Zero Carbon Analytics

Briefing

Technical appendix for Climate extremes in the UK are driving up lamb prices for consumers

Joanne Bentley-McKune
March 2026

Contents

Background

This technical appendix documents the full model specification, data sources and robustness checks underlying the consumer impact estimates reported in the main briefing.

The core analysis uses a rolling-window ordinary least squares (OLS) regression to estimate how rainfall anomalies in Great Britain transmit to farmgate lamb prices, controlling for feed costs, domestic supply, temperature and seasonal patterns. The model is estimated on monthly data from January 2019 to February 2026 using publicly available Met Office climate series and AHDB market data.

Because the model relies on a relatively short time series and the climate variables are correlated by construction, we subject the results to a number of checks. These include Newey-West corrected inference to address serial autocorrelation, variance inflation diagnostics for multicollinearity, a pork falsification test (as a negative control) to confirm the rainfall signal operates through the grass-pathway mechanism specific to lamb, a lead-lag placebo test to verify the causal timing, CPI deflation to rule out inflationary confounding, and coefficient stability tests under alternative window lengths, global price controls and a Ukraine war dummy. A counterfactual decomposition is used to separate the 2025 drought's contribution from the lingering effects of the preceding wet winter.

The analysis period (2019–2026, 72 monthly observations) is short relative to the number of model parameters. With limited data, there is a risk that the model captures noise rather than genuine structure. We address this in two ways: by testing whether the model predicts accurately over a withheld period it was not trained on (‘Out-of-sample prediction’, below), and by using autocorrelation-robust standard errors throughout to avoid overstating statistical significance. Both checks suggest the estimated relationships are genuine rather than artefacts of a short sample.

Together, these checks establish that the estimated climate transmission is statistically robust, economically plausible, and specific to the grass-fed lamb supply chain.

All analyses were conducted in R 4.5.0 (2025-04-11). Please contact the author of the report for the R scripts.

Core lamb model

The OLS regression model asks: does GB rainfall, averaged over 6 and 12 months, predict farmgate lamb prices after controlling for hay costs, sheep slaughter volumes and seasonal patterns? The 6-month rainfall window has a negative coefficient

(-0.00644, Newey-West (NW) $P = 0.006$), meaning drier conditions push prices up over a 6-month horizon (consistent with a supply squeeze as grass and forage deteriorate). The 12-month window has a positive coefficient (+0.00852, NW $P = 0.010$), meaning prolonged wet conditions also push prices up (consistent with waterlogging, disease pressure, and sustained feed cost increases). The attribution logic follows from the signs: drought produces a negative anomaly, which interacts with the negative 6-month coefficient to produce a price increase. Sustained wet conditions produce a positive anomaly, which interacts with the positive 12-month coefficient to produce a price increase.

Hay price is strongly significant ($P < 0.001$), and slaughter volume has the expected negative sign (-0.71, NW $P < 0.001$): i.e. more animals coming to market depresses price. The Adjusted R^2 is 0.67. This model generates all the consumer impact numbers in the report.

Figure 1. Autocorrelation in model residuals. The autocorrelation function (ACF, left panel) measures how strongly each month's model error correlates with errors in previous months: a spike at lag 1 means this month's error tends to resemble last month's. The partial autocorrelation function (PACF, right panel) isolates the direct relationship at each lag after removing the influence of intervening months. Spikes crossing the dashed blue lines indicate statistically significant correlation. The left panel shows significant correlation at lags 1 and 2, confirming that model residuals are not independent from one month to the next. This is expected in monthly price data and is the reason we use Newey-West-corrected standard errors throughout, which account for this when assessing whether coefficients are statistically significant. The right panel shows a sharp spike at lag 1 only, suggesting the autocorrelation is primarily a first-order effect: i.e. each month's error carries over to the next, but the carryover does not compound further. This is consistent with the Newey-West lag of 3 being sufficient to correct for the dependence structure.

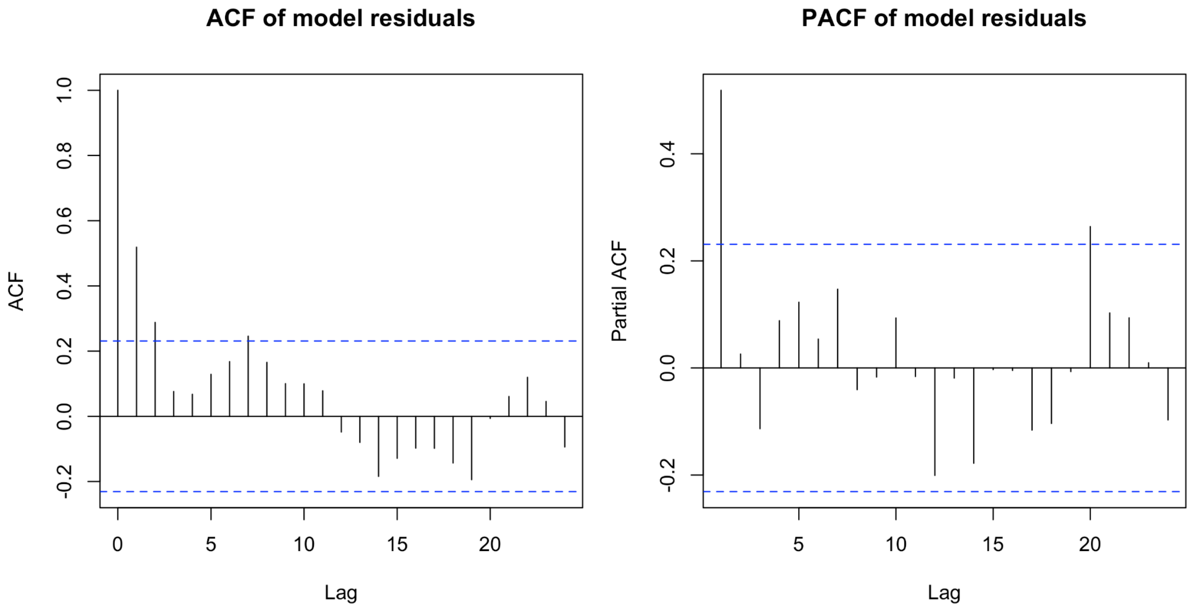

Figure 2: Climate anomalies and their effects on farmgate prices. The top panel shows the two rainfall signals the model uses: a 6-month average (red) that captures shorter shocks such as drought, and a 12-month average (blue) that captures longer disruptions such as the wet winter. When either line is away from zero, the model predicts a climate effect on prices. The bottom panel shows the result. The black line is the model's predicted farmgate price using actual climate conditions; the blue line is what they would have been under normal weather, holding everything else (hay prices, slaughter volumes, seasonal patterns) at actual values. The blue line trends upwards even without climate because those non-climate costs were rising too, so prices would have increased regardless, just not by as much. The gap between the two lines (red shaded area) is the climate premium. The key takeaway is that from 2024 onwards the black line pulls well above the blue, and stays there, reflecting first the wet winter (visible as the blue 12-month line rising sharply in the top panel) and then the 2025 drought (the red 6-month line dropping again). The two events overlap, and the gap between the price lines never closes because the next shock arrives before the previous one fades. The brief period around 2021 where the two lines converge does not reflect normal weather conditions but rather a phase where the model's short-term and long-term rainfall channels produced offsetting effects on price.

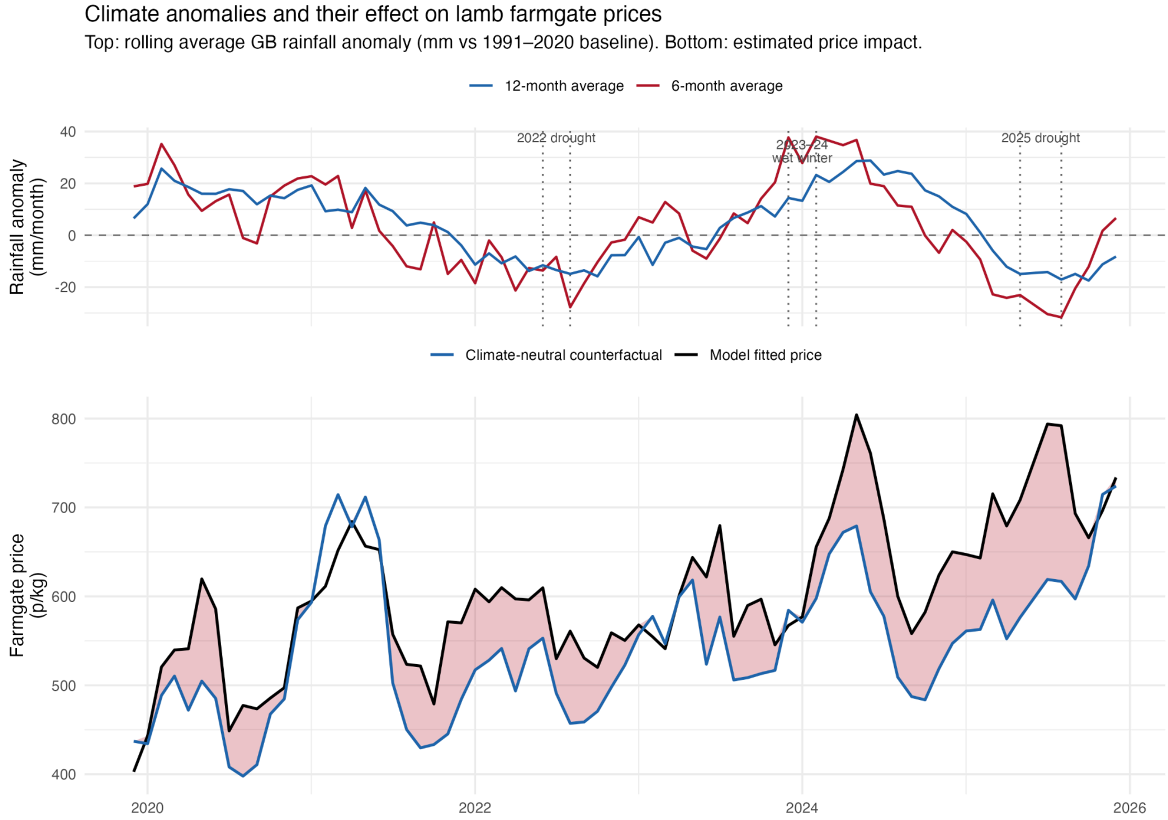

FAO-controlled lamb model

This model adds the [FAO global ovine meat price index](#) to the core lamb model to answer the question of whether the UK rainfall signal is actually just proxying for global meat market movements. The model indicates that the FAO index is non-significant (P = 0.50), and the rainfall coefficients barely move: rain_6m goes from -0.00644 to -0.00669, rain_12m from +0.00852 to +0.00833. This is consistent with the UK lamb market being largely self-contained (net exporter, >100% self-sufficiency), so domestic climate conditions dominate.

The FAO-controlled model comparison was repeated with reported Newey-West corrected inference specifically. Rain_6m under NW correction is -0.00669 (P = 0.002) and rain_12m is +0.00833 (P = 0.011), with both remaining significant. The FAO ovine index itself is confirmed at P = 0.50, carrying no additional explanatory power.

Pork falsification tests

These are the key negative controls. Pig production runs on soy feed, not grass, so if the lamb rainfall signal is genuinely operating through the grass/hay pathway, rainfall should have no predictive power for pork.

We tested five pork model specifications (using AHDB data) to verify that the rainfall signal is specific to lamb (Figure 3). Across all specifications that included rainfall (whether alone, alongside soy prices, or alongside both soy and hay) no rainfall coefficient was statistically significant. By contrast, rolling soy price windows were consistently the dominant predictor of pork prices. This is consistent with the underlying biology: pig production relies on feeds derived from soy rather than on grazed grass, so UK rainfall conditions should not (and do not) transmit to pork prices. The fact that the same rainfall variables are highly significant for lamb but carry no explanatory power for pork provides evidence that the lamb result reflects a genuine grass-pathway mechanism rather than a spurious correlation with broader economic conditions.

Figure 3. Rainfall coefficients for lamb vs pork. For lamb, the 6-month and 12-month rainfall coefficients are clearly separated from zero, confirming that UK rainfall conditions significantly predict farmgate lamb prices. For pork, all three rainfall coefficients are indistinguishable from zero, with confidence intervals centred on the zero line.

Rainfall coefficients: lamb vs pork

Newey-West 95% CIs. Significant for lamb, indistinguishable from zero for pork.

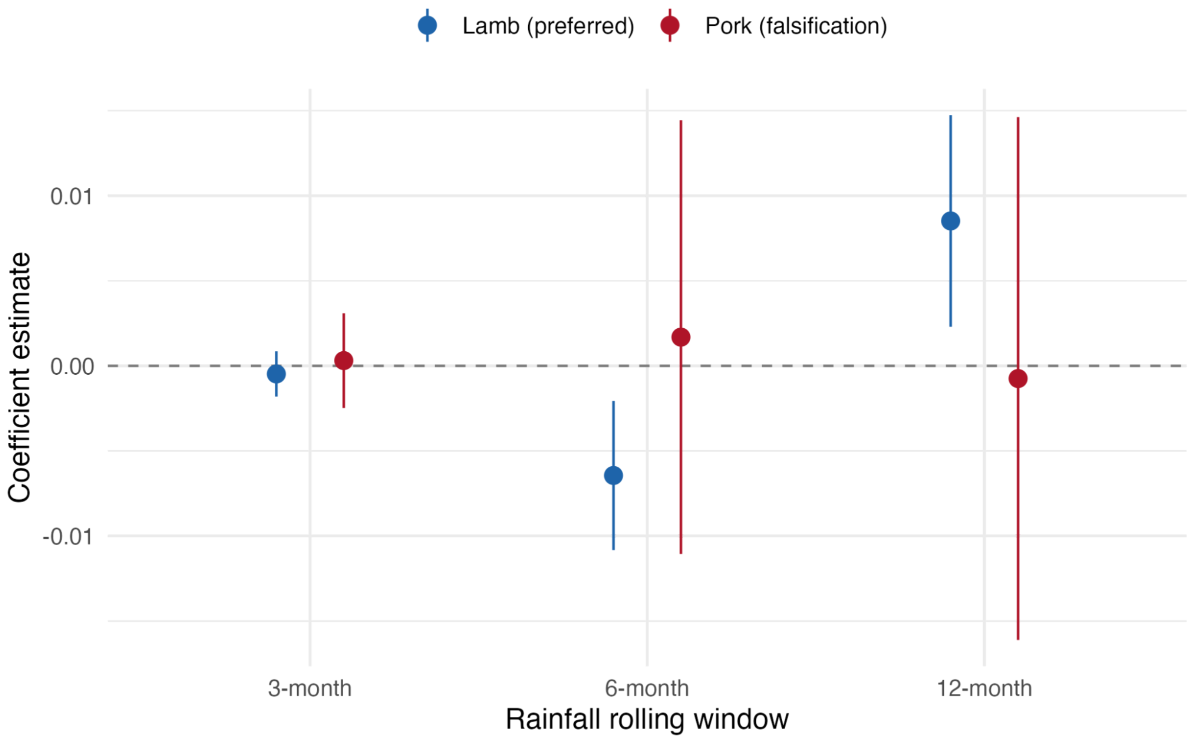

Grass growth pathway models

These models replace hay with observed GrassCheck GB grass growth data to test whether the mechanism works through actual forage conditions, not just the hay market. The model with grass and without hay finds log_grass is negative and significant (-0.24, OLS *P* = 0.030, NW *P* = 0.011), which suggests that higher pasture availability reduces lamb prices. Rain_12m remains positive and significant (OLS *P* = 0.040). The adjusted R^2 is lower (0.31 vs 0.67) because grass data is restricted to the growing season (April–October) and has fewer observations (i.e. we have fewer data points in the model). This independently corroborates that the hay-based model is likely capturing actual forage dynamics, not a statistical artefact. Under Newey-West correction, log_grass is -0.24 (*P* = 0.011): a 1% increase in grass growth reduces farmgate lamb prices by 0.24%, confirming the biological supply channel.

Slaughter mediation test

This is a two-stage test. First we ask: does rainfall predict sheep slaughter volumes? If yes, slaughter would mediate the climate effect (i.e., climate-> destocking-> price), and including both in the model would be problematic. We found that all rainfall coefficients on log_sheep_slaughter were nonsignificant (rain_3m *P* = 0.79, rain_6m *P* = 0.92, rain_12m *P* = 0.61). Climate and slaughter operate through separate channels: climate works through feed costs, slaughter through throughput volume. Stage 2 confirms that in the preferred model with both included, the climate coefficients remain significant under Newey-West correction. This justifies including slaughter as an independent supply control.

Stationarity tests: Augmented Dickey-Fuller (ADF) and Kwiatkowski-Phillips-Schmidt-Shin (KPSS)

These confirm the variables are suitable for OLS regression. log_lamb and log_hay are stationary by ADF (*P* = 0.037 and 0.033). Rainfall and temperature anomalies are ambiguous under ADF (*P* = 0.28 and 0.04) but confirmed stationary by KPSS (*P* = 0.10 and 0.07, both above the 0.05 rejection threshold). No variable is I(2), which is the key precondition.

Ukraine invasion dummy

We added a binary dummy for March–December 2022 (the commodity price shock from the invasion of Ukraine). The question is whether the rainfall coefficients are just absorbing a war-driven food price spike. Rain_6m moves from -0.00644 to -0.00654, rain_12m from +0.00852 to +0.00932, indicating the same signs and similar magnitudes. The Ukraine dummy itself is nonsignificant (*P* = 0.386). The Adjusted R^2 is unchanged (0.673 vs 0.672).

Variance inflation factors (VIF)

Rolling rainfall averages are correlated by construction (rain_6m/rain_12m correlation = 0.80), so this checks whether multicollinearity is distorting the coefficients. In the full model with month fixed effects, rain_6m has VIF = 7.73 (moderate) and rain_12m has VIF = 5.64 (moderate). Stripping month fixed effects drops these to 4.92 and 3.88. None exceeds 10, the conventional threshold for severe collinearity. The opposing signs on rain_6m and rain_12m are stable across all alternative window specifications, which further confirms collinearity isn't driving the sign structure.

CPI deflation

Between 2019 and 2025, UK food prices rose substantially. If lamb prices rose in nominal terms (the raw pound-per-kilogram price), part of that increase might simply reflect the general inflationary environment rather than climate-specific pressure. We deflated lamb prices by ONS food CPI to test whether the climate signal survives rather than being an artefact of the common inflationary trend. Rain_6m goes from -0.00644 (nominal, $P < 0.001$) to -0.00434 (real, $P = 0.017$). Rain_12m goes from +0.00852 ($P < 0.001$) to +0.00434 ($P = 0.032$). Both remain significant with the same signs. The coefficients are reduced, which is expected: CPI absorbs some of the nominal price variation, but the climate signal is not inflationary noise.

Serial autocorrelation and heteroscedasticity

Monthly price series are typically autocorrelated, in that this month's price partly predicts next month's. Standard regression assumes independent errors, and when that assumption is violated, the estimated precision of our coefficients (the *P*-values and confidence intervals) can be overstated. We also tested for heteroscedasticity, which measures whether the variance of the errors was consistent across the sample or whether it changed.

The Breusch-Godfrey test detects significant serial correlation (LM = 25.2, $P < 0.001$ at order 3), which is expected in monthly price data. The Breusch-Pagan test finds no heteroscedasticity (BP = 18.0, $P = 0.46$). The Newey-West correction (lag = 3) addresses the autocorrelation directly: under corrected standard errors, rain_6m remains significant at $P = 0.006$ and rain_12m at $P = 0.010$. This is why Newey-West *P*-values are used as the primary inference (Figure 1).

Rolling window sensitivity

This tests whether the opposing-sign two-stage structure (negative on the medium window, positive on the long window) is an artefact of the specific 3/6/12-month choice. Four specifications are compared: 3/6/12 (base), 4/6/18, 3/6/9, and 2/6/9. In all four, rain_6m is negative and significant (NW $P = 0.001$ to 0.006), and the long window is positive and significant (NW $P = 0.004$ to 0.010). The sign_consistent flag is TRUE for all four. The 3/6/9 specification has the highest adjusted R^2 (0.723). The structure is therefore robust to window choice.

Figure 4. Rainfall coefficient stability. The medium-term rainfall coefficient (blue) is consistently negative and the long-term coefficient (red) consistently positive across

all four window specifications, with all confidence intervals excluding zero. This confirms that the two-stage transmission structure (a supply squeeze over the medium term followed by sustained price elevation over the longer term) is a robust feature of the data rather than an artefact of the specific window lengths chosen.

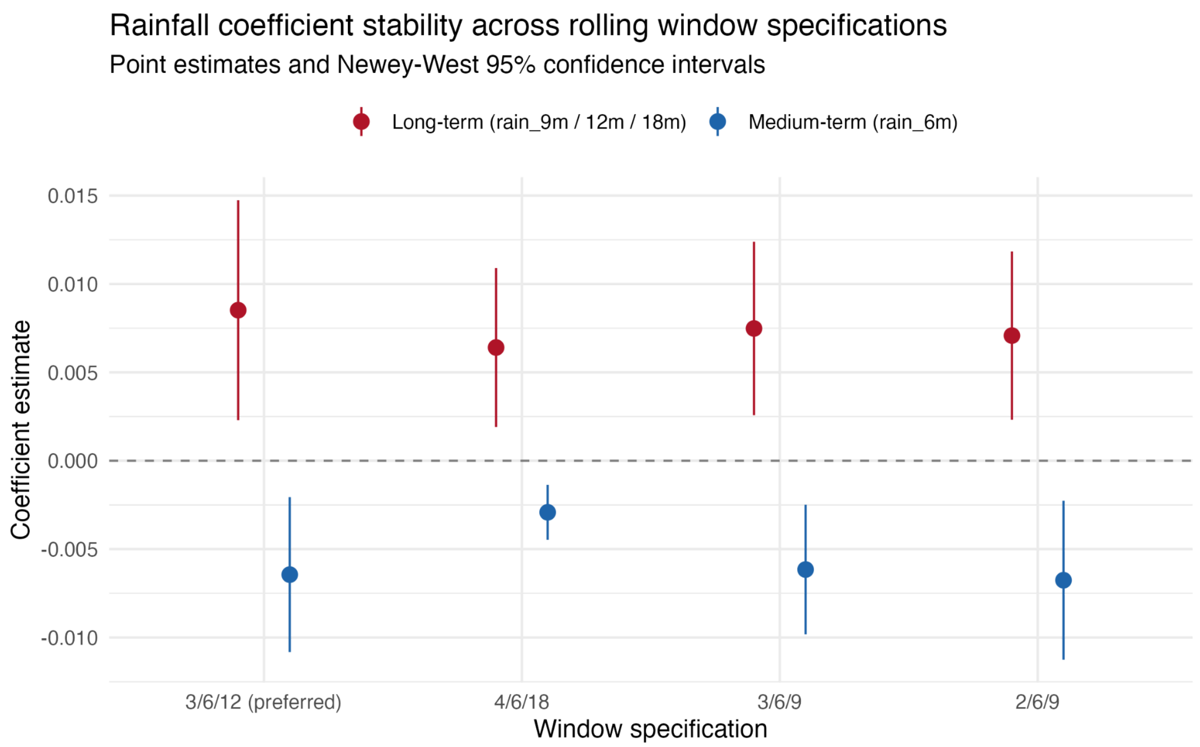

Out-of-sample prediction

A model that fits its training data well may not predict reliably beyond it, particularly with a short time series where overfitting is a real risk. We re-estimated the model using only data from 2019 to 2023 and used it to predict lamb prices over 2024 and 2025. This tests whether the climate transmission structure identified in the training period can anticipate price movements during subsequent climate extremes it has never seen. The model is re-estimated on 2019–2023 data only, then used to predict 2024–2025 lamb prices (which are withheld from fitting). Mean prediction error is 96.6p/kg (13.4%), with RMSE of 126.2p/kg. The model systematically underpredicts 2024 prices (errors of 14–28%), which is consistent with the 2023–2024 wet winter pushing prices above what the training period's climate variation could fully anticipate. Errors are small in magnitude by late 2025, suggesting the model captures the structural relationships correctly even if the magnitude of the 2024 shock was partially out-of-training-sample range. Direction accuracy is 73%, which indicates that the model correctly predicts whether prices go up or down in roughly three out of four months.

Figure 5. Out-of-sample prediction. The model, trained only on 2019–2023 data, systematically underpredicts farmgate prices in early 2024, consistent with the wet winter's effects reaching the market faster than the 12-month rolling window could absorb. By mid-2025, predictions converge closely with actual prices, suggesting the model's climate transmission structure is genuine and capable of tracking price movements outside the period it was estimated on.

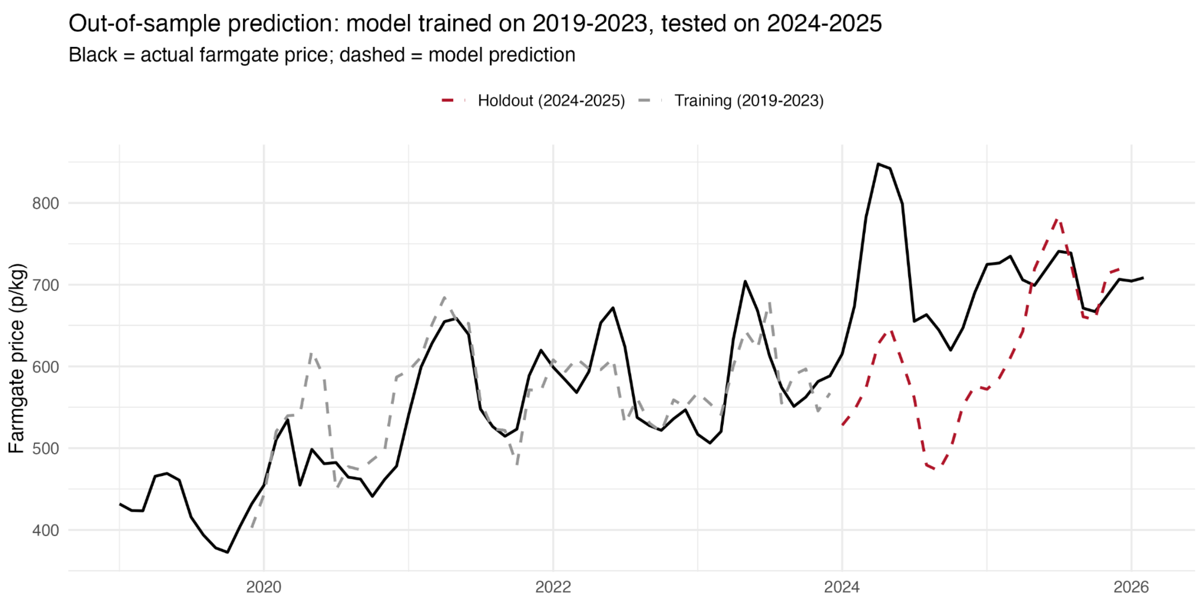

Lead-lag falsification (temporal placebo)

If the relationship between rainfall and lamb prices is genuine, it should only work in one direction: past rainfall should predict current prices, because weather conditions in previous months affect grass growth, feed costs and lamb supply in ways that take time to reach the market. Future rainfall, being conditions that have not yet occurred, should have no predictive power over today's prices, because farmers and markets cannot respond to weather that hasn't happened yet. If we found that rainfall three or six months in the future was a significant predictor of current prices, that would indicate the model is picking up a coincidental statistical pattern rather than a real economic mechanism.

This tests individual monthly lags and leads of rainfall_anomaly (at 0, 3 and 6 months in each direction) rather than rolling windows. Rain_lead3 (P = 0.14) and rain_lead6 (P = 0.44) are both nonsignificant, indicating that the temporal placebo passes. Rain_lag6 is significant (P = 0.028), independently recovering the 6-month lagged transmission structure without imposing rolling mean smoothing. Rain_lag0 (contemporaneous) is also significant (*P* = 0.021), consistent with immediate feed cost impacts. This provides independent corroboration of the lag structure.

Decomposing the 2025 drought

By 2025, rain_12m was simultaneously encoding the wet winter legacy and the emerging drought. Therefore, summing both would double-count. This analysis builds a counterfactual: set 2025 rainfall anomalies to zero in rain_3m and rain_6m (the drought channels) while leaving rain_12m at actual values (preserving the wet winter legacy). The difference between predicted prices at actual versus counterfactual values isolates the drought's additional contribution. Over March–October 2025 (the clean window before autumn rains re-enter), the mean drought premium is 92p/kg farmgate, translating to GBP 2.47 per household over that 8-month period.

Cumulative degradation of grass and hay stocks

Three consecutive climate extremes raise the question of whether the system is progressively weakening, i.e., whether the same rainfall anomaly produces worse grass growth or higher hay prices in later years. For grass growth, the rainfall × year interaction is nonsignificant (*P* = 0.63), meaning there's no evidence of worsening sensitivity. For the hay market, the rainfall × year interaction is significant (*P* = 0.007), suggesting the hay price response to rainfall anomalies has strengthened over time.