The Geography of Consumer Sentiment: Evidence from Extreme Heat*

Sankalp Mathur[†] Ronit Mukherji [‡]

September 8, 2026

Abstract

Why sentiment predicts spending is contested: it may summarize fundamentals or transmit disturbances through beliefs and demand. We study this distinction using extreme heat, an independently measured physical shock, and India's Consumer Pyramids Household Survey matched to hourly temperature data. Unusual exposure to $40–45°C$ predicts worse sentiment, weaker purchase intentions, and lower subsequent expenditure. The central finding is geographic: common heat exposure across the local area carries substantially more information about all three outcomes than exposure at the household's own location. Own exposure adds little once common conditions are included. A joint one-hour increase in common and own heat anomalies predicts 1.5 percent lower subsequent expenditure. We develop a model linking sentiment, intended purchases, and expenditure through expected local resources. Shared heat can reduce local demand and earnings, creating feedback that amplifies the initial contraction. The model derives spatial restrictions explaining how shared physical disturbances can coordinate local demand.

Keywords: Temperature Shocks, Economic Sentiment, Household Consumption, Subjective Expectations, Climate Economics

JEL Codes: E21, Q54, D84, R12, D91

*This research uses proprietary data maintained by the Centre for Monitoring Indian Economy Pvt. Ltd. (CMIE) and made available through the Centre for Economic Data and Analysis (CEDA) at Ashoka University. Corresponding Author: Ronit Mukherji (ronit.mukherji@ashoka.edu.in)

[†]Centre for Advanced Financial Research and Learning (CAFRAL), Reserve Bank of India (RBI), Mumbai, India; Email: sankalp.mathur@cafral.org.in

[‡]Ashoka University; E-mail: ronit.mukherji@ashoka.edu.in

1 Introduction

Consumer sentiment predicts household spending, yet why it does so remains contested. One view is that sentiment summarizes information about income, employment, and future financial circumstances: households become pessimistic and reduce spending because their prospects have deteriorated (Barsky and Sims, 2012). Another is that changes in outlook help transmit shocks to demand. Households may postpone purchases when they expect weaker conditions, and a coordinated reduction in spending can lower the sales, earnings, and employment opportunities on which those expectations depend (Angeletos and La'O, 2013; Benhabib et al., 2015). Understanding this relationship requires connecting the shocks households face to their sentiment, purchase plans, and subsequent spending. We analyse these household responses and ask whether they primarily reflect conditions at their own location or shocks shared across the surrounding economy.

We study these questions using extreme heat, an independently measured physical disturbance with a clearly observed geographic distribution. We match India's Consumer Pyramids Household Survey (CPHS) to ERA5, which provides hourly air-temperature estimates on a fine geographic grid. The household panel records sentiment, purchase intentions, income, and detailed expenditure. We construct an ownership-adjusted intention index and follow the same households from contemporaneous sentiment and purchase plans to expenditure in the next survey wave. The geographic detail allows us to compare exposure at the household's own location with heat shared across its local economy.

Our central finding is that extreme heat lowers sentiment, intentions and subsequent demand. The spatial dimension is critical: common heat exposure matters substantially more than own-location exposure for the observed responses. When both measures enter together, common exposure has a much larger negative coefficient for sentiment, purchase intentions, and subsequent expenditure, while own exposure adds little. We measure common exposure by averaging heat anomalies across other surveyed cells in the household's district, giving each cell equal weight. The economic distinction matters because households depend on employers, customers, shops, and services beyond their own cell. Heat across this wider area can affect their earnings and spending opportunities even when conditions at their own location are unchanged. An estimate based on own-location heat alone can capture a demand response while obscuring the role of conditions across the local economy.

We follow the same households over time and relate changes in their sentiment, purchase plans, and subsequent expenditure to unusual heat on their interview dates. For each temperature range, we measure unusual exposure as interview-day hours minus the average daily hours in the same grid cell and calendar month of the preceding year. Household fixed effects account for persistent household characteristics, while separate month-of-year effects account for common seasonal patterns. The specification retains year-to-year heat variation shared across locations because these common movements are part of the disturbance we study. Identification requires the heat anomalies to be unrelated to other unobserved changes affecting sentiment and demand, conditional on the fixed effects. Comparisons across outcomes use the same household-wave observations, and expenditure regressions additionally control for heat at the succeeding interview.

We first establish a nonlinear response of sentiment to temperature. The largest negative coefficient appears between 40 and $45°\mathrm{C}$; coefficients between 25 and $40°\mathrm{C}$ are small and statistically indistinguishable from zero. In the baseline joint-bin specification, an additional hour in the upper range relative to its benchmark predicts a 0.018 standard-deviation decline in sentiment. Occupational exposure strengthens this response. In the expanded-control estimates, the decline among households whose heads work in high-physical-intensity occupations is roughly twice that in lower-intensity occupations, and the difference is statistically significant. The significant response in the lower-intensity group also establishes that the sentiment result is not confined to households with the greatest direct workplace exposure.

The same heat episode is accompanied by weaker intended and realized demand. Extreme heat predicts lower sentiment, lower ownership-adjusted purchase intentions, and lower next-wave expenditure. An additional anomalously hot hour predicts approximately 0.8 percent lower subsequent expenditure in the household-level specification. The decline extends to both durable and nondurable spending. Its larger magnitude for durables is consistent with greater scope to postpone infrequent purchases, while the nondurable response shows that the contraction also reaches routine expenditure. Following the same households across these outcomes links the deterioration in sentiment and purchase plans to observed spending, rather than inferring household behaviour from confidence alone.

Sentiment and purchase intentions also contain distinct information about subsequent spending. When entered jointly, both predict next-wave expenditure after conditioning on heat at the interview and expenditure dates, current income, and current expenditure. This relationship appears across total, durable, and nondurable spending. An instrumental-variables estimate using extreme heat to instrument sentiment is positive and statistically significant: a one-unit fall in sentiment corresponds to a 1.32-unit fall in the purchase-intention index. Interpreting this coefficient as the effect of sentiment requires heat to influence purchase plans through sentiment. Together, these results show that households' economic assessments and stated plans have a spending counterpart beyond the interview itself.

We develop a model of consumer sentiment and demand to explain why this geographic distinction has economic meaning. Households form sentiment using information about their own circumstances, heat exposure, and expected local resources. Sentiment shapes desired purchases, which help determine subsequent expenditure. Heat can also disrupt work, travel, and exchange directly. The key interaction is that household earnings and opportunities depend partly on local spending. A shared heat disturbance can weaken expected sales and earnings, worsen sentiment, and reduce desired demand. The resulting fall in expenditure then lowers the local resources households anticipated when forming their sentiments. This feedback connects an initially physical disturbance to a coordinated demand response and yields an equilibrium multiplier separating the initial contraction from its amplification.

The model explains why common exposure can matter more than heat at the household's own location. Own heat changes the household's immediate conditions; shared heat can also change the local demand and earnings on which its prospects depend. This distinction delivers a testable benchmark: if own-location conditions fully explain household responses, common exposure should add nothing once own heat is held fixed. Departures from this benchmark reveal the importance of the wider economic environment. The model derives these restrictions for sentiment, purchase intentions, and expenditure, and shows that the sum of the common and own coefficients measures the response to a joint local increase in heat. It therefore gives economic meaning to both the spatial comparison and the overall demand contraction, allowing direct disruption and feedback through expected local activity to operate together.

We next bring the spatial implications to the data and quantify the importance of common relative to own exposure. A one-standard-deviation increase in common heat predicts declines of 0.027 standard deviations in sentiment, 0.014 in purchase intentions, and 0.024 in subsequent expenditure. The corresponding incremental own-exposure estimates are only 0.003, 0.001, and 0.005 standard deviations in magnitude. Equality of the common and own coefficients per hour of exposure is rejected for sentiment and expenditure; intentions display the same point-estimate ordering. When both exposure anomalies increase by one hour, predicted next-wave expenditure falls by approximately 1.5 percent, equivalent to about INR 205 at the sample's weighted median expenditure. The relative importance of common exposure therefore survives rescaling, and its contribution to the joint response has a tangible spending counterpart.

We assess robustness across inference, exposure measurement, timing, and household composition. Cell clustering is primary; district clustering preserves the significant common-versus-own differences for sentiment and expenditure and the negative joint responses for all three outcomes. Fixed-radius measures covering 40 and 60 kilometres reproduce the broader common-exposure pattern across administrative boundaries. Common exposure also retains its importance when the own-cell regressor is replaced by a 3×3 or 5×5 neighbourhood average, or when assigned-cell heat enters flexibly across temperature bins. In a separate exact-date 5×5 check, we compare the assigned cell with its 24 neighbours, all measured on the focal interview date. Their joint response is negative and significant for all three outcomes. The district and fixed-radius designs address our main question—whether common conditions add information beyond the household's own exposure—while the exact-date design checks that local heat measured on the interview day predicts the same set of outcomes. The common sentiment and expenditure responses also appear among urban nonagricultural households, and controlling for both common and own heat at the succeeding interview preserves the initial joint expenditure response.

1.1 Contribution to the literature

Our first contribution is to the literature on consumer sentiment and household demand. Existing work examines the forecasting content of confidence (Carroll et al., 1994; Bram and Ludvigson, 1998; Ludvigson, 2004), uses political events to study changes in sentiment (Gillitzer and Prasad, 2018; Benhabib and Spiegel, 2019; Mian et al., 2023; Ghosh et al., 2026), and investigates expectations and spending through observational and experimental evidence (Bachmann et al., 2015; Roth and Wohlfart, 2020; Lagerborg et al., 2023; Chopra et al., 2025). We combine an independently observed change in physical circumstances with within-household evidence on economic assessments, ownership-adjusted purchase plans, and subsequent expenditure. We complement information experiments by studying how household sentiment and demand respond to a naturally occurring physical disturbance. The contribution is both longitudinal and geographic: we follow the household's response across outcomes and establish that shared local exposure contains substantially more incremental information than conditions at the assigned location alone. This connects the study of sentiment to the economic environment in which households form plans and carry them out.

Our second contribution links expectation formation to the spatial transmission of disturbances. Personal experience and attention shape beliefs about economic conditions (Malmendier and Nagel, 2016; Kuchler and Zafar, 2019; Binder and Makridis, 2022; Andre et al., 2022; Kamdar and Ray, 2024; Taubinsky et al., 2026), while models of sentiment study how information and beliefs interact with endogenous activity (Lorenzoni, 2009; Barsky and Sims, 2012; Angeletos and La'O, 2013; Benhabib et al., 2015; Angeletos et al., 2018; Fève and

Guay, 2019; Acharya et al., 2021; Angeletos and Lian, 2022). We make the spatial extent of an observed disturbance a source of empirical discipline for these questions. The model derives the common, own-location, and joint responses across three household outcomes, and links them to direct disruption and feedback through expected local resources. This brings the distinction between direct exposure and adjustment across connected markets into the measurement and interpretation of consumer sentiment (Huber, 2023; Adão et al., 2026).

Our third contribution is to the economics of temperature. Research documents consequences of heat for output, labour supply, and productivity (Dell et al., 2012; Burke et al., 2015; Graff Zivin and Neidell, 2014; Somanathan et al., 2021), alongside effects on purchasing decisions, mood, and economic expectations (Conlin et al., 2007; Busse et al., 2015; Baylis, 2020; Makridis and Schloetzer, 2023; Georgarakos et al., 2025). We bring together the demand-side responses of economic sentiment, purchase plans, and later spending, and show why their measurement requires attention to the surrounding local economy. The common-exposure result adds a distinction between the household's assigned thermal conditions and the broader heat episode in which it makes economic decisions. Together with the model, it provides a framework for understanding how a physical disturbance can be accompanied by a geographically shared contraction in household demand.

The remainder of the paper proceeds as follows. Section 2 describes the data. Section 3 presents the empirical design, and Section 4 reports the household-level results. Section 5 develops the model. Section 6 examines common and own exposure and their economic interpretation. Section 7 concludes.

2 Data

The analysis combines the Consumer Pyramids Household Survey (CPHS) panel with hourly temperature data from ERA5. The unit of observation is a household interview, which links contemporaneous economic assessments and purchase plans to cell-level temperature on the realized interview date and, where available, to subsequent expenditure.

2.1 The CPHS panel

The sample covers survey waves 25–34, spanning 2022–2025. CPHS is organized on an approximately four-month interview schedule. The resulting panel is unbalanced because households may enter, leave, or miss individual waves, and we retain that structure. A next-wave outcome is defined only when household i is observed in wave $w+1$ immediately after wave w; a later observation following a missed wave is not treated as consecutive.[1] Weighted estimates use the CPHS household survey weight. Household locations are mapped to $0.25^{\circ} \times 0.25^{\circ}$ ERA5 cells. Figure 2.1 shows the geographic distribution of observations across cells and districts. The sample spans much of India, although survey density varies across locations. The CPHS sampling frame and demographic composition differ from Census and official-survey benchmarks, including undercoverage of poorer and more mobile households and differences in age, gender, education, and rural–urban composition (Somanchi, 2021; Pais and Rawal, 2021). Our estimand is therefore the weighted within-household response among households represented in the panel, rather than a national population mean. Household fixed effects remove time-invariant compositional differences, and Appendix Table A.1 shows that neither assigned-cell nor common exposure predicts re-interview in the following wave.

2.2 Consumer sentiment and purchase plans

The CPHS consumer-sentiment module asks four questions about households' current and prospective economic circumstances:

1. Compared with a year ago, how is your family faring financially these days?
   Responses: Better, Same, Worse.

2. How do you think your family will be faring financially one year from now?
   Responses: Better, Same, Worse.

3. How would you describe financial and business conditions in the country over the next 12 months?
   Responses: Good, Uncertain, Bad.

4. What do you think financial and business conditions in the country will be over the next five years?
   Responses: Continuously good times, Uncertain with ups and downs, Continuously bad times.

These items closely parallel the personal-finance and national-outlook questions in the University of Michigan Survey of Consumers and the Westpac–Melbourne Institute survey (Gillitzer and Prasad, 2018). We orient responses so that higher values denote a more favourable outlook, standardize the four components, and combine them using the inverse-covariance

[1] Households are interviewed three times a year, with each interview collecting data for the preceding quadrimester. The interviews occur in different seasons and hence span different temperature ranges. The usual schedules are January–May–September, February–June–October, March–July–November, and April–August–December. These interview months repeat across years.

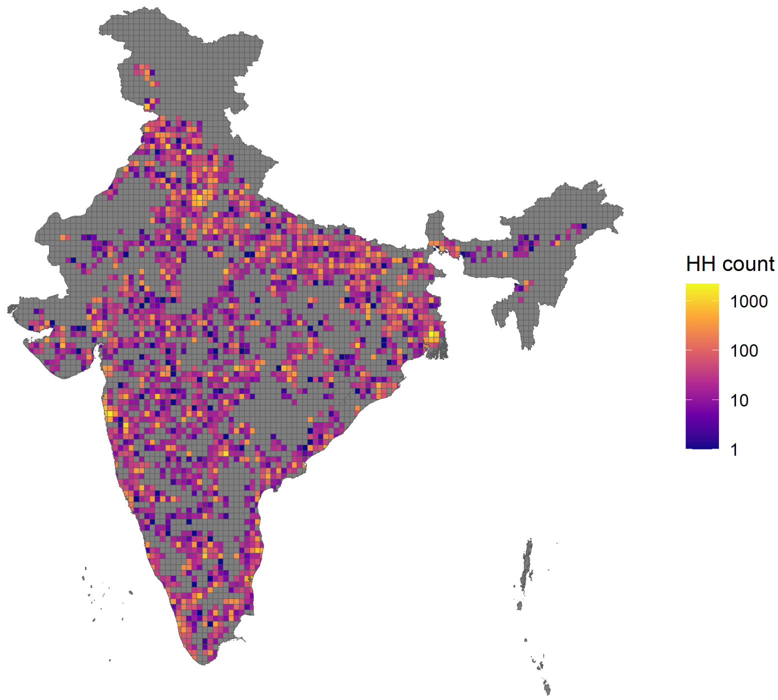

(a) Households per grid cell

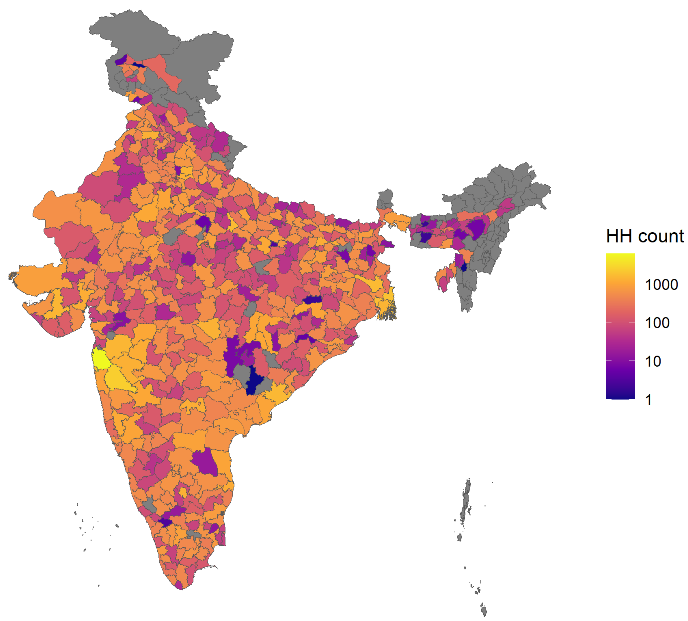

(b) Households per district

Figure 2.1: Geographic Coverage of the CPHS Sample

Notes: Panel A reports the number of CPHS household observations in each ERA5 grid cell. Panel B reports the corresponding count by district. Grey areas contain no observations in the analysis data.

weighting procedure of Anderson (2008), following Ghosh et al. (2026). The resulting primary sentiment index is normalized to have mean zero and unit variance. A broader index that includes the survey's purchase-timing question is used as an alternative measure. The four components have a standardized Cronbach's alpha of 0.676. Correlations are strongest within the two personal-finance items and within the two national-outlook items, consistent with interpreting the index as a broad measure of households' economic sentiment rather than a single narrow survey construct. Appendix Table A.2 reports the component diagnostics.

We measure intended demand using the CPHS questions asking whether the household plans to acquire each of twelve assets within the next 120 days: a house, refrigerator, air-conditioner, cooler, washing machine, television, computer, car, two-wheeler, genset or inverter, tractor, and cattle. We code each affirmative plan as one and each negative plan as zero. The ownership adjustment is applied asset by asset: a category already owned by the household is excluded from that household's acquisition set, so the measure captures planned first acquisition rather than mechanically treating current owners as households with no purchase plan. We combine the eligible components and standardize the resulting household-wave measure, which we call the all-asset ownership-adjusted purchase-intention index. A general good-time-to-buy response and an analogously constructed index restricted to consumer durables and vehicles provide complementary outcomes.

2.3 Expenditure, income, and occupational exposure

The CPHS records household expenditure across more than 150 categories. Following Mitra and Mukherji (2026), we aggregate these items into total, durable, and nondurable expenditure. The main realized-demand outcome is the inverse hyperbolic sine of total expenditure in the next survey wave. This transformation retains zero values and limits the influence of the upper tail.[2] Durable and nondurable expenditure describe the composition of the response.

The survey also records household income and member-level time use. Some expanded sensitivity specifications use current income and its square, next-wave income, and the household head's time spent working for an employer, travelling, and in indoor entertainment. Because next-wave income and contemporaneous time use may respond to heat, specifications containing these variables are not treated as the preferred reduced-form or IV estimates. The

[2]In the expenditure-composition sample before estimation-specific singleton removal, only six of $613{,}048$ observations report zero total expenditure. The corresponding counts are 465 for durable expenditure and 13 for nondurable expenditure. The transformation is therefore predominantly log-like while retaining zero-valued observations.

predictive expenditure specifications use the inverse hyperbolic sine of current income and current expenditure.

We classify the household head's occupation using the Occupational Information Network (O*NET), which records the physical and environmental characteristics of occupations on a one-to-five scale. We use the measures "Exposed to High Temperatures" and "Outdoors, Exposed to All Weather" and match CPHS occupations through the National Classification of Occupations crosswalk. Appendix Table A.3 reports the matched scores; its upper panel defines the high-physical-intensity group used in the heterogeneity analysis.

2.4 Temperature exposure

ERA5 (Hersbach et al., 2020) reports hourly two-metre air temperature on a $0.25° \times 0.25°$ grid. We assign each household the hourly temperature path in its grid cell on the realized interview date. Let τ_{it} denote that date, $c(i,t)$ the corresponding cell, and b a temperature bin. Interview-day exposure to bin b is

$$H_{c(i,t),\tau_{it},b} = \sum_{\ell=1}^{24} \mathbf{1}\left\{T_{c(i,t),\tau_{it},\ell} \in b\right\}. \tag{2.1}$$

The bins are below $0°\text{C}$, successive five-degree intervals from $0°\text{C}$ to $45°\text{C}$, and above $45°\text{C}$.

We measure unusual exposure relative to the cell's recent seasonal benchmark. For calendar month m and year y, define

$$\overline{H}_{c,m,y-1,b} = \frac{1}{D_{m,y-1}} \sum_{r\in(m,y-1)} H_{c,r,b}, \tag{2.2}$$

where $D_{m,y-1}$ is the number of days in that month. The temperature anomaly assigned to household i at interview t is

$$h^b_{it} = H_{c(i,t),\tau_{it},b} - \overline{H}_{c(i,t),m(\tau_{it}),y(\tau_{it})-1,b}. \tag{2.3}$$

One unit of h^b_{it} is therefore one additional interview-day hour in bin b relative to the preceding year's average day in the same cell and calendar month. The primary exposure, $h_{it} \equiv h^{40–45}_{it}$, is the anomaly in hours between $40°\text{C}$ and $45°\text{C}$. We refer to h^b_{it} as a temperature shock throughout. Although interview-day exposure is recorded in hours, subtracting the preceding-year monthly average makes the shock continuous and potentially negative.

CPHS records the interview date but not the interview time. The 24-hour interview-day count can therefore include hours occurring after a household reports its sentiment and purchase plans. We interpret interview-day exposure as a reduced-form marker of a persistent local heat episode surrounding the interview. Fully pre-interview exposure remains predictive: in exact-date local specifications, exposure during days 8–30 before the interview is associated with lower sentiment and purchase intentions. Future temperature windows also predict some current outcomes, consistent with strong temporal persistence rather than a conventional placebo. Appendix Table A.4 reports these temperature-window patterns. Within active cell-month risk sets, hotter assigned-cell and local days do not reduce successful-interview incidence or counts, providing no evidence that completed interviews are selected away from hot days.

Temperature is assigned using the realized interview date, while the month-of-year fixed effect is indexed by the CPHS survey month recorded for the observation.[3]

Table 2.1 describes the principal estimation samples. Extreme heat is concentrated in the upper tail of the exposure distribution. In the spatial sample, the standard deviations of the assigned-cell and common $40–45°\mathrm{C}$ anomalies are 1.114 and 0.831 hours, respectively. Actual exposure to this temperature range occurs in the assigned cell in 5.3 percent of interviews. At least one other CPHS observation used to construct the mixed-date survey-schedule proxy records such exposure in 13.3 percent of interviews.

3 Empirical Strategy

3.1 Baseline reduced forms

The baseline design combines repeated observations of the same households with the spatial and temporal resolution of ERA5. For each temperature bin, the exposure shock compares the hours observed on the household's realized interview date with the average day in the same ERA5 cell and calendar month one year earlier. The estimates therefore ask whether a household's outlook changes when its location experiences unusually greater exposure in a particular temperature range, after removing persistent household differences and the common seasonal cycle. This design retains the interannual variation in spatially coherent heat that is central to the paper's analysis of common local exposure.

[3]Field interviews can spill into an adjacent calendar month. In the 636,767-observation temporal spatial sample, 2.19 percent of observations (3.29 percent using analytic survey weights) have an actual interview month different from the recorded CPHS month. Replacing recorded month-of-year effects with actual interview month-of-year effects produces nearly identical spatial estimates. Exposure always follows the realized interview date.

Table 2.1: Summary Statistics

Panel A. Survey outcomes and transformed expenditure						
	Mean	SD	p25	Median	p75	Observations
Consumer-sentiment index	−0.027	0.720	−0.329	0.334	0.366	624,652
Purchase-intention index	0.016	1.050	−0.165	−0.154	−0.139	624,652
Current IHS total expenditure	10.151	0.464	9.851	10.148	10.439	601,791
Next-wave IHS total expenditure	10.202	0.466	9.902	10.196	10.490	624,652
Panel B. Household resources in levels						
	Mean	SD	p25	Median	p75	Observations
Current total expenditure (INR)	14,306	10,293	9,492	12,768	17,076	601,791
Next-wave total expenditure (INR)	15,174	116,681	9,988	13,404	17,974	624,652
Current income (INR)	24,255	19,515	13,350	19,500	29,000	601,791
Next-wave income (INR)	25,413	19,795	14,000	20,350	30,350	624,652
Panel C. Extreme-heat exposure						
	Mean	SD	p90	p95	p99	Observations
Assigned-cell anomaly	0.078	1.114	0.000	0.000	6.000	584,632
Leave-one-cell-out common anomaly	0.076	0.831	0.000	0.859	3.911	584,632
Panel D. Sample composition						
	Share					Observations
Urban households	0.385					624,652
Physically intensive occupations	0.632					624,652

Notes: Statistics use analytic survey weights; observation counts are unweighted. Panels A, B, and D use the exact post-singleton common-outcome sample. Current resources are not required for inclusion in that sample and are therefore observed for fewer households. Panel C uses the exact spatial sample. The sentiment and purchase-intention indices are standardized in their construction samples, so their means and standard deviations need not equal zero and one in the restricted estimation sample. IHS denotes the inverse hyperbolic sine.

Let S_{it} denote the primary sentiment index, I_{it} the all-asset ownership-adjusted purchase-intention index, and h_{it} the 40–$45°\mathrm{C}$ temperature shock. For the two outcomes measured at the interview date, we estimate

$$Y_{it} = \beta_Y h_{it} + \alpha_i + \mu_{m(t)} + u^Y_{it}, \qquad Y_{it} \in \{S_{it}, I_{it}\}. \tag{3.1}$$

The household fixed effect α_i absorbs all persistent differences across households, while the month-of-year fixed effect $\mu_{m(t)}$, with $m(t) \in \{1, \ldots, 12\}$, absorbs seasonality common to the panel. Household and recorded CPHS month-of-year effects enter separately.[4] Calendar-year effects would remove the economy-wide component of interannual heat variation and require the coefficient to be identified only from relative exposure differences within a common year.[5] Widespread interannual heat movements are part of the common disturbance studied in this paper. Appendix Table A.5 reports that alternative estimand and shows that the upper-tail coefficient is attenuated.

Identification requires the remaining temperature anomaly to be orthogonal to other time-varying determinants of household sentiment and demand, conditional on the fixed effects. All regressions use analytic survey weights. Standard errors are clustered by ERA5 cell, the level at which assigned temperature is measured; the spatial analysis additionally reports district-clustered inference because common exposure is shared within districts.

The 40–$45°\mathrm{C}$ bin is the primary exposure, but we do not impose that choice without first characterizing the response over the temperature distribution. We estimate

$$S_{it} = \sum_{b \in \mathcal{B}} \beta_b h^b_{it} + \alpha_i + \mu_{m(t)} + u^S_{it}, \tag{3.2}$$

where $\mathcal{B} = \{15\text{–}20, 20\text{–}25, 25\text{–}30, 30\text{–}35, 35\text{–}40, 40\text{–}45\}$. Since the temperature bins exhaust the 24 hours of the day, their anomalies sum to zero and one category must serve as the reference. In the baseline profile, the six bins in $\mathcal{B}$ enter jointly, while exposure below $15°\mathrm{C}$ and above $45°\mathrm{C}$ forms a pooled reference category. Each β_b measures the change in sentiment associated with shifting one interview-day hour from the reference range into bin b, relative to the preceding-year benchmark and holding exposure in the other included bins fixed.

[4] If households move, they are not followed in our sample, therefore, household fixed effects give us the narrowest location fixed effects. Results also go through with cell fixed effects.

[5] Large-scale climate modes such as El Niño and La Niña generate spatially coherent, though geographically heterogeneous, year-to-year temperature variation across tropical locations (Hsiang and Meng, 2015). Panel time effects remove the component of global disturbances that is common across locations (Dell et al., 2014, pp. 781–782); more generally, fixed-effects choices determine which components of interannual weather variation identify the response (Kolstad and Moore, 2020).

Differences between coefficients on included bins are invariant to this normalization. The profile identifies where in the temperature distribution the response emerges, motivates the upper-tail exposure used in the subsequent analysis, and allows us to test whether occupational heat intensity amplifies the response. Appendix Table A.6 enters the complete temperature distribution with $25\text{–}30°\mathrm{C}$ as the reference; the upper-tail result is unchanged.

3.2 Next-wave expenditure and predictive relationships

Let

$$C_{i,w+1} = \operatorname{asinh}\left(\text{Total expenditure}_{i,w+1}\right), \tag{3.3}$$

denote expenditure in the immediately succeeding survey wave. Equation (3.3) defines the outcome rather than an estimating equation. We construct it only when household i is observed in the immediately succeeding survey wave; an observation following a missed wave is not treated as $w+1$. Its reduced-form equation is

$$C_{i,w+1} = \beta_C h_{iw} + \rho h_{i,w+1} + \alpha_i + \mu_{m(i,w)} + u^C_{i,w+1}. \tag{3.4}$$

Each observation is indexed by the household's interview in wave w. The regressor h_{iw} measures heat at that interview, while $h_{i,w+1}$ controls for heat when expenditure is subsequently observed. The fixed effect $\mu_{m(i,w)}$ corresponds to the recorded CPHS survey month-of-year of the wave-w observation. We estimate equation (3.4) by weighted least squares using the wave-w analytic survey weight, absorb separate household and month-of-year fixed effects, and cluster standard errors by the household's assigned ERA5 cell. For direct comparison across outcomes, the sentiment, purchase-intention, and expenditure reduced forms use the common sample on which all three outcomes are observed.

We next relate the survey measures at wave w to expenditure in wave $w+1$:

$$\begin{aligned} C_{i,w+1} = {} & \gamma_S S_{iw} + \gamma_I I_{iw} + \widetilde{\beta}_C h_{iw} + \widetilde{\rho} h_{i,w+1} + \alpha_i + \mu_{m(i,w)} \\ & + \xi_C C_{iw} + \xi_Y \operatorname{asinh}\left(\text{Income}_{iw}\right) + v_{i,w+1}. \end{aligned} \tag{3.5}$$

We estimate γ_S and γ_I separately and jointly. The resource-adjusted specification includes current IHS expenditure and income; the baseline predictive specification omits these controls. All columns reported in Appendix Table A.7 use the same complete sample, so differences across specifications do not reflect changes in sample composition. The coefficients γ_S and γ_I measure the incremental predictive content of current sentiment and purchase intentions for subsequent expenditure.

The equations without potentially heat-responsive controls define the baseline reduced forms. Expanded specifications containing next-wave income, current income and its square, time-use measures, and occupational physical intensity are reported only as sensitivity checks because some of these variables may respond to heat.

4 Empirical Results

The empirical analysis connects extreme heat to three stages of household demand: economic beliefs or sentiment, intended purchases, and subsequent expenditure. We first establish where the sentiment response appears in the temperature distribution and how it varies with occupational exposure. We then follow the same households from their reported sentiment and plans to spending in the next survey wave. Finally, we examine the relationship between the two survey measures and their separate predictive content for expenditure. The economic question is whether heat is associated only with a less favourable survey response or also with changes in what households plan to buy and ultimately spend.

All estimates use the weights and separate household and month-of-year fixed effects described in Section 3. Single-outcome regressions use the available sample for that specification; comparisons across the three outcomes use identical household-wave observations. This holds sample composition fixed when moving from sentiment to plans and spending. Expanded specifications add income, time-use, and occupational controls as stability checks. The baseline leaves these potential responses to heat outside the conditioning set. Appendix Table A.8 reports the same results with a broader sentiment index that includes a good-time-to-buy question.

4.1 Extreme heat and consumer sentiment

The six-bin estimates reveal a sharply nonlinear sentiment response. In the specification that enters all six temperature shocks jointly, one additional 40–$45°\mathrm{C}$ hour relative to the preceding-year cell-month benchmark predicts a 0.01820 standard-deviation decline in sentiment (SE 0.00342). Figure 4.1 plots the expanded-control specification, which produces the same shape and a corresponding coefficient of -0.01713. The coefficients between $25°\mathrm{C}$ and $40°\mathrm{C}$ are small and statistically indistinguishable from zero. The upper-tail response is therefore not the continuation of a linear relationship between warmer temperatures and sentiment; it emerges when exposure reaches extreme levels.

The 15–25°C coefficients are also negative, but considerably smaller than the 40–45°C coefficient. Economically, the sharp upper-tail response matters because extreme heat can make work, travel, and ordinary transactions more difficult even when modest temperature changes have little effect on a household's outlook. The estimates locate this response without imposing a constant effect of each additional degree. Appendix Table A.6 enters the complete temperature distribution relative to the 25–30°C bin and retains the pronounced negative response between 40°C and 45°C. Exposure above 45°C is too rare to yield a precise coefficient. Appendix Table A.9 reports both joint-bin specifications.

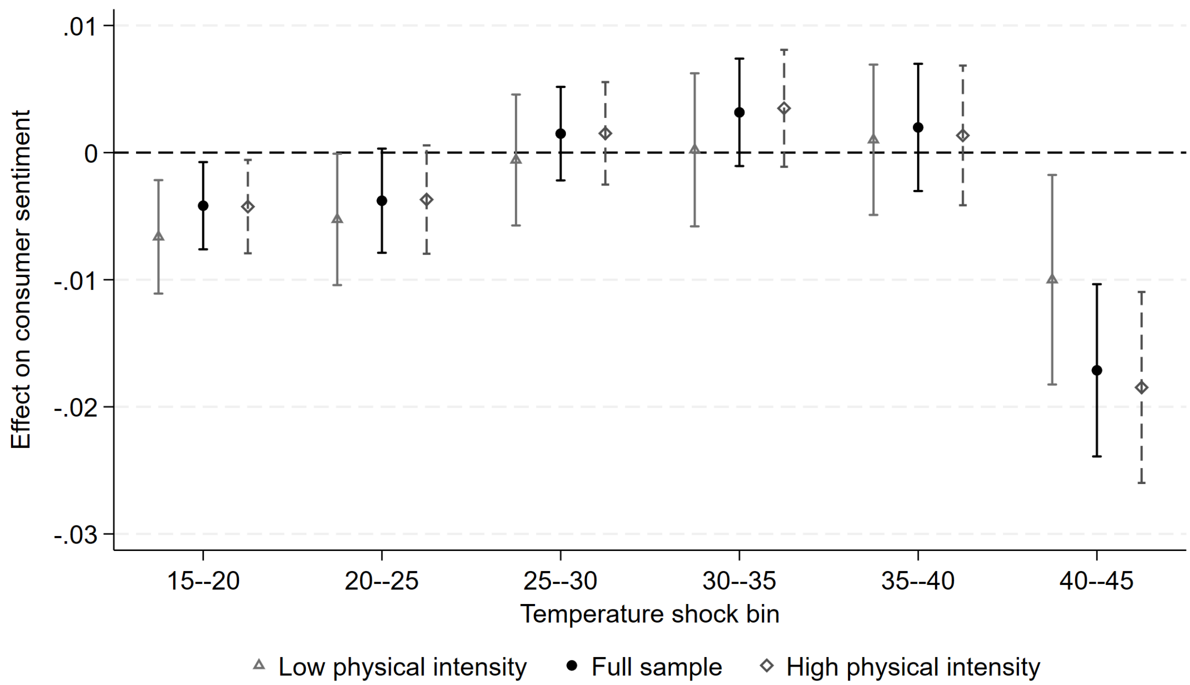

Figure 4.1: Temperature Shocks and Consumer Sentiment

Notes: The figure reports coefficients and 95 percent confidence intervals from regressions of the primary consumer-sentiment index on shocks to the six temperature bins shown on the horizontal axis. Higher values of the index denote more favourable sentiment. Black circles use the full estimation sample; triangles and diamonds use households in low- and high-physical-intensity occupations, respectively. All bins enter each regression jointly. The specifications absorb separate household and month-of-year fixed effects and condition on next-wave income, a quadratic in current income, and time spent working for an employer, travelling, and in indoor entertainment. The full-sample regression also includes physical occupational intensity. Regressions use analytic survey weights and standard errors clustered by geographic cell.

The occupational comparison helps interpret this pattern. With income and time-use controls included, the coefficient is -0.01822 among households whose heads work in high-physical-intensity occupations and -0.00856 among the lower-intensity group. The difference is -0.00966 (SE 0.00441). A larger response in physically intensive work is consistent with greater exposure to heat-related strain and interruptions to earning opportunities. Yet the lower-intensity response is also significant: a household need not face the greatest direct workplace exposure for its economic outlook to deteriorate. This motivates the later comparison between heat at the household's assigned location and conditions in the wider local economy. Appendix Table A.10 reports the interaction estimates.

4.2 Extreme heat and household demand

We next ask whether the deterioration in sentiment has a counterpart in household demand. The ownership-adjusted intention index records plans to acquire assets the household does not already own, distinguishing intended acquisitions from existing holdings. Realized demand is measured by the inverse hyperbolic sine of total expenditure in the following wave. A household facing weaker economic prospects may delay a planned purchase or reduce other spending; following both measures shows whether the change in outlook extends to decisions with a financial consequence.

Table 4.1 reports the three reduced forms on the same $624{,}652$ household-wave observations. The coefficients are -0.01183 for sentiment, -0.00940 for purchase intentions, and -0.00792 for next-wave expenditure, each precisely estimated. An additional anomalously hot hour therefore predicts approximately 0.8 percent lower subsequent expenditure. The spending equation controls for heat when expenditure is observed, separating the initial episode from conditions at the later interview. The result links weaker contemporaneous sentiment and plans to a spending decline recorded approximately four months later.

The estimates are similar with expanded controls and alternative purchase-intention measures. Neither assigned-cell nor common heat significantly predicts observation in the following wave, providing no evidence of heat-related attrition on this margin. Appendix Tables A.11 and A.1 report these checks.

Sentiment and purchase intentions: IV evidence. To examine the link between sentiment and purchase plans, we instrument the primary sentiment index with the 40–$45°\mathrm{C}$ shock. Table 4.2 reports a baseline 2SLS coefficient of 1.321 (SE 0.513): a one-unit decline in sentiment corresponds to a 1.321-unit decline in the intention index. The estimate is the ratio of the negative heat coefficient for purchase intentions, -0.01187, to the negative first-stage coefficient for sentiment, -0.00898. Its positive sign has a direct economic interpretation: the variation in sentiments associated with extreme heat is accompanied by weaker, rather than

Table 4.1: Extreme Heat, Consumer Sentiment, Purchase Plans, and Expenditure

	Sentiment index	Purchase intentions	Next-wave expenditure
	(1)	(2)	(3)
40–45°C shock	-0.01183^{***}	-0.00940^{***}	-0.00792^{***}
	(0.00243)	(0.00312)	(0.00156)
Heat at expenditure date	No	No	Yes
Household fixed effects	Yes	Yes	Yes
Month-of-year fixed effects	Yes	Yes	Yes
Observations	624,652	624,652	624,652
Geographic-cell clusters	1,819	1,819	1,819

Notes: Each column reports a separate reduced-form regression on an identical sample. The sentiment index excludes the survey's "good time to buy" question. Purchase intentions are adjusted for current ownership across all asset categories. Next-wave expenditure is transformed using the inverse hyperbolic sine. All specifications use analytic survey weights and standard errors clustered by geographic cell. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

stronger, plans to acquire assets.

Interpreting this slope as the effect of sentiment requires the exclusion restriction that heat changes purchase plans through sentiment. The IV interpretation therefore rests on that restriction, whereas the negative heat responses of sentiment and intentions do not. The baseline Kleibergen–Paap statistic is 16.96, and the Anderson–Rubin test rejects a zero coefficient with $p = 0.0002$. District clustering leaves the coefficient unchanged, raises its standard error to 0.756, and gives a Kleibergen–Paap statistic of 8.61 and an Anderson–Rubin p-value of 0.008. The latter provides inference robust to a weak first stage. Appendix A.12 sets out the specification; Appendix Tables A.12 and A.13 report the alternative intention measures and instrument diagnostics.

4.3 Sentiment, purchase plans, and subsequent expenditure

Two questions remain about spending. Which purchases account for the decline, and do sentiment and plans each predict expenditure after observed heat and household resources are taken into account? Table 4.3 answers both questions on a common sample of 600,237 observations. Panel A shows that extreme heat predicts lower total, durable, and nondurable

Table 4.2: Heat-Induced Sentiment and Purchase Intentions

	OLS	Instrumental variables	
	Expanded controls	Baseline	Expanded controls
	(1)	(2)	(3)
Sentiment index	-0.01575	1.32141^{**}	0.53278^{**}
	(0.02563)	(0.51284)	(0.25157)
First stage			
40–45°C shock	—	-0.00898^{***}	-0.01658^{***}
		(0.00218)	(0.00308)
Reduced form			
40–45°C shock	—	-0.01187^{***}	-0.00883^{**}
		(0.00314)	(0.00388)
Kleibergen–Paap F-statistic	—	16.96	28.94
Anderson–Rubin p-value	—	0.0002	0.0230
Endogeneity-test p-value	—	—	0.0155
Expanded controls	Yes	No	Yes
Household fixed effects	Yes	Yes	Yes
Month-of-year fixed effects	Yes	Yes	Yes
Observations	360,995	907,403	360,995
Geographic-cell clusters	1,757	2,008	1,757

Notes: The dependent variable is the all-asset ownership-adjusted purchase-intention index constructed from CPHS responses. Column (1) reports the conditional OLS relationship with both the expanded controls and the heat shock included. Column (2) is the baseline IV specification because it excludes next-wave income and time-use variables that may respond to heat; column (3) is an expanded-control sensitivity. All specifications use analytic survey weights and cell-clustered standard errors. The Anderson–Rubin test is robust to weak identification. The text states the exclusion restriction for interpreting the IV estimates. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

expenditure in the following wave. The durable coefficient, -0.00916, is larger in magnitude than the nondurable coefficient, -0.00416. Durables often involve infrequent purchases whose timing can be adjusted when prospects worsen. Routine expenditure offers less scope for such postponement, yet it also falls. The contraction therefore reaches beyond the asset purchases recorded in the intention index.

Table 4.3: Extreme Heat, Sentiment, and Subsequent Expenditure

	Total expenditure	Durable expenditure	Nondurable expenditure
	(1)	(2)	(3)
Panel A. Reduced-form effect of extreme heat			
40–45°C shock	-0.00758^{***}	-0.00916^{***}	-0.00416^{***}
	(0.00157)	(0.00293)	(0.00134)
Heat at expenditure date	Yes	Yes	Yes
Current income and expenditure	No	No	No
Panel B. Predicting next-wave expenditure			
Consumer sentiment	0.03689^{***}	0.07389^{***}	0.02606^{***}
	(0.00318)	(0.00547)	(0.00318)
Purchase intentions	0.00397^{**}	0.01109^{***}	0.00375^{**}
	(0.00170)	(0.00245)	(0.00163)
40–45°C shock	-0.00706^{***}	-0.00816^{***}	-0.00404^{***}
	(0.00149)	(0.00288)	(0.00126)
Heat at expenditure date	Yes	Yes	Yes
Current income and expenditure	Yes	Yes	Yes
Household fixed effects	Yes	Yes	Yes
Month-of-year fixed effects	Yes	Yes	Yes
Observations	600,237	600,237	600,237
Geographic-cell clusters	1,806	1,806	1,806

Notes: All expenditure outcomes are transformed using the inverse hyperbolic sine. Panel A reports separate reduced-form regressions. Panel B enters sentiment and purchase intentions jointly and additionally controls for current IHS income and the current IHS value of the corresponding expenditure category. Both panels condition on heat at the expenditure date. All specifications use the same sample, separate household and month-of-year fixed effects, analytic survey weights, and cell-clustered standard errors. Panel B reports conditional predictive associations. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Panel B enters sentiment and purchase intentions jointly and conditions on current income, the corresponding current expenditure category, and heat at both dates. Both survey measures predict subsequent spending across all three categories. The total-expenditure coefficients are 0.03689 for sentiment and 0.00397 for intentions. The two measures answer different economic questions: sentiment records how households assess their prospects, while intentions record specific acquisition plans. Their separate predictive content indicates that a general outlook contains information beyond the purchases a household currently says it intends to make.

The heat coefficient remains negative after sentiment, plans, and current resources are included. This is economically useful: heat can affect spending through disruptions to earnings or transactions as well as through households' assessments and plans. The evidence does not require one route to displace the other. Taken together, the results establish an upper-tail sentiment response, weaker intended and realized demand associated with the same initial heat episode, and separate predictive information in sentiment and purchase plans. Section 5 brings these findings into a model of local demand. Section 6 then asks whether the relevant disturbance is captured by the household's own location or is shared across its wider economic environment.

5 A Model of Consumer Sentiment and Local Demand

Section 4 establishes that extreme heat moves household outlooks, purchase plans, and subsequent expenditure together. The household-level results motivate a further question: can heat shared across a local economy have a larger effect on sentiment and demand than heat at the household's own location? We develop a model of consumer sentiment and local demand in which outlooks respond to expected local resources, purchase plans respond to outlooks, and expenditure reflects both intended demand and direct disruption. The model brings the three empirical outcomes into a common equilibrium framework.

Consider a household whose income depends partly on local customers. Heat that discourages purchases elsewhere in the area can weaken its expected earnings even if its own location experiences little additional heat. The household may then postpone a purchase, reducing another household's sales or employment opportunities. This interaction gives a shared disturbance a reach beyond the households most directly exposed. The model formalizes that feedback and derives distinct responses to own-location heat, common heat, and an increase in both.

This feedback builds on the link between local spending and earnings (Mian and Sufi, 2014; Nakamura and Steinsson, 2014) and the interaction of expectations with economic activity (Angeletos and La'O, 2013; Angeletos and Lian, 2022). The contribution here is to connect those forces to an observed physical disturbance and derive spatial predictions across sentiment, intended demand, and expenditure. Sentiment can respond to genuine changes in prospects, and physical disruption can coexist with demand feedback; amplification does not require households to become pessimistic for no economic reason.

5.1 Setup and spatial exposure

The model is a first-order representation around a local steady state. Sentiment and purchase intentions use the scales of their observed survey indices. Household expenditure, $c_{i,t+1}$, corresponds to the inverse hyperbolic sine of next-wave expenditure, and $C_{d,t+1}$ is its local-market mean. Thus, C averages the household-level measure rather than transforming aggregate expenditure. These units make the model's responses comparable to the empirical outcomes; the coefficients linking them measure pass-through on the chosen scales.

Shared local exposure. There is a continuum of households $i \in [0,1]$ in local market d. The model period runs from the survey interview at date t to the next expenditure observation, approximately four months later. Let h_{it} denote the extreme-heat anomaly assigned to household i's cell, and define a latent shared local component as

$$x_{dt} = \int_0^1 h_{it}\, di. \tag{5.1}$$

The empirical common measures approximate this shared state in different ways. The survey-schedule proxy uses other observed cells in the district and survey month; the exact-date measure uses the 24 neighbouring cells on the focal interview date. Section 6.1 distinguishes their construction and purpose, and in the continuum benchmark, individual exposure can be written as

$$h_{it} = x_{dt} + u_{it}, \qquad \int_0^1 u_{it}\, di = 0,$$

where u_{it} is a cell-specific deviation. A normalized uniform local increase raises both the common and assigned components one-for-one, so the coefficient sum describes the joint change under the chosen partition.

Households need not observe temperatures in every other cell. Weather information, workplace conditions, service interruptions, and their own experience can inform them about the wider heat episode. The common state summarizes the conditions behind those signals, not a temperature average that households must calculate themselves.

Measurement. Latent sentiment and desired demand are linked to the observed survey indices by

$$S_{it}^{\text{obs}} = s_{it} + \eta_{it}^{S}, \qquad I_{it}^{\text{obs}} = I_{it} + \eta_{it}^{I}, \tag{5.2}$$

where the scales of s_{it} and I_{it} are normalized to the observed consumer-sentiment and ownership-adjusted purchase-intention indices, and η_{it}^{S} and η_{it}^{I} capture reporting noise and other differences between the latent variables and their survey measures. Conditional on the fixed effects, the measurement disturbances are assumed to be mean independent of the heat innovations.

Timing. The sequence of events is as follows.

1. At time t, households observe heat and other information available by the interview, including signals about local conditions and their own circumstances. The empirical daily measure represents the surrounding heat episode, as explained in Section 2.4.

2. Households form expectations about future local expenditure $C_{d,t+1}$. Based on these expectations and their information set, they form an economic outlook s_{it} and desired demand I_{it}, which are reported through the corresponding survey indices.

3. Desired demand maps imperfectly into observed expenditure $c_{i,t+1}$. Own and common heat may also affect expenditure directly.

4. Household expenditure aggregates into local expenditure $C_{d,t+1}$, which affects the earnings and economic opportunities that households anticipated when forming their sentiment.

5.2 Sentiment, purchase plans, and realized expenditure

The link from local demand to household sentiment operates through the resources and economic opportunities that depend on local activity. Let

$$y_{i,t+1}^{L} = r_{it} + aC_{d,t+1} + \varepsilon_{i,t+1}^{y}, \qquad a > 0, \tag{5.3}$$

where $y^L_{i,t+1}$ denotes the component of household resources that is sensitive to local economic activity, r_{it} collects household fundamentals known or included in the household's information set at t, and $C_{d,t+1}$ is local expenditure. The parameter a measures the sensitivity of household resources to local demand. A decline in local expenditure can reduce working hours, labour earnings, small-business profits, informal employment, or other economic opportunities. The variable $y^L_{i,t+1}$ isolates this locally sensitive component rather than total annual household income.

Economic sentiment. Households place weight $\omega \geq 0$ on expected local resources when forming their economic outlook, s_{it}:

$$s_{it} = g_{it} - \kappa_h h_{it} - \kappa_x x_{dt} + \omega \mathbb{E}_{it} y^L_{i,t+1}, \tag{5.4}$$

where g_{it} collects fundamentals that enter the household's outlook directly. Assuming $\mathbb{E}_{it}\varepsilon^y_{i,t+1} = 0$, define $f_{it} \equiv g_{it} + \omega r_{it}$ and $q \equiv \omega a$. Household sentiment can then be written as

$$s_{it} = f_{it} - \kappa_h h_{it} - \kappa_x x_{dt} + q\mathbb{E}_{it} C_{d,t+1}, \qquad \kappa_h, \kappa_x, q \geq 0. \tag{5.5}$$

Equation (5.5) gives heat two ways to affect sentiment. It can change the household's assessment directly, through κ_h and κ_x, or change what the household expects to earn from local activity. The latter route is stronger when resources depend more on local spending (a) and households place more weight on those resources (ω); their product is q. Even with $\kappa_x = 0$, heat that reduces local demand can therefore worsen sentiment when $q > 0$. Sentiment records an assessment of economic prospects, not simply current realized income.

Purchase plans. A household's economic outlook shapes its desired purchase demand:

$$I_{it} = b s_{it} + \varepsilon^I_{it}, \qquad b > 0, \tag{5.6}$$

where b measures how desired demand changes with the household's outlook, and ε^I_{it} collects other influences on purchase plans. A household that expects weaker prospects may defer an acquisition even before its current income changes. The benchmark routes the heat response of plans through sentiment; this is the exclusion restriction underlying the IV interpretation in Section 4.2. Reporting noise and other influences on both indices mean that the OLS relationship need not equal b.

Next-period expenditure. Stated purchase plans need not translate fully into realized expenditure. Moreover, heat may affect realized expenditure without first operating through sentiment or desired demand. We therefore write next-period expenditure as

$$c_{i,t+1} = \phi I_{it} - \delta h_{it} - \zeta x_{dt} + \varepsilon^c_{i,t+1}, \qquad \phi > 0, \quad \delta, \zeta \geq 0. \tag{5.7}$$

The parameter ϕ translates intended demand into expenditure in the following wave. It is not a probability or a fraction of plans fulfilled: the intention index and expenditure have different units, so ϕ need not be below one. The error $\varepsilon^c_{i,t+1}$ collects other influences on spending, including household resources and expenditure needs.

The direct terms distinguish disruption at the household's own location (δ) from disruption shared across the market (ζ). Difficulties working or shopping can affect the household directly; electricity outages, transport interruptions, and weaker trading conditions can affect many households together. Because the outcome is next-wave expenditure, these terms capture consequences of the initial episode that carry over into later spending. The model assumes their net effects are contractionary, $\delta, \zeta \geq 0$: lost resources and disrupted purchases outweigh any subsequent catch-up or replacement spending. These sign restrictions underlie the pass-through bound below.

Pass-through. After substituting equation (5.6), the composite sentiment-to-expenditure pass-through is

$$\lambda \equiv \phi b > 0. \tag{5.8}$$

A one-unit increase in sentiment raises desired demand by b, and each unit of desired demand raises expenditure by ϕ. Their product λ measures the spending response through this planning channel. It is economically important because it determines how strongly a change in outlook feeds into local demand, and hence how much demand can feed back into sentiment. The equilibrium depends on this product rather than on b and ϕ separately.

5.3 Local equilibrium and demand feedback

Household sentiment depends on expected local expenditure, while local expenditure depends on household sentiment and desired demand. Closing this feedback loop therefore requires households' expectations to be consistent with the expenditure that is eventually realized in the local market.

Households may hold different forecasts of local expenditure. We write:

$$\mathbb{E}_{it}C_{d,t+1} = C_{d,t+1} + \varepsilon^E_{it}, \qquad \int_0^1 \varepsilon^E_{it}\, di = 0.$$

Individual forecasts can be wrong, but the model assumes that their errors cancel across the local market. This aggregate consistency condition is

$$\int_0^1 \mathbb{E}_{it}C_{d,t+1}\, di = C_{d,t+1}.$$

We also assume that the idiosyncratic disturbances in the intention and expenditure equations integrate to zero. Household-specific forecast errors and idiosyncratic disturbances are mean independent of own and common heat and enter the regression disturbance. Substituting equations (5.5) and (5.6) into equation (5.7), using $\lambda = \phi b$, and averaging across households gives

$$C_{d,t+1} = \lambda\left[f_{dt} - (\kappa_h + \kappa_x)x_{dt} + qC_{d,t+1}\right] - (\delta + \zeta)x_{dt}, \tag{5.9}$$

where $f_{dt} \equiv \int_0^1 f_{it}\, di$ is the local mean of household fundamentals. In deriving this expression, we use $\int_0^1 h_{it}\, di = x_{dt}$: under a uniform increase in local heat, households experience both greater own exposure and greater common exposure.

To distinguish the initial effect of heat from the subsequent feedback, define

$$A \equiv \delta + \zeta + \lambda(\kappa_h + \kappa_x).$$

The term A is the decline in local expenditure caused by a one-unit uniform increase in heat before any feedback through expected local activity occurs. It combines direct expenditure disruptions in households' own locations (δ), disruptions shared across the local market (ζ), and the reduction in expenditure operating through the effects of own and common heat on sentiment and desired demand $\big[\lambda(\kappa_h + \kappa_x)\big]$.

Define the feedback gain and equilibrium multiplier as

$$\gamma \equiv q\lambda, \qquad \mathcal{M} \equiv \frac{1}{1-\gamma}. \tag{5.10}$$

The feedback gain γ measures the strength of the equilibrium loop. A one-unit decline in local expenditure lowers sentiment by q, and the resulting decline in sentiment reduces expenditure by λ. Each unit of the initial expenditure contraction therefore generates an additional contraction of $q\lambda = \gamma$.

For a one-unit uniform increase in heat, the initial expenditure decline is $-A$. The resulting decline in expected local activity generates a further response of $-\gamma A$, followed by $-\gamma^2 A$, $\gamma^2 A$, and so on. The stability condition $0 \leq \gamma < 1$ requires each successive round to be smaller than the preceding one. The total heat-induced contraction is therefore:

$$-A(1+\gamma+\gamma^2+\cdots) = -\frac{A}{1-\gamma} = -\mathcal{M}A.$$

Equilibrium. Under $0 \leq q\lambda < 1$, solving equation (5.9) gives the unique stable equilibrium:

$$C_{d,t+1} = \mathcal{M}\left[\lambda f_{dt} - \{\delta + \zeta + \lambda(\kappa_h + \kappa_x)\}x_{dt}\right]. \tag{5.11}$$

The equilibrium response separates two economic forces. The initial contraction A reflects what heat does to sentiment, plans, and transactions before local demand feeds back into household prospects. The multiplier $\mathcal{M}$ captures the additional decline arising from that feedback. A large spending response can therefore reflect a large initial disruption, strong propagation, or both.

Response to a uniform local heat increase. Let

$$s_{dt} \equiv \int_0^1 s_{it}\, di, \qquad I_{dt} \equiv \int_0^1 I_{it}\, di$$

denote mean sentiment and desired demand in the local market. Let υ index a uniform increase in local exposure, so that

$$\frac{dh_{it}}{d\upsilon} = 1, \qquad \frac{dx_{dt}}{d\upsilon} = 1$$

for every household. The derivatives below are total responses to this uniform local increase, not partial responses to common exposure holding assigned exposure fixed.

The responses of mean sentiment, desired demand, and expenditure are

$$\frac{ds_{dt}}{d\upsilon} = -\mathcal{M}\big[\kappa_h + \kappa_x + q(\delta+\zeta)\big], \tag{5.12}$$

$$\frac{dI_{dt}}{d\upsilon} = b\frac{ds_{dt}}{d\upsilon}, \tag{5.13}$$

$$\frac{dC_{d,t+1}}{d\upsilon} = -\mathcal{M}\{\delta + \zeta + \lambda(\kappa_h + \kappa_x)\}. \tag{5.14}$$

The expenditure response combines direct disruption, $\delta + \zeta$, with the response operating through sentiment and desired demand, $\lambda(\kappa_h + \kappa_x)$; the equilibrium multiplier scales both components and, when $\gamma > 0$, strictly amplifies any positive component of the initial contraction. The sentiment response similarly combines the immediate effect of own and common heat, $\kappa_h + \kappa_x$, with the deterioration in outlook generated when direct disruptions reduce local expenditure and expected local resources, $q(\delta + \zeta)$. Desired demand inherits the sentiment response through b. The model therefore allows physical disruption and sentiment-based propagation to operate simultaneously. Appendix B.1 provides the derivation.

5.4 Common versus own-cell exposure

We now connect the model to the spatial decomposition estimated below. For each outcome, let β_{zh} denote the partial response to assigned-cell exposure holding common exposure fixed, and let β_{zx} denote the partial response to common exposure holding assigned-cell exposure fixed. Here $z \in \{S, I, C\}$ denotes sentiment, desired demand, or subsequent expenditure. Linking these model responses to the empirical coefficients requires alignment in exposure definitions, timing, weights, scales, and response slopes.

The distinction between the two coefficients is economically important. The decomposition uses a continuum approximation in which the focal assigned-location unit has negligible mass in the local market. A change in that unit's exposure, holding common exposure fixed, therefore does not affect aggregate local expenditure and does not generate feedback through expected local activity. A change in common exposure, by contrast, affects conditions across the local market. It can consequently influence the household both directly and through the response of local expenditure.

Proposition (Common–own-cell decomposition). The model implies

$$\beta_{Sx} = -\mathcal{M}\left[\kappa_x + q(\delta + \zeta + \lambda\kappa_h)\right], \qquad \beta_{Sh} = -\kappa_h, \tag{5.15}$$

$$\beta_{Ix} = b\beta_{Sx}, \qquad \beta_{Ih} = b\beta_{Sh}, \tag{5.16}$$

$$\beta_{Cx} = \lambda\beta_{Sx} - \zeta, \qquad \beta_{Ch} = \lambda\beta_{Sh} - \delta. \tag{5.17}$$

Equation (5.15) contains the central spatial implication of the model. The own-exposure coefficient is $\beta_{Sh} = -\kappa_h$. Holding common exposure fixed, an increase in the focal assigned-location unit's heat exposure directly worsens the household's outlook through κ_h. Because that unit has negligible mass in the continuum approximation, the change does not move local expenditure and therefore does not generate an additional equilibrium effect on sentiment. The common-exposure coefficient contains both a direct outlook response and equilibrium feedback. The parameter κ_x captures the direct response of the household's outlook to heat shared across the local market, holding expected local expenditure fixed. The remaining terms arise because common heat lowers local expenditure and the resulting contraction feeds back into the household's outlook through q.

The component $q(\delta + \zeta)$ captures feedback generated by direct expenditure disruptions. Common heat increases own-location disruptions among other households, summarized by δ, and disruptions shared across the market, summarized by ζ. Both reduce local expenditure and thereby weaken the focal household's expected economic prospects. The component $q\lambda\kappa_h$ captures a second route: heat worsens other households' sentiment through their own exposure, causing them to reduce desired and realized demand. The multiplier $\mathcal{M}$ incorporates the successive rounds through which lower local expenditure further weakens sentiment and demand.

Equation (5.16) follows from the desired- demand equation. Because heat affects desired demand only through sentiment in the benchmark model, the common- and own-exposure responses of desired demand are proportional to the corresponding sentiment responses, with proportionality factor b. These restrictions need not hold if heat directly changes purchase plans independently of sentiment.

Equation (5.17) separates the expenditure response into a component operating through sentiment and desired demand and a component reflecting direct disruption. Common exposure affects expenditure through the sentiment channel, $\lambda\beta_{Sx}$, and through direct market-wide disruption, $-\zeta$. Own exposure similarly affects expenditure through the sentiment channel, $\lambda\beta_{Sh}$, and through direct own-location disruption, $-\delta$.

For a normalized uniform local increase, common and own exposure both rise one-for-one. Its total effect is therefore the sum of the corresponding partial coefficients, as derived in equations (5.12)–(5.14):

$$\beta_S^{\text{total}} = \beta_{Sx} + \beta_{Sh}, \qquad \beta_I^{\text{total}} = \beta_{Ix} + \beta_{Ih}, \qquad \beta_C^{\text{total}} = \beta_{Cx} + \beta_{Ch}.$$

Within the model, the coefficient sums measure the response to a normalized local heat increase, while the separate coefficients describe its spatial incidence under the chosen exposure partition. In the data, the sums instead refer to a change in which both constructed regressors rise by one unit. Appendix B.4 proves the model identities.

An exact-cell-only mechanism imposes $\kappa_x = \zeta = q = 0$, with at least one of κ_h or δ strictly positive. It predicts that heat in the household's assigned cell reduces sentiment or expenditure but that exposure outside that cell has no effect:

$$\beta_{Sx} = \beta_{Ix} = \beta_{Cx} = 0.$$

A statistically significant negative surrounding-exposure coefficient tests this benchmark's zero restriction under the mapping conditions stated above. When the household's immediate environment is defined more broadly, some locally shared exposure is reassigned from x_{dt} to h_{it}. The coefficient sum remains the response to a uniform local increase, while the separate coefficients describe the spatial partition associated with the chosen boundary.

5.5 Empirical content of the model

The model organizes the evidence around three questions. First, does heat at the household's own location matter after shared conditions are held fixed? Second, does heat elsewhere in the local economy add information after own exposure is held fixed? Third, what happens when the household and its surrounding economy become hotter together? The own coefficient, common coefficient, and their sum answer these questions, respectively. They distinguish an isolated disturbance from a shared episode capable of changing local earnings and demand.

The economic importance of common exposure follows from this distinction. A household's own conditions need not deteriorate further for weaker spending elsewhere to affect its prospects. Direct common disruption and the local-demand feedback can both make the common coefficient larger than the own coefficient. The model consequently directs attention to the relative responses across all three outcomes, rather than to sentiment alone. It also supplies a benchmark: when only the assigned location matters, all three common coefficients are zero. Section 6 examines these spatial implications and their relationship to the observed proxies.

The economic scale of sentiment pass-through. The model also disciplines how much expenditure can move through sentiment and purchase plans. Equation (5.17) gives $\beta_{Cx} = \lambda\beta_{Sx} - \zeta$. If common exposure lowers sentiment and expenditure and direct common disruption is nonnegative, it follows that

$$0 < \lambda \leq \frac{\beta_{Cx}}{\beta_{Sx}}. \tag{5.18}$$

To see the intuition, first set direct common disruption to zero. The entire expenditure response then passes through sentiment, so the ratio equals λ. With direct disruption present, part of the spending decline occurs outside that channel, making the ratio an upper limit rather than an estimate of sentiment pass-through. The reported common coefficients have the ratio

$$\frac{-0.01359}{-0.02316} = 0.587 \approx 0.59.$$

This number illustrates the scale of the coefficient ratio; applying it as an empirical bound additionally requires a common conditioning set. Table 6.1 uses the same sample and weights, but only the expenditure equation controls for heat at the succeeding interview. The displayed ratio is therefore not an established numerical bound on λ. Independently of a numerical calibration, λ remains central to the mechanism: together with the response of sentiment to local activity, q, it determines the feedback gain $q\lambda$.

6 The Spatial Structure and Economic Content of Sentiment

The household-level results establish that extreme heat predicts weaker sentiment and demand. We now ask where the relevant exposure lies. Does heat at the household's own location account for the relationship, or do conditions across the local economy add information? The model makes this an economic distinction: a household can be affected by changes in the prospects of its customers, employers, and trading partners as well as by its own immediate conditions.

6.1 Constructing common exposure

Measurement of survey-schedule common exposure. Let c index the household's ERA5 cell, d its district, y the recorded CPHS survey-cycle year, and m the recorded CPHS survey month. Let $\bar{h}_{cdym}$ denote the unweighted mean interview-day 40–$45°\mathrm{C}$ anomaly among sampled households assigned to observed cell c in district–year–month dym. Define leave-one-cell-out common exposure as

$$\widetilde{x}_{-c,dym} = \frac{1}{N_{dym} - 1} \sum_{c' \neq c} \bar{h}_{c'dym}, \tag{6.1}$$

where N_{dym} is the total number of CPHS-observed cells in that district–year–month, including c. The measure is defined when $N_{dym} \geq 2$. Equal cell weights prevent respondent counts from determining the measure, and leaving out c prevents the focal cell from entering mechanically. Because households in different cells can be interviewed on different dates, $\widetilde{x}_{-c,dym}$ is a survey-schedule proxy for shared local conditions. It is not a same-day average over every physical cell in the district. The assigned-cell regressor h_{iym} is the anomaly assigned to the focal household on its realized interview date.

We use two spatial robustness checks. First, we broaden own exposure to the average heat anomaly across a 3×3 or 5×5 neighbourhood while retaining the district common measure. This tests whether common exposure remains important when own exposure covers a wider area. Second, we jointly include own-cell exposure and average exposure in the surrounding cells of a 5×5 neighbourhood, both measured on the household's exact interview date. This checks the combined response to local heat measured on the same day. The district is the baseline area for common exposure because it brings together multiple cells and provides a setting in which infrastructure, retail activity, and labour-market conditions can be shared. It is a useful starting point rather than an assumed exact boundary of household economic connections. Fixed-radius measures allow those connections to cross district borders; neighbourhood measures check whether the assigned cell is too narrow a description of immediate exposure.

Estimating equations. The estimating equations are

$$\begin{aligned} S^{\text{obs}}_{iym} &= \alpha^S_i + \mu^S_m + \beta_{Sx}\widetilde{x}_{-c,dym} + \beta_{Sh} h_{iym} + u^S_{iym}, \\ I^{\text{obs}}_{iym} &= \alpha^I_i + \mu^I_m + \beta_{Ix}\widetilde{x}_{-c,dym} + \beta_{Ih} h_{iym} + u^I_{iym}, \\ c_{i,w+1} &= \alpha^C_i + \mu^C_m + \beta_{Cx}\widetilde{x}_{-c,dym} + \beta_{Ch} h_{iym} + \rho_h h_{i,w+1} + u^C_{iym}, \end{aligned} \tag{6.2}$$

where w indexes the initial survey wave. The household fixed effects α^z_i absorb persistent differences across households, while the recorded month-of-year effects μ^z_m absorb common seasonality. The expenditure equation conditions on assigned-cell exposure at the succeeding interview. Table 6.1 reports this matched-sample benchmark. Appendix Table C.1 re-estimates the equations on outcome-specific samples and adds the succeeding-interview common proxy $\widetilde{x}_{-c',d',y',m'}$ to the expenditure specification, where (c', d', y', m') describe the household's spatial and recorded survey-cycle position at wave $w + 1$.

Identifying variation. The common coefficient uses changes in the survey-schedule proxy, conditional on assigned-cell exposure. The assigned coefficient captures incremental variation after the common proxy is held fixed. Identification requires these residual heat measures to be orthogonal to other time-varying determinants of the outcomes. All specifications use analytic survey weights. Cell-clustered inference is primary; district-clustered estimates are reported prominently because the common regressor is shared within districts. Under the mean-independence condition in equation (5.2), the coefficients estimate responses of the normalized observed indices to the constructed proxies; Section 6.5 relates these responses to the model objects.

Assigned-cell and common exposure naturally co-move. After removing the separate household and month-of-year fixed effects, their weighted correlation is 0.632 and the partial R^2 is 0.399. Thus, own exposure explains 39.9 percent of the remaining variation in common exposure, leaving 60.1 percent for the conditional comparison. The two regressors are related, but they are not interchangeable measures of the same variation.

6.2 Common and assigned-cell exposure estimates

The two exposure regressors are measured in anomalously hot hours. Their coefficients therefore compare responses to the same unit increase, holding the other measure fixed. Their sum gives the fitted response when both rise by one hour. This distinction follows the economic logic of separating own exposure from adjustment across connected locations (Huber, 2023; Adão et al., 2026).

Main estimates. Table 6.1 shows a clear and consistent spatial ordering. Conditional on assigned-cell heat, the survey-schedule common coefficient is -0.02316 for sentiment, -0.01824 for purchase intentions, and -0.01359 for next-wave expenditure. Each is negative and statistically significant. The corresponding incremental assigned-cell coefficients are -0.00169, -0.00081, and -0.00188 and are statistically indistinguishable from zero. Equality between the common- and assigned-exposure coefficients is rejected for sentiment ($p = 0.002$) and expenditure ($p = 0.005$); purchase intentions display the same point-estimate ordering, although the difference is less precisely estimated ($p = 0.112$).

The fitted responses when common and assigned exposure each rise by one unit are -0.02485 for sentiment, -0.01905 for purchase intentions, and -0.01547 for expenditure. Common exposure accounts for nearly all of each fitted joint response. These results give the geographic comparison economic context and intuition: the exposure measure that matters most for households' sentiment also matters most for their purchase plans and later spending. The broader heat environment therefore connects the sentiment response to a corresponding contraction in household demand.

Table 6.1: Survey-Schedule Common and Assigned-Cell Extreme-Heat Exposure

	Consumer sentiment (1)	Purchase intentions (2)	Next-wave expenditure (3)
Panel A. Exposure coefficients			
Survey-schedule common exposure	-0.02316^{***}	-0.01824^{**}	-0.01359^{***}
	(0.00490)	(0.00788)	(0.00304)
Assigned-cell exposure	-0.00169	-0.00081	-0.00188
	(0.00297)	(0.00400)	(0.00173)
Heat at expenditure date	—	—	0.00445^{***}
			(0.00160)
Panel B. Spatial tests and coefficient sums			
p-value: common = assigned cell	0.002	0.112	0.005
Coefficient sum (joint exposure increase)	-0.02485^{***}	-0.01905^{***}	-0.01547^{***}
	(0.00405)	(0.00602)	(0.00261)
Household fixed effects	Yes	Yes	Yes
Month-of-year fixed effects	Yes	Yes	Yes
Observations	584,632	584,632	584,632
Geographic-cell clusters	1,796	1,796	1,796

Notes: The common proxy is the equal-cell-weighted leave-one-cell-out anomaly among other CPHS-observed cells in the same district, recorded survey-cycle year, and recorded CPHS survey month. Interviews in those cells may occur on different dates. All specifications absorb separate household and recorded month-of-year fixed effects, use analytic survey weights, and cluster standard errors by geographic cell. Column (3) uses IHS total expenditure and controls for assigned-cell heat at the expenditure date; Appendix Table C.1 reports the specification with both common and assigned exposure at the succeeding interview and district inference. The estimation sample contains 115,086 households and 1,796 geographic cells after singleton removal. Panel B reports equality tests and the coefficient sum for a joint one-unit increase in the two constructed regressors. Standard errors are in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels, respectively.

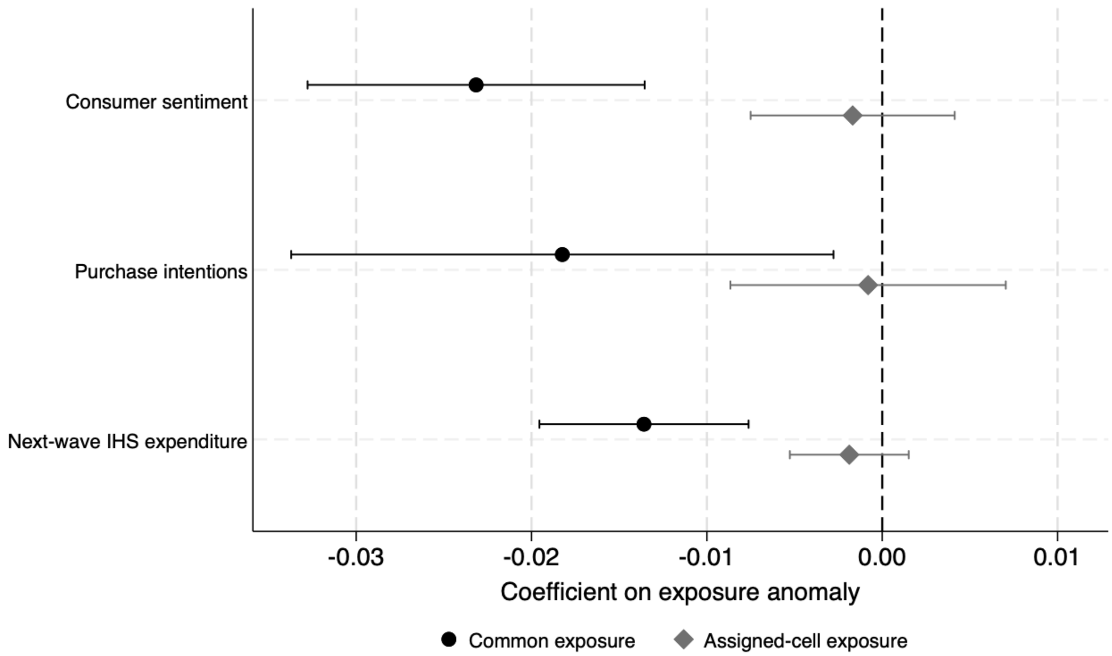

Figure 6.1: Survey-Schedule Common and Assigned-Cell Extreme-Heat Exposure

Notes: The figure plots the common- and assigned-cell exposure coefficients reported in Panel A of Table 6.1. Horizontal bars are 95 percent confidence intervals based on standard errors clustered by geographic cell. All specifications absorb separate household and month-of-year fixed effects and use analytic survey weights. The expenditure specification additionally controls for heat at the expenditure date. Coefficients are expressed in the original units of each outcome.

6.3 Economic magnitude and demand content

Common exposure is less dispersed than assigned-cell exposure: their weighted standard deviations are 0.831 and 1.114 hours, respectively. Rescaling both the exposures and outcomes preserves the spatial ordering. A one-standard-deviation increase in common exposure predicts declines of 0.0269 standard deviations in sentiment, 0.0142 in purchase intentions, and 0.0243 in next-wave expenditure. The corresponding incremental assigned-cell estimates are -0.0026, -0.0008, and -0.0045 standard deviations. Appendix Table C.2 reports the standardized component estimates. At the weighted median next-wave expenditure of INR $13{,}389$, the expenditure response when both regressors rise by one unit corresponds to approximately INR 205, or 1.5 percent.

The expenditure response matters for the local economy because one household's spending generates sales and income for others. A shared heat episode can affect the customers, employers, services, and markets on which a household depends, even when additional heat at its assigned location contributes little to the regression. Section 4 also shows that sentiment predicts subsequent expenditure after heat and current resources are taken into account. Together, the findings connect the geography of the initial disturbance to household assessments with predictive content for later demand. This is the setting in which the model's feedback can operate: weaker local prospects lead households to plan less spending, and lower spending can in turn weaken local prospects.

6.4 Robustness and complementary spatial designs

Inference and outcome timing. The point-estimate ordering is unchanged when standard errors are clustered by district. Under district clustering, the common coefficients for sentiment and expenditure remain statistically significant, their assigned-cell counterparts remain small, and equality is rejected with p-values of 0.0029 and 0.0181. Purchase intentions retain the same point-estimate ordering, although their coefficient difference is not statistically distinguishable. The joint coefficient sums remain negative and statistically precise for all three outcomes. In the expenditure equation, additionally controlling for both common and assigned exposure at the succeeding interview leaves the initial joint response at -0.01505 (cell-clustered SE 0.00268; district-clustered SE 0.00433). Appendix Tables C.3 and C.1 report these estimates.

Temporal alignment with fixed geography. We evaluate the original set of other eligible cells on each household's exact interview date, holding geographic coverage, cell weights, the estimation sample, and controls fixed (Appendix Table C.4). The joint responses remain negative and statistically significant for sentiment, purchase intentions, and subsequent expenditure under both clustering levels, although their magnitudes are smaller. The joint expenditure response is approximately 0.9 percent; common exposure remains negative and significant, while own exposure is close to zero. After household and month-of-year fixed effects, the correlation between common and own exposure rises from 0.632 to 0.944.

Alternative economic geography. Replacing districts with fixed-radius local areas produces closely similar estimates. When common exposure is constructed from other observed cells within 40 or 60 kilometres, including cells across district boundaries, the common coefficients remain negative for all three outcomes and the joint sums remain precisely estimated. We also partition observed district exposure, at each of 25-, 40-, and 60-kilometre radii, into cells inside the radius and the remaining cells outside it. The component coefficients reallocate as the boundary is redrawn, while their sums remain stable. Appendix Tables C.5 and C.6 report the two exercises.

Exact-date 5×5 local exposure. This check replaces the survey-schedule common proxy with the mean anomaly in the 24 cells surrounding the assigned cell. Both regressors are evaluated on the focal household's exact interview date, regardless of interview activity elsewhere. It therefore asks whether the same-day local heat episode predicts the outcomes. The surrounding-plus-assigned sums are -0.00965 for sentiment, -0.01209 for intentions, and -0.01005 for next-wave expenditure, negative and significant under both cell and district clustering. Adjacent-cell temperatures move closely together: separating their coefficients is difficult, especially for sentiment and intentions, even though the joint response is precisely estimated. The design supplies a temporal-alignment check on the local heat finding, not a precise ranking of the two components. A second partition of the same grid into an inner 3×3 area and a nonoverlapping outer 16-cell ring also produces negative, significant sums for all three outcomes under district clustering. Appendix Tables C.7 and C.8 report the estimates.

Broadened personal exposure. A different concern is that a household works, travels, or shops outside its assigned cell, so the common coefficient might simply compensate for a narrow own-exposure measure. We replace own-cell heat with the complete-grid average over either 3×3 or 5×5 cells, retaining the original district common proxy. The common coefficient remains materially unchanged and the broadened personal coefficient remains small. This is the 5×5 check directly relevant to that concern about the common-versus-own comparison. Unlike the exact-date surrounding-cell regression above, it compares the district proxy with a broader personal measure; some cells can enter both, so it is not a disjoint spatial partition. Appendix Table C.9 reports the estimates.

Functional form and sample composition. Allowing assigned-cell exposure to enter flexibly across all six temperature bins preserves the spatial ordering for all three outcomes. Among urban nonagricultural households, the common response remains negative for sentiment and expenditure, extending the pattern beyond agricultural households. The estimates are also stable to equal-cell versus survey-weighted common exposure and to winsorizing or trimming the upper tail of expenditure. Requiring at least five observed cells preserves the common-exposure results for all three outcomes; the much smaller ten-cell sample retains the intention and expenditure effects but is imprecise for sentiment. Appendix Tables C.10, C.11, C.12, and C.13 report these checks. Appendix C.16 separately enters current raw heat and the preceding-year benchmark.

6.5 Economic interpretation

The central lesson is that household exposure has an economic geography. In the district and fixed-radius decompositions, shared conditions contain information about sentiment and demand that own-location heat does not capture. That information remains when personal exposure is broadened to nearby cells. The exact-date local estimates add a different result: a heat episode measured across neighbouring cells on the household's interview day predicts weaker sentiment, plans, and later expenditure. This combination connects the common-exposure finding to a same-day local response without treating the two constructions as the same measure.

The model connects this geographic pattern to household expectations and local demand. A household can be affected both by disruption to its own work and by weaker conditions among its customers and employers. As households revise their outlooks and reduce planned purchases, lower expenditure can further weaken local incomes. Common exposure can therefore influence sentiment and demand through the initial shared disruption and the feedback it generates. This gives the spatial comparison an economic role: it connects the reach of an observed physical shock to the conditions that can coordinate household outlooks and spending.

Taken together, the spatial results bring the household's wider economic environment into the study of consumer sentiment. Common exposure accounts for most of the joint heat response in outlooks, purchase plans, and subsequent expenditure. The geographic pattern therefore extends from what households report to what they later spend. Connecting these responses within the same household panel makes the reach of the disturbance an economic object: it reveals the importance of shared local conditions for understanding how household sentiment and demand move together.

7 Conclusion

This paper links an independently measured physical disturbance to household sentiment and demand. Within the same households, unusual upper-tail heat predicts worse economic sentiment, weaker ownership-adjusted purchase intentions, and lower next-wave expenditure. Sentiment and purchase intentions also retain distinct predictive content for spending after accounting for heat and current household resources.

The central contribution is geographic. In the survey-schedule decomposition, common exposure across other observed cells carries most of the incremental association, while own-cell exposure adds little once common conditions are included. The common–own differences for sentiment and expenditure survive district clustering and fixed-radius definitions. A joint one-hour increase in common and own heat anomalies predicts approximately 1.5 percent lower next-wave expenditure. The exact-date 5×5 check also confirms negative joint local responses for all three outcomes.

The model gives these findings economic content. It connects sentiment to expected local resources, allows intended purchases to shape realized expenditure, and accommodates direct disruption to work and exchange. Shared heat can reduce local spending and earnings, generating feedback that amplifies the initial contraction. The equilibrium delivers distinct common, own-location, and joint-exposure responses, together with a theoretical bound on sentiment-to-expenditure pass-through under stated sign restrictions. This links the geography of the disturbance to the coordination of household outlooks and demand.

By tracing heat from economic assessments and purchase plans to subsequent expenditure, the paper extends the measured consequences of extreme temperature beyond production, labour supply, and health. The sentiment and expenditure responses among urban nonagricultural households, alongside the sentiment response in less physically intensive occupations, show that this demand margin extends beyond agriculture and the most directly heat-exposed work. Maintaining electricity, transport, payment systems, and retail activity during extreme heat can support the local economic connections that sustain household resources and demand.

References

Acharya, S., Benhabib, J., and Huo, Z. (2021). The anatomy of sentiment-driven fluctuations. *Journal of Economic Theory*, 195:105280.

Adão, R., Arkolakis, C., and Esposito, F. (2026). General equilibrium effects in space: Theory and measurement. *American Economic Journal: Macroeconomics*. Forthcoming.

Anderson, M. L. (2008). Multiple inference and gender differences in the effects of early intervention: A reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects. *Journal of the American Statistical Association*, 103(484):1481–1495.

Andre, P., Pizzinelli, C., Roth, C., and Wohlfart, J. (2022). Subjective models of the macroeconomy: Evidence from experts and representative samples. *Review of Economic Studies*, 89(6):2958–2991.

Angeletos, G.-M., Collard, F., and Dellas, H. (2018). Quantifying confidence. *Econometrica*, 86(5):1689–1726.

Angeletos, G.-M. and La'O, J. (2013). Sentiments. *Econometrica*, 81(2):739–779.

Angeletos, G.-M. and Lian, C. (2022). Confidence and the propagation of demand shocks. *Review of Economic Studies*, 89(3):1085–1119.

Bachmann, R., Berg, T. O., and Sims, E. R. (2015). Inflation expectations and readiness to spend: Cross-sectional evidence. *American Economic Journal: Economic Policy*, 7(1):1–35.

Barsky, R. B. and Sims, E. R. (2012). Information, animal spirits, and the meaning of innovations in consumer confidence. *American Economic Review*, 102(4):1343–1377.

Baylis, P. (2020). Temperature and temperament: Evidence from Twitter. *Journal of Public Economics*, 184:104161.

Benhabib, J. and Spiegel, M. M. (2019). Sentiments and economic activity: Evidence from U.S. states. *Economic Journal*, 129(618):715–733.

Benhabib, J., Wang, P., and Wen, Y. (2015). Sentiments and aggregate demand fluctuations. *Econometrica*, 83(2):549–585.

Binder, C. and Makridis, C. (2022). Stuck in the seventies: Gas prices and consumer sentiment. *Review of Economics and Statistics*, 104(2):293–305.

Bram, J. and Ludvigson, S. C. (1998). Does consumer confidence forecast household expenditure? A sentiment index horse race. *Federal Reserve Bank of New York Economic Policy Review*, 4(2):59–78.

Burke, M., Hsiang, S. M., and Miguel, E. (2015). Global non-linear effect of temperature on economic production. *Nature*, 527(7577):235–239.

Busse, M. R., Pope, D. G., Pope, J. C., and Silva-Risso, J. (2015). The psychological effect of weather on car purchases. *Quarterly Journal of Economics*, 130(1):371–414.

Carroll, C. D., Fuhrer, J. C., and Wilcox, D. W. (1994). Does consumer sentiment forecast household spending? If so, why? *American Economic Review*, 84(5):1397–1408.

Chopra, F., Roth, C., and Wohlfart, J. (2025). Home price expectations and spending: Evidence from a field experiment. *American Economic Review*, 115(7):2267–2305.

Conlin, M., O'Donoghue, T., and Vogelsang, T. J. (2007). Projection bias in catalog orders. *American Economic Review*, 97(4):1217–1249.

Dell, M., Jones, B. F., and Olken, B. A. (2012). Temperature shocks and economic growth: Evidence from the last half century. *American Economic Journal: Macroeconomics*, 4(3):66–95.

Dell, M., Jones, B. F., and Olken, B. A. (2014). What do we learn from the weather? the new climate–economy literature. *Journal of Economic Literature*, 52(3):740–798.

Fève, P. and Guay, A. (2019). Sentiments in SVARs. *Economic Journal*, 129(618):877–896.

Georgarakos, D., Kenny, G., Meyer, J., and van Rooij, M. (2025). How do rising temperatures affect inflation expectations? Working Paper 843, De Nederlandsche Bank. Revised March 2026 as ECB Working Paper No. 3132.

Ghosh, A., Mitra, A., and Mukherji, R. (2026). Identity politics, partisan sentiment and household spending. Working paper, January 2026.

Gillitzer, C. and Prasad, N. (2018). The effect of consumer sentiment on consumption. *American Economic Journal: Macroeconomics*, 10(4):234–269.

Graff Zivin, J. and Neidell, M. (2014). Temperature and the allocation of time: Implications for climate change. *Journal of Labor Economics*, 32(1):1–26.

Hersbach, H., Bell, B., Berrisford, P., et al. (2020). The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730):1999–2049.

Hsiang, S. M. and Meng, K. C. (2015). Tropical economics. *American Economic Review*, 105(5):257–261.

Huber, K. (2023). Estimating general equilibrium spillovers of large-scale shocks. *Review of Financial Studies*, 36(4):1548–1584.

Kamdar, R. and Ray, W. (2024). Attention-driven sentiment and the business cycle. CEPR Discussion Paper 18984, Centre for Economic Policy Research.

Kolstad, C. D. and Moore, F. C. (2020). Estimating the economic impacts of climate change using weather observations. *Review of Environmental Economics and Policy*, 14(1):1–24.

Kuchler, T. and Zafar, B. (2019). Personal experiences and expectations about aggregate outcomes. *Journal of Finance*, 74(5):2491–2542.

Lagerborg, A., Pappa, E., and Ravn, M. O. (2023). Sentimental business cycles. *Review of Economic Studies*, 90(3):1358–1393.

Lorenzoni, G. (2009). A theory of demand shocks. *American Economic Review*, 99(5):2050–2084.

Ludvigson, S. C. (2004). Consumer confidence and consumer spending. *Journal of Economic Perspectives*, 18(2):29–50.

Makridis, C. A. and Schloetzer, J. D. (2023). Extreme local temperatures lower expressed sentiment about U.S. economic conditions with implications for the stock returns of local firms. *Journal of Behavioral and Experimental Finance*, 37:100710.

Malmendier, U. and Nagel, S. (2016). Learning from inflation experiences. *Quarterly Journal of Economics*, 131(1):53–87.

Mian, A. and Sufi, A. (2014). What explains the 2007–2009 drop in employment? *Econometrica*, 82(6):2197–2223.

Mian, A., Sufi, A., and Khoshkhou, N. (2023). Partisan bias, economic expectations, and household spending. *Review of Economics and Statistics*, 105(3):493–510.

Mitra, A. and Mukherji, R. (2026). Class, caste and conspicuous consumption in India. *Economica*, 93(370):439–467.

Nakamura, E. and Steinsson, J. (2014). Fiscal stimulus in a monetary union: Evidence from US regions. *American Economic Review*, 104(3):753–792.

Pais, J. and Rawal, V. (2021). CMIE's Consumer Pyramids Household Surveys: An assessment. The India Forum.

Roth, C. and Wohlfart, J. (2020). How do expectations about the macroeconomy affect personal expectations and behavior? *Review of Economics and Statistics*, 102(4):731–748.

Somanathan, E., Somanathan, R., Sudarshan, A., and Tewari, M. (2021). The impact of temperature on productivity and labor supply: Evidence from Indian manufacturing. *Journal of Political Economy*, 129(6):1797–1827.

Somanchi, A. (2021). Missing the poor, big time: A critical assessment of the Consumer Pyramids Household Survey. SocArXiv preprint, Center for Open Science.

Taubinsky, D., Butera, L., Saccarola, M., and Lian, C. (2026). Beliefs about the economy are excessively sensitive to household-level shocks: Evidence from linked survey and administrative data. *Quarterly Journal of Economics*. Advance article, qjag032.

Appendix

A Additional Data and Empirical Evidence

This appendix reports the measurement details and supporting specifications for the data and empirical results in the order in which they arise in the main text.

A.1 Observation in the succeeding survey wave

Table A.1 examines whether extreme heat predicts observation in the immediately succeeding survey wave. In these data, expenditure is observed whenever the household completes that next-wave interview, so the re-interview and expenditure-observation indicators coincide. Neither common nor assigned-cell heat predicts subsequent observation.

Table A.1: Extreme Heat and Next-Wave Observation

	Observed in next consecutive wave
Leave-one-cell-out common exposure	0.00464
	(0.00668)
Assigned-cell exposure	0.00314
	(0.00276)
Joint change: common + assigned cell	0.00778
	(0.00626)
p-value: common = assigned cell	0.852
Household fixed effects	Yes
Month-of-year fixed effects	Yes
Observations	783,766
Geographic-cell clusters	1,987

Notes: The dependent variable equals one when the household is observed in the immediately succeeding CPHS wave. Regressions use analytic survey weights, absorb separate household and month-of-year fixed effects, and cluster standard errors by geographic cell.

A.2 Internal structure of the sentiment index

The primary sentiment index combines two assessments of family finances and two assessments of national economic conditions. Table A.2 reports their weighted response distributions and internal-consistency diagnostics. The four-item standardized Cronbach's alpha is 0.676. This statistic summarizes pooled internal consistency; it is not a reliability ratio for the within-household sentiment innovations used in the fixed-effects regressions. The two personal-finance items have a weighted correlation of 0.611, while the two national-outlook items have a weighted correlation of 0.555. Cross-domain correlations range from 0.184 to 0.197. The index therefore combines related but distinct personal and aggregate dimensions of households' economic sentiment.

Table A.2: Distribution and Internal Structure of the Consumer-Sentiment Index

	Weighted response share			Item–rest correlation	Alpha if omitted
Component	Adverse	Neutral	Favourable		
Family finances compared with previous year	0.195	0.550	0.255	0.468	0.602
Expected family finances one year ahead	0.187	0.577	0.236	0.474	0.598
National conditions over next 12 months	0.208	0.553	0.240	0.441	0.620
National conditions over next five years	0.173	0.618	0.208	0.444	0.618
Four-item standardized alpha				0.676	

Notes: Response shares use the exact common-outcome sample (N = 624,652) and analytic survey weights. "Adverse," "neutral," and "favourable" correspond to the three ordered responses to each question. Item–rest correlations and Cronbach's alpha calculations use the full available sample of 1,152,348 observations and are unweighted. The primary index is constructed using the inverse-covariance weighting procedure described in Section 2.2.

A.3 O*NET occupational exposure classification

Table A.3 reports the occupational exposure scores used to classify the household head's occupation. The upper panel defines the high-physical-intensity group.

Table A.3: O*NET Occupational Exposure Scores

Nature of occupation	Heat exposure		Outdoor exposure	
	Mean	Median	Mean	Median
High-physical-intensity occupations				
Agricultural labourer	3.83	3.85	4.41	4.45
Small farmer	3.73	3.81	3.90	4.28
Organised farmer	3.67	3.85	3.60	4.37
Industrial worker	2.92	3.13	2.44	2.12
Wage labourer	2.88	2.86	2.84	2.21
Support staff	2.70	2.76	2.70	2.38
Small trader or hawker	2.53	2.44	2.00	1.62
Low-skilled entrepreneur	2.48	2.44	2.35	1.95
Other occupations				
Student	2.35	2.34	2.33	2.12
Non-industrial technical employee	2.34	2.00	2.30	1.84
Homemaker	2.24	1.75	2.35	2.05
Manager	2.07	2.17	2.45	2.71
Home-based worker	2.03	1.73	1.68	1.63
Businessperson	2.00	2.00	2.66	2.66
White-collar professional employee	1.95	1.67	2.19	1.83
White-collar clerical employee	1.70	1.39	1.87	1.63
High-skilled entrepreneur	1.69	1.59	1.83	1.69
Qualified self-employed professional	1.64	1.52	1.78	1.65
Politician, social worker, or activist	1.44	1.44	1.83	1.83

Notes: Entries are occupation-level scores constructed from the O*NET measures "Exposed to High Temperatures" and "Outdoors, Exposed to All Weather," matched to the CPHS occupational categories through the National Classification of Occupations crosswalk. Scores range from one to five. The upper panel defines the high-physical-intensity group used in the heterogeneity analysis.

A.4 Temporal windows and interview timing

As interview timestamps are unavailable, we examine complete-grid heat before and after the interview. Panel A of Table A.4 enters nonoverlapping pre-interview windows jointly. Heat during days 8–30 before the interview is negatively associated with both sentiment and purchase intentions. Panel B shows that future local temperature windows also predict current outcomes. They demonstrate persistence in the local temperature regime and qualify an interpretation based on an isolated interview-day event. Within active cell-month risk sets, the probability of any successful interview is weakly higher, rather than lower, on hotter assigned-cell and local days: the coefficients are 0.00115 (SE 0.00061, $p = 0.060$) and 0.00138 (SE 0.00069, $p = 0.046$). Successful-interview counts are unrelated to assigned or local heat. Failed contacts and rescheduling are not observed, so these diagnostics do not establish random timing.

A.5 Alternative within-year specification

Table A.5 holds the exact headline sample fixed and compares the baseline to specifications with recorded survey-year and actual interview-year effects. The two year measures produce virtually identical estimates; only three observations in the 377,980-observation headline sample before singleton removal have different recorded and actual years. Adding year effects attenuates the 40–$45°\mathrm{C}$ coefficient substantially in both the six-bin profile and the single-bin first stage. Joint profile tests can remain significant because lower-temperature bins contribute, but neither the extreme-heat coefficient nor its upper-tail contrasts remains significant at the 5 percent level. Dropping month-of-year effects does not restore the headline result.

Year-effect specifications answer a narrower question than the baseline: they partial out common annual movements and identify from relative heat variation within each year. The baseline instead retains spatially coherent interannual heat variation because that variation forms part of the common exposure of interest. The attenuation therefore establishes that the headline coefficient is not identified exclusively from within-year relative deviations; it cannot distinguish removal of aggregate confounding from removal of economically relevant common heat variation. We report this alternative estimand transparently but do not use it as the baseline.

Table A.4: Temporal Alignment of Local Heat, Sentiment, and Purchase Plans

	Consumer sentiment	Purchase intentions
Panel A. Exact 5×5 *pre-interview windows*		
Interview-date local anomaly	-0.01017^{**}	-0.00961
	(0.00470)	(0.00986)
Previous-day local anomaly	0.00818^{*}	0.01619
	(0.00465)	(0.01604)
Mean local anomaly, days 2–7 before	-0.00646	-0.01381
	(0.00525)	(0.01378)
Mean local anomaly, days 8–30 before	-0.01725^{**}	-0.03145^{**}
	(0.00754)	(0.01334)
Panel B. Exact 5×5 *future-weather diagnostics*		
Interview-date local anomaly	-0.00552	-0.00049
	(0.00409)	(0.00647)
Mean local anomaly, days 1–30 after	-0.01409^{*}	-0.03435^{***}
	(0.00842)	(0.01027)
Mean local anomaly, days 31–90 after	-0.01554^{*}	0.00233
	(0.00806)	(0.00884)
Joint p-value: future windows	0.004	0.002
Observations	856,898	856,898
District clusters	406	406

Notes: Local exposure is the complete-grid 5×5 mean evaluated relative to the preceding-year cell-month benchmark. Panel A enters the four pre-interview terms jointly. Panel B is a separate regression containing the interview-date term and two future windows. All specifications absorb separate household and recorded month-of-year fixed effects, use analytic survey weights, and cluster by district. Future exposure is interpreted as a persistence diagnostic, not a treatment or conventional placebo. Standard errors are in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Table A.5: Extreme Heat and Sentiment: Sensitivity to Survey-Year Effects

	No year FE	Recorded-year FE	Actual-year FE	Recorded-year FE, no month FE
Six-bin, no expanded controls	−0.01820	−0.00444	−0.00444	−0.00566
	(0.00342)	(0.00345)	(0.00345)	(0.00340)
Expanded plotted six-bin profile	−0.01713	−0.00448	−0.00448	−0.00566
	(0.00345)	(0.00344)	(0.00344)	(0.00339)
Expanded single-bin first stage	−0.01658	−0.00076	−0.00076	−0.00375
	(0.00308)	(0.00303)	(0.00303)	(0.00297)
Household fixed effects	Yes	Yes	Yes	Yes
Month-of-year fixed effects	Yes	Yes	Yes	No
Observations	360,995	360,995	360,995	360,995
Geographic-cell clusters	1,757	1,757	1,757	1,757

Notes: Each entry is the coefficient on the 40–$45°\mathrm{C}$ anomaly. The first two rows include six temperature-bin anomalies jointly; the plotted profile also includes the paper's expanded controls. The third row uses the expanded single-bin first-stage specification. Recorded-year effects use the CPHS survey-cycle year; actual-year effects are decoded from the interview date. Standard errors are clustered by geographic cell. The corresponding p-values in columns (1)–(4) are below 0.001, 0.193, 0.193, and 0.095 for the plotted profile and below 0.001, 0.801, 0.801, and 0.207 for the single-bin specification.

A.6 Complete temperature distribution

The upper-tail result does not depend on the pooled omitted category in the baseline profile. Table A.6 enters the complete temperature distribution relative to the 25–$30°\mathrm{C}$ bin. The 40–$45°\mathrm{C}$ coefficient remains negative and significant for sentiment, purchase intentions, and next-wave expenditure. Actual exposure above $45°\mathrm{C}$ occurs in only 0.074 percent of weighted interviews in the profile sample, leaving that coefficient imprecisely estimated.

Table A.6: Complete Temperature Distribution, Sentiment, and Demand

	Expanded-sample sentiment	Common-sample sentiment	Common-sample intentions	Common-sample next-wave expenditure
	(1)	(2)	(3)	(4)
Below $0°$C	0.00992^{**}	0.00756	-0.00291	-0.00552^{**}
	(0.00493)	(0.00642)	(0.00451)	(0.00258)
0–$5°$C	-0.00451	-0.00659^{*}	-0.00749	-0.00613
	(0.00351)	(0.00347)	(0.00624)	(0.00464)
5–$10°$C	-0.00264	-0.00137	-0.00646^{*}	-0.00091
	(0.00291)	(0.00296)	(0.00341)	(0.00176)
10–$15°$C	-0.00132	-0.00526^{***}	-0.00231	-0.00104
	(0.00194)	(0.00194)	(0.00542)	(0.00121)
15–$20°$C	-0.00571^{***}	-0.00968^{***}	-0.00727	-0.00468^{***}
	(0.00182)	(0.00170)	(0.00486)	(0.00106)
20–$25°$C	-0.00530^{***}	-0.00663^{***}	0.00251	-0.00457^{***}
	(0.00188)	(0.00164)	(0.00233)	(0.00067)
30–$35°$C	0.00167	-0.00067	-0.00188	-0.00126^{*}
	(0.00129)	(0.00116)	(0.00369)	(0.00073)
35–$40°$C	0.00047	-0.00076	-0.01208^{***}	-0.00299^{***}
	(0.00208)	(0.00171)	(0.00279)	(0.00111)
40–$45°$C	-0.01897^{***}	-0.01383^{***}	-0.00688^{**}	-0.00872^{***}
	(0.00314)	(0.00247)	(0.00326)	(0.00156)
Above $45°$C	0.03840	0.01650	0.02451	-0.00777
	(0.03522)	(0.01727)	(0.02096)	(0.01277)
Expanded controls	Yes	No	No	No
Complete expenditure-date bin controls	No	No	No	Yes
Household fixed effects	Yes	Yes	Yes	Yes
Month-of-year fixed effects	Yes	Yes	Yes	Yes
Observations	360,995	624,652	624,652	624,652
Geographic-cell clusters	1,757	1,819	1,819	1,819

Notes: Every column includes the complete interview-date temperature distribution, with the 25–$30°$C bin omitted. Column (1) uses the expanded controls defined in the text. Column (4) additionally includes the complete temperature-bin vector at the expenditure date, omitting the same reference category. Actual exposure above $45°$C occurs in 0.074 percent of weighted observations in the expanded profile sample. All specifications use analytic survey weights and cell-clustered standard errors. ***, **, and * denote significance at the 1, 5, and 10 percent levels, respectively.

A.7 Predictive content of sentiment and purchase plans

Table A.7 reports the predictive specifications described in Section 3.2 on a common sample.

Table A.7: The Predictive Content of Sentiment and Purchase Plans

	Dependent variable: next-wave IHS total expenditure			
	(1)	(2)	(3)	(4)
Consumer sentiment	0.04966***	—	0.04976***	0.03689***
	(0.00346)		(0.00346)	(0.00318)
Purchase intentions	—	0.00635***	0.00660***	0.00397**
		(0.00242)	(0.00221)	(0.00170)
40–45°C shock	−0.00699***	−0.00752***	−0.00693***	−0.00706***
	(0.00153)	(0.00157)	(0.00153)	(0.00149)
Heat at expenditure date	Yes	Yes	Yes	Yes
Current income and expenditure	No	No	No	Yes
Household fixed effects	Yes	Yes	Yes	Yes
Month-of-year fixed effects	Yes	Yes	Yes	Yes
Observations	600,237	600,237	600,237	600,237
Geographic-cell clusters	1,806	1,806	1,806	1,806

Notes: Each column reports a separate weighted regression. The sentiment index excludes the purchase-timing question. Purchase intentions are adjusted for ownership across all asset categories. All specifications condition on the heat shock at the interview and heat at the expenditure date. Column (4) additionally controls for current IHS income and expenditure. All specifications use analytic survey weights and cell-clustered standard errors. The coefficients are conditional predictive associations. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

A.8 Alternative sentiment index

The primary sentiment index excludes the purchase-timing question. Table A.8 compares it with the broader index that includes that item.

Table A.8: Extreme Heat and Alternative Sentiment Indices

	Primary index (1)	Broader index (2)
40–45°C shock	-0.00898^{***}	-0.00737^{***}
	(0.00218)	(0.00217)
Household fixed effects	Yes	Yes
Month-of-year fixed effects	Yes	Yes
Observations	907,403	907,403
Geographic-cell clusters	2,008	2,008

Notes: The primary index excludes the "good time to buy" item; the broader index includes it. Before singleton removal, their correlation weighted using the analytic survey weights is 0.9735 ($N = 910{,}649$). Both regressions use analytic survey weights and cell-clustered standard errors, reported in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels, respectively.

A.9 Additional temperature-profile estimates

Table A.9 reports the estimates underlying Figure 4.1.

Table A.9: Consumer Sentiment Across the Temperature Distribution

	Full sample		Physical intensity	
	Baseline	Expanded	Low	High
	(1)	(2)	(3)	(4)
15–20°C shock	-0.00519^{***}	-0.00418^{**}	-0.00663^{***}	-0.00425^{**}
	(0.00176)	(0.00175)	(0.00227)	(0.00187)
20–25°C shock	-0.00496^{**}	-0.00378^{*}	-0.00526^{**}	-0.00369^{*}
	(0.00209)	(0.00209)	(0.00263)	(0.00217)
25–30°C shock	0.00089	0.00149	-0.00058	0.00151
	(0.00187)	(0.00188)	(0.00262)	(0.00206)
30–35°C shock	0.00247	0.00317	0.00022	0.00349
	(0.00216)	(0.00216)	(0.00307)	(0.00235)
35–40°C shock	0.00166	0.00198	0.00101	0.00135
	(0.00255)	(0.00255)	(0.00301)	(0.00280)
40–45°C shock	-0.01820^{***}	-0.01713^{***}	-0.01000^{**}	-0.01847^{***}
	(0.00342)	(0.00345)	(0.00420)	(0.00383)
Expanded controls	No	Yes	Yes	Yes
Household fixed effects	Yes	Yes	Yes	Yes
Month-of-year fixed effects	Yes	Yes	Yes	Yes
Observations	360,995	360,995	77,606	272,563
Geographic-cell clusters	1,757	1,757	923	1,738

Notes: All six shocks enter each regression jointly. Because the anomalies across the exhaustive temperature bins sum to zero, exposure below 15°C and above 45°C forms the pooled reference category. The expanded controls are next-wave income, current income and its square, time spent working for an employer, travelling, and in indoor entertainment, and, in column (2), occupational physical intensity. All regressions use analytic survey weights and cell-clustered standard errors, reported in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels, respectively.

A.10 Occupational heterogeneity in the heat response

Table A.10 reports the interaction estimates for the occupational gradient in the response to extreme heat.

Table A.10: Extreme Heat, Occupational Exposure, and Consumer Sentiment

	Dependent variable: consumer sentiment	
	(1)	(2)
40–45°C effect: low physical intensity	-0.00807^{**}	-0.00856^{**}
	(0.00393)	(0.00391)
40–45°C effect: high physical intensity	-0.01867^{***}	-0.01822^{***}
	(0.00335)	(0.00340)
High minus low	-0.01060^{**}	-0.00966^{**}
	(0.00439)	(0.00441)
Income and time-use controls	No	Yes
Household fixed effects	Yes	Yes
Month-of-year fixed effects	Yes	Yes
Observations	360,995	360,995
Geographic-cell clusters	1,757	1,757

Notes: The table reports marginal effects from regressions that interact the 40–45°C shock with an indicator for a physically intensive occupation. Higher values of the dependent variable denote more favourable sentiment. Column (2) additionally controls for next-wave income, a quadratic in current income, and time spent working for an employer, travelling, and in indoor entertainment. Both columns use the same sample, analytic survey weights, and standard errors clustered by geographic cell. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

A.11 Expanded controls and alternative intended-demand measures

Table A.11 reports the heat reduced forms with the expanded conditioning set and alternative intended-demand outcomes.

Table A.11: Expanded-Control Reduced Forms

	Consumer sentiment (1)	Good time to buy (2)	All assets, ownership adjusted (3)	Durables/vehicles, ownership adjusted (4)	Next-wave expenditure (5)
40–45°C shock	-0.01658^{***}	-0.01027^{**}	-0.00883^{**}	-0.00888^{**}	-0.01126^{***}
	(0.00308)	(0.00432)	(0.00388)	(0.00388)	(0.00158)
Heat at expenditure date	—	—	—	—	0.00363^{**}
					(0.00155)
Household fixed effects	Yes	Yes	Yes	Yes	Yes
Month-of-year fixed effects	Yes	Yes	Yes	Yes	Yes
Observations	360,995	360,995	360,995	360,995	360,995
Geographic-cell clusters	1,757	1,757	1,757	1,757	1,757

Notes: The expanded controls are next-wave income, current income and its square, time spent working for an employer, travelling, and in indoor entertainment, and occupational physical intensity. Column (2) uses the standardized general good-time-to-buy response; column (3) uses the all-asset ownership-adjusted index; column (4) uses the ownership-adjusted index for durables and vehicles. All regressions use analytic survey weights and cell-clustered standard errors, reported in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels, respectively.

A.12 Heat-induced sentiment and purchase plans: IV specification

The sharp upper-tail break in Figure 4.1 motivates using the 40–45°C temperature shock as an instrument for the primary sentiment index in the purchase-intention equation. The first and second stages are

$$S_{it} = \pi h_{it} + \mathbf{W}_{it}'\mathbf{\Gamma}_S + \alpha_i + \mu_{m(t)} + u_{it}, \tag{A.1}$$

$$I_{it} = \theta \widehat{S}_{it} + \mathbf{W}_{it}'\mathbf{\Gamma}_I + \alpha_i + \mu_{m(t)} + \eta_{it}, \tag{A.2}$$

where the baseline specification omits $\mathbf{W}_{it}$, leaving next-wave income, contemporaneous time use, and other possible responses to heat outside the conditioning set. Expanded-control estimates provide a sensitivity check. The 2SLS coefficient is the ratio of the purchase-intention reduced form to the sentiment first stage. Its interpretation as the effect of sentiment requires the heat instrument to be unrelated to η_{it} after the fixed effects and included controls are removed. In economic terms, heat must affect purchase plans through the measured outlook rather than through an additional route such as a change in purchase costs or access. Section 4.2 discusses this exclusion restriction. Both stages use the baseline weights and fixed effects. We report cell-clustered estimates, district-clustered sensitivity, the Kleibergen–Paap first-stage statistic, and Anderson–Rubin weak-instrument-robust inference. Table 4.2 in

Section 4 reports the headline estimates. Alternative purchase-intention measures and a common-exposure instrument produce directionally similar estimates and are reported as diagnostics in Appendix Tables A.12 and A.13.

A.13 Instrumental variables with alternative intended-demand measures

Table A.12 applies the expanded IV specification to the two alternative intended-demand outcomes. The corresponding estimate for the primary all-asset ownership-adjusted index appears in Table 4.2.

Table A.12: Instrumental-Variables Estimates for Alternative Intended Demand

	Good time to buy (1)	Durables/vehicles, ownership adjusted (2)
Sentiment index (2SLS)	0.61978***	0.53557**
	(0.22360)	(0.25556)
First stage: dependent variable is sentiment		
40–45°C shock	−0.01658***	−0.01658***
	(0.00308)	(0.00308)
Reduced form: dependent variable is intended demand		
40–45°C shock	−0.01027**	−0.00888**
	(0.00432)	(0.00388)
Kleibergen–Paap F-statistic	28.94	28.94
Anderson–Rubin p-value	0.0175	0.0224
Household fixed effects	Yes	Yes
Month-of-year fixed effects	Yes	Yes
Observations	360,995	360,995
Geographic-cell clusters	1,757	1,757

Notes: The excluded instrument is the 40–45°C shock. Both columns include next-wave income, current income and its square, the three time-use measures, and occupational physical intensity. Regressions use analytic survey weights and cell-clustered standard errors. The Anderson–Rubin p-values test a zero coefficient on sentiment. A structural interpretation requires the exclusion restriction stated in Appendix A.12; the coefficients are reported as conditional scale calibrations. ***, **, and * denote significance at the 1, 5, and 10 percent levels, respectively.

Table A.13: Common Exposure as an Instrument: Interview-Condition Diagnostic

	Assigned-cell temperature controls	
	40–45°C exposure	All six temperature bins
	(1)	(2)
Sentiment index (2SLS)	0.62292**	0.40977
	(0.31254)	(0.28441)
First stage: dependent variable is sentiment		
Leave-one-cell-out common exposure	−0.03165***	−0.03369***
	(0.00572)	(0.00571)
Reduced form: dependent variable is purchase intentions		
Leave-one-cell-out common exposure	−0.01972**	−0.01381
	(0.00956)	(0.00959)
Kleibergen–Paap F-statistic	30.60	34.83
Anderson–Rubin p-value	0.0393	0.1502
Expanded controls	Yes	Yes
Household fixed effects	Yes	Yes
Month-of-year fixed effects	Yes	Yes
Observations	338,529	338,529
Geographic-cell clusters	1,730	1,730

Notes: The dependent variable in the second stage is the all-asset ownership-adjusted purchase-intention index. The excluded instrument is equal-cell-weighted leave-one-cell-out common exposure. Column (1) includes assigned-cell exposure in the 40–45°C bin. Column (2) includes assigned-cell exposure in the 15–20, 20–25, 25–30, 30–35, 35–40, and 40–45°C bins jointly. Both columns include next-wave income, a quadratic in current income, the three time-use measures, and occupational physical intensity. Specifications absorb separate household and month-of-year fixed effects, use analytic survey weights, and cluster standard errors by geographic cell. The exercise assesses an explanation based on conditions specific to the respondent's interview location; common exposure may also influence purchase plans through shared local conditions. ***, **, and * denote significance at the 1, 5, and 10 percent levels, respectively.

B Model Derivation

This appendix proves the local-equilibrium result used in the main text. Household mass in each local market is normalized to one, and idiosyncratic disturbances and forecast errors integrate to zero.

B.1 Equilibrium and the response to uniform exposure

Substituting equations (5.5) and (5.6) into equation (5.7) and using $\lambda = \phi b$ gives

$$c_{i,t+1} = \lambda \left[f_{it} - \kappa_h h_{it} - \kappa_x x_{dt} + q\mathbb{E}_{it} C_{d,t+1} \right] - \delta h_{it} - \zeta x_{dt} + \widetilde{\varepsilon}^c_{i,t+1}, \tag{B.1}$$

where $\widetilde{\varepsilon}^c_{i,t+1} = \phi \varepsilon^I_{it} + \varepsilon^c_{i,t+1}$. Define

$$f_{dt} = \int_0^1 f_{it}\, di, \qquad C_{d,t+1} = \int_0^1 c_{i,t+1}\, di, \qquad x_{dt} = \int_0^1 h_{it}\, di. \tag{B.2}$$

Aggregate consistency of household forecasts implies $\int_0^1 \mathbb{E}_{it} C_{d,t+1}\, di = C_{d,t+1}$. Averaging equation (B.1) therefore yields

$$(1 - q\lambda) C_{d,t+1} = \lambda f_{dt} - \{\delta + \zeta + \lambda(\kappa_h + \kappa_x)\} x_{dt}. \tag{B.3}$$

Under $0 \leq q\lambda < 1$, define

$$\mathcal{M} = \frac{1}{1 - q\lambda}. \tag{B.4}$$

The unique stable equilibrium is

$$C_{d,t+1} = \mathcal{M} \left[\lambda f_{dt} - \{\delta + \zeta + \lambda(\kappa_h + \kappa_x)\} x_{dt} \right]. \tag{B.5}$$

Let υ index a uniform increase in exposure, with $dh_{it}/d\upsilon = dx_{dt}/d\upsilon = 1$ for every household. Differentiating equation (B.5) gives

$$\frac{dC_{d,t+1}}{d\upsilon} = -\mathcal{M}\{\delta + \zeta + \lambda(\kappa_h + \kappa_x)\}. \tag{B.6}$$

Average sentiment satisfies

$$s_{dt} = f_{dt} - (\kappa_h + \kappa_x) x_{dt} + q C_{d,t+1}. \tag{B.7}$$

Using equation (B.6) and $1 + q\lambda\mathcal{M} = \mathcal{M}$,

$$\begin{aligned}\frac{ds_{dt}}{d\upsilon} &= -(\kappa_h + \kappa_x) + q\frac{dC_{d,t+1}}{d\upsilon} \\ &= -\mathcal{M}\{\kappa_h + \kappa_x + q(\delta + \zeta)\}.\end{aligned} \tag{B.8}$$

Since $I_{dt} = bs_{dt}$ after aggregation of the mean-zero intention disturbance,

$$\frac{dI_{dt}}{d\upsilon} = -b\mathcal{M}\{\kappa_h + \kappa_x + q(\delta + \zeta)\}. \tag{B.9}$$

B.2 Pass-through and direct disruption

Under homogeneous response slopes, alignment of empirical and model weights, timing, and proxy scales, the negligible-location continuum approximation, mean-independent measurement and structural disturbances, and the exclusion restriction that heat affects desired demand through sentiment rather than through an additional direct channel, the population coefficients imply, for $\beta_{Sx} \neq 0$,

$$b = \frac{\beta_{Ix}}{\beta_{Sx}}, \qquad \kappa_h = -\beta_{Sh}, \qquad d_h \equiv \lambda\kappa_h + \delta = -\beta_{Ch}. \tag{B.10}$$

The common expenditure equation can be written as

$$\zeta = \lambda\beta_{Sx} - \beta_{Cx}. \tag{B.11}$$

If common exposure lowers sentiment and expenditure and $\zeta \geq 0$, then

$$0 < \lambda \leq \frac{\beta_{Cx}}{\beta_{Sx}}. \tag{B.12}$$

If $\kappa_h > 0$ and $\delta \geq 0$, the own-cell equation supplies the additional restriction

$$\lambda \leq \frac{-\beta_{Ch}}{-\beta_{Sh}}. \tag{B.13}$$

These cross-equation mappings require coefficients defined on a common target population with compatible conditioning sets. The displayed ratio $(-0.01359)/(-0.02316) \approx 0.59$ uses the same observations and weights, but the expenditure regression in Table 6.1 controls for assigned heat at the succeeding interview while the sentiment regression does not.

Matching the sample alone therefore does not establish that this ratio satisfies equation (5.18). An empirical application requires an aligned sentiment regression or a justified invariance condition showing that adding the succeeding-interview control leaves its common-exposure slope unchanged. We report 0.59 as an illustration of the coefficient ratio, not as an established numerical upper bound. Even with alignment, uncertainty about a ratio requires the joint cross-equation covariance in addition to the model assumptions above. The specification with both common and assigned exposure at the succeeding interview uses different outcome-specific samples and is likewise not combined into a structural ratio.

B.3 Equilibrium feedback

Let

$$A = \delta + \zeta + \lambda(\kappa_h + \kappa_x). \tag{B.14}$$

The uniform expenditure response satisfies

$$\beta_C^{\text{total}} = -\mathcal{M}A. \tag{B.15}$$

Within the model, the reduced form determines this equilibrium product but does not separately determine the initial impulse and the feedback multiplier. For any nonpositive total response $\beta_C^{\text{total}} \leq 0$ and any $\gamma = q\lambda \in [0, 1)$, the pair

$$\mathcal{M}(\gamma) = \frac{1}{1-\gamma}, \qquad A(\gamma) = -(1-\gamma)\beta_C^{\text{total}} \geq 0, \tag{B.16}$$

reproduces the same response. The quantitative analysis therefore uses the equilibrium reduced-form effect and does not assign a numerical value to the feedback multiplier.

B.4 Derivation of the common–assigned decomposition

Consider first a change in common exposure that leaves the focal assigned-location unit's exposure fixed. Under the continuum approximation, that unit has negligible mass in the local market, so the aggregate response remains equation (B.6). Equation (5.5) then implies

$$\begin{aligned} \beta_{Sx} &= -\kappa_x + q\frac{\partial C_{d,t+1}}{\partial x_{dt}} \\ &= -\mathcal{M}\{\kappa_x + q(\delta + \zeta + \lambda\kappa_h)\}. \end{aligned} \tag{B.17}$$

Holding common exposure fixed, a change in the focal assigned-location unit's exposure has no aggregate effect under the same approximation, so

$$\beta_{Sh} = -\kappa_h. \tag{B.18}$$

The desired demand coefficients follow directly from equation (5.6):

$$\beta_{Ix} = b\beta_{Sx}, \qquad \beta_{Ih} = b\beta_{Sh}. \tag{B.19}$$

Finally, equation (5.7) gives

$$\beta_{Cx} = \lambda\beta_{Sx} - \zeta, \qquad \beta_{Ch} = \lambda\beta_{Sh} - \delta. \tag{B.20}$$

A normalized uniform local increase raises common and own exposure one-for-one. Adding the two sentiment coefficients gives

$$\beta_{Sx} + \beta_{Sh} = -\mathcal{M}\{\kappa_h + \kappa_x + q(\delta + \zeta)\}, \tag{B.21}$$

and adding the two expenditure coefficients gives

$$\beta_{Cx} + \beta_{Ch} = -\mathcal{M}\{\delta + \zeta + \lambda(\kappa_h + \kappa_x)\}. \tag{B.22}$$

These expressions coincide with equations (B.8) and (B.6), respectively.

C Spatial Exposure: Construction, Inference, and Evidence

C.1 Leave-one-cell-out exposure

Let c index ERA5 cells, d districts, y recorded CPHS survey-cycle years, and m recorded CPHS survey months. Households in the same cell can receive different shocks because interviews occur on different dates. The first step therefore forms the unweighted mean household-assigned shock within each cell–district–year–month:

$$\bar{h}_{cdym} = \frac{1}{n_{cdym}} \sum_{i \in (c,d,y,m)} h_{iym}. \tag{C.1}$$

Let N_{dym} denote the total number of CPHS-observed cells in district d in recorded survey-cycle year y and recorded survey month m, including cell c. For a household in cell c, survey-schedule common exposure is

$$\widetilde{x}_{-c,dym} = \frac{1}{N_{dym} - 1} \sum_{c' \neq c} \bar{h}_{c'dym}, \qquad N_{dym} \geq 2. \tag{C.2}$$

The measure is missing when the district–recorded-year–recorded-month group contains fewer than two observed cells. Interviews in different cells within a group can occur on different dates, so this object is not a same-day physical district average. Each eligible cell receives equal weight, and the entire own cell is omitted before the average is assigned to its households. This construction prevents survey density and the own-cell realization from entering the common measure mechanically.

The model aggregate uses household weights. Equation (C.2), by contrast, uses equal cell weights. Mapping the separate empirical coefficients to model primitives requires alignment in weights, timing, proxy scale, and response slopes; weight alignment alone is not sufficient. For a uniform shift in the model, every normalized leave-one-out mean and every own-cell exposure rises by the same amount. In the empirical application, the common-plus-own coefficient is instead the fitted response to a joint one-unit change in the two constructed regressors.

C.2 Estimation sample and linear combinations

The common spatial sample requires the primary sentiment index, the all-asset ownership-adjusted purchase-intention index, next-wave total expenditure, own and common exposure, heat at the expenditure date, a positive survey weight, and the identifiers used for fixed effects and clustering. Next-wave outcomes are defined only for consecutive survey waves. The resulting sample contains 596,543 observations before singleton removal and 584,632 observations after absorption of the separate household and month-of-year fixed effects. The estimation sample contains 115,086 households and 1,796 geographic cells. No year fixed effects are included.

The matched-sample benchmark in Table 6.1 includes only the future assigned-cell term shown in equation (6.2) and is estimated with analytic survey weights and cell-clustered standard errors. The fully controlled equation, including future common and assigned exposure, is estimated on outcome-specific samples in Table C.1. For outcome $z \in \{S, I, C\}$, collect the common and own coefficients and their covariance matrix as

$$\widehat{\beta}_z = \begin{pmatrix} \widehat{\beta}_{zx} \\ \widehat{\beta}_{zh} \end{pmatrix}, \qquad \widehat{V}_z = \mathrm{Var}(\widehat{\beta}_z). \tag{C.3}$$

With $R_+ = (1, 1)$ and $R_- = (1, -1)$, the coefficient sum and the test of equal common and own coefficients are

$$\widehat{\beta}_z^{\text{total}} = R_+ \widehat{\beta}_z, \qquad \mathrm{se}(\widehat{\beta}_z^{\text{total}}) = \sqrt{R_+ \widehat{V}_z R_+'}, \tag{C.4}$$

$$H_0 : R_- \beta_z = 0. \tag{C.5}$$

The coefficient-sum confidence intervals consequently incorporate the estimated covariance between the two exposure coefficients within each outcome equation.

C.3 Outcome-specific spatial estimates with future common and assigned exposure

The matched-sample benchmark in Table 6.1 controls for assigned-cell heat at the succeeding interview. Table C.1 re-estimates the survey-schedule decomposition on outcome-specific samples and includes both common and assigned exposure at the succeeding interview in the expenditure equation. The initial common-plus-assigned response remains negative for all three outcomes under cell and district clustering. The separate common purchase-intention coefficient is less precise with district clustering, but its initial coefficient sum remains significant at the five-percent level.

Table C.1: Survey-Schedule Spatial Estimates with Future Common and Assigned Exposure

	Consumer sentiment	Purchase intentions	Next-wave expenditure
Panel A. Geographic-cell clusters			
Initial survey-schedule common exposure	-0.01726^{***}	-0.03526^{***}	-0.01281^{***}
	(0.00473)	(0.01297)	(0.00329)
Initial assigned-cell exposure	-0.00104	0.00532	-0.00223
	(0.00291)	(0.00497)	(0.00188)
Future common exposure	—	—	0.00809^{**}
			(0.00348)
Future assigned-cell exposure	—	—	0.00084
			(0.00233)
Initial common + assigned	-0.01830^{***}	-0.02993^{***}	-0.01505^{***}
	(0.00367)	(0.00912)	(0.00268)
Panel B. District clusters			
Initial survey-schedule common exposure	-0.01726^{***}	-0.03526^{*}	-0.01281^{***}
	(0.00583)	(0.02008)	(0.00458)
Initial assigned-cell exposure	-0.00104	0.00532	-0.00223
	(0.00232)	(0.00660)	(0.00183)
Future common exposure	—	—	0.00809^{*}
			(0.00432)
Future assigned-cell exposure	—	—	0.00084
			(0.00246)
Initial common + assigned	-0.01830^{***}	-0.02993^{**}	-0.01505^{***}
	(0.00585)	(0.01455)	(0.00433)
Observations	847,761	847,761	577,509
Geographic-cell clusters	1,989	1,989	1,791
District clusters	385	385	365

Notes: The common proxy is constructed from other CPHS-observed cells in the same district, recorded survey-cycle year, and recorded survey month. All regressions absorb separate household and recorded month-of-year fixed effects and use analytic survey weights. The expenditure specification includes both common and assigned exposure at the succeeding interview. The joint tests of the two future-exposure terms have $p = 0.0009$ with cell clustering and $p = 0.0707$ with district clustering. Standard errors are in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

C.4 Standardized spatial effects and purchase-intention sample stability

Table C.2 expresses the spatial coefficients in standard-deviation units. Panel A uses the common-outcome sample. Panel B compares the purchase-intention estimates across the principal sample restrictions.

Table C.2: Standardized Common and Assigned-Cell Effects

	Sentiment	Purchase intentions	Next-wave expenditure
Panel A. Common-outcome sample			
Common exposure	-0.02692^{***}	-0.01421^{**}	-0.02427^{***}
	(0.00570)	(0.00614)	(0.00543)
Assigned-cell exposure	-0.00263	-0.00084	-0.00450
	(0.00463)	(0.00418)	(0.00414)
Observations	584,632	584,632	584,632
Geographic-cell clusters	1,796	1,796	1,796

Panel B. Purchase-intention estimates across samples

	Purchase-intention effect		Observations	Clusters
Sample	Common	Assigned cell		
Common-outcome sample	-0.01421^{**}	-0.00084	584,632	1,796
	(0.00614)	(0.00418)		
Maximum intentions sample	-0.02937^{***}	0.00586	847,761	1,989
	(0.01081)	(0.00548)		
At least five observed cells	-0.02262^{***}	0.00528	288,177	1,367
	(0.00753)	(0.00521)		
At least ten observed cells	-0.11524^{***}	0.01321	48,372	391
	(0.04411)	(0.02433)		

Notes: Within each estimation sample, the dependent variable, common exposure, and assigned-cell exposure are standardized using means and standard deviations weighted by the analytic survey weights. Coefficients therefore report outcome standard deviations per standard deviation of the corresponding exposure measure. The standard deviations of common and assigned-cell exposure in Panel A are 0.831 and 1.114 hours. The expenditure equation controls for heat at the expenditure date. All specifications absorb separate household and month-of-year fixed effects, use analytic survey weights, and cluster standard errors by geographic cell. ***, **, and * denote significance at the 1, 5, and 10 percent levels, respectively.

C.5 Alternative spatial specifications

Table C.3 reports two changes to the spatial design. Panel A retains equal-cell common exposure and clusters by district. Panel B returns to cell-clustered inference and constructs common exposure using the survey weights of the other cells.

Table C.3: Alternative Common–Own Exposure Specifications

	Sentiment	Purchase intentions	Next-wave expenditure
	(1)	(2)	(3)
Panel A. Equal-cell exposure; district clustering			
Common exposure	-0.02316^{***}	-0.01824	-0.01359^{***}
	(0.00610)	(0.01113)	(0.00427)
Incremental own-cell exposure	-0.00169	-0.00081	-0.00188
	(0.00279)	(0.00427)	(0.00167)
p-value: common = own	0.0029	0.2160	0.0181
Coefficient sum	-0.02485^{***}	-0.01905^{**}	-0.01547^{***}
	(0.00622)	(0.00928)	(0.00420)
Observations	584,632	584,632	584,632
District clusters	369	369	369
Panel B. Survey-weighted exposure; cell clustering			
Common exposure	-0.02362^{***}	-0.01611^{**}	-0.01250^{***}
	(0.00481)	(0.00818)	(0.00293)
Incremental own-cell exposure	-0.00124	-0.00164	-0.00226
	(0.00292)	(0.00409)	(0.00171)
p-value: common = own	0.0012	0.2041	0.0119
Coefficient sum	-0.02486^{***}	-0.01776^{***}	-0.01476^{***}
	(0.00400)	(0.00613)	(0.00255)
Observations	584,632	584,632	584,632
Geographic-cell clusters	1,796	1,796	1,796

Notes: Each column reports a separate regression on common and own-cell exposure. The reported coefficient sum adds the two estimates. Expenditure regressions also control for heat at the expenditure date. All specifications absorb separate household and month-of-year fixed effects and use analytic survey weights. Standard errors are in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels, respectively.

C.6 Temporal alignment with fixed geographic coverage

We retain the original set of eligible CPHS-observed cells within each district, recorded survey-cycle year, and recorded survey month, and evaluate their heat anomalies on each focal household's exact interview date. The common measure remains an equal-weighted average across the other cells, excluding the household's assigned cell. This changes temporal alignment while preserving geographic membership and cell weights. Daily weather is available for every required cell-date, and the original and aligned specifications use the same 584,632 household-wave observations, fixed effects, analytic survey weights, and outcome-specific controls. Table C.4 reproduces the survey-schedule estimates and reports the aligned comparison. After household and month-of-year fixed effects, the correlation between common and own exposure rises from 0.632 to 0.944. The joint responses are smaller in magnitude but remain negative and statistically significant for all three outcomes under both clustering levels.

Table C.4: Common and Own Heat Exposure with Fixed Geography and Alternative Timing

	Sentiment	Purchase intentions	Next-wave expenditure
Panel A. Original survey-schedule common exposure			
Common exposure	-0.02316^{***}	-0.01824^{**}	-0.01359^{***}
Cell-clustered SE	(0.00490)	(0.00788)	(0.00304)
District-clustered SE	[0.00610]	[0.01113]	[0.00427]
Own exposure	-0.00169	-0.00081	-0.00188
Cell-clustered SE	(0.00297)	(0.00400)	(0.00173)
District-clustered SE	[0.00279]	[0.00427]	[0.00167]
Common + own	-0.02485^{***}	-0.01905^{***}	-0.01547^{***}
Cell-clustered SE	(0.00405)	(0.00602)	(0.00261)
District-clustered SE	[0.00622]	[0.00928]	[0.00420]
p: common = own, cell	0.0023	0.1118	0.0054
p: common = own, district	0.0029	0.2160	0.0181
Panel B. Same cells evaluated on the focal interview date			
Common exposure	0.00086	-0.00424	-0.00912^{**}
Cell-clustered SE	(0.00767)	(0.00949)	(0.00433)
District-clustered SE	[0.00873]	[0.01036]	[0.00449]
Own exposure	-0.01329^{*}	-0.00550	0.00001
Cell-clustered SE	(0.00728)	(0.00856)	(0.00402)
District-clustered SE	[0.00777]	[0.00900]	[0.00368]
Common + own	-0.01243^{***}	-0.00973^{***}	-0.00911^{***}
Cell-clustered SE	(0.00263)	(0.00345)	(0.00173)
District-clustered SE	[0.00384]	[0.00476]	[0.00256]
p: common = own, cell	0.3364	0.9434	0.2635
p: common = own, district	0.3793	0.9466	0.2425
Observations in each panel	584,632	584,632	584,632
Households	115,086	115,086	115,086
Geographic-cell clusters	1,796	1,796	1,796
District clusters	369	369	369
Heat at expenditure date	No	No	Yes
Household fixed effects	Yes	Yes	Yes
Month-of-year fixed effects	Yes	Yes	Yes

Notes: Both panels use the same observations and the same original eligible cell sets. Panel A first averages nonmissing interview-day heat anomalies within each cell and recorded district–year–month group, then averages equally across other eligible cells. Panel B evaluates those same other cells on the focal household's interview date, using the daily climate file independently of interview activity in those cells. Own exposure is unchanged. All regressions use analytic survey weights and separate household and recorded month-of-year fixed effects. Expenditure is the inverse hyperbolic sine of total expenditure in the next consecutive survey wave. Both expenditure specifications control for own-cell heat at that succeeding interview. Parentheses report geographic-cell-clustered standard errors; brackets report district-clustered standard errors. Coefficient sums and equality tests use the full estimated covariance matrix. Stars refer to cell-clustered inference: ***, **, and * indicate significance at the 1, 5, and 10 percent levels. For Panel B's joint responses, district-clustered p-values are 0.0013, 0.0418, and 0.0004, respectively.

C.7 Fixed-radius local exposure

The district definition is replaced with fixed-radius neighbourhoods centred on the household's coordinates. Common exposure is the equal-weighted mean shock among other CPHS-observed ERA5 cells within the stated radius and the same recorded survey-cycle year and survey month, including cells across district boundaries. Interviews can occur on different dates within those groups. The household's assigned-cell shock enters separately. Table C.5 shows that the common coefficients and coefficient sums are close to their district-based counterparts.

Table C.5: Extreme Heat in Fixed-Radius Local Areas

	Sentiment	Purchase intentions	Next-wave expenditure
Panel A. Other cells within 40 kilometres			
Fixed-radius common exposure	-0.02384^{***}	-0.02058^{**}	-0.01315^{***}
	(0.00506)	(0.00891)	(0.00322)
Assigned-cell exposure	-0.00075	0.00208	-0.00112
	(0.00301)	(0.00454)	(0.00175)
Coefficient sum	-0.02458^{***}	-0.01850^{***}	-0.01427^{***}
	(0.00419)	(0.00649)	(0.00276)
p-value: common = assigned cell	0.001	0.072	0.006
Observations	505,343	505,343	505,343
Geographic-cell clusters	1,671	1,671	1,671
Panel B. Other cells within 60 kilometres			
Fixed-radius common exposure	-0.02312^{***}	-0.02040^{**}	-0.01424^{***}
	(0.00526)	(0.00874)	(0.00336)
Assigned-cell exposure	-0.00262	0.00025	-0.00117
	(0.00307)	(0.00389)	(0.00173)
Coefficient sum	-0.02574^{***}	-0.02015^{***}	-0.01540^{***}
	(0.00430)	(0.00696)	(0.00285)
p-value: common = assigned cell	0.006	0.075	0.004
Observations	555,383	555,383	555,383
Geographic-cell clusters	1,775	1,775	1,775

Notes: Fixed-radius common exposure is the mean heat shock in other CPHS-observed ERA5 cells within the stated distance of the household and the same recorded survey-cycle year and survey month. Interview dates can differ within these groups. The measure requires at least two other observed cells and crosses district boundaries. Expenditure regressions control for heat at the expenditure date. All specifications absorb separate household and month-of-year fixed effects, use analytic survey weights, and cluster standard errors by geographic cell. ***, **, and * denote significance at the 1, 5, and 10 percent levels, respectively.

C.8 Expanding the household's immediate thermal environment

For each radius, cells inside the buffer form the household's expanded immediate environment and the remaining observed cells in its district form the surrounding component. Table C.6 uses an identical sample across all three radii. Expanding the buffer shifts the allocation between the two coefficients, while their sum remains stable.

C.9 Exact-interview-date local exposure

The survey-schedule common measure pools heat anomalies observed on potentially different interview dates. We therefore construct a complementary measure from the complete ERA5 grid on each focal household's exact interview date. The surrounding regressor is the equal-weighted mean 40–$45^{\circ}\mathrm{C}$ anomaly in the 24 cells surrounding the assigned cell in a 5×5 grid. This construction is available independently of CPHS interview activity in neighbouring cells.

Table C.7 reports the surrounding and assigned-cell coefficients and their sum. The sums are negative and statistically precise for sentiment, purchase intentions, and next-wave expenditure under both cell and district clustering. Same-date temperatures move closely together across adjacent cells, making the separate sentiment and intention coefficients imprecise while their sums are well estimated. This exact-date 5×5 exercise therefore checks the joint local response. It differs from the broadened-personal-exposure check in Appendix C.11, which averages all 25 cells into one regressor and retains the district common proxy. The district and fixed-radius decompositions provide the principal common-versus-own comparison.

C.10 Nonoverlapping exact-date spatial partition

To examine the spatial allocation of the exact-date joint response using disjoint components, we partition the complete 5×5 grid into a nine-cell inner area and a nonoverlapping 16-cell outer ring. The separate coefficients are imprecise, but their sums are negative for all three outcomes under district clustering. This confirms the joint local response under an alternative disjoint spatial partition. The exercise measures geographic exposure rather than household-specific mobility patterns.

Table C.6: Immediate and Surrounding Exposure Across Buffer Radii

	25 km	40 km	60 km
Panel A. Consumer sentiment			
Exposure outside buffer	-0.01691^{*}	-0.00879	-0.00500
	(0.00876)	(0.00822)	(0.00814)
Exposure inside buffer	-0.01235	-0.02068^{**}	-0.02441^{***}
	(0.00869)	(0.00808)	(0.00811)
Coefficient sum	-0.02927^{***}	-0.02947^{***}	-0.02941^{***}
	(0.00656)	(0.00663)	(0.00692)
p-value: outside = inside	0.778	0.425	0.187
Panel B. Purchase intentions			
Exposure outside buffer	-0.05116^{***}	-0.03893^{***}	-0.01924^{*}
	(0.01820)	(0.01304)	(0.01059)
Exposure inside buffer	-0.00465	-0.01790	-0.03655^{***}
	(0.01253)	(0.01174)	(0.01350)
Coefficient sum	-0.05581^{***}	-0.05683^{***}	-0.05580^{***}
	(0.01258)	(0.01304)	(0.01302)
p-value: outside = inside	0.104	0.319	0.398
Panel C. Next-wave expenditure			
Exposure outside buffer	-0.00961^{**}	-0.00952^{**}	-0.00514
	(0.00463)	(0.00427)	(0.00439)
Exposure inside buffer	-0.00481	-0.00559	-0.01010^{**}
	(0.00428)	(0.00464)	(0.00476)
Coefficient sum	-0.01442^{***}	-0.01511^{***}	-0.01524^{***}
	(0.00299)	(0.00291)	(0.00303)
p-value: outside = inside	0.568	0.641	0.566
Observations	118,303	118,303	118,303
Geographic-cell clusters	999	999	999

Notes: Inside-buffer exposure is the equal-weighted mean shock among observed cells in the household's district within the stated radius. Outside-buffer exposure is the corresponding mean among the remaining observed cells in the district. The sample requires both measures at all three radii and therefore consists primarily of geographically larger districts. Expenditure regressions control for heat at the expenditure date. All specifications absorb separate household and month-of-year fixed effects, use analytic survey weights, and cluster standard errors by geographic cell. ***, **, and * denote significance at the 1, 5, and 10 percent levels, respectively.

Table C.7: Exact-Interview-Date Local Exposure

	Consumer sentiment	Purchase intentions	Next-wave expenditure
Panel A. Geographic-cell clusters			
Mean exposure in 24 surrounding cells	-0.00481	-0.00157	-0.01595^{***}
	(0.00844)	(0.01024)	(0.00548)
Assigned-cell exposure	-0.00484	-0.01052	0.00591
	(0.00761)	(0.00930)	(0.00476)
Surrounding + assigned	-0.00965^{***}	-0.01209^{***}	-0.01005^{***}
	(0.00247)	(0.00347)	(0.00184)
Panel B. District clusters			
Mean exposure in 24 surrounding cells	-0.00481	-0.00157	-0.01595^{**}
	(0.00921)	(0.01271)	(0.00693)
Assigned-cell exposure	-0.00484	-0.01052	0.00591
	(0.00777)	(0.01015)	(0.00540)
Surrounding + assigned	-0.00965^{***}	-0.01209^{**}	-0.01005^{***}
	(0.00361)	(0.00537)	(0.00281)
Observations	907,403	907,403	624,652
Geographic-cell clusters	2,008	2,008	1,819
District clusters	406	406	393

Notes: The surrounding measure is the equal-weighted anomaly in the 24 ERA5 cells around the focal cell, evaluated on the focal household's exact interview date. All regressions absorb separate household and recorded month-of-year fixed effects and use analytic survey weights. The expenditure specification includes surrounding and assigned exposure at the succeeding interview. The coefficient sum uses its full covariance. Standard errors are in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Table C.8: Nonoverlapping Inner and Outer Exact-Date Exposure

	Consumer sentiment	Purchase intentions	Next-wave expenditure
Inner 3×3 exposure	-0.00901	-0.00649	0.00147
	(0.01377)	(0.01521)	(0.00995)
Outer 16-cell exposure	-0.00063	-0.00619	-0.01137
	(0.01514)	(0.01842)	(0.01154)
Inner + outer	-0.00964^{**}	-0.01267^{**}	-0.00990^{***}
	(0.00372)	(0.00566)	(0.00296)
Observations	907,403	907,403	624,652
District clusters	406	406	393

Notes: Exposures are evaluated on the focal household's exact interview date. The inner and outer areas are nonoverlapping. All regressions absorb separate household and recorded month-of-year fixed effects, use analytic survey weights, and cluster by district. The expenditure equation includes both inner and outer exposure at the succeeding interview. Standard errors are in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

C.11 Broadened measures of personal thermal exposure

The baseline decomposition measures personal exposure using the household's assigned ERA5 cell. To allow experienced thermal conditions to extend beyond that cell, we construct two broader measures from the complete ERA5 grid. The first averages the shock across the assigned cell and its eight immediately surrounding cells; the second averages across the assigned cell and its 24 surrounding cells, evaluated on the focal interview date. These 3×3 and 5×5 measures replace the assigned-cell regressor, while the leave-one-cell-out district measure continues to capture common exposure. Table C.9 reports the resulting decomposition on the same sample used in Table 6.1. Because some cells in these grids may also enter the district proxy, this is a robustness check that controls for broader nearby temperature, not a disjoint spatial decomposition; Appendix C.10 provides the nonoverlapping exact-date partition.

Table C.9: Common Exposure and Broadened Personal Exposure

	Consumer sentiment	Purchase intentions	Next-wave expenditure
Panel A. Assigned cell and eight surrounding ERA5 cells			
Common district exposure	-0.02371^{***}	-0.01980^{**}	-0.01282^{***}
	(0.00497)	(0.00815)	(0.00311)
Broadened personal exposure	-0.00106	0.00103	-0.00284
	(0.00317)	(0.00454)	(0.00190)
p-value: common = personal	0.002	0.076	0.025
Joint change: common + personal	-0.02477^{***}	-0.01877^{***}	-0.01566^{***}
	(0.00406)	(0.00600)	(0.00262)
Panel B. Assigned cell and 24 surrounding ERA5 cells			
Common district exposure	-0.02402^{***}	-0.01992^{**}	-0.01262^{***}
	(0.00501)	(0.00804)	(0.00314)
Broadened personal exposure	-0.00071	0.00122	-0.00318
	(0.00333)	(0.00472)	(0.00204)
p-value: common = personal	0.002	0.072	0.040
Joint change: common + personal	-0.02474^{***}	-0.01870^{***}	-0.01580^{***}
	(0.00406)	(0.00602)	(0.00262)
Household fixed effects	Yes	Yes	Yes
Month-of-year fixed effects	Yes	Yes	Yes
Heat at expenditure date	No	No	Yes
Observations	584,632	584,632	584,632
Geographic-cell clusters	1,796	1,796	1,796

Notes: Broadened personal exposure is the equal-weighted mean extreme-heat shock in the household's assigned ERA5 cell and either its eight surrounding cells (Panel A) or its 24 surrounding cells (Panel B). Common exposure is the equal-cell-weighted leave-one-cell-out district measure. All specifications absorb separate household and month-of-year fixed effects, use analytic survey weights, and cluster standard errors by geographic cell. Column (3) uses IHS total expenditure and additionally controls for heat at the expenditure date. Standard errors are reported in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels, respectively.

C.12 Flexible controls for assigned-cell temperature exposure

Table C.10 augments the spatial specifications with the household's assigned-cell exposure in all six temperature bins. This allows the direct response to local thermal conditions to vary flexibly across the temperature distribution.

Table C.10: Spatial Exposure with Flexible Assigned-Cell Temperature Controls

	Consumer sentiment (1)	Purchase intentions (2)	Next-wave expenditure (3)
Leave-one-cell-out common exposure	-0.02478^{***}	-0.03155^{**}	-0.01338^{***}
	(0.00492)	(0.01319)	(0.00311)
Assigned-cell 40–45°C exposure	0.00160	0.00371	-0.00184
	(0.00349)	(0.00602)	(0.00208)
Heat at expenditure date	—	—	0.00451^{***}
			(0.00161)
Joint change: common + assigned cell	-0.02317^{***}	-0.02783^{***}	-0.01522^{***}
	(0.00444)	(0.00921)	(0.00292)
p-value: common = assigned cell	0.0003	0.0543	0.0088
All assigned-cell temperature-bin controls	Yes	Yes	Yes
Household fixed effects	Yes	Yes	Yes
Month-of-year fixed effects	Yes	Yes	Yes
Observations	584,632	847,761	584,632
Geographic-cell clusters	1,796	1,989	1,796

Notes: Common exposure is the equal-cell-weighted leave-one-cell-out 40–45°C heat shock among other cells in the district–year–month. Every specification includes assigned-cell heat shocks for the 15–20, 20–25, 25–30, 30–35, 35–40, and 40–45°C bins. Columns (1) and (3) use the common spatial sample. Column (2) uses the maximum sample available for ownership-adjusted purchase intentions and therefore does not require next-wave expenditure to be observed. Column (3) additionally controls for heat at the expenditure date. All specifications use analytic survey weights, absorb separate household and month-of-year fixed effects, and cluster standard errors by geographic cell. Standard errors are reported in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels, respectively.

C.13 Urban nonagricultural households

Table C.11 restricts the spatial analysis to urban households whose reported occupation is nonagricultural. The common sentiment and expenditure responses remain negative, while the assigned-cell coefficients remain small.

Table C.11: Common Exposure Among Urban Nonagricultural Households

	Sentiment	Next-wave expenditure
Leave-one-cell-out common exposure	-0.01896^{***}	-0.01506^{***}
	(0.00581)	(0.00438)
Assigned-cell exposure	-0.00470	-0.00257
	(0.00372)	(0.00203)
Coefficient sum	-0.02366^{***}	-0.01763^{***}
	(0.00609)	(0.00424)
p-value: common = assigned cell	0.062	0.020
Household fixed effects	Yes	Yes
Month-of-year fixed effects	Yes	Yes
Observations	407,354	407,354
Geographic-cell clusters	592	592

Notes: The sample excludes rural and agricultural households. The expenditure regression controls for heat at the expenditure date. Both specifications absorb separate household and month-of-year fixed effects, use analytic survey weights, and cluster standard errors by geographic cell. ***, **, and * denote significance at the 1, 5, and 10 percent levels, respectively.

C.14 Spatial support within district–year–months

The baseline common-exposure measure is defined whenever at least two cells are observed in a district–year–month. The weighted median is five cells. Table C.12 imposes more restrictive support requirements. Requiring at least five cells preserves the common-exposure result for all three outcomes. Requiring ten cells retains significant effects for intentions and expenditure but leaves only 48,372 observations and 391 geographic cells; the sentiment coefficient is then imprecisely estimated.

C.15 Influence of extreme expenditure observations

Table C.13 examines the influence of the upper tail of next-wave expenditure. The spatial estimates are essentially unchanged after either winsorizing or trimming expenditure at its weighted 99.9th percentile.

Table C.12: Common Exposure and the Number of Observed Cells

	Sentiment	Purchase intentions	Next-wave expenditure
Panel A. At least five observed cells			
Common exposure	-0.02044^{***}	-0.03129^{***}	-0.01766^{***}
	(0.00665)	(0.01041)	(0.00390)
Assigned-cell exposure	-0.00232	0.00536	-0.00060
	(0.00375)	(0.00529)	(0.00233)
Coefficient sum	-0.02276^{***}	-0.02593^{***}	-0.01826^{***}
	(0.00544)	(0.00826)	(0.00317)
p-value: common = assigned cell	0.052	0.011	0.002
Observations	288,177	288,177	288,177
Geographic-cell clusters	1,367	1,367	1,367
Panel B. At least ten observed cells			
Common exposure	-0.00223	-0.09537^{***}	-0.01655^{**}
	(0.01130)	(0.03650)	(0.00764)
Assigned-cell exposure	-0.00185	0.00848	0.00345
	(0.00808)	(0.01562)	(0.00505)
Coefficient sum	-0.00408	-0.08688^{***}	-0.01311^{**}
	(0.01080)	(0.02415)	(0.00539)
p-value: common = assigned cell	0.982	0.041	0.090
Observations	48,372	48,372	48,372
Geographic-cell clusters	391	391	391

Notes: Each panel restricts the baseline common spatial sample according to the number of observed cells in the district–year–month. Expenditure regressions control for heat at the expenditure date. All specifications absorb separate household and month-of-year fixed effects, use analytic survey weights, and cluster standard errors by geographic cell. ***, **, and * denote significance at the 1, 5, and 10 percent levels, respectively.

Table C.13: Spatial Expenditure Effects and Upper-Tail Observations

	Baseline (1)	Winsorized (2)	Trimmed (3)
Common exposure	-0.01359^{***}	-0.01353^{***}	-0.01333^{***}
	(0.00304)	(0.00304)	(0.00302)
Assigned-cell exposure	-0.00188	-0.00187	-0.00181
	(0.00173)	(0.00172)	(0.00171)
Heat at expenditure date	0.00445^{***}	0.00445^{***}	0.00444^{***}
	(0.00160)	(0.00160)	(0.00160)
Coefficient sum	-0.01547^{***}	-0.01541^{***}	-0.01514^{***}
	(0.00261)	(0.00262)	(0.00261)
p-value: common = assigned cell	0.0054	0.0054	0.0057
Household fixed effects	Yes	Yes	Yes
Month-of-year fixed effects	Yes	Yes	Yes
Observations	584,632	584,632	583,787
Geographic-cell clusters	1,796	1,796	1,796

Notes: Column (1) reports the baseline IHS expenditure specification. Column (2) winsorizes next-wave expenditure in levels at its 99.9th percentile weighted using the analytic survey weights, INR 77,389, before applying the IHS transformation. Column (3) removes observations above the same threshold. All specifications control for heat at the expenditure date, absorb separate household and month-of-year fixed effects, use analytic survey weights, and cluster standard errors by geographic cell. ***, **, and * denote significance at the 1, 5, and 10 percent levels, respectively.

C.16 Raw-current heat and the prior benchmark

The anomaly subtracts the preceding-year cell-month benchmark from current raw exposure. As a sensitivity, we enter current raw heat and the prior benchmark separately for both survey-schedule common and assigned exposure. Under district clustering, the current common coefficients remain negative but are less precise, and the equal-and-opposite anomaly restrictions are not uniformly supported. The p-values for the current common-plus-assigned raw response are 0.200 for sentiment, 0.109 for purchase intentions, and 0.054 for expenditure; the joint anomaly restrictions are rejected for sentiment and expenditure. We therefore treat the anomaly as a transparent normalization around a local seasonal benchmark, not as a structurally validated restriction that isolates a pure current-heat primitive.