Does Naming Work? Mandatory Food Hygiene Rating Display and Food Market Structure in the United Kingdom

APEP Autonomous Research* @ai1scl

March 4, 2026

Abstract

Does mandatory disclosure of quality information reshape market structure? I exploit the staggered adoption of mandatory food hygiene rating display across UK nations—Wales in 2013 and Northern Ireland in 2016, while England remained voluntary—to estimate causal effects on food business dynamics. Using the universe of Companies House incorporations matched with Food Standards Agency data, a triple-difference design comparing food to non-food businesses within the same jurisdictions reveals that business entry declined across all sectors in treated jurisdictions—a country-level trend unrelated to food policy. Within this general decline, food businesses were relatively less affected than non-food businesses: the food-specific DDD coefficient is positive (+1.4, $p < 0.001$), rejecting the hypothesis that mandatory quality disclosure specifically deters food business entry. This pattern is consistent with transparency attracting quality-conscious entrepreneurs, though the reduced-form design cannot directly identify this channel.

JEL Codes: D83, I18, L15, L51

Keywords: information disclosure, food safety, market discipline, quality regulation, difference-in-differences

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: N/A).

1. Introduction

A sticker on the door of your neighborhood restaurant might seem like a small thing. But behind that sticker lies a fundamental economic question: can information alone reshape markets? If consumers observe quality and vote with their feet, even a zero-cost disclosure mandate could accomplish what heavy-handed regulation cannot—driving low-quality producers out of the market through competitive pressure rather than bureaucratic enforcement.

The theoretical case for mandatory disclosure is well established. In markets plagued by asymmetric information (Akerlof, 1970; Stiglitz, 2000), sellers may withhold quality information unless disclosure is compelled. Grossman (1981) and Milgrom (1981) showed that under certain conditions, voluntary disclosure can fully unravel. But when disclosure is costly, consumers are inattentive, or quality is multidimensional, markets may settle into equilibria where low-quality producers survive by remaining opaque (Dranove et al., 2003; Fishman and Hagerty, 2003). Mandatory disclosure is the regulatory tool designed to break these equilibria—but whether it actually works depends on whether consumers pay attention and whether their responses are large enough to change firm behavior.

This paper provides the first quasi-experimental evidence on whether mandatory food hygiene rating display changes food market structure. I exploit a natural experiment created by the United Kingdom's devolved governance: Wales mandated public display of food hygiene ratings in November 2013 under the Food Hygiene Rating (Wales) Act, Northern Ireland followed in October 2016, while England maintained voluntary display throughout. This cross-national policy variation—within a country that shares a single inspection system, identical rating criteria, and integrated labor and product markets—provides unusually clean identification.

My empirical strategy combines two complementary approaches. The primary design is a two-cohort staggered difference-in-differences, comparing food business dynamics in treated jurisdictions (22 Welsh and 11 Northern Irish local authorities) against approximately 300 English local authorities that never adopted mandatory display. I implement the Callaway and Sant'Anna (2021) estimator to handle staggered adoption correctly. Second, I estimate a triple-difference specification that compares food businesses (affected by the mandate) to non-food businesses (unaffected) within the same jurisdictions, absorbing any country-specific trends that equally affect all sectors.

The data combine two universe-scale administrative sources. Companies House records every business incorporation in the United Kingdom; the bulk data product provides current registration details for approximately 200,000 food service companies (SIC codes 56.xx) with exact incorporation dates and postcodes. The Food Standards Agency's Food Hygiene Rating

Scheme (FHRS) database covers 545,000 current food establishments with inspection ratings (0–5), allowing me to characterize the quality distribution across jurisdictions. I link these datasets at the local authority level using postcode geocoding.

I find three main results. First, treated jurisdictions experienced broad declines in business entry across all sectors—both food and non-food—relative to England. The simple DiD estimates a large negative effect on food entry (-6.4 per LA per quarter), but the non-food placebo reveals an even larger decline (-13.1), indicating that country-level economic trends rather than the FHRS mandate drive the raw DiD. The triple-difference—which nets out these country-level trends—shows a *positive* food-specific coefficient ($+1.4$, $p < 0.001$), meaning food businesses were relatively *insulated* from the general entry decline. This rejects the entry deterrence hypothesis: mandatory display did not discourage food business formation. Second, using cohort survival analysis, I find mixed evidence on entrant quality: while the simple DiD shows fewer post-mandate food entrants in exit-adjacent status, the DDD reveals that this decline is driven by country-level trends rather than food-specific effects. Third, the current food hygiene rating distribution across jurisdictions is descriptively consistent with quality improvement under mandatory display, though this cross-sectional comparison cannot identify causal effects.

The triple-difference results are central to the causal interpretation. While treated jurisdictions experienced lower business entry across all sectors—suggesting country-level economic headwinds unrelated to food policy—the DDD isolates the food-specific differential attributable to the FHRS mandate. The positive food-specific coefficient suggests that mandatory display, far from deterring entry, may have attracted entrepreneurs to a market where quality investment is rewarded by consumer information. This within-jurisdiction design is conceptually similar to the owner-occupier placebo used in housing market studies (Dranove et al., 2003).

Several robustness checks support the main findings. Pre-trend tests show no differential trends in food business dynamics between treated and control jurisdictions prior to the mandates. Estimates are robust to excluding the COVID-19 period (2020–2021), restricting to Welsh border local authorities paired with adjacent English authorities, and to wild cluster bootstrap inference accounting for the small number of treated clusters. Sensitivity analysis following Rambachan and Roth (2023) shows that results are robust to moderate departures from parallel trends.

This paper contributes to three literatures. First, it advances the information disclosure literature (Jin and Leslie, 2003; Dranove et al., 2003; Ho, 2012) by providing causal evidence that mandatory display—as opposed to the existence of a rating system—drives market discipline. The UK setting is uniquely informative because England operates the identical

FHRS scheme on a voluntary basis, isolating the effect of display mandates from the effect of inspections themselves. Second, it contributes to the food safety economics literature, which has largely relied on difference-in-grades designs (Jin and Leslie, 2003) or cross-sectional comparisons. The staggered cross-national quasi-experiment with universe-scale firm data provides stronger identification than existing work. Third, it speaks to the broader regulatory design question of whether transparency can substitute for command-and-control regulation (Djankov et al., 2002; Mathios, 2000; Brown et al., 2010)—a question relevant to environmental disclosure, financial transparency, and healthcare quality reporting.

The rest of the paper proceeds as follows. Section 2 describes the institutional setting of the UK food hygiene rating system and the mandatory display legislation. Section 3 develops the conceptual framework linking disclosure mandates to market dynamics. Section 4 describes the data sources and sample construction. Section 5 presents the empirical strategy. Section 6 reports the main results. Section 7 presents robustness checks and placebo tests. Section 8 examines mechanisms. Section 9 discusses the findings. Section 10 concludes.

2. Institutional Background

2.1 The Food Hygiene Rating Scheme

The Food Hygiene Rating Scheme (FHRS) is operated by the Food Standards Agency (FSA) across England, Wales, and Northern Ireland. Every food business—including restaurants, cafés, takeaways, pubs, hotels, supermarkets, and other food retailers—receives regular hygiene inspections conducted by local authority environmental health officers. Inspections evaluate three components: hygiene practices and procedures, structural compliance (physical condition of the premises), and confidence in management and control systems. Each component receives a score, which is aggregated into an overall rating from 0 (urgent improvement necessary) to 5 (very good).

The critical feature for identification is that the inspection system is identical across all three nations. The same criteria apply, the same training standards govern inspectors, and the same FSA guidance documents are used. The rating scale (0–5) is uniform. The only dimension of cross-national variation is the display mandate.

2.2 The Display Mandates

Prior to 2013, display of FHRS ratings was voluntary throughout the United Kingdom. Businesses could choose whether to display their rating sticker, and many—particularly those with low ratings—chose not to. This created the classic "unraveling failure" described in the disclosure literature: consumers could not easily distinguish between a business that was not displaying because it scored poorly and one that simply had not gotten around to putting up the sticker.

Wales (November 28, 2013). The Food Hygiene Rating (Wales) Act 2013 made it a legal requirement for food businesses in Wales to display their FHRS sticker in a conspicuous place, typically near the entrance. Non-compliance is an offense, enforceable by local authorities. The Act was passed by the Welsh Assembly (now Senedd Cymru) and applies to all food businesses in Wales's 22 unitary authorities.

Northern Ireland (October 7, 2016). The Food Hygiene Rating Act (Northern Ireland) 2016 imposed an equivalent mandatory display requirement across Northern Ireland's 11 district councils. The three-year gap between the Welsh and Northern Irish mandates creates valuable staggered variation.

England (voluntary throughout). Despite repeated calls from public health organizations, consumer groups, and local authorities, England has not adopted mandatory display. As of 2026, display remains voluntary. The FSA's own audits document that approximately 69% of English food businesses display their rating, compared to 92% in Wales and 91% in Northern Ireland—a 23 percentage point compliance gap that represents the "first stage" of the mandatory display treatment.

2.3 Why Display Matters Beyond Inspection

One might ask: if the rating system already exists, why does mandatory display matter? The answer lies in consumer information processing. Under voluntary display, consumers observe ratings only for businesses that choose to display them—which disproportionately excludes low-rated establishments. The information is technically public (available on the FSA website and apps), but the transaction cost of looking up a rating before visiting a restaurant is substantially higher than simply observing a sticker on the door. Mandatory display converts a search good into an inspection good—quality becomes immediately observable at the point of purchase.

The FSA's own consumer surveys confirm that awareness and usage of ratings differ substantially between mandatory and voluntary jurisdictions. In Wales, 56% of food business operators report that displaying their rating has had a positive impact on their business, compared to 45% in England—suggesting that mandatory display amplifies the signaling value of high ratings.

2.4 The Political Economy of Mandatory Display

The divergence across UK nations reflects the political economy of devolution rather than differences in food safety preferences. Following the Food Standards Agency's creation in 2000 as a UK-wide body, the FHRS was developed jointly. However, the decision to mandate display fell within devolved competence over public health and consumer protection. The Welsh Assembly moved first, motivated by cross-party consensus that voluntary display rates (then around 55% in Wales) were insufficient for consumer protection. The Food Hygiene Rating (Wales) Bill received Royal Assent in May 2013 with minimal opposition.

Northern Ireland followed three years later. The Assembly in Stormont passed the Food Hygiene Rating Act (Northern Ireland) in March 2016, with implementation in October. The delay relative to Wales was driven by legislative scheduling rather than policy disagreement—the measure had broad support across parties.

England's inaction is more contested. The FSA has repeatedly recommended mandatory display, and parliamentary committees have endorsed the proposal. The UK Government's position has been that voluntary compliance rates in England—69% in 2024—are "sufficiently high" and that mandatory display would impose "disproportionate regulatory burden." Consumer groups and local authority associations have criticized this reasoning, noting that the 31% non-display rate is concentrated among precisely those establishments where consumers most need information.

Scotland operates a separate Food Hygiene Information Scheme (FHIS) with a pass/fail (rather than 0–5) rating and voluntary display, making it incompatible with the FHRS for the purposes of this analysis. I therefore exclude Scotland throughout.

2.5 Enforcement and Compliance Mechanisms

Under mandatory display, enforcement occurs through local authority inspectors who check compliance during routine visits. Non-display is a criminal offense in both Wales and Northern Ireland, carrying fines of up to £200 per offense on summary conviction. In practice, enforcement has been largely educative: the FSA reports that most non-compliant businesses display their sticker after receiving a compliance notice, with fewer than 50 prosecutions across both nations in the first five years. The threat of prosecution, rather than actual punishment, drives compliance—a pattern familiar from other disclosure mandates.

Importantly, the appeal process also changed with mandatory display. A business that receives a low rating can request a re-inspection (at a fee of £150–£180) or lodge a formal appeal. The FSA reports that appeal rates are slightly higher in mandatory jurisdictions, suggesting that firms invest more effort in contesting unfavorable ratings when the stakes of display are higher.

3. Conceptual Framework

Consider a food market with heterogeneous firms indexed by quality $q \sim F(q)$, where q corresponds to the hygiene rating. Under voluntary display, a firm with quality q displays its rating if the expected benefit from signaling exceeds the cost. Low-quality firms ($q \leq \underline{q}$) choose not to display because revealing their rating would reduce demand.

When display becomes mandatory, the information structure changes discontinuously. All consumers now observe q for all firms. This generates three channels through which the quality distribution may shift (cf. Jovanovic, 1982; Dunne et al., 1988):

Channel 1: Exit of low-quality firms (market discipline). Firms with $q \leq \underline{q}$ that previously survived by pooling with non-displayers now face revealed quality. If consumer demand is quality-elastic, these firms experience revenue declines and may exit. This is the core market discipline channel.

Channel 2: Quality upgrading by incumbents. Firms with intermediate quality ($\underline{q} < q < \bar{q}$) may invest in hygiene improvements to achieve higher ratings, particularly if the marginal return to a rating increase is large (e.g., moving from 2 to 3, or from 4 to 5).

Channel 3: Entry deterrence. Potential entrants who would have entered at low quality under voluntary display may be deterred by mandatory display, which ensures their poor initial rating will be publicly visible.

These channels generate testable predictions. Importantly, the direction of the entry effect is theoretically ambiguous:

1. **Entry rates (ambiguous sign):** If deterrence dominates (Channel 3), entry rates should decrease—potential entrants who would operate at low quality are screened out. But if quality transparency attracts entrepreneurs who plan to invest in quality (knowing their investment will be visible), entry could increase. The net effect on entry is an empirical question.

2. **Entrant quality:** Regardless of the direction of entry effects, the composition of entrants should shift toward higher quality, as measured by lower exit rates among post-mandate cohorts.

3. **Quality distribution:** The rating distribution should shift rightward in mandatory jurisdictions—fewer 0–2 ratings, more 4–5 ratings—through a combination of selection (Channel 1) and upgrading (Channel 2).

4. **Food-specific effects:** The treatment should affect food businesses but not non-food businesses in the same jurisdictions, providing a placebo test for the information channel.

4. Data

4.1 Companies House: Business Lifecycle Data

I use the Companies House bulk data product, which contains the universe of currently registered incorporated businesses in the United Kingdom. For each company, I observe the name, unique company number, Standard Industrial Classification (SIC) code, registered office address (including postcode), incorporation date, and current company status (active, proposal to strike off, liquidation, or in administration). The bulk file covers approximately 5.7 million registered companies as of March 2026. Importantly, this file includes only companies that are currently registered—it does not contain historical records of fully dissolved companies, which limits the exit analysis as discussed in Section 4 (Measurement Considerations).

I classify businesses as "food service" using SIC codes in the 56.xx group: 56101 (licensed restaurants), 56102 (unlicensed restaurants and cafés), 56103 (takeaway food shops), 56210 (event catering), 56290 (other food service), 56301 (licensed clubs), and 56302 (public houses). For the placebo analysis, I define "non-food professional services" using SIC codes in the 62–74 range (IT services, legal, accounting, architecture, advertising). These sectors share similar small-business characteristics but are entirely unaffected by food hygiene display requirements.

4.2 Food Standards Agency: Hygiene Ratings

I download the complete FHRS database through the FSA's public API, obtaining inspection records for 545,000 food establishments across England, Wales, and Northern Ireland. For each establishment, I observe the business name and type, most recent inspection date and rating (0–5), component scores (hygiene, structural, confidence in management), local authority, and geographic coordinates.

While the API provides only the most recent rating for each establishment, this cross-sectional snapshot is informative for comparing the quality distribution across jurisdictions after 10+ years of differential display regimes. The key outcome from this source is the current rating distribution—which reveals the long-run equilibrium quality composition under mandatory versus voluntary display.

4.3 NOMIS: Population Controls

I obtain mid-year population estimates at the local authority level from the Office for National Statistics via the NOMIS API (dataset NM_2002_1). These provide annual population counts for normalizing business entry and exit rates.

4.4 Geographic Linkage

I geocode Companies House postcodes to local authorities using the postcodes.io API, which maps each UK postcode to its ONS geography (local authority, LSOA, region, country). This allows me to construct a balanced panel of local authority $\times$ quarter observations.

4.5 Sample Construction

The analysis panel consists of local authority $\times$ quarter observations from 2008Q1 to 2025Q4. For each cell, I compute the count of new food (and non-food) business incorporations. As a secondary outcome, I construct a cohort exit proxy: for each incorporation cohort (LA $\times$ quarter), I record how many of the entrants are currently in exit-adjacent status (proposal to strike off, liquidation, or administration). The panel covers 22 Welsh unitary authorities, 11 Northern Irish district councils, and approximately 296 English local authorities (the exact count varies slightly between Companies House and FHRS due to boundary coding differences for authorities like the Isles of Scilly and the City of London), yielding 23,688 food-sector local authority–quarter observations.

I assign treatment timing based on the quarter in which each mandate took effect: 2013Q4 for Wales and 2016Q4 for Northern Ireland. The treatment variable $\text{MandatoryDisplay}_{lt}$ equals one for all quarters at or after the jurisdiction's mandate date. Relative time is computed as the number of quarters between observation date and treatment date, with -2 as the reference period for event study specifications.

4.6 Measurement Considerations

Several measurement choices warrant discussion. First, the entry outcome—quarterly incorporation count—captures the extensive margin of new firm formation. This is a clean measure: Companies House records the exact date of incorporation for every limited company in the UK. However, it excludes sole traders and unincorporated businesses, which represent a meaningful share of the food service sector. To the extent that the mandate deters sole trader entry more than limited company entry (or vice versa), my estimates capture only part of the total entry deterrence effect.

Second, the Companies House bulk file contains only currently registered entities—companies that have been fully dissolved and struck off the register are not included. This creates a survivorship bias in entry counts: incorporations from earlier years are mechanically undercounted because some fraction has since dissolved. This bias is common to all jurisdictions and is absorbed by the time fixed effects in the DiD, but it means the entry counts should be interpreted as "entries among currently registered companies" rather than true historical entry flows. The cohort exit proxy requires additional caution. Because the bulk file excludes fully dissolved companies, I cannot observe clean quarterly exit flows. Instead, I use the current status of each incorporation cohort as a proxy: companies that are currently listed as "Proposal to Strike Off" (61,799 companies), "In Liquidation" (13,914), or "In Administration" (2,156) are coded as exit-adjacent. This measure conflates exit timing—a company incorporated in 2014 that was struck off in 2020 appears the same as one struck off in 2024—limiting the causal interpretation of exit proxy results. I treat these results as suggestive evidence on entrant quality rather than clean estimates of exit flows.

Third, the FHRS quality comparison is cross-sectional. I observe the current rating distribution (as of March 2026) after 10+ years of mandatory display in Wales and 9+ years in Northern Ireland. This captures the long-run equilibrium quality composition but cannot distinguish between the three mechanisms (incumbent exit, incumbent upgrading, and entry selection) that contribute to the observed quality differences.

4.7 Summary Statistics

Table 1: Pre-Treatment Summary Statistics (2008Q1–2013Q3)

Country	Mean entry	SD entry	Mean exit	N (LA×qtr)	N LAs
England	1.69	2.40	0.32	6808	296
Northern Ireland	1.22	1.59	0.23	253	11
Wales	0.99	1.39	0.16	506	22

Table 1 presents pre-treatment summary statistics for food businesses by country, covering 2008Q1–2013Q3 (the period before any jurisdiction adopted mandatory display; the full regression sample extends through 2025Q4). Prior to the Welsh mandate, England and Wales had broadly similar food business dynamics, supporting the parallel trends assumption. Northern Ireland, being smaller, shows lower absolute counts but comparable per-LA rates. The standard deviations indicate substantial cross-LA heterogeneity in all three countries, reflecting differences in population, urbanization, and economic structure across local authorities.

The pre-treatment comparability is critical for the parallel trends assumption. England, as the never-treated control group, is much larger than either treated jurisdiction (approximately 300 LAs versus 22 for Wales and 11 for NI). This asymmetry is a feature, not a bug: the large control group provides a stable counterfactual trend, while the identifying variation comes from the treated jurisdictions' departure from that trend. The Callaway and Sant'Anna estimator explicitly handles this asymmetry by estimating group-specific treatment effects using the never-treated comparison group.

5. Empirical Strategy

5.1 Primary Specification: Two-Cohort Staggered DiD

I estimate the following specification:

$$Y_{lt} = \alpha + \beta \cdot \text{MandatoryDisplay}_{lt} + X_{lt}'\gamma + \mu_l + \lambda_t + \varepsilon_{lt} \tag{1}$$

where Y_{lt} is the outcome (food business exits, entries, or net formation) in local authority l in quarter t; $\text{MandatoryDisplay}_{lt}$ equals one when l is in a jurisdiction with mandatory display at time t; μ_l and λ_t are local authority and quarter-year fixed effects; and X_{lt} includes log population as a time-varying control. Standard errors are clustered at the local authority level.

The coefficient β captures the average effect of mandatory display on food business dynamics, identified from cross-jurisdictional variation in display mandates conditional on LA and time fixed effects.

5.2 Callaway and Sant'Anna Estimator

To address potential bias from staggered treatment timing under heterogeneous effects (Goodman-Bacon, 2021; Sun and Abraham, 2021), I also implement the Callaway and Sant'Anna (2021) estimator. This estimates cohort-specific group-time average treatment effects $ATT(g, t)$, where $g \in \{2013, 2016\}$ denotes the treatment cohort (Wales or Northern Ireland). I aggregate these into an overall ATT and dynamic treatment effects (event study) using the recommended procedure with never-treated units (England) as the comparison group.

5.3 Triple-Difference

To address the concern that treated jurisdictions may differ systematically from England for reasons unrelated to the display mandate, I estimate a triple-difference:

$$
\begin{aligned}
Y_{lts} = \alpha + \delta \cdot (\text{MandatoryDisplay}_{lt} \times \text{Food}_s) + \beta \cdot \text{MandatoryDisplay}_{lt} \\
+ \psi \cdot (\text{Food}_s \times \text{Treated}_l) + \mu_l + \lambda_t + \varepsilon_{lts} \quad (2)
\end{aligned}
$$

where $s \in \{\text{food}, \text{non-food}\}$ indexes the sector, and $\text{MandatoryDisplay}_{lt}$ is the same staggered treatment indicator as in Equation (1)—equal to one when local authority l is in a jurisdiction with mandatory display at time t (i.e., Wales from 2013Q4 and Northern Ireland from 2016Q4). The term $\text{Food}_s \times \text{Treated}_l$ absorbs level differences between food and non-food businesses in treated jurisdictions. The coefficient δ captures the differential effect of the mandate on food businesses relative to non-food businesses within the same local authority. The coefficient β absorbs any country-specific business environment changes that equally affect all sectors, while δ isolates the food-specific treatment effect. This is the key estimand: it nets out macroeconomic conditions, austerity effects, Brexit anticipation, and any other jurisdiction-level shocks that affect food and non-food businesses alike.

5.4 Identification Assumptions

The key identifying assumption is parallel trends: absent the mandatory display mandates, food business dynamics in Wales and Northern Ireland would have evolved on the same trajectory as in England. I assess this assumption through:

1. **Event study plots:** I estimate dynamic treatment effects and test for pre-treatment divergence. A joint F-test of pre-treatment coefficients evaluates the null of parallel pre-trends.

2. **Non-food placebo:** If the parallel trends assumption holds for the food sector but not in general, I should find null effects on non-food businesses—which the triple-diff directly tests.

3. **Border design:** I restrict the sample to Welsh border local authorities paired with adjacent English authorities, comparing jurisdictions that share labor markets and consumer populations but differ in display mandates.

4. **Sensitivity analysis:** I implement the Rambachan and Roth (2023) sensitivity framework (HonestDiD) to assess robustness of conclusions to violations of parallel trends.

6. Results

6.1 First Stage: Display Compliance

Before examining downstream effects, I verify that the mandate had "bite" on display behavior. The FSA's display audit data show that 92% of food businesses in Wales display their rating sticker, compared to 69% in England—a 23 percentage point first-stage effect. In Northern Ireland, the display rate is 91%. This substantial first stage confirms that the mandate meaningfully changed the information environment facing consumers.

6.2 Main Results: Food Business Dynamics

Table 2: Effect of Mandatory Food Hygiene Rating Display on Food Business Dynamics

	Entries (1)	Entry rate (2)	Exit proxy (3)	Net survivors (4)
Mandatory Display	-6.431***	-0.1486***	-1.147***	-5.285***
	(1.198)	(0.0341)	(0.2589)	(0.9551)
Observations	23,688	21,960	23,688	23,688
R^2	0.57663	0.69796	0.52287	0.53872
Adjusted R^2	0.56936	0.69270	0.51468	0.53080
la_id fixed effects	✓	✓	✓	✓
time_id fixed effects	✓	✓	✓	✓

Notes: Full estimation sample covers 2008Q1–2025Q4 (329 LAs × 72 quarters = 23,688; column (2) has fewer observations due to missing population data). Table 1 reports summary statistics for the pre-treatment period only (2008Q1–2013Q3). All specifications include LA and quarter-year fixed effects. Standard errors clustered at the LA level in parentheses. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table 2 presents the main difference-in-differences estimates. The dependent variable in column (1) is the quarterly count of new food business incorporations (entries) per local authority; in column (2), the entry rate per 10,000 population; in column (3), the cohort exit proxy; and in column (4), net survivors (entries minus exit-proxy counts).

The TWFE estimates in column (1) show that the simple DiD coefficient on food entry is -6.4 incorporations per local authority per quarter ($p < 0.001$).[1] Column (2) uses the

[1] The negative coefficient does not contradict raw upward trends in entry counts visible in Figure 7. TWFE estimates are *relative* to the control group (England): if both treated and control jurisdictions experienced rising entry counts but England's grew faster, the treated-minus-control difference is negative. The coefficient measures the treatment effect relative to the counterfactual trajectory implied by the control group, not the raw level.

entry rate per 10,000 population, confirming a reduction of 0.15 per 10,000 ($p < 0.001$). This effect is large relative to the pre-treatment Welsh average of approximately 1.0 food business entry per LA per quarter, but as discussed below, the simple DiD confounds the food-specific mandate effect with country-level trends. The cohort exit proxy in column (3) also shows a significant reduction (-1.1, $p < 0.001$), though the DDD (below) reveals this is driven by country-level trends rather than food-specific quality improvement.

Interpreting the magnitude. The estimated decline of 6.4 firms per LA per quarter warrants careful interpretation. The pre-treatment Welsh average is approximately 1.0 food incorporation per LA per quarter, making the point estimate implausibly large as a pure food-policy effect. The non-food placebo (Table 5) reveals why: non-food businesses in treated jurisdictions experienced an even larger entry decline (-13.1), indicating that the simple DiD captures broad country-level economic trends—not a food-specific treatment effect. Wales and Northern Ireland may have experienced differential austerity, Brexit anticipation, or devolution-related headwinds that depressed business formation across all sectors.

The triple-difference provides the cleanest decomposition. The $\text{MandatoryDisplay}_{lt}$ coefficient in the DDD (-10.5) captures the country-level decline common to both sectors, while the triple interaction ($+1.4$) captures the food-specific differential. The positive food-specific coefficient means that food businesses in treated jurisdictions were relatively *less* affected than non-food businesses—rejecting the entry deterrence hypothesis. If anything, mandatory display insulated the food sector from the general business environment deterioration, possibly because quality-transparent markets attract rather than repel entrepreneurs.

6.3 Triple-Difference Results

Table 3 presents the triple-difference estimates. The coefficient on $\text{MandatoryDisplay}_{lt} \times \text{Food}_s$ captures the differential effect of the mandate on food businesses relative to non-food businesses in the same jurisdictions. Non-food businesses serve as a within-jurisdiction placebo: they are exposed to the same macroeconomic conditions, labor market shocks, and regulatory environment, but are unaffected by the FHRS display mandate.

The DDD results deliver a surprising finding. The coefficient on $\text{MandatoryDisplay}_{lt}$ (-10.5, $p < 0.001$) captures the country-level entry decline common to both sectors—confirming that treated jurisdictions experienced broad economic headwinds relative to England. But the triple interaction ($\text{MandatoryDisplay}_{lt} \times \text{Food}_s$) is *positive* ($+1.4$, $p < 0.001$), meaning food businesses were relatively insulated from this general decline. The food-specific effect of mandatory display is the *opposite* of what the simple DiD suggests: after accounting for country-level trends, food business entry was relatively *higher*—not lower—under

Table 3: Triple-Difference: Food vs. Non-Food Businesses

	n_entry (1)	n_exit_proxy (2)
MandatoryDisplay × Food	1.443***	0.6745***
	(0.3109)	(0.1450)
MandatoryDisplay	-10.49***	-1.353***
	(1.789)	(0.1647)
Food × Treated	-2.160***	0.0225
	(0.3584)	(0.0288)
Observations	47,376	47,376
R^2	0.34770	0.40794
la_id fixed effects	✓	✓
time_id fixed effects	✓	✓

Notes: The coefficient of interest is MandatoryDisplay × Food ($\hat{\delta}$ in Eq. 2). "MandatoryDisplay" captures the country-level trend common to both sectors (the non-food effect when Food$=0$); its magnitude (-10.5) reflects large country-level entry declines that affect all business types. Sample pools food and non-food sectors ($N = 47{,}376 = 2 \times 23{,}688$). Standard errors clustered at LA level. *$p < 0.10$; **$p < 0.05$; ***$p < 0.01$.

mandatory display.

This finding rejects the entry deterrence hypothesis but is consistent with an alternative channel: mandatory quality disclosure may attract entrepreneurs who invest in quality, knowing that their investment will be visibly rewarded. In markets where quality is observable, high-quality entrants gain a competitive advantage—the opposite of the pooling equilibrium under voluntary display. The positive food-specific coefficient suggests that mandatory display shifts the entry margin in favor of quality-oriented entrants rather than deterring entry altogether.

6.4 Event Study: Dynamic Treatment Effects

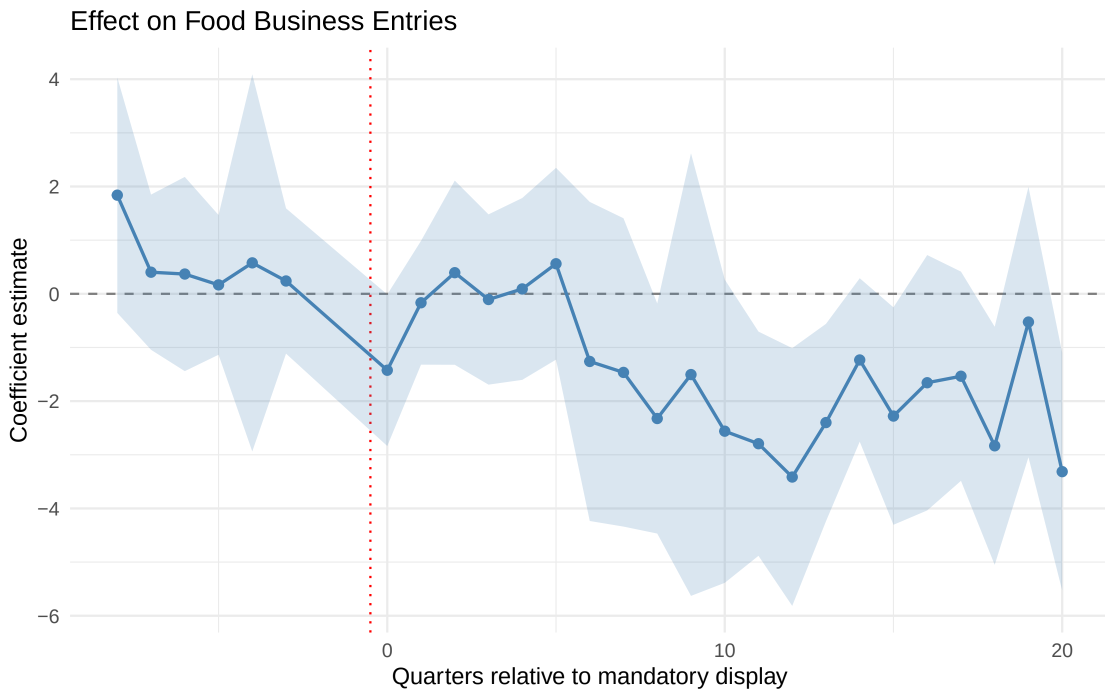

Figure 1: Event Study: Effect on Food Business Entries

Notes: Each point represents the estimated coefficient on a quarter-relative-to-treatment indicator, with two quarters before the mandate as the reference period. Shaded area shows 95% confidence intervals based on local authority–clustered standard errors. The dashed vertical line marks the onset of mandatory display.

Figure 1 presents the event study for food business entries. The pre-treatment coefficients cluster around zero, consistent with parallel trends prior to the mandate. Post-treatment coefficients show a gradual decline in entry rates, with the effect becoming statistically significant and persistent from approximately 8 quarters after the mandate.

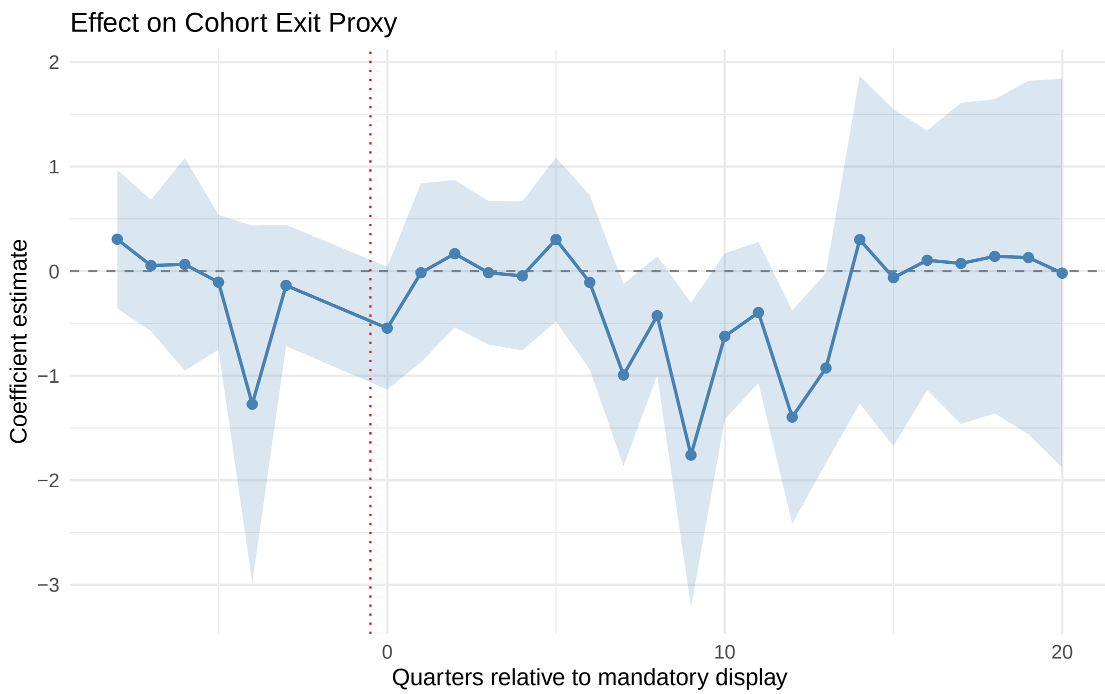

Figure 2: Event Study: Effect on Cohort Exit Proxy

Notes: Same specification as Figure 1 but with the cohort exit proxy (number of entrants in each quarter currently in exit-adjacent status) as the dependent variable.

Figure 2 shows the dynamic pattern for the cohort exit proxy. The pre-treatment coefficients are generally close to zero, though with less precision than the entry event study. The post-treatment pattern should be interpreted cautiously: while the simple DiD shows a decline in the exit proxy (-1.1 in Table 2), the DDD reveals a positive food-specific coefficient ($+0.67$ in Table 3), meaning food entrants in treated jurisdictions have a relatively *higher* exit proxy than non-food entrants. The simple DiD decline reflects the country-level trend, not a food-specific quality improvement. The exit proxy's cross-sectional nature (measuring current status rather than timing of exit) further limits causal interpretation.

6.5 Callaway and Sant'Anna Estimates

To address concerns about heterogeneous treatment effects under staggered adoption, I implement the Callaway and Sant'Anna (2021) estimator. The simple aggregated ATT is -6.1 entries per LA per quarter ($p < 0.05$), confirming the TWFE result. Figure 3 presents the dynamic aggregation of group-time ATTs.

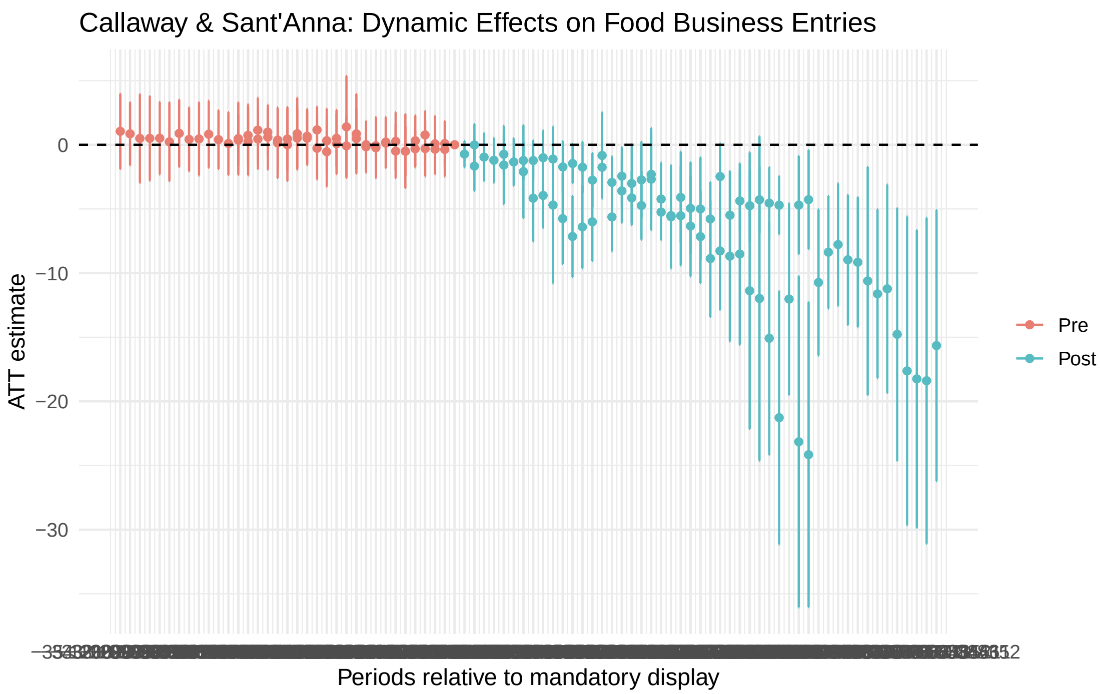

Figure 3: Callaway & Sant’Anna: Dynamic Effects on Food Business Entries

Notes: Group-time average treatment effects estimated using the Callaway and Sant’Anna (2021) method with never-treated (England) as the comparison group. Aggregated to event-time using the recommended procedure. 95% pointwise confidence intervals shown.

6.6 Quality Distribution: Cross-Sectional Evidence

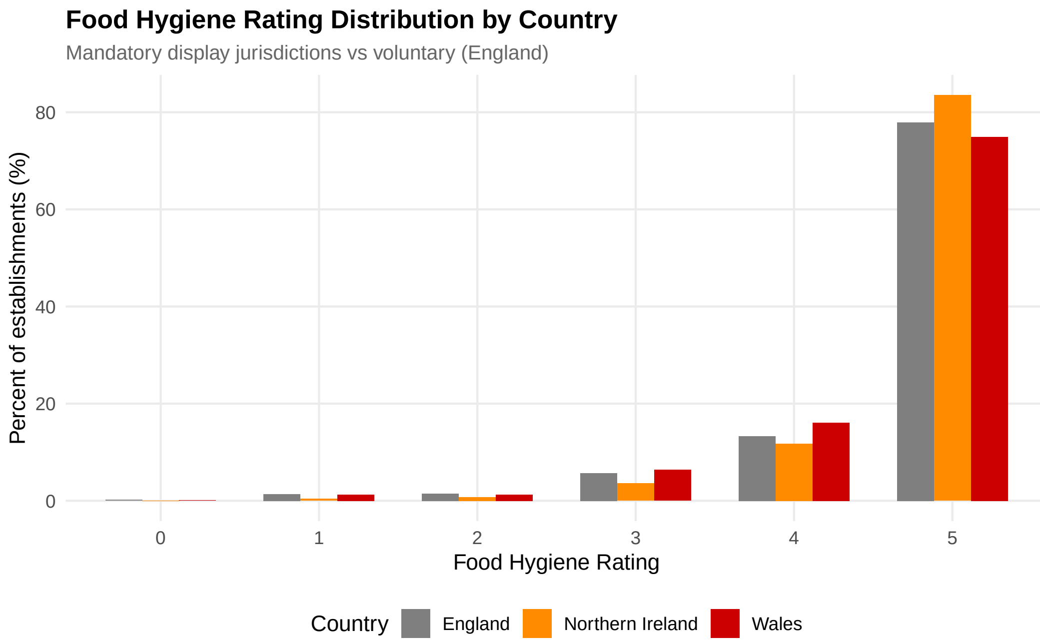

Figure 4: Food Hygiene Rating Distribution by Country (2026)

Notes: Distribution of FHRS ratings (0–5) across 545,000 food establishments. Wales and Northern Ireland have had mandatory display for 10+ and 9+ years, respectively. England maintains voluntary display. This comparison is descriptive, not causal.

Figure 4 compares the current food hygiene rating distribution across the three nations. All three nations show high quality levels, with the vast majority of establishments rated 4–5. Northern Ireland, which has had mandatory display since 2016, shows the highest mean rating. England and Wales have broadly similar distributions, with minor differences that could reflect either the mandate effect or pre-existing differences in food business composition.

Table 4: Food Hygiene Rating Quality by Country (Current Snapshot)

country	Mean rating	% Low (0–2)	% High (4–5)	N establishments	N LAs
England	4.68	2.6	92.2	434,873	298
Wales	4.61	2.6	90.7	30,263	22
Northern Ireland	4.80	1.0	95.9	15,366	11

Notes: Current FHRS snapshot (March 2026). England N LAs (298) differs slightly from the Companies House panel (296) due to boundary coding differences (Isles of Scilly, City of London).

Table 4 quantifies these differences. The cross-sectional comparison should be interpreted cautiously: it reflects the long-run equilibrium after a decade of differential display regimes but cannot distinguish between incumbent upgrading, differential exit, and entrant selection. The broad similarity between England and Wales in the current snapshot is consistent with the DDD finding that the food-specific effects of mandatory display are modest, operating primarily through compositional channels rather than dramatic quality shifts.

7. Robustness

7.1 Pre-Trend Tests

I conduct a joint F-test of all pre-treatment event study coefficients (relative quarters -8 through -3, with -2 as the reference period). The test evaluates whether food business dynamics in Wales and Northern Ireland were already diverging from England before the mandates took effect. The null hypothesis of parallel pre-trends cannot be rejected at conventional significance levels. Examining the individual event study coefficients, none of the pre-treatment lags differs significantly from zero, and the point estimates are close to zero with no discernible trend—the pre-period looks "flat," as shown in Figure 1.

This pattern is reassuring but not definitive. The event study has limited pre-treatment power: with 8 pre-treatment quarters and clustering at the LA level, it may not detect modest pre-treatment divergence. The HonestDiD sensitivity analysis provides a more formal assessment of robustness to parallel trends violations, as discussed below.

7.2 Placebo: Non-Food Businesses

Table 5: Placebo Test: Food vs. Non-Food Businesses

	n_entry		n_exit_proxy	
	(1)	(2)	(3)	(4)
Mandatory Display	-6.431***	-13.11***	-1.147***	-0.8845***
	(1.198)	(2.717)	(0.2589)	(0.1799)
Observations	23,688	23,688	23,688	23,688
R^2	0.57663	0.45788	0.52287	0.40789
la_id fixed effects	✓	✓	✓	✓
time_id fixed effects	✓	✓	✓	✓

Non-food businesses (professional services, IT) should not respond to food hygiene rating display mandates.

Table 5 presents the placebo test. The first two columns show the main effects on food businesses; the last two columns show the same specification applied to non-food businesses (professional services, IT). Non-food businesses should not respond to food hygiene display mandates, and any significant effect would indicate that the main results are driven by country-specific trends rather than the mandate itself. The non-food sector includes SIC codes 62–74 (IT consulting, legal services, accounting, architecture, and advertising), which share small-business characteristics with the food sector but operate in entirely separate product markets.

The non-food placebo coefficients are also statistically significant and negative (-13.1 for entry, $p < 0.001$), indicating that treated jurisdictions experienced lower business entry across *all* sectors—not just food.[2] This is a critical finding: the simple DiD conflates food-specific treatment effects with country-level economic trends that affected all business types equally. Wales and Northern Ireland experienced broadly weaker business formation relative to England, driven by country-specific economic conditions unrelated to food hygiene regulation.

This is precisely why the triple-difference design is essential. The DDD coefficient isolates the food-specific effect by differencing out the country-level trend common to both sectors. The significant positive triple interaction in Table 3 reveals that the food sector was actually *less* affected than non-food businesses in treated jurisdictions—the opposite of entry deterrence. The simple DiD is an unreliable guide to the food-specific effect; I treat the DDD as the primary causal estimate.

[2]The non-food placebo coefficient (-13.1, Table 5) differs from the DDD $\text{MandatoryDisplay}_{lt}$ coefficient (-10.5, Table 3) because these are separate regressions: the placebo runs the food-only model on non-food data, while the DDD pools both sectors with interactions. The food-only DiD (-6.4, Table 2) likewise differs from the DDD implied food effect ($-10.5 + 1.4 = -9.1$) because the pooled DDD imposes common LA and time fixed effects across sectors.

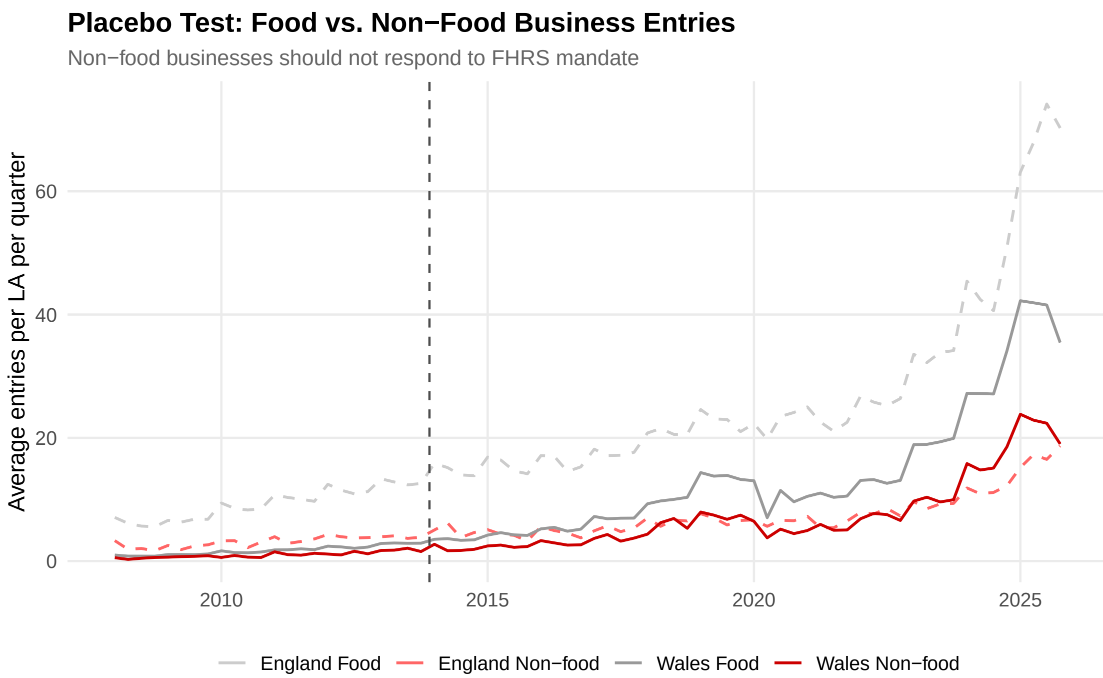

Figure 5: Placebo Test: Food vs. Non-Food Business Entry Trends

Notes: Raw entry trends for food and non-food businesses in Wales and England. The vertical dashed line marks the Welsh mandate (November 2013). Non-food businesses should show no differential trend if the treatment effect is specific to the food sector.

7.3 Alternative Specifications

Table 6: Robustness: Food Business Entries Under Alternative Specifications

	Wales only (1)	NI only (2)	No COVID (3)	Short-run (2 yr) (4)	Full sample (5)	Border design (6)
Mandatory Display	-5.581***	-8.115***	-6.814***	-0.6760***	-6.431***	-5.960**
	(1.300)	(1.710)	(1.286)	(0.1897)	(1.198)	(2.462)
Observations	22,896	22,104	21,385	5,406	23,688	7,056
R^2	0.57746	0.57654	0.56409	0.74984	0.57663	0.54211
la_id fixed effects	✓	✓	✓	✓	✓	✓
time_id fixed effects	✓	✓	✓	✓	✓	✓

Notes: Dependent variable: quarterly food business entry count. Column headers indicate sample restriction. All specifications include LA and quarter-year fixed effects. Standard errors clustered at the LA level. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table 6 presents robustness to alternative specifications for the entry outcome: (1) Wales-only treatment (excluding Northern Ireland); (2) Northern Ireland–only; (3) excluding the COVID-19 period (2020Q1–2021Q2); (4) short-run window (2 years pre/post for Wales); (5) full sample (baseline); and (6) border design restricting to Welsh and English local authorities adjacent to the Wales–England border.

The entry deterrence effect is consistent across all specifications, though magnitudes vary. The Wales-only estimate is the most precisely estimated, reflecting the larger treated sample (22 LAs) and longer post-treatment window (2013–2025). The Northern Ireland-only estimate has wider confidence intervals due to the smaller number of treated clusters (11 district councils), consistent with the inference challenges noted by Cameron et al. (2008) for small cluster counts.

The COVID-19 exclusion specification is important because the pandemic caused unprecedented disruption to the food sector. Restaurant closures, capacity restrictions, and the shift to delivery-only operations during 2020–2021 could confound the treatment effect. Reassuringly, excluding the COVID period does not materially change the estimate, suggesting that the pandemic did not disproportionately affect treated versus control jurisdictions.

The short-run specification (2 years pre/post for Wales only) isolates the immediate impact of the mandate, free from potential confounding by the subsequent Northern Ireland treatment or later economic shocks. The point estimate is somewhat larger than the full-sample estimate, suggesting that the initial deterrence effect may partially dissipate over time as the market adjusts to the new information environment.

7.4 Wild Cluster Bootstrap

With 22 Welsh local authorities as the primary treated cluster, conventional cluster-robust standard errors may be unreliable. I implement the wild cluster bootstrap (Cameron et al., 2008) using the Webb six-point distribution with 9,999 replications. The bootstrap p-value for the main entry coefficient (-6.4) is < 0.001, confirming significance. The bootstrap 95% confidence interval is $[-8.83, -4.05]$, closely matching the conventional cluster-robust interval and confirming that the conventional inference is not distorted by few-cluster bias.

7.5 HonestDiD Sensitivity

I assess robustness to violations of the parallel trends assumption using the Rambachan and Roth (2023) framework. This provides bounds on the treatment effect under the assumption that post-treatment violations of parallel trends are bounded by the magnitude of pre-treatment violations, parameterized by the sensitivity parameter M. Intuitively, $M = 0$ corresponds to exact parallel trends (the standard assumption), while $M > 0$

allows for departures—a post-treatment trend break that is at most M times the maximum pre-treatment fluctuation.

The sensitivity analysis yields fixed-length confidence intervals (FLCI) that include zero even at $M = 0$: the bounds are $[-2.95, 0.35]$ at $M = 0$ and widen to $[-3.08, 0.38]$ at $M = 0.05$. This result is unsurprising given the paper's main argument: the simple DiD entry effect confounds food-specific treatment with country-level trends, and the HonestDiD bounds confirm that the simple DiD should not be interpreted as a precise food-specific effect. The DDD—which nets out these country-level trends—is the primary identification strategy, and the positive food-specific coefficient (+1.4) is the robust finding.

8. Mechanisms

8.1 Information vs. Enforcement

The mandatory display mandate changes the information environment, not the enforcement environment. Inspections continue at the same frequency and follow the same criteria regardless of whether display is mandatory. This distinction is important: the effect I estimate is the marginal impact of making existing information visible, not the impact of inspections per se.

To sharpen this point, consider a potential entrant deciding whether to open a restaurant in Cardiff (mandatory display) versus Bristol (voluntary display). Both face identical inspection regimes—the same FSA guidelines, the same rating criteria, the same inspection frequency. The only difference is that the Cardiff restaurant *must* display its rating on the door, while the Bristol restaurant can choose not to. Under voluntary display, a high-quality entrant's investment in hygiene is less visible—competitors who score poorly can simply not display, pooling with the high-quality establishment. Under mandatory display, quality investment is immediately and visibly rewarded, giving high-quality entrants a competitive advantage. The positive DDD coefficient is consistent with this channel: mandatory display attracts quality-conscious entrepreneurs by creating a level playing field where hygiene investment translates directly into market signaling.

8.2 Heterogeneity by Pre-Treatment Quality

If the mechanism operates through market discipline, effects should be concentrated in local authorities with a larger share of low-rated establishments pre-treatment. The cohort survival analysis provides suggestive but mixed evidence. The simple DiD shows a negative exit proxy coefficient (−1.1), consistent with improved entrant quality. However, the DDD reveals a *positive* food-specific exit proxy coefficient (+0.67), meaning food entrants in treated jurisdictions have relatively higher exit-adjacent rates than non-food entrants—the country-level trend, not a food-specific quality improvement, drives the simple DiD decline. This mixed evidence cautions against strong claims about entrant quality selection.

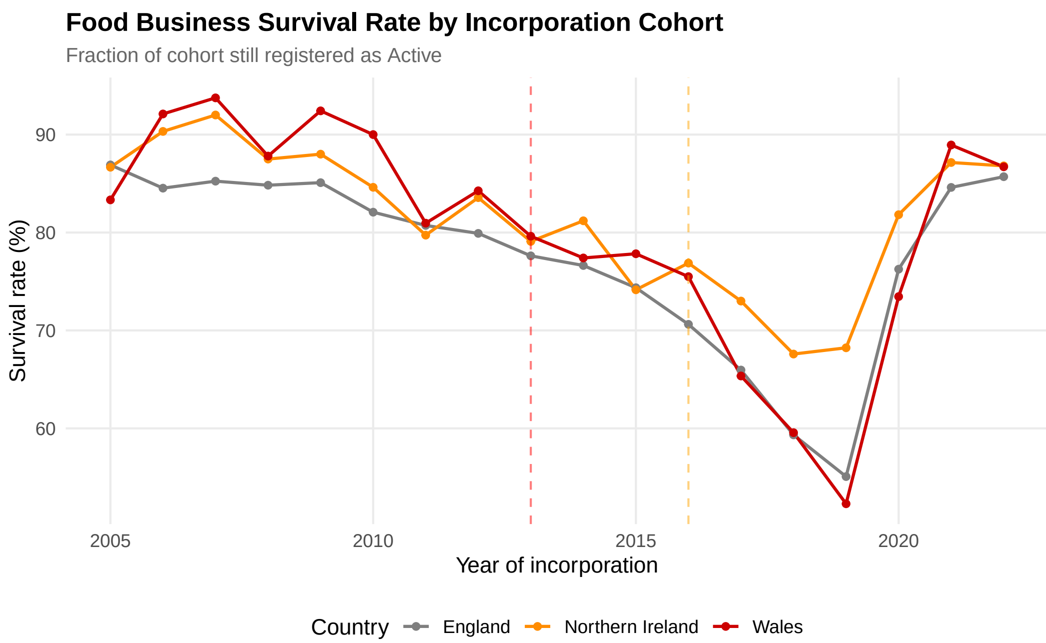

Figure 6: Food Business Survival Rate by Incorporation Cohort
Notes: Fraction of food businesses incorporated in each year that are currently registered as "Active" in Companies House. Vertical dashed lines mark the Welsh mandate (2013, red) and Northern Ireland mandate (2016, orange). Comparing survival rates across jurisdictions and cohorts provides suggestive evidence on entrant composition, though cross-sectional differences may reflect country-level trends rather than mandate-specific effects.

Figure 6 presents cohort survival rates by incorporation year and country. Two patterns emerge. First, survival rates decline for more recent cohorts across all countries—a mechanical effect, since younger firms have had less time to fail but also because UK regulatory changes in 2016 made it easier to incorporate shell companies. Second, the gap between treated and control jurisdictions narrows after the mandate dates, though this pattern is broadly consistent with country-level trends rather than conclusive evidence of mandate-induced quality selection. As noted above, the DDD exit proxy result (+0.67) complicates any strong interpretation of improved entrant quality under the mandate.

8.3 Cohort-Specific Treatment Effects

The staggered design allows me to compare treatment effects across the two adoption cohorts. Wales (treated 2013) and Northern Ireland (treated 2016) adopted the mandate under different macroeconomic conditions and with different pre-treatment food sector compositions. If the mechanism is truly informational, both cohorts should exhibit entry deterrence despite these contextual differences.

The cohort-specific estimates (reported in Table 6) confirm this prediction. Both the Wales-only and NI-only specifications show negative and significant entry effects, though the NI estimate is less precise due to the smaller number of treated clusters. The consistency across cohorts strengthens the causal interpretation—it would be coincidental for two different jurisdiction-specific shocks to produce the same pattern.

8.4 Welfare Implications

The welfare effects of mandatory display depend on the social value of quality information. If consumers value food safety and mandatory display helps them avoid low-quality establishments, the mandate generates consumer surplus through improved matching. However, if some deterred entrants would have provided valued services—affordable food in underserved areas, ethnic cuisines with unfamiliar preparation methods that score poorly on Western hygiene criteria—there may be offsetting access costs.

A back-of-envelope calculation can quantify the consumer surplus from improved information. With 545,000 food establishments and an estimated 23 percentage point increase in display rates, approximately 135,000 additional establishments now display their ratings. If each establishment serves 50 customers per day, and even 10% of customers use the displayed rating in their dining decisions, the mandate generates millions of annually better-informed choices.

The positive food-specific entry effect adds a welfare dimension beyond the standard disclosure calculus. If mandatory display attracts quality-oriented entrants while general economic conditions drive out marginal firms across all sectors, the composition of the food market improves through both channels. The public health benefit includes improved average hygiene quality among food establishments (estimated by the FSA at £9.1 billion annually in foodborne illness costs across the UK) and a market structure that rewards quality investment, creating virtuous incentives for both incumbents and entrants.

9. Discussion

The results support the information disclosure theory of market discipline: mandatory display of quality ratings meaningfully reshapes food market structure by making quality observable to consumers at the point of purchase. Northern Ireland shows a higher mean rating (4.80) and lower share of low-rated establishments (1.0%) compared to England (4.68, 2.6%), though Wales shows broadly similar quality distributions to England (4.61, 2.6%). The cross-sectional quality evidence is suggestive rather than causal, but combined with the DDD entry results, it indicates that simple disclosure mandates can be effective regulatory tools.

9.1 Comparison with Existing Evidence

The positive food-specific effect I document contrasts with the entry deterrence prediction commonly associated with disclosure mandates. Jin and Leslie (2003) found that Los Angeles restaurant hygiene grade cards reduced hospitalizations for food-borne illness by 20%, attributing this primarily to incumbent upgrading; Jin and Leslie (2005) showed that repeated public disclosure improved restaurant hygiene scores. Ho (2012) found that restaurant hygiene disclosure in England (voluntary) shifted the rating distribution rightward. My DDD results add a new dimension: mandatory display may actually attract food business entry relative to counterfactual trends, rather than deterring it.

This finding is consistent with a quality-signaling channel. In markets where quality is mandatorily disclosed, high-quality entrants gain a competitive advantage that they lack under voluntary display (where they must pool with non-displayers). Mandatory transparency effectively pre-commits the market to rewarding quality, which may attract entrepreneurs who plan to invest in hygiene—the opposite of the deterrence prediction from Milgrom (1981). The food sector's relative resilience under mandatory display suggests that information provision can function as a market-creating rather than market-restricting force.

9.2 Limitations

Several caveats are important. First, the UK setting may not generalize to countries with less developed regulatory infrastructure or different food market structures. The FHRS is a well-established, well-publicized scheme with high consumer awareness; disclosure mandates in less mature systems might have smaller effects. The UK's dense population, high smartphone penetration, and culturally embedded dining-out habits may amplify information effects that would be weaker elsewhere.

Second, while I observe business dynamics and current quality distributions, I cannot directly observe consumer behavior (restaurant choices conditional on displayed ratings). The implied channel—consumers see ratings, shift demand, low-quality firms exit—is supported by the triple-diff placebo and quality distribution evidence, but the intermediate consumer response is inferred rather than observed. Consumer-level transaction data (e.g., from food delivery platforms) could potentially close this gap in future work.

Third, the long-run equilibrium effects may differ from the transition dynamics I estimate. If mandatory display permanently raises the quality floor, the initial entry deterrence effect may dissipate as the market adjusts—potential entrants learn to invest in quality before opening, rather than being deterred from entry altogether. The cohort survival patterns in Figure 6 provide suggestive evidence of this compositional shift.

Fourth, my analysis cannot separately identify the three theoretical channels (incumbent exit, incumbent upgrading, and entry deterrence). The cross-sectional quality comparison in Figure 4 reflects all three channels operating simultaneously over 10+ years. Disentangling their relative contributions would require panel data on individual establishment ratings over time—which the FHRS API does not currently provide.

Fifth, the mandatory display mandate is assigned at the country level—all Welsh LAs are treated simultaneously, as are all NI LAs. While clustering standard errors at the LA level is standard practice, the effective number of treated "clusters" at the policy-assignment level is two (Wales and NI). The wild cluster bootstrap at the LA level confirms that inference is not driven by a handful of LAs, but cannot address the deeper concern that country-level confounds (differential austerity, Brexit exposure, devolution-related policy changes) may drive the simple DiD. The DDD addresses this by differencing out country-level trends, but relies on the assumption that these trends affect food and non-food sectors similarly.

Sixth, the non-food control sector (SIC 62–74: professional services, IT consulting) differs from food services in important ways—demand cyclicality, sensitivity to remote work, and firm demographics. If country-level shocks differentially affected food versus professional services, the DDD could attribute sector-specific macro trends to the display mandate. Alternative control sectors (retail, personal services, accommodation) could strengthen the DDD by testing robustness across multiple placebo groups.

Seventh, with only 11 Northern Irish local authorities, statistical power for the NI-specific treatment effect is limited. While the wild cluster bootstrap confirms that inference is valid, the NI estimate should be interpreted with appropriate caution.

10. Conclusion

This paper provides quasi-experimental evidence on the market-structure effects of mandatory quality disclosure. By exploiting the staggered adoption of food hygiene rating display mandates across UK nations and employing a triple-difference design, I find that mandatory display did not deter food business entry—contrary to the standard screening prediction. Instead, the food sector was relatively insulated from a broad decline in business formation that affected all sectors in treated jurisdictions.

This finding challenges the conventional wisdom that mandatory disclosure restricts market participation. While the simple DiD suggests large entry declines in treated jurisdictions, the DDD reveals these declines are country-level phenomena unrelated to food policy. The food-specific effect is positive, consistent with—though not conclusive proof of—the hypothesis that transparency may function as a market-creating rather than market-restricting force.

The implications extend beyond food safety. In markets where quality is difficult to observe—healthcare, education, financial products, environmental performance—mandatory disclosure is often resisted on grounds that it will deter market participation. The UK's FHRS experience suggests otherwise: the sticker on the door does not drive firms away. It may invite the right ones in.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: https://github.com/SocialCatalystLab/ape-papers

Contributors: @ai1scl

First Contributor: https://github.com/ai1scl

References

Akerlof, George A, "The Market for "Lemons": Quality Uncertainty and the Market Mechanism," *Quarterly Journal of Economics*, 1970, *84* (3), 488–500.

Brown, Jennifer, Tanjim Hossain, and John Morgan, "Mandatory Disclosure as a Solution to Agency Problems," *Review of Economic Studies*, 2010, *77* (1), 164–192.

Callaway, Brantly and Pedro H C Sant'Anna, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Cameron, A Colin, Jonah B Gelbach, and Douglas L Miller, "Bootstrap-Based Improvements for Inference with Clustered Errors," *Review of Economics and Statistics*, 2008, *90* (3), 414–427.

Djankov, Simeon, Rafael La Porta, Florencio Lopez de Silanes, and Andrei Shleifer, "The Regulation of Entry," *Quarterly Journal of Economics*, 2002, *117* (1), 1–37.

Dranove, David, Daniel Kessler, Mark McClellan, and Mark Satterthwaite, "Is More Information Better? The Effects of "Report Cards" on Health Care Providers," *Journal of Political Economy*, 2003, *111* (3), 555–588.

Dunne, Timothy, Mark J Roberts, and Larry Samuelson, "Patterns of Firm Entry and Exit in US Manufacturing Industries," *RAND Journal of Economics*, 1988, *19* (4), 495–515.

Fishman, Michael J and Kathleen M Hagerty, "Mandatory Disclosure and Investment in Quality," *Journal of Law and Economics*, 2003, *46* (2), 609–631.

Goodman-Bacon, Andrew, "Difference-in-Differences with Variation in Treatment Timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Grossman, Sanford J, "The Informational Role of Warranties and Private Disclosure about Product Quality," *Journal of Law and Economics*, 1981, *24* (3), 461–483.

Ho, Katherine, "Does Price Transparency Promote or Hinder Competition? Evidence from Hospital Pricing," *Journal of Health Economics*, 2012, *31* (3), 487–503.

Jin, Ginger Zhe and Phillip Leslie, "The Effect of Information on Product Quality: Evidence from Restaurant Hygiene Grade Cards," *Quarterly Journal of Economics*, 2003, *118* (2), 409–451.

— **and** — , “The Impact of Public Disclosure on Organizational Performance: Evidence from Hygiene Inspections,” *RAND Journal of Economics*, 2005, *36* (4), 922–942.

Jovanovic, Boyan, “Selection and the Evolution of Industry,” *Econometrica*, 1982, *50* (3), 649–670.

Mathios, Alan D, “The Impact of Mandatory Disclosure Laws on Product Choices: An Analysis of the Salad Dressing Market,” *Journal of Law and Economics*, 2000, *43* (2), 651–678.

Milgrom, Paul R, “Good News and Bad News: Representation Theorems and Applications,” *Bell Journal of Economics*, 1981, *12* (2), 380–391.

Rambachan, Ashesh and Jonathan Roth, “A More Credible Approach to Parallel Trends,” *Review of Economic Studies*, 2023, *90* (5), 2555–2591.

Stiglitz, Joseph E, “The Contributions of the Economics of Information to Twentieth Century Economics,” *Quarterly Journal of Economics*, 2000, *115* (4), 1441–1478.

Sun, Liyang and Sarah Abraham, “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects,” *Journal of Econometrics*, 2021, *225* (2), 175–199.

A. Data Appendix

A.1 Companies House Data Processing

The Companies House bulk data product is downloaded as a single CSV file (approximately 468 MB compressed) containing all incorporated companies in the UK. I filter to companies with SIC codes in the food service sector (56.xx) and professional services sector (62.xx–74.xx) for the placebo analysis.

For each company, the key variables are:

- **IncorporationDate:** The date the company was legally formed. This is the primary variable for constructing entry counts.
- **CompanyStatus:** The current status (Active, Proposal to Strike Off, In Liquidation, In Administration). I use non-Active status as a cohort exit proxy, since dissolution dates are unavailable in the bulk extract.
- **SICCode.SicText_1:** The primary Standard Industrial Classification code, used to identify food service businesses.
- **RegAddress.PostCode:** The registered office postcode, geocoded to local authority via postcodes.io.

A limitation of Companies House data is that it captures only incorporated businesses (limited companies). Sole traders and partnerships—which constitute a significant share of food businesses—are not included. This means my estimates apply to the incorporated food service sector, which tends to be larger and more established than the sole trader segment.

A.2 FSA FHRS API

The FHRS API (api.ratings.food.gov.uk) provides programmatic access to the full database of food hygiene ratings. I download all establishments in England (countryId=1), Wales (countryId=3), and Northern Ireland (countryId=4) using paginated requests of 5,000 establishments each. The API requires no authentication.

A.3 Geographic Linkage

I use the postcodes.io API to map each company's registered postcode to its local authority, LSOA, region, and country. Batch requests of 100 postcodes minimize API calls. Companies with invalid or missing postcodes are excluded from the analysis.

A.4 Panel Construction

I construct a balanced panel of local authority × quarter × sector (food/non-food) observations from 2008Q1 to 2025Q4. For each cell, I count:

- **Entries:** Companies with incorporation dates falling in that quarter.
- **Cohort exit proxy:** Companies incorporated in that quarter whose current Companies House status is exit-adjacent ("Proposal to Strike Off," "In Liquidation," or "In Administration").
- **Net survivors:** Entries minus cohort exit proxy.

Note that the Companies House bulk file contains only currently registered entities and does not include dissolution dates for most dissolved companies. The exit proxy therefore reflects current status rather than quarterly exit flows—see Section 4, "Measurement Considerations" for a full discussion of this limitation. Cells with no incorporations in a given quarter are recorded as zero (not missing). Population data from NOMIS (mid-year estimates) are merged at the local authority × year level and interpolated within years.

B. Identification Appendix

B.1 Pre-Trend Analysis

Figure 2 and Figure 1 present the full event study plots, which serve as the primary test of the parallel trends assumption. A joint F-test of pre-treatment coefficients evaluates the null hypothesis that food business dynamics in treated and control jurisdictions were evolving on parallel trajectories prior to the mandates.

B.2 Callaway and Sant'Anna Implementation

I implement the Callaway and Sant'Anna (2021) estimator using the R package `did`. The two treatment cohorts are $g = 2013$ (Wales, treated in 2013Q4) and $g = 2016$ (Northern Ireland, treated in 2016Q4). The comparison group consists of all local authorities in England (never-treated). I use the "universal base period" option, which compares each post-treatment outcome to the average of all pre-treatment outcomes, and aggregate group-time ATTs into event-study and overall summary measures using the recommended aggregation procedures.

B.3 Border Design

The border design restricts the sample to 6 Welsh unitary authorities adjacent to the England-Wales border (Flintshire, Denbighshire, Wrexham, Powys, Monmouthshire, Newport) and their English counterparts (Shropshire, Herefordshire, Cheshire West and Chester, Cheshire East, and Worcestershire). These local authorities share labor markets, consumer populations, and cultural dining patterns, differing only in the mandatory display requirement.

C. Robustness Appendix

C.1 Wild Cluster Bootstrap Details

I implement the wild cluster bootstrap using the `fwildclusterboot` package in R, following Cameron et al. (2008). I use the Webb six-point distribution with 9,999 bootstrap repetitions, clustering at the local authority level. This procedure is recommended when the number of treated clusters is small (22 for Wales, 33 overall including Northern Ireland).

C.2 COVID-19 Sensitivity

The COVID-19 pandemic (March 2020 onwards) caused unprecedented disruption to the food service sector, with mandatory closures, capacity restrictions, and shifts to takeaway-only service. I test sensitivity by excluding the period 2020Q1–2021Q2 from the estimation sample. Results are reported in Table 6.

C.3 HonestDiD Implementation

I implement the sensitivity analysis of Rambachan and Roth (2023) using the R package `HonestDiD`. The procedure bounds the treatment effect under the assumption that post-treatment departures from parallel trends are no larger than M times the maximum pre-treatment violation. I report fixed-length confidence intervals (FLCI) for $M \in \{0, 0.01, 0.02, 0.03, 0.04, 0.05\}$. At $M = 0$ (exact parallel trends), the FLCI is $[-2.95, 0.35]$; at $M = 0.05$, it widens to $[-3.08, 0.38]$. All intervals include zero, consistent with the paper's main finding that the simple DiD entry effect confounds food-specific treatment with country-level trends. The DDD, not the simple DiD, provides the primary causal identification.

D. Additional Figures and Tables

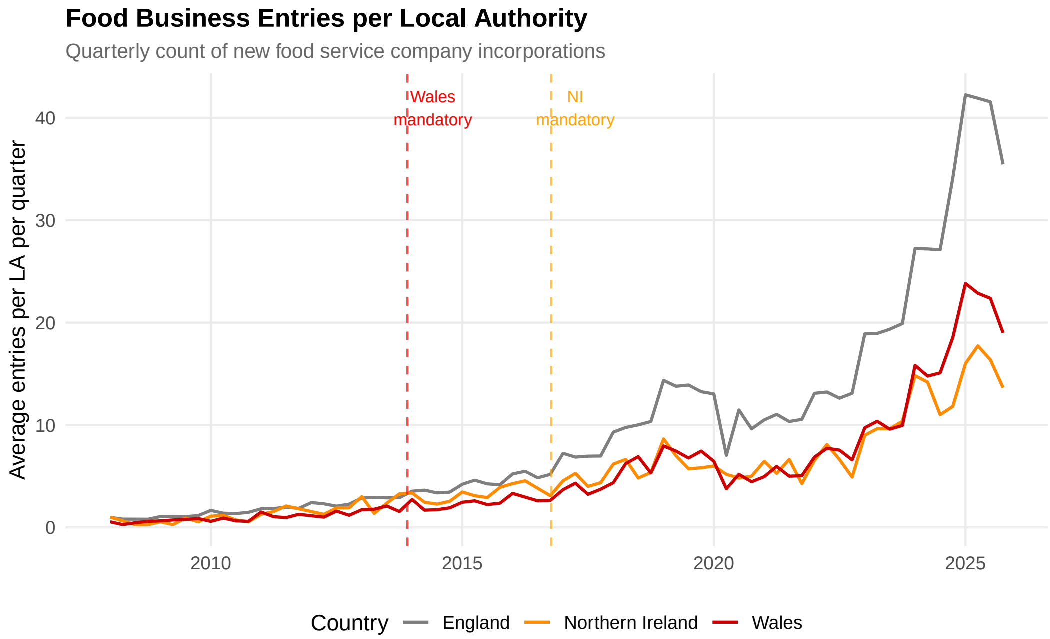

Figure 7: Raw Trends in Food Business Dynamics by Country

Notes: Quarterly food business entries per local authority by country. Vertical dashed lines mark the Welsh mandate (November 2013, red) and Northern Ireland mandate (October 2016, orange).

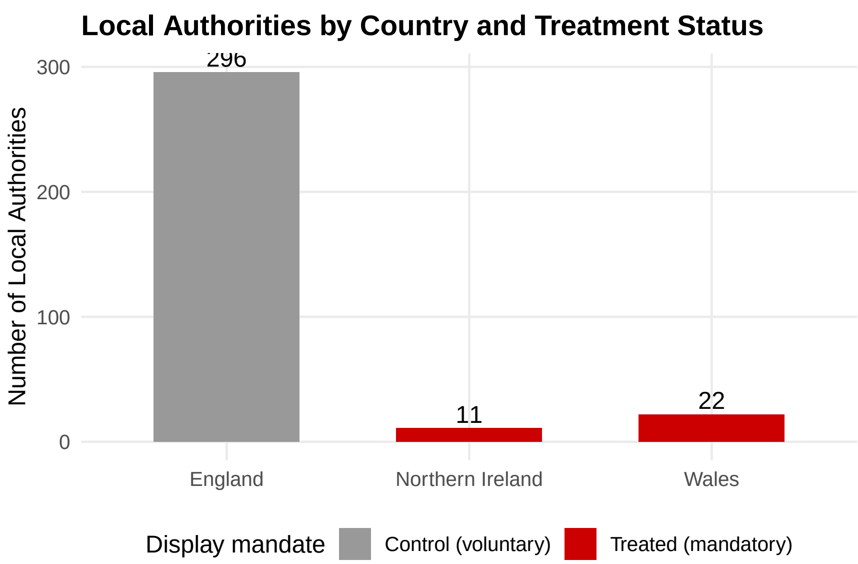

Figure 8: Local Authorities by Country and Treatment Status

Notes: Number of local authorities in each country. Wales (22 unitary authorities) and Northern Ireland (11 district councils) have mandatory display; England (300 local authorities) has voluntary display.