Demand-Driven Stagnation: Customer Acquisition and Firm Growth *

Thomas Lazarowicz†

This version: November 2, 2026

Click here for the latest version

Abstract

What are the effects of a fiscal consolidation on firm growth, and do they depend on its distributional incidence? I study the UK's 2010–2015 austerity programme, whose welfare cuts fell mostly on low-income households. Young firms in more exposed local authorities cut customer acquisition by up to 15% per standard deviation relative to less exposed peers, while their employment, capacity utilisation and markups were unaffected. Young firms' value added fell and remained depressed through 2015, even after spending on customer acquisition recovered, suggesting a durable loss of customer base. Old firms showed no comparable adjustment. To explain why regressive shocks disproportionately scar young firms, I build a general equilibrium model in which firms accumulate customers from a heterogeneous-income population. Low-income customers churn faster, so young firms inherit a poor-biased customer base. This creates a "churn trap": young firms depend most on the households a regressive consolidation hits hardest. The effects are especially strong for firms that operate in 'luxury' sectors, whose demand is the most income-elastic. In the calibrated model, the 2010 austerity programme destroys about half as much again of young firms' customer capital as a progressive consolidation of the same size, and costs about nine pence more output per pound withdrawn than a progressive design, £11bn over the decade; the cohort entering during the shock remains persistently smaller than in the absence of the consolidation. A flat design would have avoided about a fifth of young firms' loss of customer capital, and a progressive design about three tenths.

Keywords: Customer capital, firm dynamics, non-homothetic preferences, fiscal austerity

JEL Codes: E23, E24, E62, L11, D12

*I am indebted to Vincent Sterk, Wei Cui, and Dajana Xhani for their guidance and support throughout this project. I am grateful to Klaus Adam, Paul Bouscasse, Wendy Carlin, Vasco Carvalho, Adrien Couturier, Ararat Gocmen, Jeremias Klaeui, Marta Morazzoni, Lukas Nord, Alan Olivi, Paula Patzelt, Franck Portier, Andrew Preston, Lukasz Rachel, Morten Ravn, Morgane Richard and seminar participants at UCL and various institutions for helpful comments and suggestions. I thank the UK Data Service for guidance and assistance with regard to the data. The UK Data Service agrees that the attached outputs are non-disclosive, and cannot be used to identify a person or organisation. The use of these data does not imply the endorsement of the data owner or the UK Data Service at the UK Data Archive in relation to the interpretation or analysis of the data. This work uses research datasets which may not exactly reproduce National Statistics aggregates. All remaining errors are my own.

†University College London, thomas.lazarowicz.16@ucl.ac.uk

1 Introduction

What are the channels and economic consequences of a contractionary demand shock? A large literature has focused on how differences in marginal propensities to consume can amplify the effects of shocks that fall disproportionately on low-income households (Auclert, 2019; Patterson, 2023). But the effects of demand shocks on firms are less well understood, even though recent work shows that demand conditions at the point of firm entry can have lasting effects on firm growth (Sedláček & Sterk, 2017). This paper shows that beyond aggregate demand conditions, the composition of demand shocks matters for firm growth by affecting firms' incentive to invest in acquiring new customers. I leverage plausibly exogenous variation in exposure to welfare cuts across UK local authorities to estimate the margins of adjustment for firms. I then build and quantify a model of heterogeneous firms and households that rationalises these empirical findings and can be used to evaluate shocks of equal size but different incidence.

The UK welfare reforms of 2010 provide a natural setting to study this question. In 2010 the United Kingdom cut welfare spending by £19 billion per year, or about 1% of GDP. These spending cuts fell disproportionately on people at the bottom of the income distribution. At the same time, and like many other advanced economies, the UK grew more slowly than before the 2008 financial crisis, with GDP per capita falling away from its pre-crisis trend path and never returning to it. The slowdown was not evenly spread: among young firms, usually the fastest-growing group and the main driver of job creation (Decker et al., 2017), output per employee stalled after 2010 and then declined, while it kept rising for older firms. The paper shows, both through novel empirical evidence and a calibrated model, that a customer-acquisition channel coupled with a regressive fiscal shock is a powerful mechanism for generating a slowdown in firm growth.

This paper makes three contributions. First, I combine administrative data covering the near-universe of UK firms with a shift-share measure of local exposure to the welfare cuts, and show that firms responded most strongly on a single margin: young firms cut investment in customer acquisition, while other margins of adjustment remained unchanged. Second, I build a general equilibrium model of firm dynamics in which heterogeneous firms grow by acquiring customers that differ in income and in the rates at which they end relationships. Third, to identify the contribution of the heterogeneous impact of this policy across the income distribution, I use the model to hold the total size of the withdrawal fixed, and vary where in the distribution the withdrawal was concentrated. I find that the actual distribution of the welfare cuts had a larger negative effect on customer capital and aggregate output than a consolidation of equal size that was concentrated on higher income households would have.

I construct plausibly exogenous fiscal shocks at the local-authority level from spatial variation in exposure to national welfare cuts. From 2010, the UK Government introduced cuts to welfare benefits that applied uniformly across the country. The share of the population claiming these benefits varies widely across the UK, so this national reform produced different losses in benefit income across different local authorities. I combine the national saving from each reform, as costed by Beatty and Fothergill (2014), with each authority's claimant shares for that benefit in 2009–10, which gives a measure of exposure that I can also follow over

time; Fetzer (2019) uses their published cross-section of losses by authority. Identification in this design follows the shift-share framework of Goldsmith-Pinkham et al. (2020), which requires the claimant shares to be exogenous. The shares were fixed before the cuts were announced, and I show that they are persistent. I additionally control for exposure to the financial crisis by including the change in each local authority's average claimant rate between 2005 and 2009, interacted with the young-firm indicator by year, so that young and old firms in districts hit harder by the Great Recession can follow different paths. I then use a triple-difference design, which compares the change after 2010 in the gap between young and old firms across local authorities whose residents lost very different amounts of benefit income, within the same sector and region. The identifying assumption is that, without the cuts, the gap between young and old firms would have evolved in the same way in more and less exposed local authorities, and before 2010 it shows no differential trend in the variables of interest.

I construct a unique dataset that combines spatial exposure to these cuts with employment, payroll, accounting and balance sheet data at the establishment and firm levels. The resulting dataset has a very good coverage of the entire firm distribution, capturing young and small firms that are absent or undersampled in more standard datasets of listed firms, making it appropriate for studying questions that relate to firm growth and life-cycle.

I find that, in response to austerity, young firms cut their spending on demand investment, leading to a persistent decline in value added. Using the quasi-experimental variation in local exposure to spending cuts, I show that young firms in more exposed local authorities cut their spending on demand investment by 15% more per standard deviation of exposure than their less exposed counterparts. A number of alternative adjustment margins, including employment, capacity utilisation (proxied by energy usage), markups, and TFP, show no response. Value added of the same firms fell and remained lower through 2015, more persistently than the decline in acquisition itself. I take this as evidence of a persistent and durable loss of customer capital that a one-off demand shock would not reproduce. In contrast, older firms in the same local authorities did not respond. A second important finding is that the response of young firms is stronger in luxury (high income elasticity) sectors, with young luxury firms displaying a negative and significant response, and young necessity firms showing no comparable decline.[1] Together, these findings suggest that austerity reached firms through the demand for their products, and that how exposed a firm was depended upon its position in the life cycle and the sector it operated in.

While the empirical evidence points to the relevant margins of adjustment, it cannot by itself explain why young firms and luxury sectors were the most exposed, and, because it comes from a single consolidation, it cannot show how the same firms would have fared had the cuts fallen on different households. To answer both questions, I build a general-equilibrium model of firm dynamics with heterogeneous firms and heterogeneous customers. This model is consistent with and quantitatively disciplined by the empirical findings. Firms differ in productivity, in the number of customers they have, and in the composition of those

[1] This pattern is consistent with household evidence from the Living Costs and Food Survey, where benefit recipients cut their spending in luxury categories relative to non-recipients, while differences in necessity categories were negligible (Section 3.6).

customers across income groups. Firms acquire customers by investing in marketing spending, and lose existing customers at rates that differ by household income. When demand falls, firms cut this investment as the value of a new customer falls, which makes the fall in their customer base persistent, most of all for young firms.

The model mechanism operates through two facts of household demand that are key to matching the empirical response. First, lower-income households search harder for suppliers (Kaplan & Menzio, 2016; Nord, 2023) and end more of their existing relationships, which I document in UK household panel data (Section 4.3), so the pool of searching households in any given period is biased towards low-income households. Second, households in the model have non-homothetic preferences over sectors, so income elasticities differ across sectors and across households. Non-homothetic preferences concentrate the loss on young firms in luxury sectors, consistent with the empirical evidence, and make young firms lose more than older ones even when every household's spending is cut in the same proportion. They also magnify a regressive cut, since the luxury elasticity is above one and highest for the poorest households. To first order, a firm's exposure is the revenue weighted average of its customers' income elasticities. Differential sensitivity between young and older firms is then a covariance between the households they sell to and how exposed those households are to income fluctuations. Intuitively, a young firm looking to grow is most exposed to this pool of searching customers, and household preferences mean that these same customers are the ones cutting back their luxury sector spending the most in the case of a demand shock. A consolidation that falls on lower income households, like the 2010 UK programme that is the focus of this paper, consequently hits young firms the hardest, and disproportionately in luxury sectors.

The household side of the model is disciplined by evidence from the Living Costs and Food Survey. The model imposes three restrictions on the rate at which households end their relationships with suppliers, namely that it falls with income, that it rises when a household's spending per supplier falls, and that it does not depend on the prices suppliers post. Appendix D shows that such restrictions can be microfounded in a household problem close to Bornstein (2025), in which finding a supplier takes time and a relationship ends when a disruption specific to that supplier costs the household more than the relationship is worth. On the firm side, the model is calibrated to the life-cycle profile of marketing intensity in the Annual Business Survey, which it reproduces along with average firm size and the aggregate luxury expenditure share. The model also reproduces the untargeted fall in firm entry and in the number of young firms in more exposed local authorities in the years immediately after austerity.

With the calibrated model in hand, I return to the question posed at the outset, whether the design of the consolidation, that is, which households bore the cuts, played a role in shaping firm growth. I implement a counterfactual experiment to isolate the contribution of the incidence of the austerity programme. I hold spending cuts fixed, and vary the incidence across the actual 2010 design, a flat design that takes the same proportion of income from every household, and a progressive design that hits households higher up the income distribution. The main finding of the model is that the design matters for firm growth even when the size of the consolidation is held fixed. Under the actual 2010 distribution of welfare cuts, young firms

lose about 1.5 times as much customer capital as under a progressive design withdrawing the same spending. For every pound withdrawn, the 2010 design also costs about 9 pence more output than the progressive design, or £11bn over 2010–19. The cohort of firms that enters during the consolidation remains 13 per cent below its no-shock path in 2030. A progressive redesign raising identical revenue, with the losses placed higher up the income distribution, would have avoided about three tenths of young firms' damage. A progressive design moves part of the revenue loss from young to old firms, but costs less output in aggregate.

The rest of the paper is organised as follows. Section 2 reviews the related literature. Section 3 presents the data and empirical evidence, including the event study results and null results for alternative margins. Section 4 develops the model. Section 5 covers equilibrium and calibration. Section 6 quantifies how the incidence of the consolidation shapes young-firm scarring. Section 7 concludes.

2 Literature

This paper contributes to a literature that establishes that there can be long-lasting effects of demand conditions on the productive capacity of the economy. Benigno and Fornaro (2018) highlight one mechanism, where weak demand leads to weak innovation and therefore long-lasting stagnation, while Anzoategui et al. (2019) show that a large portion of the post-2008 slowdown was an endogenous response of technology adoption and R&D to weak demand conditions. Bai et al. (2026) show that demand shocks can appear as productivity shocks even with technology unchanged, when output requires matched customers and lower search by households leaves capacity unsold. In this paper, a firm's customer base is directly accumulated by firms through costly acquisition, rather than through household search. Fiscal austerity that falls on low income households lowers the value of investing in customer acquisition and consequently lowers output for years after the cuts end. The damage is transmitted to the supply side of the economy through firms' customer bases rather than through technology or R&D, and its size and persistence depend on which households bear the cuts as well as on its aggregate size. Here, the persistent cost is the customers that young firms did not acquire. A firm that underinvests during a downturn remains smaller than it would otherwise have been, and remains so long after the cuts end.

Another branch of literature focuses on the distributional effects of fiscal policy and how these affect demand. Auclert (2019) documents how redistribution of policy shocks affects aggregate spending through heterogeneous household spending responses, while Patterson (2023) finds that the differential impact of recessionary shocks across workers magnifies the overall shock effect. Here, who bears a shock matters to the extent that households differ in their marginal propensities to consume, so that a withdrawal from high-MPC households has a very different effect from one concentrated in low-MPC households. I offer a different perspective, in which a withdrawal from low income households has additional effects through their interactions with firms. Low income households separate from existing suppliers more often, and so make up a disproportionately large share of the customers a firm that wants to grow can reach. Consequently, a consolidation that falls on these households removes demand from the firms that are least able to replace it. The cost then falls on the customer capital

of young firms and outlasts the period of the consolidation, which would not be delivered by differences in MPCs alone. A fall in local spending that reached all firms in the same manner would be expected to show up in old firm and young firm behaviour, which I show is not the case in Section 3.

The paper also contributes to a literature that examines the cost of the fiscal consolidations that followed the financial crisis. Fatás and Summers (2018) find strong hysteresis effects of fiscal policy on output, and House et al. (2020) attribute three quarters of the variation in GDP across advanced economies in 2010–14 to differences in austerity. Both studies measure a consolidation through the total amount withdrawn. I add a firm-level mechanism to these analyses and show that for a given size, the loss depends upon which households the cuts fell on. The empirical setting exploits the design of the UK consolidation, which fell disproportionately on low income households (Beatty & Fothergill, 2014; De Agostini et al., 2018).

The paper is closely related to work that examines young firm scarring over the business cycle. Sedláček and Sterk (2017) document that startups born in recessions remain persistently smaller than they otherwise would have, while Sedláček (2020) shows that a decline in entry during economic slowdowns leaves a scarring effect on the aggregate economy through a "missing generation" of high growth young firms contributing to aggregate employment growth. While the mechanisms in these papers operate through entry and frictions in the labour market, in my paper the mechanism operates through the demand investment decisions of firms that do enter. Similarly, Moreira (2016) finds that cohorts born in slowdowns enter smaller and remain smaller, with persistence arising from selection at entry and slow demand accumulation. In my setting, the size of a given cohort's deficit depends upon which households the downturn falls upon.

The mechanism builds on the literature that treats a firm's customer base as a slow-moving intangible asset accumulated through costly investment (Ravn et al., 2006; Arkolakis, 2010; Drozd & Nosal, 2012; Gourio & Rudanko, 2014; Ignaszak & Sedláček, 2026). A large literature provides evidence that supports the idea that customer acquisition is a key margin through which firms grow. Hottman et al. (2016) attribute the bulk of firm-size dispersion to heterogeneity in demand appeal rather than through costs; Einav et al. (2021) and Argente et al. (2025) show that entrants grow by adding customers rather than by selling more to existing ones, and Foster et al. (2016) and Fitzgerald et al. (2024) provide evidence for the same quantity driven accumulation of demand in manufacturing and export markets respectively. Fitzgerald et al. (2024) further show that customer bases respond to costly marketing rather than to dynamic pricing, which I also find evidence for in Section 3, and which motivates the model of Section 4.

A complementary strand studies the pricing and non-price activities through which firms build demand (Rachel, 2021; Afrouzi et al., 2023; Cavenaile et al., 2025, 2026). Roldan-Blanco and Gilbukh (2021) show that young firms increase their customer base by reducing markups in the early years of operation to attract new customers, while Hubmer and Nord (2025) embed competition for demand in a general equilibrium framework, where greater marketing by firms has the effect of raising customer stealing. My contribution relative to this literature

is to introduce rich heterogeneity on the demand side of the economy. The model introduced in Section 4 features households that differ in income, in baskets consumed, and in responsiveness to shocks. As most existing models of customer capital feature homogeneous buyers, there is consequently no role for the distribution of shocks to matter for the firm side of the economy.

Closest to the household side of the model are papers in which households choose how many suppliers to hold, or in which customers differ in how attached they are to their suppliers. Michelacci et al. (2022) model the set of varieties a household buys, to which it adds at a cost, and show that this extensive margin amplifies the effect of fiscal transfers on consumption by more than 30 per cent. Paciello et al. (2019) make the separation of a customer endogenous, with the seller's price as the trigger, and Bornstein (2025) lets inertia differ with the consumer's age, showing population ageing lowers the share of young firms and raises profits. In my model the separation rate differs with the household's income, which is the dimension along which a fiscal consolidation is designed.

The paper is related to a large literature on the macroeconomic role of demand composition and its aggregate consequences (Boppart, 2014; Handbury, 2021; Sangani, 2022; Nord, 2023; Olivi et al., 2023; Auer et al., 2024; Mongey & Waugh, 2025). Specifically, the paper provides evidence for and adopts the non-homothetic CES demand system of Comin et al. (2021). Jaimovich et al. (2019) utilise a similar mechanism to show that during downturns, households switch to lower quality products, and so demand for income elastic goods contracts disproportionately. I provide evidence for the same behaviour in UK household data, and embed this feature in a model of firm dynamics, linking the composition of demand to the firms that are most exposed to it.

3 Empirical Evidence

This section asks whether and how the decline in household demand after 2010 affected young firms, and shows that this occurred primarily through the margin they spend most intensively on relative to older firms, customer acquisition and growth. The welfare cuts studied in this paper affected firms through two steps. First, the welfare cuts withdrew more income from more exposed districts, and second, households affected by these measures reduced their spending in response. I show evidence from household surveys that both stages are present in the UK, which constitutes a first stage for the demand side effects of the spending cuts. On the firm side, I exploit variation in the treatment intensity to isolate the differential effects on younger and older firms in a given location. In this sense, the focus of the empirical work here identifies slopes across local authorities rather than an aggregate response. Section 3.1 describes the data, Section 3.2 the stylised facts, Section 3.3 the shock measure and the income withdrawal, Section 3.4 the firm-level design and results, Section 3.5 the other margins and Section 3.6 the split by sector type.

3.1 Data

To address the question, I construct a novel dataset by combining five sources: administrative data covering the quasi-universe of UK firms, survey accounts with detailed financial

and balance sheet data (importantly including expenditure on marketing and promotional activities), payroll records, household expenditure microdata and local authority level measures of exposure to welfare reforms. I describe each data source in turn, further details on linking the datasets are in Appendix A.

Business Structure Database (BSD). BSD provides quasi-universal[2] yearly snapshots of firms in the UK, recording the date of creation, location, size and industry classification. I define a firm as young if it is two years old or younger in the BSD.

Annual Business Survey (ABS). The ABS is a size stratified random sample of approximately 60,000 firms per year. The ABS contains extensive financial and balance sheet data, for instance turnover, intermediate good costs, capital expenditure and importantly for this paper expenditure on "advertising, marketing, and promotional costs", which includes advertising campaigns, direct marketing, trade fairs, and related customer-facing activities. I measure firm output as the logarithm of real gross value added, deflated to 2010 prices using the ONS producer price indices. Because the linked data is drawn from the population firm size distribution, it contains the young, small, privately held firms that datasets of listed firms omit, which is the part of the distribution where the mechanism documented in this paper is concentrated.

Annual Survey of Hours and Earnings (ASHE). ASHE is a random sample of employees drawn from HMRC PAYE records, and provides the measure of labour costs used in the empirical results below.

Living Costs and Food Survey (LCFS). The LCFS is a household survey that contains detailed household expenditure by category together with income sources, including income from social security benefits and which benefits the household receives. I use the LCFS for two purposes in what follows. First, it provides the household and local level evidence that benefit recipients cut expenditure following the welfare cuts, and second it disciplines the non-homothetic demand system used in the model. I use the LCFS to estimate sector specific income elasticities that determine whether a sector is classified as a luxury or necessity sector. [3]

Austerity exposure. I construct the treatment measure used below by disaggregating the national welfare spending reductions in a financial loss per working age adult for each UK local authority following the methodology of Beatty and Fothergill (2014). As benefit policy is set nationally, variation in exposure reflects differences in claimant caseloads across authorities, whose geography was in place well before the reforms. Section 3.3 details the construction and measures realised income withdrawal from DWP benefit expenditure and HMRC tax credit statistics.

[2]The BSD consists of firms that are registered for Value Added Tax and/or Pay As You Earn, covering 99% of firms in the UK.

[3]Sectors are classified as luxury or necessity by their income elasticities, and the resulting classification coincides with that of Olivi et al. (2023).

Samples. The empirical analysis uses two samples. The firm-level event studies and pooled regressions use the linked BSD–ABS panel, 2005–2015, excluding non-market sectors and trimming outliers.[4] As ABS financial variables are item-specific, estimation samples differ slightly across outcomes: the output regressions require non-missing GVA and employment ($N = 165{,}314$), while the marketing regressions require non-missing advertising expenditure ($N = 171{,}052$).[5] The household-level analysis uses the LCFS, 2002–2016, in which households with state pension income are kept and counted among the non-recipients. Summary statistics for both samples are reported in Table B.6.

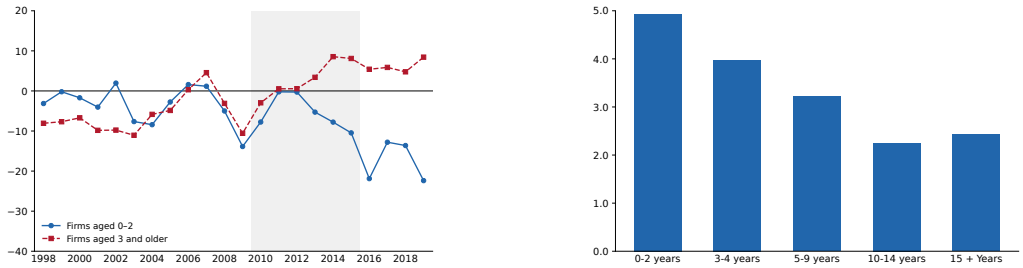
3.2 Descriptive Evidence

Before turning to the causal analysis, I document two stylised facts to motivate the focus upon young firms and their response to demand shocks.

Fact 1: After 2010 output per employee kept growing among older firms but stagnated among young firms.

Fact 2: Customer acquisition spending follows a strong life-cycle pattern: young firms allocate over twice as much revenue to marketing as old firms (5% vs. 2%).

Figure 1a plots log output per employee by young versus old firms. The age bins move together until 2010. After 2010 output per employee rises in older firms, while the youngest firms stall and fall back, so by the end of the period the youngest firms are well below the oldest. Figure 1b plots marketing intensity, defined as expenditure on customer acquisition divided by total turnover, across the firm age distribution.



(a) Output per employee by age bin

(b) Marketing intensity by age bin

Figure (1) Firm Dynamics by Age

The two facts motivate the question set out at the start of this section.

3.3 Austerity Exposure

This subsection details the construction of the shock and provides evidence for the first household stage, the withdrawal of income from exposed districts. Section 3.6 shows the fall in spending that followed.

[4]The non-market sectors are finance and insurance (SIC 64–66), agriculture (SIC 1–3), mining (SIC 5–9), utilities (SIC 35–39), public administration (SIC 84), education and health (SIC 85–88) and real estate (SIC 68). Each year, observations in the top and bottom 0.1 per cent of the distribution of output growth, marketing intensity and input factor shares are dropped.

[5]The estimation window ends in 2015 to avoid confounding from Brexit-related effects, since Fetzer (2019) documents a strong association between this austerity measure and subsequent support for Brexit; the BSD extends further, but declining ABS coverage of financial variables after 2019 limits the usable panel in any case.

The shock measure. The UK's 2010 austerity programme reduced welfare spending by £19 billion annually through cuts to housing benefit, disability allowances, tax credits, and below-inflation uprating of working-age benefits. Local authorities' exposure to this programme differed depending upon the extent to which households in the area were affected by the spending reductions. To construct a local authority specific measure of treatment intensity, I define local authority exposure as:

$$T_d = \frac{\sum_{b=1}^{B} S_{b,d} \times L_b}{W_d} \tag{1}$$

where B denotes the number of reforms, L_b is the national saving from reform b in the Treasury costing, $S_{b,d}$ is local authority d's share of that saving as allocated by the methodology of Beatty and Fothergill (2014), and W_d is the number of working-age adults in local authority d. For all but the three Housing Benefit reforms, $S_{b,d}$ is the authority's share of the national caseload that the reform acts on.[6]

Figure 2 maps the distribution of T_d across Great Britain. Treatment exposure is highest in the former industrial areas of South Wales, Merseyside and the North East. The most exposed local authority, Blackpool, has an estimated loss of £914 per working-age adult; whereas the lowest exposed area, the City of London, has an estimated exposure of £177. Figure 2 shows exactly the cross sectional variation used in the section below, namely whether or not young firms in more exposed (darkly shaded) areas responded differently from young firms in less exposed (lightly shaded) areas, relative to older firms in their respective locations.

Realised income withdrawal. To demonstrate that the predicted exposure measure does capture a reduction in household income, I measure realised benefit expenditure at the local authority level by combining DWP benefit expenditure with HMRC tax credit entitlement at the local authority level, covering fiscal years 2005/06 to 2014/15.[7] I use the combined data to construct a measure of annual benefit and tax credit income per working age adult, in pounds, the same units as T_d.

Benefit income per working-age adult was rising faster in more exposed districts before 2010, as caseloads grew in the recession. After removing that pre-reform trend, £0.68 had been withdrawn per pound of predicted loss by 2014/15 (Appendix A.4.4), which I interpret as evidence that the measure captures actual income withdrawal.

3.4 Event Study

Specification. Motivated by the preceding evidence, I estimate the following event study specification:

[6] For the three Housing Benefit reforms (the Local Housing Allowance, the under-occupation charge and non-dependant deductions) the loss per claimant also varies with local rents, so for them $S_{b,d}$ differs from the caseload share and the variation in T_d comes through the loss per claimant as well as through the caseload.

[7] Expenditure in the DWP accounts is benefit income received by households in the district; the outcome here is income, not consumption spending. Expenditure is divided by the working-age population in 2010 and matched to 2016 local authority boundaries, and the estimation sample is the 315 English districts observed in all ten years.

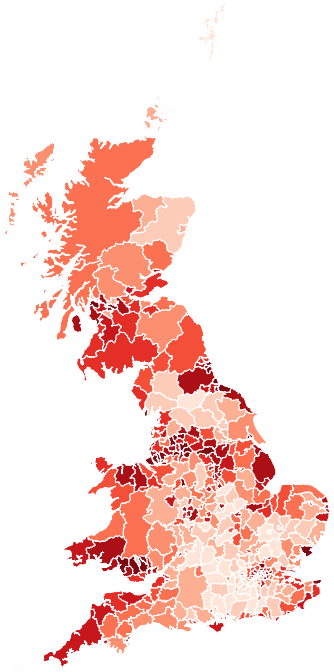


Figure (2) Austerity Intensity. Annual financial loss per working-age adult (T_d) by local authority district, in national deciles; darker shading indicates greater exposure.

$$y_{i,d,t} = \alpha_d + \beta_{s(i)} + \gamma_{r(d),t} + \omega_t \cdot \text{Young}_{i,t} + \upsilon \cdot T_d \cdot \text{Young}_{i,t} + \sum_{\substack{\tau=2005 \\ \tau \neq 2009}}^{2015} \delta_\tau \cdot \mathbf{1}[t=\tau] \cdot T_d \cdot \text{Young}_{i,t} + \sum_{\substack{\tau=2005 \\ \tau \neq 2009}}^{2015} \zeta_\tau \cdot \mathbf{1}[t=\tau] \cdot T_d + X'_{i,d,t}\Gamma + \epsilon_{i,d,t} \quad (2)$$

where $y_{i,d,t}$ is the outcome for firm i in local authority d at time t. The fixed effects are: local authority (α_d), two-digit SIC sector ($\beta_{s(i)}$, where $s(i)$ is firm i's sector), and region-by-year ($\gamma_{r(d),t}$), where $r(d)$ maps each local authority to its NUTS-1 region.[8] The indicator $\text{Young}_{i,t}$ equals one if firm i is aged two years old or less in year t.

The coefficient of interest δ_τ captures the differential effect of a one-standard-deviation increase in exposure on young firms relative to old firms in the same local authority, in year τ, relative to the base year 2009, and υ is the difference between young and old firms in the slope on T_d in the base year itself. The year effects ω_t are specific to young firms, so they absorb the national response of young firms relative to old ones along with anything else common to young firms in every district in a year. $\zeta_\tau \cdot \mathbf{1}[t=\tau] \cdot T_d$ absorb any time varying effect of higher austerity exposure on all firms. In this sense the specification does not rule out effects on the whole firm population in a given location, but the term δ_τ should be viewed as the additional effect on young firms relative to old firms. In equation (2) T_d is the exposure

[8]The region-by-year fixed effects absorb broad regional trends but do not absorb local-authority-level variation in T_d. $\mathbf{1}[t=\tau] \cdot T_d$ captures the effect of the austerity programme on all firms by year, implying that δ_τ captures the additional effect on young firms. Local authority fixed effects absorb persistent spatial differences between those regions.

measure defined above expressed in standard deviations, so one unit is £121, and $\mathbf{1}[t = \tau]$ is an indicator for calendar year τ. The control vector $X'_{i,d,t}$ includes an indicator for whether the firm operates in a tradeable or non-tradeable sector, the local-authority-level unemployment rate, the Herfindahl–Hirschman index (measured at the region by two-digit SIC by year level) and whether the firm is owned in the UK or elsewhere. I also include as an additional control the change in the local authority average of the claimant rates for the benefits in the sample from 2005-2009. As the period of interest to this paper followed immediately from the Great Recession, I additionally interact this control variable with the young firm dummy variable by year to control for any differential effects that could arise from young and old firms being differentially affected by the Great Recession in a given local authority. Standard errors are clustered at the local authority level and regressions are weighted by ABS sampling weights.

Identification then relies on within region and within sector variation in austerity exposure across local authorities. δ_τ then measures the additional movement in young firms' outcomes, relative to older firms in the same district, relative to 2009, in a district that was exposed to one more standard deviation of treatment.

Identification. The design uses spatial variation in a nationally legislated fiscal shock, in the tradition of Mian and Sufi (2014), and it requires three assumptions, which I take in turn together with what the data says about each.

The national savings L_b in equation (1) are common across districts, so all of the variation in T_d comes from the claimant shares $S_{b,d}$ relative to the working-age population W_d. In that sense, identification relies on exogeneity of shares in the Goldsmith-Pinkham et al. (2020) sense. I use claimant shares from the fiscal year 2009-2010, which were therefore fixed before the welfare cuts were announced. As Section 3.3 shows, T_d explains 90 per cent of the variation in district benefit spending as early as 2002, and 96 per cent of districts in the top tercile on 2010 claimant rates were already there on 2005 rates.

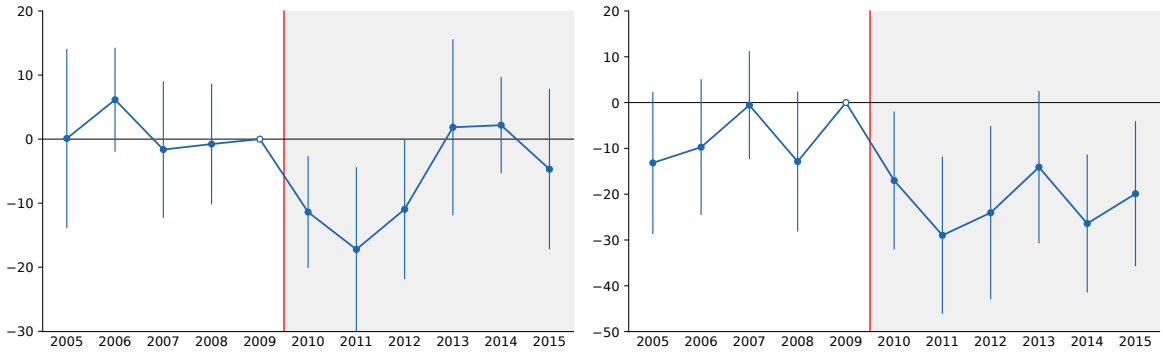
The claimant shares reflect the historical geography of benefit dependency, and they are correlated with much else about a district. Conditional on region effects, a one-standard-deviation increase in T_d comes with (i) a higher jobseeker's allowance claimant rate in 2005; (ii) a lower employment rate in 2009; and (iii) a higher public-sector share of employment. It also comes with a larger rise in the claimant rate between 2008 and 2010, and each of these four differences is between 0.42 and 0.75 standard deviations (Table B.8, Panel A).[9] Note that the design does not require these characteristics to be balanced, since the design here does not compare the levels of more against less exposed districts. Such differences are absorbed by the local authority fixed effects of equation (2). Similarly, any channel that affected younger and old firms together is captured by the year fixed effects. What the design requires is that there was no district level change after 2010 that was correlated with claimant shares that differentially affected young and old firms.

Some district characteristics in the firm data differ with exposure before 2010, notably employment, output per worker and the tradeable share, but not marketing, value added, TFP, markups or the young-firm share, and the 2005–09 changes in marketing and output

[9]Region effects absorb 27 per cent of the variance of T_d, so 85 per cent of its standard deviation, or £103 of £121, is within region.

per worker are unrelated to exposure (Table B.8, Panel B). Panel C makes the comparison that equation (2) makes, between young and old firms in the same district. Pooling the firm-years from 2005 to 2009, the gap between young and old firms did not vary with exposure in any of marketing expenditure, output per worker, employment, markups, TFP, ownership or market concentration (Table B.8, Panel C). Importantly, the design does not require these differences to be zero, since the term $\upsilon\, T_d\, \text{Young}_{i,t}$ absorbs a difference in levels between young and old firms. What it does require is that the two groups were on parallel paths before 2010, which is the second assumption below.

Identification requires parallel trends, namely that the gap between young and old firms followed parallel paths in more and less exposed areas before the welfare cuts were announced. I report the tests for GVA and marketing expenditure directly in the event study figures. Identification additionally relies upon the usual Stable Unit Treatment Value Assumption (SUTVA), namely that outcomes in a district depend only on its own exposure and not on the exposure of other districts. Effects common to all districts, such as a fall in national demand, do not violate it and are absorbed by the year effects, so the concern is spillovers between districts. I test for these spillovers at the end of this section by adding the exposure of neighbouring local authorities to equation (2).



(a) Marketing expenditure

(b) Value Added

Figure (3) Event Study Results. Coefficients on exposure T_d interacted with year, estimated separately for young firms (aged two years or less) and old firms, scaled to a one-standard-deviation increase in T_d and multiplied by 100, with 95% confidence intervals. Base year 2009. Panel (a): marketing expenditure; Panel (b): Value Added. The pooled estimates of δ_τ from equation (2) are in Table B.7 (171,052 and 165,314 firm-years). The joint test that the pooled pre-period interactions for 2005–2008 are zero gives $F = 0.68$, $p = 0.61$ for marketing and $F = 1.8$, $p = 0.13$ for value added. Standard errors clustered by local authority (340 clusters).

Results. Young firms in more exposed districts cut marketing after the shock (Figure 3, Panel (a)). At its peak in 2011 the cut is about 15 log points per standard deviation of exposure, and the coefficients for 2010 to 2012 are significant at the five per cent level in 2011 and 2012 (Appendix Table B.7). The four pre-period coefficients are close to zero and jointly insignificant. The figure estimates the exposure coefficients separately for young and old firms, and in the young-firm regression the coefficients are about -11, -15 and -12 log points in the three years after the shock, each excluding zero. The pre-period for marketing passes the parallel-trends test above. The differential response between young and old firms fades from

2013 onwards[10], which I interpret as evidence for marketing investment as a forward-looking investment that responds to news today, influencing one timing assumption in the model. Additionally, I find that, of the input margins examined in Section 3.5, customer acquisition is the one that responds.

Panel (b) shows that young firms in more exposed districts had lower value added after 2010 than before. The pooled coefficients relative to 2009 run from −0.14 to −0.29 over 2010 to 2014 and are still −0.20 in 2015. In the young-firm regression of Panel (b), where the district effects absorb the difference in levels, none of the four pre-period coefficients is individually significant and the post-2010 coefficients exclude zero in every year through 2015.

Spatial spillovers. The demand of a firm is not restricted by local authority boundaries, so a firm's demand shock may depend on the cuts in neighbouring local authorities as well as its own. Such a spillover would be a concern for the SUTVA assumption detailed above. Table 1 adds to a pooled version of equation (2), with Post equal to one in 2010–2012, the mean exposure of local authorities in successive distance bands, interacted as the treatment variable is with the young firm indicator. The own coefficient is stable across bands and significant at the ten per cent level or better in three of the four bands for each outcome. Within 30 kilometres, where own and neighbouring exposure are most correlated, the neighbouring coefficient is close to zero for both outcomes. Neighbouring exposure is negative at 30 to 90 kilometres, which more likely reflects sub-regional conditions than cross-boundary shopping, but holding it fixed leaves the own coefficient little changed.

Table (1) Own and Neighbouring Exposure

	Marketing expenditure		Output per worker		
Band (km)	Own	Neighbours	Own	Neighbours	Corr.
0–30	-0.087^{*}	−0.012	-0.103^{**}	0.030	0.53
	(0.046)	(0.050)	(0.044)	(0.040)	
30–60	−0.070	-0.069^{**}	−0.056	-0.078^{**}	0.30
	(0.045)	(0.030)	(0.042)	(0.036)	
60–90	-0.080^{*}	-0.080^{***}	-0.075^{*}	-0.059^{*}	0.18
	(0.043)	(0.030)	(0.042)	(0.031)	
90–120	-0.085^{**}	−0.041	-0.076^{*}	−0.040	0.24
	(0.042)	(0.038)	(0.041)	(0.032)	

Notes: Each row reports one regression per outcome of the outcome on Post × Young × own exposure and Post × Young × the mean exposure of the local authorities whose centroids lie within the distance band, with Post equal to one in 2010–2012. Both exposures are in standard deviations, so the two columns are comparable in magnitude. The last column is the correlation between own and neighbouring exposure. 155,113 to 159,533 firm-years for marketing and 150,072 to 154,353 for output per worker. Standard errors clustered by local authority in parentheses, 340 clusters; $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

3.5 Alternative Channels

The sceptical reader may argue that the output and demand investment declines documented above could reflect margins other than customer acquisition. This section discusses other

[10]Note that the structure of the data and the classification of firms as young or old means that a firm born in 2010 would not be classified as a young firm in 2013. From 2013 the comparison group therefore includes the firms that were young in 2010, which pulls the differential towards zero.

Table (2) Young Firms and Exposure: Other Outcomes

	Employment (1)	Tradeable (2)	UK-owned (3)	HHI (4)	Markup (5)	TFP (6)	Net capex (7)	Energy (8)
2005	0.041	0.057*	0.011	34.3	−0.129	0.120*	−0.071	8.1
	(0.055)	(0.032)	(0.013)	(52.5)	(0.169)	(0.068)	(0.099)	(8.4)
2006	−0.057	0.039	0.010	−55.8	−0.225	−0.021	0.118	5.1
	(0.056)	(0.030)	(0.013)	(42.7)	(0.178)	(0.100)	(0.097)	(8.5)
2007	0.042	0.087**	0.008	−46.5	0.219	0.118*	−0.019	3.5
	(0.048)	(0.036)	(0.014)	(33.6)	(0.226)	(0.062)	(0.099)	(9.4)
2008	0.004	0.015	0.001	11.1	−0.006	0.063	0.070	4.0
	(0.057)	(0.045)	(0.015)	(41.6)	(0.059)	(0.040)	(0.100)	(11.0)
2010	−0.013	0.029	−0.015	−15.0	0.015	−0.005	0.026	−0.4
	(0.055)	(0.035)	(0.014)	(46.0)	(0.051)	(0.039)	(0.103)	(10.3)
2011	−0.015	0.021	0.007	−1.1	−0.079	0.001	−0.045	−2.7
	(0.064)	(0.035)	(0.016)	(31.9)	(0.050)	(0.023)	(0.103)	(11.5)
2012	0.060	0.037	0.009	−9.6	−0.159	−0.044	0.039	−3.8
	(0.067)	(0.034)	(0.015)	(30.9)	(0.106)	(0.037)	(0.089)	(12.7)
2013	0.051	0.046	−0.003	34.1	−0.111**	−0.045	−0.171	−2.2
	(0.059)	(0.032)	(0.013)	(58.8)	(0.045)	(0.041)	(0.110)	(11.8)
2014	0.061	0.049	0.009	2.0	0.076	0.013	−0.061	5.3
	(0.052)	(0.036)	(0.014)	(28.2)	(0.204)	(0.043)	(0.098)	(8.3)
2015	−0.075	−0.030	0.015	−5.6	−0.042	−0.053	−0.077	3.8
	(0.067)	(0.033)	(0.014)	(34.5)	(0.090)	(0.034)	(0.107)	(8.5)
Young $\times T_d$, 2009 level	−0.048	−0.037	−0.010	17.2	0.047	0.008	0.004	6.03
	(0.044)	(0.027)	(0.013)	(23.2)	(0.029)	(0.021)	(0.034)	(27.5)
Pre-period F	1.81	1.84	0.51	1.27	0.78	1.87	1.43	0.85
p-value	0.133	0.127	0.730	0.286	0.539	0.122	0.23	0.36
Observations	171,059	171,059	171,059	171,059	143,170	143,146	171,059	171,059
Clusters	340	340	340	340	340	340	340	340

Notes: Coefficients on Young $\times T_d \times$ year (δ_τ), exposure in standard deviations, base year 2009. Columns (1), (5) and (6) are in logs; columns (2) and (3) are indicators for a tradeable sector and UK ownership, and column (4) is the Herfindahl–Hirschman index. Columns (7) and (8) are net capital expenditure and energy expenditure. Standard errors in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

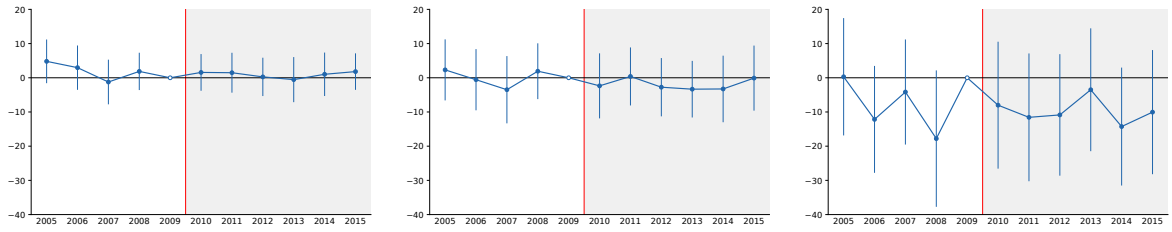
potential channels in turn and provides evidence as to why the demand investment channel is the appropriate one to consider in this setting. Table 2 reports the event study coefficients for employment, markups and firm-level TFP (columns 1, 5 and 6), three measures of the composition of young firms (columns 2–4), net capital expenditure (column 7) and energy expenditure, a proxy for capacity utilisation (column 8).

Production, pricing and utilisation. Young firms in more exposed districts did not reduce their workforce, become less productive, or price differently from older firms; the markup coefficients are insignificant in every year except 2013. Nor did they use less energy. Following Burnside et al. (1995), I treat energy use as a proxy for capacity utilisation: if young firms in more exposed districts were producing further below capacity, their energy use would fall relative to old firms', and column (8) shows no such fall in any year. Similarly, the composition of young firms by tradeable sector, domestic ownership or market concentration, in columns (2) through (4), does not depend upon exposure after 2010. I estimate markups following Kariel et al. (2024), who follow De Loecker and Warzynski (2012) in estimating markups in the UK firm population.[11] The markup here is the ratio of the estimated sector

[11]Note that the production function here is estimated using OLS for each two digit sector - the structure of the ABS data which uses a rotating panel prevents the use of lagged inputs and estimation via the control

output elasticity to the firm's materials revenue share. The markup result has important implications for the modelling choice in section 4. I interpret the null differential response of the markup and the significant differential response of marketing activity as evidence against a channel in which firms build their customer base through pricing (Ravn et al., 2006; Paciello et al., 2019; Bornstein, 2025), and in favour of one in which they acquire customers through marketing (Arkolakis, 2010; Gourio & Rudanko, 2014), consistent with Fitzgerald et al. (2024).

Local labour market. An additional margin could be through the local labour market responding to the shock, and changing input costs, with firms reacting to changed intermediate costs rather than demand conditions. Figure 4 addresses this point directly and shows that neither district level employment nor the district level wage bill respond to an increase in exposure, suggesting that the results are not picking up confounding effects through the local labour market. Panel C repeats the event study specification of equation (2) with the firm wage bill as the dependent variable, and finds no significant difference between young and old firms. Together with the firm level employment null of Table 2, I interpret this as evidence against the mechanism operating through the local labour market or firm level labour hiring or hoarding.



(a) District employment (b) District wage bill (c) Young-firm wage bill

Figure (4) Labour Market Responses. Coefficients on exposure T_d interacted with year, scaled to a one-standard-deviation increase in T_d and multiplied by 100, with 95% confidence intervals. Base year 2009. Panel (a): district employment; Panel (b): district wage bill; Panel (c): the additional response of the young-firm wage bill, Young $\times T_d \times$ year from equation (2). The shaded area is 2010–2015.

Financial frictions. Column (7) of Table 2 addresses the financial frictions channel over and above the baseline specification, which controls for firm size as a measure of size-based financial frictions. It reports the event study specification of equation (2) with net capital expenditure as the dependent variable, and again finds no significant differential response between young and old firms. I view this as evidence that the shock is not transmitted to firm level behaviour through financing constraints binding differentially for young firms, where cuts to marketing investment would reflect financial constraint driven reductions in non production activity.

3.6 Sectoral Heterogeneity: The Role of Non-Homothetic Demand

The results presented above compare young and old firms across all sectors. Here I split the household and firm survey data into luxury and necessity sectors, and show that the

function method of Ackerberg et al. (2015).

mechanism on both sides of the economy is stronger in luxury sectors, where spending is more income elastic. Sectors are classified according to the estimated income elasticities from the LCFS.

Household spending. A decline in household income only matters for firms to the extent that households pass it through to spending and consumption. To demonstrate that households that were exposed to these spending reductions did indeed reduce their consumption, I estimate an event study of household expenditure for benefit recipients relative to non-recipients with household size and region by year fixed effects. Figure 5 plots the coefficients,[12] measured against 2010.[13]

I also use the same panel of households to decompose the spending response by sector type. I find that recipients' spending in luxury categories fell relative to non-recipients, while their spending on necessity categories did not fall significantly relative to the pre spending cut period. This pattern is consistent with non-homothetic household preferences. When income falls, households disproportionately cut expenditure on more income elastic sectors. Figure 5 shows the fall in total spending relative to 2010. Spending in luxury sectors also declines while spending in necessity sectors remains flat. I interpret this as evidence that the total spending reduction was concentrated in luxury sectors.

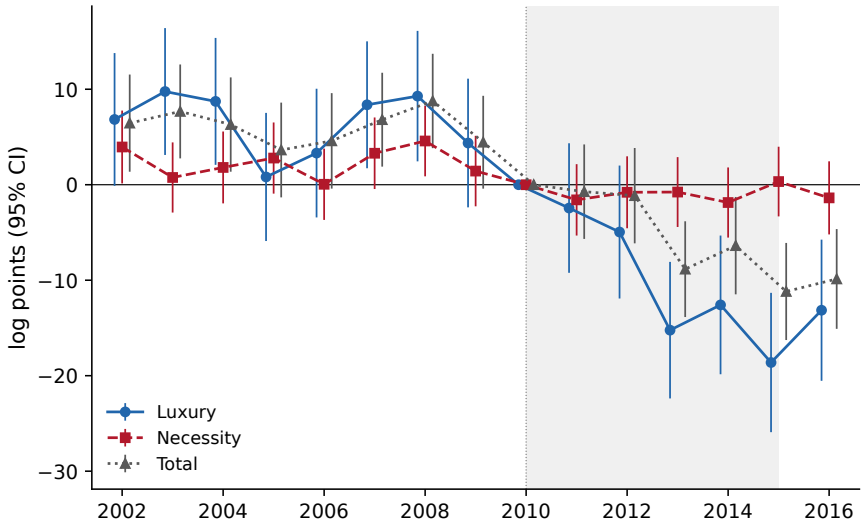


Figure (5) Relative Expenditure of Benefit Recipients vs. Non-Recipients (LCFS). Coefficients on benefit-recipient status interacted with year from regressions of log household expenditure on household-size indicators and region-by-year effects, Living Costs and Food Survey 2002–2016, in log points, with heteroskedasticity-robust 95 per cent confidence intervals. Luxury is COICOP divisions 3, 5, 9, 11 and 12 and necessity is divisions 1, 4, 6 and 8. UKDS microdata.

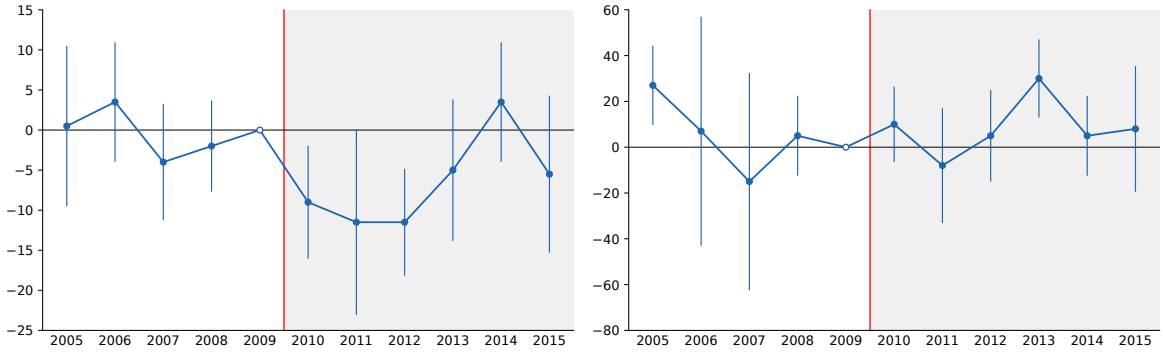
Taking stock, the household evidence supports the link between the policy and demand withdrawal. As the spending cuts took place, more affected households reduced consumption to a greater extent, and this reduction was concentrated in income elastic sectors, which I use

[12] A household is a recipient if it reports positive income from any social security benefit, including child benefit and tax credits, and a household with any state pension income is counted as a non-recipient.

[13] Recipient status is measured contemporaneously, so households that lost eligibility as a direct result of the reforms are coded as non-recipients, contaminating the control group and biasing the estimated gap towards zero, so I view the estimated gap as a lower bound.

as evidence for the introduction of non-homothetic consumer preferences in the model.

Firms' marketing. To examine whether the household level sectoral spending reductions that were concentrated in luxury sectors are passed through to firms, I classify firms into luxury and necessity sectors using the income elasticities estimated from the LCFS demand system used in the model calibration. Figure 6 estimates the event study specification (2) separately for firms in each category, with marketing expenditure as the dependent variable. The results confirm the same sectoral split seen in the household data is present in the firm population. Young firms in luxury sectors cut marketing expenditure relative to older firms in 2010, 2011 and 2012. On the other hand, young necessity firms show no comparable decline relative to their older counterparts. The young firm response of Section 3.4 is therefore stronger in luxury sectors, in line with the household evidence above, which is reflected in the non-homothetic preferences in the model.



(a) Luxury sectors

(b) Necessity sectors

Figure (6) Marketing Response by Sector Type. Coefficients δ_τ on Young $\times T_d \times$ year from equation (2) with marketing expenditure as the outcome, estimated separately for firms in luxury and necessity sectors, scaled to a one-standard-deviation increase in T_d and multiplied by 100, with 95% confidence intervals. Base year 2009. The shaded area is 2010–2015. The two panels have different vertical scales.

3.7 Taking Stock

The empirical evidence presented above is a set of facts about consumer and firm level heterogeneity that a model of the firm must be disciplined by. Following the 2010 welfare cuts which withdrew demand from households, young firms in more exposed areas cut their investment in customer acquisition, while older more established firms did not (Figure 3). This differential is not apparent in other firm level inputs: employment, TFP, capital formation, capacity utilisation and markups do not respond in the same manner (Table 2), nor does the local labour market (Figure 4). On both sides of the market, the response is concentrated in luxury sectors, where recipient households cut their spending the most and where young firms cut their marketing expenditure (Figures 5 and 6). Young firms also exhibit lower value added after 2010, which I interpret as evidence of a demand channel operating through a smaller customer base rather than a fall in production efficiency, given the absence of any differential response in TFP.

The design compares young and old firms within a given local authority, and so cannot say why young firms were more exposed, how much all firms lost, or how a consolidation with

a different distributional incidence would have affected firms and the aggregate economy. The following sections develop a model, disciplined by this evidence, to answer these questions.

4 Model

This section develops a dynamic general equilibrium model of heterogeneous households and firms that is disciplined by the empirical findings and provides a framework for conducting counterfactual exercises. Households have to be matched to a firm in order to buy from them, and firms can invest in customer acquisition, which is the key dynamic channel of the model. Because of household heterogeneity in income (expenditure), the composition of a firm's customer base will be a key state variable that will evolve over the life cycle of the firm as it acquires new customers. It also generates heterogeneous firm exposure to income shocks that hit households differentially, like the austerity measures. Furthermore, the model will also generate differential effects across sectors by allowing for non-homotheticities in demand, so that income elasticities will play a key role in determining how the value of matching to a customer changes due to a shock. Section 4.1 illustrates the mechanism, and shows why young firms are more exposed to a consolidation that falls primarily on low income households. Section 4.2 presents the household problem and solves for demand, Section 4.3 describes how matches are formed and how they evolve, and Section 4.4 lays out the firm's static and dynamic problem. Finally, Section 4.5 defines general equilibrium in steady state and along a transition path.

4.1 Key Channel and Economic Intuition

The model's central prediction is that young firms' revenue, and with it their customer capital, falls by more than older firms' when a consolidation falls on low-income households. Consequently, firms also cut their investment in customer acquisition as the value of a new customer falls, which makes this loss persistent, most of all for young firms who are most exposed to such a shock.

A firm's revenue is the spending of its customers, $R = \int_\Theta n(\theta)\, v(\theta)\, d\theta$, with $n(\theta)$ its relationships with households of type θ and $v(\theta)$ the spending each carries.[14] A shock moves both:

$$d\ln R = \int_\Theta r(\theta)\, d\ln v(\theta)\, d\theta + \int_\Theta r(\theta)\, d\ln n(\theta)\, d\theta, \qquad r(\theta) \equiv \frac{n(\theta)\, v(\theta)}{R} \tag{3}$$

The first term is the impact, when customers are fixed; the second is the customer base, which is where persistence comes from. At impact a firm loses more than the economy only through a covariance between who it sells to, $r(\theta)$ relative to the economy's shares, and how hard those customers are hit.

One Sector Case. Let a consolidation cut the expenditure of households of type θ by the proportion $e(\theta) \equiv -d\ln E(\theta)$, so that $d\ln v(\theta) = -e(\theta)$ on impact and $\frac{R_0}{R} - 1 = -\int_\Theta r(\theta)\, e(\theta)\, d\theta$, with R_0 revenue at impact. The differential between a young firm with revenue shares $r^Y(\theta)$, and an old one, with $r^O(\theta)$ is then:

[14] Up to a firm-specific price factor common to all types, as under constant markups (Section 4.4).

$$\frac{R_0^Y}{R^Y} - \frac{R_0^O}{R^O} \;=\; -\int_\Theta \left[r^Y(\theta) - r^O(\theta)\right]\left[e(\theta) - \bar{e}\right] d\theta, \tag{4}$$

where $\bar{e}$ is a cut common to every type. The differential is therefore a covariance between the composition gap $r^Y(\theta) - r^O(\theta)$ and the cut. Young and old firms' revenues respond differently only if they draw their revenue from the types in different proportions and the incidence of the cut differs across types, and young firms lose more when they draw more of their revenue from the types the cut hits hardest[15].

Here, this composition gap can arise from differential churn. Poorer households end their existing supplier relationships more often, so the pool of searching households is thus poor biased and the customers that firms investing in demand are most likely to acquire are poorer than the population. A young firm whose customers are mostly recent acquisitions from that pool has a poorer and faster churning customer base than an older firm whose customer base is more weighted towards the wealthier households whose relationships last. Here, if all households ended their relationships at the same rate, young and old firms would hold customer bases of the same composition and the differential would be zero even if incidence differed across type. I call this exposure of young firms to the households a regressive cut reaches a "churn trap".

Two sectors and non-homothetic demand. With multiple sectors that differ in income elasticities, a cut $e(\theta)$ in household expenditure lowers its spending on sector j, to first order, by $\eta_j(\theta)e(\theta)$, where $\eta_j(\theta)$ is the elasticity of a type-θ household's spending on sector j with respect to its total expenditure. For a firm in sector j with revenue shares $r_j(\theta)$:

$$\frac{R_0}{R} - 1 \;\approx\; -\int_\Theta r_j(\theta)\,\eta_j(\theta)\,e(\theta)\,d\theta, \tag{5}$$

$$\frac{R_0^Y}{R^Y} - \frac{R_0^O}{R^O} \;\approx\; -\int_\Theta \left[r_j^Y(\theta) - r_j^O(\theta)\right]\left[\eta_j(\theta)e(\theta) - \bar{e}\right] d\theta. \tag{6}$$

Non-homothetic preferences matter for the mechanism in three ways. First, with homothetic demand the cut would be $e(\theta)$ in both sectors, whereas non-homotheticity makes it larger in the luxury sector, so for a given composition gap the differential is larger among luxury firms, which is the pattern of Section 3.6. Second, the mechanism now operates even when the incidence of the cut is equal. For a cut that is the same proportion e for every type, the differential is $-e\int_\Theta \left[r_j^Y(\theta) - r_j^O(\theta)\right]\left[\eta_j(\theta) - \bar{\eta}_j\right] d\theta$ for any common $\bar{\eta}_j$, which is negative in the luxury sector because the income elasticity of luxury spending falls with income, and which is zero if either differential churn or non-homothetic preferences is absent. Third, the luxury elasticity is above one and highest for the poorest households, so non-homotheticity magnifies a regressive cut. In the necessity sector the elasticity is below one and also falls with income, but less steeply, so necessities do not reverse the mechanism, and under a regressive design their elasticity below one dampens the cut.

[15]The two conditions are necessary for a differential, and they are sufficient when the composition gap changes sign once and the cut changes monotonically with income.

4.2 Households

The economy is populated by a unit mass of households, indexed by their labour efficiency θ and distributed on $\Theta \subseteq \mathbb{R}_+$ with density $\phi(\theta)$. Each period, households decide how much labour h_t to supply, and choose how much to consume of each variety $c_{j\omega,t}$ in order to maximise their total utility. The consumption aggregator, which will be discussed in more detail in the next section, can be written in its most general form as $\mathcal{U}(c_{j\omega,t})$. The disutility of labour is given by $\xi_\theta(h)$, where the subscript θ allows for it to potentially depend on the household type[16]. There is an aggregator $u(C, \xi_\theta(h))$ that combines consumption and labour into flow utility. Households also own the firms, which pay a dividend of Π_t, so their utility maximising problem can be written as

$$\begin{aligned} \max_{\{c_{j\omega,t}(\theta),\, h_t(\theta)\}} \quad & \sum_{t\geq 0} \beta^t \, u\Big(\mathcal{U}\big(\{c_{j\omega,t}(\theta)\}\big), \xi_\theta(h_t(\theta))\Big) \\ \text{s.t.} \quad & \sum_{j=1}^{J} \int_{\mathcal{M}_{j,t}(\theta)} p_{j\omega,t}\, c_{j\omega,t}(\theta)\, d\omega \;=\; \theta w_t h_t(\theta) + \Pi_t + \tau_t(\theta), \end{aligned} \tag{7}$$

where w_t is the wage per efficiency unit, and $\tau_t(\theta)$ denotes any transfers that the government implements.[17] Since I do not allow for a saving technology, the household problem essentially collapses to solving a sequence of static problems for the allocation of expenditure and labour supply in each period. Finally, note that because at any point in time households are only matched with a subset of firms $\mathcal{M}_{j,t}(\theta)$, they can only buy from these firms. Here, the household takes the process of matching and separation as given when it allocates its spending, and it follows the law of motion of Section 4.3, so the household treats its potential suppliers as given, though I allow for that set to vary by household type. Appendix D writes down a household problem that delivers this process.

Consumption Aggregator. Preferences have two tiers: a non-homothetic CES aggregator across sectors j, and a standard CES aggregator over varieties ω within each sector. I describe these in turn, and define the corresponding sectoral price indices, as well as the relevant elasticities that enter the firm's problem. For ease of notation, I drop the time t subscript and define the steady state expressions for these objects, where one can write everything in terms of the household type θ. I denote by $E(\theta)$ the total expenditure of the household of type θ, which together with the vector of prices determines the demand function for any variety that the household can buy from.

Letting $c_j(\theta)$ denote sectoral consumption, I assume that the aggregate consumption index is defined by a non-homothetic CES aggregator as in Comin et al. (2021) which is given by

$$C(\theta) \;=\; \left[\sum_{j=1}^{J} \zeta_j^{\frac{1}{\sigma}} \Big(C(\theta)^{\gamma_j}\Big)^{\frac{1}{\sigma}} c_j(\theta)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}, \tag{8}$$

[16] Note that under the model parametrisation here, the labour supply decision does not affect expenditure allocation, and so I do not detail the choice here. Over transitions the labour supply choice will become an important margin that maps to aggregate outcomes, I discuss in more detail in Section 6.

[17] These are set to zero in steady state, but will be used to match the fiscal consolidation that was implemented in the UK, as well as conduct counterfactual exercises.

where $\sigma < 1$ is the elasticity of substitution across sectors, $\zeta_j > 0$ are demand shifters, and γ_j governs the non-homotheticities across sectors. A higher γ_j implies that a sector is more of a luxury, so that its budget share will increase with the expenditure level. This specification also nests the standard homothetic CES when $\gamma_j = 0$ for all sectors. I normalise the necessity sector's γ_j to zero, so that γ_1 in Section 5.2 is the non-homotheticity of luxury spending. Note that the composite is defined as an implicit function of the consumption level $C(\theta)$, which is exactly why a household's effective cost of living depends upon its income level.

Let $p_j(\theta)$ be the price index of sector j, so that the budget constraint of the upper problem is given by $\sum_{j=1}^{J} p_j(\theta)\, c_j(\theta) \;=\; E(\theta)$. The sectoral shares can be written as

$$s_j(\theta) \;=\; \frac{\zeta_j\, p_j(\theta)^{1-\sigma}\, C(\theta)^{\gamma_j}}{\sum_{j'=1}^{J} \zeta_{j'}\, p_{j'}(\theta)^{1-\sigma}\, C(\theta)^{\gamma_{j'}}}, \tag{9}$$

which depend on the household's consumption level $C(\theta)$, and hence vary systematically with income. What matters for firms is not the level of those shares but how far they move when income moves, which is summarised by the elasticity of a type-θ household's spending on sector j, $E_j(\theta) \equiv s_j(\theta)E(\theta)$, with respect to its total expenditure $E(\theta)$.

$$\eta_j(\theta) \;=\; 1 + \frac{1-\sigma}{1-\sigma+\bar{\gamma}(\theta)}\big(\gamma_j - \bar{\gamma}(\theta)\big), \tag{10}$$

where $\bar{\gamma}(\theta) \;=\; \sum_{j=1}^{J} s_j(\theta)\gamma_j$ is the average of the sector parameters γ_j weighted by the household's expenditure shares[18]. Equivalently, the marginal budget share of sector j is $s_j(\theta)\eta_j(\theta)$, so again with homothetic preferences it equals the average budget share and all households respond identically to income changes.

The luxury elasticity falls with income, which is central to the mechanism of Section 4.1. In the LCFS it falls monotonically with income[19]. Consider a poor and a rich household that each reduce expenditure by 1%, disproportionately so for luxuries. The poor household has a smaller luxury budget share to begin with, and so has a larger percentage fall in luxury spending. To see this, define $\eta_C(\theta) = \frac{(1-\sigma)}{(1-\sigma+\bar{\gamma}(\theta))}$ as the elasticity of the consumption composite with respect to expenditure so that $\eta_j(\theta) = 1 + \eta_C(\theta)\big(\gamma_j - \bar{\gamma}(\theta)\big)$. The term $\bar{\gamma}(\theta)$ is lower for a poorer household, because poorer households spend a larger proportion of their budget on necessity goods, so the term $\big(\gamma_j - \bar{\gamma}(\theta)\big)$ is larger for poorer households. The same argument holds for the elasticity of consumption with respect to expenditure η_C, where a lower $\bar{\gamma}(\theta)$ means that a given cut in expenditure lowers the consumption of the composite good of the poor households by more than the wealthier household[20]. It follows that a firm's demand response to an aggregate income shock depends upon the composition of its customer base $\pi(\cdot)$ in Section 4.4, as its different customer types respond at different rates.

[18] Appendix C.1 derives it by differentiating the sector shares and the expenditure function in $\ln C$ at given prices.

[19] With total expenditure instrumented by income, it is 1.29 in the bottom quintile and 1.19 in the top.

[20] As σ approaches one, $\eta_j(\theta) - 1$ becomes proportional to $\big(\gamma_j - \bar{\gamma}(\theta)\big)/\bar{\gamma}(\theta)$, the percentage gap between the sector's γ_j and the household's average. Note that at $\sigma = 1$ the term is undefined.

Sectoral Aggregator. Within each sector, varieties are aggregated with a standard CES function

$$c_j(\theta) \;=\; \left[\int_{\mathcal{M}_j(\theta)} c_{j\omega}(\theta)^{\frac{\varepsilon-1}{\varepsilon}}\, d\omega\right]^{\frac{\varepsilon}{\varepsilon-1}}, \tag{11}$$

where ε is the elasticity of substitution across varieties within a sector. For a given level of expenditure $E_j(\theta)$, the household maximises this composite subject to the budget constraint

$$\int_{\mathcal{M}_j(\theta)} p_{j\omega}\, c_{j\omega}(\theta)\, d\omega \;=\; E_j(\theta). \tag{12}$$

This gives the usual CES demand and sectoral price index,

$$c_{j\omega}(\theta) = \left(\frac{p_{j\omega}}{p_j(\theta)}\right)^{-\varepsilon} c_j(\theta), \qquad p_j(\theta) = \left[\int_{\mathcal{M}_j(\theta)} p_{j\omega}^{1-\varepsilon}\, d\omega\right]^{\frac{1}{1-\varepsilon}}. \tag{13}$$

Consequently, the sector price index each household faces is a household specific index over the suppliers it holds. As a household holds a continuum of suppliers, each individual supplier is of measure zero in the index, so no one firm's price affects it and each firm charges a constant markup. One can show as in Arkolakis (2010) that in this setting although the price index differs across household types θ(types have different suppliers). Within a type, as acquisition is undirected, each household's price index depends only on its type and spending, so each household within a given type holds suppliers drawn from the same distribution and by applying the law of large numbers, faces the same index $p_j(\theta)$.

One can show as in Arkolakis (2010) that the monopolistic competition structure means that while a single relationship is a measure zero component of the set, the index itself does differ across households reflecting the mass of firms it is aware of. In this sense and consistent with CES preferences in general, here households exhibit love of variety,

$$d\ln p_j(\theta) \;=\; -\frac{1}{\varepsilon-1}\, d\ln \int_{\mathcal{M}_j(\theta)} p_{j\omega}^{1-\varepsilon}\, d\omega, \tag{14}$$

so a household that loses suppliers accounting for a share ς of its sector spending faces an effective price index higher by the factor $(1-\varsigma)^{-1/(\varepsilon-1)}$, at unchanged posted prices. With a common price the index falls with the size of the set, $\partial \ln p_j/\partial m_j \;=\; -1/((\varepsilon-1)m_j)$, which is the love of variety. The own index enters the household's problem through the non-homothetic aggregator. The composite the household reaches from $E(\theta)$ is the solution of $\mathcal{E}(C;\mathbf{p}(\theta)) = E(\theta)$, where $\mathcal{E}(C;\mathbf{p})$ is the expenditure needed to reach a composite C at prices $\mathbf{p}$ (Appendix C.1), at its own indices $\mathbf{p}(\theta) = (p_1(\theta),\ldots,p_J(\theta))$, and its shares (9) are evaluated at the same indices, so a household that loses suppliers reaches a smaller composite at unchanged spending, as if the sector's posted prices had risen. Section 4.3 describes how relationships in the economy end and how they are formed.

4.3 Customer–Firm Matching and Separation

Households are not aware of every firm in the economy, and they can only consume varieties from the firms they hold a relationship with. I describe the relationships that a household of type θ holds in sector j by the set $\mathcal{M}_j(\theta)$, and denote its size by $m_j(\theta)$. The particular firms in a set differ across households of the same type, since each household holds the firms whose contacts happened to reach it, and only the number of suppliers $m_j(\theta)$ and the sector shares $s_j(\theta)$ are common to the type.

For each household, the number of suppliers evolves according to:

$$m_{j,t+1}(\theta) \;=\; \big(1-\bar{\psi}_{j,t}(\theta)\big)\big(1-\lambda_{j,t}(\theta)\big)\, m_{j,t}(\theta) + A_{j,t}(\theta). \tag{15}$$

The first term is the set the household keeps from last period. Each relationship ends with probability $\lambda_{j,t}(\theta) = \lambda_j\big(\theta, q_{j,t}(\theta)\big)$, which is household specific and depends upon its type and its spending per supplier, $q_{j,t}(\theta) \equiv E_{j,t}(\theta)/m_{j,t}(\theta)$. A relationship is also dissolved if the supplier exits, with probability $\bar{\psi}_{j,t}(\theta)$, the exit rate averaged over the suppliers that a type θ household holds, which Section 4.5 expresses in terms of the firm exit rate $\psi(a)$. A relationship therefore survives any given period with probability $\big(1-\bar{\psi}_{j,t}(\theta)\big)\big(1-\lambda_{j,t}(\theta)\big)$. The second term, $A_{j,t}(\theta)$, is the number of new suppliers a household takes on. Each relationship that ends, through separation or exit, leaves a place open. A household with a place open is in the pool that firms' marketing reaches, and one with none open is out of it. Contacts come from firms in proportion to their marketing effort and are matched to open places at random. The separation rate summarises how relationships end, and rather than derive it from an explicit household decision problem I impose the following restrictions on the function λ_j.

Assumption 1 (Separation)**.** Let $E_j(\theta)$ denote the sector spending of a type-θ household in the stationary equilibrium, which is increasing in θ, $m_j(\theta)$ the number of suppliers it holds, $q_j(\theta) \equiv E_j(\theta)/m_j(\theta)$ its spending per supplier, and $\lambda_j(\theta) \equiv \lambda_j(\theta, q_j(\theta))$ its stationary separation rate. (i) Along the stationary cross-section the separation rate falls with type, $\lambda_j'(\theta) < 0$. (ii) For a given type the separation rate falls with spending per supplier, $\partial\lambda_j(\theta,q)/\partial q < 0$. (iii) The separation rate does not depend on the prices the household's suppliers post.

The first restriction is taken across households in the stationary equilibrium, where spending $E_j(\theta)$ rises with θ, and says that poorer households end more of their relationships each period than richer ones do. Section 4.4 uses it to derive the composition of the pool from which firms acquire and of firms' customer bases. The second restriction says that a household that spends less ends more of its relationships, reflecting the evidence of Kaplan and Menzio (2016) and Nord (2023). A consolidation therefore ends relationships, rather than only changing the pounds each of them carries, and it ends most of them where the cut in spending falls. The third restriction means that a household does not adjust its spending to keep its suppliers, and that a firm cannot hold a customer by pricing, which Section 4.4 uses.[21]

[21] This is consistent with the null response of markups in Table 2, which I read as evidence against differential dynamic pricing.

Assumption 1 restricts signs. The separation rate used in the numerical exercises of a type-θ household is determined by two objects. Namely, the stationary separation rate $\lambda_j(\theta)$, and how the separation rate changes as expenditure changes. The elasticity is a constant $\kappa > 0$, the same for every type and sector, so that along a transition the rate follows:

$$\lambda_j(\theta, q) \;=\; \lambda_j(\theta) \left(\frac{q}{q_j(\theta)} \right)^{-\kappa}. \tag{16}$$

Terms without a time subscript are stationary values throughout. In the quantitative model, where θ is binned into income quintiles (Section 5.1), the separation rate is therefore summarised by six numbers, the five stationary rates in Table 4, one for each income quintile and common to both sectors, and κ.

Discussion. Kaplan and Menzio (2016) and Nord (2023) show that the unemployed and households with lower expenditure spend more time and effort shopping and pay less for the same goods. Understanding Society asks at each annual interview whether the household has satellite or cable television, a landline, a mobile telephone and internet access at home. Between 2009 and 2015, households in the bottom quintile of equivalised income stopped having one of these items between one interview and the next 2.20 times as often as those in the top quintile. Holding fixed the age of the reference person, the type of household, the region and the item, the ratio is 1.95 (Appendix B.7).[22] The same panel supports the second restriction, since a household whose own income falls ends more of its relationships.

Appendix D writes down a household problem that delivers the restrictions of Assumption 1, the path rule (16) and how the number of suppliers a household holds varies with its type. It takes the Pareto-distributed cost of Bornstein (2025) and the costly adoption of varieties of Michelacci et al. (2022), with time in place of money; in Bornstein the cost is one of switching, while here it is one of keeping a supplier. In that problem finding a supplier takes time, so a household holds a finite set, and a relationship ends when a disruption specific to that supplier costs the household more than the relationship is worth. When the cost of a disruption has a Pareto tail the path rule holds exactly, with κ the index of the tail, and a scale of the cost that differs across types fits any profile of stationary rates $\lambda_j(\theta)$.

4.4 Firms

Each sector j is populated by a continuum of monopolistically competitive firms, each producing a single variety with a linear technology using labour only $y_{j\omega} = z_{j\omega} l_{j\omega}$. Firms differ in their productivity $z_{j\omega}$ which is drawn at entry from a lognormal distribution and fixed thereafter. Every period, conditional on not being hit by an exit shock, firms have to set a price and choose how much to spend on advertising to attract new customers. Because of the assumptions made, the pricing decision of firms is static and collapses to a constant markup, and all the interesting action happens on the dynamic choices of customer accumulation. I

[22] The survey records whether a household has an item and not who supplies it, so a change of supplier within an item counts as continuing. In the British Household Panel Survey for 1996 to 2008 the two ratios are 1.79 and 2.09.

now describe both those problems in turn.

Static Problem. Let $\boldsymbol{n}_{j\omega}(\cdot)$ on Θ be the mass of customers that the firm is currently matched to, where I am suppressing the time t for ease of notation. This is the firm's key state variable, and it summarises everything that the firm needs to keep track of on the household side. For example, the customer base of any firm is given by $N_{j\omega} = \int_\Theta n_{j\omega}(\theta)\, d\theta$, while $\pi_{j\omega}(\theta) = \frac{n_{j\omega}(\theta)}{N_{j\omega}}$ describes its composition.
Using the household's individual demand for any given variety given in equation (13), one can write the total demand that the firm faces as

$$D_{j\omega}(p) = \int_\Theta n_{j\omega}(\theta) \left(\frac{p}{p_j(\theta)}\right)^{-\varepsilon} \frac{E_j(\theta)}{p_j(\theta)}\, d\theta, \tag{17}$$

so that $D_{j\omega}(p) = p^{-\varepsilon}\tilde{D}_{j\omega}$ with $\tilde{D}_{j\omega} \equiv \int_\Theta n_{j\omega}(\theta)\, p_j(\theta)^{\varepsilon-1} E_j(\theta)\, d\theta$, which does not depend on the firm's own price.

With homothetic preferences, the sectoral shares $s_j(\theta)$ would be common across types and the total demand can be written as a function of the customer base of the firm. Given non-homotheticities in the sectoral aggregator, $s_j(\theta)$ will vary systematically with the household's type (expenditure) so the composition of customers will matter for determining the total demand curve. Nonetheless, because the price elasticity ε is common across households, the curvature of total demand is fixed by this parameter, even though its level shifts with the composition of the firm's customers. The pricing problem is therefore the standard static one,

$$\max_p \ \Big(p - \frac{w}{z_{j\omega}}\Big) p^{-\varepsilon} \tilde{D}_{j\omega} \qquad \Longrightarrow \qquad p \ = \ \frac{\varepsilon}{\varepsilon - 1}\, \frac{w}{z_{j\omega}}, \tag{18}$$

giving rise to a constant markup over marginal cost, independent of the customers' sets $\boldsymbol{n}_{j\omega}(\cdot)$. Three things make the constant markup the firm's optimum, one a result and two assumptions. The firm is of measure zero in every customer's set, which is a result of the continuum of firms (as in Arkolakis, 2010), so its price leaves its customers' indices unchanged and the price elasticity of demand is the within-sector elasticity ε for every customer. Retention does not depend on the firm's price and acquisition is undirected, so a lower price does not attract customers either, and there is no customer-market motive for pricing.

Note that this formulation makes clear that conditional on productivity, two firms in the same sector of different ages will charge the same price and thus extract the same markup. In that sense, the model here reflects the empirical evidence in Section 3.5 that found no differential markup response between young and old firms. Consequently there is no role for dynamic pricing over the life cycle. A firm that loses demand cannot respond by repricing to attract new customers. It can only adjust its spending on acquisition.

Given prices, a type-θ relationship yields revenue $v_j(\theta, z) = p(z)^{1-\varepsilon} p_j(\theta)^{\varepsilon-1} E_j(\theta)$, and flow profit is

$$\Pi_j\big(z, n(\cdot)\big) \ = \ \int_\Theta n(\theta)\, \varpi_j(\theta, z)\, d\theta, \qquad \varpi_j(\theta, z) \ = \ \frac{v_j(\theta, z)}{\varepsilon}, \tag{19}$$

with $p(z)$ from (18). Profit depends on how many customers a firm has and on who they are, since richer households buy more of every variety, and it is linear in the customer measure. Quantity is the only dimension through which firms extract different profit levels for identical units per customer, since consumer preferences imply a homogeneous markup across households.

Customer dynamics and the acquisition pool. A firm attracts new customers by spending on marketing. Its marketing reaches households without regard to their type, and a household that is reached takes the firm on if it is accepting new suppliers, that is, if it has an open place. Acquisition is undirected and allocated at random across the acquiring pool of households, not the population at large. Consequently firms that are acquiring more customers are more exposed to any fluctuations in the pool of households looking for a new supplier.

In a stationary equilibrium every household refills exactly the relationships it ends and the ones it loses when a supplier exits, so the pool has density

$$F_j(\theta) \;=\; \frac{\left[\bar{\psi}_j(\theta) + (1-\bar{\psi}_j(\theta))\lambda_j(\theta)\right] m_j(\theta)\,\phi(\theta)}{\int_\Theta \left[\bar{\psi}_j(\theta') + (1-\bar{\psi}_j(\theta'))\lambda_j(\theta')\right] m_j(\theta')\,\phi(\theta')\,d\theta'}, \tag{20}$$

which over-weights households whose relationships end more often and households that hold more suppliers. In the rest of this subsection I suppress the sector index on F, λ, n and V. Here $\bar{\psi}_j(\theta)$ is the supplier exit rate of (15), written out in (28), which includes certain exit at the terminal age of Section 4.4.1, and note that $F = \phi$ exactly when $[\bar{\psi}_j(\theta)+(1-\bar{\psi}_j(\theta))\lambda(\theta)]\,m(\theta)$ is constant across types. Households that spend more hold more varieties (Michelacci et al., 2022), which pushes the pool towards richer types. Comparing types, the pool is poorer than the population only when $m(\theta)$ rises with income more slowly than its turnover $\bar{\psi}_j(\theta) + (1-\bar{\psi}_j(\theta))\lambda(\theta)$ falls (Appendix C.2), as it does at the calibration. Because a household spreads its spending in a sector over the suppliers it holds, the number it holds sets how many relationships each type adds to the pool but not the total spending those relationships carry. The revenue firms acquire from type θ is proportional to $[\bar{\psi}_j(\theta) + (1-\bar{\psi}_j(\theta))\lambda(\theta)]\,E_j(\theta)\,\phi(\theta)$, so relative to spending the pool over-weights the types whose relationships turn over fastest, which at the calibration are the poorer ones.

Choosing marketing x_t at cost wx_t^ν with $\nu > 1$, a firm's customer measure evolves as

$$n_{t+1}(\theta) \;=\; \big(1-\lambda_t(\theta)\big)\,n_t(\theta) + x_t^\alpha F(\theta), \tag{21}$$

where $n_t(\theta)$ is the measure of type-θ customers the firm holds in period t, $n_{t+1}(\theta)$ the measure it carries into period $t+1$, and $\alpha \in (0,1)$ governs the concavity of acquisition. Acquisition is undirected; a firm can choose how much marketing investment to undertake but not which households it is directed at. This means that new customers arrive with the composition of the pool rather than of the firm's existing base. Separation, instead, is type specific, so a given customer base changes its composition at a rate governed by the separation rate λ. Entrants draw their initial level of customers $\bar{n}_0$, whose composition is exactly that of the pool, so that

at age $a = 0$ a firm holds

$$n(\theta) \;=\; \bar{n}_0\, F(\theta), \tag{22}$$

so every entering cohort inherits the pool by construction. Because poorer households' relationships end sooner, a firm's composition then drifts away from the pool towards $F(\theta)/\lambda(\theta)$, the composition of the relationships that survive, which is richer than the pool whatever the pool is. A firm of any age therefore holds a richer base than an entrant. An older firm also holds a richer base than a younger one under a condition on contacts, and the two together give the composition gap in (6).[23]

4.4.1 Dynamic customer acquisition

Firms discount at β and exit exogenously at rate $\psi(a)$. The probability that a firm of age a exits falls with its age until it reaches a floor,

$$\psi(a) \;=\; \max\{\psi_0 - \psi_1\, a,\; \underline{\psi}\} \text{ for } a < \bar{a}, \qquad \psi(\bar{a}) = 1, \tag{23}$$

and a firm that reaches the terminal age $\bar{a}$ exits after producing. A surviving firm solves

$$\begin{aligned} V\big(a, z, n(\cdot)\big) = \max_{x \geq 0} \Big\{ \Pi_j\big(z, n(\cdot)\big) - w\, x^{\nu} + \beta\big(1 - \psi(a)\big)\, V\big(a+1, z, n'(\cdot)\big) \Big\} \\ \text{s.t.} \quad n'(\theta) = \big(1 - \lambda(\theta)\big)\, n(\theta) + x^{\alpha} F(\theta) \quad \forall\, \theta \in \Theta \end{aligned} \tag{24}$$

with $V(\bar{a}, z, n(\cdot)) = \Pi_j(z, n(\cdot))$, where a prime marks next period's value of a state, so that $n'(\cdot)$ is the measure $n_{t+1}(\cdot)$ of (21) at the stationary rates.

Age is a state of the problem because a firm exits for certain at the terminal age $\bar{a}$ and its exit rate falls with age until it reaches a floor, as in (23). The value of a new customer, and with it the firm's marketing, therefore depends on how long the firm can expect to keep the customers it acquires. Optimal marketing satisfies

$$\begin{aligned} x^{*} \;&=\; \left[\frac{\beta\big(1 - \psi(a)\big)\, \alpha}{\nu\, w}\, V_{\text{new}} \right]^{\frac{1}{\nu - \alpha}}, \\ V_{\text{new}} \;&\equiv\; \int_{\Theta} F(\theta)\, \frac{\partial V\big(a+1, z, n'(\cdot)\big)}{\partial n'(\theta)}\, d\theta \;=\; \int_{\Theta} F(\theta)\, \mathcal{V}_j(\theta; a+1, z)\, d\theta, \end{aligned} \tag{25}$$

the pool-weighted marginal value of a new customer. Here a type-θ relationship of a firm of age $a < \bar{a}$ survives into the next period with probability $\varrho_j(\theta; a) = (1 - \psi(a))(1 - \lambda_j(\theta))$, and its value to the firm is $\mathcal{V}_j(\theta; a, z) = \varpi_j(\theta, z) + \beta \varrho_j(\theta; a) \mathcal{V}_j(\theta; a+1, z)$, with $\mathcal{V}_j(\theta; \bar{a}, z) = \varpi_j(\theta, z)$ (Appendix C.3). A consolidation lowers the profit $\varpi_j(\theta, z)$ that a firm earns from the types whose spending it cuts, and V_{new} weights each type's value by the pool. The marketing of every firm therefore falls with it, which is the response previewed in Section 4.1.

Entry. Entry is sector-specific and elastic. An entrant to sector j pays f_j units of labour, draws its productivity z from G_j, and begins with $\bar{n}_0$ relationships in the composition of the

[23] The condition is that contacts are constant across ages and that an entrant starts with no more relationships of each type than it would hold in the long run (Appendix C.3).

pool, the base (22). The mass of entrants in sector j responds to the ratio of the expected value of entry to its cost,

$$M_{e,j} \;=\; \bar{M}_{e,j}\Big(\frac{V_j^e}{w f_j}\Big)^{\nu_e}, \qquad V_j^e \;=\; \int V_j\big(0, z, \bar{n}_0 F(\cdot)\big)\, dG_j(z), \tag{26}$$

where $\bar{M}_{e,j}$ is the stationary entrant mass, so entry is exactly free in the stationary equilibrium ($V_j^e = w f_j$) and ν_e governs the elasticity of the entry response along transitions. In the limit $\nu_e \to \infty$ entry is infinitely elastic and adjusts instantaneously to ensure free entry holds every period. A consolidation that lowers the value of entry therefore lowers the mass of firms that start from the pool. In the stationary equilibrium these are the firms with the poorest bases, since an entrant's customers are poorer than those of a firm of any older age (Appendix C.3).

4.5 General Equilibrium

In the stationary equilibrium the firm distribution, household incomes and the sets households hold are determined jointly: the labour market clears, entry is free and prices are consistent with firms' pricing. The government appears only through the consolidation of Section 6.1.

In this section I restore the firm index ω suppressed in Section 4.4. Sector j has a set Ω_j of active firms, of mass M_j. The firm producing variety $\omega \in \Omega_j$ is in state $\big(a_{j\omega}, z_{j\omega}, n_{j\omega}(\cdot)\big)$, posts the price $p_{j\omega}$ of (18), earns $v_j(\theta, z_{j\omega})$ from each type-θ relationship and the flow profit $\Pi_{j\omega} \equiv \Pi_j\big(z_{j\omega}, n_{j\omega}(\cdot)\big)$ of (19), and spends $w\, x^{\nu}_{j\omega}$ on marketing. I write μ_j for the distribution of these states across Ω_j.

Adding up. Households spend their income every period, and firms sell only to households that hold a relationship with them. The relationships the firms of sector j hold with type θ are therefore the suppliers type-θ households hold, and the revenue those relationships carry is what the households spend in the sector:

$$\begin{aligned} \int_{\Omega_j} n_{j\omega}(\theta)\, d\omega \;&=\; m_j(\theta)\, \phi(\theta), \\ \int_{\Omega_j} n_{j\omega}(\theta)\, v_j(\theta, z_{j\omega})\, d\omega \;&=\; E_j(\theta)\, \phi(\theta), \end{aligned} \qquad \forall\, \theta \in \Theta. \tag{27}$$

In the same terms, a type's own sector index (13) and the exit rate of its suppliers in (15) are

$$\begin{aligned} p_j(\theta) \;&=\; \Big[\frac{1}{\phi(\theta)} \int_{\Omega_j} n_{j\omega}(\theta)\, p_{j\omega}^{1-\varepsilon}\, d\omega\Big]^{\frac{1}{1-\varepsilon}}, \\ \bar{\psi}_j(\theta) \;&=\; \frac{1}{m_j(\theta)\, \phi(\theta)} \int_{\Omega_j} \psi(a_{j\omega})\, n_{j\omega}(\theta)\, d\omega, \end{aligned} \tag{28}$$

with $\psi(\bar{a}) = 1$. These hold as identities in every period (dated t along a transition): the relationships firms hold and the suppliers households hold are two slices of one measure (Section 4.3 and Appendix C.2), and period-t spending goes to period-t relationships on both sides.

Matching. Contacts are matched to open places at random, and in the stationary equilibrium every open place is filled. For every type θ, the relationships its households open in sector j each period therefore equal those that incumbents' marketing and entrants' initial bases take from the type,

$$\big[\bar{\psi}_j(\theta) + \big(1 - \bar{\psi}_j(\theta)\big)\lambda_j(\theta)\big]\, m_j(\theta)\, \phi(\theta) \;=\; F_j(\theta)\Big[\int_{\Omega_j} x_{j\omega}^{\alpha}\, d\omega \;+\; \bar{n}_0\, M_{e,j}\Big], \tag{29}$$

which, with $\int_\Theta F_j(\theta)\, d\theta = 1$, gives the pool (20).

Prices and the wage. The market index of sector j, taken over all of its firms rather than over one household's suppliers, is consistent with firms' pricing,

$$P_j \;=\; \Big[\int_{\Omega_j} p_{j\omega}^{1-\varepsilon}\, d\omega\Big]^{\frac{1}{1-\varepsilon}} \;=\; \frac{\varepsilon}{\varepsilon - 1}\, w \Big[M_j\, \mathbb{E}_{G_j}\big(z^{\varepsilon-1}\big)\Big]^{\frac{1}{1-\varepsilon}}, \qquad j = 1, \ldots, J, \tag{30}$$

where the second equality uses that exit does not depend on productivity, so the productivity of active firms is distributed as G_j. One sector's market index is normalised to one and serves as the numéraire. With prices a constant markup over w/z, that sector's condition determines the wage from its productivity distribution and its mass of firms, and the conditions of the other sectors determine their indices.

Profits and labour. Firms' operating profits, net of what they spend on acquisition and entry, are rebated equally to households,

$$\bar{\Pi} \;=\; \sum_{j=1}^{J} \int_{\Omega_j} \Big[\Pi_{j\omega} - w\, x_{j\omega}^{\nu}\Big]\, d\omega \;-\; \sum_{j=1}^{J} w\, f_j\, M_{e,j}. \tag{31}$$

Along a transition the rebate is Π_t, with $\bar{\Pi}$ its stationary value. Labour is used in production, marketing and entry, and the labour market clears when

$$\int_\Theta \theta\, h(\theta)\, \phi(\theta)\, d\theta \;=\; \sum_{j=1}^{J} \int_{\Omega_j} \Big[\frac{y_{j\omega}}{z_{j\omega}} + x_{j\omega}^{\nu}\Big]\, d\omega \;+\; \sum_{j=1}^{J} f_j\, M_{e,j}, \tag{32}$$

where $h(\theta)$ is the labour supply of type θ from (7) and $y_{j\omega} = D_{j\omega}(p_{j\omega})$ is the output of firm ω, from (17).

In the stationary equilibrium, given (27), (31) and the household budgets, (32) holds exactly by Walras' law at any positive wage, because the whole of each household's expenditure is spent on goods.

Equilibrium. A stationary equilibrium consists of a wage w, market indices P_j, a rebate $\bar{\Pi}$ and entrant masses $M_{e,j}$; for every type $\theta \in \Theta$, labour supply $h(\theta)$, expenditure $E(\theta)$, sector spending $E_j(\theta)$, own indices $p_j(\theta)$, numbers of suppliers $m_j(\theta)$ and separation rates $\lambda_j(\theta)$; pools F_j; firm values V_j and marketing policies x_j; and invariant distributions μ_j of firms over $(a, z, n(\cdot))$. In it,

(i) households choose $h(\theta)$ and spend $E(\theta) = \theta\, w\, h(\theta) + \bar{\Pi}$ as in (7), allocate it by (9) and (13), and hold sets that are stationary under (15) at the rates $\lambda_j(\theta)$ of (16);

(ii) firms price by (18), with marketing policies x_j, and value functions V_j that solve (24);

(iii) every open place is filled, (29), and the pool of potential customers is stationary and given by (20);

(iv) entry is free, $V_j^e = w\, f_j$;

(v) relationships and revenue add up by type, (27), and own indices and supplier exit rates satisfy (28);

(vi) market indices are consistent with pricing, (30), with one of them the numéraire, and the rebate is (31); and

(vii) each μ_j is invariant under (21), entry (26) and exit (23).

A transition is a sequence of these objects, dated t, given a path of transfers $\tau_t(\theta)$, which Section 6.1 sets from the consolidation. Along it the wage is held at its staionary value, and the rebate Π_t is solved period by period from (31). Households enter a given period with their sets taken as given, and choose hours and spending in each period. The separation rates $\lambda_{j,t}(\theta)$ then follow as a function of the household's expenditure and consequently its average spending with each current supplier. A firm's contacts are drawn from the acquisition pool F_j, and (27) and (28) hold period by period, so each household's own indices move with the suppliers it holds. As the wage is held fixed over the transition, the labour market need not clear. The labour that firms stop using, in production as well as in acquisition and entry, is idle and is not paid. Households spend their labour income and the rebate less the withdrawal, while firms pay only the labour they use, so the wages of the idle labour equal the pounds withdrawn in every period. The withdrawal is therefore the wages that households of each type lose, in the proportions of the design, and no government or foreign agent receives it. Every household spends the whole of its income, so a pound withdrawn from any type cuts spending by a pound. With a realistic distribution of marginal propensities to consume, lower-income households would cut spending by more per pound withdrawn than higher income households, and the 2010 design would withdraw more spending than the progressive design for the same pounds. Section 5.1 describes the computation, in which Θ is discretised into income quintiles.

5 Quantitative Analysis

I calibrate the model to the firm-level data of Section 3.1 and to the household evidence of Section 3.6. Table 3 links the model objects to their empirical counterparts. Output in the model maps to total turnover, and marketing expenditure wx^{ν} corresponds to the ABS question on advertising and marketing services. Consequently the model produces the same marketing intensity variable as plotted in Figure 1b. The model's customer composition π does not have a direct empirical counterpart, and Section 6.2 tests it on moments the calibration does not target. Section 5.1 describes the computation, Section 5.2 the calibration and the fit.

Table (3) Mapping between model objects and empirical measures

Model Object	**Empirical Measure**	**Data Source**
Firm-level variables		
Output $y = z \cdot l$	Total turnover	ABS
Marketing expenditure wx^{ν}	Advertising & marketing spend	ABS Question 8.5
Marketing intensity wx^{ν}/R	Marketing spend / turnover	ABS
Labour l	Number of employees	BSD
Firm age a	Years since incorporation	BSD
Productivity z	(Latent; identified from size distribution)	ABS/BSD
Household-level variables		
Expenditure E_k	Total weekly spending by income quintile	LCFS
Luxury expenditure share s_{1k}	Non-essential spending share	LCFS
Income elasticity η_{jk}	Engel curve slopes	LCFS
Aggregate / geographic variables		
Resources withdrawn	Austerity shock per working-age adult, by income group	DWP/HMRC; De Agostini et al. (2018)
Markup $\varepsilon/(\varepsilon - 1)$	Estimated markups	De Loecker and Eeckhout (2018); Aquilante et al. (2019)

5.1 Stationary Equilibrium and Computation

The stationary equilibrium is as defined in Section 4.5. The firm's state includes a measure of customers, so the state space is in principle infinite-dimensional. In the numerical solution I discretise the customer distribution into K income bins[24]. Since the customer distribution is determined by heterogeneous labour efficiency θ, this amounts to partitioning a log-normal distribution into K bins, whose bin means are available in closed form. Instead of the measure $n(\theta)$, the firm then tracks a stock n with a histogram $\boldsymbol{\pi} = (\pi_1, \ldots, \pi_K)$, and the separation function $\lambda_j(\theta)$ becomes λ_{jk}, one rate per income bin evaluated at the bin's spending per relationship. More generally, a subscript k on a function of θ, such as E_k or η_{jk}, denotes its value at θ_k, the mean of bin k, whose mass is $1/K$. The full state space is $(a, n, z, \boldsymbol{\pi})$: age, total customer count, productivity, and the K-dimensional customer composition vector. The household side is static given (E_k) and is solved once per period at the market indices.

The equilibrium is computed as an outer fixed point on the acquisition pool around a solve for the remaining objects. Given a pool, the solver guesses the wage and luxury price, the entrant masses, the profit rebate and the own indices, solves for the aggregate outputs from the household block as described above. It then solves the firm problem for each sector, computing optimal marketing policies $x^*(a, n, z, \boldsymbol{\pi})$. The stationary firm distribution follows by fixed-point iteration on the distribution operator, which incorporates ageing, per-bin customer

[24] With $\ln\theta \sim \mathcal{N}(\mu_\theta, \sigma_\theta^2)$, the bins are the K equal-mass partitions $[q_{k-1}, q_k]$, where $q_k = \exp\{\mu_\theta + \sigma_\theta\,\Phi^{-1}(k/K)\}$ and Φ is the standard normal CDF. Each bin is represented by its conditional mean, which for the log-normal is available in closed form:

$$\theta_k \;=\; \mathbb{E}\big[\theta \,\big|\, q_{k-1} \le \theta < q_k\big] \;=\; K\, e^{\mu_\theta + \sigma_\theta^2/2}\Big[\Phi\big(\Phi^{-1}(k/K) - \sigma_\theta\big) - \Phi\big(\Phi^{-1}((k-1)/K) - \sigma_\theta\big)\Big],$$

with population shares $1/K$.

evolution, composition updating, entry and exit.[25]

5.2 Calibration

The model parameters are set in two stages. Parameters with direct empirical counterparts or established values in the literature are fixed externally. The remaining eight parameters are calibrated jointly to match moments of UK firm-level data from the Annual Business Survey via simulated method of moments. The objective minimises the weighted distance between model and data over twelve moments. These are average firm size, the mean marketing-to-revenue ratio across firms, five age-binned marketing-to-revenue ratios, the aggregate luxury expenditure share, the shares of firms that survive to ages five and fifteen, and the event study's two coefficients for 2010, namely the differential marketing expenditure response of young and old firms to a given demand shock (Section 3.4).

5.2.1 Externally Set Parameters

Table 4 reports the externally set parameters. $\sigma = 0.5$ is a standard value in the literature (Comin et al., 2021). With $\sigma < 1$ the sectors are complements, so a rise in a sector's relative price raises its budget share, while the fall in the luxury share with income comes from γ_1. The within sector elasticity is set at $\varepsilon = 3.5$ to generate a markup of 1.40, consistent with recent UK estimates (De Loecker & Eeckhout, 2018; Aquilante et al., 2019). The necessity demand shifter is normalised to $\zeta_0 = 0.80$.

Three further parameters are fixed from the household data rather than calibrated to firm moments. The luxury non-homotheticity parameter $\gamma_1 = 1.22$ and the dispersion of log household labour efficiency $\sigma_\theta = 0.50$ are set to match the LCFS evidence on income elasticities and the UK income distribution, so that the demand system firms face is disciplined by household data; the mean log efficiency $\mu_\theta = 2.865$ is a normalisation of the income level. $(\mu_\theta, \sigma_\theta)$ parameterise the log-normal distribution of household labour efficiency: θ_k is the mean of quintile k of a log-normal with these parameters, generating expenditure $E_k = w\theta_k\bar{h} + \bar{\Pi}$, with hours normalised to $\bar{h} = 1$ in the steady state.

The discount factor $\beta = 0.96$ implies a 4% annual real interest rate. Firms' productivity is drawn from G_j, log-normal with mean zero and standard deviation $\sigma_z = 0.50$ in both sectors. The entry elasticity is set to $\nu_e = 2$, which governs only the response of entry along transitions, since entry is exactly free in the steady state (26).

The stationary separation rates λ_k are set to fall with income, so that the rate of the bottom quintile is 1.8 times that of the top, at the lower end of the household panel ratios of Section 4.3. Since the panel does not record who supplies an item, it sets the ratio of the rates but not their level. Their mean of 0.15 is close to the annual rate at which households in the US Consumer Expenditure Survey discontinue an insurance policy (U.S. Bureau of Labor Statistics, 2020). The elasticity of the separation rate to spending per relationship is set to $\kappa = 0.3$, consistent with the Understanding Society estimates of Section 4.3.

[25]The model is solved in JAX and runs on a single NVIDIA GeForce RTX 4090 GPU. The firm distribution over $(a, n, z, \boldsymbol{\pi})$ is cast as a single tensor and the Bellman sweep, distribution updating and aggregation can then be vectorised across the entire state space. A single solve takes 20 seconds which makes the simulated method of moments and transition solution feasible.

Table (4) Externally Set Parameters

Parameter	Description	Value	Source
σ	Cross-sector CES elasticity	0.5	Comin et al. (2021)
ε	Within-sector elasticity	3.5	40% markup
γ_1	Luxury non-homotheticity	1.22	LCFS GMM estimates
ζ_0	Necessity demand shifter	0.80	Normalisation
ϵ_h	Frisch elasticity of hours	1	Labour supply by income
J	Number of sectors	2	Necessity + luxury
σ_θ	S.d. of log household efficiency	0.50	LCFS income dispersion
μ_θ	Mean log household efficiency	2.865	Income normalisation
K	Income bins (quintiles)	5	UK income distribution
β	Discount factor	0.96	Standard
(μ_z, σ_z)	Firm entry productivity (log-normal)	(0, 0.50)	Normalisation; firm-size dispersion
ν_e	Entry elasticity	2	Transitions only
$\underline{\psi}$	Floor of the firm exit rate	0.05	Set
$\bar{a}$	Firm ages	45	Certain exit at the last age
λ_k	Stationary separation rates	$0.197 \to 0.109$	Household panels; Understanding Society
κ	Separation elasticity to spending	0.3	Understanding Society

5.2.2 Internally Calibrated Parameters

Eight parameters are calibrated to the firm moments. Table 5 presents the calibrated values. The resulting stationary equilibrium has a wage of $w = 1.58$ and a luxury price index of $P_1 = 0.81$.

Table (5) Internally Calibrated Parameters

Parameter	Description	Value
ζ_1	Luxury demand shifter	0.0154
ψ_0	Exit rate at entry	0.1965
ψ_1	Fall in the exit rate per year of age	0.0086
f_0	Entry cost (necessity)	20.19
f_1	Entry cost (luxury)	8.43 ($= 0.417 \times f_0$)
ν	Marketing cost curvature	5.71
α	Marketing returns	0.68
$\bar{n}_0$	Entrants' mean customer base (relationships)	0.54

The demand shifter ζ_1 is calibrated to match the aggregate luxury share. Average firm size is pinned down by the level of entry costs, relative sector masses by the ratio of entry costs, and the exit rate and its fall with age by the survival of firms to ages five and fifteen. With α constant across ages, the marketing technology enters the firm's problem only through the ratio $\frac{\nu}{\alpha}$, which the life-cycle profile of marketing intensity pins down. With $\nu = 5.71$ marketing costs are sharply convex, $c(x) = wx^{5.71}$, and $\alpha = 0.68$ makes the acquisition technology concave in effort, so the concavity requirement $\nu > \alpha$ holds with room. The pool from which firms acquire (20) is poorer than the population in both sectors, only slightly in the luxury sector, where 21 per cent of a firm's contacts reach the bottom income quintile

and 19 per cent the top, and by more in the necessity sector, at 29 and 12 per cent.

5.2.3 Model Fit

Table 6 reports the fit on the moments of the objective at the steady state, with the weight each carries.

Table (6) Model Fit

Moment	Data	Model
Average firm size (employees)	8.0	8.4
Marketing/revenue, mean over firms	3.5%	2.9%
Marketing/revenue, age 1–2	4.9%	4.2%
Marketing/revenue, age 3–4	3.9%	3.5%
Marketing/revenue, age 5–9	3.2%	2.4%
Marketing/revenue, age 10–15	2.2%	2.1%
Marketing/revenue, age 16+	2.4%	1.8%
Survival to age five	41.5%	37.2%
Survival to age fifteen	10.9%	11.0%
Event study 2010, exposure × young[26]	−8.9 (4.8)	−7.1
Event study 2010, exposure	−0.2 (2.3)	−1.3
Luxury expenditure share	42%	41.1%

The model reproduces the decline in marketing intensity with firm age, though not quite the level for the youngest firms, at 4.2 against 5.0 per cent, so that the mean intensity across firms is also somewhat below its target. The model matches average firm size and the luxury expenditure share well, and also the share of entrants that survive to age fifteen, while it understates survival to age five somewhat. The model's luxury income elasticity, averaged over quintiles with spending weights, is 1.36, slightly above the LCFS estimates instrumented by income. Finally the event study coefficients for 2010 lie within one standard error of the estimates. The marketing of young firms falls by 7.1% per standard deviation of exposure relative to that of older firms, against 8.7% in the data.

Figure 7 illustrates the fit. Panel (a) shows the declining marketing intensity over the firm life cycle. Marketing rises with the firm's age, since an older firm is less likely to exit, but revenue grows faster, so intensity falls with age. Panel (b) shows the non-homothetic expenditure pattern: the luxury budget share rises from 0.30 in the bottom quintile to 0.48 in the top, which is the Engel-curve variation that transmits income shocks differentially across sectors. Panel (c) plots the number of relationships a firm holds by age in each sector.

6 The Incidence of Austerity and Young-Firm Scarring

This section uses the calibrated model to study what the consolidation did to firms and how much of that depended on which households bore the cuts. Section 6.1 details the counterfactuals. Section 6.2 compares the event study from a simulated panel of firms in the

[26]The two coefficients are those of the event study of Section 3.4 for 2010, in percent of marketing per standard deviation of exposure, with standard errors in brackets. The first is the response of firms aged up to two years relative to older firms and the second the response of firms aged three and over.

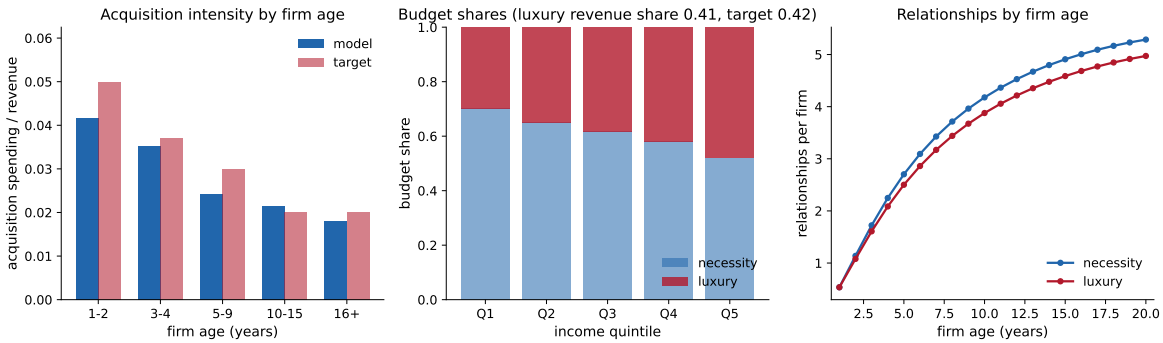


Figure (7) Calibration fit. Panel (a): acquisition spending over revenue by firm-age bin, the marketing intensity of Table 6, model against the ABS targets. Panel (b): budget shares by income quintile under the non-homothetic CES preferences of Section 4.2, at the calibration; the luxury revenue share is 0.41 against the target of 0.42. Panel (c): relationships per firm by firm age in the two sectors.

model against the event study in the data, and shows why young firms were more exposed and how much all firms lost. Section 6.3 holds the spending withdrawn fixed and alters the households that it is withdrawn from, and Section 6.4 discusses the aggregate cost of each policy.

6.1 The experiment

I implement the consolidation as a sequence of cuts to household spending. In period t a design lowers the spending of each household of type k by the proportion $\tau_t \ell_k$,

$$E_{k,t} \;=\; E_k\big(1-\tau_t\ell_k\big), \qquad \sum_k E_k\ell_k \;=\; \sum_k E_k, \tag{33}$$

where τ_t is the spending withdrawn as a share of aggregate spending and ℓ_k is the incidence of the design, normalised so that every design withdraws the same spending. In the notation of (7) the transfer to a household of type θ is $\tau_t(\theta) = -\tau_t\,\ell(\theta)\,E(\theta)$, with $\ell(\theta)$ the incidence at its quintile. The 2010 design reconstructs the incidence of the reforms from the loss intensities of De Agostini et al. (2018). It withdraws 21 per cent of the spending of the bottom quintile in 2010, 13 per cent of that of the second quintile, and 0.4 per cent or less of that of the top two quintiles. The withdrawal is 5.0 per cent of aggregate spending in 2010 and fades by a tenth a year, so that by 2019 it is about two fifths of its initial size.[27]

I compare the 2010 design with two others. A flat design withdraws the same share of every household's expenditure. A progressive design is the mirror image of the 2010 one, with its losses at the top of the distribution.

Along a transition households also adjust their hours. I give the disutility of labour in (7) the form of Greenwood et al. (1988), $\xi_\theta(h) = b(\theta)\,h^{1+1/\epsilon_h}/(1+1/\epsilon_h)$ with $u(C,\xi) =$

[27] In the notation of (33) the withdrawal is $\tau_t = \tau \times 0.9^t$ from 2010. The reforms themselves came into effect between 2010 and 2015. All transition paths are computed from the full nonlinear system.

$u\big(\mathcal{E}(C;\mathbf{p}(\theta))-\xi\big)$, so that a household of type θ solves

$$\begin{aligned}\max_{\{c_{j\omega,t}(\theta),\,h_t(\theta)\}}\ &\sum_{t\geq 0}\beta^t\,u\Big(\mathcal{E}\big(C_t(\theta);\mathbf{p}(\theta)\big)-b(\theta)\,\frac{h_t(\theta)^{1+1/\epsilon_h}}{1+1/\epsilon_h}\Big)\\ \text{s.t.}\ &\sum_{j=1}^{J}\int_{\mathcal{M}_{j,t}(\theta)}p_{j\omega,t}\,c_{j\omega,t}(\theta)\,d\omega\ =\ w_t\,\theta\,h_t(\theta)+\Pi_t-\tau_t\ell(\theta)E(\theta),\end{aligned}\tag{34}$$

period by period given its sets, where $\tau_t\ell(\theta)E(\theta)$ is the withdrawal of (33), and consumption is valued in pounds at the household's stationary own indices $p(\theta)$. The Frisch elasticity ϵ_h is set to one for every type (Table 4), and $b(\theta)$ so that every type works $\bar{h}$ hours in the stationary equilibrium. Given its hours this is the static problem (7) of Section 4.2, so the stationary equilibrium of Section 4.5 is unchanged. The first-order condition for hours equates the disutility of an hour to the wage valued at the household's own prices, so that, to first order,

$$\frac{h_t(\theta)}{\bar{h}}\ =\ \Big[\frac{w_t}{w}\prod_j\Big(\frac{p_{j,t}(\theta)}{p_j(\theta)}\Big)^{-s_j(\theta)\eta_j(\theta)}\Big]^{\epsilon_h},\tag{35}$$

where the weights are the marginal budget shares $s_j(\theta)\eta_j(\theta)$. To first order, hours do not depend on what the household spends. A household works fewer hours only when it loses suppliers, which raises its own indices by (14), and since households hold the sets they inherited at impact, hours first move in the second period. The labour income a household loses lowers its spending below (33) by $w\theta(\bar{h}-h_t(\theta))$, and the rebate moves it by $\Pi_t-\bar{\Pi}$. The labour that firms stop using for marketing and entry is idle (Section 4.5), and its lost wages equal, pound for pound, the profits firms save.

6.2 Model Validation

The evidence of Section 3 is an event study, so to validate the model, I run its regression on a panel of firms simulated from the model under the 2010 design. An authority with exposure T_d faces the national design scaled by $T_d/\bar{T}$, where $\bar{T}$ is the national loss of £470 per working-age adult in Beatty and Fothergill (2014) and the standard deviation of exposure is the £121 of Section 3.3.[28] As in Section 3.4, young firms in this subsection are those aged up to two years in each year. Figure 8 sets the model's coefficients against the estimates. In 2010 young firms in the model cut their marketing relative to older firms by 8.7 log points per standard deviation of exposure, against an estimate of 8.9. The model therefore produces almost all of the estimated cut in the first year, which is a targeted moment of the calibration (Section 5.2), and about one half in the two years after, which are not targeted. In 2011 and 2012 the model's coefficients are below the estimates, but inside their 95 per cent intervals.

A firm's revenue changes, to first order and as shown in (5), by the revenue-weighted average of its customers' spending cuts on its sector. As a result, the firms hit hardest are those that are most exposed to the households most affected under a given design. By the first

[28]The simulated panel has firms of both sectors in 340 local authorities, whose exposures are drawn from a distribution with the mean, standard deviation and range of the measure in Section 3.3. The regression is that of Section 3.4. Nothing national enters households' expenditure in the model, so the local response is the aggregate one.

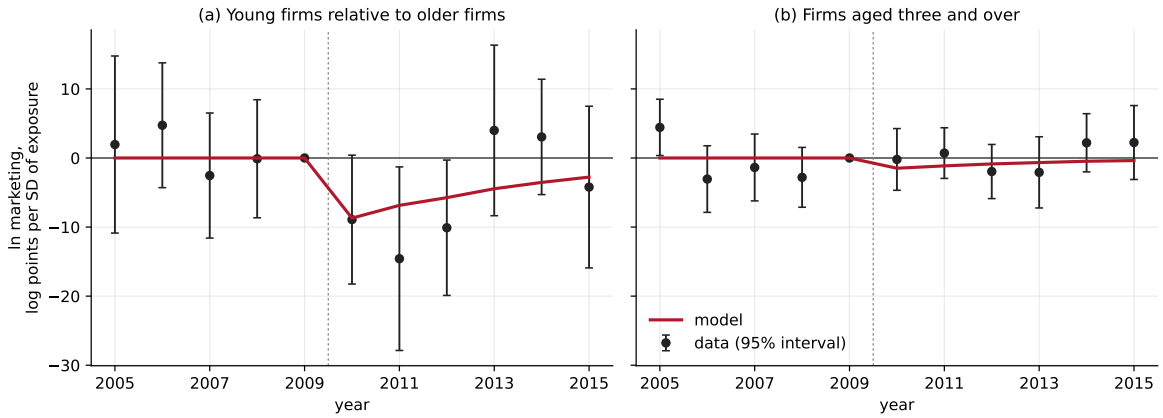


Figure (8) The event study of Section 3.4 and the same regression on the simulated panel, firms of both sectors. Log marketing, coefficients per standard deviation of exposure, base year 2009. The points are the estimates with their 95 per cent intervals and the line is the model under the 2010 design.

restriction of Assumption 1, lower income households end their relationships with suppliers more often, so, as in Section 4.4, the average relationship that survives in an older firm is richer than those of a young firm. Panel (a) of Figure 9 shows the result in the stationary equilibrium. The two poorest quintiles account for 29 per cent of the revenue of a firm in its first year, and for less the older the firm. The same two quintiles bear 81 per cent of the spending the 2010 design withdraws. Panel (b) shows the impact term of (3), revenue in 2010 by the age of the firm, and the fall is largest for the youngest firms, at 5.8 per cent in their first year. As in the empirical work I consider firms as young if they are two years old or younger.

Young firms lose 5.6 per cent of their revenue in 2010 and old firms 4.7 per cent. Writing the cut in sector spending as the common cut plus the part due to the sector's income elasticity, the decomposition (6) assigns about seven tenths of the gap in the luxury sector to the exposure of young firms' customers to the cut and the rest to their higher income elasticities. The dashed lines repeat the experiment at a stationary equilibrium in which every quintile ends its relationships at the same rate. Young and old firms then hold customer bases of the same composition, and revenue falls by 5.0 per cent at every age, which is the fall in aggregate spending.

Figure 10 follows the two groups over time. Revenue recovers as the cut fades, though for young firms only from 2011. Over 2010 to 2019 young firms lose 6.1 per cent of the revenue of stationary firms of the same ages and old firms 4.2 per cent, and both groups are still below in 2020. Customer capital, which is the number of customer relationships firms hold, does not follow spending. A cut in a household's spending raises the rate at which its relationships end.[29] In 2010 the separation rate of the bottom quintile rises by 12 per cent in the luxury sector and by half as much in the necessity sector, while the rates of the top two quintiles do not move. The customer capital of all firms, including those that enter after 2010, therefore goes on falling for six years, to 4.5 per cent below its stationary level in 2016. The loss comes soonest for young firms, whose customer capital is 4.3 per cent below that of stationary firms

[29]The elasticity of the separation rate to a household's spending per relationship is set to 0.3 (Section 5.2).

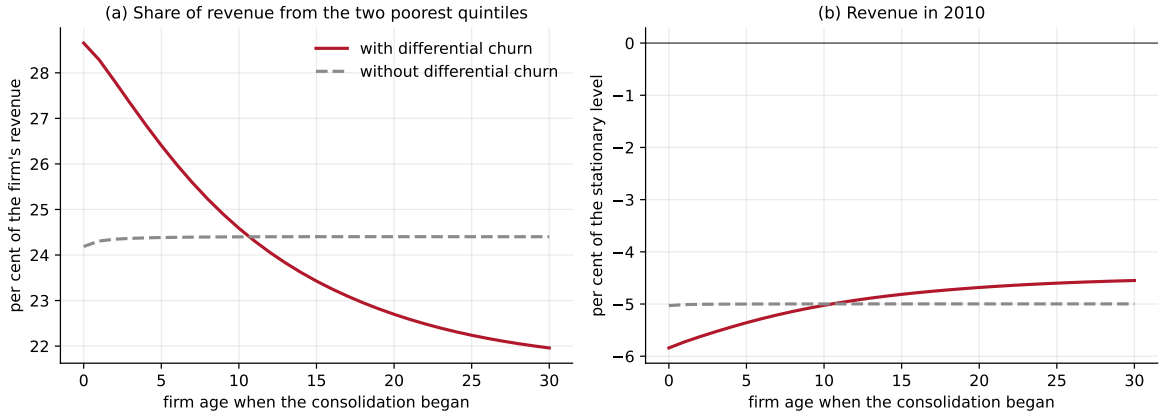


Figure (9) The customer base and the loss of revenue by firm age under the 2010 design, both sectors. (a) The share of a firm's revenue that comes from the two poorest quintiles in the stationary equilibrium. (b) Revenue in 2010, per cent of the stationary level. The dashed lines are the same objects without differential churn.

of the same ages in 2012, more than twice the trough of old firms, which comes in 2015. The fall in spending is itself persistent, but customer capital returns even more slowly.

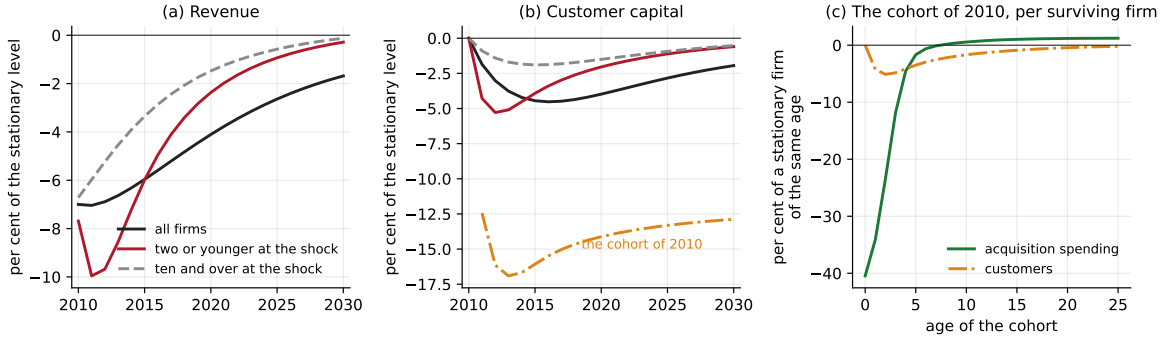


Figure (10) Firms under the 2010 design, 2010 to 2030, both sectors. (a) Revenue and (b) customer capital, per cent of the stationary level; each group is set against stationary firms of the same ages, the revenue of all firms is household spending, and panel (b) adds the cohort of 2010. (c) The cohort of 2010 per surviving firm and by its age: acquisition spending and customers, per cent of a stationary firm of the same age.

The largest shortfall is that of the first cohort whose entry responds to the consolidation, which I call the cohort of 2010. Entry falls with the value of entering, so the cohort has 12 per cent fewer firms than a stationary cohort, and its firms are also smaller. Panel (c) follows them per surviving firm. They spend 40 per cent less on acquisition in their first year than a stationary entrant, and from age seven have almost returned to the amount the same firm would spend in the steady state. The number of customers they serve reflects this decline in acquisition expenditure, which is lower from birth and remains lower even at age 15.

This reflects that, as in Sedláček and Sterk (2017), conditions at entry stay with a cohort. This cohort is the group primarily affected by the churn trap, since it draws its first customers from the pool the cut affects the most. Note that the trap is a shortfall that the cohort carries along the transition rather than a second equilibrium, since the economy in time returns to the stationary equilibrium it started from.

Figure 11 compares these paths with district data. In the more exposed local authorities

firm births per working-age adult fall relative to the rest of their region in 2010 and stay lower until 2012. [30] In the data five-year survival is no lower in the more exposed authorities, as the model assumes.

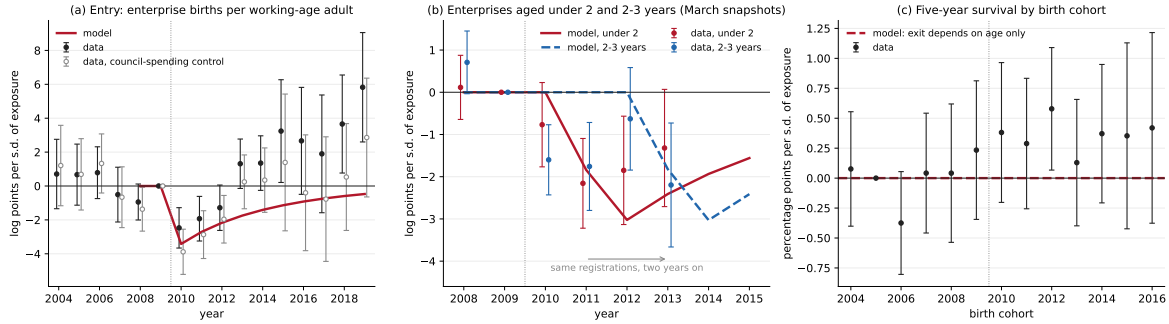


Figure (11) Local authority firm data against the model. (a) Births per working age adult and the equivalent model result. (b) Firms aged under two and aged two to three years old. (c) Five year survival by birth cohort. Source: ONS business demography and business register by local authority

6.3 Who bears the cut

To see how much of this depends on which households bear the cut, I compare the 2010 design with the flat and progressive designs. The three designs withdraw the same pounds in every year, the difference being which households are reducing spending. Aggregate revenue, equal to measured output here by construction, falls in 2010 by those pounds and by the change in the rebate.[31] The change in the rebate differs across designs with what firms save on acquisition and entry, so the impact fall in revenue is not the same under each.

Table 7 and Figure 12 report the results.

Table (7) Losses over 2010 to 2019 under the three designs

	Revenue			Customer capital		
	2010	Flat	Progressive	2010	Flat	Progressive
With differential churn						
Young firms	6.11	5.56	5.31	3.02	2.35	2.08
Old firms	4.17	4.42	4.60	1.40	0.92	0.77
All firms	5.99	5.79	5.72	3.49	2.89	2.66
Without differential churn						
Young firms	5.36	5.32	5.32	1.78	1.80	1.87
Old firms	4.07	4.09	4.12	0.28	0.27	0.30
All firms	5.58	5.59	5.61	2.30	2.36	2.42

Notes. Loss over 2010 to 2019 as per cent of the stationary level of firms of the same ages, both sectors, under the 2010, flat and progressive designs, each scaled to withdraw the same spending in every year. Young firms are those younger than two when the consolidation began and old firms those of ten and over; all firms include those that enter after 2010. Customer capital is the number of customer relationships. Each block is a separate stationary equilibrium and set of transitions. In the second block every quintile ends its relationships at the same rate, and that rate does not respond to spending.

[30] The stock of enterprises aged two to three is already lower in the more exposed authorities in March 2010, before the June 2010 Budget.

[31] In 2010 households keep their suppliers and their hours, so spending falls by the pounds withdrawn and by the change in the rebate, which is the profit lost on the revenue less the acquisition and entry costs saved (Section 6.1).

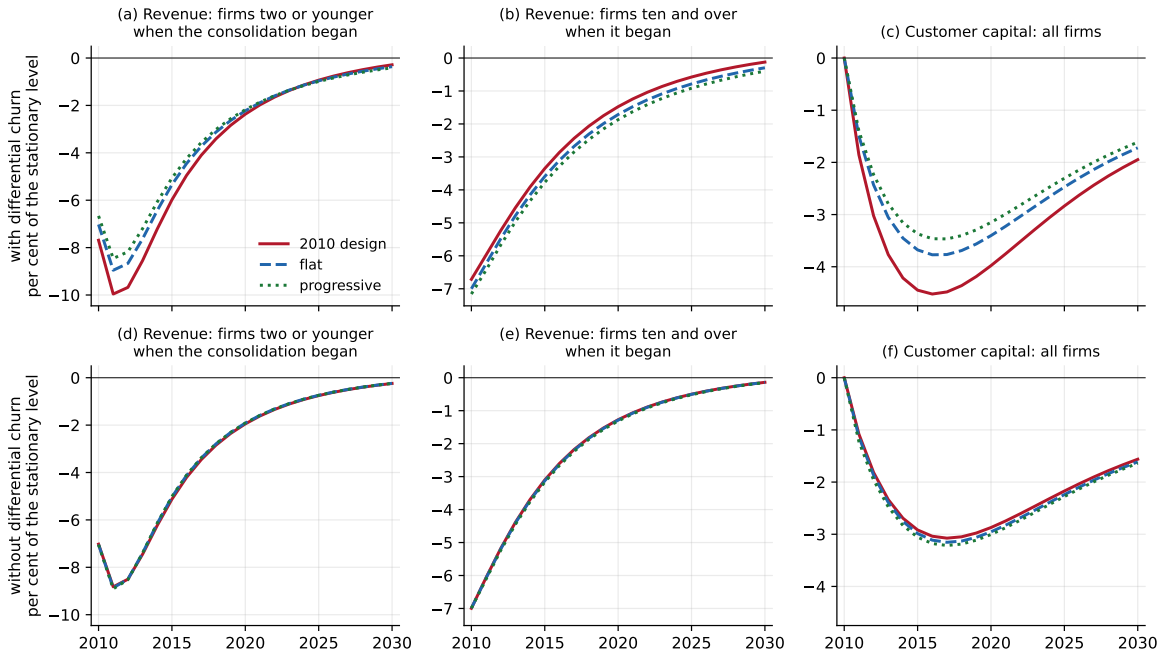


Figure (12) The three designs when they withdraw the same spending, 2010 to 2030, per cent of the stationary level, both sectors. Top row, with differential churn: the revenue of young firms and of old firms, each against stationary firms of the same ages, and the customer capital of all firms. Bottom row: the same without differential churn.

Each design leads to a different loss across given firms. Under the flat design young firms lose less revenue over 2010 to 2019 than under the 2010 design, and old firms more (Table 7). Under the flat design, which cuts every household's spending in the same proportion, young firms still lose more revenue than old firms, which is the second way non-homothetic preferences enter in Section 4.1. Intuitively, a more regressive design moves the relative revenue loss from older to younger firms. This result is even more pronounced in customer capital, directly because the households more exposed to the 2010 design are those households who separate most readily from existing suppliers. Over 2010 to 2019 the 2010 design costs young firms about half as much again of their customer capital as the progressive design and about three tenths more than the flat design. A flat design would therefore have avoided about a fifth of young firms' loss and a progressive design about three tenths. The gap widens through the decade as the cohorts that enter during the consolidation carry their smaller customer bases forward. Young firms lose more customer capital than old firms in each design.

The second block of the table removes the first two restrictions of Assumption 1, the higher rate of poorer households and the response of the rate to spending, and I refer to it as the economy without differential churn. When every quintile ends its relationships at the same rate, and that rate does not respond to spending, there is no role for the distributional design of the policy: all firms lose almost identical revenue and customer capital under each design, and young firms lose no more under the 2010 design than under the others. Young firms still lose more than old firms over the decade, since they are still building their bases and so lose more when every firm acquires less, but the gap is the same under each design. Young firms lose more than under a flat design only when a shock falls more on low-income households,

as the 2010 design did. Non-homothetic preferences add to that loss in the luxury sector, where the spending of those households is the most income-elastic, as the decomposition in Section 6.2 shows for the 2010 design.

6.4 The aggregate cost

The three designs withdraw the same spending, but they differ in how much further output falls, since a household that loses suppliers works fewer hours, and since the profits firms rebate to households fall with their revenue.Firms also stop using labour, in production as well as in acquisition and entry, and that labour is idle, but its wages are the pounds withdrawn themselves (Section 6.1), so they add nothing to the loss beyond them. Per pound withdrawn over 2010 to 2019, output falls by 9pence more under the 2010 design than under the progressive design and by 6 pence more than under the flat design. The 2010 design costs more since it withdraws its pounds from the households whose relationships end most readily, so that their own price indices rise most and they work fewer hours, which accounts for about seven tenths of the gap against the progressive design. The rest is the rebate from firms profits, which fall further under it as revenue falls further. Without differential churn the three designs cost the same. Scaled to the welfare cuts of 2010 to 2019, the 2010 design cost about £11bn more output than a progressive design withdrawing the same money would have, and about £8bn more than a flat design, and the difference goes on increasing for as long as the cuts remain in place (Figure 13).[32]

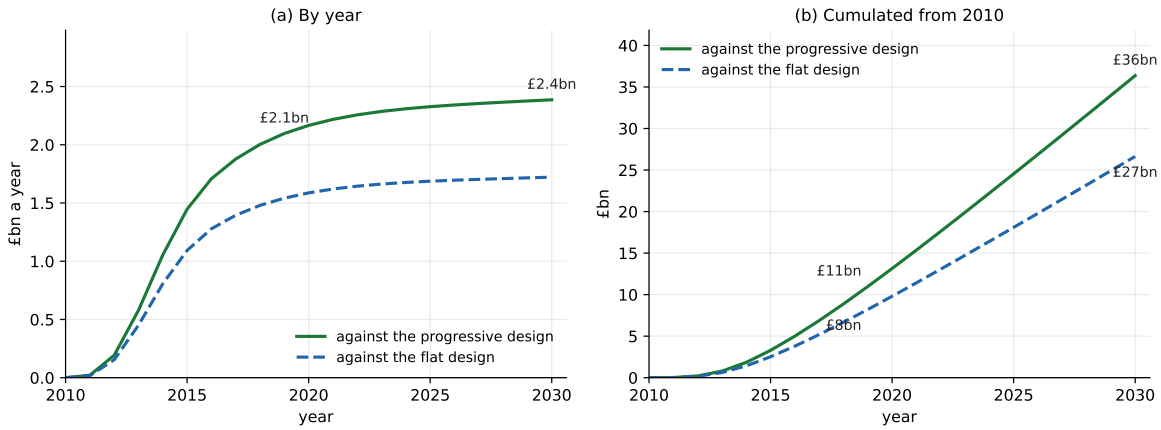


Figure (13) The extra output lost under the 2010 design. Output lost beyond the pounds withdrawn under the 2010 design less the same loss under the progressive design and under the flat design, in £bn at 2010 prices, on the calendar of the reforms (in full effect from 2015/16 and held at £19bn a year thereafter): (a) by year; (b) cumulated from 2010.

Output does not record the suppliers that households lose. Households have non-homothetic preferences and value variety across the suppliers they hold (11), so a household whose relationships end pays more, at its own price index, for what it still buys (14). This is the term through which a household values a relationship in (14). Valued in pounds over

[32]The simulations withdraw spending along a path that fades by a tenth a year. The figures in the text instead apply the model to the path of the cuts in the OBR's policy costings, in full effect from 2015/16 and held at £19bn a year thereafter, at 2010 prices, taking the cost as linear in the size of the cut. On the model's own path the extra cost against the progressive design is also £11bn by 2019, but £18bn rather than £36bn by 2030, since that path fades.

ten years, this loss adds a further 5 pence per pound withdrawn to the extra cost of the 2010 design against the progressive design, and 4 pence against the flat design. Figure 14 shows who bears the loss beyond the pounds withdrawn. It falls on the households whose relationships ended: in 2015 the own price index of the bottom quintile is 2.5 per cent above its stationary level under the 2010 design, against 1.4 per cent under the flat design, and the same households work 4.0 per cent fewer hours. Beyond the pounds withdrawn from them, households in the poorest quintile therefore lose about half as much again under the 2010 design as under the progressive design, while households in the top two quintiles lose the same under every design. A productivity index does not record the loss either, since firm productivity is drawn at entry and is independent of age, so the model is consistent with the finding that firm-level TFP shows no differential response (Section 3.5).

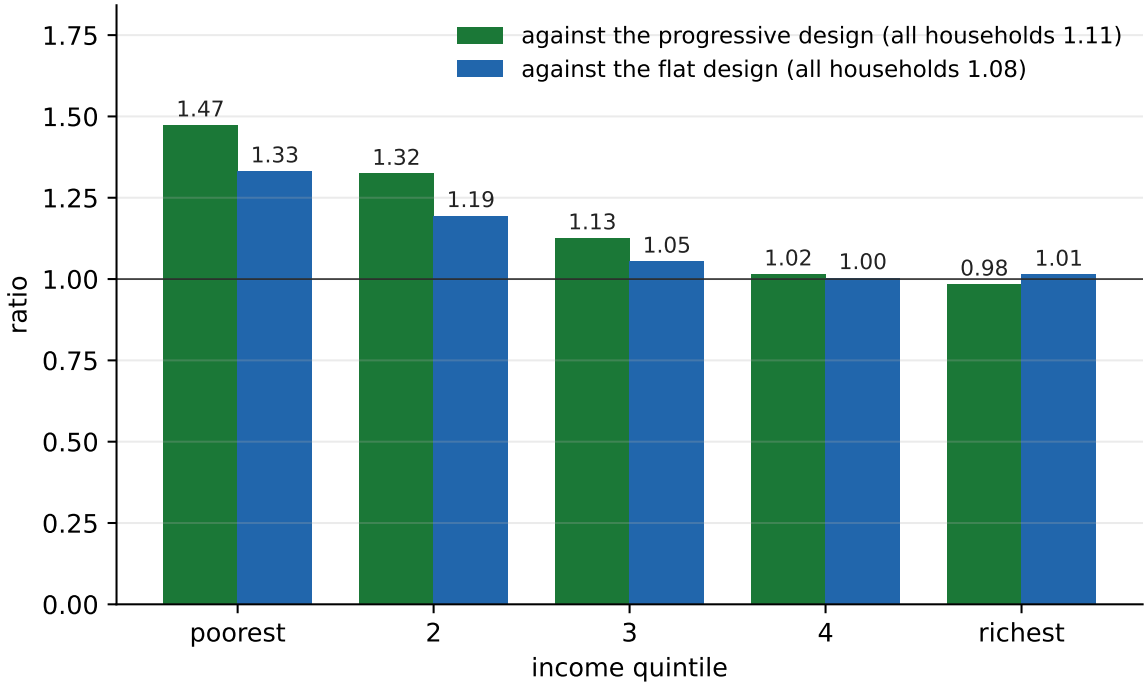


Figure (14) Who bears the extra loss of the 2010 design. The loss per household over 2010 to 2019 beyond the pounds withdrawn from it, under the 2010 design relative to the progressive design and relative to the flat design, by income quintile; the legend gives the same ratio for all households together.

The loss beyond the pounds withdrawn outlasts the cut. Under the 2010 design it peaks in 2015 and halves by 2027, and in 2030 it is nearly three times the spending withdrawn in that year. Without differential churn it is the same under every design.

7 Conclusion

What are the effects of a fiscal consolidation on firm growth, and do they depend on its distributional incidence? I study the UK's 2010–2015 austerity programme, whose welfare cuts fell mostly on low-income households. Young firms in more exposed local authorities cut customer acquisition by around 15% per standard deviation relative to less exposed peers, while their employment, capacity utilisation and markups were unaffected. Young firms' value

added fell and remained depressed through 2015, even after spending on customer acquisition recovered, suggesting a durable loss of customer base. Old firms showed no comparable adjustment. To explain why regressive shocks disproportionately scar young firms, I build a general equilibrium model in which firms accumulate customers from a heterogeneous-income population. Low-income customers churn faster, so young firms inherit a poor-biased customer base. This creates a "churn trap": young firms depend most on the households a regressive consolidation hits hardest. The effects are especially strong for firms that operate in 'luxury' sectors, whose demand is the most income-elastic. In the calibrated model, the 2010 austerity programme destroys about half as much again of young firms' customer capital as a progressive consolidation of the same size, and costs about nine pence more output per pound withdrawn than a progressive design, £11bn over the decade; the cohort entering during the shock remains persistently smaller than in the absence of the consolidation. A flat design would have avoided about a fifth of young firms' loss of customer capital, and a progressive design about three tenths.

References

Ackerberg, D. A., Caves, K., & Frazer, G. (2015). Identification Properties of Recent Production Function Estimators. *Econometrica*, *83*(6), 2411–2451.

Afrouzi, H., Drenik, A., & Kim, R. (2023). *Concentration, Market Power, and Misallocation: The Role of Endogenous Customer Acquisition* (tech. rep.). National Bureau of Economic Research.

Aguiar, M., & Hurst, E. (2007). Life-Cycle Prices and Production. *American Economic Review*, *97*(5), 1533–1559.

Anzoategui, D., Comin, D., Gertler, M., & Martinez, J. (2019). Endogenous Technology Adoption and R&D as Sources of Business Cycle Persistence. *American Economic Journal: Macroeconomics*, *11*(3), 67–110.

Aquilante, T., Chowla, S., Dacic, N., Haldane, A., Masolo, R., Schneider, P., Seneca, M., & Tatomir, S. (2019). *Market Power and Monetary Policy* (Staff Working Paper No. 798). Bank of England.

Argente, D., Fitzgerald, D., Moreira, S., & Priolo, A. (2025). How Do Entrants Build Market Share? The Role of Demand Frictions. *American Economic Review: Insights*, *7*(2), 160–176.

Arkolakis, C. (2010). Market Penetration Costs and the New Consumers Margin in International Trade. *Journal of Political Economy*, *118*(6), 1151–1199.

Auclert, A. (2019). Monetary Policy and the Redistribution Channel. *American Economic Review*, *109*(6), 2333–2367.

Auer, R., Burstein, A., Lein, S., & Vogel, J. (2024). Unequal Expenditure Switching: Evidence from Switzerland. *Review of Economic Studies*, *91*(5), 2572–2603.

Bai, Y., Rios-Rull, J.-V., & Storesletten, K. (2026). Demand Shocks as Technology Shocks. *Review of Economic Studies*, *93*(2), 798–832.

Beatty, C., & Fothergill, S. (2014). The Local and Regional Impact of the UK's Welfare Reforms. *Cambridge Journal of Regions, Economy and Society*, *7*(1), 63–79.

Benigno, G., & Fornaro, L. (2018). Stagnation Traps. *The Review of Economic Studies*, *85*(3), 1425–1470.

Boppart, T. (2014). Structural Change and the Kaldor Facts in a Growth Model with Relative Price Effects and Non-Gorman Preferences. *Econometrica*, *82*(6), 2167–2196.

Bornstein, G. (2025). *Entry and Profits in an Aging Economy: The Role of Consumer Inertia* (tech. rep.). National Bureau of Economic Research.

Burnside, C., Eichenbaum, M., & Rebelo, S. (1995). Capital Utilization and Returns to Scale. *NBER Macroeconomics Annual*, *10*, 67–110.

Cavenaile, L., Celik, M. A., Perla, J., & Roldan-Blanco, P. (2026). *A Theory of Dynamic Product Awareness and Targeted Advertising*. BSE, Barcelona School of Economics.

Cavenaile, L., Celik, M. A., Roldan-Blanco, P., & Tian, X. (2025). Style over Substance? Advertising, Innovation, and Endogenous Market Structure. *Journal of Monetary Economics*, *149*, 103683.

Comin, D., Lashkari, D., & Mestieri, M. (2021). Structural Change with Long-Run Income and Price Effects. *Econometrica*, *89*(1), 311–374.

De Agostini, P., Hills, J., & Sutherland, H. (2018). Were We Really All in It Together? The Distributional Effects of the 2010–15 UK Coalition Government's Tax-Benefit Policy Changes. *Social Policy & Administration*, *52*(5), 929–949.

De Loecker, J., & Eeckhout, J. (2018). *Global Market Power* (Working Paper No. 24768). National Bureau of Economic Research.

De Loecker, J., & Warzynski, F. (2012). Markups and Firm-Level Export Status. *American Economic Review*, *102*(6), 2437–2471.

Decker, R. A., Haltiwanger, J., Jarmin, R. S., & Miranda, J. (2017). Declining Dynamism, Allocative Efficiency, and the Productivity Slowdown. *American Economic Review*, *107*(5), 322–326.

Drozd, L. A., & Nosal, J. B. (2012). Understanding International Prices: Customers as Capital. *American Economic Review*, *102*(1), 364–395.

Einav, L., Klenow, P. J., Levin, J. D., & Murciano-Goroff, R. (2021). *Customers and Retail Growth* (tech. rep.). National Bureau of Economic Research.

Fatás, A., & Summers, L. H. (2018). The Permanent Effects of Fiscal Consolidations. *Journal of International Economics*, *112*, 238–250.

Fetzer, T. (2019). Did Austerity Cause Brexit? *American Economic Review*, *109*(11), 3849–3886.

Fitzgerald, D., Haller, S., & Yedid-Levi, Y. (2024). How Exporters Grow. *Review of Economic Studies*, *91*(4), 2276–2306.

Foster, L., Haltiwanger, J., & Syverson, C. (2016). The Slow Growth of New Plants: Learning about Demand? *Economica*, *83*(329), 91–129.

Goldsmith-Pinkham, P., Sorkin, I., & Swift, H. (2020). Bartik Instruments: What, When, Why, and How. *American Economic Review*, *110*(8), 2586–2624.

Gourio, F., & Rudanko, L. (2014). Customer Capital. *Review of Economic Studies*, *81*(3), 1102–1136.

Greenwood, J., Hercowitz, Z., & Huffman, G. W. (1988). Investment, Capacity Utilization, and the Real Business Cycle. *The American Economic Review*, 402–417.

Handbury, J. (2021). Are Poor Cities Cheap for Everyone? Non-Homotheticity and the Cost of Living across US Cities. *Econometrica*, *89*(6), 2679–2715.

Hottman, C. J., Redding, S. J., & Weinstein, D. E. (2016). Quantifying the Sources of Firm Heterogeneity. *The Quarterly Journal of Economics*, *131*(3), 1291–1364.

House, C. L., Proebsting, C., & Tesar, L. L. (2020). Austerity in the Aftermath of the Great Recession. *Journal of Monetary Economics*, *115*, 37–63.

Hubmer, J., & Nord, L. (2025). Investment in Demand and Dynamic Competition for Customers. *Available at SSRN 5491506*.

Ignaszak, M., & Sedláček, P. (2026). Customer Acquisition, Business Dynamism, and Aggregate Growth. *Review of Economic Studies*.

Jaimovich, N., Rebelo, S., & Wong, A. (2019). Trading Down and the Business Cycle. *Journal of Monetary Economics*, *102*, 96–121.

Kaplan, G., & Menzio, G. (2016). Shopping Externalities and Self-Fulfilling Unemployment Fluctuations. *Journal of Political Economy*, *124*(3), 771–825.

Kariel, J., Savagar, A., & Hwang, K. I. (2024). Aggregate Market Power with Small Firms: Evidence from the UK. *Working Paper*.

Mian, A., & Sufi, A. (2014). What Explains the 2007–2009 Drop in Employment? *Econometrica*, *82*(6), 2197–2223.

Michelacci, C., Paciello, L., & Pozzi, A. (2022). The Extensive Margin of Aggregate Consumption Demand. *The Review of Economic Studies*, *89*(2), 909–947.

Mongey, S., & Waugh, M. E. (2025). *Pricing Inequality* (tech. rep.). National Bureau of Economic Research.

Moreira, S. (2016). Firm Dynamics, Persistent Effects of Entry Conditions, and Business Cycles. *Persistent Effects of Entry Conditions, and Business Cycles (October 1, 2016)*.

Nord, L. (2023). Shopping, Demand Composition, and Equilibrium Prices. *Available at SSRN 4178271*.

Olivi, A., Sterk, V., & Xhani, D. (2023). *Optimal Monetary Policy during a Cost of Living Crisis* (tech. rep.). mimeo, UCL.

Paciello, L., Pozzi, A., & Trachter, N. (2019). Price Dynamics with Customer Markets. *International Economic Review*, *60*(1), 413–446.

Patterson, C. (2023). The Matching Multiplier and the Amplification of Recessions. *American Economic Review*, *113*(4), 982–1012.

Rachel, L. (2021). Leisure-Enhancing Technological Change. *Job Market Paper, London School of Economics*.

Ravn, M., Schmitt-Grohé, S., & Uribe, M. (2006). Deep Habits. *The Review of Economic Studies*, *73*(1), 195–218.

Roldan-Blanco, P., & Gilbukh, S. (2021). Firm Dynamics and Pricing under Customer Capital Accumulation. *Journal of Monetary Economics*, *118*, 99–119.

Sangani, K. (2022). Markups across the Income Distribution: Measurement and Implications. *Available at SSRN 4092068*.

Sedláček, P. (2020). Lost Generations of Firms and Aggregate Labor Market Dynamics. *Journal of Monetary Economics*, *111*, 16–31.

Sedláček, P., & Sterk, V. (2017). The Growth Potential of Startups over the Business Cycle. *American Economic Review*, *107*(10), 3182–3210.

U.S. Bureau of Labor Statistics. (2020). Consumer Expenditure Surveys: Interview Survey Public-Use Microdata, 2010–2019.

A Data

This appendix documents the data behind Section 3, the household evidence on separation rates used in Sections 4.3 and 5.2, and the incidence of the reforms and the district data of Section 6. It describes (i) the eight sources and how they are linked (Appendix A.1); (ii) the construction of the firm sample and the firm-level variables (Appendix A.2); (iii) the classification of sectors into luxuries and necessities (Appendix A.3); (iv) the exposure measure T_d, the reforms behind it and the panel of realised benefit income (Appendix A.4); (v) the incidence of the reforms across households (Appendix A.5); (vi) the household data (Appendix A.6); and (vii) the district data on firm births and survival (Appendix A.7).

A.1 Sources and linkage

I construct a novel dataset by combining a number of sources. The Business Structure Database (BSD) provides quasi-universal[33] yearly snapshots of firms in the UK, recording the date of creation, location, size and industry classification. The Annual Business Survey (ABS) is a size stratified random sample of approximately 60,000 firms per year. The ABS contains extensive financial and balance sheet data, for instance turnover, intermediate good costs, capital expenditure and importantly for this paper expenditure on "advertising, marketing, and promotional costs", which includes advertising campaigns, direct marketing, trade fairs, and related customer-facing activities. The Annual Survey of Hours and Earnings (ASHE) provides the payroll and wage bill data used in the empirical work. The Living Costs and Food Survey (LCFS) is a household survey that contains detailed household expenditure by category together with income sources, including income from social security benefits and which benefits the household receives. I use the LCFS to estimate the non-homothetic demand system used in Section 4. I use the Understanding Society dataset to provide evidence for household separation rates across the business cycle. The ONS Business Demography database is used to calculate the response of entry and exit in local authorities with differing exposure to the austerity shock. The austerity shock is a panel of local authority exposure to the aggregate saving from the 2010 welfare reforms. The BSD, ABS, and ASHE are restricted Office for National Statistics (ONS) microdata, which are accessed in the UK Data Service SecureLab, where I link the local authority exposure to austerity measure to study the effect on firms. The remaining datasets are publicly available and are not accessed in the UKDS SecureLab. Table A.1 summarises the sources and relevant variables used from each.

Firms. The BSD and the ABS are both drawn from the Inter-Departmental Business Register, the ONS register of businesses registered for VAT or PAYE, so the two share the register's reference numbers. The BSD is a set of yearly snapshots of the register, recording for each firm its date of creation, location, size and industry classification, and the ABS is a size-stratified random sample of approximately 60,000 firms per year drawn from the same register. The unit of the firm panel is the business as the ABS surveys it, the reporting unit, which for most firms is the enterprise itself, and the BSD and the ABS are linked on the ONS

[33] The BSD consists of firms that are registered for Value Added Tax and/or Pay As You Earn, covering 99% of firms in the UK.

Table (A.1) Data sources

Source	Unit	Years used	Link	Variables
BSD	Enterprise	2005–2015	Reporting unit reference	Birth date, postcode and local authority, two-digit SIC sector
ABS	Reporting unit; about 60,000 a year	2005–2015	Reporting unit reference	Turnover, gross value added, employment, marketing, capital expenditure, energy, materials, ownership, tradeable sector
ASHE	Employee job; 1 per cent of PAYE jobs	2005–2015	Enterprise reference of the employer	Labour costs
LCFS	Household in a survey quarter	2002–2016	Not linked	Spending by COICOP division, benefit income, household size
Exposure	Local authority (December 2016 boundaries)	Costings 2013; claimant shares 2009–10	Local authority code	T_d; claimant rates and benefit expenditure
Understanding Society; BHPS	Household; pairs of annual interviews	2009–2015; 1996–2008	Not linked	Satellite or cable television, landline, mobile telephone and internet access at home; household income
District statistics	Local authority (April 2023 boundaries)	2004–2019	Local authority code	Enterprise births and five-year survival, population aged 16–64, council service spending

Notes. The BSD, ABS and ASHE are accessed in the UK Data Service SecureLab. The exposure measure applies the 2013 Treasury costings of the ten reforms, which Beatty and Fothergill (2014) also start from, to each authority's own claimant shares, and is described in Appendix A.4; the claimant rates and benefit expenditure are those of Appendices A.4.3 and A.4.4. Understanding Society, the BHPS and the CE are described in Appendix A.6, and the district statistics in Appendix A.7.

reporting unit reference (Appendix A.2). ASHE is a stratified sample of employees drawn from HMRC PAYE records, and can be linked to the BSD through the enterprise reference.

Location. Firms are linked to a local authority through the postcode of its reporting unit, and the exposure measure T_d is then linked to each local authority through local authorities on their 2016 boundaries, which is the level of spatial variation used in the paper. In the UK, each local authority is then a sub unit of a NUTS-1 region (East England, Scotland etc), which is used to set the region-by-year fixed effects of equation (2). Note that Northern Ireland is not contained in the data so the estimation sample consequently consists solely of Great Britain, and the estimation sample covers 346 local authorities as in (Table B.6).

Households and exposure. The household data is sourced from the LCFS, which is a national survey. The household evidence described in Section 3.6 compares benefit recipients against non-recipients across the country. Similarly, the demand-system estimates used in the paper and detailed in Appendix A.3 are estimated at the national level. I construct the exposure to austerity measure by taking the aggregate losses from Beatty and Fothergill (2014) and interacting them with the local authority averages of lagged claimant shares to construct a time series of predicted exposure that predates the announcement of the welfare package. I use Understanding Society to calibrate the separation rates and elasticities used in Sections 4.2 and 5.2 (Appendix B.7).

A.2 Firm sample and variables

The firm level data is drawn from the BSD, the ABS and ASHE, and is accessed through the UK Data Service SecureLab. I clean and prepare the data in the following steps. The initial sample constitutes all enterprises that appear in the BSD between 2005 and 2020, I drop all firms with missing birth dates or industry classifications. I link the BSD to the subsample ABS on the reported unit reference. The ABS is stratified by size and so a firm that is not sampled by the ABS in a given year is excluded from the analysis in that year. Following Kariel et al. (2024), I drop non-market sectors, namely finance and insurance (SIC 64–66), agriculture (SIC 1–3), mining (SIC 5–9), utilities (SIC 35–39), public administration (SIC 84), education and health (SIC 85–88) and real estate (SIC 68). The sample is winsorised the top 0.1% of the distribution of output per worker, marketing intensity, and input factor shares. Firms are then linked to the local authority exposure measure using the postcode of the reporting unit, and matched to the exposure measure. Note that the postcode is no longer included as standard in the ABS, fortunately this project predates it, an alternative linking measure using super output areas is available. For estimation, I restrict the sample to end in 2015 to avoid picking up Brexit related effects. Fetzer (2019) documents a strong association between Beatty and Fothergill's measure and subsequent support for Brexit.

Because ABS financial variables are item-specific, estimation samples differ across outcomes. The output per worker regressions require non-missing GVA and employment (165,314 firm-years), the marketing regressions require non-missing advertising expenditure (171,052), and the regressions for employment, the three composition outcomes, net capital

expenditure and energy in Table 2 have 171,059. The markup and TFP regressions also require the inputs of the production function, which leaves 143,170 and 143,146 firm-years. Before item non-response the linked sample has 13,326 firm-years of young firms and 256,087 of older firms (Table B.6).

The ABS is a rotating sample. The largest firms are surveyed every year, while smaller firms move in and out of the sample, so only firms sampled in consecutive years have lagged values, and in practice these are mostly the larger firms. This is why the production function behind the markup and TFP measures is estimated without lagged instruments (Appendix B.4). Table A.2 defines each variable.

Table (A.2) Firm-level variables

Variable	Definition	Source
Outcomes		
Value added	Log of real gross value added, deflated to 2010 prices with two-digit SIC deflators from the ONS producer price indices	ABS
Output per worker	Log of real gross value added per employee, deflated to 2010 prices with two-digit SIC deflators from the ONS producer price indices	ABS
Marketing expenditure	Advertising, marketing and promotional costs (ABS question 8.5), in logs	ABS
Marketing intensity	Marketing expenditure over total turnover	ABS
Employment	Log number of employees	ABS
Capital expenditure	Net capital expenditure	ABS
Energy expenditure	Expenditure on energy, a proxy for capacity utilisation (Section 3.5)	ABS
Markup	Output elasticity of materials over the materials share of revenue; the elasticity is the materials coefficient of a Cobb–Douglas production function in deflated output, materials, the wage bill and capital, estimated by OLS within each two-digit SIC sector[a]	ABS
TFP	Residual of the same production function	ABS
Wage bill	Labour costs of the firm (Figure 4)	ASHE
Firm characteristics		
Firm age	Years since first appearance in the BSD	BSD
Young firm	Aged two years or less[b]; every other firm is in the comparison group	BSD
Sector	Two-digit SIC 2007	BSD
Luxury sector	Two-digit SIC sector classified as luxury by its income elasticity, estimated from the LCFS[c]	LCFS
Tradeable	Tradeable-sector indicator as defined in the ABS	ABS
UK owned	Indicator for UK ownership	ABS
Concentration	Herfindahl–Hirschman index by region, two-digit SIC and year	ABS
Treatment		
T_d	Anticipated annual loss per working-age adult from the reforms, £; in standard deviations (£121) in equation (2)	Treasury costings; DWP and HMRC

Notes. (a) Section 3.5 describes the markups as following Kariel et al. (2024), who apply the production approach of De Loecker and Warzynski (2012), with the production function estimated by OLS as here; Appendix B.4 sets out the estimation. (b) The definition of Section 3.1; a firm leaves the group in the year it turns three. (c) The classification of Appendix A.3.

A.3 Sector classification

Sectors are classified as luxury or necessity by their income elasticities, estimated from the LCFS, and the resulting classification coincides with that of Olivi et al. (2023). The

classification is made on the twelve COICOP divisions of household spending and carried to firms through their two-digit SIC sector.

Table A.3 lists the divisions in each group, together with the GMM estimates of the non-homothetic CES demand system of Comin et al. (2021) on the LCFS, which Appendix Table B.14 reports in full. Each estimate measures a division's non-homotheticity relative to clothing, the base category, so a value above one means that a division is more income-elastic than clothing. Food and household fuel, the two necessity divisions in the estimation, have estimates of 0.50 and 0.52, while furnishings and recreation have estimates of 1.21 and 1.23, and restaurants and hotels and miscellaneous goods and services of 0.98 and 0.86. Moreover, in each of the four specifications of Appendix Table B.14, food and household fuel lie below every division of the luxury group. Health and communication, the other two necessity divisions, are not part of the estimation, and alcohol and tobacco, transport and education are in neither group.

Firms are classified as luxury or necessity producers by their two-digit SIC 2007 sector, each of which is assigned to the luxury group, the necessity group or neither, and the split regressions of Figure 6 use the firms in sectors assigned to one of the two groups.

Table (A.3) COICOP divisions, sector groups and demand-system estimates

	COICOP division	Group	Estimate relative to clothing
1	Food and non-alcoholic beverages	Necessity	0.50 (0.02)
2	Alcoholic beverages and tobacco	Neither	–
3	Clothing and footwear	Luxury	1 (base)
4	Electricity, gas and other fuels	Necessity	0.52 (0.02)
5	Furnishings and household equipment	Luxury	1.21 (0.05)
6	Health	Necessity	–
7	Transport	Neither	0.90 (0.04)
8	Communication	Necessity	–
9	Recreation and culture	Luxury	1.23 (0.05)
10	Education	Neither	–
11	Restaurants and hotels	Luxury	0.98 (0.04)
12	Miscellaneous goods and services	Luxury	0.86 (0.03)

Notes. Groups as in Figure 5. Division 4 enters both the classification and household spending as electricity, gas and other fuels only, since housing costs are excluded from expenditure (Appendix A.6). Estimates are column (1) of Appendix Table B.14, the GMM estimates on the 2019 survey (3,164 observations), with standard errors in parentheses; – denotes a division that is not part of the estimation.

A.4 The exposure measure

A.4.1 The ten reforms

The welfare reforms studied in this paper are the ten that Beatty and Fothergill (2014) cost for every local authority in Great Britain, listed with their 2013 costings in Appendix Table B.9. Most of them were introduced by the Coalition government, largely through the Welfare Reform Act 2012. Part of the incapacity benefit reform, however, is a Labour measure that predates the 2010 general election and took full effect only in the following years. Once in full effect the ten reforms were expected to save £18.9 billion a year, about £470 for every adult of working age. Beatty and Fothergill allocate each saving on caseloads taken between August 2011 and February 2013, while T_d uses the authority's own shares of 2009–10, and the

two measures have a correlation > 0.95 (Appendix A.4.2).

A.4.2 Construction

The exposure measure T_d of equation (1) applies the Treasury costing of each of the ten reforms, the national saving L_b that Beatty and Fothergill (2014) start from, to the authority's own share $S_{b,d}$ of the reform's caseload in 2009–10, and divides by its working-age population.

A.4.3 Persistence of claimant shares

The claimant shares behind T_d are those of 2009–10, and the spatial distribution they describe was in place well before the reforms. Section 3.3 reports one check of this, which is constructed as follows. First, I regress T_d on district spending per working-age adult on the five main working-age benefits (Housing Benefit, Disability Living Allowance, incapacity benefits, Income Support and Jobseeker's Allowance), taken from the DWP tables of benefit expenditure by local authority. The R^2 is 0.90 on the spending of 2002, 0.92 on that of 2007 and 0.95 on that of 2010.

A.4.4 Realised benefit income

The outcome in Section 3.3 combines two published sources for the fiscal years 2005/06 to 2014/15. DWP benefit expenditure by local authority and benefit (2000/01–2014/15) provides working-age benefit spending: Housing Benefit, which includes pensioner claims, working-age Disability Living Allowance, Employment and Support Allowance, Incapacity Benefit, Income Support, Jobseeker's Allowance, working-age Severe Disablement Allowance, Carer's Allowance and bereavement benefits. HMRC's finalised annual award statistics provide tax credit entitlement by local authority from 2005/06; total entitlement is the average number of benefiting families multiplied by the average annualised award. Expenditure is divided by the working-age population in 2010 and matched to 2016 local authority boundaries. The working-age population is not held for the Scottish and Welsh districts or for Shepway, and the Isles of Scilly has no published T_d. Nine further English districts lack a complete tax credit series, four of them unitary authorities created in the 2009 reorganisation. This leaves 315 English districts observed in all ten years.

The estimates are from the event study

$$b_{d,t} \;=\; \alpha_d \;+\; \gamma_{r(d),t} \;+\; \sum_{\tau \neq 2009} \beta_\tau \cdot \mathbf{1}[t = \tau] \cdot T_d \;+\; \varepsilon_{d,t}, \tag{36}$$

where $b_{d,t}$ is benefit and tax credit income per working-age adult in district d in fiscal year t, α_d is a district effect, $\gamma_{r(d),t}$ a region-by-year effect, and β_τ is measured in pounds realised per pound of predicted loss, so full incidence corresponds to $\beta_\tau = -1$. Standard errors are clustered by local authority.

Table A.4 reports the estimates with and without detrending by removing the pre austerity trend. I construct the pre trend by taking the line of best fit through prereform years and subtracting from all observations.

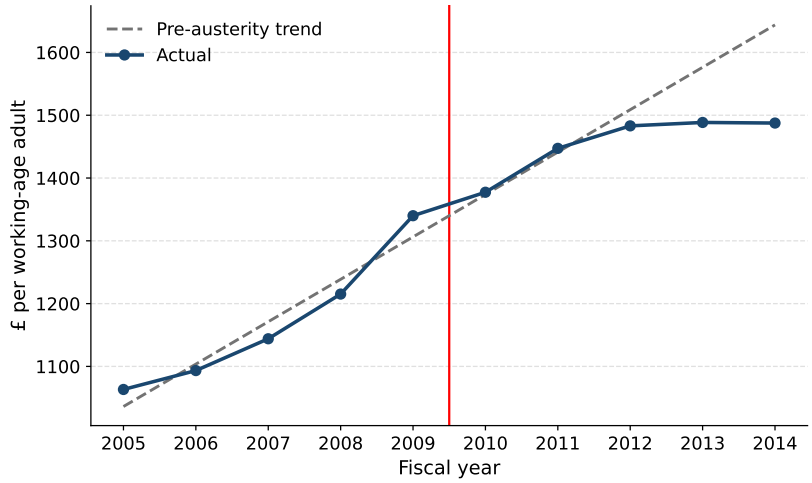


Figure (15) Benefit and tax credit income per working-age adult. Mean over the 315 districts of Appendix A.4.4, fiscal years 2005/06 to 2014/15, actual and on its 2005/06 to 2009/10 trend. The levels estimates of Table A.4 rise before 2010 with cash uprating and recession-driven caseload growth, both proportional to the local claimant base; their pre-trend slope is 0.24 per year (s.e. 0.01).

Table (A.4) Realised benefit income on exposure, by fiscal year

Fiscal year	Levels		Pre-trend adjusted	
2005/06	−0.978	[−1.055, −0.902]	−0.012	[−0.020, −0.004]
2006/07	−0.850	[−0.915, −0.785]	−0.125	[−0.142, −0.108]
2007/08	−0.672	[−0.728, −0.616]	−0.189	[−0.220, −0.158]
2008/09	−0.391	[−0.427, −0.354]	−0.149	[−0.177, −0.121]
2009/10	reference			
2010/11	0.159	[0.129, 0.189]	−0.083	[−0.110, −0.056]
2011/12	0.389	[0.344, 0.434]	−0.094	[−0.133, −0.055]
2012/13	0.516	[0.451, 0.581]	−0.209	[−0.262, −0.156]
2013/14	0.534	[0.462, 0.607]	−0.432	[−0.497, −0.367]
2014/15	0.526	[0.442, 0.611]	−0.682	[−0.763, −0.600]

Notes. Estimates of β_τ from equation (36), pounds of benefit and tax credit income per working-age adult per pound of T_d, with 95 per cent confidence intervals in brackets; 315 districts, standard errors clustered by district. The adjusted column removes a linear trend fitted to the 2005/06 to 2009/10 coefficients. Fiscal years are labelled by their starting year in equation (36).

A.5 The incidence of the reforms across households

I construct the vector of household incidence following the analysis of De Agostini et al. (2018), who estimate the change in household net income from the 2010-2015 UK government policies by income decile. Table A.5 reports the exposure of the three designs before normalisation to ensure the aggregate withdrawal is comparable across all designs. Losses from the 2010 design fall with income, indeed three of the top four deciles show a gain across the period. Note that this reflects a change in income tax thresholds rather than positive benefit transfers to these households. To construct the vector used in the paper, all estimates are bounded below by a zero so that the estimated loss is non-negative. I average deciles into quintiles, largely driven by computational concerns in the numerical implementation, and then normalise the loss so that it withdraws the same spending under each design, which is the basis of the counterfactual examined in the paper.

$$\ell_k \;=\; \tilde{\ell}_k \,\frac{\sum_{k'} E_{k'}}{\sum_{k'} \tilde{\ell}_{k'} E_{k'}}, \tag{37}$$

where E_k is the stationary spending of a household in quintile k. In period t the design then lowers the spending of a household in quintile k by $\tau_t \ell_k E_k$, so a quintile's spending falls in proportion to its loss intensity, which in the data is a share of net income. Section 6 reports the cut in the bottom quintile's spending that it implies and the share of the spending withdrawn that falls on the two poorest quintiles.

The flat design sets $\tilde{\ell}_k = 1$ for every quintile, which the normalisation leaves unchanged, so that it withdraws the same share of every household's expenditure. The progressive design is the inverse of the 2010 vector before normalisation, scaled in the same way. The 2010 vector puts no weight on the fourth quintile and a small weight on the fifth. Its mirror image therefore puts a small weight on the bottom quintile and none on the second, so the design is not monotone at the bottom. Note that the inverse design assumes positive withdrawals from higher income households, in practice, this would require a different set of fiscal instruments to achieve than in this setting.

Table (A.5) The incidence of the three designs before normalisation

	Loss intensity $\tilde{\ell}_k$ before normalisation		
Quintile	2010 design	flat	progressive
1 (poorest)	2.45	1	0.05
2	1.55	1	0
3	0.50	1	0.50
4	0	1	1.55
5 (richest)	0.05	1	2.45

Notes. Column 2: the losses in household net income from the tax and benefit changes of 2010–15 by decile in De Agostini et al. (2018), with the three gains set to zero, averaged over the two deciles in each quintile. Column 3: one for every quintile. Column 4: column 2 in reverse order. Each vector is then scaled as in equation (37), at the model's stationary expenditure, so that every design withdraws the same spending.

A.6 Household data

Living Costs and Food Survey. The household sample pools the Expenditure and Food Survey and its successor from 2008, the Living Costs and Food Survey, for the calendar years 2002 to 2016.[34] The unit of observation is the household in a survey quarter, and the sample has 87,081 household-quarters and 81,284 households, counting a case number within a year as one household, as Panel B of Table B.6 reports.

A household is a benefit recipient if it reports positive income in the quarter from Income Support, Housing Benefit, other social security benefits, or social security benefits in total, which include child benefit and tax credits. A household that reports any state pension income is coded as a non-recipient, so households with state pension income are kept in the sample and counted among the non-recipients. Recipient status is measured contemporaneously, and records whose recipient status is fractional are dropped. Recipients make up 39 per cent of the sample.

Spending on each COICOP division is the division's share of total expenditure multiplied by total expenditure. Total expenditure excludes housing costs, and of division 4 (housing, water, electricity, gas and other fuels) only electricity, gas and other fuels are kept, which is the necessity division called household fuel in Table A.3. Luxury spending is the sum over the five divisions of the luxury group and necessity spending the sum over the four divisions of the necessity group, while alcohol and tobacco (2), transport (7) and education (10) are in neither group.

For the event study of Figure 5 the quarterly records of a household are averaged to one observation per household and year. The outcome is the natural logarithm of weekly expenditure (total, luxury or necessity), on the observations on which it is positive, and the coefficients are reported in log points. Each regression includes indicators for household size (one to five or more members) and region-by-year effects, the interactions of recipient status with year are measured against 2010, and the confidence intervals are heteroskedasticity-robust. Panel (b) of Figure 5 expresses the same coefficients against their 2005 to 2009 mean, with standard errors from the corresponding linear combinations.

Understanding Society and the British Household Panel Survey. The evidence on how often households end their relationships with suppliers is from Understanding Society for 2009 to 2015 and from its predecessor, the British Household Panel Survey, for 1996 to 2008. At each annual interview one member reports for the household whether it has satellite or cable television, a landline, a mobile telephone and internet access at home, and an item ends when the household has it at one interview and not at the next.[35] A household is followed from one interview to the next through its reference person, which gives 147,008 linked pairs of interviews with the first in 2009 to 2015. The British Household Panel Survey gives 90,069 linked pairs. Appendix B.7 reports the ending rates by income quintile.

[34] UK Data Service study numbers 5003, 5210, 5375, 5688, 5986 and 6118 (Expenditure and Food Survey, 2002–03 to 2007) and 6385, 6655, 6945, 7272, 7472, 7702, 7992, 8210 and 8351 (Living Costs and Food Survey, 2008 to 2016–17).

[35] The British Household Panel Survey asks about mobile telephones from 2006, and about internet access in 1996 and from 2000.

A.7 District data on firm births and survival

The district data that Section 6 sets against the model are published statistics for local authorities on their April 2023 boundaries. Firm births and survival are from the ONS business demography reference tables, which count the births of enterprises on the Inter-Departmental Business Register in each local authority and year and the number of each year's births that survive one to five years. Births in each year from 2008 to 2019 are taken from the edition of the tables for two years later, and those of 2004 to 2007 from the 2009 edition. The five-year survival of a birth cohort is the share of its births still active five years later, and each cohort from 2004 to 2016 is taken from the edition for five years after it. Births are divided by the population aged 16 to 64 from the ONS mid-year estimates, published on Nomis, and the outcome is the logarithm of births per working-age adult.

The exposure of an April 2023 authority is the average of T_d over the December 2016 authorities it contains, weighted by their working-age population in 2009, and it is measured in standard deviations of T_d across the 378 authorities of Appendix A.4.2. The sample is the 348 authorities of Great Britain with a published T_d, which leaves out the Isles of Scilly and Merthyr Tydfil.

The council-spending control is the change in the logarithm of council service spending per head between 2009/10 and 2014/15, standardised and interacted with year. Service spending is net current expenditure on services in the Revenue Outturn Summary of the Department for Communities and Local Government, less education, police, fire and rescue and, in 2014/15, public health, which councils took over in 2013. The spending of each shire district includes a share of its county's spending in proportion to population, and spending is divided by the population of all ages from the mid-year estimates. The Revenue Outturn covers English councils, so the specification with this control has the 295 English authorities.

B Additional empirical results

This appendix reports further detail on the empirical evidence of Sections 3.1 to 3.7, and the district and household evidence that Sections 4 to 6 draw on. Section B.1 gives summary statistics for the two samples and the exposure measure, and Section B.2 covers identification, including the balance tests. Section B.3 covers the remaining outcomes and sample splits. Section B.4 describes how markups and total factor productivity are estimated, and Section B.5 reports firm births and the survival of birth cohorts in public district data. Section B.6 reports the household evidence from the Living Costs and Food Survey, and Section B.7 the household evidence on the separation rates of Assumption 1 and on their level.

B.1 Summary statistics

Table B.6 describes the two samples of Section 3 and the exposure measure.

B.2 Identification

As in the usual case when comparing spatial exposure to an aggregate shock in these designs, the coefficient δ_τ in equation (2) measures a slope across local authorities rather than the effect on the level of young firms' outcomes in the aggregate. To control for factors affecting all local authorities in a given NUTS-1 region in a year, particularly the response of all firms to the aggregate consolidation, I include region-by-year fixed effects in the regression. Consequently, ζ_τ absorbs district level responses that are common to young and old firms alike. δ_τ then measures the additional effect on young firms' outcomes, relative to 2009 and older firms in the same local authority, when that local authority is one standard deviation more exposed. Note that the structure of the ABS means that the estimation sample is best thought of as a panel of repeated cross sections with some repeated firms. A firm that was young in 2010 has consequently left the group defined as young by 2013, so the coefficients best describe cohorts of young firms rather than a given firm followed through time. For this reason the estimated event study does not include firm fixed effects. Note that this repeated cohort structure can be thought of as biasing estimated gaps toward zero, if conditions at entry are persistent, a firm that was exposed in 2010 and enters the comparison group in 2013 pulls down the comparison average.

Coefficients and pre-trends. Table B.7 reports the pooled event study of equation (2) for marketing expenditure and value added, the young-firm differential δ_τ that Figure 3 plots and the all-firm term ζ_τ, which is the response of firms older than two years. The joint test that the four pre-period interactions are zero gives $F = 0.68$, $p = 0.61$ for marketing and $F = 1.8$, $p = 0.13$ for value added. For value added the 2005 and 2008 coefficients are negative at the ten per cent level, at about half the size of the coefficients from 2011 on, which are larger and individually significant in every year through 2015.

Table (B.6) Summary statistics

Panel A: Firm sample (BSD–ABS linked panel, 2005–2015)	
Firm-years, output regressions (GVA per employee)	165,314
Firm-years, marketing regressions	171,052
Firm-years, employment regressions	171,059
Firm-years, young (aged 0–2) / old, full sample	13,326 / 256,087
Firms per year (approx.)	60,000
Local authorities	346
Marketing / revenue, age 1–2	5.0%
Marketing / revenue, age 3–4	3.7%
Marketing / revenue, age 5–9	3.0%
Marketing / revenue, age 10–15	2.0%
Marketing / revenue, age 16+	2.0%
Marketing / revenue, all firms	3.5%
Average employment	8.0

Panel B: Household sample (LCFS, 2002–2016)

	All		Recipients	Non-recipients
	Mean	S.d.	Mean	Mean
Total expenditure (£/week, excl. housing)	362.5	308.2	398.0	340.2
Total expenditure per household member (£/week)	165.9	137.0	127.3	190.2
Luxury-category expenditure (£/week)	191.0	207.4	208.7	179.9
Necessity-category expenditure (£/week)	89.3	58.7	98.8	83.2
Luxury share of expenditure	0.48	0.17	0.48	0.48
Necessity share of expenditure	0.32	0.17	0.31	0.32
Gross household income (£/week)	625.7	457.0	652.1	609.0
Benefit income, non-pension (£/week)	45.5	75.9	93.3	15.4
Household size	2.37	1.27	3.24	1.82
Number of children	0.56	0.96	1.34	0.06
Benefit recipient (share)	0.39	0.49	1.00	0.00
Household-quarters (observations)	87,081		33,652	53,429
Households (case by survey year)	81,284			

Panel C: Austerity exposure T_d, £ per working-age adult per year, by local authority

	Mean	S.d.	P10 / P90	Min / Max
T_d (N = 378 local authorities)	447	121	300 / 614	177 / 914

Notes. Panel A: the linked BSD–ABS panel used in Section 3.4. The marketing-intensity and size rows are data moments targeted in Table 6. Panel B: the LCFS analysis sample of Section 3.6, unweighted as in the estimation. Expenditure excludes housing; benefit income excludes the state pension; recipients are households with positive income from any social security benefit (Income Support, Housing Benefit, other benefits, or total benefit income, which includes child benefit and tax credits), measured contemporaneously, and households with any state pension income are counted as non-recipients. Recipient households are larger, which is why their total spending and income exceed non-recipients' while spending per member is lower. Panel C: the anticipated financial loss per working-age adult, the shock measure of Section 3.3, across the 378 local authorities for which the claimant counts are held; the estimation sample uses the 346 with linked firm data.

Table (B.7) Event study coefficients: young and all firms

	Marketing expenditure		Value added	
Year	Young $\times T_d$ (δ_τ)	T_d (ζ_τ)	Young $\times T_d$ (δ_τ)	T_d (ζ_τ)
2005	0.019	0.044^{**}	-0.132^{*}	0.024
	(0.065)	(0.021)	(0.078)	(0.038)
2006	0.047	−0.030	−0.097	−0.001
	(0.046)	(0.025)	(0.074)	(0.039)
2007	−0.025	−0.014	−0.005	−0.005
	(0.046)	(0.025)	(0.059)	(0.019)
2008	−0.001	−0.028	-0.128^{*}	0.032
	(0.044)	(0.022)	(0.077)	(0.036)
2010	-0.089^{*}	−0.002	-0.170^{**}	−0.016
	(0.048)	(0.023)	(0.075)	(0.036)
2011	-0.146^{**}	0.007	-0.289^{***}	0.018
	(0.068)	(0.019)	(0.086)	(0.036)
2012	-0.101^{**}	−0.020	-0.240^{**}	−0.008
	(0.050)	(0.020)	(0.095)	(0.035)
2013	0.040	−0.021	-0.140^{*}	−0.007
	(0.063)	(0.026)	(0.083)	(0.030)
2014	0.030	0.022	-0.264^{***}	0.032
	(0.043)	(0.021)	(0.076)	(0.041)
2015	−0.042	0.022	-0.199^{***}	0.039
	(0.060)	(0.027)	(0.079)	(0.042)
Young $\times T_d$, 2009 level	0.012		0.110^{**}	
	(0.029)		(0.056)	
Observations	171,052		165,314	

Notes: Coefficients on exposure T_d (in standard deviations) interacted with year, for all firms (ζ_τ) and the additional effect on young firms (δ_τ), base year 2009. The ζ_τ column is the response of old firms. The 2009 level row is the coefficient on Young $\times T_d$. Local authority and two-digit sector fixed effects, region by year effects, and the controls of Section 3.4: tradeable sector, UK ownership, the Herfindahl–Hirschman index, the local unemployment rate, and the 2005–09 change in claimant rates interacted with the young-firm indicator by year. Standard errors in parentheses, 340 clusters. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Balance on pre-reform characteristics. Table B.8 reports the balance tests discussed in Section 3.4. Panel A regresses seventeen pre-reform district characteristics from public data on exposure, conditional on region effects, and Panel B regresses the 2005–2009 means of the firm data in the 346 local authorities with linked firms on exposure. Panel C pools the firm-years from 2005 to 2009 and regresses each outcome on exposure and on its interaction with the young-firm indicator, which is the comparison that equation (2) makes, and Section 3.4 discusses its estimates. No interaction is significant, while exposure itself loads on employment and market concentration, as in Panel B.

The shift-share structure. As discussed in the main body of the paper, the national savings L_b used in equation (1) are an aggregate measure common to all districts, consequently the variation in T_d across local authorities is determined by variation in the claimant shares. Identification is thus share based in the spirit of Goldsmith-Pinkham et al. (2020).

Table B.9 decomposes Beatty and Fothergill's measure by reform. The weight of reform k is $\alpha_k = \operatorname{cov}(\ell_{dk}, T_d)/\operatorname{var}(T_d)$, where ℓ_{dk} is the loss per working-age adult from reform k in district d. These are then the Rotemberg weights recommended in Goldsmith-Pinkham et al. (2020) for an unweighted regression of district level outcomes on the treatment measures T_d across local authorities. Note that these are not the weights used in the firm-level estimates of Section 3.4, which weight each district by firm years and show coefficients from treatment T_d interacted with the young firm dummy variable and year.

The sickness and disability component carries 45 per cent of the variance of their measure and the incapacity reform alone 32 per cent. The four components are positively correlated with one another across districts, so none offsets another in its variance. Removing any single reform leaves a measure correlated with theirs at 0.958 or more, the lowest being without the incapacity reform, and removing the whole sickness and disability component leaves a correlation of 0.907. Measuring the shares before the recession changes the measure little. T_d is correlated with their measure at 0.973 across the 310 districts with both, and the same construction on 2007 claimant shares at 0.966.

Panel A of Table B.8 regresses pre-reform district characteristics from public sources on T_d, with region effects, and Table B.10 repeats these regressions on each component of Beatty and Fothergill's measure. On the employment rate, the claimant rate in 2005 and its rise over 2008 to 2010, and the public-sector share, the sickness and disability, tax credit and uprating components load as T_d does. The housing component loads weakly on these and most on the private-rented share of households, at 0.54, and on the growth of disposable income and house prices between 2005 and 2009, at 0.41 and 0.24, on which the tax credit and uprating components do not load.

B.3 Other outcomes and sample splits

All outcomes. Table 2 in the main text reports δ_τ for employment, markups, TFP, the three composition outcomes, net capital expenditure and energy expenditure. Estimated without the tradeable, ownership and concentration controls, the marketing coefficient is -0.109 for 2010 and -0.139 for 2011, against -0.089 and -0.146 with them in Table B.7 of Section 3.4.

Table (B.8) Balance on Pre-Reform Characteristics

Panel A: Exposure and pre-reform district characteristics, public data, region effects

Characteristic	Coef. on T_d	$[t]$	N
Employment rate, ages 16–64, 2009	−0.63	[−12.0]	352
Log median full-time pay, 2009	−0.54	[−11.7]	339
Log disposable income per head, 2009	−0.55	[−5.7]	301
Log median house price, 2009	−0.47	[−9.8]	293
JSA claimant rate, 2005	0.75	[14.5]	303
Public-sector employment share, 2009	0.42	[6.7]	328
Manufacturing employment share, 2009	0.08	[1.3]	328
Finance, professional and ICT share, 2009	−0.32	[−3.7]	328
Micro-firm share (under ten employees), 2010	−0.56	[−8.8]	328
Enterprises per 1,000 adults, 2010	−0.27	[−1.5]	353
Jobs density, 2009	−0.20	[−1.1]	328
Private-rented share of households, 2011	0.35	[5.5]	340
Log working-age population, 2009	0.32	[4.2]	353
Change in JSA claimant rate, 2008–10	0.74	[13.7]	303
Change in log jobs, 2005–09	−0.22	[−3.7]	328
Change in log income per head, 2005–09	0.18	[2.3]	301
Change in log house price, 2005–09	0.13	[1.8]	293

Panel B: Firm data, district means before 2010

Pre-2010 characteristic	Coef. on T_d	(s.e.)	N
Log employment	0.039***	(0.014)	346
Log output per worker	−0.034**	(0.016)	346
Log TFP	0.002	(0.010)	346
Markup	0.014	(0.036)	346
Log marketing expenditure	−0.007	(0.018)	346
Log GVA	0.006	(0.023)	346
Share of young firms	0.005	(0.005)	346
Share in tradeable sectors	0.026***	(0.007)	346
Share in necessity sectors	−0.001	(0.001)	346
Share UK-owned	0.001	(0.001)	346
Herfindahl–Hirschman index	11.6	(8.2)	346
Log number of firms	0.077*	(0.042)	346
Log employment, 2007–08 average	0.040	(0.031)	346
Log output per worker, 2007–08 average	−0.017	(0.025)	346
Change in log marketing, 2005 to 2009	0.073	(0.102)	255
Change in log output per worker, 2005 to 2009	0.054	(0.096)	240

Panel C: Firm-years before 2010, by exposure and age

	T_d		Young × T_d		
Outcome	Coef.	(s.e.)	Coef.	(s.e.)	N
Log employment	0.055***	(0.017)	−0.023	(0.031)	57,324
Log output per worker	−0.005	(0.032)	0.054	(0.041)	55,452
Log marketing expenditure	0.021	(0.043)	−0.039	(0.039)	57,321
UK-owned	0.006	(0.004)	−0.011	(0.007)	57,324
Markup	−0.008	(0.017)	0.014	(0.028)	57,324
Log TFP	−0.015	(0.010)	0.009	(0.029)	57,318
Herfindahl–Hirschman index	40.9***	(13.7)	−2.6	(21.2)	57,324

Notes: Panel A: each cell is a separate regression across local authorities of the standardised characteristic on the standardised exposure measure T_d, with region fixed effects; heteroskedasticity-robust t-statistics in brackets. Rows below the inner line are changes over the pre-reform period; JSA is jobseeker's allowance. Sources: Annual Population Survey, ASHE, ONS regional accounts, House Price Statistics for Small Areas, NOMIS claimant count, BRES, UK Business Counts and the 2011 Census. Panel B: each row is a regression across local authorities of the 2005–2009 average of the characteristic, unless the row states otherwise, on exposure T_d in standard deviations. Panel C: firm-years from 2005 to 2009; each row regresses the outcome on exposure T_d in standard deviations and its interaction with the young-firm indicator, with standard errors clustered by local authority. Standard errors in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table (B.9) Exposure by Reform: Weights and Leave-One-Out Correlations

Component or reform	Mean loss (£)	Weight α_k Raw	Weight α_k Within region	Correlation with T_d when removed
Sickness and disability	158	0.452	0.428	0.907
Incapacity benefit to ESA	103	0.320	0.305	0.958
DLA to PIP	36	0.083	0.079	0.998
Under-occupation charge	11	0.031	0.024	0.999
Non-dependant deductions	8	0.018	0.019	1.000
Housing	41	0.168	0.184	0.954
Local housing allowance	36	0.143	0.159	0.971
Benefit cap	5	0.025	0.025	0.998
Uprating and council tax	88	0.219	0.230	0.994
One per cent uprating	80	0.192	0.197	0.998
Council tax support	7	0.028	0.033	0.998
Tax credits	87	0.174	0.179	0.992
Child benefit	71	−0.011	−0.008	0.998
Local authorities			370	

Notes: Components of Beatty and Fothergill's measure in the 2013 costing, reconstructed from their local authority estimates (Beatty & Fothergill, 2014) and their 2016 update by rescaling each reform's local geography to its 2013 national total; the reconstructed sum has a correlation of 0.993 with their measure. Mean loss is in pounds per working-age adult. The weight is $\alpha_k = \operatorname{cov}(\ell_{dk}, T_d)/\operatorname{var}(T_d)$, computed raw and conditional on region effects. The last column is the correlation of their measure with itself less the component or reform. ESA is employment and support allowance, DLA disability living allowance and PIP personal independence payment. Public data only.

Capital expenditure. Column (7) of Table 2 in Section 3.5 reports the event study for net capital expenditure on 171,059 firm-years. No coefficient is significant. The point estimates are close to zero in 2010 to 2012, the years in which the marketing cut is largest, and reach about −17 per cent in 2013.

Table (B.10) Pre-Reform District Characteristics and the Components of Exposure

Characteristic	Components of T_d				N
	Sickness and disability	Housing	Tax credits	Uprating and council tax	
Employment rate, ages 16–64, 2009	−0.68	−0.31	−0.55	−0.58	352
	[−12.7]	[−2.5]	[−11.3]	[−11.6]	
Log median full-time pay, 2009	−0.56	−0.20	−0.63	−0.56	339
	[−12.5]	[−1.3]	[−18.9]	[−16.5]	
Log disposable income per head, 2009	−0.53	−0.03	−0.71	−0.63	301
	[−10.3]	[−0.2]	[−8.2]	[−7.8]	
Log median house price, 2009	−0.52	−0.02	−0.58	−0.52	293
	[−17.0]	[−0.2]	[−18.8]	[−15.5]	
JSA claimant rate, 2005	0.80	0.37	0.66	0.70	303
	[16.9]	[2.0]	[13.5]	[15.7]	
Public-sector employment share, 2009	0.49	0.17	0.38	0.44	328
	[9.3]	[1.1]	[5.8]	[8.1]	
Manufacturing employment share, 2009	0.08	−0.06	0.17	0.10	328
	[1.2]	[−1.6]	[2.9]	[1.7]	
Finance, professional and ICT share, 2009	−0.27	−0.08	−0.45	−0.40	328
	[−4.5]	[−0.6]	[−5.1]	[−5.2]	
Micro-firm share (under ten employees), 2010	−0.66	−0.27	−0.50	−0.49	328
	[−13.5]	[−3.7]	[−7.2]	[−8.2]	
Enterprises per 1,000 adults, 2010	−0.18	−0.11	−0.32	−0.27	353
	[−2.2]	[−0.7]	[−1.5]	[−1.6]	
Jobs density, 2009	−0.10	−0.11	−0.25	−0.20	328
	[−1.1]	[−0.7]	[−1.2]	[−1.2]	
Private-rented share of households, 2011	0.29	0.54	0.21	0.28	340
	[4.9]	[4.6]	[3.3]	[4.9]	
Log working-age population, 2009	0.30	0.18	0.28	0.30	353
	[4.2]	[1.9]	[3.5]	[4.1]	
Change in JSA claimant rate, 2008–10	0.74	0.29	0.73	0.73	303
	[15.1]	[1.8]	[16.4]	[15.5]	
Change in log jobs, 2005–09	−0.20	−0.05	−0.27	−0.26	328
	[−3.2]	[−0.7]	[−4.9]	[−4.7]	
Change in log income per head, 2005–09	0.19	0.41	−0.01	0.07	301
	[2.5]	[5.2]	[−0.2]	[1.0]	
Change in log house price, 2005–09	0.22	0.24	−0.02	0.01	293
	[3.2]	[3.0]	[−0.3]	[0.2]	

Notes: Each cell is a separate regression across local authorities of the standardised characteristic on the standardised component of exposure in the column heading, with region fixed effects; heteroskedasticity-robust t-statistics in brackets. The components are those of Beatty and Fothergill's measure in the 2013 costing; the coefficients on T_d itself are in Panel A of Table B.8. Rows below the line are changes over the pre-reform period; JSA is jobseeker's allowance. Sources: Annual Population Survey, ASHE, ONS regional accounts, House Price Statistics for Small Areas, NOMIS claimant count, BRES, UK Business Counts and the 2011 Census.

B.4 Markups and total factor productivity

The markup and productivity outcomes in Table 2 of Section 3.5 come from a revenue production function estimated separately for each two-digit SIC sector s by ordinary least squares, with standard errors clustered by firm,

$$\ln y_{it} = \beta_m^s \ln M_{it} + \beta_w^s \ln W_{it} + \beta_k^s \ln K_{it} + \varepsilon_{it}. \tag{38}$$

Output y_{it} is deflated by a two-digit industry price index and materials M_{it} by a producer price index. The wage bill W_{it} is deflated by the consumer price index, and K_{it} is capital. The markup follows the production approach of De Loecker and Warzynski (2012), which Kariel et al. (2024) apply to UK firm data. With materials as the flexible input, the markup is the output elasticity of materials divided by the share of materials in revenue Y_{it},

$$\mu_{it} = \frac{\hat{\beta}_m^s}{M_{it}/Y_{it}}, \tag{39}$$

and total factor productivity is the residual $\ln y_{it} - \widehat{\ln y}_{it}$ of equation (38).

This is not the estimator the literature prefers. Ordinary least squares ignores the correlation between input choices and productivity that the control-function estimator of Ackerberg et al. (2015) is designed to remove. That estimator needs lagged inputs, and the ABS is a rotating stratified sample in which mostly large firms are observed in consecutive years, so it can only be estimated on a selected subsample. Output is also deflated with industry rather than firm prices, so the residual measures revenue productivity and still contains firm-level variation in prices and demand.

B.5 Firm births and survival in public district data

The firm panel of Section 3.4 follows the firms that the ABS samples, so it does not measure how many firms enter. Public district data on enterprise births measure entry directly. Figure 11 in Section 6.2 plots the district event study of births per working-age adult and of the five-year survival of each birth cohort on exposure T_d, with local authority and region-by-year effects and standard errors clustered by local authority.[36] Enterprise births per working-age adult in more exposed districts were 2.5 log points per standard deviation of exposure lower in 2010 than in 2009 (standard error 0.6) and 1.9 lower in 2011 (0.7). In this regression, on all 348 authorities, the five pre-period coefficients have a joint p-value of 0.140, and the coefficients are positive from 2013. Controlling for the change in council service spending per head between 2009/10 and 2014/15, interacted with year, deepens the fall in births in 2010 to 3.9 log points (0.7). The council data cover England only, so this regression uses the 295 English authorities, and its five pre-period coefficients have a joint p-value of 0.014. Figure 11 sets both regressions against the model's entry under the 2010 design. The same regression on the logarithm of the number of enterprises aged under two, and of those aged two to three,

[36] The sample is the 348 local authorities on April 2023 boundaries, and the 295 English authorities for the regression with the council-spending control. Births and five-year survival are from the ONS business demography, enterprises by age from the ONS *UK Business: Activity, Size and Location*, and the population aged 16 to 64 from the ONS mid-year estimates (Appendix A.7).

in March of each year from 2008 to 2013 gives panel (b) of Figure 11, with March 2009 as the base. Enterprises aged under two in more exposed districts were 2.2 log points per standard deviation of exposure fewer in March 2011 than in March 2009 (standard error 0.5). Their coefficient for March 2008 is 0.1 (0.4). Enterprises aged two to three were already 1.6 log points fewer in March 2010 (0.4), before the June 2010 Budget. By March 2013 this group holds the survivors of the enterprises that were under two in March 2011, and it is 2.2 log points lower (0.7). Five-year survival of the enterprises born in 2010 and 2011 is no lower in more exposed districts than that of the cohort of 2005, the reference cohort, and the coefficient for the cohort of 2010 is 0.4 percentage points per standard deviation (standard error 0.3).[37]

[37] Five-year survival is the share of a birth cohort still active five years later, so the 2004 cohort is the only one whose window closes by 2009, and the cohort of 2005, whose window closes in 2010, is the reference.

B.6 Household evidence from the Living Costs and Food Survey

Appendix A.6 describes the household sample, the definition of benefit recipients and the spending measures, Appendix A.3 lists the COICOP divisions in the luxury and necessity groups, and Panel B of Table B.6 gives summary statistics.

B.6.1 Event study

Figure 5 in the main text plots the coefficients on benefit-recipient status interacted with year from regressions of log household expenditure on one observation per household and year, with household-size indicators and region-by-year effects and heteroskedasticity-robust standard errors, and Appendix A.6 describes the estimation. Table B.11 reports the coefficients and their standard errors, against 2010 and against the 2005 to 2009 mean. Against 2010, recipients' luxury spending relative to non-recipients' was 15.2 log points lower in 2013 and 18.6 lower in 2015, their total spending about half as much lower, and their necessity spending unchanged.

Table (B.11) Coefficients Plotted in Figure 5

	(a) Against 2010			(b) Against the 2005–09 mean		
Year	Luxury	Necessity	Total	Luxury	Necessity	Total
2002	6.8 (3.5)	4.0 (1.9)	6.4 (2.6)	1.6 (2.7)	1.5 (1.5)	0.8 (2.0)
2003	9.8 (3.4)	0.8 (1.9)	7.7 (2.5)	4.5 (2.5)	−1.7 (1.4)	2.0 (1.9)
2004	8.7 (3.4)	1.8 (1.9)	6.3 (2.5)	3.5 (2.5)	−0.6 (1.4)	0.6 (1.9)
2005	0.8 (3.4)	2.8 (1.9)	3.6 (2.5)	−4.4 (2.1)	0.4 (1.2)	−2.0 (1.6)
2006	3.3 (3.4)	0.0 (1.9)	4.6 (2.6)	−1.9 (2.1)	−2.4 (1.2)	−1.1 (1.6)
2007	8.4 (3.4)	3.3 (1.9)	6.8 (2.5)	3.1 (2.0)	0.9 (1.2)	1.2 (1.5)
2008	9.3 (3.5)	4.6 (1.9)	8.8 (2.5)	4.1 (2.1)	2.2 (1.1)	3.1 (1.5)
2009	4.4 (3.4)	1.4 (1.9)	4.5 (2.5)	−0.9 (2.1)	−1.0 (1.1)	−1.2 (1.5)
2010	ref.	ref.	ref.	−5.2 (2.7)	−2.4 (1.5)	−5.7 (2.0)
2011	−2.4 (3.5)	−1.6 (1.9)	−0.7 (2.5)	−7.7 (2.6)	−4.0 (1.4)	−6.4 (1.9)
2012	−5.0 (3.6)	−0.8 (1.9)	−1.1 (2.5)	−10.2 (2.7)	−3.2 (1.4)	−6.8 (1.9)
2013	−15.2 (3.6)	−0.8 (1.9)	−8.8 (2.6)	−20.5 (2.8)	−3.2 (1.4)	−14.5 (1.9)
2014	−12.6 (3.7)	−1.9 (1.9)	−6.4 (2.6)	−17.8 (2.9)	−4.3 (1.4)	−12.0 (2.0)
2015	−18.6 (3.7)	0.3 (1.9)	−11.2 (2.6)	−23.8 (2.9)	−2.1 (1.4)	−16.8 (2.0)
2016	−13.1 (3.8)	−1.4 (2.0)	−9.9 (2.7)	−18.4 (3.0)	−3.8 (1.5)	−15.5 (2.1)

Notes. Coefficients on benefit-recipient status interacted with year, in log points, with heteroskedasticity-robust standard errors in parentheses, from the regressions of Figure 5: log household expenditure on household-size indicators and region-by-year effects, Living Costs and Food Survey 2002–2016, 81,244 (luxury), 81,268 (necessity) and 81,284 (total) household-years.

Figure 16 repeats the comparison under twelve combinations of recipient definition and base period. The four estimators are the pooled difference-in-differences, the event study averaged over 2011 to 2016 against 2010 and against the 2005 to 2009 mean, and the spending index. The luxury-minus-necessity gap is significantly negative in all 31 regression estimates, at between 7.5 and 26.2 log points, and negative in all 9 index versions.[38] Luxury spending is significantly negative in 30 of the 31 and total spending in 28. The exceptions are the estimates against 2010 that leave out tax-credit recipients, for luxury and total spending, and those that drop households with a top earner, for total spending. Against 2010 the gap was

[38] Grouping recipients instead by the benefit that a reform acted on, against two comparison groups, the gap is negative in all 42 estimates and significantly so in 36. In those groups necessity spending is significantly higher in 10 of the 42 estimates and total spending significantly lower in 21.

not flat before the reforms either, since its event-study coefficients are significantly positive in one to four of the eight years from 2002 to 2009, depending on the definition. Part of the gap measured against 2005 to 2009 or in the pooled comparison therefore predates 2010.

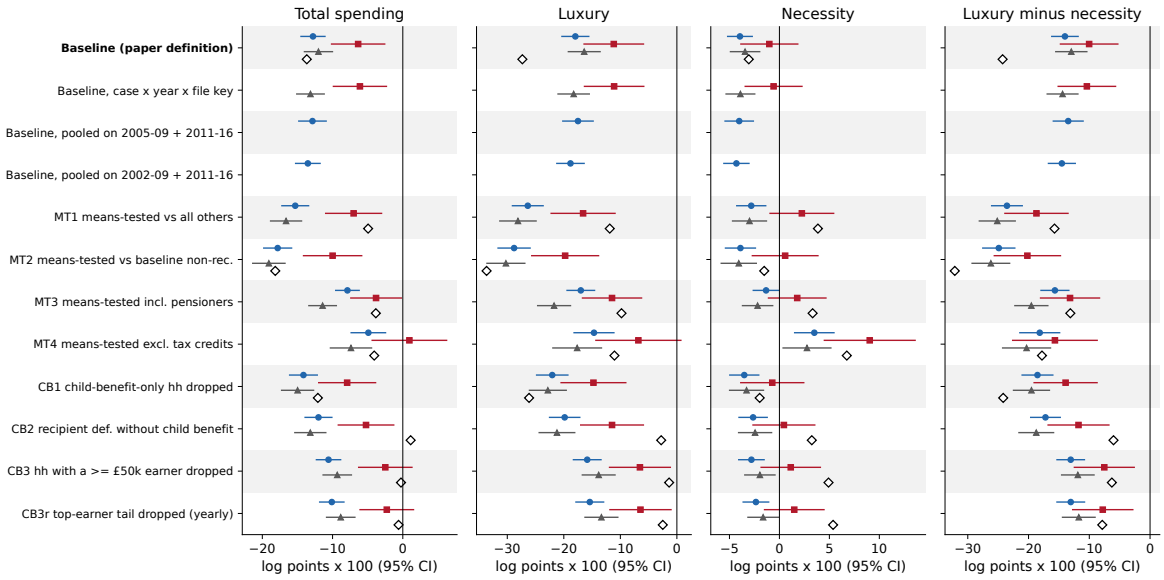


Figure (16) Household Spending Response across Recipient Definitions and Base Periods (LCFS). Post-2010 response of benefit recipients' spending relative to comparison households, in log points with 95 per cent confidence intervals, for total, luxury and necessity spending and the luxury-minus-necessity gap. Estimators: the pooled difference-in-differences of 2011–2016 against 2002–2010 (circles); the event study of Figure 5 averaged over 2011–2016, against 2010 (squares) and against the 2005–2009 mean (triangles); and the equal-weighted index averaged over 2013–2015 (2010 = 100), which has no interval (open diamonds). Rows: the paper definition; the same with households keyed by case, year and survey file; the pooled comparison on 2005–2009 and 2011–2016, or on 2002–2009 and 2011–2016; means-tested recipients against all other households (MT1), against the paper's non-recipients (MT2), including pensioner households (MT3) and excluding tax-credit recipients (MT4); and the paper definition without households whose only benefit is child benefit (CB1), with child benefit not counted (CB2), without households with an earner on £50,000 or more (CB3) and without the top-earner tail defined year by year (CB3r).

B.6.2 Difference-in-differences

Table B.12 reports the pooled difference-in-differences. In columns (1) and (2) the outcome is $\ln(1+y/100)$, where y is weekly spending on the category in pounds, regressed on recipient status, its interaction with an indicator for 2011 onwards, and year, region and household-size effects, with standard errors clustered by LCFS case number. The coefficient on the interaction is -0.097 for luxury spending and -0.013 for necessity spending, both significant at the one per cent level. Controlling for log total expenditure, which leaves only the change in how recipients divide a given budget, the luxury coefficient stays negative and the necessity coefficient turns positive, both significant at the one per cent level.[39] With spending in pounds inside the logarithm, $\ln(1+y)$, where the added one matters little, and standard errors clustered by household, the same regression gives a coefficient of -0.180 for luxury spending (column (3)). That coefficient is about 4.5 times the necessity coefficient in column

[39] Like any pooled comparison, these coefficients average over the whole pre-period, and panel (a) of Figure 5 shows that recipients' spending relative to non-recipients was above its 2010 level in most years before 2010.

(4), and both are significant at the one per cent level. Total spending, which enters in logs, is 12.8 log points lower for recipients after 2010 (column (5)).

Table (B.12) Household Spending After 2010: Difference-in-Differences

	(1)	(2)	(3)	(4)	(5)
	Luxury	Necessity	Luxury	Necessity	Total
Outcome	$\ln(1+y/100)$		$\ln(1+y)$		$\ln y$
Recipient × Post	-0.097^{***}	-0.013^{***}	-0.180^{***}	-0.040^{***}	-0.128^{***}
	(0.007)	(0.003)	(0.013)	(0.007)	(0.009)
Year, region and household-size effects	Yes	Yes	Yes	Yes	Yes
Clusters	Case number			Household	
Observations			87,081		

Notes: Living Costs and Food Survey 2002–2016, household-quarters. y is weekly spending on the category in pounds, and in column (5) total weekly expenditure; Post equals one from 2011. Standard errors in parentheses, clustered by LCFS case number in columns (1) and (2) and by household, a case number in a survey year, in columns (3) to (5). $^{***}p < 0.01$.

B.6.3 Pass-through from income to spending

Spending also follows income closely for these households. Relative to non-recipients, recipients' gross income fell by 8.8 per cent after 2010 and their total spending by 13.2 per cent, which is a marginal propensity to spend out of the gross-income loss of 0.94 (standard error 0.09).[40] Of this, 0.65 (0.06) comes out of luxury categories and 0.04 (0.01) out of necessities (Table B.13).[41]

Table (B.13) Pass-Through from Gross Income to Spending

	Recipient × Post (£ a week)	Recipients' mean, 2002–2010 (£)	Per cent of mean	Propensity out of gross income
Gross household income	−55.0	627.5	−8.8	
	(5.9)			
Total spending	−51.7	393.1	−13.2	0.94
	(4.1)			(0.09)
Luxury spending	−35.5	210.2	−16.9	0.65
	(2.8)			(0.06)
Necessity spending	−2.3	92.3	−2.5	0.04
	(0.8)			(0.01)
Observations		87,081		

Notes: Living Costs and Food Survey 2002–2016, household-quarters. The first column is the coefficient on recipient status interacted with Post, which equals one from 2011, in a regression of the outcome in pounds a week on recipient status, that interaction and year, region and household-size effects. Gross household income is normal weekly gross household income, top-coded. The propensity is the ratio of the spending coefficient to the gross-income coefficient, with delta-method standard errors from the joint covariance of the two. Standard errors clustered by household in parentheses.

[40] Measured instead as a ratio of log changes, the fall in total spending relative to the fall in gross income is 1.68 (0.17).

[41] The propensity is measured out of gross income rather than out of measured benefit income. Recipient status is defined contemporaneously, and the households the reforms pushed off entitlement leave the recipient group, so measured benefit income among remaining recipients rises after 2010 by £23 per week. The propensity is also a relative object, the recipient response net of the non-recipient one, on the same pooled comparison as the difference-in-differences above.

B.6.4 Income elasticities

Table B.14 reports the non-homothetic CES demand system of Comin et al. (2021) estimated by the generalised method of moments. For each category i the estimating equation relates log spending on i to log total expenditure E, log prices, and log spending on a base category b, clothing,

$$\ln x_i = \tilde{\sigma}(1-\vartheta_i)\ln E + (1-\tilde{\sigma})\ln p_i - \vartheta_i(1-\tilde{\sigma})\ln p_b + \vartheta_i \ln x_b + \text{controls}, \tag{40}$$

where ϑ_i is the non-homotheticity parameter of category i relative to the base. The elasticity of substitution across categories, $\tilde{\sigma}$, is fixed at 0.1. The controls are quarter effects, household type and household age group, with year effects added in column (4). The instruments are log income, the controls, quarter and expenditure-quintile effects and a price instrument for each category, observations carry the survey weights, and standard errors are clustered by household. Columns (1) to (3) use the 2019 survey, with (2) on winsorised data and (3) adding region to the instruments, and column (4) pools all survey years. Categories with ϑ_i above one are more income-elastic than clothing. In column (1) the parameters of furniture and recreation are above 1.2, while those of electricity and gas and of food are about 0.5.

Table (B.14) Non-Homothetic CES Demand: GMM Estimates of ϑ_i

	(1) 2019	(2) 2019, winsorised	(3) 2019, region instruments	(4) All years
Food	0.50	0.46	0.49	0.37
	(0.02)	(0.02)	(0.02)	(0.00)
Electricity and gas	0.52	0.47	0.51	0.30
	(0.02)	(0.02)	(0.02)	(0.00)
Furniture	1.21	1.12	1.19	1.20
	(0.05)	(0.05)	(0.05)	(0.01)
Transport	0.90	0.89	0.88	1.10
	(0.04)	(0.04)	(0.04)	(0.01)
Recreation	1.23	1.15	1.20	1.11
	(0.05)	(0.04)	(0.04)	(0.01)
Restaurants and hotels	0.98	0.97	0.96	0.99
	(0.04)	(0.04)	(0.03)	(0.01)
Miscellaneous	0.86	0.82	0.84	0.90
	(0.03)	(0.03)	(0.03)	(0.01)
N	3,164	2,815	3,164	56,538

Notes: Estimates of ϑ_i in equation (40), the non-homotheticity parameter of each category relative to the base category, clothing, with $\tilde{\sigma} = 0.1$. Living Costs and Food Survey, survey-weighted; standard errors clustered by household in parentheses.

Table B.15 reports Engel slopes, the slopes of log category spending on log total expenditure, within quintiles of total expenditure over 2002 to 2010.

The instrumented luxury elasticity also falls with income, and Table B.16 reports the estimates by quintile that Section 4.2 summarises, together with their least-squares counterparts and equation (10) evaluated at the luxury budget share of each income quintile.

Table (B.15) Engel Slopes by Quintile of Total Expenditure

	Luxury		Necessity		Mean weekly	
	Slope	(s.e.)	Slope	(s.e.)	expenditure (£)	N
Bottom quintile	1.30	(0.042)	0.67	(0.020)	85	11,078
Second	1.30	(0.024)	0.46	(0.024)	175	11,104
Third	1.25	(0.029)	0.42	(0.030)	273	11,100
Fourth	1.16	(0.027)	0.36	(0.029)	405	11,103
Top quintile	1.01	(0.016)	0.29	(0.015)	771	11,104
All households	1.23	(0.005)	0.44	(0.004)	342	55,489
Benefit recipients	1.20	(0.005)	0.46	(0.005)	393	22,187
Non-recipients	1.24	(0.008)	0.43	(0.005)	307	33,302

Notes: Slopes of log category expenditure on log total expenditure, Living Costs and Food Survey 2002–2010, within quintiles of total expenditure defined within year. Controls for household size and composition, age of the household reference person, and year and region effects. Standard errors clustered by household. N is for the luxury regressions; the necessity samples differ by at most 12 household-quarters.

Table (B.16) The Luxury Income Elasticity by Quintile

	Bottom	2	3	4	Top	Bottom − top
Quintiles of equivalised income						
Instrumented	1.287	1.253	1.230	1.212	1.190	0.097
	(0.016)	(0.007)	(0.003)	(0.006)	(0.013)	(0.028)
Least squares	1.271	1.245	1.227	1.215	1.200	0.071
	(0.004)	(0.003)	(0.002)	(0.002)	(0.003)	(0.006)
Luxury budget share	0.426	0.461	0.489	0.507	0.533	
Equation (10) at these shares	1.344	1.309	1.284	1.269	1.248	0.096
Quintiles of equivalised expenditure						
Instrumented	1.344	1.265	1.227	1.199	1.166	0.179
	(0.030)	(0.009)	(0.003)	(0.009)	(0.019)	(0.049)
Least squares	1.314	1.254	1.225	1.205	1.182	0.132
	(0.008)	(0.003)	(0.002)	(0.003)	(0.004)	(0.011)

Notes: Living Costs and Food Survey 2002–2016, 86,888 household-quarters with positive income, unweighted. One quadratic Engel curve for the luxury budget share, $w = a + b \ln E + c(\ln E)^2$, with E total expenditure excluding housing divided by the square root of household size; the elasticity at quintile q is $1 + (b + 2c \ln E_q)/w_q$ at the quintile's mean log expenditure and mean luxury share. Instrumented: $\ln E$ and its square instrumented by log equivalised disposable household income and its square. Controls: region-by-year and interview-quarter effects, household size, the number of children aged 0–1, 2–4 and 5–17, and the age band of the household reference person. Quintiles are formed from deciles defined within year and household size. Standard errors clustered by household, by the delta method, in parentheses. The luxury budget share is the mean share in each income quintile, and the equation row evaluates equation (10) at that share with $\gamma_1 = 1.22$ and $\sigma = 0.5$.

B.7 Household evidence on separation rates

Restrictions 1 and 2 of Assumption 1 relate to how the separation rate varies across households and across the business cycle. Restriction 1 says that lower income households separate from existing suppliers at a higher rate, and the second that a given household ends more relationships when it spends a lower amount on each. This section reviews the household evidence from Understanding Society that details whether a household stops having a service, as well as similar estimates from Ofcom's affordability tracker and from the LCFS survey for spending on existing services.

Understanding Society asks interviewees at an annual frequency whether or not a household has a given telecommunications service. An item is classed as ending when a household reports having it in the prior interview and not in the subsequent one. Across 2009-2015, households in the lower quintile of income report stopping one of these services 2.2 times as often as those in the top quintile. Panel A of Table B.17 reports the ratio with controls for working age households as well as from the Ofcom tracker.

The ending rate additionally falls with income in Understanding Society, both across the households in the survey and when looking at within household variation across years (Panel B of Table B.17).

The elasticity κ is disciplined by the second set of these estimates. With household fixed effects, a 1% fall in income increases the household ending rate by 0.11%. Using the estimated income elasticities from the LCFS, gives an ending elasticity of $\kappa = 0.3$ with a standard error of 0.13, which is the value used in the model. I do not reject the hypothesis of an equal ending rate across income quintiles, and so set one κ for all households. Note that the estimates from the survey data identifies only the ratio across income quintiles and not within quintiles. In the model I set the mean separation rate as 0.15, reflecting that relationships end either because households switch suppliers, or because firms exit and cease to produce. This estimate is consistent with the ranges reported in the literature (Gourio & Rudanko, 2014; Foster et al., 2016; Roldan-Blanco & Gilbukh, 2021; Afrouzi et al., 2023).

Table (B.17) How Often Households End Relationships, by Income

Panel A: Ending rate in the bottom income group relative to the top			
Survey and years	Relationship and ending	Ratio	With controls
Understanding Society, 2009–2015	Four household services, gone at the next annual interview	2.20 (0.06)	1.95 (0.05)
	working-age households in the bottom quintile	2.26 (0.06)	2.03 (0.06)
Understanding Society, 2016–2024	Landline, mobile telephone and internet access at home	2.10 (0.09)	1.95 (0.08)
British Household Panel Survey, 1996–2008	Same items	1.79 (0.09)	2.09 (0.12)
Ofcom Affordability Tracker, 2022–2026	Cancelled a service it could no longer afford, in the previous month	1.89 [1.53, 2.33]	2.57
	respondents under 65	2.51 [1.99, 3.12]	2.65
Panel B: A household's own finances			
Survey and years	Change	Effect on endings	
Understanding Society, 2009–2015	Fall of one log point in household income, with controls	0.0197 (0.0017) [0.38 of the mean rate]	
	with household fixed effects	0.0056 (0.0017) [0.11 of the mean rate]	
	Rise of one log point, with household fixed effects	0.0017 (0.0014) [0.03 of the mean rate]	
Ofcom Affordability Tracker, 2020–2021	Finances worse against unchanged: per cent who cancelled a service	7.9 (0.8) against 3.0 (0.3) ratio 2.7 [2.0, 3.6]	

Notes: Panel A: ratios of the ending rate of households in the bottom quintile of equivalised income to that of the top quintile, with standard errors in parentheses. Panel B: linear probability estimate of discontinuing a relationship against income

Table (B.18) The Elasticity of the Separation Rate to Spending

Ending response	Spending on the services	Income elasticity	κ
Within household, fall in income 0.108 (0.033)	Expenditure classes	0.259	0.42 (0.13)
	Model sector	0.485	0.22 (0.07)
	Item, among households that have it	0.123	0.88 (0.27)
Within household, rise in income −0.032 (0.028)	Expenditure classes	0.259	−0.12 (0.11)
Across households, income level 0.333 (0.015)	Sector, per expenditure division bought	0.419	0.79 (0.04)
	Item, among households that have it	0.123	2.70 (0.15)

Notes: Understanding Society and Living Costs and Food Survey, 2009 to 2015. κ is the response of the ending rate in Understanding Society, as a share of its mean and with its standard error below each row label, divided by the income elasticity of spending in the Living Costs and Food Survey on log equivalised income with year, region, household-size and age-band controls; standard errors by the delta method, treating the two surveys as independent. The within-household responses are those of Panel B of Table B.17, per log point of income lost (for a rise, the fall in the ending rate per log point gained), and the across-household response is the elasticity of the ending rate to the level of income. The expenditure classes are telephone and internet services and cultural services, which contain television subscriptions; the model sector is luxury for pay television and necessity for communication; across households, spending is per expenditure division or item bought; each denominator weights the services by their shares of endings.

C Model derivations

This appendix derives the equations of Section 4 from the household's problem (7) and the firm's problem (24), and the hours rule of Section 6.1. It takes Assumption 1 and the path rule (16) as given; Appendix D sets out a household problem consistent with them and the literature.

C.1 Households

The household's problem. A type-θ household solves (7), and spends in period t

$$E_t(\theta) \;\equiv\; \theta w_t h_t + \bar{\Pi}_t + \tau_t(\theta). \tag{C1}$$

The problem is time separable: households do not save, taking on a supplier does not involve a resource or utility cost, and by the third restriction of Assumption 1 the household takes the rate at which its relationships end as given.

Suppressing the time index for ease of notation, the static problem is

$$\max_{\{c_{j\omega}\},\,h} \; u\big(\mathcal{U}(c_{j\omega}),\, \xi_\theta(h)\big) \qquad \text{s.t.} \qquad \sum_j \int_{\mathcal{M}_j(\theta)} p_{j\omega}\, c_{j\omega}\, d\omega \;=\; \theta w h + \bar{\Pi} + \tau(\theta), \tag{C2}$$

which the household solves by two-stage budgeting: within each sector it solves its cost minimisation problem, which defines its own index $p_j(\theta)$, and then across sectors it allocates expenditure E to attain the largest composite $C(E;p)$ at those indices. Hours satisfy

$$u_C\, \frac{\partial C}{\partial E}\, \theta w \;=\; -\,u_\xi\, \xi'_\theta(h), \tag{C3}$$

the value of an hour's pay in the composite against its disutility. In the stationary equilibrium every type works $\bar{h}$ hours, so $E(\theta) = w\theta\bar{h} + \bar{\Pi}$.

Sector shares and the expenditure function. Minimising $\sum_j p_j c_j$ subject to (8) at their own indices, with multiplier Λ on the aggregator, gives the solution to sectoral demand as $c_j = \Lambda^\sigma \zeta_j^\sigma C^{\gamma_j} p_j^{-\sigma}$, so the budget shares can be expressed as

$$s_j(\theta) \;=\; \frac{p_j c_j}{\sum_i p_i c_i} \;=\; \frac{\zeta_j^\sigma\, p_j(\theta)^{1-\sigma}\, C(\theta)^{\gamma_j}}{\sum_i \zeta_i^\sigma\, p_i(\theta)^{1-\sigma}\, C(\theta)^{\gamma_i}},$$

which is (9), and the cost of the composite is then

$$\mathcal{E}(C;p) \;=\; \Big[\sum_j \zeta_j^\sigma p_j^{1-\sigma} C^{(1-\sigma)+\gamma_j}\Big]^{\frac{1}{1-\sigma}}, \tag{C4}$$

which rises in C, so the household's composite is the unique solution of $\mathcal{E}(C;p(\theta)) = E(\theta)$. With $\bar{\gamma} = \sum_j s_j\gamma_j$ and $\eta_C = \frac{(1-\sigma)}{(1-\sigma+\bar{\gamma})}$,

$$\frac{\partial \ln C}{\partial \ln E} \;=\; \eta_C, \qquad \left.\frac{\partial \ln C}{\partial \ln p_j}\right|_E \;=\; -\eta_C\, s_j, \qquad \eta_j \;=\; \frac{\partial \ln E_j}{\partial \ln E} \;=\; 1+\eta_C(\gamma_j-\bar{\gamma}), \tag{C5}$$

the last of which is (10). η_C is the elasticity of the composite to expenditure, below one whenever $\bar{\gamma} > 0$ because a further unit of spending is disproportionately spent on sectors with a higher γ_j.

The income elasticity falls with income. A household's elasticity in sector j depends on the household only through $\bar{\gamma}$, and at given indices

$$\frac{d\eta_j}{d\bar{\gamma}} \;=\; -\frac{(1-\sigma)(1-\sigma+\gamma_j)}{(1-\sigma+\bar{\gamma})^2} \;<\; 0, \qquad\qquad \frac{\partial\bar{\gamma}}{\partial \ln E} \;=\; \eta_C \sum_j s_j(\gamma_j-\bar{\gamma})^2 \;>\; 0. \tag{C6}$$

A richer household allocates more of its budget in high-γ_j sectors, so its $\bar{\gamma}$ is higher and its elasticity lower in every sector. A poorer household's lower $\bar{\gamma}$ both widens the gap $\gamma_j - \bar{\gamma}$ and raises η_C, the two parts that Section 4.2 describes. With two sectors, $\gamma_0 = 0 < \gamma_1$ and $\bar{\gamma} = \gamma_1 s_1$, so

$$\eta_1 - 1 \;=\; \eta_C\gamma_1 s_0 \;>\; 0, \qquad 1-\eta_0 \;=\; \eta_C\gamma_1 s_1 \;>\; 0, \qquad \eta_1-\eta_0 \;=\; \eta_C\gamma_1, \tag{C7}$$

and the necessity elasticity changes across households by $\frac{(1-\sigma)}{(1-\sigma+\gamma_1)}$ times the change in the luxury elasticity.

Demand within a sector and the own index. Given its set $\mathcal{M}_j(\theta)$ and posted prices, the household reaches the level $c_j(\theta)$ of (11) through cost minimisation, which gives the demand (13) with the own index (14) as the multiplier, and variety ω takes the share $\varsigma_\omega = (p_{j\omega}/p_j(\theta))^{1-\varepsilon}$ of sector spending. The own index is a CES price index over the suppliers the household holds, so losing suppliers carrying a share ς of its sector spending, or a marginal supplier, raises it by

$$\frac{p_j'(\theta)}{p_j(\theta)} \;=\; (1-\varsigma)^{-\frac{1}{\varepsilon-1}}, \qquad\qquad \frac{\partial \ln p_j(\theta)}{\partial m_j} \;=\; -\frac{1}{(\varepsilon-1)\,m_j(\theta)}, \tag{C8}$$

at unchanged posted prices, and with a common price $\bar{p}$, $p_j(\theta) = \bar{p}\,m_j(\theta)^{-1/(\varepsilon-1)}$. A single firm has no weight in the integral, so neither its price nor whether the household holds it moves the household's index, composite or shares (Arkolakis, 2010). Losing a supplier has the effect of increasing a household's price index.

C.2 Matching and separation

The law of motion of the sets. A relationship ends with probability $\lambda_{j,t}(\theta)$ and loses its supplier to exit with probability $\bar{\psi}_{j,t}(\theta)$, independently, so a mass $(1-\bar{\psi}_{j,t})(1-\lambda_{j,t})m_{j,t}$ survives the period, and the household takes on every supplier whose contact reaches an open place, since doing so costs nothing and lowers its own index, which gives (15). In a stationary equilibrium it takes on

$$A_j(\theta) \;=\; \big[\bar{\psi}_j(\theta) + \big(1-\bar{\psi}_j(\theta)\big)\lambda_j(\theta)\big]\, m_j(\theta) \tag{C9}$$

suppliers each period, that is, exactly the relationships it loses. The exit rate it faces is that of its suppliers, averaged over their ages. Writing $n_{j\omega,t}(\theta)$ for the type-θ relationships of firm ω, of age $a_{j\omega}$, the stock of relationships between type θ and sector j is $\int_{\Omega_j} n_{j\omega,t}(\theta)\,d\omega = m_{j,t}(\theta)\phi(\theta)$ by (27), and

$$\bar{\psi}_{j,t}(\theta) \;=\; \frac{1}{m_{j,t}(\theta)\,\phi(\theta)} \int_{\Omega_j} \psi(a_{j\omega})\, n_{j\omega,t}(\theta)\,d\omega, \tag{C10}$$

which is (28). Richer households' relationships last longer and consequently are more often with older firms, so $\bar{\psi}_{j,t}(\theta)$ varies with θ although $\psi(a)$ is common to all firms of an age. On the firm's side a type-θ relationship of a firm of age a survives with probability $(1-\psi(a))(1-\lambda_t(\theta))$, the $\varrho_k(a)$ of Section 4.4.1. With one separation rate for every type, every firm's customers keep the composition of the pool and the firm's state reduces to its number of customers, as in Gourio and Rudanko (2014) and Drozd and Nosal (2012).

Assumption 1 at impact. The dynamic separation rate (16) is given by a first-order expansion of $\ln \lambda_j(\theta, q)$ in $\ln q$ about the stationary point, and with $q_{j,t}(\theta) = E_{j,t}(\theta)/m_{j,t}(\theta)$,

$$\frac{\lambda_{j,t}(\theta)}{\lambda_j(\theta)} \;=\; \Big(\frac{E_j(\theta)}{E_{j,t}(\theta)}\Big)^{\kappa} \Big(\frac{m_{j,t}(\theta)}{m_j(\theta)}\Big)^{\kappa}. \tag{C11}$$

At impact the household holds its stationary set, so the rate moves with spending alone, with elasticity $-\kappa\eta_j(\theta)$ to expenditure: a consolidation raises the rate most where it cuts sector spending most. After impact a household that has lost suppliers spends more on each it keeps, and this second term lowers the rate. Note that under the model parameterisation the first term dominates the second.

The pool. Firms' contacts in sector j in period t are $\int_{\Omega_j} x^{\alpha}_{j\omega,t}\,d\omega + \bar{n}_0 M_{e,j,t}$, those of the survivors and the bases of the entrants, and $F_{j,t}(\theta)$ is the share of them that reaches type θ. The stock of relationships follows

$$m_{j,t+1}(\theta)\phi(\theta) \;=\; \big(1-\bar{\psi}_{j,t}(\theta)\big)\big(1-\lambda_{j,t}(\theta)\big)m_{j,t}(\theta)\phi(\theta) \;+\; F_{j,t}(\theta)\Big[\int_{\Omega_j} x^{\alpha}_{j\omega,t}\,d\omega + \bar{n}_0 M_{e,j,t}\Big]. \tag{C12}$$

In a stationary equilibrium contacts are allocated at random across open places and every open place is filled, so a type's share of contacts is its share of the relationships lost,

$$F_j(\theta) \;\propto\; \big[\bar{\psi}_j(\theta) + \big(1-\bar{\psi}_j(\theta)\big)\lambda_j(\theta)\big]\, m_j(\theta)\,\phi(\theta), \tag{C13}$$

which is (20), and (C12) at $m_{j,t+1} = m_{j,t}$ is the matching condition (29). A type is over-represented among the households looking for a supplier in proportion to how often it replaces one, so the pool is poorer than the population when

$$-\frac{d\ln\big[\bar{\psi}_j(\theta) + (1-\bar{\psi}_j(\theta))\lambda_j(\theta)\big]}{d\ln\theta} \;>\; \frac{d\ln m_j(\theta)}{d\ln\theta}, \tag{C14}$$

that is, when turnover falls with income faster than the number of suppliers rises, which the first restriction of Assumption 1 delivers at the calibration in both sectors. By

revenue, a unit of contacts brings from type θ an amount proportional to $[\bar{\psi}_j(\theta) + (1 - \bar{\psi}_j(\theta))\lambda_j(\theta)]E_j(\theta)\phi(\theta)/\bar{d}_j(\theta)$, with $\bar{d}_j(\theta)$ the mean of $p_{j\omega}^{1-\varepsilon}$ over the set.

C.3 Firms

Demand, pricing and profit. By (13), a type-θ customer of a firm posting p buys

$$c_j(\theta; p) \;=\; \Big(\frac{p}{p_j(\theta)}\Big)^{-\varepsilon}\frac{E_j(\theta)}{p_j(\theta)}, \tag{C15}$$

so a relationship yields the revenue $v_j(\theta, z) = p^{1-\varepsilon}p_j(\theta)^{\varepsilon-1}E_j(\theta)$, worth more the more the household spends in the sector and the fewer other suppliers it spreads that spending over. Adding up over the firm's customers,

$$D(p) \;=\; p^{-\varepsilon}\tilde{D}, \qquad \tilde{D} \;=\; \int_\Theta n(\theta)\, p_j(\theta)^{\varepsilon-1}E_j(\theta)\, d\theta, \tag{C16}$$

which is (17): the level moves with the composition of the firm's customers, but every customer has the same price elasticity, so the firm faces ε whatever its customers. A firm's price leaves its customers' own indices unchanged, since any one supplier is atomistic in them, and by the third restriction of Assumption 1 it moves neither the relationships the firm keeps nor, since acquisition is undirected, those it gains, consistent with the null markup response of Section 3. The price therefore maximises this period's profit $(p - w/z)p^{-\varepsilon}\tilde{D}$ alone, which gives the constant markup (18), and the profit of a relationship is $\varpi(\theta, z) = v_j(\theta, z)/\varepsilon$, which is (19).

The customer law of motion and composition by age. A surviving firm keeps $(1 - \lambda(\theta))n(\theta)$ customers of type θ and adds $x^\alpha F(\theta)$, the contacts its marketing makes that reach the type, which gives (21), and an entrant's $\bar{n}_0$ relationships are contacts of the same kind, (22). A firm that made x_i^α contacts at age i therefore holds at age a

$$n(\theta; a) \;=\; F(\theta)\, Q_a\big(1 - \lambda(\theta)\big), \qquad Q_a(y) \;=\; \bar{n}_0 y^a + \sum_{i<a} x_i^\alpha y^{a-1-i}, \tag{C17}$$

the pool weighted by how much of each past cohort has survived.

The value of a firm and marketing. The value (24) is affine in the customer measure, as Section 4.4.1 states. Let $\mathcal{V}_j(\theta; a, z)$ be the value of a type-θ relationship to a firm of age a and productivity z, the recursion of Section 4.4.1 with the survival probability $\varrho_j(\theta; a) = (1 - \psi(a))(1 - \lambda_j(\theta))$, and $V_{\text{new}}(a, z) = \int_\Theta F(\theta)\, \mathcal{V}_j(\theta; a + 1, z)\, d\theta$ the value of a contact drawn from the pool, as in (25). Then at every age

$$\begin{aligned} V\big(a, z, n(\cdot)\big) \;&=\; \mathcal{X}(a, z) + \int_\Theta n(\theta)\, \mathcal{V}_j(\theta; a, z)\, d\theta, \\ \mathcal{X}(a, z) \;&=\; \max_{x \geq 0}\Big\{ - wx^\nu + \beta\big(1 - \psi(a)\big)x^\alpha V_{\text{new}}(a, z)\Big\} + \beta\big(1 - \psi(a)\big)\mathcal{X}(a + 1, z), \qquad \mathcal{X}(\bar{a}, z) = 0, \end{aligned} \tag{C18}$$

where $\mathcal{X}(a,z)$ is the value of the firm's future acquisition, net of its cost. The decomposition holds by backward induction from $\bar{a}$, where the value is the profit $\Pi_j(z,n(\cdot)) = \int_\Theta n(\theta)\varpi_j(\theta,z)\,d\theta$: if it holds at $a+1$, substituting the law of motion of (24) into the continuation value separates the terms in n, which do not involve x, from the terms in x, which do not involve n. Along a transition the same recursions hold with dated profits and rates. Because the value is affine, the marginal value of a contact is $V_{\text{new}}(a,z)$ whatever the firm's base, and the first-order condition for marketing gives

$$x^*(a,z) \;=\; \left[\frac{\beta\big(1-\psi(a)\big)\,\alpha}{\nu\,w}\,V_{\text{new}}(a,z)\right]^{\frac{1}{\nu-\alpha}}, \tag{C19}$$

which is (25). Marketing costs the share α/ν of its gross return, is proportional to $z^{(\varepsilon-1)/(\nu-\alpha)}$, and does not depend on the firm's customers, since what a contact is worth depends on who is looking for a supplier and not on whom the firm already serves. At a given wage the response along a transition,

$$\ln\frac{x_t^*(a,z)}{x^*(a,z)} \;=\; \frac{1}{\nu-\alpha}\,\ln\frac{V_{\text{new},t}(a,z)}{V_{\text{new}}(a,z)}, \tag{C20}$$

is the same for every firm of an age in a sector, and differs across ages because $V_{\text{new},t}$ discounts future profits with age-specific survival. The expected value of entry in (26) is $V_j^e \;=\; \int\big[\mathcal{X}(0,z) + \bar{n}_0\int_\Theta F(\theta)\,\mathcal{V}_j(\theta;0,z)\,d\theta\big]\,dG_j(z)$, and in the stationary equilibrium $M_{e,j} = \bar{M}_{e,j}$ exactly when $V_j^e = wf_j$.

C.4 General equilibrium

Adding up. A relationship is counted once, as a supplier in a household's set and as a customer in a firm's base, so $\int_{\Omega_j} n_{j\omega}(\theta)\,d\omega = m_j(\theta)\phi(\theta)$ and the own index can be written over the firms' bases, $p_j(\theta)^{1-\varepsilon} = \phi(\theta)^{-1}\int_{\Omega_j} n_{j\omega}(\theta)p_{j\omega}^{1-\varepsilon}\,d\omega$, as in (28). Hence

$$\int_{\Omega_j} n_{j\omega}(\theta)\,v_j(\theta,z_{j\omega})\,d\omega \;=\; E_j(\theta)\,\phi(\theta), \tag{C21}$$

which is (27): the revenue firms earn from a type is what that type spends, in every period, whatever the sets. With bins, $\phi(\theta)$ is replaced by $1/K$.

Prices and the wage. The productivity distribution of the firms in Ω_j is G_j in every period, so with $p_{j\omega} = \frac{\varepsilon}{\varepsilon-1}w/z_{j\omega}$ the market index is (30), the numéraire $P_0 = 1$ sets $w = \frac{\varepsilon-1}{\varepsilon}[M_0\mathbb{E}_{G_0}z^{\varepsilon-1}]^{1/(\varepsilon-1)}$, and when the two sectors draw from the same distribution and share the exit rule, as at the calibration,

$$P_1 \;=\; \Big(\frac{M_1}{M_0}\Big)^{-\frac{1}{\varepsilon-1}} \;=\; \Big(\frac{M_{e,1}}{M_{e,0}}\Big)^{-\frac{1}{\varepsilon-1}}, \tag{C22}$$

so the relative price of luxuries is set by relative entry alone.

Walras' law. With $L = \int\theta\bar{h}\phi(\theta)d\theta$ the labour supplied, $L^{\text{prod}} = \sum_j\int_{\Omega_j} y_{j\omega}/z_{j\omega}\,d\omega$ the labour used in production and $L^{\text{mkt}} = \sum_j\int_{\Omega_j} x_{j\omega}^{\nu}\,d\omega$ the labour used in acquisition, revenue

counted from the household side by (C21) and from the firm side by (31) is

$$R \;=\; wL + \bar{\Pi} \;=\; wL^{\text{prod}} + wL^{\text{mkt}} + w\sum_j f_j M_{e,j} + \bar{\Pi},$$

so the labour market (32) clears at every positive wage if revenue holds, the rebate is consistent and every household spends its income.

Spending and revenue along a transition. Along a transition the wage is held at the steady state level, the rebate Π_t is the rebate (31) of the date-t distribution, and a type-θ household spends $E_t(\theta) = w\theta h_t(\theta) + \Pi_t - \tau_t \ell(\theta) E(\theta)$, the budget of (34) with the withdrawal of (33). By (C21) revenue, which is measured output, is what households spend,

$$R_t \;=\; \frac{1}{K}\sum_k E_{k,t} \;=\; (1-\tau_t)\, R \;-\; \frac{w}{K}\sum_k \theta_k\big(\bar{h} - h_{k,t}\big) \;+\; \big(\Pi_t - \bar{\Pi}\big), \tag{C23}$$

using $\sum_k E_k \ell_k = \sum_k E_k$ from (33), with R stationary revenue. Revenue falls by the share of spending withdrawn, by the labour income households lose as they work fewer hours, and by the fall in the rebate as profits fall with revenue, less the acquisition and entry costs firms save, which raise the rebate. The wages of the labour firms stop using are the pounds withdrawn (Section 6.1). At impact households hold their stationary sets, own indices and hours, so $R_0 = (1-\tau_0)R + (\Pi_0 - \bar{\Pi})$. A firm's problem along the path depends on households only through the paths of their spending $E_{jk,t}$, own indices $p_{jk,t}$ and separation rates $\lambda_{jk,t}$, and the labour market (32) is not imposed along the path.

C.5 Hours along a transition

With the preferences (34), a household values the composite it buys at today's own indices $p_t(\theta)$ at its stationary indices $p(\theta)$, and with $b(\theta) = w\theta\bar{h}^{-1/\epsilon_h}$ chooses h to maximise

$$\Upsilon(h) \;=\; \mathcal{E}\Big(C\big(w_t\theta h + \Pi_t - \tau_t \ell E;\, p_t\big);\, p\Big) - b\,\frac{h^{1+1/\epsilon_h}}{1+1/\epsilon_h},$$

so that its hours satisfy

$$\frac{h_t(\theta)}{\bar{h}} \;=\; \Big[\frac{w_t}{w}\,\Delta_t(\theta)\Big]^{\epsilon_h(\theta)}, \qquad \Delta_t(\theta) \;\equiv\; \frac{\mathcal{E}_C\big(C_t(\theta); p(\theta)\big)}{\mathcal{E}_C\big(C_t(\theta); p_t(\theta)\big)}, \tag{C24}$$

where $\mathcal{E}_C \equiv \partial\mathcal{E}/\partial C$ and Δ_t is what an hour's pay buys at today's indices valued at the stationary ones, equal to one at the stationary indices whatever the household spends. To first order in the indices, $\ln \Delta_t(\theta) = -\sum_j s_j(\theta)\eta_j(\theta)\ln(p_{j,t}(\theta)/p_j(\theta))$, with weights that sum to one, which is (35). At the held wage a household works fewer hours only when its own indices rise, which happens when it loses suppliers, so hours first move in the second period, and to first order they do not depend on its spending, on the withdrawal or on the rebate.

C.6 Revenue and relationships at impact

Let a design cut the expenditure of type θ at $t = 0$ by the proportion $e(\theta) = \tau_0 \ell(\theta) - (\Pi_0 - \bar{\Pi})/E(\theta)$, since hours do not move at impact and the rebate moves by the same pounds for every type. Posted prices, sets and the firm distribution are the stationary ones at $t = 0$, so a relationship's revenue moves with the type's spending alone. With $r_j(\theta)$ the share of a firm's revenue from type θ, as in Section 4.1,

$$\frac{R_0}{R} - 1 \;=\; \int_\Theta r_j(\theta)\Big(\frac{E_{j,0}(\theta)}{E_j(\theta)} - 1\Big)d\theta \;\approx\; -\int_\Theta r_j(\theta)\eta_j(\theta)e(\theta)\,d\theta, \tag{C25}$$

which is (5): a firm loses the revenue-weighted average of its customers' cuts. For two firms with revenue shares r_j^Y and r_j^O, both of which integrate to one, any constant $\bar{e}$ can be subtracted from the cut, and

$$\frac{R_0^Y}{R^Y} - \frac{R_0^O}{R^O} \;\approx\; -\int_\Theta \big[r_j^Y(\theta) - r_j^O(\theta)\big]\big[\eta_j(\theta)e(\theta) - \bar{e}\big]\,d\theta, \tag{C26}$$

which is (6). With $g(\theta) = r_j^Y(\theta)/r_j^O(\theta)$, whose mean under r_j^O is one, the right side is minus the covariance of g with $\eta_j e$ under r_j^O: the young firm loses more when the types it is over-exposed to are the types whose spending falls most. The covariance is positive when g falls with income and $\eta_j e$ is higher for poorer households, crossing its mean once. For single firms g is proportional to the ratio of their compositions, so under the first restriction of Assumption 1 it falls with income for an entrant against any older firm. For a cut common to every type the differential is $-e$ times the covariance of g with η_j, zero if g is constant, as with one separation rate, or if $\eta_j \equiv 1$, as with homothetic preferences, and negative when both fall with income: without differential churn young and old firms serve the same customers, and without non-homothetic preferences every customer cuts in the same proportion.

The split of the differential. At impact a household's budget share moves only with its expenditure, so the exact differential splits into the exposure of the two firms' customers to the cut and the change in their budget shares,

$$\frac{R_0^Y}{R^Y} - \frac{R_0^O}{R^O} \;=\; -\int_\Theta \big(r_j^Y - r_j^O\big)\,e\,d\theta \;-\; \int_\Theta \big(r_j^Y - r_j^O\big)\,(1-e)\Big(1 - \frac{s_{j,0}}{s_j}\Big)d\theta. \tag{C27}$$

The first term is what the young firm loses because its customers lose more income, the second what it loses because its poorer customers move their budget away from the sector by more; The cut $e(\theta)$ is the fall in the household's spending, the pounds the design withdraws less its share of the change in the rebate (Appendix C.4). The first term therefore carries the change in the rebate as well as the design's incidence, and under the flat design the cut is not the same proportion of spending for every type. In the luxury sector the gap between the young and old firms of Section 6.2 in 2010 is 1.07 points under the 2010 design, and the first term accounts for about seven tenths of it. Non-homothetic preferences make the differential negative under an equal cut, add to the young firm's loss in the luxury sector and narrow it in the necessity sector under a regressive cut, and since $\eta_1 - \eta_0 = \eta_C \gamma_1$ make the differential

larger in the luxury sector than in the necessity sector at the same revenue shares by γ_1 times the covariance of g with $\eta_C e$, which is positive under a regressive cut because η_C falls with income.

Relationships lost at impact. At impact $\lambda_0(\theta)/\lambda(\theta) = (E_j(\theta)/E_{j,0}(\theta))^\kappa$ by (C11). A surviving firm loses $\int_\Theta \lambda(\theta)n(\theta)\,d\theta$ relationships to separation at the stationary rates, so with the loss shares $\tilde{r}(\theta) = \lambda(\theta)n(\theta)/\int_\Theta \lambda n\,d\theta$, its separations in period 0 rise by

$$\frac{\int_\Theta \lambda_0(\theta)n(\theta)\,d\theta}{\int_\Theta \lambda(\theta)n(\theta)\,d\theta} - 1 \;=\; \int_\Theta \tilde{r}(\theta)\Big[\Big(\frac{E_j(\theta)}{E_{j,0}(\theta)}\Big)^\kappa - 1\Big]d\theta \;\approx\; \kappa \int_\Theta \tilde{r}(\theta)\eta_j(\theta)e(\theta)\,d\theta. \tag{C28}$$

Between a young and an old firm the difference in this rise is the covariance of (6) with loss shares in place of revenue shares and the opposite sign, since $\tilde{r}^Y/\tilde{r}^O$ is proportional to g, so the young firm's separations rise in at least as large a proportion as the old firm's. The relationships that end in period 0 are missing from the base from period 1 on: the firm's revenue falls at impact with its customers' spending and stays lower afterwards because it has fewer customers, which is where the persistence of the loss comes from.

D The household's problem of keeping and finding suppliers

Section 4 summarises a type-θ household's separation decisions by the rate $\lambda_j(\theta, q)$ at which its relationships in sector j end when it spends q per supplier, and restricts that rate by Assumption 1. This section sets out a household problem that satisfies the three restrictions set out in (16), and how the number of suppliers $m_j(\theta)$ varies by household type. Here, adding an additional supplier involves the household incurring a cost reflecting the time taken to find one. So households add suppliers until the marginal value of an additional one is equal to the cost of acquiring it. I assume following the retention and separation structure in Bornstein (2025), that a relationship ends when a household draws a supplier specific disruption cost greater than the value of keeping it. Note that unlike those papers here I model the disruption cost as a cost of staying rather than a cost of switching. Define the following notation for use in what follows. $E_{j,t}(\theta) = s_{j,t}(\theta)E_t(\theta)$ is spending on sector j by household type θ. Terms without time subscripts denote the stationary values. $q_j(\theta) = \frac{E_j(\theta)}{m_j(\theta)}$ is stationary spending per supplier, $\mathcal{M}_{j,t}(\theta)$ is the set of mass $m_{j,t}(\theta)$ the household holds, and $\bar{\psi}_{j,t}(\theta)$ is the exit rate of its suppliers. As in the main text, elasticities, budget shares and composite consumption are evaluated at household specific indices.

The environment. I extend the decision problem of the household in the main text in the following ways. First, time spent on finding an additional supplier is time that the household cannot spend on work or leisure, so flow utility then becomes $u(X_t(\theta))$, where:

$$X_t(\theta) \;=\; \mathcal{E}\big(C_t(\theta); p(\theta)\big) - b(\theta)\Big[\frac{h_t(\theta)^{1+1/\epsilon_h}}{1+1/\epsilon_h} + \bar{h}^{1/\epsilon_h}\, t_f \sum_j A_{j,t}(\theta)\Big] - \sum_j D_{j,t}(\theta), \qquad \text{(D1)}$$

where $C_t(\theta)$ is the composite (8), $A_{j,t}(\theta)$ is the number of suppliers the household takes on, t_f the hours each one takes, and $D_{j,t}$ the sum of the disruption costs the household faces from its existing suppliers, paid as a utility cost. The time spent on acquiring a supplier is valued at the marginal disutility of labour, so an hour spent looking for a supplier is worth one unit of the household efficiency wage. For a household to find an additional supplier, that supplier must also be trying to contact a household, since a firm can only enter the household awareness set if it contacts the household as in (Section 4.4).

The contacts that reach a type-θ household in sector j have mass $\Gamma_{j,t}(\theta) = \frac{F_{j,t}(\theta)\left[\int_{\Omega_j} x^{\alpha}_{j\omega,t}\, d\omega + \bar{n}_0 M_{e,j,t}\right]}{\phi(\theta)}$. This is the mass of contacts that firms make, distributed according to the acquisition pool, per household of a given type, which importantly is taken as given by the households. In the stationary equilibrium, by definition all open places are filled so the share of contacts is equal to the share of additional suppliers households want to acquire. Finally, over and above the model in the main text, in each period every relationship that does not dissolve because of firm exit, faces a disruption cost that is supplier specific, a cost δ drawn from G_δ on $[0, \infty)$, which is independent across relationships, households, and time. Its scale may differ by type, as below.

Lower-income households end more relationships because the disruption cost does not depend on the household's income or labour supply, whereas the cost of replacing a relationship

is time, valued at the household's wage.[42] The timing is as follows. Within a period, the household enters with given sets. It chooses hours worked and solves its expenditure allocation problem across the measure of firms it is aware of. Its suppliers then exit with the probability $\bar{\psi}_{j,t}(\theta)$, surviving relationships then draw the relationship specific disruption cost and the household chooses whether or not to separate. The household then accepts contacts to fill its open places. This then forms the supplier set for the following period, namely $\mathcal{M}_{j,t+1}(\theta)$.

Adding dm relationships from suppliers in sector j lowers the household's own index according to $d\ln p_j(\theta) = \frac{-(p_{j\omega}/p_j(\theta))^{1-\varepsilon} dm}{(\varepsilon-1)}$ following (14). This follows from the usual love of variety argument in CES preferences. Applying Shephard's lemma, $\partial\mathcal{E}/\partial \ln p_j = p_j c_j = E_j$, the household would need $S_{j\omega}(\theta)\, dm$ fewer pounds to reach the same composite without the additional suppliers, where

$$S_{j\omega}(\theta) \;=\; \frac{1}{\varepsilon - 1}\, E_j(\theta) \left(\frac{p_{j\omega}}{p_j(\theta)} \right)^{1-\varepsilon} \tag{D2}$$

is the surplus of a relationship with a supplier posting $p_{j\omega}$. The shares $(p_{j\omega}/p_j(\theta))^{1-\varepsilon}$ integrate to one over the set, so the average surplus of the household's relationships in the sector is its spending per supplier over $\varepsilon - 1$,

$$S_j(\theta) \;=\; \frac{E_j(\theta)}{(\varepsilon - 1)\, m_j(\theta)}. \tag{D3}$$

A particular supplier is worth more or less than this average through its price. I take the household to value every relationship in a sector at the average, so that no price enters its decisions. This is the third restriction of Assumption 1, and it reflects the null markup response of Section 3: a household that valued each supplier at its own surplus would keep cheaper suppliers more often, and a firm could hold customers by cutting its price.

Keeping suppliers. The household keeps a surviving relationship when its disruption cost is below the value of the relationship, $\delta_\omega \leq S_{j,t}(\theta)$. The costs are independent draws, so its separation rate is

$$\lambda_{j,t}(\theta) \;=\; 1 - G_\delta\!\left(\frac{E_{j,t}(\theta)}{(\varepsilon - 1)\, m_{j,t}(\theta)} \right), \tag{D4}$$

which depends on spending and on the set only through spending per supplier $q = E_{j,t}/m_{j,t}$, falls with q, which is the second restriction of Assumption 1, and contains no posted price, as the third requires. A household that spends less in a sector draws less surplus from each relationship and absorbs fewer disruptions. As in Bornstein (2025), where the cost of switching brands is drawn from a Pareto distribution whose median differs by age, I take the disruption cost to have a Pareto tail over the values a relationship takes, with an index $\kappa > 0$ common to every type and sector and a scale $\delta_0(\theta)$ that plays the part of his median,

$$1 - G_\delta(\delta) \;=\; \left(\frac{\delta}{\delta_0(\theta)} \right)^{-\kappa}, \tag{D5}$$

[42] In Bornstein (2025) the cost of switching brands is proportional to consumption, and a household's spending plays no part in whether it keeps a brand.

so that by (D4)

$$\frac{\lambda_{j,t}(\theta)}{\lambda_j(\theta)} \;=\; \left(\frac{E_{j,t}(\theta)/m_{j,t}(\theta)}{E_j(\theta)/m_j(\theta)}\right)^{-\kappa}, \tag{D6}$$

which is the path rule (16) at $q = E_{j,t}/m_{j,t}$. Following the shock there are two channels at play. A negative expenditure shock ends more relationships with elasticity $-\kappa\eta_j(\theta)$, however, after impact, each household spends more on each surviving relationship, which lowers its separation rate.

Taking on suppliers. Taking a supplier on takes t_f hours, the same for every household, and an hour is worth $w\theta$ to a type-θ household, its earnings from the hour, so a supplier costs $\chi(\theta) = w\theta t_f$ pounds at the stationary own indices: the opportunity cost of shopping time, which rises with the wage as in Aguiar and Hurst (2007). The household accepts a contact while a further supplier is worth at least this, $S_j \geq \chi(\theta)$, and since the value of a supplier falls as the set grows, it wants

$$m^*_{j,t}(\theta) \;=\; \frac{E_{j,t}(\theta)}{(\varepsilon-1)\,w\theta t_f} \tag{D7}$$

suppliers and takes on

$$A_{j,t}(\theta) \;=\; \min\Big\{\Gamma_{j,t}(\theta),\; \max\big\{0,\; m^*_{j,t}(\theta) - \hat{m}_{j,t}(\theta)\big\}\Big\}, \qquad \hat{m}_{j,t}(\theta) = \big(1-\bar{\psi}_{j,t}(\theta)\big)\big(1-\lambda_{j,t}(\theta)\big)m_{j,t}(\theta), \tag{D8}$$

every contact that reaches it while its set is below the target, which is the second term of (15). A household that spends more wants more suppliers, with elasticity $\eta_j(\theta)$, and a household below its target accepts every contact, so the suppliers it adds are a random draw from firms' contacts.

The stationary household. In a stationary equilibrium every open place is filled, so a type-θ household holds $m_j(\theta) = m^*_j(\theta)$ suppliers and takes on $[\bar{\psi}_j(\theta) + (1-\bar{\psi}_j(\theta))\lambda_j(\theta)]\, m_j(\theta)$ contacts each period, the relationships it loses, which is its weight in the pool (20). The value of a relationship then equals the cost of replacing it, $S_j(\theta) = w\theta t_f$, so by (D4)

$$\lambda_j(\theta) \;=\; 1 - G_\delta\big(w\theta t_f\big). \tag{D9}$$

A poorer household's time is worth fewer pounds, so a relationship is worth less to it and the same disruption costs end more of its relationships, which is the first restriction of Assumption 1.

The problem set out above delivers the three restrictions of Assumption 1 and the path rule (16): the separation rate falls with spending per supplier by (D4), contains no supplier's price because the household values every relationship at the sector average, falls with type in the stationary equilibrium by (D9), and moves with spending per supplier at the elasticity κ by (D6).

E Computation

The code is written in JAX and runs on a single GPU. A subscript k is the income bin, $j \in \{0, 1\}$ the necessity and luxury sectors, and λ_k the stationary separation rate of bin k (Table 4).

E.1 State space and grids

Section 5.1 details the discretisation of a firm's customer measure into a count n and composition $\boldsymbol{\pi}$ over $K = 5$ income bins, so that a firm's state variables are given by $(a, n, z, \boldsymbol{\pi})$. The bins are evaluated at the quintiles of the log-normally distributed labour efficiency parameters. The model discussed in the paper uses as grids income bins with $K = 5$ points, an age grid from 0 to an upper age of 45 where firms exit with probability one. The customer grid ranges from 0 to 30, which is normalised by the amount of labour available in the steady state. Time invariant productivity is on a geometrically spaced grid, the logarithm of which is distributed normally with 5 points. The firm's state space then features approximately 1.7 million states across the two sectors. It is cast as a single tensor so that the solution to firm and household policies, distribution updated and aggregation can be solved as one vectorised operation over the entire state space. Doing so means that each step can be solved independently across states offering large advantages over a multi core CPU. The bottleneck here is bandwidth due to the large state space rather than computational time.

E.2 The stationary equilibrium

The stationary equilibrium detailed in Section 4.5 is solved as two nested fixed points, the outer of which is a fixed point on the sectors' acquisition pools around a solve for the aggregate variables. The acquisition pool, as is the income distribution is truncated into K equal mass bins. Consequently the acquisition pool for sector j is given by:

$$F_{jk} \;\propto\; \left[\psi_{jk} + (1 - \psi_{jk})\lambda_k\right] m_{jk}, \tag{E10}$$

where ψ_{jk} is the exit rate of a type-k household's suppliers in sector j, which depends on the distribution of firms, and the sets m_{jk} are set in proportion to $\eta_{C,k}\, s_{jk}$. Given these pools, the remaining fifteen unknowns are the wage, the luxury price index (necessity goods are the numeraire), the mass of entrants in each sector, profits, and the ten price indices across household types p_{jk}. The labour market then clears by Walras' law. The computational algorithm is as follows:

1. Guess a set of the 15 unknowns, solve for the implied acquisition pools and separation rates
2. Find entrant and acquisition policies to determine the transitions of the distribution
3. Given the solutions to the previous steps, solve for the solution to the household block, solve for the firm block by backward iteration from terminal age, and iterate over the stationary distribution.

4. Update exit rates and pools from step 1
5. Continue until the pool share finds a fixed point

E.3 Transitions

I solve the transitions as a perfect foresight path for $T = 40$ years from the stationary equilibrium, defining $t = 0$ as 2010. A household of type k then spends $E_{k,t} = w\theta_k h_{k,t} + \Pi_t - \tau_t \ell_k E_k$, where $\tau_t =$ is the withdrawal of (33). The wage is held fixed at its stationary value throughout the transition, so hours evolve as (35) with Frisch elasticity $\epsilon_h = 1$. Separation rates evolve according to (16). Note that the shock here is implemented as a persistent but decaying income withdrawal, which is at odds with the actual structure of the UK welfare reforms, which were introduced as permanent reforms. The decay returns the economy towards its initial stationary equilibrium within the horizon, which the algorithm takes as its terminal condition. The algorithm to solve for the transitions is as follows:

1. Start with the unknowns at their steady state levels and the rate of expenditure reduction known
2. Compute hours, spending, and demand for each income bin K
3. Start from the terminal period and use a backward pass to solve for the firm's problem in $t = T - 1, \ldots, 0$ with period-t profits and withdrawal rates, record policies and entry masses.
4. Forward pass: From the stationary distribution in $t = 0$, apply the operator on the distribution given each period's policies, entrant masses and the known expenditure withdrawals.
5. Continue until convergence holds across the transition (in practice this takes between 30 and 40 iterations).

The homogeneous churn rate transitions are solved the same way but with one separation rate constant across time and households. The transition for one design takes approximately 30 minutes on a single NVIDIA GeForce RTX 4090 GPU (24 GB).

E.4 Event Study in the Model

To match the empirical event study to the model event study, I simulate a panel of 340 local authorities, with heterogeneous exposure according to the empirical distribution of treatment T_d. Note that I allow for no interaction between local authorities and so these are best thought of as differing levels of exposure, reflecting the argument in Section 3 that the fixed effects absorb any aggregate effects. I then repeat the event study in the data on the simulated firm panel.