29] Vector $\mathbf{X}_{ht}$ is a set of control variables comprising geographic factors (region fixed effects and municipality-size dummies in five categories), demographic characteristics (age in three groups, gender, and education level in four categories), household structure (number of children, household size, number of workers adjusted for part-time employment, and the partnership status of the primary earner), job characteristics (sector in four groups, occupation in six categories, temporary contract status, and part-time status), and wealth proxies (homeownership, ownership of a second residence, mortgage status, and residential-area dummies in seven groups). Consistent with the existing literature (e.g., Campos and Reggio, 2015), variables that affect the level of consumption enter the model in first differences, whereas variables shaping a household's consumption profile enter in levels measured at $t-1$. Variables affecting the consumption level include household size, the number of children, the part-time–weighted number of workers, and the partnership status of the primary earner. All remaining controls enter in levels at $t-1$ to capture the household's consumption profile; partnership status enters in both first differences and lagged levels, as it influences both the level and the profile of consumption.

5.2 Results

Panel A of Table 2 presents population-weighted main results.[30] In our preferred specification with the full set of controls, individuals affected by the reform increased their

[29] Although the EPF contains information on individual income, income data are often missing, and the survey does not indicate whether other household members are wage earners or self-employed. Accordingly, we define treatment based on the income range of the primary earner. For our income measure, we use per-capita household income (OECD scale), an imputed variable provided in the EPF to reduce missing values. We exclude households whose primary earner is older than 60, since minimum-wage policies primarily affect active labor-force participants and older individuals are more likely to be retired or weakly attached to the labor market. We also exclude households for which the reported gender of the primary earner changes across waves, as such changes likely reflect a change in the identity of the primary earner or a substantial shift in household composition, reducing comparability over time.

[30] Table A.9 in Appendix A reports consumption composition by ECOICOP 2-digit category for the household sample. While these categories are not directly comparable to the aggregated groupings used in Table A.2 for the transaction data, some broad patterns are worth noting. Expenditure in restaurants and hotels is similarly represented in both datasets, accounting for approximately 13 percent of the total. However, other categories exhibit notable differences. In the household data, combined spending on food, alcohol, and clothing represents around 22 percent of total consumption, compared to roughly 33 percent for the aggregated "Food and clothing" category in the transaction data. Transportation accounts for a larger share in the household data (18 vs. 10 percent), as does "Home" at 13 vs. 6 percent. Conversely, furnishings have a larger weight in the transaction data (14 vs. 5 percent). These differences are consistent with heterogeneity in payment methods across product types, but should be interpreted with caution given that the category definitions are not directly aligned across datasets.

consumption by 10.2 percent in 2019 relative to the control group. When consumption is measured in levels, treated households increased their consumption by approximately €969 more between 2018 and 2019 than comparable control households. Omitting population weights yields similar estimates—9.4 percent and €857 (see Table A.10 in Appendix A).[31]

Table 2: Consumption: Household Budget Survey 2018-2019

	(1)	(2)	(3)	(4)	(5)	(6)
			Panel A: $\Delta \ln C_{ht}$			
Treat	0.081	0.086	0.082	0.093	0.090	0.102
	(0.030)***	(0.031)***	(0.032)***	(0.032)***	(0.031)***	(0.034)***
Mean outcome	9.40	9.40	9.40	9.40	9.40	9.40
			Panel B: ΔC_{ht}			
Treat	711.684	813.221	733.207	843.809	823.721	968.617
	(347.227)**	(352.041)**	(375.867)*	(379.822)**	(379.011)**	(457.223)**
Mean outcome	13981.11	13981.11	13981.11	13981.11	13981.11	13981.11
Observations	3009	3009	3009	3009	3009	3009
Geographic controls	No	Yes	Yes	Yes	Yes	Yes
Wealth controls	No	No	Yes	Yes	Yes	Yes
Demographic controls	No	No	No	Yes	Yes	Yes
Household structure controls	No	No	No	No	Yes	Yes
Job characteristics controls	No	No	No	No	No	Yes

Notes: Robust standard errors in parentheses. The dependent variable is the change in real consumption per capita, adjusted using the OECD equivalence scale, winsorized at the 1st and 99th percentiles. Panel A reports estimates in logs while Panel B uses levels. A household is considered treated if its primary earner is a wage earner with monthly earnings below €1,000 in 2018. Column (1) includes no controls. Column (2) adds geographic controls (region fixed effects and municipality-size dummies). Column (3) further includes wealth controls (homeownership, ownership of a second residence, mortgage status, and residential-area dummies). Column (4) adds demographic controls (age, gender, and education level). Column (5) further includes household structure controls (number of children, household size, number of workers, and partnership status). Column (6) adds job characteristics controls (sector, occupation, temporary contract status, and part-time status). Mean outcome refers to the average level of the dependent variable in period $t-1$. ${}^{*}p < 0.10$, ${}^{**}p < 0.05$, ${}^{***}p < 0.01$. *Source:* EPF.

An important advantage of using household-level data is that it allows us to estimate the MPCs, which provide a direct measure of how additional income translates into household spending. Unlike reduced-form consumption effects, MPCs are informative about the strength of consumption smoothing, liquidity constraints, and potential aggregate demand implications of income policies.

Using exposure to the minimum wage reform as an instrument for income changes, Table 3 confirms that the instrument is both relevant and strong. In the log specification, treated households experienced income growth that was 14.1 percent higher than that of control households. The level specification tells a consistent story: treated households experienced an increase in income of approximately €1,215 more than comparable control households. The resulting IV estimates imply a consumption elasticity of 0.72 and a marginal propensity to consume of 0.80 for affected households. This suggests that roughly 80 percent of reform-induced income gains are translated into additional consumption.

[31] In placebo regressions using pre-reform data from 2017–2018, differential consumption arises only marginally in log specifications without controls and is absent in our preferred specification with the full set of controls (see Table A.11 in Appendix A)

The magnitude is consistent with estimates in the literature for low-income and liquidity-constrained households, which typically report MPCs well above those of higher-income groups and often in the range of 0.6 to 0.9 following income shocks or transfers (Johnson et al., 2006; Parker et al., 2013; Jappelli and Pistaferri, 2014).[32]

Table 3: First Stage and IV Estimates on Consumption

	Log Outcomes		Level Outcomes	
	First Stage	IV	First Stage	IV
	(1)	(2)	(3)	(4)
Treat	0.141	0.721	1215.478	0.797
	(0.022)***	(0.251)***	(241.140)***	(0.386)**
Mean outcome	9.43	9.40	14298.51	13981.11
Observations	3009	3009	3009	3009
F-stat (first-stage)		40.68		25.41

Notes: Robust standard errors in parentheses. Columns (1) and (3) report the first-stage estimates, where the dependent variable is the change in real household income per capita. Columns (2) and (4) report the IV estimates, where the dependent variable is the change in real household consumption per capita. Per capita values are obtained using the OECD equivalence scale, and dependent variables are winsorized at the 1st and 99th percentiles. Columns (1) and (2) use changes in logs while columns (3) and (4) use changes in levels. All specifications include geographic, wealth, demographic, household structure, and job characteristic controls. The F-statistic refers to the first-stage test of instrument relevance. Mean outcome refers to the average level of the dependent variable in period $t-1$. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. *Source:* EPF.

Table 4: Consumption components: Household Budget Survey (Levels)

	Food	Alcohol	Clothing	Home	Furniture	Health
	(1)	(2)	(3)	(4)	(5)	(6)
Treat	148.991	104.065	129.020	28.533	65.799	90.746
	(91.131)	(37.798)***	(96.757)	(81.569)	(63.519)	(87.855)
Mean outcome	2114.78	455.44	1041.13	1972.46	782.59	651.90
Observations	2990	1739	2321	2987	2892	1910

	Transportation	Communications	Leisure	Education	Restaurants	Other
	(7)	(8)	(9)	(10)	(11)	(12)
Treat	201.181	–8.670	213.775	74.345	46.612	150.587
	(304.901)	(18.080)	(114.786)*	(89.870)	(117.514)	(60.112)**
Mean outcome	2787.12	523.87	1234.98	737.09	2042.13	1217.30
Observations	2789	2955	2550	1005	2627	2996

Notes: Robust standard errors in parentheses. The dependent variable is the change in real consumption per capita, adjusted using the OECD equivalence scale, in each ECOICOP category, winsorized at the 1st and 99th percentiles. Each coefficient corresponds to the treatment estimate from equation (5), estimated separately for each consumption category. All specifications include geographic, wealth, demographic, household structure, and job characteristic controls. Mean outcome refers to the average level of the dependent variable in period $t-1$. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. *Source:* EPF.

Finally, Table 4 presents disaggregated results by 2-digit ECOICOP consumption categories. Each coefficient corresponds to β_1 from equation (5), estimated using consumption

[32] In a recent study, González Salmerón (2025) estimates a similar MPC exploiting the 2019 MW increase using a different methodology.

in a given ECOICOP category as the dependent variable. Although these categories are not fully comparable to those in Section 4.4—as the Banc Sabadell data do not follow the ECOICOP classification—we again find consumption responses concentrated in discretionary items, such as leisure and alcohol. Further disaggregation at the 3-digit level shows that the increase in leisure consumption is primarily driven by recreational activities. Similarly, the positive response in the "other" category is largely explained by higher spending on nonessential personal care (see table A.13 in Appendix A). In contrast, estimates for non-discretionary consumption categories are not statistically different from zero.[33]

6 The magnitude of local aggregate effects

Local demand spillovers arise when an income shock generates local spending beyond the direct response of the households that receive the initial income gain. A key advantage of the municipality-level estimates in Section 4 is that they capture the total local spending response to the reform, including any amplification operating through local demand channels. By contrast, the household-level consumption responses estimated in Section 5 primarily reflect the spending response of affected households. Comparing the two therefore provides a way to gauge the extent of local demand amplification generated by the reform.

Let τ_L and τ_H denote the estimated municipality-level and household-level consumption responses, respectively. We relate the two through

$$\tau_L = m(s_T - s_C)\tau_H, \tag{6}$$

where s_T and s_C denote the share of consumption accounted for by affected households in treated and control municipalities, respectively. The parameter m measures local demand amplification. A value of $m = 1$ implies no amplification beyond the direct spending response of affected households, whereas $m > 1$ indicates the presence of local demand spillovers.

Connecting the two estimates is nevertheless challenging. The transaction data cap-

[33] In contrast to the transaction-level results, the household survey data yield no statistically significant effect on spending in restaurants and hotels. One possible explanation is differences in measurement in the EPF: restaurant expenditures are recorded in a two-week expenditure diary and may therefore be imperfectly captured if spending in this category is relatively infrequent, while accommodation expenditures are collected using a three-month retrospective recall window. Both features may attenuate estimated effects relative to transaction data that observe spending continuously. Importantly, the transaction-data result does not appear to be driven by differential selection into card use among treated households: restricting the survey sample to individuals below age 35, who are more likely to rely on card payments within low-income groups, also yields no statistically significant effect.

ture only card-based expenditures, whereas the household survey measures total consumption. Because both estimates are expressed as percentage changes, differences in the underlying consumption basket directly affect their magnitudes. In particular, expenditure categories such as housing services, utilities, and insurance enter the denominator of the household consumption measure but are largely absent from the transaction data. Since these categories are unlikely to respond in the short run, the household estimate mechanically understates the spending response for the subset of consumption categories captured in the transaction data. A second challenge is that card-based spending may respond differently from total spending even after adjusting for differences in expenditure coverage. Consequently, the observed amplification factor m^{obs} combines both local demand amplification and differences in the responsiveness of card-based and total consumption. Letting ϕ denote the latter, we have $m^{obs} = m\phi$, so recovering the true amplification factor requires an assumption about ϕ.

We begin by adjusting the household consumption response to account for housing services, utilities, and other differences in consumption coverage between the household and transaction data.[34] This adjustment increases the estimated household consumption response τ_H from 10.2% to 13.9%. We then scale this response by the difference in exposure across municipalities. Affected households account for 14.8% of consumption in treated municipalities and 6.7% in control municipalities.[35] Meanwhile, our municipality-level estimates of τ_L imply a spending increase of 4.5%. Finally, recovering m requires an assumption about ϕ. We calibrate $\phi \in [1.25, 1.5]$, implying that card-recorded expenditures respond more strongly than total consumption within the adjusted expenditure basket. Because we do not observe this relationship directly, this range should be interpreted only as a plausible calibration. The assumption is motivated by evidence that card payments are disproportionately used for larger and more discretionary purchases (Klee, 2008; Shy, 2023), together with our finding that the spending response to the minimum wage increase is concentrated in discretionary consumption categories. Under these assumptions, we obtain an amplification factor of

$$m \in [2.6, 3.2]. \tag{7}$$

[34] Specifically, we exclude ECOICOP Group 4 (housing, water, electricity, gas, and other fuels), Subgroup 7.1 (purchase of vehicles, but not vehicle operation or other transportation expenditures), Group 10 (education, primarily enrollment fees), and Subgroup 12.5 (insurance services). This adjustment is conservative, as the household survey includes additional expenditure categories that are unlikely to be observed in card transaction data.

[35] Because municipality-level consumption shares are not observed, we approximate s_T and s_C using the average share of affected individuals in treated and control municipalities (22.75% and 10.34%, respectively), scaled by the ratio of average consumption among affected households to average household consumption in the EPF (65.02%).

The amplification factor can be mapped into a conventional spending multiplier under two additional assumptions. First, the marginal propensity to consume estimated from the household data is representative of the first-round spending response generated by the reform at the local level. Second, we interpret the amplification factor m as measuring how much total local spending is ultimately generated for each euro of first-round spending by affected households. This interpretation relies on the adjustments described above, which allow the ratio of percentage spending responses to serve as an approximation to the corresponding ratio of spending changes in levels. Under these assumptions, the local spending multiplier can be approximated as

$$\mu \approx m \times \text{MPC} \tag{8}$$

Using our estimate of MPC $= 0.8$ and the range for m in (7), we obtain

$$\mu \in [2.1, 2.5] \tag{9}$$

These values should be interpreted as back-of-the-envelope calculations rather than precise estimates. In particular, the reported range reflects only the calibration band on ϕ, and it does not incorporate sampling uncertainty in τ_L, τ_H, or the MPC, each of which is estimated with error. With that caveat, the calculation suggests that the local spending effects of the reform were substantially larger than the first-round response of affected households alone.[36]

This range is somewhat above estimates of local output multipliers for government spending shocks, which are often found to lie between 1.5 and 2 (e.g., Nakamura and Steinsson, 2014; Suárez Serrato and Wingender, 2016; Chodorow-Reich, 2019; Auerbach et al., 2020; Dupor et al., 2023). Several factors may explain why our estimates are larger. First, the income gains accrue directly to low-wage workers with high marginal propensities to consume, generating particularly strong local demand shocks in areas where high-MPC households are geographically concentrated. Second, unlike many government spending shocks, the reform transfers purchasing power directly to households rather than through intermediated channels. Moreover, because the MW increase is expected to be permanent, households may perceive it as an increase in lifetime resources and increase current spending by borrowing against higher expected future earnings. This mechanism can strengthen the initial demand response to the reform, consistent with the debt-financed spending responses to MW increases documented by Aaronson et al. (2012), as well as broader evidence that financially constrained households exhibit strong

[36]The case $\phi = 1$, which would arise if we assume the same response from cards and cash, implies an amplification factor of approximately 3.9 and a spending multiplier of 3.1. For the reasons discussed above, we regard values of $\phi > 1$ as more plausible.

consumption responses to changes in wealth (Mian et al., 2013). If local demand responses are nonlinear, a stronger initial demand impulse may also contribute to larger multiplier effects.[37]

Finally, our estimates should be interpreted as capturing general-equilibrium effects at the local level rather than the national level. Quantifying a national multiplier would require a fully specified general-equilibrium model with endogenous price and quantity adjustments, which lies beyond the scope of this paper.

7 Firm-side adjustments

Minimum wage increases affect household consumption through multiple adjustment margins. While the primary effect is an increase in real earnings for covered workers, firms may respond by adjusting employment, prices, profits, or reorganizing production to absorb higher labor costs. When firms mainly adjust profits or improve efficiency through production reorganization, wage gains translate into higher real consumption. By contrast, if firms respond by increasing prices, part of the minimum wage gain is lost through an erosion of purchasing power. If the response instead takes the form of employment reductions, these may offset total income gains and dampen aggregate consumption. As a result, the net effect on consumption reflects the combined incidence of these channels. In this section, we examine firm-side adjustments to better understand how firms absorbed the increase in labor costs resulting from the minimum wage increase.

7.1 The role of prices

From a theoretical perspective, firms may pass through part of a minimum wage increase into higher prices, either to offset higher production costs or in response to increased demand from workers earning higher wages.

As discussed above, our benchmark analysis relies on real consumption measures obtained by deflating nominal expenditures using province-, month-, and ECOICOP-specific CPIs from INE. This deflation strategy absorbs all price variation common to municipalities within a given province–month–category cell, thereby accounting for substantial price movements potentially induced by the minimum wage reform. However, it does not capture residual price differences across municipalities within the same province.

If such within-province price variation were economically meaningful and systematically related to minimum-wage exposure, our baseline estimates could partly reflect local

[37] We find some evidence of changes in hours worked and employment that may occur only when the initial demand shock is sufficiently large (see Table A.8 and Footnote 26), consistent with the possibility of nonlinear local demand propagation.

Figure 6: Event study: The role of prices

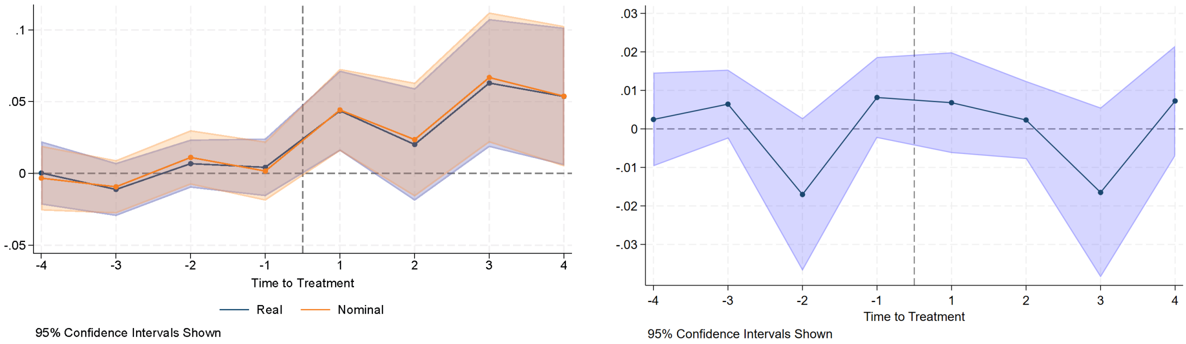

(a) Real/nominal consumption

(b) Prices

Notes: Coefficients are estimated from a dynamic DiD specification. The horizontal axis represents quarters relative to the reform ($q = 1$ is the first post-reform quarter). Shaded areas indicate 95% confidence intervals, with standard errors clustered at the municipality level in panel (a) and at the sector level in panel (b). In Figure 6.a, we estimate a version of equation (2) where the dependent variable is the log of real and nominal per-capita consumption. Regressions include municipality and region-by-time fixed effects, and the sample is restricted to municipalities with a balanced panel of 24 monthly observations (1,999 municipalities; 47,976 observations). In Figure 6.b, we estimate a version of equation (4) where the dependent variable is the log of the quarterly price index, with sectors as the unit of analysis. The regression includes sector and time fixed effects, along with baseline controls interacted with time dummies, and is restricted to sectors with at least 100 workers in the MCVL (70 sectors; 560 observations). *Source:* IEF, AEAT, MCVL, INE, Banc Sabadell, and the Business Margins Observatory.

price adjustments rather than genuine changes in real consumption. Because municipality-level price indices are unavailable, we assess the relevance of this concern through two complementary exercises.

First, we re-estimate the effect of the 2019 minimum wage reform using nominal rather than real consumption expenditures. This specification leaves all price variation—both across and within provinces—in the outcome variable. The logic is the following. If municipalities more exposed to the reform experienced relatively larger price increases, estimated treatment effects should be larger when using nominal consumption. In that case, part of the baseline effect could plausibly be driven by unobserved within-province price variation not captured by province-level CPIs.

Figure 6.a compares the event-study estimates based on real consumption with those obtained using nominal consumption. The estimates are nearly identical across specifications, with no meaningful differences in either magnitude or dynamics. This similarity suggests that price adjustments associated with cross-municipality differences in minimum-wage exposure across provinces are not quantitatively important for our results. While this exercise cannot directly test for price adjustments occurring within provinces, this concern is further mitigated by the province-level analysis in Section 4.3.2, in which both consumption and prices are measured at the province level. The fact that we detect a similar increase in consumption in that setting supports the interpretation that our baseline results reflect real consumption changes rather than local price effects.

To further investigate the potential role of prices, we conduct a second exercise using sector-level price data from the Business Margins Observatory, which provides quarterly price indices for 72 sectors.[38] Using this sector-level variation, we examine whether sectors with higher exposure to the minimum wage increase exhibit systematically different price dynamics relative to less exposed sectors. Following Section 4.3.2 and using the MCVL, we compute the intensity of the wage shock at the sector level using equation (3). Empirically, we estimate a two-way fixed effects model for 2018–2019 in which the dependent variable is the log of the price index, and the treatment variable is defined as above or below the median exposure value across sectors.[39] Figure 6.b presents the event-study results, which show no effect of the reform on sectoral prices.[40]

Although this analysis is conducted at a different level of aggregation than the municipal specifications, sectoral variation is particularly well suited to capture short-run cost-push price effects. Increases in labor costs should translate most immediately into sectoral prices, whereas demand-driven price pressures are likely to materialize more gradually. The absence of detectable sector-level price responses therefore suggests limited short-run price pass-through following the reform.

Overall, these findings reinforce the robustness of our baseline findings and indicate that the observed consumption responses are driven by real changes in spending rather than by price adjustments.

7.2 The labor market

The labor market effects of minimum wage increases have been extensively studied, yet theoretical predictions remain ambiguous, depending on factors such as the degree of labor market competition, labor demand elasticities, and income–consumption elasticities. Although our primary focus is on the impact of the minimum wage on consumption, we examine labor market outcomes for completeness. Specifically, we analyze the effect of Spain's 2019 minimum wage reform on employment and unemployment using two complementary exercises.

First, we apply the same empirical strategy as in our baseline consumption analysis (Equation (2)), using monthly unemployment rates at the municipal level as the dependent variable. The unemployment rate is computed as the ratio of registered unemployed

[38] Sector classification is based on NACE (Nomenclature of Economic Activities) codes at varying levels of aggregation. A complete list is provided in Table A.14 in Appendix A.

[39] Table A.15 in Appendix A presents the 10 sectors that were most and least affected by the reform. The regression controls for the share of women, immigrants, and workers with tertiary education, as well as average age and occupation rank in 2018, all obtained from the MCVL and interacted with calendar-time dummies. Our final sample excludes two sectors—Tobacco and Water Transport—for which fewer than 100 workers are observed in the MCVL to compute exposure. Including these sectors does not affect the results.

[40] The coefficient (standard error) of the static differences-in-differences regression is -0.0007 (0.0049).

individuals from SEPE to the working-age population obtained from INE.[41]

Second, we use administrative data from the MCVL to construct monthly measures of employment at a broader geographic level. In this exercise, we compute the number of individuals employed in each month and geographic unit and replicate the specification used in Sections 4.3.2 and 4.5 to study the effects of the reform on consumption and income using administrative data (Equation (4)).

Both exercises yield consistent results, reported in Figure A.3 in Appendix A. We find no evidence of differential pre-trends in either outcome. Following the reform, unemployment rates remain essentially unchanged in more-exposed municipalities relative to less-exposed ones. The MCVL employment estimates, if anything, point to a small positive effect on employment.[42]

Comparing Figures 3 and A.3, we conclude that while the minimum wage increase led to an immediate and sizable increase in consumption, it had no discernible short-run impact on unemployment. Several mechanisms could rationalize this finding, and we remain agnostic about their relative importance. One possibility is that employment effects are muted due to monopsony power in local labor markets. Alternatively, any negative employment effects may have been offset by general equilibrium forces: increased consumption among low-wage workers may have stimulated local demand, an effect captured by our consumption analysis. Finally, employment adjustments may occur more gradually than consumption responses and therefore fall outside the time horizon of our study. In any case, for the purposes of our analysis, the absence of short-run employment effects supports the interpretation that the observed consumption response is not driven by changes in labor market attachment.

7.3 Other margins

Given the absence of aggregate effects on prices and employment, a natural question arises: who bore the cost of the reform? We investigate this by examining sectoral adjustments using firm-level data from the Bank of Spain's *Central de Balances Integrada* (CBI).[43]

We employ the same empirical specification used for the sectoral price analysis in Section 7.1. Sector-level exposure to the wage shock is computed using the MCVL, with

[41] We use the working-age population because labor force data are only available at the provincial level. Using province-level labor force data instead yields very similar results.

[42] In a recent study analyzing the interaction between collective agreements and minimum wages, Riudavets (2025) finds a similar result by exploiting province-sector exposure to the reform. Results are unchanged when using hours worked instead of employment as the dependent variable.

[43] The CBI contains financial information for non-financial companies obtained from annual accounts filed in the *Registro Mercantil*. Since the data are consolidated at the company level and do not include establishment-level details, geographic variation cannot be exploited.

sectors classified as treated or control based on median exposure.[44] As the CBI provides annual data, we focus on the period 2015–2019. This yields four pre-reform years to test for parallel trends, while limiting the post-reform window to 2019 to avoid contamination from the COVID-19 pandemic.

We quantify the reform's impact using sector-level aggregates for value added, profits, number of firms, and number of workers—variables that capture the overall scale of economic activity.[45] We additionally examine wages, labor productivity, and profitability per worker, as well as estimated markups and markdowns.[46]

The event-study results are reported in Figure A.4 in Appendix A.[47] The results reveal a distinct pattern, though estimates are generally imprecise and warrant cautious interpretation. We find no statistically significant effect on total sectoral value added or profits, suggesting aggregate economic activity in more exposed sectors remained stable. However, these sectors experienced contractions in both the number of firms (-1.7 percent) and the number of workers (-3 percent), alongside an increase in average wages of approximately 2.7 percent, consistent with the policy's intent. The combination of stable aggregate output with fewer firms and workers implies rising productivity and profitability per worker—up 2.0 and 2.2 percent, respectively—a pattern consistent with the exit of lower-productivity marginal firms unable to absorb higher wage costs. Markups and markdowns show no meaningful changes, indicating that the reform did not substantially alter the pricing power or cost structures of surviving firms.

Taken together, these sectoral results point to a reallocation mechanism as the primary channel of adjustment. The cost of the reform appears to have been absorbed through the exit of less productive firms and a contraction in employment within the most exposed sectors. The displaced workers and economic activity were likely reallocated elsewhere in the economy, consistent with our earlier findings of stable aggregate employment at the geographic level. Thus, the burden of the reform appears to have fallen disproportionately on marginal firms, contributing to an increase in average productivity without reducing

[44] All regressions include sector and time fixed effects, along with time-interacted controls for the 2018 share of women, immigrants, and workers with tertiary education, as well as average age and worker rank. Pre-reform characteristics of low- and high-exposure sectors are compared in Appendix A Table A.16.

[45] Using means instead of aggregates would primarily reflect composition changes driven by firm exits rather than total sectoral activity. To mitigate the influence of errors, outliers, and compositional effects, we discard firm-year observations with non-positive values in sales, value added, workers, or labor costs. We then winsorize the top and bottom 1 percent of value added, profits, and labor costs within each sector. Finally, we restrict the sample to firms present for at least three years within the 2015–2019 window.

[46] Wages, labor productivity, and profitability per worker are computed as the sectoral sum of labor costs, value added, or profits divided by the number of workers in the sector. Markups and markdowns are estimated following Yeh et al. (2022), based on the "production function approach" of Loecker and Warzynski (2012).

[47] Table A.17 in Appendix A shows the corresponding static DiD estimates for the 2018–2019 period.

total sectoral output or affecting competitive pricing.

8 Conclusions

This paper provides new evidence on how minimum wage increases affect household consumption, exploiting the large minimum wage increase implemented in Spain in 2019. Using high frequency municipal level transaction data from credit cards and point of sale records, complemented by household survey data, we assess the extent to which higher wages translate into increased consumer spending.

Using transaction data, we find that the reform increased aggregate local consumption substantially, by about 4.5 percent, with estimates ranging between 3.6 and 5.3 percent, depending on the specification. This is consistent with evidence of sizable income gains at the bottom of the wage distribution. The impact on consumption is concentrated in discretionary spending categories such as electronics, leisure, restaurants, and hotels. This evidence suggests that low-wage workers responded to higher incomes by reallocating spending toward income-elastic goods and services. Using household survey data, we find evidence pointing in the same direction: consumption rises overall and shifts toward discretionary categories. This response is consistent with nonhomothetic preferences, whereby the composition of spending shifts as income rises.

Although our data do not allow us to precisely estimate a consumption multiplier, a back-of-the-envelope calculation provides suggestive evidence of local aggregate effects, with a multiplier exceeding the multipliers estimated in the literature for government spending shocks. At the same time, we find no evidence of significant price pass through or aggregate employment losses. Instead, sectoral patterns point to adjustment through firm dynamics. Sectors more exposed to the reform experienced reductions in the number of firms and workers, while total value added and profits remained broadly stable, a pattern consistent with the exit of lower productivity firms. Taken together, these findings suggest that Spain's 2019 minimum wage increase expanded the consumption set of low wage households without measurable employment losses in the short run, implying potential welfare gains for workers. Assessing the medium run implications, including the extent to which economic activity and employment reallocated across sectors or locations, remains an important avenue for future research.

References

Aaronson, D., Agarwal, S., and French, E. (2012). The spending and debt response to minimum wage hikes. *American Economic Review*, 102(7):3111–39.

AEA (2023). Anuario estadístico de la banca en España 2023. Asociación Española de Banca. Accessed: February 27, 2025. Available at: https://s1.aebanca.es/wp-content/uploads/2024/07/anuario-2023.pdf.

Alonso, C. (2022). Beyond labor market outcomes: The impact of the minimum wage on nondurable consumption. *Journal of Human Resources*, 57(5):1690–1714.

Ashenfelter, O. and Jurajda, S. (2022). Minimum wages, wages, and price pass-through: The case of Mcdonald's restaurants. *Journal of Labor Economics*, 40(S1):S179–S201.

Attanasio, O. P. and Davis, S. J. (1996). Relative wage movements and the distribution of consumption. *Journal of Political Economy*, 104(6):1227–1262.

Auerbach, A. J., Gorodnichenko, Y., and Murphy, D. (2020). Local fiscal multipliers and fiscal spillovers in the USA. *IMF Economic Review*, 68:195–229.

Autor, D. H., Dorn, D., and Hanson, G. H. (2013). The China syndrome: Local labor market effects of import competition in the United States. *American Economic Review*, 103(6):2121–68.

Autor, D. H., Manning, A., and Smith, C. L. (2016). The contribution of the minimum wage to US wage inequality over three decades: A reassessment. *American Economic Journal: Applied Economics*, 8(1):58–99.

Azar, J., Huet-Vaughn, E., Marinescu, I., Taska, B., and von Wachter, T. (2024). Minimum wage employment effects and labor market concentration. *The Review of Economic Studies*, 91(4):1843–1883.

Bartik, A. W., Currie, J., Greenstone, M., and Knittel, C. R. (2019). The local economic and welfare consequences of hydraulic fracturing. *American Economic Journal: Applied Economics*, 11(4):105–155.

BELab (2025). Banco de España/CORPME, Colegio de Registradores de la Propiedad y Mercantiles de España. CBI Database, DOI: 10.48719/BELab.CBI9522_01. Accessed: 2026-01-09.

Black, D., McKinnish, T., and Sanders, S. (2005). The economic impact of the coal boom and bust. *The Economic Journal*, 115(503):449–476.

Bonhomme, S. and Hospido, L. (2017). The cycle of earnings inequality: Evidence from Spanish social security data. *The Economic Journal*, 127(603):1244–1278.

Caliendo, M., Fedorets, A., Preuss, M., Schröder, C., and Wittbrodt, L. (2018). The short-run employment effects of the German minimum wage reform. *Labour Economics*, 53:46–62.

Caliendo, M. and Wittbrodt, L. (2022). Did the minimum wage reduce the gender wage gap in Germany? *Labour Economics*, 78:102228.

Callaway, B., Goodman-Bacon, A., and Sant'Anna, P. H. (2024). Event studies with a continuous treatment. In *AEA Papers and Proceedings*, volume 114, pages 601–605.

Campos, R. G. and Reggio, I. (2015). Consumption in the shadow of unemployment. *European Economic Review*, 78:39–54.

Card, D. (1992). Using regional variation in wages to measure the effects of the federal minimum wage. *ILR Review*, 46(1):22–37.

Card, D. and Krueger, A. B. (1994). Minimum wages and employment: A case study of the fast-food industry in New Jersey and Pennsylvania. *The American Economic Review*, 84(4):772–793.

Cengiz, D., Dube, A., Lindner, A., and Zipperer, B. (2019). The effect of minimum wages on low-wage jobs. *The Quarterly Journal of Economics*, 134(3):1405–1454.

Chodorow-Reich, G. (2019). Geographic cross-sectional fiscal spending multipliers: What have we learned? *American Economic Journal: Economic Policy*, 11(2):1–34.

Clemens, J. (2021). How do firms respond to minimum wage increases? Understanding the relevance of non-employment margins. *Journal of Economic Perspectives*, 35(1):51–72.

Clemens, J. and Wither, M. (2019). The minimum wage and the Great Recession: Evidence of effects on the employment and income trajectories of low-skilled workers. *Journal of Public Economics*, 170:53–67.

Colella, F., Lalive, R., Sakalli, S. O., and Thoenig, M. (2019). Inference with arbitrary clustering. Technical Report IZA Discussion Paper 12584, Institute of Labor Economics (IZA).

Cooper, D., Luego-Prado, M. J., and Parker, J. A. (2020). The local aggregate effects of minimum wage increases. *Journal of Money, Credit and Banking*, 52(1):5–35.

Dautović, E., Hau, H., and Huang, Y. (2026). Consumption response to minimum wages: Evidence from Chinese households. *The Review of Economics and Statistics*, 108(3):737–754.

de Chaisemartin, C., D'Haultfœuille, X., and Vazquez-Bare, G. (2024). Difference-in-difference estimators with continuous treatments and no stayers. In *AEA Papers and Proceedings*, volume 114, pages 610–613.

Deaton, A. (1992). *Understanding Consumption*. Clarendon Lectures in Economics. Oxford University Press, Oxford.

Dickens, R., Machin, S., and Manning, A. (1999). The effects of minimum wages on employment: Theory and evidence from Britain. *Journal of Labor Economics*, 17(1):1–22.

Dube, A. (2019). Minimum wages and the distribution of family incomes. *American Economic Journal: Applied Economics*, 11(4):268–304.

Dube, A., Lester, T. W., and Reich, M. (2010). Minimum wage effects across state borders: Estimates using contiguous counties. *The Review of Economics and Statistics*, 92(4):945–964.

Dupor, B., Karabarbounis, M., Kudlyak, M., and Mehkari, M. S. (2023). Regional consumption responses and the aggregate fiscal multiplier. *Review of Economic Studies*, 90(6):2982–3021.

Dustmann, C., Lindner, A., Schönberg, U., Umkehrer, M., and vom Berge, P. (2021). Reallocation effects of the minimum wage. *The Quarterly Journal of Economics*, 137(1):267–328.

Engbom, N. and Moser, C. (2022). Earnings inequality and the minimum wage: Evidence from Brazil. *American Economic Review*, 112(12):3803–47.

European Central Bank (2022). Minimum wages and their role for euro area wage growth. *ECB Economic Bulletin*, 3. Focus – Box 4.

Ferrando, L. and Posada, D. (2023). Uso del efectivo y otros medios de pago:? Cómo está cambiando la forma en que pagamos? *Boletín Económico, Banco de España*, (2023/T1).

Fitzgerald, J., Gottschalk, P., and Moffitt, R. (1998). An analysis of sample attrition in panel data: The Michigan Panel Study of Income Dynamics. *Journal of Human Resources*, 33(2):251–299.

Geary, R. C. (1950). A note on "A constant utility index of the cost of living". *The Review of Economic Studies*, 18(1):65–66.

Gonzálvez Salmerón, R. M. (2025). *Consumption Responses to the Minimum Wage*. Ph.d. dissertation, Universidad Complutense de Madrid, Madrid, Spain. Chapter 5, coauthored with Daniel Fernández, Paloma Villanueva and Luis Cárdenas. Defended 18 November 2025. Available at https://hdl.handle.net/20.500.14352/132140.

Gorjon, L., Martinez de LaFuente, D., and Romero, G. (2024). Employment effects of the minimum wage: Evidence from the Spanish 2019 reform. *SERIEs*, 15:1–55.

Harasztosi, P. and Lindner, A. (2019). Who pays for the minimum wage? *American Economic Review*, 109(8):2693–2727.

INE (2024). Porcentaje de trabajadores en función de su ganancia con respecto al salario mínimo interprofesional (SMI) por tipo de jornada y sexo. Accessed: 2025-05-17. Available at: https://www.ine.es/jaxiT3/Tabla.htm?t=28182.

Jappelli, T. and Pistaferri, L. (2014). Fiscal policy and MPC heterogeneity. *American Economic Journal: Macroeconomics*, 6(4):107–36.

Jiménez, B. (2023). The political economy of the minimum wage. *Labour Economics*, 85:102463.

Johnson, D. S., Parker, J. A., and Souleles, N. S. (2006). Household expenditure and the income tax rebates of 2001. *American Economic Review*, 96(5):1589–1610.

Klee, E. (2008). How people pay: Evidence from grocery store data. *Journal of Monetary Economics*, 55(3):526–541.

Leung, J. H. (2021). Minimum wage and real wage inequality: Evidence from pass-through to retail prices. *The Review of Economics and Statistics*, 103(4):754–769.

Link, S. (2024). The price and employment response of firms to the introduction of minimum wages. *Journal of Public Economics*, 239:105236.

Loecker, J. D. and Warzynski, F. (2012). Markups and firm-level export status. *American Economic Review*, 102(6):2437–2471.

Meer, J. and West, J. (2016). Effects of the minimum wage on employment dynamics. *Journal of Human Resources*, 51(2):500–522.

Meyer, B. D., Mok, W. K. C., and Sullivan, J. X. (2015). Household surveys in crisis. *Journal of Economic Perspectives*, 29(4):199–226.

Mian, A., Sufi, A., and Rao, K. (2013). Household balance sheets, consumption, and the economic slump. *Quarterly Journal of Economics*, 128(4):1687–1726.

Nakamura, E. and Steinsson, J. (2014). Fiscal stimulus in a monetary union: Evidence from US regions. *American Economic Review*, 104(3):753–92.

Parker, J. A., Souleles, N. S., Johnson, D. S., and McClelland, R. (2013). Consumer spending and the economic stimulus payments of 2008. *American Economic Review*, 103(6):2530–2553.

Rao, N. L. and Risch, M. (2025). Who's afraid of the minimum wage? Measuring the impacts on independent businesses using matched U.S. tax returns. *The Quarterly Journal of Economics*, 141(1):373–427.

Renkin, T., Montialoux, C., and Siegenthaler, M. (2022). The pass-through of minimum wages into U.S. retail prices: Evidence from supermarket scanner data. *The Review of Economics and Statistics*, 104(5):890–908.

Riudavets, M. (2025). *The Interaction between Minimum Wages and Collective Bargaining: Explaining Wage Spillovers*. Ph.d. dissertation, University of Helsinki, Finland. Chapter 5, coauthored with Ferran Elias. Defended 20 August 2025. Available at http://hdl.handle.net/10138/599131.

Shy, O. (2023). Cash is alive: How economists explain holding and use of cash. *Journal of Economic Literature*, 61(4):pp. 1465–1520.

Stewart, M. B. (2002). Estimating the impact of the minimum wage using geographical wage variation. *Oxford Bulletin of Economics and Statistics*, 64:583–605.

Stone, R. (1954). Linear expenditure systems and demand analysis: An application to the pattern of British demand. *The Economic Journal*, 64(255):511–527.

Suárez Serrato, J. C. and Wingender, P. (2016). Estimating local fiscal multipliers. Working Paper 22425, National Bureau of Economic Research.

Yagan, D. (2019). Employment hysteresis from the Great Recession. *Journal of Political Economy*, 127(5):2505–2558.

Yeh, C., Macaluso, C., and Hershbein, B. (2022). Monopsony in the US labor market. *American Economic Review*, 112(7):2099–2138.

APPENDIX A: Supplementary tables and figures

Table A.1: Correlation between province-level under-coverage of low-income tax declarations and worker characteristics, 2018

	(1)	(2)	(3)	(4)
% Women	0.086			
	(0.099)			
% Young		0.122		
		(0.075)		
% Low skilled			0.008	
			(0.051)	
% Low education				0.022
				(0.038)
Mean outcome	10.69	10.69	10.69	10.69
Observations	46	46	46	46

Notes: Standard errors clustered at the province level in parentheses. The dependent variable is the ratio of non-declared, non-mandatory tax declarations to the total number of tax declarations presented at the province level. This ratio serves as a proxy for the "relative under-coverage of low-income individuals" in the tax data, i.e., the share of potentially missing low-wage workers relative to the observed filing population. We are not able to compute a voluntary filing ratio because the tax declaration data does not include information on the total number of payers, which determines the income threshold for filing voluntarily. The table shows the correlation between this ratio and four proxies for groups of workers likely to be more affected by the minimum wage reform. Data on tax declarations come from the IEF. The sample of non-declared, non-mandatory declarations includes 42,520 observations, compared to 3,011,866 declarations in the main sample. The relatively small number of observations makes it unfeasible to use this information at the municipality level in our main analysis. To construct province-level ratios, we apply the IEF elevation factor to make the two samples comparable. The sample covers 46 provinces; Ceuta and Melilla are excluded, and data for the Basque Country and Navarra are unavailable due to separate regional tax rules. Proxies for worker characteristics are obtained from the Labour Force Survey microdata: *Women* indicates the share of female workers; *Young* refers to the share of workers under 35; *Low skilled* refers to the share of workers in services (restoration, personal services, protection, or retail) or in elementary/unskilled occupations; *Low education* refers to the share of workers with, at most, lower secondary education. * p<0.1, ** p<0.05, *** p<0.01. *Source:* IEF, AEAT, and Labor Force Survey.

Table A.2: Consumption composition

Components	% Total consumption 2018	% Total consumption 2019
Food and clothing	32.1	34.7
Home and car	6.1	5.9
Furnishing	14.0	14.2
Health and education	8.8	8.7
Transportation	10.9	9.6
Travel	4.3	3.7
Communication	0.6	0.5
Electronics	3.8	3.8
Leisure	5.0	4.8
Restaurants and hotels	13.4	13.2
Other	1.0	1.0
Avg. per capita consumption	€38.2	€41.0

Notes: Consumption composition using monthly per capita averages by year from our econometric sample (1,999 municipalities). Food and clothing include expenditures on food products and apparel. Home and car covers housing-related expenses and vehicle repairs. Furnishing includes expenditures on home improvement, furniture, houseware, and other retail goods. Health and education encompass expenses related to health and education. Transportation includes spending on petrol stations, parking, various modes of transportation, and tolls. Travel considers expenditures on plane tickets and travel-related lodging. Communication covers telecommunications expenses. Electronics covers spending on electronic devices and photography equipment. Leisure comprises expenditure on products related to sports, toys, lotteries, pets, and other recreational items. Restaurants and hotels include spending at bars, restaurants, hotels. Other incorporates expenses on insurance, donations, taxes, and ATMs. *Source:* Banc Sabadell and INE.

Table A.3: Pre-reform characteristics (2018)

	Low exposure		High exposure		
	Mean	Median	Mean	Median	Difference
Unemployment rate	8.67	8.21	10.10	9.63	-1.42***
	(3.52)		(4.05)		
Population	26,074.25	4628	10,127.23	3528	15,947.01***
	(128,525.33)		(21,896.74)		
Share 16-65	65.29	65.90	64.92	65.54	0.37*
	(3.63)		(4.23)		
Share women	49.54	49.91	49.11	49.35	0.43***
	(2.05)		(2.03)		
Gross income pc	14,655.63	14,357.50	12,288.76	12,052.00	2,366.87***
	(3,703.03)		(2,410.57)		
Gini	30.08	29.80	28.96	28.90	1.12***
	(3.38)		(2.65)		
Municipalities	1000		999		

Notes: Pre-reform characteristics of municipalities with low and high exposure to the 2019 MW reform. Low (high) exposure includes municipalities below (above) the median exposure to the MW 2019 reform. Standard deviations in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. *Source:* SEPE and INE.

Table A.4: Geographic areas most and least exposed (in %)

Top 10 Most Exposed Areas: Area	%		Top 10 Least Exposed Areas: Area	%
La Orotava	0.895	1	Álava (prov)	0.032
Cáceres (prov)	0.877	2	Boadilla del Monte	0.035
Arona	0.831	3	Pozuelo de Alarcón	0.036
Badajoz (prov)	0.814	4	Basauri	0.044
Santa Cruz de Tenerife (prov)	0.757	5	Pamplona	0.050
Adeje	0.701	6	San Cugat del Vallès	0.054
Ourense (prov)	0.666	7	San Sebastián	0.056
Granadilla de Abona	0.629	8	El Prat de Llobregat	0.058
Telde	0.611	9	Gipuzkoa (prov)	0.061
Ourense	0.608	10	Las Rozas de Madrid	0.070

Notes: Intensity gap (in %) relative to the new 2019 minimum wage, as defined in equation (3). *Source:* MCVL.

Table A.5: Pre-reform characteristics (2018): MCVL

	MCVL groups					Province groups				
	Low exposure		High exposure			Low exposure		High exposure		
	Mean	Median	Mean	Median	Difference	Mean	Median	Mean	Median	Difference
Unemployment rate	8.92	8.84	13.13	12.58	-4.21^{***}	8.45	8.23	12.78	13.57	-4.33^{***}
	(2.36)		(2.85)			(1.52)		(2.09)		
Population	226,918	96,405	194,214	99,433	32,704	1,076,823	580,229	785,235	676,376	291,589
	(392,302)		(214,066)			(1,602,576)		(510,809)		
Share 16–65	65.94	65.78	67.10	67.11	-1.15^{**}	64.94	64.75	65.77	66.29	-0.83
	(2.07)		(3.01)			(1.47)		(2.69)		
Share women	51.12	51.15	50.75	50.63	0.37^{*}	50.66	50.66	50.57	50.55	0.09
	(1.36)		(1.34)			(0.92)		(0.66)		
Gross income pc	16,536.53	15,500.00	12,119.34	12,051.00	$4,417.19^{***}$	15,459.70	14,909.04	11,995.03	11,708.18	$3,464.66^{***}$
	(4,166.14)		(1,760.97)			(1,545.90)		(958.08)		
Gini	31.51	30.90	32.12	32.00	-0.61	30.47	30.09	31.26	30.68	-0.79
	(3.02)		(2.47)			(1.57)		(1.59)		
Observations	111		110			25		25		

Notes: Pre-reform characteristics of geographic units with low and high exposure to the 2019 MW reform. Low (high) exposure includes areas below (above) the median exposure to the MW 2019 reform. Standard deviations in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. *Source:* SEPE and INE.

Table A.6: DiD: Consumption components

	Total	Food and clothing	Home and car	Furnishing	Health and education	Transportation
	(1)	(2)	(3)	(4)	(5)	(6)
Treat x Post 2019	0.045	–0.067	0.042	0.010	0.017	0.065
	(0.017)***	(0.026)***	(0.029)	(0.029)	(0.024)	(0.040)
Mean outcome	2.62	0.93	–0.31	–0.33	–0.64	0.45
Observations	3998	3722	3258	3314	3450	3388
Municipalities	1999	1861	1629	1657	1725	1694

	Travel	Communications	Electronics	Leisure	Restaurants and hotels	Other
	(7)	(8)	(9)	(10)	(11)	(12)
Treat x Post 2019	0.014	–0.035	0.202	0.117	0.087	0.150
	(0.064)	(0.078)	(0.065)***	(0.038)***	(0.024)***	(0.077)*
Mean outcome	–1.48	–3.06	–1.21	–1.04	0.43	–3.18
Observations	1518	1256	2430	2718	3822	1212
Municipalities	759	628	1215	1359	1911	606

Notes: Standard errors clustered at the municipality level in parentheses. The dependent variable is the logarithm of real per capita consumption in each category. Each column reports the static DiD estimate from equation (1), using the yearly average log per capita consumption in the indicated category as the dependent variable. All regressions include municipality and region-by-time fixed effects.The number of municipalities varies across categories due to missing values in some consumption categories for certain municipalities. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. *Source:* IEF, AEAT, INE, and Banc Sabadell.

Table A.7: First-stage estimates: Hourly wage

	Average	Pct 5	Pct 10	Pct 30	Pct 50	Pct 70	Pct 90
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
				Hourly wage			
Treat x Post 2019	0.005	0.048	0.025	0.003	0.003	0.003	0.007
	(0.001)***	(0.001)***	(0.001)***	(0.001)***	(0.001)***	(0.001)***	(0.002)***
Mean outcome	2.57	1.92	2.00	2.22	2.39	2.62	3.06
Observations	5304	5304	5304	5304	5304	5304	5304
Groups	221	221	221	221	221	221	221

Notes: Standard errors clustered at the geographic unit level in parentheses. The dependent variable is the natural logarithm of the real hourly wage at the indicated percentile. Treatment is defined as above-median exposure to the 2019 MW reform, based on the hourly wage gap measure computed from the MCVL (see equation (3)). All specifications include geographic-area and time fixed effects as well as baseline controls (i.e., the 2018 values of the logarithm of population, the share of working-age population, and the share of women) interacted with time fixed effects. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. *Source:* MCVL and INE.

Table A.8: First-stage estimates: Wage bill and decomposition

	Total	Pct 0–5	Pct 5–10	Pct 20–30	Pct 40–50	Pct 60–70	Pct 80–90
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Wage bill							
Treat × Post 2019	0.012	0.086	0.046	0.023	0.010	0.009	0.006
	(0.002)***	(0.004)***	(0.004)***	(0.003)***	(0.003)***	(0.003)***	(0.002)***
Mean outcome	14.74	10.92	11.02	11.96	12.20	12.45	12.82
Panel B: Decomposition							
(i) Employment							
Treat × Post 2019	0.008	0.009	0.008	0.008	0.008	0.009	0.009
	(0.002)***	(0.002)***	(0.002)***	(0.002)***	(0.002)***	(0.002)***	(0.002)***
(ii) Hours per worker							
Treat × Post 2019	0.001	0.030	0.003	0.010	–0.002	–0.003	–0.008
	(0.001)	(0.003)***	(0.003)	(0.002)***	(0.002)	(0.001)**	(0.002)***
(iii) Average hourly wage							
Treat × Post 2019	0.003	0.047	0.035	0.004	0.003	0.003	0.006
	(0.001)***	(0.001)***	(0.001)***	(0.001)***	(0.001)***	(0.001)***	(0.001)***
Observations	5304	5304	5304	5304	5304	5304	5304
Groups	221	221	221	221	221	221	221

Notes: Standard errors clustered at the geographic unit level in parentheses. Panel A reports the effect on the total wage bill accruing to workers in each percentile segment of the wage distribution. Panel B decomposes the wage bill effect using the accounting identity $\ln(\text{wage bill}) = \ln(N) + \ln(h/N) + \ln(Wh)$, where N denotes employment, h/N hours worked per employee,and Wh the average hourly wage within each segment. Since percentile groups are assigned by recalculating the wage distribution after accounting for employment changes, the employment component in row (i) is similar across groups by construction; the informative margin is rows (ii) and (iii). Treatment is defined as above-median exposure to the 2019 MW reform, based on the hourly wage gap measure computed from the MCVL (see equation (3)). All specifications include geographic-area and time fixed effects as well as baseline controls (i.e., the 2018 values of the logarithm of population, the share of working-age population, and the share of women) interacted with time fixed effects.$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. *Source*: MCVL and INE.

Table A.9: Consumption composition: Household Budget Survey

Components	% Total consumption 2018	% Total consumption 2019
Food	15.9	16.1
Alcohol	2.1	2.1
Clothing	6.4	6.8
Home	12.8	13.6
Furniture	5.4	4.7
Health	3.6	3.9
Transportation	18.5	17.7
Communications	3.7	3.6
Leisure	7.6	7.3
Education	2.4	2.4
Restaurants and hotels	12.9	12.9
Other	8.8	8.9
Avg. consumption per capita	€13981.11	€13834.26

Notes: Consumption composition by ECOICOP 2-digit category for the household econometric sample (3,009 households). Each row reports the share of total consumption accounted for by each category in 2018 and 2019. Average consumption per capita is obtained using the OECD equivalence scale and expressed in real terms. *Source:* EPF.

Table A.10: Consumption: Household Budget Survey, robustness

	Consumption		Income		IV	
	Logs	Levels	Logs	Levels	Logs	Levels
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Without weights						
Treat	0.094	857.023	0.131	1119.806	0.720	0.765
	(0.028)***	(349.520)**	(0.018)***	(183.494)***	(0.227)***	(0.329)**
Mean outcome	9.37	13467.81	9.43	14049.80	9.37	13467.81
Observations	3009	3009	3009	3009	3009	3009
F-stat (first stage)					51.10	37.24
Panel B: Trimmed bottom and top 5%						
Treat	0.101	964.866	0.154	1236.101	0.659	0.781
	(0.027)***	(295.515)***	(0.023)***	(225.338)***	(0.180)***	(0.263)***
Mean outcome	9.39	12766.17	9.44	13683.76	9.39	12766.17
Observations	2709	2709	2709	2709	2709	2709
F-stat (first stage)					43.78	30.09

Notes: Robust standard errors in parentheses. Columns (1) and (2) report reduced-form estimates where the dependent variable is the change in real household consumption per capita. Columns (3) and (4) report first-stage estimates where the dependent variable is the change in real household income per capita. Columns (5) and (6) report IV estimates where the dependent variable is the change in real household consumption per capita. Per capita values are obtained using the OECD equivalence scale. Panel A replicates the baseline estimates without population weights, with dependent variables winsorized at the 1st and 99th percentiles. Panel B trims the bottom and top 5% of the dependent variable. Odd columns use changes in logs while even columns use changes in levels. All specifications include geographic, wealth, demographic, household structure, and job characteristic controls. The F-statistic refers to the first-stage test of instrument relevance. Mean outcome refers to the average level of the dependent variable in period $t-1$. $^{*}p<0.10$, $^{**}p<0.05$, $^{***}p<0.01$. *Source:* EPF.

Table A.11: Consumption: Household Budget Survey 2017-2018

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: $\Delta \ln C_{ht}$						
Treat	0.056	0.049	0.044	0.043	0.040	0.048
	(0.024)**	(0.024)**	(0.026)	(0.028)	(0.028)	(0.031)
Mean outcome	9.40	9.40	9.40	9.40	9.40	9.40
Panel B: ΔC_{ht}						
Treat	337.369	262.418	172.368	193.396	115.570	207.470
	(291.717)	(302.558)	(337.881)	(363.339)	(364.164)	(411.301)
Mean outcome	13930.40	13930.40	13930.40	13930.40	13930.40	13930.40
Observations	3302	3302	3302	3302	3302	3302
Geographic controls	No	Yes	Yes	Yes	Yes	Yes
Wealth controls	No	No	Yes	Yes	Yes	Yes
Demographic controls	No	No	No	Yes	Yes	Yes
Household structure controls	No	No	No	No	Yes	Yes
Job characteristics controls	No	No	No	No	No	Yes

Notes: Robust standard errors in parentheses. The dependent variable is the change in real consumption per capita, adjusted using the OECD equivalence scale, winsorized at the 1st and 99th percentiles. Panel A reports estimates in logs while Panel B uses levels. The table replicates the baseline specification of Table 2 using pre-reform data from 2017–2018 as a placebo test. A household is considered treated if its primary earner is a wage earner with monthly earnings below €1,000 in 2017. Column (1) includes no controls. Column (2) adds geographic controls (region fixed effects and municipality-size dummies). Column (3) further includes wealth controls (homeownership, ownership of a second residence, mortgage status, and residential-area dummies). Column (4) adds demographic controls (age, gender, and education level). Column (5) further includes household structure controls (number of children, household size, number of workers, and partnership status). Column (6) adds job characteristics controls (sector, occupation, temporary contract status, and part-time status). Mean outcome refers to the average level of the dependent variable in period $t-1$. $^{*}p<0.10$, $^{**}p<0.05$, $^{***}p<0.01$. Source: EPF.

Table A.12: Consumption components: Household Budget Survey (Logs)

	Food	Alcohol	Clothing	Home	Furniture	Health
	(1)	(2)	(3)	(4)	(5)	(6)
Treat	0.043	0.055	0.055	0.007	0.199	0.144
	(0.059)	(0.127)	(0.136)	(0.035)	(0.134)	(0.166)
Mean outcome	7.49	5.29	6.45	7.25	5.85	5.64
Observations	2990	1739	2321	2987	2892	1910

	Transportation	Communications	Leisure	Education	Restaurants	Other
	(7)	(8)	(9)	(10)	(11)	(12)
Treat	0.186	–0.013	0.370	–0.115	0.067	0.186
	(0.128)	(0.043)	(0.139)***	(0.249)	(0.102)	(0.061)***
Mean outcome	7.22	6.10	6.43	5.89	7.13	6.79
Observations	2789	2955	2550	1005	2627	2996

Notes:Robust standard errors in parentheses. The dependent variable is the change in the log of real consumption per capita in each ECOICOP category, adjusted using the OECD equivalence scale, winsorized at the 1st and 99th percentiles. Each coefficient corresponds to the treatment estimate from equation (5), estimated separately for each consumption category. All specifications include geographic, wealth, demographic, household structure, and job characteristic controls. Mean outcome refers to the average level of the dependent variable in period $t-1$. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. *Source:* EPF.

Table A.13: Consumption components: Household Budget Survey (Levels, ECOICOP 3-digit)

	09.1 A-V equip.	09.2 Major durables	09.3 Recreational items	09.4 Cultural services	09.5 Books & press	09.6 Package holidays	
	(1)	(2)	(3)	(4)	(5)	(6)	
Treat	7.282	–208.702	–23.404	225.455	55.584	–207.391	
	(47.516)	(172.756)	(102.684)	(97.857)**	(43.247)	(155.954)	
Mean outcome	259.49	148.53	533.51	622.24	249.55	740.69	
Observations	1,048	69	1,356	1,202	1,077	541	

	12.1 Personal care	12.3 Personal effects	12.4 Social protection	12.5 Insurance	12.6 Financial services	12.7 Other services	12.8 Remittances
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treat	99.843	–20.233	–2519.816	11.710	4.447	8.944	112.509
	(42.425)**	(59.631)	(2300.093)	(18.145)	(11.633)	(87.484)	(276.618)
Mean outcome	482.72	211.98	974.13	536.21	54.08	245.94	329.74
Observations	2,693	855	48	2,871	200	444	100

Notes: Robust standard errors in parentheses. The dependent variable is the change in the level of real consumption per capita in each ECOICOP 3-digit category, adjusted using the OECD equivalence scale, winsorized at the 1st and 99th percentiles. Each coefficient corresponds to the treatment estimate from equation (5), estimated separately for each consumption category. The top panel disaggregates ECOICOP 09 (Recreation and culture) into: 09.1 audio-visual, photographic and information processing equipment; 09.2 other major durables for recreation and culture; 09.3 other recreational items, equipment, gardens and pets; 09.4 recreational and cultural services; 09.5 newspapers, books and stationery; 09.6 package holidays. Bottom panel disaggregates ECOICOP 12 (Miscellaneous goods and services) into: 12.1 personal care; 12.3 personal effects n.e.c.; 12.4 social protection; 12.5 insurance; 12.6 financial services n.e.c.; 1 12.7 other services n.e.c.; 12.8 remittances. Category 12.2 is not collected in the Household Budget Survey. Estimates for categories 09.2 ($N = 69$) and 12.4 ($N = 48$) are based on very few observations and should be interpreted with caution. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. *Source:* EPF.

Table A.14: Sector Classification with 2 and 3-digit NACE Codes

#	Sector	NACE	#	Sector	NACE	#	Sector	NACE	#	Sector	NACE
1	Agriculture	01+02+03	19	Electronics	26	37	Wholesale: household	464	55	Accommodation	55+56
2	Mining	05-09	20	Electrical equip.	27	38	Wholesale: ICT	465	56	Publishing	58
3	Food	10	21	Machinery	28	39	Wholesale: machinery	466	57	Film/TV/sound	59
4	Beverages	11	22	Motor vehicles	29	40	Other wholesale	467	58	Telecom	61
5*	Tobacco	12	23	Other transport	30	41	Non-special wholesale	469	59	IT services	62
6	Textiles	13	24	Furniture	31	42	Retail: general	471	60	Information services	63
7	Apparel	14	25	Other manufacturing	32	43	Retail: food/drinks	472	61	Real estate	68
8	Leather/footwear	15	26	Machinery repair	33	44	Retail: fuel	473	62	Legal/accounting	69
9	Wood/cork	16	27	Electricity/gas	35	45	Retail: ICT	474	63	Architecture/engineering	71
10	Paper	17	28	Water/waste	36-39	46	Retail: household goods	475	64	Advertising	73
11	Printing	18	29	Construction	41-43	47	Retail: cultural	476	65	Other professional	74
12	Coke/petroleum	19	30	Vehicle sales	451	48	Retail: other	477	66	Veterinary	75
13	Chemical	20	31	Vehicle maintenance	452	49	Retail: markets	478	67	Rental/leasing	77
14	Pharmaceuticals	21	32	Vehicle parts	453	50	Retail: other channels	479	68	Employment	78
15	Rubber/plastic	22	33	Motorcycles	454	51	Land transport	49	69	Travel agencies	79
16	Non-metallic minerals	23	34	Wholesale intermediaries	461	52*	Water transport	50	70	Security	80
17	Basic metals	24	35	Wholesale: agriculture	462	53	Air transport	51	71	Building services	81
18	Metal products	25	36	Wholesale: food/drinks	463	54	Warehousing	52	72	Office support	82

Notes: The table presents the sector classification used in the sectoral analysis, along with the corresponding 2 and 3-digit NACE codes. Sectors marked with an asterisk are excluded from the empirical analysis due to fewer than 100 workers observed in the MCVL. *Source:* Business Margins Observatory.

Table A.15: Sectors Most and Least Exposed (in %)

Top 10 Most Exposed Sectors			Top 10 Least Exposed Sectors	
Sector	%		Sector	%
Retail: markets	2.560	1	Pharmaceuticals	0.001
Veterinary	2.018	2	Coke/petroleum	0.003
Retail: food/drinks	1.214	3	Other transport	0.005
Office support	0.801	4	Basic metals	0.007
Building services	0.593	5	Electricity/gas	0.012
Real estate	0.565	6	Motor vehicles	0.018
Other manufacturing	0.497	7	Mining	0.023
Agriculture	0.496	8	Machinery	0.029
Retail: cultural	0.473	9	Chemical	0.030
Apparel	0.438	10	Air transport	0.030

Notes: Intensity gap (in %) relative to the new 2019 minimum wage, as defined in equation (3), aggregated at the sector level. *Source:* MCVL.

Table A.16: Pre-reform characteristics (2018): CBI

	Low exposure		High exposure		
	Mean	Median	Mean	Median	Difference
Value added	3,407,853.39	2,728,471.37	2,128,314.48	1,493,035.81	1,279,538.91*
	(2,598,764.23)		(2,457,171.42)		
Profits	1,120,330.84	862,188.27	572,043.67	284,879.97	548,287.17**
	(880,088.00)		(810,700.57)		
Number of companies	4,885.49	2,189.00	6,343.77	3,817.00	-1,458.29
	(9,784.77)		(6,988.76)		
Number of workers	80,453.05	52,708.99	91,736.35	37,512.50	-11,283.30
	(86,038.08)		(115,683.37)		
Wage	32.59	31.48	21.30	23.62	11.29***
	(12.92)		(7.04)		
Productivity pw	58.07	46.67	30.34	30.90	27.73**
	(52.57)		(13.57)		
Profits pw	25.18	14.08	8.53	6.30	16.64*
	(42.38)		(8.13)		
Markup	1.06	1.01	1.15	1.04	-0.09
	(0.23)		(0.31)		
Markdown	2.08	1.68	1.41	1.28	0.67**
	(1.10)		(0.87)		
Average worker age	42.37	42.36	39.86	39.09	2.51***
	(1.78)		(3.11)		
Average worker-rank	6.09	6.39	6.86	6.82	-0.77*
	(1.33)		(1.16)		
Men (%)	73.08	75.60	52.32	52.42	20.76***
	(13.47)		(17.19)		
Spanish (%)	94.00	95.07	90.38	90.59	3.63***
	(3.81)		(4.07)		
Tertiary education (%)	24.44	19.29	19.79	15.06	4.65
	(14.30)		(13.49)		
Observations	35		35		70

Notes: Pre-reform characteristics of sectors with low and high exposure to the 2019 MW reform. Low (high) exposure includes sectors below (above) the median exposure to the MW 2019 reform. Standard deviations in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Monetary values are expressed in 2019 thousands of euros, using a sector-specific deflator. Central de Balances Integrada (BELab 2025) and MCVL.

Table A.17: DiD: Other margins

	Value added	Profits	N. companies	N. workers	Wage	Productivity pw	Profits pw	Markup	Markdown
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Treat x Post 2019	–0.010	–0.008	–0.017	–0.030	0.027	0.020	0.022	–0.001	0.007
	(0.016)	(0.025)	(0.010)	(0.026)	(0.020)	(0.020)	(0.022)	(0.007)	(0.016)
Mean outcome	14.42	13.05	7.87	10.84	3.20	3.58	2.22	0.08	0.37
Observations	140	140	140	140	140	140	140	140	140
Sectors	70	70	70	70	70	70	70	70	70

Notes: Standard errors clustered at the sector level in parentheses. The dependent variables are the logs of value added, profits, number of companies, number of workers, wages, productivity per worker, profits per worker, markups, and markdowns. The table replicates the analysis of Figure A.4 using a static DiD specification restricted to the 2018–2019 period, evaluating the effect of the reform relative to the year immediately preceding it. The regression includes sector and time fixed effects, along with baseline controls interacted with time dummies. Only sectors with at least 100 workers in the MCVL are included (70 sectors; 140 observations). $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. *Source:* Central de Balances Integrada (BELab, 2025) and MCVL.

Figure A.1: Exposure to the MW 2019 reform: Econometric sample

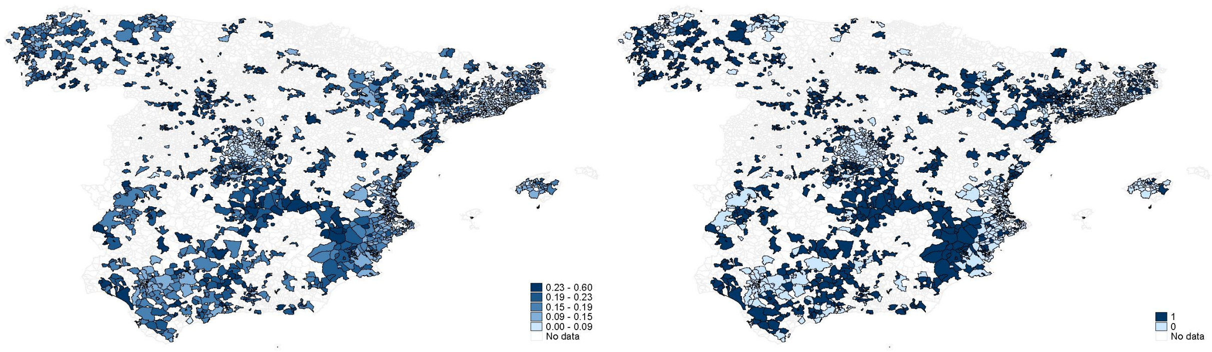

(a) Continuous (b) Dummy

Notes: Share of tax declarations from individuals aged 16–65 with an annual income in the range €8,000–€12,600. The econometric sample is reduced from 5,520 to 1,999 municipalities after requiring a balanced panel of monthly consumption observations and applying additional sample restrictions to mitigate measurement error and the influence of outliers. The sample covers 78% of the total population (approximately 36 million people). *Source:* IEF, AEAT, and INE.

Figure A.2: Event study: Consumption

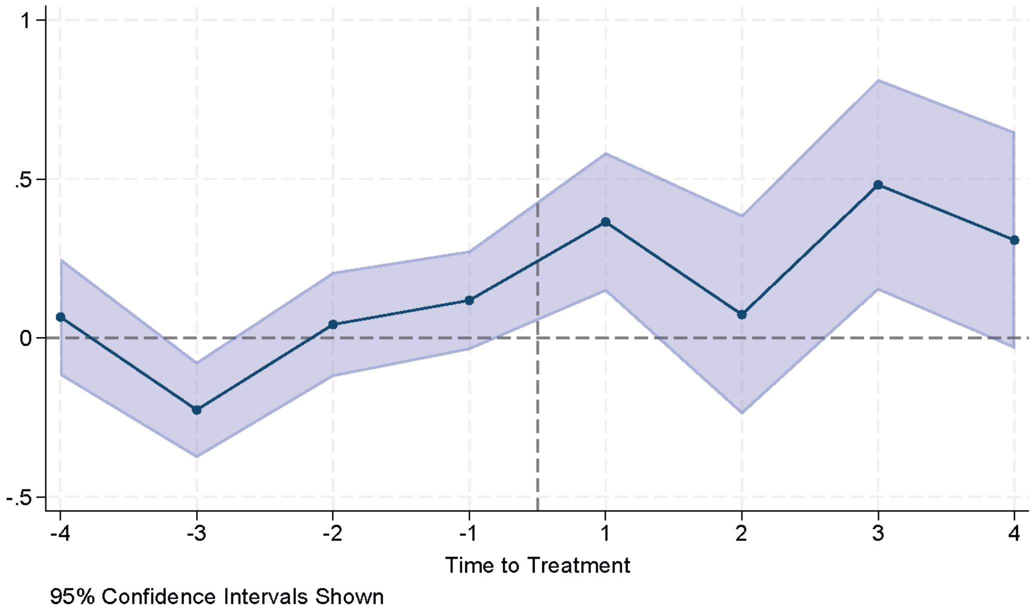

Notes: Dynamic DiD coefficients from a version of equation (2) using the continuous treatment variable. The horizontal axis shows quarters relative to the reform ($q = 1$ is the first post-reform quarter). Shaded areas are 95% confidence intervals, with standard errors clustered at the municipality level. Regression includes municipality and region-by-time fixed effects. The sample is restricted to municipalities with a balanced panel of 24 monthly observations. 1,999 municipalities; 47,976 observations. *Source:* IEF, AEAT, INE, and Banc Sabadell.

Figure A.3: Event study: Labor market

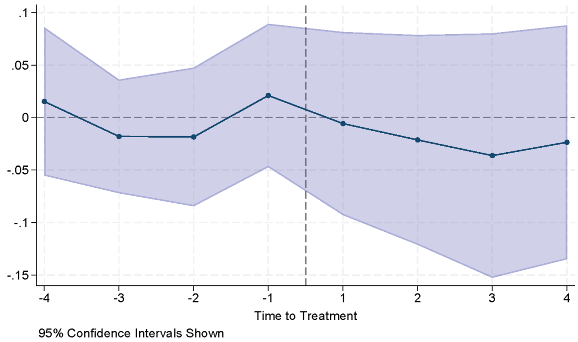

(a) Unemployment rate

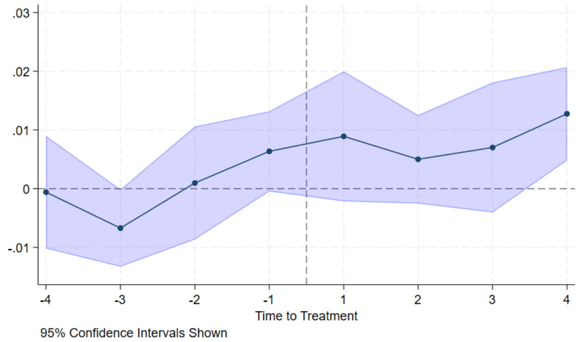

(b) Employment

Notes: Coefficients are estimated from a dynamic DiD specification. The horizontal axis represents quarters relative to the reform ($q = 1$ is the first post-reform quarter). Shaded areas indicate 95% confidence intervals, with standard errors clustered at the municipality level in panel (a) and at the geographic unit level in panel (b). In Figure A.3.a, we estimate a version of equation (2) where the dependent variable is the unemployment rate. The regression includes municipality and region-by-time fixed effects, and the sample is restricted to municipalities with a balanced panel of 24 monthly observations (1,999 municipalities; 47,976 observations). In Figure A.3.b, we estimate a version of equation (4) where the dependent variable is the log of employment, with MCVL geographic units as the unit of analysis. The regression includes geographic-unit and time fixed effects, along with baseline controls interacted with time dummies (221 geographic units; 5,304 observations). The corresponding static DiD estimates are -0.022 (0.019) for panel (a) and 0.008 (0.002) for panel (b), with standard errors in parentheses. *Source:* IEF, AEAT, MCVL, and SEPE.

Figure A.4: Event study: Effects on firms

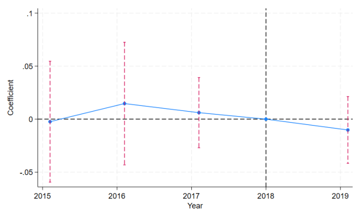

(a) Log value added

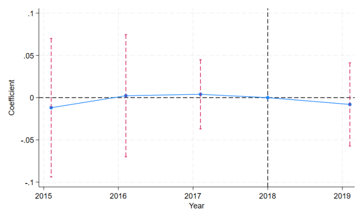

(b) Log profits

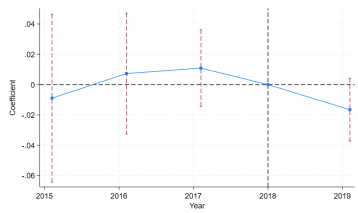

(c) Log number of companies

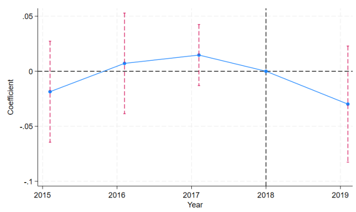

(d) Log number of workers

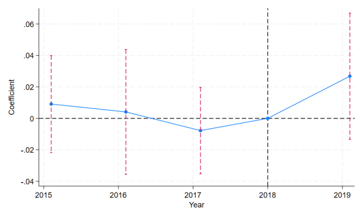

(e) Log wage

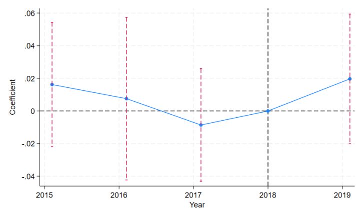

(f) Log productivity pw

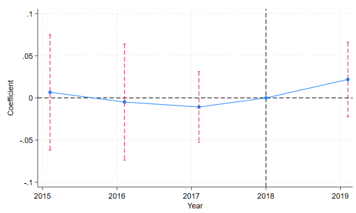

(g) Log profits pw

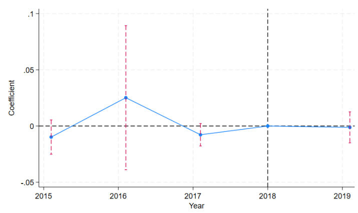

(h) Log markup

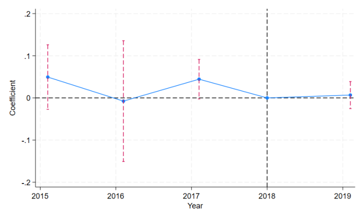

(i) Log markdown

Notes: Coefficients are estimated from a dynamic DiD specification. The horizontal axis represents years relative to the reform, with the dashed vertical line indicating 2018, the last year before the reform. Red-dashed lines indicate 95% confidence intervals, with standard errors clustered at the sector level. We estimate a version of equation (4) where the dependent variables are the logs of value added, profits, number of companies, number of workers, wages, productivity per worker, profits per worker, markups, and markdowns, with sectors as the unit of analysis and year as the unit of time. Regressions include sector and time fixed effects, along with baseline controls interacted with time dummies, and are restricted to sectors with at least 100 workers in the MCVL (70 sectors; 350 observations). *Source:* Central de Balances Integrada (BELab, 2025) and MCVL.

APPENDIX B: Representativeness of Banc Sabadell data and payment instruments

Consumption coverage

We benchmark our transaction data against National Accounts data from the INE. Banc Sabadell's credit cards and POS transactions account for 3.7 percent of total private consumption in 2018 and 4.0 percent in 2019 (Table B.1). Coverage is highly stable both across regions and over time, with an average year-to-year fluctuation of approximately 0.25 percentage points.

Crucially for our identification strategy, Figure B.1 shows that, at the provincial level, changes in coverage are uncorrelated with our measure of exposure to the minimum wage reform.[48] This alleviates concerns that shifts in Banc Sabadell's market presence, such as differential geographic expansion, could confound our estimates of interest.

Figure B.1: Exposure to the 2019 MW increase and changes in Banc Sabadell's coverage

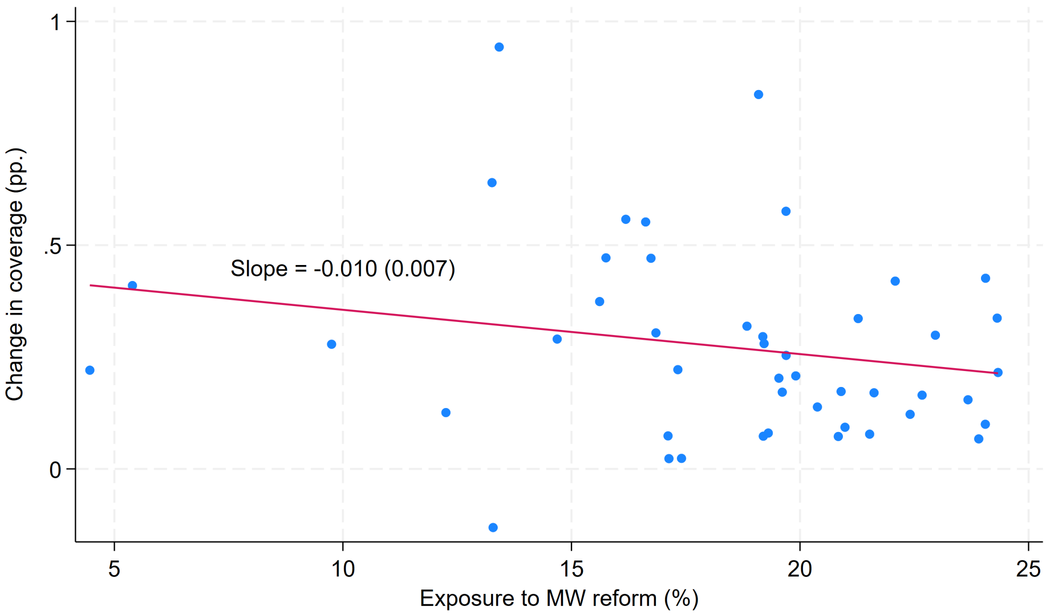

Note: The figure plots the change in Banc Sabadell's coverage of private consumption between 2018 and 2019 (in percentage points) against a measure of provincial exposure to the 2019 minimum wage reform, as defined in Section 4.1. Coverage is measured as the ratio of Banc Sabadell credit card and POS transaction volumes to total private consumption. Private consumption data are taken from the Regional Accounts and imputed at the provincial level using the region-level ratio of private consumption to value added, assumed constant within a region-year. Each dot represents a province, and the fitted line corresponds to a linear regression of changes in coverage on exposure. *Source:* Data from the IEF, AEAT, INE, and Banc Sabadell.

[48]National Accounts data report value added at the provincial level, whereas private consumption is only available at the regional level. To increase the number of observations, we impute provincial private consumption using the region-level ratio of private consumption to value added, assuming this ratio is constant within a region-year. Using private consumption measured directly at the regional level yields similar results, with a regression coefficient (standard error) of 0.005 (0.008).

Table B.1: Regional coverage of consumption in Banc Sabadell transaction data (%)

Region	2018	2019
Andalucía	1.6	1.8
Aragón	4.0	4.5
Asturias, Principado de	5.1	5.4
Balears, Illes	3.6	3.5
Canarias	1.6	1.6
Cantabria	1.9	2.2
Castilla y León	2.1	2.3
Castilla-La Mancha	2.0	2.2
Cataluña	7.1	7.5
Comunitat Valenciana	4.1	4.5
Extremadura	1.5	1.7
Galicia	2.9	3.6
Madrid, Comunidad de	4.5	4.5
Murcia, Región de	3.6	4.2
Navarra, Comunidad Foral de	1.5	1.5
País Vasco	2.7	2.9
Rioja, La	2.0	2.1
Spain (Total)	3.7	4.0

Notes: The table reports the share of total private consumption covered by Banc Sabadell credit card and POS transaction data for each Spanish region in 2018 and 2019. Coverage is computed as the ratio of Banc Sabadell transaction volumes to total private consumption from the Regional Accounts. *Source:* INE and Banc Sabadell.

Geographic coverage of POS devices

Banc Sabadell is the largest POS provider in Spain, with approximately 20 percent of the national market according to the information provided by the Bank. Its POS network is geographicaly diversified and spans, to different extend, all Spanish regions.

Table B.2 compares the regional distribution of Banc Sabadell POS transactions with that of total POS spending.[49] The two distributions align closely, indicating that Banc Sabadell's footprint is both stable and broadly representative of geographic spending patterns.

Table B.2: Regional distribution of total POS spending and Banc Sabadell POS transactions (%)

	POS Spending		POS BS	
	2018	2019	2018	2019
Andalucía	11.9	12.4	7.1	7.2
Aragón	2.2	2.3	2.9	3.0
Asturias, Principado de	1.2	1.4	2.9	2.9
Balears, Illes	4.4	4.9	3.7	3.3
Canarias	5.8	5.4	2.2	2.0
Cantabria	1.0	1.0	0.7	0.7
Castilla y León	3.7	3.4	2.7	2.7
Castilla-La Mancha	2.6	2.3	1.9	1.9
Cataluña	23.9	24.4	33.8	33.2
Comunitat Valenciana	9.8	9.9	11.8	12.2
Extremadura	1.1	1.1	0.7	0.7
Galicia	3.9	3.4	4.0	4.6
Madrid, Comunidad de	19.9	20.2	18.7	18.4
Murcia, Región de	1.9	2.1	2.6	2.8
Navarra, Comunidad Foral de	2.3	2.0	0.6	0.6
País Vasco	3.8	3.1	3.6	3.5
Rioja, La	0.6	0.5	0.3	0.3
Spain (Total)	99.8	99.8	100.0	100.0

Note: The table compares the regional distribution of Banc Sabadell POS transactions with that of total POS spending in Spain for 2018 and 2019. Each entry reports the share of total POS transaction volume accounted for by each region. The Spanish total for overall POS spending does not sum exactly to 100 due to the exclusion of Ceuta and Melilla. *Source:* Sistemapay and Banc Sabadell.

[49] Data on the regional distribution of total POS spending are obtained from Sistemapay, *Informe Anual 2018* and *Informe Anual 2019*. Available at: https://www.sistemapay.com/wp-content/uploads/2023/12/Informe-Anual-2018-Sistema-de-Tarjetas-y-Medios-de-Pago_03072019.pdf and https://www.sistemapay.com/wp-content/uploads/2023/12/Informe-Anual-2019-Sistema-de-Tarjetas-y-Medios-de-Pago.pdf . Last accessed January 8, 2026.

Payment instruments

Information on payment instruments in Spain is scarce. To provide context on their use across income groups, we rely on the 2019 ECB Study on the Payment Attitudes of Consumers in the Euro Area (SPACE), which documents the payment habits of consumers across euro area countries. The survey was first conducted in 2019 and has since been followed by two additional waves (2022 and 2024), with the next wave scheduled for publication in 2026.[50]

Table B.3: Payment instruments used in physical stores by level of income, 2019.

Income group	Cash (%)	Card (%)	Mobile (%)	Other (%)	Pop. share (%)	Cons. share (%)
<=€750	66.3	22.7	9.0	2.0	9.2	8.6
€751–€1500	66.2	30.6	2.2	1.1	29.7	27.2
€1501–€2500	67.5	28.1	0.5	3.9	30.0	30.6
€2501–€4000	63.6	33.4	1.6	1.3	22.1	23.7
>€4000	66.5	28.2	4.3	1.1	9.0	9.9

Notes: Shares computed as the ratio of weighted spending by instrument to total weighted POS spending within each income group. Population and consumption shares are computed over the sample of respondents with at least one valid POS payment, excluding individuals not reporting income. *Source:* ECB SPACE (2019).

Table B.3 reports population-weighted results for Spain. Consumption is concentrated in the middle of the income distribution; households earning between €751 and €2,500 account for 57.8 percent of POS spending (27.2 percent in the €751–€1,500 bracket and 30.6 percent in the €1,501–€2,500 bracket), and those earning €2,501–€4,000 for an additional 23.7 percent. Households at the top of the distribution (above €4,000) and at the bottom (below €750) contribute 9.9 and 8.6 percent, respectively. Consumption shares track population shares closely within each bracket, suggesting that while consumption rises with income, the variation across groups is relatively modest.[51]

Turning to payment instruments, cash was the dominant method in 2019, accounting for 66.0 percent of total POS transaction value, while card payments represented 29.6 percent. As shown in Table B.3, cash usage is remarkably stable across income groups, ranging from 63.6 to 67.5 percent, with no clear income gradient. Card usage is somewhat lower among the lowest-income households (below €750), at 22.7 percent, compared to a range of 28.1–33.4 percent for all other groups, again without a systematic pattern.[52] Other payment methods remain marginal across most groups. The main exception is

[50] A previous version of this paper used information from the same source directly provided by Ferrando and Posada (2023), from the Bank of Spain. The results are not directly comparable because they used different income brackets, likely due to access to more detailed microdata than those publicly available, and because they do not explicitly state whether their results are survey-weighted. Nonetheless, the findings are qualitatively robust. The main reason for the update is the inclusion of population shares, which were not available in Ferrando and Posada (2023).

[51] This applies to in-person POS spending; higher-income households may shift a larger share of consumption to other channels.

[52] Unweighted results yield a qualitatively different picture for the lowest-income group, with a cash share of only 52.7 percent and a card share of 32.5 percent, in line with the rest of the distribution. These figures are closer to those reported by Ferrando and Posada (2023).

mobile payments among the lowest-income group, which account for 9.0 percent of their POS spending, substantially higher than for any other income bracket.[53]

[53] Mobile payments include banking apps, digital wallets, PayPal, Apple Pay, and similar solutions.

APPENDIX C: Robustness of the local consumption estimates

This appendix provides the detailed results of the robustness exercises summarized in Section 4.3.1. The first part assesses the sensitivity of the baseline estimate to alternative specifications, exposure measures, and sample definitions; the second part addresses geographic spillovers.

Alternative specifications

Table C.1 reports the baseline static DiD estimate in column (1)—a 4.5 percent relative increase in consumption in more exposed municipalities—together with a battery of robustness checks. Column (2) adds predetermined controls (the 2018 logarithm of population, the share of working-age population, and the share of women) interacted with time fixed effects, leaving the coefficient essentially unchanged. Column (3) replaces the binary treatment with the original continuous exposure measure; scaling the coefficient by the average treatment intensity (0.165) yields an implied ATT of 5.1 percent, closely aligned with the baseline.

Columns (4) and (5) use alternative exposure measures. Column (4) defines exposure as the share of individuals with annual income below €12,600 irrespective of age. While this introduces noise by including groups such as pensioners who are not directly affected by the reform, it draws on a larger sample of tax declarations; the estimated effect is smaller (2.6 percent) but remains highly significant. Column (5) classifies treatment relative to the median exposure within each region, ensuring balanced treatment and control groups within regions; the coefficient is slightly smaller, likely reflecting comparisons between more similar municipalities, but again highly significant.

Columns (6) and (7) explore the potential influence of tourism on our results. Column (6) excludes the peak tourist months of July and August, showing no significant change in the coefficient of interest. Column (7) relaxes the threshold for excluding municipalities with high seasonal consumption, increasing it from 50 to 100 percent. Again, we observe no significant changes in the coefficient, suggesting that our results are not driven by seasonal tourism effects.

Finally, columns (8) to (10) examine the robustness of our findings to alternative sample definitions and exposure thresholds. Column (8) includes all municipalities with available exposure measures, without discarding outliers, and applies the inverse hyperbolic sine transformation to account for zeros in consumption. The estimated coefficient remains positive and significant but is smaller (1.5 percent), likely reflecting additional measurement noise and potential spillovers, particularly among smaller municipalities. Columns (9) and (10) redefine the treatment and control groups to test the robustness of the baseline median-split specification. Column (9) compares municipalities in the top quartile of exposure with those below the median, yielding an effect of 4.3 percent, consistent with the baseline. Column (10) contrasts the bottom and top quartiles, producing a larger coefficient of 8.0 percent, consistent with stronger treatment effects when comparing municipalities at the extremes of the exposure distribution.

Table C.1: Baseline and Robustness Results for Consumption

	Baseline (1)	Controls (2)	Cont. (3)	Alt. exp. (4)	Reg. median (5)
Treat x Post 2019	0.045	0.046	0.307	0.026	0.022
	(0.007)***	(0.007)***	(0.058)***	(0.007)***	(0.006)***
Mean outcome	2.62	2.62	2.62	2.62	2.62
Observations	47976	47976	47976	47976	47976
Municipalities	1999	1999	1999	1999	1999

	No summer (6)	Sum. ratio > 2 (7)	Full sample (8)	$P50-P75$ (9)	$P25-P75$ (10)
Treat x Post 2019	0.043	0.043	0.015	0.043	0.080
	(0.008)***	(0.006)***	(0.005)***	(0.011)***	(0.015)***
Mean outcome	2.61	2.53	2.01	2.67	2.65
Observations	39980	62592	132480	35976	24000
Municipalities	1999	2608	5520	1499	1000

Notes: Standard errors clustered at the municipality level in parentheses. The dependent variable is the logarithm of real per capita consumption. The "Baseline" column represents the benchmark estimation from the main analysis. Column (2) includes baseline controls (i.e., the 2018 values of the logarithm of population, the share of working-age population, and the share of women) interacted with time fixed effects. Column (3) uses a continuous measure of exposure. The "Alt. exp." column defines exposure as the share of individuals with an annual income below €12,600. Column (5) separates municipalities into high- and low-exposure groups within each region. Column (6) omits data from July and August. Column (7) excludes municipalities where summer consumption is, on average, at least twice the consumption in other months. Column (8) includes all municipalities for which we have an exposure measure, without discarding outliers, and applies the inverse hyperbolic sine transformation to allow for zeros. Column (9) defines the control group as municipalities below the median of the exposure distribution, while the treatment group includes only those in the top quartile. Column (10) uses the bottom quartile of exposure as the control group and the top quartile as the treatment. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. *Source:* IEF, AEAT, INE, and Banc Sabadell.

Geographic spillovers

Working with municipal-level data introduces the possibility of spatial interdependence in consumption patterns, since economic activity in one municipality may spill over into neighboring areas, potentially violating the SUTVA. To address these potential spatial spillovers, we conduct several robustness checks that progressively account for geographic relationships.

Table C.2 summarizes the results. We begin by restricting the sample to exclude smaller municipalities. Column (1) omits municipalities with fewer than 100 income tax declarations, which also helps reduce measurement noise in the exposure variable. Column (2) restricts the sample to municipalities with more than 1,000 inhabitants in 2018. The estimated effects—5.3 percent and 4.3 percent, respectively—remain close to the baseline result.

Column (3) redefines the geographic unit as large urban areas, as classified by the INE, reducing the sample to 73 cross-sectional units. This specification implicitly accounts for stronger economic integration within metropolitan regions and mitigates potential spillover concerns by aggregating at a broader spatial scale. Despite the smaller sample size, the estimated impact of the 2019 minimum wage increase remains positive and consistent with the baseline (3.6 percent).

In columns (4) through (7), we explicitly account for spatial dependence using spatial econometric models.[54] Column (4) reports a Spatial Error Model (SEM) in which

[54] Estimating a spatial panel model with high-dimensional fixed effects is computationally intensive and

Table C.2: Consumption Responses and Robustness to Geographic Spillovers

	$N > 100$	$\text{Pop} > 1,000$	Urban Areas	SEM	Contiguity	50km	200km
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treat x Post 2019	0.053	0.043	0.036	0.045	0.045	0.045	0.045
	(0.007)***	(0.007)***	(0.009)***	(0.015)***	(0.017)***	(0.017)***	(0.017)***
Spatial spillover					–0.005	–0.002	–0.004
					(0.014)	(0.014)	(0.013)
Mean outcome	2.76	2.67	3.60	2.62	2.62	2.62	2.62
Observations	31920	38136	1752	1999	1999	1999	1999
Municipalities	1330	1589	73	1999	1999	1999	1999

Notes: Standard errors in parentheses. The dependent variable is the logarithm of real per capita consumption. Columns (1)–(3) report standard errors clustered at the municipality level. Columns (4)–(7) report spatial econometric estimates using Generalized Spatial Two-Stage Least Squares (GS2SLS) with heteroskedasticity-robust standard errors. Columns (1) and (2) restrict the sample to municipalities with more than 100 income tax declarations and populations exceeding 1,000, respectively. Column (3) uses large urban areas as the geographic unit. Column (4) presents a Spatial Error Model based on a row-standardized contiguity matrix. Columns (5)–(7) incorporate a spatial spillover variable defined as the logarithm of the spatially weighted neighboring exposed population relative to a municipality's own working-age population. The spatial weight matrices are: (5) a row-standardized contiguity matrix, (6) an inverse-distance matrix truncated at 50 km, and (7) an inverse-distance matrix truncated at 200 km. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. *Source:* IEF, AEAT, INE, and Banc Sabadell.

Table C.3: Robustness: Spatially correlated standard errors and province clustering

	Cutoff 20 km	Cutoff 50 km	Cutoff 100 km	Cutoff 200 km	Province cluster
	(1)	(2)	(3)	(4)	(5)
Treat x Post 2019	0.045	0.045	0.045	0.045	0.045
	(0.007)***	(0.008)***	(0.009)***	(0.009)***	(0.019)**
Mean outcome	2.62	2.62	2.62	2.62	2.62
Observations	47976	47976	47976	47976	47976
Municipalities	1999	1999	1999	1999	1999

Notes: Columns (1)–(4) report standard errors adjusted for spatial correlation following Colella et al. (2019). Spatial correlation is assumed to decline linearly with distance and vanish beyond the specified cutoff. Results are shown for four distance thresholds: 20 km, 50 km, 100 km, and 200 km. Column (5) reports standard errors clustered at the province level. All regressions include municipality and region-time fixed effects. * p<0.1, ** p<0.05, *** p<0.01. *Source:* IEF, AEAT, INE, and Banc Sabadell.

spatial correlation in the error term is modeled using a row-standardized contiguity matrix. Columns (5) to (7) introduce a spatial weighted neighboring exposure variable to test for demand-side spillovers—that is, whether local consumption is influenced by the reform's exposure in neighboring municipalities. The spillover measure is defined as the logarithm of the spatially weighted neighboring exposed population relative to a municipality's own working-age population. The weights are determined by the corresponding spatial matrix, which is row-standardized so that weights sum to one across neighbors. This construction captures the idea that potential spillovers depend on the scale of reform exposure in surrounding areas relative to the size of the local population. We evaluate the robustness of this channel using three spatial weighting matrices: a contiguity matrix and inverse-distance matrices truncated at 50 km and 200 km.

Across all specifications, the coefficient on the neighboring exposure variable is small and statistically insignificant, indicating limited evidence of demand spillovers across municipal boundaries. At the same time, the coefficient on direct municipal exposure remains positive and highly significant in all spatial specifications, suggesting that the observed increase in consumption primarily reflects local effects rather than cross-municipal propagation.[55]

Finally, spatial correlation across municipalities could understate the standard errors. Table C.3 re-estimates the baseline with spatially corrected standard errors following Colella et al. (2019) under several distance thresholds, and with standard errors clustered at the province level; the coefficient remains statistically significant throughout.

often infeasible with large cross-sectional units. To overcome this challenge, we collapse the consumption data into pre- and post-reform averages (2018–2019) and estimate the model in first differences. This transformation eliminates municipality fixed effects while preserving the spatial structure of the data, allowing for the identification of spatial error and spillover components in a tractable way.

[55] The results are virtually unchanged if we instead use the total neighboring exposed population relative to own working-age population without row-standardization, or if we replace the logarithmic specification with a linear neighboring exposure measure. In all cases, the spillover coefficient remains small and statistically insignificant, while the coefficient on direct municipal exposure is essentially unaffected. These results are available upon request.

APPENDIX D: Construction of the MCVL-based exposure measure

This appendix describes how we construct the Social Security–based measure of local exposure to the 2019 minimum wage increase used in Section 4.3.2. The measure relies on pre-reform (2018) hourly wages computed from the MCVL.

Sample. We retain employees aged 18–65 affiliated with the General Regime of the Social Security throughout 2018, excluding workers who were self-employed or affiliated with any special regime at any point in the year, since contribution rules differ across regimes and complicate the construction of comparable hourly wages. We further keep only standard employment spells, discarding those corresponding to unemployment benefits or training contracts.

Hourly wages. For each worker–month we observe the Social Security contribution base, CB_{im}, and impute hours worked, H_{im}, assuming a full-time month of 160 hours scaled by the share of days worked in the month and the worker's part-time coefficient (which we allow to change within a spell). The monthly hourly wage is

$$w_{im} = \frac{CB_{im}}{H_{im}}.$$

We compute wages at the monthly level, uncensor them as described below, and only then aggregate to an annual hourly wage as the hours-weighted average of uncensored monthly wages.

Censoring. Contribution bases are bounded by statutory monthly floors and ceilings that vary by year and contribution group, so observed wages are censored at both tails. We convert each monthly threshold to an hourly value using the 160-hour full-time benchmark and flag a worker–month as bottom-censored when its hourly wage sits at the floor and as top-censored when it sits close to the ceiling. About 1.8 percent of worker–month observations are bottom-censored and 10.1 percent top-censored.

Following Bonhomme and Hospido (2017), we recover the censored wages with cell-specific Tobit models. Cells c are defined by age group, calendar month, and broad contribution category, and within each cell we estimate

$$\log(w_{im}) = \mu_c + \varepsilon_{im}, \qquad \varepsilon_{im} \sim N(0, \sigma_c^2),$$

by censored maximum likelihood. Uncensored observations keep their observed wage. For censored observations we draw the latent wage from the estimated within-cell distribution, truncated to the censored region:

$$\log(w_{im}) \sim N(\mu_c, \sigma_c^2) \,\big|\, \log(w_{im}) \leq \log(\underline{w}_{im})$$

for bottom-censored observations, where $\underline{w}_{im}$ is the hourly floor, and

$$\log(w_{im}) \sim N(\mu_c, \sigma_c^2) \,\big|\, \log(w_{im}) \geq \log(\overline{w}_{im})$$

for top-censored observations, where $\overline{w}_{im}$ is the hourly ceiling. For each censored worker–month we take ten draws from the relevant truncated distribution and replace the censored wage with their average. We then aggregate the uncensored monthly wages and hours to annual hourly wages.

Geographic units. The MCVL reports municipality of residence only for individuals living in municipalities with more than 40,000 inhabitants; smaller localities are identified only at the province level. We therefore compute exposure for the 171 municipalities above the threshold and assign workers in smaller localities to their province, yielding 221 geographic units. As a further robustness check, we aggregate all workers to the 50 Spanish provinces.

Exposure measure. For each geographic unit g we define exposure as the wage-bill shortfall relative to the new minimum:

$$\mathrm{GAP}_g = \frac{\sum_{i \in g} h_i \max\{0,\ \mathrm{MW} - w_i\}}{\sum_{i \in g} h_i \, w_i},$$

where h_i and w_i are worker i's 2018 annual hours and hourly wage and MW is the 2019 statutory minimum. The numerator is the additional pay needed to bring every worker in g up to the minimum; dividing by the unit's wage bill and multiplying by 100 gives the percentage increase in the aggregate wage bill required to comply with the reform. Areas with more workers below the minimum, and workers further below it, receive higher values.

APPENDIX E: Conceptual framework

Our empirical findings show that the consumption response was concentrated in discretionary goods. Here, we show that this pattern can be rationalized by nonhomothetic preferences. A natural framework for this purpose is the Stone–Geary utility function, originally proposed by Geary (1950) and Stone (1954). This utility function captures the idea that utility depends on consumption in excess of subsistence levels.[56] With Stone-Geary utility function, a representative MW consumer has preferences over N goods:

$$U(C_1, \ldots, C_N) = \prod_{i=1}^{N} (C_i - \bar{C}_i)^{\alpha_i},$$

where C_i denotes consumption of good i, $\bar{C}_i$ is the subsistence level of good i, and $\alpha_i > 0$ reflects preference weights that satifiy the Cobb-Douglas property $\sum_{i=1}^{N} \alpha_i = 1$. If prices p_i, consumption levels C_i, and disposable income MW satisfy the budget constraint $\sum_{i=1}^{N} p_i C_i = \text{MW}$, one can easily derive the well-known demand function in this setup:

$$C_i = \bar{C}_i + \frac{\alpha_i}{p_i} \left(\text{MW} - \sum_{j=1}^{N} p_j \bar{C}_j \right),$$

This demand function shows that the consumption level of good i, C_i, consists of two parts. The first part, $\bar{C}_i$, is the minimum or subsistence amount of good i. The second part depends on the consumer's income available after covering the total subsistence costs of all goods: $\text{MW} - \sum_{j=1}^{N} p_j \bar{C}_j$. Therefore, additional income is first allocated toward meeting the subsistence level of each good, where subsistence requirements may differ across goods. The remaining income is then allocated across goods according to the preference weight α_i, adjusted for the price p_i of good i.

For simplicity, we can group all goods into two categories: B, which includes non-discretionary goods, and D, which includes discretionary goods. Then, the demand for each category is:

$$C_B = \bar{C}_B + \frac{\alpha_B}{p_B} \left(\text{MW} - p_B \bar{C}_B - p_D \bar{C}_D \right)$$
$$C_D = \bar{C}_D + \frac{\alpha_D}{p_D} \left(\text{MW} - p_B \bar{C}_B - p_D \bar{C}_D \right)$$

Therefore, as disposable income MW increases, consumption beyond subsistence rises in proportion to $\frac{\alpha_i}{p_i}$.

Our empirical results are consistent with parameter values in which non-discretionary goods have relatively high subsistence levels, $\bar{C}_B$, and relatively low preference weights, α_B, for consumption beyond subsistence. In contrast, discretionary goods have low subsistence levels $\bar{C}_D$, but relatively high preference weights α_D. As a result, the additional income generated by the minimum wage increase translates into a relatively larger increase

[56] For simplicity, we use a static framework to illustrate how the composition of the consumption bundle adjusts in response to a minimum wage increase. We abstract from changes in total consumption, which are straightforward to characterize under the assumption of a permanent income increase, the scenario we study.

in discretionary consumption.

This framework provides a simple interpretation for why the increase in consumption following a minimum wage increase is concentrated in discretionary goods and services, such as leisure activities, electronics, and food away from home.