

Thiemo Fetzer
Christina Palmou
Jakob Schneebacher
Ivan Yotzov

Working paper

26/44

How do firms cope with economic shocks in real time?

How do firms cope with economic shocks in real time?

Thiemo Fetzter

Christina Palmou

Jakob Schneebacher

Ivan Yotzov*

First version: September 18, 2024

This version: July 21, 2026

Abstract

In a turbulent world, effective policy requires granular estimates of firm responses to shocks as they unfold. We develop a replicable framework that combines high-frequency administrative and survey microdata with a pre-registered shift-share design to estimate multi-margin firm responses in real time. We show that the resulting coefficient vectors can be interpreted as first-order perturbations of firms' optimal-response functions. We apply the framework to the 2021–2023 UK energy price shock and find that energy-intensive firms pass on cost increases, build cash reserves and shift towards homeworking. We find little evidence of aggregate employment losses or firm exits. Responses are highly heterogeneous: small firms drive price pass-through; large firms, capital investment. We show that our real-time estimates are consistent with the UK's structural business survey, released two years later, and across survey instruments. We apply the framework out of sample and in real time to the energy price shock triggered by the 2026 US-Israeli strikes on Iran and discuss implications for the design of energy support schemes and environmental policy.

JEL Classification: C23; D22; D24; H23; L11

Keywords: real-time evaluation; firm dynamics; economic shocks; energy prices

*Fetzter is based at the University of Warwick and Bonn. Palmou is a PhD candidate at King's College London (KCL) and the Economic Statistics Centre of Excellence (ESCoE). Schneebacher is based at the Institute for Fiscal Studies and affiliated with KCL and ESCoE. Yotzov is based at the Bank of England (BoE). We are grateful to Monique Ebell, Kai Fischer, Guido Franco, Etienne Guige, Peter Lambert, Ludovic Panon, Sophie Piton, Joel Stiebale and audiences at the CEPR-Bank of Italy workshop 2023, the Kent workshop on firm dynamics 2024, the MaCCI, RES, ESCoE, EEA-ESEM annual conferences all in 2024, as well as the DMP workshop on "Business Conditions, Expectations and Uncertainty" 2025 and the EU-OECD conference on "Reigniting Growth" 2025 for comments. Any views expressed are solely those of the authors and do not necessarily represent those of the BoE. This paper uses Office for National Statistics research datasets via the Secure Research Service. Outputs may not exactly reproduce National Statistics aggregates. Fetzter acknowledges support from the CAGE Research Centre, the Leverhulme Prize in Economics, the European Research Council Starting Grant (ERC, MEGEO, 101042703) and the Deutsche Forschungsgemeinschaft (EconTribute, DFG, EXC 2126/1-390838866). Schneebacher acknowledges support from ESRC grant ES/Z504634/1. Corresponding author: Jakob Schneebacher (jakob.schneebacher@ifs.org.uk).


1 Introduction

Economic shocks arrive without warning: the Covid-19 pandemic shut down entire sectors overnight; Russia’s invasion of Ukraine quadrupled European wholesale energy prices within months; the 2026 US-Israeli strikes on Iran triggered an unforeseen disruption to oil shipments. In each case, policymakers needed to understand how firms responded not just *on average* but across the whole size and industry distribution, and across many adjustment margins, to design targeted support. However, the data infrastructure to provide these answers in real time rarely exists, and statistics based on structural business surveys are typically published with lags of one to two years (R. A. Decker & Haltiwanger, 2022; Fetzer et al., 2024; Feld & Fetzer, 2024). This lag creates space for anecdotes, media narratives and aggregate indicators that mask firm heterogeneity to influence the public policy discussion in the interim.

In this paper, we build a replicable framework for estimating how firms respond to economic shocks in near-real time. Three building blocks make this possible. First, we link high-frequency administrative and survey microdata from the UK Office for National Statistics (ONS), combining the universal coverage of administrative records with the richness of qualitative survey data. Second, we pre-register the analysis plan to discipline hypothesis generation across a large number of outcome margins. Third, we show that the coefficient vectors we estimate can be seen as entries of the comparative-statics Jacobian of the firm’s optimal-response function. This gives the multi-margin analysis a clear economic interpretation: each estimated coefficient corresponds to a first-order perturbation of the firm’s re-optimised choice along a specific margin.

We developed and applied the framework for the 2021–2023 UK energy price shock triggered by Russia’s invasion of Ukraine, which led to a large and sudden increase in energy prices. Figure 1 shows that within a few months, wholesale energy prices in the UK quadrupled. Differentially exposed firms, based on their pre-shock energy intensity, provide the identifying variation for a shift-share design (Bartik, 1991; Borusyak et al., 2022). But the methodology is not specific to this event. Any shock that generates differential firm-level exposure can be studied within this framework, as we demonstrate

Figure 1: System Average Price for natural gas in Great Britain



Notes: The figure displays the System Average Price (SAP) of all gas traded through the balancing market. The SAP aggregates all trades conducted on the On-The-Day Commodity Market. The horizontal grey line denotes November 2021.

with the oil price shock triggered by the 2026 US-Israeli strikes on Iran.

We then subject the real-time estimates to three validation exercises. First, we compare estimates from the Business Insights and Conditions Survey (BICS) to those from the Bank of England’s Decision Maker Panel (DMP), testing whether the two main high-frequency survey instruments available to UK policymakers give consistent signals. Second, we compare our survey-based estimates to the 2022 and 2023 Annual Business Survey (ABS), which provides a high-quality but slow-moving benchmark. Third, we apply the same methodology out of sample to the energy price shock triggered by the 2026 US-Israeli strikes on Iran. This last exercise is a proof of concept for generalisation: we built the framework for the Ukraine shock, and yet it produced sensible, real-time estimates for a different shock in a different geopolitical context.

Using quarterly administrative data, we find that on average firms exposed to the 2021–2023 energy shock do not reduce labour inputs. Only small firms are slightly more likely to exit the market. For the smaller sample of firms in the BICS, we additionally observe rich qualitative outcomes at high frequency. On average, energy-intensive firms

increase their output prices as they see the prices of their inputs rise. Turnover expectations adjust downwards, and firms shift towards working from home. We also find some evidence of precautionary liquidity management, including statistically significant changes to cash reserves, debt repayments and reported confidence in firms' ability to meet their debt obligations. We find important differences in how large and small firms react. The estimated price pass-through response of small firms is roughly twice that of large firms, consistent with imperfect pass-through by firms with market power. Only larger firms adjust their capital to a statistically significant degree. When we estimate coefficients for small industry cells and cluster across response margins, we see that firms' adjustments are highly heterogeneous by industry and correlated across margins. Applied out of sample to the 2026 oil shock, the framework delivers credible estimates within two months of the survey reference period. Specifically, we find that higher materials prices, output prices and stock levels have been the most significant adjustment margins for energy-intensive firms.

We make three contributions. The first and primary contribution is methodological. We show that combining real-time administrative and survey microdata with a pre-registered shift-share design can recover multi-margin firm responses to economic shocks with little delay. The approach exploits data products that already exist in many national statistical systems but are rarely linked or analysed at this frequency. Pre-registering quasi-experimental research is still uncommon in economics,¹ but [Brodeur et al. \(2024\)](#) find that a posted pre-analysis plan reduces the likelihood of p-hacking and publication bias.² Our theoretical framework links the estimated coefficients to entries of the comparative-statics Jacobian of the firm's optimal-response function (Section 2). [Chetty et al. \(2024\)](#), who build a public database to track consumer spending, employment and business revenues during Covid-19, are closest to us in real-time economic tracking. Our framework extends this approach in two directions: from descriptive tracking to causal estimation, and from aggregate or consumer outcomes to firm-level, multi-margin adjustment. [Almuzara & Sancibrián \(2024\)](#) provide the econometric foun-

¹For a recent exception, see [Clemens & Lewis \(2026\)](#).

²[C. Decker & Ottaviani \(2023\)](#) find that pre-registering trials reduces p-hacking for medical studies.

dations for the class of panel regressions we estimate, showing how to interpret heterogeneous coefficients and obtain valid inference when micro outcomes are regressed on aggregate shocks interacted with unit-level exposure. [Schneebacher et al. \(2025\)](#) apply a similar real-time, multi-margin evaluation framework to London’s Ultra Low Emission Zone, providing further evidence that the approach generalises across policy settings.

Our second contribution is to the literature on the impact of the recent energy price shocks. Research has explored fiscal ([Bachmann et al., 2022](#); [Auclert et al., 2023](#)), trade ([Itskhoki & Mukhin, 2022](#); [Babina et al., 2026](#)) and inflationary ([Lafrogne-Joussier et al., 2023](#)) effects of Russia’s war in Ukraine. We are among the first to investigate the real effects on firm behaviour across a variety of margins. Much existing evidence on firm responses to energy price shocks comes from the oil crises of the 1970s ([Kilian, 2008, 2014](#)). Even recent exceptions, including [Fontagné et al. \(2024\)](#) on French manufacturing and [Alpino et al. \(2023\)](#) on Italian firms, stop short of the recent shocks we analyse.

Third, our estimates have implications for the design of environmental policy. A proportional carbon tax on energy inputs operates through a similar channel as energy price shocks, though important differences remain: environmental taxation allows firms to prepare over a longer time horizon, and the associated uncertainty is more within policymakers’ control. Subject to these caveats, our short- and medium-run estimates are informative of the behavioural responses to targeted taxes and subsidies on the path to Net Zero ([Gillingham & Stock, 2018](#); [Glennerster & Jayachandran, 2023](#)). Our results are consistent with large firms acting as “shock absorbers” by incompletely passing cost increases through, possibly because they have greater market power.³

The paper proceeds as follows. Section 2 outlines our theoretical framework and our pre-analysis plan. Section 3 describes the data construction and main variables of interest, Section 4 our empirical strategy, and Section 5 our empirical results. Section 6 validates our real-time estimates and applies the framework to an out-of-sample shock. Section 7 discusses implications for policy design and a brief final section concludes.

³In finding evidence of incomplete pass-through our paper is consistent with a large literature. For instance, see [Gron & Swenson \(2000\)](#), [Nakamura & Zerom \(2010\)](#) and [Hong & Li \(2017\)](#).

2 Theoretical framework and pre-analysis plan

In this section, we make explicit what the estimated coefficient vector represents economically. The empirical object in the paper is a first-order response of multiple firm margins to an exposure-weighted energy-price shock, rather than a direct estimate of structural primitives.

Define the firm-specific effective shock as

$$\tau_{it} = \omega_i P_t^E, \quad (1)$$

where P_t^E is the common energy-price component and ω_i is pre-shock energy intensity. The identifying variation in the difference-in-differences (DiD) design comes from differential ω_i exposure to the same aggregate shock path.

At the firm level, consider a dynamic optimisation problem over controls $u_{it} = (p_{it}, q_{it}, m_{it}, \ell_{it}, k_{i,t+1}, b_{i,t+1}, \dots)$ and state x_{it} , subject to exogenous constraints and market conditions r_t . Let the optimal-response function be⁴

$$u_{it}^* = h(x_{it}, r_t, \tau_{it}), \quad x_{i,t+1} = g(x_{it}, u_{it}^*, \tau_{it}). \quad (2)$$

A compact flow-profit representation is

$$\pi_{it} = p_{it}q_{it} - w_t\ell_{it} - p_t^m m_{it} - \tau_{it}e_{it} - I(k_{i,t+1}, k_{it}) - \Phi(b_{i,t+1}), \quad (3)$$

with standard transition equations for inventories, liquidity, leverage and continuation. Appendix A provides more detail on this mapping.

The margins we estimate are measurements of this optimal system:

$$y_{it} = \mathcal{M}(x_{it}, u_{it}^*, \tau_{it}), \quad (4)$$

where y_{it} stacks turnover expectations, output and input prices, inventories, cash buffers,

⁴We write the optimum as a function for expositional clarity. The argument generalises to a set-valued correspondence under standard regularity conditions; see, for instance, [Milgrom & Segal \(2002\)](#).

capital mix, output-price expectations, working practices, debt, repayment capacity and survival. A local expansion around the pre-shock equilibrium gives

$$dy_{it} = \Gamma_{it} d\tau_{it}, \quad \Gamma_{it} \equiv \frac{dy_{it}}{d\tau_{it}}. \quad (5)$$

Define the sector-size cell average

$$\bar{\Gamma}_{s\ell} \equiv \mathbb{E}[\Gamma_{it} \mid i \in (s, \ell)]. \quad (6)$$

The DiD coefficient vector identifies reduced-form approximations to $\bar{\Gamma}_{s\ell}$. The coefficients are best read as first-order perturbations of the firm's *optimal-response function* around the pre-shock optimum (that is, local equilibrium shifts after re-optimisation), rather than derivatives of primitive objective functions with controls held fixed.

Applying the chain rule to profits:

$$\frac{d\pi_{it}}{d\tau_{it}} = \nabla_y \pi_{it} \cdot \frac{dy_{it}}{d\tau_{it}} + \pi_{\tau, it}, \quad (7)$$

with expanded form

$$\begin{aligned} \frac{d\pi_{it}}{d\tau_{it}} = & q_{it} \frac{dp_{it}}{d\tau_{it}} + p_{it} \frac{dq_{it}}{d\tau_{it}} - m_{it} \frac{dp_t^m}{d\tau_{it}} - p_t^m \frac{dm_{it}}{d\tau_{it}} - w_t \frac{d\ell_{it}}{d\tau_{it}} - e_{it} \\ & - \tau_{it} \frac{de_{it}}{d\tau_{it}} - I_{k'} \frac{dk_{i,t+1}}{d\tau_{it}} - \Phi_b \frac{db_{i,t+1}}{d\tau_{it}}. \end{aligned} \quad (8)$$

The direct effect $-e_{it}$ captures the mechanical cost increase at the pre-shock level of energy use; $-\tau_{it} \frac{de_{it}}{d\tau_{it}}$ captures the endogenous adjustment in energy quantity. Pricing, quantities, substitution, labour, adaptation and financing channels jointly determine the profit derivative with respect to the shock. Repayment and survival margins are continuation-value derivatives of the same system. Appendix A derives the relevant expressions.

Table 1 summarises the compact mapping from reduced-form estimates and the intuition behind it. Energy becomes more expensive for high- ω_i firms, and adjustment can occur across prices, volumes, input mix, processes, balance sheet and survival.

Table 1: Mapping from derivatives to empirical estimates

Margin	Theoretical object	Empirical proxy	Directional prior
Output pricing	$\frac{dp_{it}}{d\tau_{it}}$ and $\frac{d\mathbb{E}_t[p_{i,t+1}]}{d\tau_{it}}$	$\hat{\beta}_{s,\text{price_sold},b},$ $\hat{\beta}_{s,\text{p_sold_e},b}$	Pass-through pressure raises prices
Input cost	$\frac{dp_{it}^m}{d\tau_{it}}$	$\hat{\beta}_{s,\text{price_mat},b}$	Input prices rise with exposure
Output	$\frac{dq_{it}}{d\tau_{it}}$ and $\frac{d\mathbb{E}_t[q_{i,t+1}]}{d\tau_{it}}$	$\hat{\beta}_{s,\text{turnover_expn},b}$	Output and expectations weaken under cost pressure
Inventory	$\frac{dn_{i,t+1}}{d\tau_{it}}$ (holding n_{it} fixed)	$\hat{\beta}_{s,\text{stocklev2},b}$	Sign depends on buffer drawdown vs. precautionary build-up
Liquidity	$\frac{dc_{i,t+1}}{d\tau_{it}}, \frac{db_{i,t+1}}{d\tau_{it}}, \frac{d\rho_{i,t+1}}{d\tau_{it}}$	$\hat{\beta}_{s,\text{cash3},b},$ $\hat{\beta}_{s,\text{repayturn2},b}$	$\hat{\beta}_{s,\text{debt},b},$ Financial stress weakens liquidity position
Investment	$\frac{dk_{i,t+1}}{d\tau_{it}}$	$\hat{\beta}_{s,\text{capital_mix},b}$	Sign depends on postponement vs. adaptation investment
Process	Component of $\frac{du_{it}^*}{d\tau_{it}}$	$\hat{\beta}_{s,\text{work_fh},b}$	Firms may shift energy costs to employees via homeworking
Continuation risk	$\frac{d\Pr(\text{survival}_{i,t+1})}{d\tau_{it}}$	$\hat{\beta}_{s,\text{survival},b}$	Cost pressure lowers survival chances

2.1 The pre-registered analysis plan

Before commencing data linking and analysis in the ONS Secure Research Service (SRS), we publicly and irreversibly posted a pre-analysis plan (PAP) on 1 December 2022 at the Open Science Foundation (OSF) using the URL <https://osf.io/5entz/>.

Our goal was to create a transparent, public record of the order of hypothesis generation and testing. As per the initial PAP, and to further prevent data mining, we built the energy intensity measure E_i separately from the outcome dataset. An update to the PAP on 1 April 2023 set out additional predictions due to a change in energy support policies for UK firms. The evaluation of this policy change remains work in progress.

3 Data sources and key variables

3.1 Data structure

Unit of analysis. The unit of analysis in this project is an individual firm i , measured at the UK Office for National Statistics (ONS) reporting unit level (RUREF).⁵ A firm might have multiple establishments or local units, denoted by the local unit reference (LUREF) and may be part of an enterprise (ENTREF) or an enterprise group, denoted by the Who-Owns-What reference (WOWREF).

Energy data. We use the Annual Purchases Survey (APS) to obtain firm-level pre-shock energy cost shares. The APS records detailed information on annual input expenditures.⁶ The next subsection describes the APS in more detail. Energy-shock exposure variables follow the energy intensity measures reported in [Office for National Statistics \(2022\)](#).

Outcome data. We use two main outcome data sources. The first is the Business Insights and Conditions Survey (BICS), a large, voluntary business survey running since the early days of the Covid-19 pandemic. The second is the Longitudinal Business Database (LBD), a new administrative data product derived from the UK’s business register. This provides sparse but near-universal information for the UK’s business population. The next subsection gives more information on specific sources.

3.2 Data sources

The Annual Purchases Survey. The Annual Purchases Survey (APS) collects detailed information on business expenditures classified as intermediate consumption. We use

⁵More precisely, the reporting unit is the economically meaningful entity at which ONS surveys are administered. Some larger firms may consist of multiple reporting units.

⁶We also produced alternative estimates using the Annual Business Survey (ABS), the UK’s structural business survey. Because the ABS reports less fine-grained and potentially more noisy input information for individual purchase items such as energy for smaller firms, we consider the APS more appropriate for the purposes of this work.

the latest pre-energy crisis year⁷ available, APS 2018, to construct our baseline energy intensity measure, E_i . We define E_i as total energy expenditure divided by total purchases for intermediate consumption. If a firm is not sampled in APS 2018, we go back to APS 2017. For robustness, we also estimate results for an expanded sample where we impute energy input intensities as average energy intensities in the same industry-region-size cell for firms not present in either wave. Alternative measures using the full breakdown of specific energy inputs (natural gas, electricity, diesel) are available on request.

The Longitudinal Business Database. The Longitudinal Business Database (LBD) is an experimental ONS data infrastructure project, aimed at allowing the quick assembly of firm-level linked datasets from a variety of sources for microdata analysis (Lemma et al., 2023). The LBD is a quarterly, linked, firm-level dataset constructed using the UK’s business register, the Inter-Departmental Business Register (IDBR). The IDBR captures the universe of UK firms that either pay value-added tax (VAT) or contribute to pay-as-you-earn (PAYE) income tax schemes. In a given quarter, the IDBR captures approximately 3 million firms. The LBD is available in the ONS Secure Research Service (SRS). We use survival, employment, and establishment count at a quarterly frequency as outcome variables in our analysis. The LBD also contains turnover (which can be used to construct labour productivity). However, since turnover (and to a lesser extent employment) is recorded on the IDBR from a variety of sources and with variable lags, we do not make use of it in this analysis.

The Business Insights and Conditions Survey. The Business Insights and Conditions Survey (BICS) is a qualitative, fortnightly, voluntary, topical business survey established during the pandemic. The BICS is sent to approximately 50,000 businesses every two weeks, with a response rate of roughly 25%. Large businesses often receive many consecutive survey waves, whereas small firms are rotated in and out to reduce survey burden. The BICS provides rich and timely information for a sample of the UK busi-

⁷The APS was re-introduced as a survey in 2015 and runs annually. Processing for the 2019 and 2020 APS was de-prioritised during the Covid-19 pandemic and data for these years have not been published. Source: [ONS Quality and Methodology Information, 2023](#).

ness population. We use two main types of variables from the BICS. First, we use a set of questions on firm input, output, pricing and innovation behaviour, as well as firms' subjective measures of financial and economic health, as the main outcome variables in this project. The exact variables of interest as well as how we construct the BICS panel can be found in the online appendix. Second, we use questions about a firm's climate attitudes and behaviour as well as their perceived exposure to the energy shock.⁸

The voluntary nature of the BICS and its 25% response rate raise a potential concern about selective non-response. If energy-intensive firms under financial stress are differentially more or less likely to respond to the survey after the shock, our treatment effect estimates could be biased. We address this concern by validating our BICS-based estimates against the UK's structural business survey, the Annual Business Survey (ABS), in Section 6.2. Because the ABS is drawn from the same register with near-universal coverage of large firms, agreement between the two sets of estimates bounds the scope for non-response bias in the average effects.

3.3 Data linking and cleaning


Figure 2 summarises the data linking process. We link BICS waves 1-89, corresponding to March 2020 to August 2023, by RUREF to our units of analysis. We harmonise responses to BICS questions on variables of interest, over time. For some variables a sufficiently long pre- and post-treatment panel exists to test anticipation effects. For others, we only have a post-treatment panel or even a single cross section.

We then link BICS respondents to the APS 2018 and 2017 using the latest available data point for each RUREF. The resulting sample consists of 9,550 firms and 199,083 observations. We strip our energy intensity measure of outliers by trimming observations below the first and above the 99th percentile in each two-digit SIC (Standard Industrial Classification) industry.⁹ This leaves 9,225 firms and 192,999 observations.

⁸As a robustness check, we also use intermittently asked BICS questions on the length and coverage of firms' energy contracts to construct an alternative energy intensity measure that interacts our baseline measure with a dummy for contracts that expire in the near term or cover energy costs only partially.

⁹By removing within-industry outliers, we effectively remove any firms with zero or near-zero energy purchases. This raises the mean energy intensity compared to the full sample (see online appendix).

Figure 2: Data linking and cleaning



Notes: Figure illustrates the data components and the associated sample sizes that we work with in the analysis as they arise when we merge different datasets together.

3.4 Key variable construction

Energy intensity measure. We measure energy intensity E_i as a share of total intermediate expenditure. This raises three concerns. First, there may be firm-specific idiosyncratic variation in the energy cost shares that is not due to the technical requirements in the production processes used by firm i . Second, firms may be exposed to the energy price shocks *indirectly*, through the costs they pay for other inputs if the production of these inputs is energy-intensive (or in turn requires energy-intensive inputs). Third, to the extent that some firms have long-term energy contracts or do not pay for their energy costs directly, their energy input may not reflect the financial cost they bear.¹⁰ While we

¹⁰Firms may also be shielded from full cost increases by government support. In the UK, the government introduced the Energy Bill Relief Scheme (EBRS) from October 2022 to March 2023, which capped the wholesale element of non-domestic energy bills at 21.1p/kWh for electricity and 7.5p/kWh for gas. It was replaced by the less generous Energy Bills Discount Scheme from April 2023. Eligibility required a

do not observe the necessary data to shed light on these hypotheses at the firm level, we investigate them here with industry-level data. Figure B.1 in the appendix shows that the distribution of energy intensities varies substantially across firm size-bands within industries, confirming that the energy intensity measure carries within-industry variation across firms of different sizes. A further concern is that E_i is measured from the APS 2018 (or 2017), three to four years before the shock. Firms may have changed their energy intensity over the intervening period through investment, output mix changes or efficiency improvements. To the extent that this introduces classical measurement error in E_i , our treatment effect estimates will be attenuated towards zero.

Validating the energy intensity measure. We validate our cost-share measure of energy intensity against four alternative industry-level measures in the appendix. Direct calorific energy quantities from the ONS environmental accounts are highly correlated with our measure (Figure B.2). Indirect energy intensity, computed by inverting the ONS input-output supply-and-use tables, is likewise highly correlated with direct exposure (Figure B.3).¹¹ Real-time energy expenditure payment flows from [Office for National Statistics \(2023a\)](#) are positively correlated with our measure (Figure B.4), confirming that firms with higher energy use are indeed more financially exposed. Finally, per-unit prices of gas and electricity rose sharply across all consumption bands (Figure B.5), suggesting that long-term contracts did not fully shield firms from the shock.

4 Empirical methodology

This section explains the empirical methodology we employ. It describes the estimating equation and assumptions required for causal inference, outlines how we analyse heterogeneity and finally explains how we explore correlations across adjustment margins by clustering on industry-specific estimates.

non-domestic contract agreed on or after 1 December 2021. We therefore measure the *net*, rather than *gross*, impact of the energy price shock.

¹¹See [ONS energy use by industry, source and fuel, 2024](#) and [ONS input-output supply-and-use tables, 2023](#).

4.1 Estimating equation

We pre-registered our main estimating equation in the initial Pre-Analysis Plan (PAP). This estimating equation takes the following form:

$$Y_{ist} = \alpha_i + \beta_{st} + \xi \cdot Post_t \cdot E_i + \nu \cdot X_{ist} + \varepsilon_{ist} \quad (9)$$

α_i captures a firm i fixed effect, absorbing any time-invariant firm characteristics. We also control for time-fixed effects β_{st} . These account for common time-varying non-linear shocks. These time-fixed effects could be specific to a group the firm belongs to, such as the industry or region that a firm i operates in, as indicated by subscript s .¹²

E_i denotes our energy intensity measure. The indicator $Post_t$ is a binary variable that captures the time period after the energy price shock. Figure 1 indicates that by November 2021 an invasion was seen as likely, leading to an anticipatory spike in the energy price. We therefore use this as our baseline treatment. Since November 2021, spot market prices for natural gas have averaged around 7 p/kWh. This is 4.3 times the average for the period from 2018 to October 2021 inclusive. For robustness, we also estimate our results with the treatment time defined as 24 February 2022, the day the invasion began. The term X_{ist} allows for the inclusion of additional control variables.

Logic of the estimation protocol. Three design choices govern how far equation (9) can be saturated: the granularity of the industry mapping $s = g(i)$,¹³ whether energy intensity is measured at the firm level or imputed at the industry level, $E_{g(i)}$; and the functional form of the exposure measure. A continuous measure assumes that the treatment effect increases monotonically in E_i ; the alternative discretises E_i relative to the within-industry median and contrasts more and less energy-intensive firms within the same industry. How demanding a specification the data can support depends on the intra-industry variation in E_i and on the size and distribution of the effects. We did not know either when we formulated our PAP. We therefore estimate the full menu of

¹²For conciseness, in the rest of the paper we refer to a firm grouping s as “industry”.

¹³The UK’s Standard Industrial Classification (SIC) distinguishes between 732 industry codes. See [Companies House, 2024](#).

Figure 3: Set of estimable two-way fixed effects models

		Unit fixed effects				
		Firm only	SIC	Region	SIC & Region	SIC $\times$ Region
Time fixed effects	Wave only	x				
	SIC1		x	x	x	x
	SIC2		x	x	x	x
	SIC3		x	x	x	x
	SIC4		x	x	x	x
	SIC5		x	x	x	x

Notes: Figure illustrates the menu of two-way fixed effects specifications that can be estimated. The full set of results across all specifications is reported in the online appendix.

specifications, from least to most demanding, and report results across all of them.

Figure 3 illustrates the set of estimable two-way fixed effects models. As the granularity of the fixed effects increases, the estimand ξ on $Post_t \cdot E_i$ may become collinear with industry-by-time fixed effects if energy intensity varies predominantly *between* rather than *within* industries. Stable estimand sizes across specifications indicate a robust result; attenuation under more demanding fixed effects may reflect collinearity or measurement error. We can also estimate industry-aggregated models, which reduce the risk of demand spillovers from treated to untreated firms that may violate the Stable Unit Treatment Value Assumption (SUTVA) (Alves et al., 2024). Such spillovers are a concern here: if energy-intensive firms raise output prices, their less-energy-intensive competitors within the same industry may gain market share, attenuating the estimated treatment effect. The progressively more demanding fixed effects specifications help assess the sensitivity of results to within-industry spillovers. The full set of results across all specifications is reported in the online appendix.

Estimation of the treatment effect. We estimate the causal effect by interacting common, unanticipated shocks to energy prices with pre-existing variation in exposure, as measured by a firm’s pre-shock energy intensity. In the empirical specification the estimated coefficient ξ on the interaction term $Post_t \cdot E_i$ captures this effect. In the notation of Section 2, E_i corresponds to the firm’s pre-shock energy intensity ω_i and ξ recovers elements of the sector-averaged Jacobian $\bar{\Gamma}_{s\ell}$. This identification strategy is in the spirit of

the shift-share approach pioneered by [Bartik \(1991\)](#) and recently characterised in much fuller detail by [Borusyak et al. \(2022\)](#). The key identifying assumption is that, conditional on firm and industry-by-time fixed effects, pre-shock energy intensity E_i affects post-shock outcomes only through the energy price channel. This would be violated if, for instance, energy-intensive firms were differentially exposed to other concurrent shocks (such as supply-chain disruptions) that are not absorbed by the fixed effects. The progressively more demanding fixed effects specifications reported in the online appendix help assess sensitivity to such confounds. For brevity we report here the treatment effect estimates obtained from the continuous exposure measure E_i , after confirming that the monotonicity assumption in the treatment effects appears to hold.

4.2 Methods for higher-dimensional analysis and clustering

Most econometric analysis focuses on a single adjustment margin. Yet as Section 2 formalises, a firm’s response to a cost shock is jointly determined across all margins. We therefore estimate treatment effects separately for small industry cells and use the resulting high-dimensional coefficient vectors to discover joint distribution patterns.

To achieve this, we estimate heterogeneous treatment effects across a broad range of industry cells resulting in a vector of estimates $\boldsymbol{\zeta}_s = (\zeta_{s1}, \dots, \zeta_{sk})$ capturing the relationship between the energy-price shock exposure of firms in industry cell s on an outcome variable j as ζ_{sj} . Each element ζ_{sj} is the analogue of the scalar ζ from equation (9), estimated separately for industry cell s and outcome j ; in the notation of the conceptual framework, $\boldsymbol{\zeta}_s$ approximates the sector-averaged Jacobian $\bar{\Gamma}_{s\ell}$. We focus on eleven different outcomes, estimating each for dozens of different industry cells. The estimated vectors $\boldsymbol{\zeta}_s$ can be thought of as capturing the joint distribution of the impact of the energy price shock on a representative firm within industry cell s . As before, we carry out the analysis at different industry granularities.

Clustering on the $\boldsymbol{\zeta}_s$ vector We employ a k-means clustering algorithm to identify common archetypes or patterns in how variables co-move across industries in response

to the energy price shock.¹⁴ We then examine this heterogeneity by industry characteristics to explain differences found, such as why some industries adjust capital while others increase stock levels. K-means clustering allows us to compute Euclidean distances between the industry-level coefficient vectors ζ_s and group industries with high intra-cluster and low inter-cluster similarity in their response profiles.¹⁵

5 Results

This section presents our main results, proceeding from pre-trends through average treatment effects and industry heterogeneity to the higher-dimensional analysis.

By way of motivation, Figures B.6 and B.7 in the appendix show the only direct, self-reported evidence on firms' mitigation behaviour. They confirm that energy-intensive firms were significantly more likely to report taking actions to reduce energy use after the price shock, including switching equipment off, adopting more energy-efficient equipment and reducing trading hours. Because the questions were only asked on one wave each, these results are descriptive rather than causal.


5.1 Pre-trends

Figure 4 reports pre-trends for our key outcome variables using an event study design in the spirit of [Sun & Abraham \(2021\)](#). For most margins (prices, capital, debt, working practices) there are no clear pre-trends, with the trend break at or after our baseline treatment date of November 2021. A small number of forward-looking margins (material input prices, cash reserves, stock levels, insolvency risk) show slight anticipation effects before November 2021. This is consistent with precautionary behaviour: firms anticipating higher energy costs may pre-emptively build cash buffers or reassess insolvency risk, whereas realised price and capital adjustments require the shock to materialise.

¹⁴The joint distributions found in the data can serve as a check on the often simple one-dimensional economic narratives that attract media attention.

¹⁵Hierarchical clustering techniques yield comparable results. We prefer k-means for its geometric interpretability in high-dimensional vector space.

Figure 4: Pre-trends for key outcome margins



Notes: Figures present event-study coefficients from the baseline specification, interacting the continuous exposure measure E_i with wave indicators and including firm and wave fixed effects. Standard errors are clustered at the industry-by-wave level. 95% confidence bands indicated.

5.2 Average treatment effects

Tables 2 to 4 report coefficients from our main specification using the linked waves of the Business Insights and Conditions Survey (BICS).¹⁶ These coefficients correspond to the entries of the sector-averaged Jacobian $\bar{\Gamma}_{s\ell}$ described in Section 2 and mapped to empirical proxies in Table 1. We find statistically significant effects on the following adjustment margins: output and input prices; capital and working practices; cash reserves; and expectations over future turnover and survival.

Panel A of Table 2 shows average treatment effects for price variables. The coefficient of interest for materials prices (which include energy) is 0.370 for all firms. At 0.326, it is slightly smaller for large firms (defined as firms with 250 or more employees) than for small and medium-sized firms (SMEs, defined as firms with fewer than 250 employees), for whom the coefficient is 0.456. On average, energy-intensive firms increase their output prices as they see the prices of their inputs rise. This average treatment effect is driven by small firms (coefficient of 0.389 and significant at the 1% level) but not large firms (coefficient of 0.185 and not statistically significant). These results are consistent with incomplete pass-through by firms with market power (proxied by size) and complete pass-through by a competitive fringe of small firms. Future price expectations are directionally similar but attenuated and only marginally significant for small firms.

As Panel B of Table 2 shows, there were no immediate turnover effects but turnover expectations varied between large and small firms. Large firms anticipated declines in revenue due to the energy cost increases (coefficient of -0.173, statistically significant at the 5% level). However, small firms did not expect decreases to the same degree. These differences may be due to differences in forecasting abilities, because small firms feel insulated through supply contracts or because they view themselves as price takers. Small firms reported higher confidence in their survival prospects despite raising stock levels and not adjusting capital. Large firms became less optimistic despite making

¹⁶The BICS outcome variables are ordinal indicators of the direction of change during the reference period, typically with three to five categories. We estimate linear models on these ordinal outcomes. This recovers average marginal effects that closely approximate those from ordered probit or logit specifications when the number of categories is small and the distribution of responses is not concentrated at the boundaries, as is generally the case in our setting.

Table 2: Impacts of the energy price shock on prices and output, for large and small firms

Panel A: Average treatment effects, firm prices

Prices		Estimate (ζ Treatment $\times$ Energy intensity)		
		all	large	sme
Price of materials	ζ	0.370***	0.326**	0.456***
	<i>se</i>	(0.103)	(0.141)	(0.0988)
	R^2	0.392	0.359	0.434
	N	94321	58584	35706
Price of goods sold	ζ	0.249**	0.185	0.389***
	<i>se</i>	(0.115)	(0.129)	(0.119)
	R^2	0.350	0.321	0.392
	N	97954	61229	36687
Prices of goods sold expectations	ζ	0.162*	0.163	0.155*
	<i>se</i>	(0.0908)	(0.112)	(0.0844)
	R^2	0.365	0.335	0.404
	N	66702	42277	24353

Panel B: Average treatment effects, firm output

Output		Estimate (ζ Treatment $\times$ Energy intensity)		
		all	large	sme
Turnover change (3 cat)	ζ	0.0781	0.0831	0.0685
	<i>se</i>	(0.112)	(0.126)	(0.190)
	R^2	0.374	0.361	0.397
	N	115592	72300	43244
Turnover expectations (3 cat)	ζ	-0.164***	-0.173**	-0.157*
	<i>se</i>	(0.0573)	(0.0694)	(0.0855)
	R^2	0.274	0.257	0.304
	N	103876	65011	38865
Export status (3 cat)	ζ	-0.0704	-0.0464	-0.100
	<i>se</i>	(0.0538)	(0.0715)	(0.0606)
	R^2	0.942	0.938	0.946
	N	75718	45647	30071

Notes: Estimated treatment effects from regressions with firm and wave fixed effects. Standard errors are clustered by wave and two-digit SIC industry. Outcome variables are ordinal indicators of the direction of change during the reference period. Standard errors are reported in parentheses. Stars denote statistical significance obtained from estimating clustered standard errors with stars indicating *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$

Table 3: Impacts of the energy price shock on inputs and processes, for large and small firms

Panel A: Average treatment effects, firm input mix **Panel B:** Average treatment effects, firm processes

Input mix		Estimate (ζ Treatment $\times$ Energy intensity)		
		all	large	sme
Capital	ζ	0.161	0.152	0.220
	<i>se</i>	(0.105)	(0.180)	(0.231)
	R^2	0.644	0.624	0.676
	N	32354	21217	11137
Capital mix	ζ	0.225**	0.262**	0.203
	<i>se</i>	(0.0943)	(0.120)	(0.163)
	R^2	0.490	0.472	0.518
	N	37758	24722	13036
Redundancies (share)	ζ	0.141	0.00461	0.413
	<i>se</i>	(0.321)	(0.241)	(0.675)
	R^2	0.230	0.214	0.254
	N	138778	87177	51601
Redundancy expectations	ζ	0.0125	-0.0158	0.0526
	<i>se</i>	(0.0337)	(0.0481)	(0.0359)
	R^2	0.501	0.479	0.542
	N	43003	25538	17465

Production process		Estimate (ζ Treatment $\times$ Energy intensity)		
		all	large	sme
Stock levels	ζ	0.0329	-0.0256	0.164
	<i>se</i>	(0.135)	(0.181)	(0.120)
	R^2	0.359	0.335	0.397
	N	68668	42271	26397
Hybrid working	ζ	-5.769**	-6.404	-5.131
	<i>se</i>	(2.816)	(4.714)	(3.126)
	R^2	0.787	0.768	0.817
	N	58129	35349	22780
Working from home	ζ	51.39***	63.11***	30.00***
	<i>se</i>	(10.96)	(15.19)	(7.288)
	R^2	0.725	0.718	0.734
	N	126124	79111	47013
Working from normal place of work	ζ	16.63**	13.07	20.34**
	<i>se</i>	(8.160)	(9.261)	(8.627)
	R^2	0.739	0.741	0.732
	N	126124	79111	47013

Notes: Estimated treatment effects from regressions with firm and wave fixed effects. Standard errors are clustered by wave and two-digit SIC industry. Outcome variables are ordinal indicators of the direction of change during the reference period. Standard errors are reported in parentheses. Stars denote statistical significance obtained from estimating clustered standard errors with stars indicating *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$

Table 4: Impacts of the energy price shock on liquidity and survival, for large and small firms

Panel A: Average treatment effects, firm liquidity **Panel B:** Average treatment effects, firm survival

Debt & liquidity		Estimate (ζ Treatment $\times$ Energy intensity)			Survival		Estimate (ζ Treatment $\times$ Energy intensity)		
		all	large	sme			all	large	sme
Cash reserve duration (5 cat)	ζ	0.263*	0.245	0.288**	Insolvency risk	ζ	-0.162**	-0.0800	-0.263**
	<i>se</i>	(0.140)	(0.177)	(0.135)		<i>se</i>	(0.0800)	(0.0967)	(0.100)
	R^2	0.787	0.786	0.786		R^2	0.666	0.664	0.669
	N	77355	48748	28607		N	81177	49908	31269
Confidence will meet obligations (5 cat)	ζ	-0.278***	-0.266**	-0.293**	Change in insolvency risk	ζ	0.0239	0.0361	0.0274
	<i>se</i>	(0.0870)	(0.101)	(0.125)		<i>se</i>	(0.0773)	(0.108)	(0.0835)
	R^2	0.663	0.645	0.690		R^2	0.500	0.475	0.542
	N	40204	24308	15896		N	40037	24954	15083
Repayments over turnover (5 cat)	ζ	-0.438***	-0.231	-0.778***	Confidence of 3m survival	ζ	0.0803	0.0369	0.169*
	<i>se</i>	(0.149)	(0.171)	(0.200)		<i>se</i>	(0.0593)	(0.0735)	(0.0877)
	R^2	0.693	0.677	0.716		R^2	0.703	0.692	0.714
	N	26521	15612	10909		N	59784	37104	22680

Notes: The cash reserve duration, confidence the firm will meet debt obligations and repayments compared to turnover variables have 5 possible responses. Estimated treatment effects from regressions with firm and wave fixed effects. Standard errors are clustered by wave and two-digit SIC industry. Outcome variables are ordinal indicators of the direction of change during the reference period. Standard errors are reported in parentheses. Stars denote statistical significance obtained from estimating clustered standard errors with stars indicating *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$

capital adjustments. This divergence between behaviours and expectations may reflect differences in information sets or planning horizons.

Panel A of Table 3 shows the average treatment effects for a firm's input mix. Overall, firms adjust planned and realised capital expenditure (capital mix coefficient of 0.225, statistically significant at the 5% level). This effect is driven by large firms. The effect for small firms is smaller and not statistically significant. This is consistent with larger firms anticipating the need to adapt their production process to the energy cost increases. Realised redundancies and redundancy expectations do not respond for either large or small firms, and in fact yield precisely estimated null effects in most cases.

Panel B of Table 3 shows the effects on firm processes: the stock levels held on premises (a measure of their efficiency and expectations about future uncertainty) and their homeworking practices. The latter is particularly relevant as homeworking may allow firms to offload energy costs onto their employees. We do not find statistically significant effects on stock levels, overall or for either size category. Working from home increases overall (coefficient of 51.39, statistically significant at the 1% level), and for both large (coefficient of 63.11, statistically significant at the 1% level) and small firms (coefficient of 30.00, statistically significant at the 1% level). This increase in working from home is accompanied by a decrease in hybrid working.¹⁷

The shift towards working from home among energy-intensive firms admits several interpretations. First, firms may be directly reducing their commercial energy consumption by closing offices or reducing operating hours; the associated heating, lighting and equipment costs are then borne by employees at home. Second, the shift may reflect a broader cost-cutting response in which firms reduce expenditure on commercial premises. Third, in principle energy-intensive firms could have been on a different working-from-home trajectory for reasons unrelated to the energy shock. However, we note that the pre-trends for working practices in Figure 4 show no evidence of such differential trends before our treatment date.

Finally, Table 4 shows the average treatment effects on firm liquidity and survival.

¹⁷Among small firms, working practices become bimodal: both fully remote and fully on-site attendance increase (office attendance coefficient of 20.34, statistically significant at the 5% level), while hybrid arrangements decline. We do not observe the same pattern for large firms.

Panel A focuses on liquidity. On average, firms built up cash reserves. This is driven predominantly by small firms (coefficient of 0.288, statistically significant at the 5% level). All firms were less confident they would meet their debt obligations (coefficient of -0.278, statistically significant at the 1% level). The estimated effects for large and small firms are similar in size and statistical significance in this regard. Affected firms increased their debt repayments relative to turnover (coefficient of -0.438, statistically significant at the 1% level), consistent with precautionary liquidity management. This overall effect is driven predominantly by small firms (coefficient of -0.778, statistically significant at the 1% level).

Panel B focuses on firm survival. We find significant decreases in the perceived level of insolvency risk (coefficient of -0.162, statistically significant at the 5% level) driven predominantly by smaller firms. This is likely related to the increase in cash reserves and debt repayment we observe among the same firms. The average treatment effects we estimate are small and statistically insignificant for firms' confidence to survive the next three months and the change in insolvency risk, although both are estimated on smaller samples.

Using quarterly data on employment, number of business sites and survival from the Longitudinal Business Database (LBD) we find similar results, shown in Table B.1 in the appendix. We find no statistically significant drops in firm-level employment for affected firms following the energy price shock. The LBD results do show an increase in firm exits among more energy-intensive small firms.

We estimate effects on approximately fifteen outcome variables, separately for all firms, large firms and small firms, raising concerns about multiple hypothesis testing. As our conceptual framework in Section 2 makes explicit, all outcome margins are entries of the same Jacobian Γ_{it} , generated by the firm's joint re-optimisation decisions in response to the shock. A firm that raises output prices ($\hat{\beta}_p > 0$) will, all else equal, also change its debt position and cash reserves. The k-means clustering analysis in Section 5 recovers this joint structure. We therefore interpret our results as correlated rather than independent responses, and focus on the consistency of the estimated pattern across

margins rather than on the statistical significance of any individual coefficient.¹⁸

For our baseline outcome variables, we also exploit the panel nature of most questions on the BICS to investigate the dynamics of firms' response to the energy price shock. Figure B.8 in the appendix plots the dynamics at quarterly frequency. The results reveal important differences between short- and long-run responses and between large and small firms. Smaller firms adjusted their prices almost instantaneously, with the effect peaking in April to June 2022 and fading by summer. Turnover expectations dropped following the peak in price rises. Debt repayments relative to turnover rose after April 2022 and peaked during the summer.


5.3 Heterogeneous industry effects

The average effects estimated in the previous subsection mask substantial heterogeneity. Firms adjust differently on different margins based on their production function, market structure or management capabilities. In addition to firm size, in our PAP we conjectured that industry is an important dimension of heterogeneity. In this subsection, we show that there is indeed large variation between fine-grained industries across all margins. Figure 5 plots the distribution of the estimated coefficients from restricted sample regressions over two-digit industries. For all margins, the variation across two-digit industries swamps the difference between estimated average effects and zero.

Not only do firms display substantial heterogeneity, but responses are correlated across response margins in predictable ways. This is consistent with the conceptual framework in Section 2: because all margins are jointly determined by the firm's re-optimisation, the entries of $\bar{\Gamma}_{s\ell}$ will in general covary across sectors. Figure 6 plots the estimated responses across response margins j , $\tilde{\zeta}_s = (\tilde{\zeta}_{s1}, \dots, \tilde{\zeta}_{sj}, \dots, \tilde{\zeta}_{sk})$, for two example sectors, manufacturing and hospitality. Even at high levels of aggregation, firms across these industries display starkly different adaptation behaviours. For instance, manufacturing sees large adaptations on indebtedness, debt repayment, turnover expectations

¹⁸Formal joint inference across margins could potentially account for cross-equation correlation, but the ordinal nature of the outcomes and the rotating panel structure of the BICS make this infeasible in our setting.

Figure 5: Coefficient distributions over industries




Notes: Figure shows the distribution of average treatment effects estimated separately for each industry division against the sample-wide average treatment effect denoted by the dotted line.

and stock levels, but a moderate capital response and no response in their cash reserves. Accommodation and food services see large impacts on their stock levels, indebtedness and survival, but no capital investment. Figure B.10 shows that these correlations across margins are not particular to the two example sectors, but apply more broadly: for instance, wholesale and retail trade sees large impacts on input and output prices, and substantial adjustment on capital and debt margins.

Adaptations likewise differ substantially by firm size. Figure B.9 in the appendix shows that large firms are more likely to adjust on output price, capital, cash, and debt repayment margins. By contrast, small firms are more likely to adjust stock levels and to see their survival prospects worsen.

Figure 6: Adjustments across margins: manufacturing and hospitality



Notes: Each axis represents a different adjustment margin. Values show the estimated average treatment effect of energy intensity on that margin for the industry sector indicated. Effects are estimated from the baseline specification with firm and wave fixed effects. The shaded area connects the estimated coefficients across margins to illustrate the response profile.

5.4 Higher-dimensional analysis and clustering

Because each firm reoptimises jointly across its choice variables (Section 2), the estimated treatment effects are systematically correlated across response margins. To uncover the structure of these joint distributions, this section reports results from a simple k-means clustering exercise. This descriptive exercise summarises the joint distribution of estimated treatment effects across margins and sectors but does not identify the structural determinants of differences across clusters.

The algorithm uncovers six distinct groupings of industries. Based on their responses, we call them “the passthrough-financed investors”, “the survivors”, “the cruisers”, “the disinvestors”, “the wrecked” and “the cash-constrained”. Appendix table B.2 and Figure B.11 give more detail about the industries in each cluster and the average treatment effect along each margin. Here, we focus on industry characteristics that may explain the differences in responses, as shown in Table 5.

“Passthrough-financed investors” are predominantly drawn from ICT, engineering and real estate and represent about 12% of economy-wide Gross Value Added (GVA). The typical firm in this cluster is small, well-managed and relatively investment-intensive.

Table 5: Characteristics of firms in each cluster

	The passthrough-financed investors	The survivors	The cruisers	The disinvestors	The wrecked	The cash constrained
Description	ICT, engineering and real estate	Complex manufacturing chemicals, metals, food products, and non-business services such as human health and education	Manufacturing of wood, textiles, electrical equipment, and business services such as wholesale and retail trade and land transport, telecoms, and waste collection	Pharmaceuticals, furniture, arts, and motion picture and other professional scientific activities	Printing and reproduction media	Scientific research and development and advertising
Energy intensity	13%	14%	10%	10%	8%	7%
Average turnover (£000)	1,821	8,800	5,269	7,453	795	3,584
Average employment	9	35	23	18	8	15
Management practices	0.54	0.50	0.49	0.48	0.47	0.53
Investment intensity (NGFCF/GVA)	15%	19%	16%	8%	10%	6%
Concentration index (Turnover HHI)	537	740	1106	635	1358	381
Average profit margin	0.24	0.18	0.12	0.26	0.20	0.22
Average IC markup	1.40	1.38	1.27	1.40	1.19	1.28
Average labour markup	2.27	1.84	2.18	3.02	2.03	4.18
EU trade intensity	3%	9.5%	10.4%	1.9%	2.2%	9.6%
Non-EU trade intensity	3.3%	9.2%	9.2%	5%	2.3%	11.8%
Number of divisions	9	27	13	5	1	2
% GVA	12.1%	32.7%	15.9%	6.7%	1.2%	2.9%

Notes: The table presents weighted averages of each industry division in each cluster. Average employment and turnover are the author's calculations using the 2021 IDBR universe, investment intensity is the author's calculations using net gross fixed capital formation over GVA at basic prices from the 2021 ABS, trade intensity is calculated as the average ratio of goods imports and exports to each divisions 2021 turnover. Management practices are taken from the 2020 Management and Expectations Survey, and turnover concentration is based on CMA calculations using the 2021 Business Structure Database. % of GVA in each cluster is calculated using the UK 2021 GVA by industry division published by the ONS in "Regional gross value added (balanced) by industry" ([Office for National Statistics, 2024](#)). The average profit margin, and mark-ups are from the ONS' dataset accompanying the publication "Trends in UK business dynamism" calculated using the 2021 ABS ([Office for National Statistics, 2023b](#)).

In line with the small firm size, industry concentration in this cluster is quite low but markups and profit margins are nonetheless on the upper end of the spectrum. “The survivors” in contrast represent the most energy-intensive cluster and make up almost 33% of GVA, including business services, health and education and some manufacturing. This cluster features on average larger firms by turnover and employment and the highest investment intensity. Concentration, profit margins and markups are among the lowest of all clusters. “The cruisers” have intermediate energy intensities of about 10% and represent a further 16% of GVA, including wholesale and retail, telecoms and waste collection. They are characterised by some concentration but low profit margins and markups, and high exposure to international trade. Firms in this cluster fall in the upper range of investment and come from the middle of the management distribution.

Finally, there are three smaller clusters of industries that adjust less well to the energy price shock, namely “the disinvestors” (7% of GVA, including pharmaceuticals, arts and some scientific activities), “the wrecked” (1% of GVA, printing and reproduction media) and “the cash constrained” (3% of GVA, scientific research and advertising). These industries feature mostly lower management quality, considerably lower investment intensity and relatively high profit margins.

6 Validating the real-time estimates

A central claim of this paper is that the combination of real-time survey and administrative data can provide reliable, timely estimates of firm responses to economic shocks. In this section, we subject this claim to three tests: consistency of the BICS estimates with the Bank of England’s Decision Maker Panel (DMP); ex-post benchmarking against the 2022 and 2023 Annual Business Survey (ABS); and an out-of-sample application to the energy price shock triggered by the 2026 US-Israeli strikes on Iran.

6.1 Comparing across high-frequency survey vehicles

The BICS, which serves as the main rapid-response business survey vehicle of HM Treasury and other central government departments, and the Decision Maker Panel (DMP), which is run by the Bank of England, are two key real-time data sources used by UK policymakers. If they produce divergent signals about how firms are responding to an economic shock, this misalignment could complicate policymakers' situational assessment and thus the ensuing policy response. In this subsection, we test whether the BICS and DMP give consistent estimates of firm responses to the 2021-2023 energy price shock.

The lack of a common sampling frame and common unique identifier between the DMP and ONS datasets prevent us from computing firm-level APS energy intensities, as we do for the BICS.¹⁹ We therefore compute them at the industry level, averaging the APS measure within SIC cells (two-digit through to five-digit). We then estimate the same shift-share specification on the DMP firms, using these industry-level energy intensities as the exposure variable. Because all firms within an industry cell share the same exposure, these are effectively between-industry regressions rather than the within-industry regressions reported in the main results.

We compare the resulting BICS and DMP coefficient vectors for each outcome margin that is available in both surveys. Table 6 provides a mapping of all the outcome variables which can be compared in the DMP and BICS. Because the BICS captures qualitative answers while the DMP measures quantitative responses, a direct comparison of the magnitudes of the estimates is not possible. Instead, for each pair we consider whether the coefficients agree or disagree on sign.

Figure 7 shows a scatterplot comparing the treatment effects across the two surveys. In the DMP specifications, we use energy intensity measured at the SIC4 industry level, and additionally include firm and time fixed effects. The BICS coefficients are the main

¹⁹Firms in the DMP are sampled from the population of UK businesses with ten or more employees in the Bureau van Dijk FAME database, which itself covers all UK incorporated employee firms. For more information on the survey design, see [Bunn et al. \(2024\)](#). The DMP also collects information at a higher level of the business hierarchy compared to ONS surveys (the enterprise reference, or “entref”, level rather than the reporting unit reference, or “ruref”, level, in ONS terminology).

Table 6: Mapping of DMP survey outcomes to their corresponding BICS variables


DMP outcome	BICS variable(s)	Sign
Annual sales growth	Turnover change (3 cat); Turnover change (6 cat)	✓
Annual real sales growth	Turnover change (3 cat); Turnover change (6 cat)	✓
Expected sales growth	Turnover expectations (5 cat); Turnover expectations (3 cat)	✓
Expected real sales growth	Turnover expectations (5 cat); Turnover expectations (3 cat)	✓
Annual output price growth	Price of goods sold	✓
Expected output price growth	Prices of goods sold expectations	✓
Annual capital expenditure growth	Capital	✓
Expected capital expenditure growth	Capital mix	✓
Annual employment growth	Redundancies (share)	✗
Expected employment growth	Share of workers expect to make redundant; Redundancy expectations	✗

Notes: Each DMP outcome is compared against one or more BICS variables. A “✗” in the Sign column indicates that the BICS variable measures the opposite direction to its mapped DMP outcome (e.g. redundancies rise when expected employment growth falls); the sign of its coefficient is reversed before comparison so that it lines up with the DMP outcome.

estimates from Section 5.2. The blue markers in the figure correspond to specifications which include all firms; green markers correspond to specifications which include large firms with 250 or more employees; and purple markers correspond to specifications for firms with fewer than 250 employees.

The results indicate strong sign alignment on treatment effects across the BICS and DMP. For specifications with all firms, 80% of coefficients agree on sign and lie in the top-right or bottom-left quadrants of the figure. The alignment is even stronger for large firms, where we find agreement in 93% of specifications. Meanwhile, for smaller firms we find agreement in around 60% of specifications. In general, agreement across data sources is consistently lower for smaller firms. One reason for this may be that larger firms are more likely to respond, on average, and therefore provide more consistent estimates. In the BICS, all firms with 250 or more employees are included in the sample, while firms with fewer than 250 employees are sampled using a random stratified design.

Figure 7: Comparison of BICS and DMP average treatment effects




Notes: Figure plots estimated BICS coefficients against DMP coefficients for the same outcome margin. Blue markers show coefficient pairs for specifications across all firms; green markers show coefficient pairs for specifications of large firms (250+ employees) only; purple markers show coefficient pairs for specifications estimated on small firms (< 250 employees). The top-right and bottom-left quadrants indicate coefficient pairs which agree on sign across the two surveys; the top-left and bottom-right quadrants contain coefficient pairs which disagree on sign. Treatment effects agree on sign in 80% of specifications with all firms, 93% of specifications with large firms, and 60% of specifications with SMEs.

We conclude that the two surveys give consistent signals about the direction of firm responses across a large majority of outcomes. This is particularly true for large firms, which are likely to have the strongest effects on aggregate outcomes. The finding is reassuring: it means that the real-time estimates available to fiscal and monetary policymakers would have pointed broadly in the same direction during the crisis, despite differences in sampling frames, survey design and response populations.

6.2 Comparing real-time estimates to the “ground truth”

The ABS is the UK’s structural business survey and, unlike the DMP, is drawn from the business register. We can therefore link the same firm-level APS energy shares to ABS respondents and estimate the shift-share specification of equation (9) exactly as

Figure 8: BICS and ABS average-effect comparison (baseline specification)



Notes: The figure plots the standardised average effect of energy intensity on each outcome margin in the BICS (horizontal) against the ABS (vertical), for the margins available in both surveys. Coefficients are estimated from the baseline shift-share specification with firm and wave fixed effects, using firm-level APS energy shares. Markers are coloured by estimation sample: all firms, large firms (250 or more employees) and SMEs. Points in the upper-right and lower-left quadrants indicate sign agreement. ABS treatment effects are standardised for comparability with the BICS coefficients. Treatment effects agree on sign in 100% of specifications with all firms, 64% of specifications with large firms, and 36% of specifications with SMEs.

we do on the BICS.²⁰ This exercise benchmarks the real-time BICS estimates against a near-universal, high-quality survey that became available only two and a half years later. Because BICS outcomes are ordinal and ABS outcomes cardinal, we standardise coefficients and compare their signs across the two surveys rather than their magnitude.

Figure 8 reports the comparison. In the pooled sample the two surveys agree on the sign of the average effect for every common outcome.²¹ Agreement weakens once we split by firm size, as we found for the BICS-DMP comparisons: the two surveys agree on roughly two-thirds of margins for large firms (64%) and roughly one-third

²⁰We restrict the comparison to BICS outcomes that have a direct counterpart in the ABS. The ABS is annual, so we estimate average effects over the 2022 and 2023 waves.

²¹The outcomes for which we are able to provide a direct comparison across the BICS and ABS are turnover, capital expenditure, stock levels, exports, imports, and redundancies.

for SMEs (36%). The survey design suggests a likely culprit for this gradient. The BICS rotates small firms in and out to limit response burden, while large firms receive many consecutive waves (Section 3). SME coefficients are therefore estimated on thinner panels, and their signs are more easily flipped by sampling noise. For the pooled and large-firm estimates, where the BICS panel is densest, the real-time survey estimates line up well with the ABS. These findings support the use of the BICS to track average firm responses in real time. For pooled and large-firm averages, the agreement also suggests that neither selective non-response in the voluntary BICS nor the categorical nature of its responses drives the estimated direction of effects. The validation is weaker for SME-specific estimates: with six comparable margins and a 36% sign-agreement rate, the SME heterogeneity results rest more heavily on the BICS alone, and we flag them as correspondingly less certain. The comparison also tests signs rather than magnitudes, so it speaks to the direction of average effects rather than to their size.²²

6.3 Out-of-sample application: the 2026 US-Israeli strikes on Iran

The 2026 US-Israeli strikes on Iran provide a real-time, out-of-sample test of the framework. The shock differs from the 2021–2023 energy crisis in its origin, its primary fuel and its transmission: where Russia’s invasion of Ukraine worked through European natural gas markets over many months, the strikes on Iran produced a sharp rise in global oil prices from 28 February 2026 onwards. We carried over the exact estimating equation (9), the firm-level APS energy intensity measure E_i and the linking pipeline unchanged, re-defining only the treatment date and refreshing the exposure measure to the latest pre-shock APS vintage (2023).²³


The estimation sample comprises BICS waves 130–157, of which waves 153–157, covering February–April 2026, fall after the shock.²⁴ For reference, BICS wave 157 had a

²²In the online appendix, we also compare the coefficient estimates between the DMP and ABS. We find stronger sign alignment for specifications with all firms (88%). Moreover, sign alignment is stronger for large firms (75% of regressions) than small firms (62% of specifications).

²³ E_i measures total energy intensity, so exposure to the primarily oil-driven 2026 shock is likely measured with error for firms whose energy mix is gas- or electricity-heavy, and the Iran coefficients may be attenuated as a result. Decomposing exposure by fuel source, and incorporating indirect exposure through input-output linkages, is the subject of ongoing work.

²⁴Wave 153’s reference period is 1–28 February 2026.

Figure 9: Firm responses to the 2021–22 and 2026 energy price shocks



Notes: Figure compares estimated treatment effects for the 2021–22 energy price shock (estimated over December 2021 to February 2022, matching the post-shock horizon available for the 2026 shock) and the 2026 oil price shock (estimated over late February to late April 2026). Bars indicate 95% confidence intervals from industry-by-wave clustered standard errors. The estimates for the 2026 oil price shock are based on data up to April 2026.

reference period of 1–30 April, was in the field 18–31 May, was made available to users on 4 June, and we obtained our cleared results on 17 June. This means we were able to obtain our elasticity estimates within two months from the survey reference period. Future estimates can be obtained on the same timescale.

As for the Ukraine shock, we use a 12-month pre-invasion period to assess pre-trends. To compare like-for-like, we benchmark the Iran estimates against Ukraine-shock coefficients estimated on a similarly-sized post-window (covering December 2021 to February 2022) rather than the full post-period used in our main estimates.

Figure 9 reports the comparison for all firms. The price margins respond consistently across the two shocks: exposed firms report rising input prices and pass a substantial part of the increase into output prices, with magnitudes similar to the equivalent horizon of the Ukraine shock. The margins that differ do so in economically interpretable ways.

Exposed firms report higher realised turnover in 2026, consistent with price-driven nominal revenue growth,²⁵ and build stock levels, consistent with precautionary stockpiling under uncertainty about supply routes. By contrast, the forward-looking margins that responded strongly to the sustained gas shock, output-price expectations and capital expenditure expectations, show no significant response at this horizon. This is consistent with firms interpreting the US-Iran conflict, and the resulting oil shock, as more transitory than Russia’s invasion of Ukraine (and the resulting gas shock). Pass-through is again concentrated among smaller firms (Figure B.12 in the appendix).

This exercise demonstrates the central claim of the paper: a framework built for one shock produced credible multi-margin estimates for a different shock, fuel source and geopolitical context, within seven weeks of the reference period and with no modification. Together with the survey cross-validation and the ex-post ABS benchmarking, this completes the case that real-time firm-level shock evaluation is feasible with existing statistical infrastructure.

7 Implications for policy design

The primary implication of this paper is that real-time, multi-margin estimation of firm responses to economic shocks is feasible and replicable. The framework requires three ingredients: high-frequency microdata with differential firm-level exposure to the shock, a pre-registered analysis plan and a conceptual framework that gives the resulting coefficients economic meaning. In the UK, these ingredients can be assembled from existing ONS data products. Many other national statistical systems collect comparable data. The validation exercises in Section 6 provide evidence that the real-time estimates are reliable: they can be validated ex post by structural business surveys, they are consistent across survey instruments and they produce sensible results when applied out of sample to a new shock. Policymakers who invest in this data infrastructure can evaluate firm responses to shocks as they unfold, rather than waiting one to two years for structural surveys to be processed. Advances in open access to real-time data sources, combined

²⁵Appendix figure B.13 shows that turnover and price coefficients are highly positively correlated.

with judicious supplementary surveys, will make this type of pre-planned, real-time evaluation easier and eventually commonplace.²⁶

The estimates themselves also have direct policy relevance. Even one year after Russia’s invasion, the direct energy price shock had not led to widespread exits for more affected firms. But debt loads increased and firms built cash buffers, threatening long-term firm survival. Large firms invested in capital to adapt to the shock and passed through less of the input price increase, acting as “shock absorbers”. Small firms relied more on stock levels, cash reserves and price pass-through. This implies that business support in an energy crisis may need to be targeted to the specific needs of different segments of the business population. When we cluster industries on their multi-margin response profiles, we see that not all firms cope equally well. The observed heterogeneity forces policymakers to make their value judgements explicit: is the aim to help the economy adapt as quickly as possible, or to provide aid to struggling businesses?

To make the heterogeneous estimates accessible, we have built an interactive web tool at <https://firmnarratives.trfetzner.com/>. The tool allows users to explore the estimated treatment effects by industry, firm size and adjustment margin, and translates the industry-level coefficient vectors ζ_s into written narratives in a variety of registers.

Our estimated responses also have implications beyond the immediate energy crisis. A proportional carbon tax on energy inputs operates through a similar price channel to the shock we study. However, three important differences may limit the direct applicability of our estimates. First, environmental taxation will generally allow firms to prepare and adapt over a longer time horizon, particularly given the gradual nature of conventionally proposed “climate ramps” (Nordhaus, 2007; Weitzman, 2007; Stern, 2008). Second, the uncertainty around future energy prices due to environmental taxation lies to a much larger extent within the control of policymakers (Bloom et al., 2007; Bloom, 2009, 2014). Third, our energy intensity measure may not perfectly capture the relevant exposure to a carbon tax: we document that direct energy expenditure is highly correlated with both calorific measures and with wider indirect energy exposure at the

²⁶In principle, any economic shock with differential firm-level exposure can be studied in this framework, provided the national statistical system collects high-frequency administrative or survey firm data.

industry level, but do not observe these measures at the firm level. Subject to these caveats, our estimated short- and medium-run effects are informative of the behavioural responses to targeted taxes and subsidies a large literature argues will be necessary on the path to Net Zero (Acemoglu et al., 2012; Aghion et al., 2016; Acemoglu et al., 2023). We provide the granular estimates to make such targeting possible.

8 Conclusion

We develop a replicable framework for estimating near-real time firm responses to economic shocks. It estimates a pre-registered shift-share design on high-frequency administrative and survey microdata and interprets the resulting multi-margin coefficient vectors as entries of the comparative-statics Jacobian of firms' optimal-response functions.

Applied to the 2021–2023 UK energy price shock, the framework reveals that affected firms passed through cost increases, built cash reserves and shifted towards homeworking. Administrative data show no aggregate employment decline; the adjustment on the extensive margin was concentrated in increased exit among energy-intensive small firms, alongside modest employment reductions in large firms. Responses were highly heterogeneous: small firms drove price pass-through while large firms drove capital investment. When we estimate coefficients for small industry cells and cluster across response margins, we uncover distinct firm-adaptation archetypes that differ systematically in production technology, market structure and management quality.

Three validation exercises support the reliability of the real-time estimates. Our estimates are consistent across the two main survey instruments available to UK policymakers. They also align with the structural business survey published 30 months later. And the framework produces sensible estimates when applied out of sample to the 2026 Iran shock. This last exercise demonstrates that the approach is not specific to the Ukraine energy crisis but generalises to new shocks in new contexts.

More broadly, we show that real-time, multi-margin, heterogeneous estimation of firm responses to economic shocks is feasible with the data infrastructure that already exists in many national statistical systems. The approach is replicable in any setting with

high-frequency administrative or survey data at the firm level and variation in shock exposure. Policymakers who invest in linking and exploiting these data products can evaluate shocks as they unfold and design targeted support with little delay. The energy-crisis estimates also provide granular estimates that are informative for the design of carbon taxes and business support on the path to Net Zero.

References

- Acemoglu, D., Aghion, P., Barrage, L., & Hémous, D. (2023). *Climate change, directed innovation, and energy transition: The long-run consequences of the shale gas revolution* (NBER Working Paper No. 31657). National Bureau of Economic Research.
- Acemoglu, D., Aghion, P., Bursztyn, L., & Hémous, D. (2012). The environment and directed technical change. *American Economic Review*, 102(1), 131–166.
- Aghion, P., Dechezleprêtre, A., Hémous, D., Martin, R., & Van Reenen, J. (2016). Carbon taxes, path dependency, and directed technical change: Evidence from the auto industry. *Journal of Political Economy*, 124(1), 1–51.
- Almuzara, M., & Sancibrián, V. (2024). *Micro responses to macro shocks* (Staff Reports No. 1090). Federal Reserve Bank of New York. (Revised July 2025) doi: 10.59576/sr.1090
- Alpino, M., Citino, L., & Frigo, A. (2023). *The effects of the 2021 energy crisis on medium-sized and large industrial firms: Evidence from Italy* (Questioni di Economia e Finanza (Occasional Papers) No. 776). Bank of Italy.
- Alves, G., Burton, W. H., & Fleitas, S. (2024). *Difference-in-differences in equilibrium: Evidence from place-based policies* (CEPR Discussion Paper No. 18916). Centre for Economic Policy Research.
- Auclert, A., Monnery, H., Rognlie, M., & Straub, L. (2023). *Managing an energy shock: Fiscal and monetary policy* (NBER Working Paper No. 31543). National Bureau of Economic Research.
- Babina, T., Hilgenstock, B., Itskhoki, O., Mironov, M., Ribakova, E., & Shapoval, N. (2026). Russian oil exports under international sanctions. *Energy Economics*, 153, 109006. doi: 10.1016/j.eneco.2025.109006
- Bachmann, R., Baqaee, D., Bayer, C., Kuhn, M., Löschel, A., Moll, B., ... Schularick, M. (2022). *What if? the economic effects for Germany of a stop of energy imports from Russia* (Policy Brief No. 028). ECONtribute.
- Bartik, T. J. (1991). *Who benefits from state and local economic development policies?* Kalamazoo, MI: W.E. Upjohn Institute for Employment Research.
- Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica*, 77(3), 623–685.
- Bloom, N. (2014). Fluctuations in uncertainty. *Journal of Economic Perspectives*, 28(2), 153–176.
- Bloom, N., Bond, S., & Van Reenen, J. (2007). Uncertainty and investment dynamics. *The Review of Economic Studies*, 74(2), 391–415.
- Borusyak, K., Hull, P., & Jaravel, X. (2022). Quasi-experimental shift-share research designs. *The Review of Economic Studies*, 89(1), 181–213.

- Brodeur, A., Cook, N., Hartley, J., & Heyes, A. (2024). Do preregistration and preanalysis plans reduce p-hacking and publication bias? Evidence from 15,992 test statistics and suggestions for improvement. *Journal of Political Economy Microeconomics*, 2(3), 527–561. doi: 10.1086/730455
- Bunn, P., Bloom, N., Crundwell, A., Khan, S., Menzies, C., Mizen, P., . . . Yotzov, I. (2024). The decision maker panel: a user’s guide. *Bank of England Staff Working Paper*(1,096).
- Chetty, R., Friedman, J. N., Stepner, M., & Opportunity Insights Team. (2024). The economic impacts of COVID-19: Evidence from a new public database built using private sector data. *The Quarterly Journal of Economics*, 139(2), 829–889. doi: 10.1093/qje/qjad048
- Clemens, M. A., & Lewis, E. G. (2026). The effect of low-skill immigration restrictions on US firms and workers: Evidence from a randomized lottery. *American Economic Journal: Applied Economics*. (Forthcoming) doi: 10.1257/app.20250049
- Decker, C., & Ottaviani, M. (2023). *Preregistration and credibility of clinical trials* (CEPR Discussion Paper No. 18180). Centre for Economic Policy Research.
- Decker, R. A., & Haltiwanger, J. (2022, may). *Business entry and exit in the COVID-19 pandemic: A preliminary look at official data* (FEDS Notes No. 2022-05-06-2). Board of Governors of the Federal Reserve System (U.S.). Retrieved from <https://doi.org/10.17016/2380-7172.3129> doi: 10.17016/2380-7172.3129
- Feld, I., & Fetzer, T. (2024). *Performative state capacity and climate (in)action* (CEPR Discussion Paper No. 18894). Centre for Economic Policy Research.
- Fetzer, T., Shaw, C., & Edenhofer, J. (2024). *Informational boundaries of the state* (CEPR Discussion Paper No. 18773). Centre for Economic Policy Research.
- Fontagné, L., Martin, P., & Orefice, G. (2024). The many channels of firm’s adjustment to energy shocks: Evidence from France. *Economic Policy*, 39(117), 5–43. doi: 10.1093/epolic/eiae011
- Gillingham, K., & Stock, J. H. (2018). The cost of reducing greenhouse gas emissions. *Journal of Economic Perspectives*, 32(4), 53–72.
- Glennester, R., & Jayachandran, S. (2023). Think globally, act globally: Opportunities to mitigate greenhouse gas emissions in low- and middle-income countries. *Journal of Economic Perspectives*, 37(3), 111–135.
- Gron, A., & Swenson, D. L. (2000). Cost pass-through in the US automobile market. *Review of Economics and Statistics*, 82(2), 316–324.
- Hong, G. H., & Li, N. (2017). Market structure and cost pass-through in retail. *Review of Economics and Statistics*, 99(1), 151–166.

- Itskhoki, O., & Mukhin, D. (2022). *Sanctions and the exchange rate* (NBER Working Paper No. 30009). National Bureau of Economic Research.
- Kilian, L. (2008). The economic effects of energy price shocks. *Journal of Economic Literature*, 46(4), 871–909.
- Kilian, L. (2014). Oil price shocks: Causes and consequences. *Annual Review of Resource Economics*, 6(1), 133–154.
- Lafrogne-Joussier, R., Martin, J., & Mejean, I. (2023). *Energy cost pass-through and the rise of inflation: Evidence from French manufacturing firms* (CEPR Discussion Paper No. 18596). Centre for Economic Policy Research.
- Lemma, Y., Lui, S., Romaniuk, L., Schneebacher, J., & Wolf, N. (2023). *The UK Longitudinal Business Database* (ESCoE Technical Report No. TR-23). Economic Statistics Centre of Excellence.
- Milgrom, P., & Segal, I. (2002). Envelope theorems for arbitrary choice sets. *Econometrica*, 70(2), 583–601.
- Nakamura, E., & Zerom, D. (2010). Accounting for incomplete pass-through. *The Review of Economic Studies*, 77(3), 1192–1230.
- Nordhaus, W. D. (2007). A review of the Stern Review on the economics of climate change. *Journal of Economic Literature*, 45(3), 686–702.
- Office for National Statistics. (2022). *Business energy spending: experimental measures from the Office for National Statistics' business surveys*. Retrieved from www.ons.gov.uk.
- Office for National Statistics. (2023a). *Industry-to-industry payment flows, UK: 2016 to 2023, experimental data and insights*. Retrieved from www.ons.gov.uk.
- Office for National Statistics. (2023b). *Summary statistics of profit margins, intermediate consumption markups and labour markups from the Annual Business Survey, UK*. Retrieved from www.ons.gov.uk.
- Office for National Statistics. (2024). *Regional gross value added (balanced) by industry: all ITL regions*. Retrieved from www.ons.gov.uk.
- Schneebacher, J., Jabbar, F., & Kariel, J. (2025). *How do commuters adapt to local pollution pricing?* (IFS Working Paper No. 25/57). Institute for Fiscal Studies.
- Stern, N. (2008). The economics of climate change. *American Economic Review*, 98(2), 1–37.
- Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2), 175–199.
- Weitzman, M. L. (2007). A review of the Stern Review on the economics of climate change. *Journal of Economic Literature*, 45(3), 703–724.

A Mapping derivatives to estimates

This appendix spells out the derivative map in more detail. Let

$$\Xi_{it} \equiv \frac{dx_{it}}{d\tau_{it}}, \quad H_{it} \equiv \frac{du_{it}^*}{d\tau_{it}}.$$

From the measurement map $y_{it} = \mathcal{M}(x_{it}, u_{it}^*, \tau_{it})$, the total derivative is

$$\Gamma_{it} \equiv \frac{dy_{it}}{d\tau_{it}} = \mathcal{M}_x \Xi_{it} + \mathcal{M}_u H_{it} + \mathcal{M}_\tau. \quad (10)$$

From the optimal-response function $u_{it}^* = h(x_{it}, r_t, \tau_{it})$ and state recursion $x_{i,t+1} = g(x_{it}, u_{it}^*, \tau_{it})$:

$$H_{it} = h_x \Xi_{it} + h_\tau, \quad \Xi_{i,t+1} = g_x \Xi_{it} + g_u H_{it} + g_\tau. \quad (11)$$

Observed reduced-form derivatives therefore combine direct effects, endogenous re-optimisation and dynamic state propagation. Take inventories as an example:

$$n_{i,t+1} = n_{it} + q_{it} - \tilde{q}_{it} \implies \frac{dn_{i,t+1}}{d\tau_{it}} = \frac{dq_{it}}{d\tau_{it}} - \frac{d\tilde{q}_{it}}{d\tau_{it}}, \quad (12)$$

where q_{it} denotes production and $\tilde{q}_{it}$ denotes sales, holding beginning-of-period inventories n_{it} fixed.

For continuation outcomes, if

$$\rho_{i,t+1} = \mathcal{R}(c_{i,t+1}, b_{i,t+1}, \dots), \quad \Pr(\text{survival}_{i,t+1}) = \mathcal{S}(c_{i,t+1}, b_{i,t+1}, k_{i,t+1}, \dots),$$

then

$$\begin{aligned} \frac{d\rho_{i,t+1}}{d\tau_{it}} &= \mathcal{R}_c \frac{dc_{i,t+1}}{d\tau_{it}} + \mathcal{R}_b \frac{db_{i,t+1}}{d\tau_{it}} + \dots, \\ \frac{d\Pr(\text{survival}_{i,t+1})}{d\tau_{it}} &= \mathcal{S}_c \frac{dc_{i,t+1}}{d\tau_{it}} + \mathcal{S}_b \frac{db_{i,t+1}}{d\tau_{it}} + \mathcal{S}_k \frac{dk_{i,t+1}}{d\tau_{it}} + \dots. \end{aligned}$$

To link this map to the stacked empirical object, define branch-specific sector vectors

$$\hat{\beta}_{sb} = (\hat{\beta}_{s1b}, \dots, \hat{\beta}_{sMb})'.$$


Under standard linearisation and branch-specific regressor scaling,

$$\hat{\beta}_{sb} \approx D_b \bar{\Gamma}_s, \quad (13)$$

where $\bar{\Gamma}_s$ is the average sector derivative vector and D_b is a diagonal matrix of scaling constants that absorbs normalisation differences across branch definitions. With continuous energy intensity, D_b is approximately the identity; with a binary indicator (above or below the within-industry median), each diagonal element rescales the corresponding derivative by the local difference in mean exposure between treated and control firms.


B Figures and tables

Figure B.1: Distribution of industry group average energy intensities by employment size-band




Notes: Figure plots the distribution of industry group average energy intensities separately by employment size-band. Energy intensity is measured as the share of energy expenditure in total purchases for intermediate consumption from the APS.

Figure B.2: Gross calorific tones of energy used compared to APS energy intensity measure by 2-digit industry




Notes: Figure plots the gross calorific tonnes of energy used by each 2-digit SIC industry, taken from the ONS environmental accounts on energy use by industry, source and fuel, against the survey-based APS energy intensity measure that we use in the analysis. Each point represents a 2-digit industry.

Figure B.3: Indirect energy intensity against direct energy intensity measure




Notes: Figure plots the direct energy intensity measure that we obtain from the APS against indirect energy intensity from input-output relationships and embodied energy intensity in intermediate good consumption.

Figure B.4: Energy intensity in the APS against BACS transaction energy intensity




Notes: Figure plots the energy intensity from the survey-based measure that we use across industries on the horizontal against a payment-flow based energy intensity measure that can be derived from industry-to-utility transaction level payment flows data. We observe that the survey based measure and the transaction-level data are highly correlated.

Figure B.5: Prices of gas and electricity purchased by non-domestic consumers in different consumption bands in the United Kingdom




Notes: The data come from "Non-Domestic Energy Prices" statistics published by the Department for Energy Security and Net Zero on 28th March 2024. The prices include the Climate Change Levy (CCL) equivalent to 0.7p/kWh for electricity and 0.4-0.7p/kWh for gas. The CCL does not affect the relative trends between the consumers of different size bands. For gas, very small consumption corresponds to annual consumption less than 278 MWh, small - 278-2,777, medium - 2,778-27,777, large - 27,778 - 277,777, very large - 277,778 - 1,111,112. For electricity: very small corresponds to annual consumption of 0-20 MWh, small - 20-499, small/medium - 500 - 1,999, medium - 2,000 - 19,999, large - 20,000 - 69,999, very large - 70,000 - 150,000, extra large - > 150,000.

Figure B.6: Relationship between energy expenditure share and direct energy mitigation measures




Notes: Figure shows the coefficient of energy intensity in a cross-sectional regression of firms' responses to the question "What actions, if any, has your business taken to reduce energy costs in the last three months?" in Wave 75. Each response is coded as zero if a business has not reported taking this action and one if it has. We cannot distinguish non-response to the question from reporting not taking this action. We control for the firm's employment size band, 5-digit industry and region. Standard errors are clustered by 2-digit industry.

Figure B.7: Relationship between energy expenditure share and self-reported impacts of the energy price rise.




Notes: Figure shows the coefficient of energy intensity in a cross-sectional regression of firms' responses to the question "How has your business been affected by recent energy price rises?" asked every 2 waves between Waves 52 - 89. Each response is coded as zero if a business has not reported taking this action and one if it has. We can't distinguish non-response to the question from reporting not taking this action. We control for the firm's employment size band, 5-digit industry and region. Standard errors are clustered by 2-digit industry.

Figure B.8: Dynamic effects for key margins of adjustment, by firm size




Notes: Each panel plots dynamic treatment effect coefficients of energy intensity for one outcome margin, estimated separately for all firms, large firms (250+ employees) and SMEs (< 250 employees). Effects are estimated from the baseline specification with firm and wave fixed effects. Standard errors are clustered by wave and two-digit SIC industry. 95% confidence bands indicated.

Figure B.9: Average coefficients by firm size: baseline outcomes



Notes: Each axis represents a different adjustment margin. Values show the standardised estimated average treatment effect of energy intensity on that margin for all firms, large firms (250+ employees) and SMEs (< 250 employees). Effects are estimated from the baseline specification with firm and wave fixed effects. The shaded areas connect the estimated coefficients across margins to illustrate the response profile of each firm size group.

Figure B.10: Adjustments across margins by industry




Notes: Each axis represents a different adjustment margin. Values show the estimated average treatment effect of energy intensity on that margin for the industry section indicated in each panel. Effects are estimated from the baseline specification with firm and wave fixed effects. The shaded area connects the estimated coefficients across margins to illustrate each industry's response profile. The corresponding profiles for manufacturing and accommodation and food services are shown in Figure 6.

Figure B.11: Cluster archetypes

The passthrough-financed investors (9 divisions)	-0.14	1.23	0.94	0.33	0.43	1.06	0.54	-0.08	0.92	-0.19
The survivors (27 divisions)	-0.10	-0.04	0.16	0.18	0.05	0.20	-0.22	0.33	-0.78	-0.09
The cruisers (13 divisions)	0.50	-0.43	0.28	0.02	-0.34	-0.65	-0.70	-1.66	0.26	0.09
The disinvestors (5 divisions)	-1.06	0.64	1.65	1.57	1.36	5.08	-1.34	-1.82	-9.37	-0.13
The wrecked (1 division)	7.08	-6.20	-0.93	-2.57	7.04	2.66	-0.10	3.08	14.57	-0.62
The cash constrained (2 divisions)	1.04	-0.42	-2.08	-0.69	3.75	-1.72	-0.56	-2.24	20.48	-0.15
	Turnover expectations	Price of goods sold	Prices of materials	Price expectations	Stock levels	Cash reserves	Capital mix	Confidence will meet debt repayments	Repayments compared to turnover	Survival


Notes: Figure shows the average treatment effect along each outcome margin for each of the six clusters identified by the k-means algorithm. Clusters are formed by grouping industry divisions on the vectors of industry-level treatment effects ξ_s , as described in Section 4. Table B.2 lists the industry divisions assigned to each cluster and Table 5 summarises their characteristics.

Figure B.12: Responses to the 2022 and 2026 shocks, by firm size



Notes: Each axis represents a different adjustment margin. Radar plots show estimated treatment effects for all firms, large firms (250 or more employees) and SMEs, for the 2021–2022 shock (matched December 2021 to February 2022 window) and the 2026 shock. The red circle marks zero. Specification as in Figure 9. The estimates for the 2026 oil price shock are based on data up to April 2026.

Figure B.13: Scatterplot of industry-level output price and turnover responses, 2026 shock



Notes: Scatterplot of industry-level output price and turnover treatment effects, from the 2026 shock. The estimates for the 2026 oil price shock are based on data up to April 2026.

Table B.1: Average treatment effects for firms' survival in the LBD

Survival (LBD)		Estimate		
		(ξ Treatment × Energy intensity)		
		all	large	sme
Log employment (LBD)	ξ	0.0186	-0.0615***	0.0422***
	se	(0.0177)	(0.0212)	(0.0149)
	R ²	0.985	0.985	0.979
	N	299650	69542	230000
Local sites (LBD)	ξ	-0.0718	-1.988	0.0439
	se	(0.497)	(2.137)	(0.0400)
	R ²	0.964	0.963	0.920
	N	285430	67684	217580
Survival (LBD)	ξ	-0.0926***	0.00806	-0.100***
	se	(0.0224)	(0.0233)	(0.0226)
	R ²	0.632	0.651	0.637
	N	299790	69542	230141

Notes: Table represents estimated treatment effects with firm and quarter fixed effects. Standard errors are reported in parentheses. Stars denote statistical significance obtained from estimating clustered standard errors by 2 digit industry with stars indicating *** p<0.01, ** p<0.05, * p<0.10

Table B.2: Division cluster assignment

Cluster 1 - Manufacture of beverages - Civil engineering - Accommodation - Computer programming, consultancy and related activities - Information service activities - Real estate activities - Activities of head offices; management consultancy activities - Architectural and engineering activities; technical testing and analysis - Office administrative, office support and other business support activities - Other personal service activities	Cluster 3 - Other mining and quarrying - Manufacture of textiles - Manufacture of wood products and cork, straw and plaiting - Manufacture of paper and paper products - Manufacture of rubber and plastic products - Manufacture of other non-metallic mineral products - Manufacture of computer, electronic and optical products - Manufacture of electrical equipment - Manufacture of motor vehicles, trailers and semi-trailers - Other manufacturing - Telecommunications - Travel agency, tour operator and other reservation service activities - Libraries, archives, museums and other cultural activities
Cluster 2 - Manufacture of food products - Manufacture of chemicals and chemical products - Manufacture of basic metals - Manufacture of fabricated metal products, except machinery and equipment - Manufacture of machinery and equipment - Manufacture of other transport equipment - Repair and installation of machinery and equipment - Water collection, treatment and supply - Waste collection, treatment and disposal activities; materials recovery - Construction of buildings - Specialised construction activities - Wholesale and retail trade and repair of motor vehicles and motorcycles - Wholesale trade, except of motor vehicles and motorcycles - Retail trade, except of motor vehicles and motorcycles - Land transport and transport via pipelines - Warehousing and support activities for transportation - Postal and courier activities - Food and beverage service activities - Publishing activities - Legal and accounting activities - Rental and leasing activities - Employment activities - Security and investigation activities - Services to buildings and landscape activities - Education - Human health activities - Sports activities and amusement and recreation activities	Cluster 4 - Manufacture of basic pharmaceutical products and preparations - Manufacture of furniture - Motion picture, video, television, sound recording and music publishing - Other professional, scientific and technical activities - Creative, arts and entertainment activities Cluster 5 - Printing and reproduction of recorded media Cluster 6 - Scientific research and development - Advertising and market research

Note: Additional tables and figures are available in the online appendix at <https://firmnarratives.trfetzter.com/>. These include: energy intensity descriptives and cross-sectional regressions; additional coefficient distributions across industries; k-means excluded industries; average treatment effects under the full menu of increasingly demanding fixed effects specifications (for all firms, large firms and small firms separately); and average treatment effects at different post-treatment windows (for all firms, large firms and small firms separately). All estimates can also be explored interactively at the same web address.