Journal of Economic Behavior and Organization 227 (2024) 106703

Contents lists available at ScienceDirect

Journal of Economic Behavior and Organization

journal homepage: www.elsevier.com/locate/jebo

Daily temperature and sales of energy-using durables ☆

Jacopo Bonan [a,b,c], Cristina Cattaneo [a,b,*], Giovanna d'Adda [d,a,b], Massimo Tavoni [a,b,e]

[a] RFF-CMCC European Institute on Economics and the Environment, Italy
[b] CMCC Foundation - Euro-Mediterranean Center on Climate Change, Italy
[c] University of Brescia, Italy
[d] University of Milan, Italy
[e] Politecnico di Milano, Italy

ARTICLE INFO

JEL classification:
D91
D12
L81

Keywords:
Projection bias
Salience
Energy appliances

ABSTRACT

Decisions with significant and long-lasting consequences can be influenced by conditions at the moment of choice, such as weather. Using administrative data from an online retailer, we examine whether temperature and other weather variables affect the search and purchase of energy-using durables, namely, air conditioners (ACs) and dryers. We observe more sales of ACs on hot days and fewer sales of dryers on hot, windy days. We find no impact for appliances whose usefulness is not affected by the weather. For AC, weather-induced searches and purchases are in lower-efficiency energy classes. Product search data allow us to look into the process leading up to purchase. Prospective AC buyers search less intensively when the temperature is higher, and the opposite holds for buyers of dryers when temperature and wind speed increase. Models of memory and attention can explain these behavioral patterns. Understanding these dynamics is important for designing adaptation and mitigation policies, given the energy needs of cooling technologies and their increased demand and usefulness in a rapidly warming world.

1. Introduction

Decisions that have future consequences are ubiquitous. Although traditional economics assumes that individuals can correctly estimate future costs and benefits, evidence shows that decisions with large and long-lasting consequences are heavily influenced by tastes, emotions, and circumstances at the moment of choice (Busse et al., 2015; Simonsohn, 2010). Recent salience models (Bordalo et al., 2022, henceforth BGS) offer a unifying framework for a wide range of deviations from standard economic theory, such as projection bias (Loewenstein et al., 2003), reference dependence (Kőszegi and Rabin, 2006; Kahneman and Tversky, 1979) and framing effects (Bordalo et al., 2013). Specifically, BGS model attention as influenced, bottom-up, by contrasting, surprising, or prominent stimuli.

We analyze a decision with long-lasting consequences, the purchase of an energy-using durable, and show how it is affected by daily weather. We consider two types of appliances with significant impacts on household energy consumption: air conditioners

☆ We are grateful to Simone Ferro, Nicola Gennaioli, seminar participants at Lund University and University of Milan-Bicocca and Conference participants at ESA 2022 and FAERE 2023 for helpful discussion and comments. Matteo Muntoni and Milica Vranic provided excellent research assistance. We acknowledge financial support from the Energy Demand changes Induced by Technological and Social innovations (EDITS) project, which is part of the initiative coordinated by the Research Institute of Innovative Technology for the Earth (RITE) and International Institute for Applied Systems Analysis (IIASA) (and funded by Ministry of Economy, Trade, and Industry (METI), Japan).

* Corresponding author at: CMCC Foundation - Euro-Mediterranean Center on Climate Change, Italy.
E-mail addresses: jacopo.bonan@eiee.org (J. Bonan), cristina.cattaneo@eiee.org (C. Cattaneo), giovanna.dadda@unimi.it (G. d'Adda), massimo.tavoni@eiee.org (M. Tavoni).

https://doi.org/10.1016/j.jebo.2024.106703
Received 14 November 2023; Received in revised form 26 July 2024; Accepted 13 August 2024
Available online 29 August 2024
0167-2681/

(ACs) and dryers. Using data from an Italian online retailer, we investigate the impact of weather on the decision to buy and the search process leading up to it. We focus on two main weather dimensions that may affect these appliances' perceived usefulness: average daily temperature and wind speed. This setting allows us to observe behavior under minimal demand effects.

We find that higher daily temperature increases purchases of ACs and decreases those of dryers; higher wind speed reduces dryer purchases but has no impact on ACs. Other dimensions of weather meaningfully correlate with purchase likelihood: an index of discomfort, capturing perceived temperature and increasing with humidity, affects purchases of ACs but not dryers. Consistent with these results, higher temperature leads to a faster search process for ACs and a slower one for dryers. Temperature also impacts the energy efficiency of ACs purchased and viewed, shifting users' attention toward lower-efficiency products. We find no effect of temperature or wind speed on the energy efficiency of dryers viewed or bought or on sales of other types of appliances, such as washing machines or dishwashers. Our results are robust to using various samples; to controlling for temporal and spatial patterns of variation in sales; and to considering non-linear effects of temperature. Finally, a survival analysis of the search process also produces consistent results.

Our findings are consistent with a model of salience, adapted from BGS, that we propose. Weather affects the perceived usefulness of the appliance and, through it, the likelihood of purchase. As the appliance's attributes related to its usefulness become more salient, other attributes, such as its energy efficiency, lose prominence. These mechanisms have implications for the search process, which becomes faster and more superficial when weather makes the perceived usefulness of an appliance more salient.

We examine potential alternative explanations for our results. Rational behavior may explain users' behavior in several ways. First, customers may have already planned to buy the appliance, and the weather could influence the timing of the purchase, prompting them to "pull the trigger". In this case, we would observe a weather-induced intertemporal substitution of sales in our data. Second, weather may lead customers to use the appliance and realize it is broken. Under this scenario, our results would be driven by appliance replacements. Third, weather may allow users to acquire new knowledge connected to the usefulness of the appliance. Fourth, the purchase may address the need to urgently deal with the weather conditions. Our findings may also be consistent with a general tendency to remain home, search, and buy under particular weather conditions. Finally, selection issues and the impact of heat on cognitive performance could also be potential drivers of the results. Through a series of ancillary exercises and analyses, we provide evidence that our data are not consistent with any of these mechanisms and instead further support salience as the key channel behind the impact of weather on actual investment decisions.

Our findings complement the evidence on how the adoption of ACs responds to expectations about future climatic conditions and energy prices (Cohen et al., 2017; Rapson, 2014). Indeed, previous studies on projection bias show that short-term effects motivated by salience co-exist with considerations based on long-term expectations (Conlin et al., 2007; Busse et al., 2015). We agree with this interpretation, which suggests that our results do not rule out the influence of long-term expectations. However, our empirical analysis, relying on short-term quasi-random variations in weather, can only identify salience effects on sales. In robustness checks, we test for the effect of weather deviations with respect to long-term trends to better isolate responses to short-term variations in weather and find consistent results.

A key limitation of our analysis lies in the representativeness of our sample. First, users on the retailer's website may not be representative of internet users. On this point, we show that our results on search are qualitatively confirmed if we use Google trends data. Second, internet users are unlikely to be representative of the general population. In 2018, 44 percent of Italians reported never using the internet for online purchases (ISTAT, 2022b). Moreover, the adoption and ownership of air conditioners, particularly expensive goods, are strongly correlated with income levels (Davis et al., 2021). For these reasons, our user sample may be more educated, affluent, and less credit-constrained than the general population. Large cities are indeed over-represented in our sample compared to the Italian population. However, we do not believe that the lack of representativeness of our sample prevents the generalizability of our results. First, existing evidence demonstrates that everyone is subject to a large number of behavioral biases, even though individuals with lower cognitive ability, education, and income are more subject to inattention, memory, and present bias (Stango and Zinman, 2023). Second, heterogeneity analysis shows that our results are similar when we focus on users from regions in different quartiles of internet usage and education. Given that high education levels and income do not rule out behavioral biases, and that the use of the internet will increase in the future, our results may represent a lower bound, at least with respect to what might be expected in the long run and over a more representative sample.

Our analysis has other limitations. First, the retailer did not provide us with data on prices. We address this limitation in two ways. We retrieve prices for a subset of products in our sample through web scraping and price tracking services. In addition, we have data on the promotions active on products and always control for them in the analysis. Second, our data does not include any information on the supply side. We do not know how the stock of appliances in different energy classes changed with the weather, nor whether the retailer placed certain products on sale more prominently on the website depending on the weather. We test for the presence of weather-related strategic behavior on the part of the retailer, but our ability to explore this channel is limited. Finally, we have no proxy for purchases' welfare impact, such as the probability of returns, as in Conlin et al. (2007).

Our results are important from a policy perspective. ACs and dryers are expensive, with impacts on residential energy use lasting up to 15 years. Space cooling and clothes drying are responsible for nearly 20 and 6 percent of residential buildings' energy consumption, respectively (IEA, 2018; Bendt, 2010). AC generates additional impacts on global warming due to its emissions of greenhouse gases. However, AC is also a critical tool in adapting to rising global temperatures, which explains its growing adoption rates (Colelli et al., 2023). Understanding whether decisions with such large and long-lasting environmental impacts are subject to bias is important to designing mitigation and adaptation policies. Effective policy response requires identifying which bias is at work. Our results suggest that efforts to make the information on AC energy efficiency salient should be highest on hot days when weather increases sales of this type of good. Policymakers should also protect consumers from marketing strategies exploiting salience biases.

This paper's results contribute to the literature on projection and salience bias in consumer purchases. Several studies show the impact of daily weather or pollution on consumers' decisions (Conlin et al., 2007; Busse et al., 2015; Acland and Levy, 2015; Buchheim and Kolaska, 2017; Chang et al., 2018; Qin et al., 2019; Liao, 2020; Lamp, 2023) and life choices with long-lasting consequences (Simonsohn, 2010).[1] We complement this literature by focusing on products that have not been studied before, and whose benefit to the consumer is directly affected by weather; by considering a novel geographical setting, Italy; and by complementing the overall sales data with information on the search process and the energy efficiency rating of the appliances viewed and purchased on the website. By exploiting search data and information on appliance energy-efficiency classes, our paper provides additional evidence on the decision-making process behind salience-induced decisions. Moreover, the available data on both AC and dryers allow us to test the broader salience-induced effects of weather on sales. As in Busse et al. (2015), we document that salience and projection bias can lead to a change in the sale probability in both directions.

He et al. (2022) specifically analyzed the effect of weather on the decision to buy a high rather than low-efficiency AC and found that deviations from comfort temperature increase the likelihood of purchasing Energy Star models. We improve on this paper by providing evidence on the effect of temperature on sales and search of ACs; by using a more precise indicator of energy efficiency, i.e., EU energy classes; and by analyzing impacts at the individual level – exploiting within-individual variation in exposure to hot and windy weather – and at finer temporal and geographical granularity.

Studies typically pit salience and projection bias as two alternative mechanisms. However, despite agreeing that both are behind the evidence they present, they cannot ultimately distinguish empirically between them. Recent theoretical developments in the salience literature overcome the distinction between these two mechanisms by making projection bias a manifestation of the broader effects of salient stimuli on attention and choice (Bordalo et al., 2022). We apply this framework to explain our findings. In addition, our analysis of search provides the first evidence, to the best of our knowledge, of how bottom-up salience affects how individuals seek and attend to information.

The remainder of the paper is organized as follows. Section 2 describes the field where the analysis is conducted, Section 3 provides a theoretical framework for our findings, Section 4 describes the empirical approach and presents results for AC and dryers, Section 5 discusses possible alternative mechanisms and Section 6 concludes.

2. Setting and data

The share of Italians using the internet has been steadily growing in the past decades in Italy, reaching 80 percent in 2023 (Appendix Figure A.1). Although only about 16 percent of appliance sales occurred online in Italy in 2018, online channels played a crucial role in purchase decisions: 74 percent of buyers of large appliances initiated their search online (Flavián et al., 2020). Restrictions due to the COVID-19 pandemic have caused a 64 percent increase in online sales in Europe.[2] The reopening of physical stores has not returned online sales to their pre-pandemic levels.

The penetration rates of ACs and dryers in Italy are still limited, though rapidly growing. In 2021, only 48.8 and 15.2 percent of Italian families owned an AC and a dryer, respectively (ISTAT, 2022a). These figures were 29.4 and 3.3 percent in 2013 (ISTAT, 2014). The life cycle is estimated to be around 15 years for ACs and 13 years for dryers. Therefore, we expect that in our data, these appliances are mainly purchased for the first time.

We use data from a large Italian online retailer. Our data comprise the full navigation history of 112,428 website users between June 1 and October 16, 2018. We identify customers primarily through their registration ID. Users making a purchase must be logged in to the website. Instead, simply navigating the website does not require users to be registered or logged in. In these cases, we identify users through cookie-based tracking. Cookies are linked to the computer's IP address and browser. This implies that we cannot identify as the same user someone who visits the website without registering from different computers or browsers or those who clear cookies. We also cannot distinguish if multiple individuals view the website from the same shared computer. These limitations primarily affect our ability to follow users' full navigation history if they are not logged in to the website when browsing, hence our analysis of search behavior.[3]

For each page viewed by a user during the study period, we have information on the type of product viewed (AC, dryer, washing machine, dishwasher, refrigerator, freezer), the type of page (e.g., product, listing, or cart), and the number of seconds spent on the page. We also know whether the user ordered the product.

We match the navigation data with product data obtained from the retailer. We have information on the energy class for each model, identified by a unique product code. The EU energy label is displayed on each product page and lists the energy consumption in kWh and the energy class (Appendix Figure A.2). Energy classes are the results of engineering estimates based on the appliance's size, energy consumption, and other parameters and range from D (least efficient) to A+++ (most efficient).[4] Energy classes are an important tool for consumers to gauge energy efficiency, given the complexity of energy consumption information expressed in

[1] Projection bias is documented not only in the field but also in experimental settings (Augenblick and Rabin, 2019).

[2] Source: Eurostat, available at https://ec.europa.eu/eurostat/statistics-explained/index.php?title=E-commerce_statistics_for_individuals.

[3] See d'Adda et al. (2022) for more information.

[4] The new EU energy label, introduced in 2021, relabeled energy classes on a scale from G (least efficient) to A (most efficient) without changing how efficiency is calculated.

kWh (d'Adda et al., 2022; Houde, 2018). Labels based on similar energy-efficiency classes are widely used worldwide, including in countries like China, India, Brazil, and South Korea.

Our data include other product information. For ACs, we know whether they are portable or fixed. For all appliances, we have information on active promotions on the day of navigation, such as free delivery or zero interest rate for payments in installments. The retailer data do not include product prices. We retrieved this information to the best of our ability through web scraping between June and July 2022. We used product codes to collect current prices from the same online retailer or other major retailers for all the products still on the market in 2022. We gathered price information for 220 AC models (out of 517 in our sample) and 168 dryer models (out of 282). These prices are used to control for the relative price of products by energy class. Even if the price levels have changed in two years, the relative prices of products in different energy classes should not differ. We provide suggestive support for this claim by collecting 2018 prices through an online tracker. They are only available for 22 AC models. Finally, for ACs, we use web scraping to collect information on their size, proxied by the number of external and internal units associated with each product code, for the same sample of models for which we collect 2022 prices.

Appendix Table A.1 provides the average prices in 2022 by energy class. Higher energy efficiency corresponds to higher prices, as expected. For the limited sample of ACs (22) for which we have both 2022 and 2018 prices, the Pearson correlation is 0.80. These statistics validate the use of 2022 prices when we specifically analyze purchases of appliances in the different energy classes.

We geolocate IP addresses to identify users' municipalities during browsing, which allows us to match users with weather data. We collect mean and maximum temperature, wind speed, and rainfall for each day and municipality. The source for meteorological data is the E-OBS Temperature and Precipitation Data Sets (Cornes et al., 2018), an ensemble dataset available on a 0.1- and 0.25-degree resolution. We downscale the gridded data to a municipal level, averaging each municipal centroid's four nearest gridded points. In addition, from Mistry (2020), we retrieve information on the thermal discomfort index, which includes different meteorological drivers of discomfort, such as temperature and humidity.[5] Finally, to conduct heterogeneous analyses, we use data on regional penetration rates of ACs (ISTAT, 2022a), regional internet use (ISTAT, 2022b), regional education level (ISTAT, 2018), and municipal income (Ministero dell'Economia e delle Finanze, 2019).

3. Theoretical framework

Recent contributions to the theoretical literature on salience attempt to interpret within the same framework results previously explained by separate models, such as projection bias, present bias, reference-point dependence, and framing effects. The critical insight of the unified model proposed by BGS is that attention is influenced, bottom-up, by salient environmental stimuli. Contrast with surroundings, surprise relative to prior experiences, and prominence within the decision context determine salience. Salience can distract decision-makers from their goals or other relevant choice attributes. In our setting, high temperature or other weather variables vary the salience of the attributes linked to the appliances' usefulness. Weather thus affects the likelihood of sales of ACs and dryers through its impact on their perceived usefulness.

We adapt the framework proposed by BGS to formally explain how weather conditions affect the users' valuation of the two appliances. We denote the $K > 1$ attributes of ACs or dryers as $(a_1, \dots, a_K)$ and distinguish between those whose salience is influenced by environmental stimuli, belonging to subset P, and those not directly influenced, belonging to subset I. The intrinsic valuation of the good is:

$$V_p = \sum_{k \in P} w_k \pi_k a_k^n + \sum_{k \in I} \tilde{w}_k \pi_k a_k^n \tag{1}$$

The first term captures the valuation of attributes $(a_k)_{k\in P}$, the second term captures the valuation of features $(a_k)_{k\in I}$, and w_k and $\tilde{w}_k$ represent the distortions to decision weights (directly and indirectly) induced by salience. a_k^n represents the database of normal attribute values from memory. In our context, the prominence of an attribute is triggered by variations in outside daily weather that affect the decision weight attached to it (π_k). A crucial attribute that is part of the subset P in the first term of Eq. (1) is the usefulness of the appliance (the ability to cool a room or dry clothes). Heat distorts upward the valuation of this attribute for AC, while cold temperature and lack of wind distort it upward for dryers. Therefore, we expect that higher temperature increases the probability of purchasing an AC and lower temperature and higher wind speed increase the probability of purchasing a dryer.

We assume weight normalization as in Bordalo et al. (2012, 2013), which implies that the attention devoted to a salient attribute is diverted from nonsalient ones.

$$\sum_{k \in P} w_k \pi_k + \sum_{k \in I} \tilde{w}_k \pi_k = 1 \tag{2}$$

When it is warmer, the high salience of the usefulness of the ACs obscures their other attributes, such as energy efficiency, that are part of the subset I in the second term of Eq. (1). A similar effect is triggered by lower temperatures and higher wind speed for dryers. Correspondingly, when weather makes the usefulness of these appliances less prominent, other attributes become relatively more relevant and receive more attention. These other attributes may concern price, energy efficiency, or other product dimensions. However, it is hard to identify and test which attributes, among the many possible ones, become more salient due to this process.

[5] The dataset on discomfort has a smaller geographical coverage compared to that of temperature, wind, and precipitation. For six municipalities, we have missing information on discomfort.

Table 1
Summary statistics.

	(1) All Appliances	(2) AC	(3) Dryers
Viewers	112,428	12,984	12,648
Viewers—multiple days	48,076	7098	6921
Buyers		2250	3424
Buyers—multiple days		1409	2090
Days from entry to purchase		3.4	5
Seconds of search—buyers		500	557
Seconds of search—nonbuyers		126	171
Products viewed—buyers		1.6	1.4
Products viewed—nonbuyers		0.9	0.8

According to our model and empirical tests, the presence of a correlation between daily weather and purchases is evidence of biased decision-making. Our analysis will, therefore attempt to rule out alternative rational explanations for a correlation between weather and sales, such as intertemporal substitution, replacement, learning or urgency of purchases. What our model does not claim is that decision-making at optimal weather is unbiased and that we can quantify the magnitude of the bias as weather moves away from it.

4. Results

4.1. Descriptive statistics

The total sample of users of the retailer's website over the study period includes 112,428 individuals who viewed at least one appliance page (AC, dryer, washing machine, dishwasher, refrigerator, freezer). A subset, 48,076, searched for products over multiple days. About 12,984 users viewed an AC page and 12,648 a dryer page at least once. Of these viewers, about 55 percent navigated the website on more than one day in the period (Table 1). We observe 2250 sales of ACs and 3424 of dryers, respectively. 1409 and 2090 of these AC and dryer sales, respectively, occur after more than one day of search.

Table 1 presents the summary statistics. On average, the time between a buyer's first entry on the website and the moment of purchase is 3.4 days for ACs and 5 days for dryers. This time interval does not correspond to the number of days of searching, as days of search may not be contiguous. On average, buyers of ACs search for 8 min daily and view 1.6 products over the search period, and buyers of dryers search for 9 min and view 1.4 products. Nonbuyers search much less intensively: average daily minutes of search shrink to 2 and 3 min, and the number of products viewed decreases to 0.9 and 0.8 for AC and dryer nonbuyers, respectively. The differences in search time and intensity between buyers and nonbuyers are statistically significant.[6]

Fig. 1 (panel a) plots the average daily temperature in Italy over the study period, and Fig. 2 (panels a, b, c) shows the average daily temperature across municipalities over three consecutive Wednesdays within our study period, as an example. These figures show that we can leverage variability across time and space. It is not uncommon for heat waves to hit the north (or the south) and then move down (or up) the peninsula. Moreover, large variations in altitude or distance to the sea over small distances imply that municipalities within the same region and near each other may be exposed to different temperatures. We find high levels of variation for wind speed as well (Fig. 1, panel b and Fig. 2, panels d, e, f).

4.2. Sales

We present the results on the impact of weather on sales. Before discussing each result, we describe the empirical strategy behind it. We estimate the effect of weather on AC/dryer sales at the municipality level using the following specification:

$$y_{mt} = \beta_1 Temperature_{mt} + h_w + g_m + \varepsilon_{mt} \tag{3}$$

y_{mt} is a count variable equal to the total number of AC/dryer sales on the retailer's website in municipality m and day t. Temperature measures the average temperature for municipality m on day t. g_m are municipality fixed effects, and h_w are week fixed effects. Standard errors are clustered at the municipality level. We also display Driscoll–Kraay standard errors, to account for spatial and temporal correlation in weather. The sample includes 3899 Italian municipalities, about half of the total number of municipalities. The analysis period is June 1, 2018 to October 16, 2018.

With this specification, we rely on short-term quasi-random variations in weather. The short period of the analysis alleviates concerns that our temperature variable expressed in levels absorbs long-term trends and seasonality. In a robustness check, we employ a different specification, where weather is measured as deviations from the average temperature on the same week over the previous 10 years.

[6] All p-values of the two-sided t-tests are < 0.000.

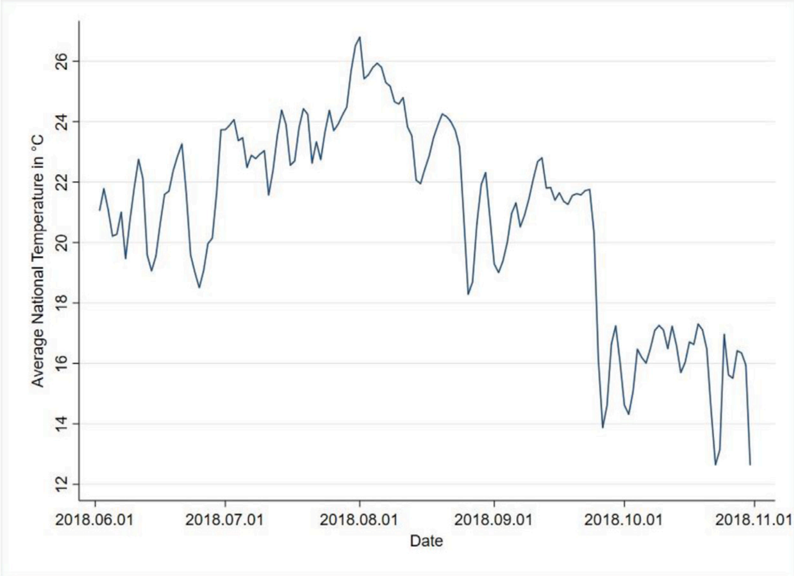

(a)

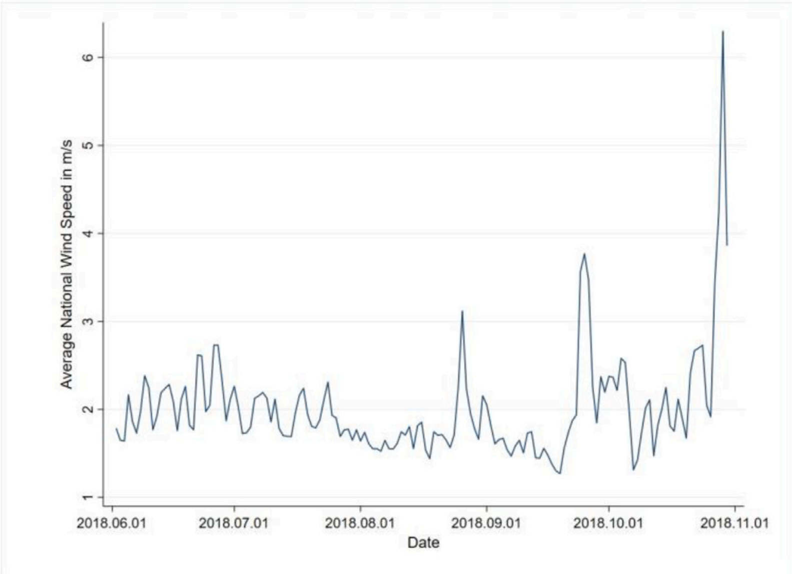

(b)

Fig. 1. Daily temperature and wind speed over the study period.
Notes: Average daily temperature (Panel a) and wind speed (Panel b) across Italian municipalities over the study period.

The results are shown in Columns 1 and 4 of Table 2. A higher mean daily temperature significantly increases the purchases of ACs and reduces those of dryers. Given that the average number of AC and dryer daily purchases in municipalities are 0.004 and 0.007, a 1 degree C increase in temperature is associated with a 7 percent increase in the number of AC purchases and a 6 percent decrease in dryer purchases. These variations are statistically significant, both when using clustered standard errors (in parentheses) and Driscoll–Kraay standard errors (in brackets). Results are consistent with Busse et al. (2015), where the temperature has the opposite effect on purchases of convertibles and four-wheel drive.[7]

[7] These results are robust to using non-linear models, such as Poisson pseudo maximum likelihood (PPML) that better deals with count data with many zeros and logit models (Appendix Table A.2). The sample of the specifications reported in Table A.2 is restricted to municipalities where at least one transaction was registered in the online retailer for ACs or dryers.

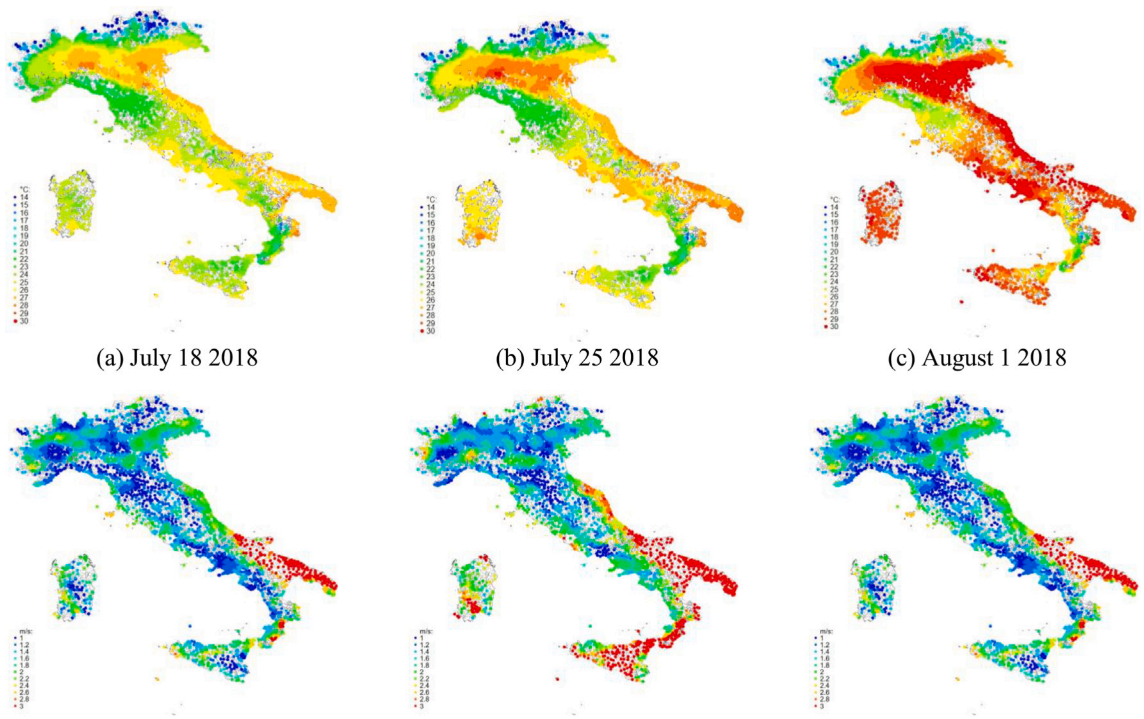

(a) July 18 2018 (b) July 25 2018 (c) August 1 2018

(d) July 18 2018 (e) July 25 2018 (f) August 1 2018

Fig. 2. Daily temperature and wind speed over the study period.
Notes: The figure depicts the average temperature (top panels) and wind speed (bottom panels) in Italian municipalities on specific days of the sample period. Each dot refers to a municipality included in the sample. The values in the legend represent the right bound of temperature and wind speed ranges starting from the previous value. For example, 15 degrees C means that the municipality experienced a daily temperature in the range of 14–15 degrees C.

Table 2
Effect of temperature on purchases.

Dep. Var.	(1) AC purchase	(2)	(3)	(4) Dryer purchase	(5)	(6)
Sample	Municipalities	All users	Viewers	Municipalities	All users	Viewers
Temperature	0.0003**	0.0005***	0.0023***	−0.0004**	−0.0004**	−0.0016**
	(0.0001)	(0.0001)	(0.0005)	(0.0002)	(0.0002)	(0.0007)
	[0.0001]	[0.0001]	[0.0005]	[0.0002]	[0.0002]	[0.0008]
Individual FE	No	Yes	Yes	No	Yes	Yes
Weekly FE	Yes	Yes	Yes	Yes	Yes	Yes
Municipality FE	Yes	No	No	Yes	No	No
Promo Dummy	No	Yes	Yes	No	Yes	Yes
Observations	538,062	193,582	39,772	538,062	193,582	38,822
Number of id	3899	48,076	7098	3899	48,076	6921
Mean Dep. Variable	0.00439	0.00786	0.0382	0.00653	0.0113	0.0562

Notes: The dependent variable in Column 1 is the number of AC purchases in municipality m and day t; in Columns 2 and (3) it is a dummy variable equal to 1 if individual i purchased an AC in day t and 0 otherwise; in Column 4, it is the number of dryer purchases in municipality m and day t; in Columns 5 and 6, it is a dummy variable equal to 1 if individual i purchased a dryer in day t and 0 otherwise; the sample in Columns 2 and 5 includes all users of the retailer's website who viewed at least one appliance page (AC, dryer, washing machine, dishwasher, refrigerator, freezer) in the period; the sample in Columns 3 and 6 includes only users who visualized at least one AC or dryer page in the period, respectively. Standard errors in parentheses are clustered at the municipality level. Driscoll-Kray standard errors in brackets. *** significance at the 1% level, ** at the 5% level, * at the 10% level. They refer to cluster robust standard errors.

Our main specification estimates the following individual-level equation:

$$y_{imt} = \beta_1 Temperature_{mt} + \beta_2 Promo_t + h_w + g_i + \varepsilon_{imt} \tag{4}$$

y_{imt} is an indicator variable equal to 1 if individual i located in municipality m purchased an AC/dryer in day t. We consider all days when user i is on the website. Temperature is again the average temperature in day t and municipality m, $Promo_t$ is an indicator variable for the presence of any promotions in day t on at least one of the ACs/dryers viewed by the user, g_i are individual fixed effects, and h_w are week fixed effects. The latter should absorb the average trend in purchases at the weekly level. We assume this seasonal variation is not area-specific but common to all municipalities in the dataset. Moreover, we believe that week fixed effects adequately absorb the temporal average pattern of purchases, given that, as indicated in Appendix Figure A.3, purchases do not follow a regular pattern within the week. We relax these assumptions in the robustness analysis. We display standard errors clustered at the municipal level and Driscoll–Kraay standard errors. The sample includes the users of the retailer's website who viewed at least one appliance page and accessed the website more than one day in the period.

This specification's results, which rely on within-individual variation in outside temperature, are displayed in Columns 2 and 5 of Table 2. They are consistent with the municipal-level results. Users are more likely to buy an AC on warmer days, and the opposite holds for dryers. This finding is confirmed when we focus on the sample of ACs and dryers viewers, i.e., on multiple-day users who viewed at least one AC or dryer page (Columns 3 and 6 of Table 2). According to this specification, if the temperature increases by 1 degree C, the incidence of ACs purchases increases by 0.23 percentage points, corresponding to 6 percent over the sample mean. The same increase in average daily temperature reduces the incidence of dryer purchases by 0.16 percentage points or 3 percent over the mean.[8] The statistical significance of the individual-level results is robust to the use of standard errors accounting for correlation across space and time.

The regression findings are confirmed in Fig. 3, where we plot the daily ratio of purchases to total views of any type of appliance on the website as temperature increases, for ACs (panel a) and dryers (panel b). This ratio grows with temperature for ACs but decreases for dryers.

As discussed in Section 3, salience theory posits that temperature may affect the likelihood that users purchase ACs or dryers because it affects the salience of their usefulness. To support this claim, we test if other dimensions of weather meaningfully affect purchases. We consider two additional variables: discomfort and wind speed. As mentioned, discomfort captures perceived temperature and depends on temperature and other factors, such as humidity. Higher perceived discomfort should make more salient the usefulness of AC. Lower wind speed is expected to increase the time that drying clothes outdoors takes, thus increasing the salience of the usefulness of dryers. In addition, we test whether users respond similarly to average and maximum daily temperatures.

We run Eq. (4), replacing temperature with the mean discomfort index at the day-municipality level (Table 3, Columns 1 and 2). Higher discomfort increases the likelihood that users purchase an AC on the website but does not affect dryer sales. We think that discomfort should not decrease dryers' perceived usefulness because it correlates with higher humidity. In Columns 3 and 4, we consider the wind speed, measured as a daily average, and find that stronger winds reduce dryer purchases but do not affect ACs. Although discomfort strongly correlates with temperature (Pearson correlation $\rho = 0.913$), wind speed and temperature do not show a strong correlation ($\rho = -0.091$). Therefore, we run a specification with both wind speed and temperature, excluding discomfort (Columns 5 and 6). Temperature alone influences AC purchases, and both higher temperatures and stronger wind negatively and significantly affect dryer purchases. In light of its explanatory power beyond that of temperature, in what follows, we include wind speed in all regressions examining dryer purchases and search.[9] In Columns 7 and 8, we replace the average daily temperature with the maximum daily temperature and find very similar results.

4.3. Search

We have seen that variations in daily weather affect sales on the website. We now investigate whether these effects extend to changes in the search process that precedes them. We first test whether temperature affects the viewing of AC or dryer pages on a given day. Consistent with the analyses on purchases, we first run a specification at the municipal level and then provide results at the individual level.

At the municipality level, we run Eq. (3), where y_{mt} is a count variable equal to the total individuals who viewed an AC/dryer in municipality m and day t. The results are displayed in Column 1 of Table 4. Panel A refers to ACs, and Panel B refers to dryers. We find that temperature increases the search intensity of ACs and reduces that of dryers, and that wind speed has a negative impact on the likelihood of searching for dryers. However, none of these effects is statistically significant.

We then run Eq. (4), where the dependent variable, y_{imt}, is an indicator equal to 1 if individual i located in municipality m viewed at least one AC/dryer page in day t. Column 2 of Table 4 shows results for the full sample of users. Consistent with the municipal level specifications, neither temperature nor wind speed significantly affects views. Although we cannot check whether the composition of viewers changes as temperature increases, the limited influence that we find of temperature and wind speed on the likelihood of search for both appliances alleviates selection concerns within our sample. We discuss this point in Section 5.

The search results concern only registered users, but we may wish to assess whether our search results are generalizable to the larger sample of internet users. We therefore use Google trends data on online search for ACs and dryers, and regress daily search volumes in each region on average daily temperature and wind speed in the region, controlling for week and region fixed-effects. The results are consistent with the ones using our sample in terms of the sign of the coefficients. However, the temperature in this analysis has a statistically significant effect on AC and dryer searches on Google (Appendix Table A.3).

To further examine the search process, we test whether daily weather affects daily search patterns. We run Eq. (4), where the dependent variable, y_{imt}, is the number of distinct ACs/dryers viewed in a day and the seconds spent viewing AC/dryer pages in a

[8] We rule out that the opposite effects of temperature on sales of ACs and dryers are due to substitution between the two appliances, as only 0.8% of the 112,428 users view both ACs and dryers, and thus may consider buying both appliances.

[9] All the results from the analysis presented in what follows are confirmed when we do not include wind speed in the dryers' specifications. Results are available upon request.

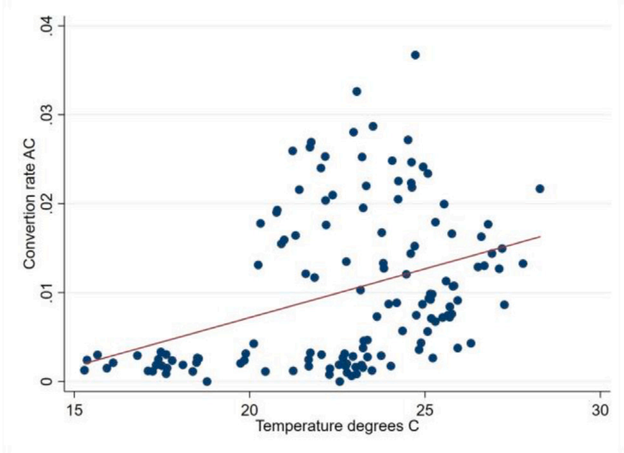

(a)

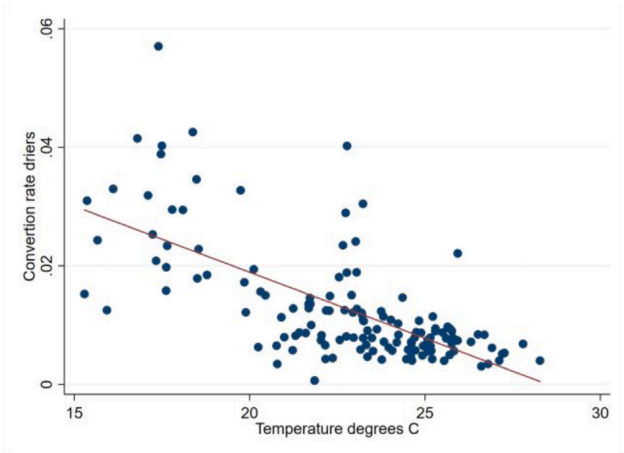

(b)

Fig. 3. Conversion rate and temperature.
Notes: The figure plots the average daily temperature and the conversion rate for ACs (panel a) and dryers (panel b). Conversion rate is the ratio between the number of users who purchased an AC/dryer and those who viewed at least one appliance in the day.

day. We focus on the sample of viewers. Daily weather has no significant impact on the number of distinct items viewed (Column 3) or seconds spent on product pages (Column 4) for both appliances.

For AC, this null effect may result from two opposite impacts of temperature on search. On the one hand, higher temperature increases the likelihood of purchasing an AC and, through this channel, the search depth. As we saw in Table 1, prospective buyers search more intensively and view more products than nonbuyers. On the other hand, higher temperatures might make the search process faster and less thorough in terms of the number of different items viewed. The opposite would hold for dryers: higher temperature and wind speed may reduce search time and intensity through their negative effect on the likelihood of purchase and increase them by inducing users to search longer and more intensively, resulting in the observed null effect.

To substantiate these claims, we provide two further pieces of evidence on the impact of weather on the speed at which users arrive at a purchase decision and the number of items they view in the process. First, we analyze the effect of daily temperature on the number of times the same AC/dryer is revisited daily. A revisit happens when users view a product that they had seen before during the day. Revisits characterize the end of the search process, when users focus on a few products they may consider buying and switch between them to compare them in more depth. We define daily revisits as the number of times a user views the same product, averaged over all products (ACs or dryers) viewed by the user on the day. The results in Column 5 indicate that a higher temperature on a day is associated with a significantly higher number of revisits among AC viewers and a marginally lower number among dryer viewers on that day. The result is consistent with the effect of temperature on purchases: for ACs, higher temperature

Table 3
Effect of weather on purchases.

Dep. Var.	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
					Purchase of			
	AC	Dryer	AC	Dryer	AC	Dryer	AC	Dryer
Sample					All users			
Discomfort	0.0006***	−0.0002						
	(0.0001)	(0.0002)						
	[0.0001]	[0.0002]						
Wind speed			−0.0001	−0.0008*	−0.0001	−0.0009*	0.0001	−0.0010**
			(0.0003)	(0.0004)	(0.0003)	(0.0005)	(0.0003)	(0.0005)
			[0.0003]	[0.0004]	[0.0003]	[0.0004]	[0.0003]	[0.0004]
Temperature (mean)					0.0005***	−0.0004**		
					(0.0001)	(0.0002)		
					[0.0001]	[0.0002]		
Temperature (Max)							0.0003***	−0.0003**
							(0.0001)	(0.0001)
							[0.0001]	[0.0001]
Observations	193,563	193,563	193,582	193,582	193,582	193,582	193,582	193,582
Number of ind.	48,072	48,072	48,076	48,076	48,076	48,076	48,076	48,076
Mean Dep	0.00786	0.0113	0.00786	0.0113	0.00786	0.0113	0.00786	0.0113

Notes: The dependent variable in odd columns is a dummy variable equal to 1 if individual i purchased an AC in day t and 0 otherwise; in even columns, it is a dummy variable equal to 1 if individual i purchased a dryer in day t and 0 otherwise; the sample includes all users. The sample size in Columns 1 and 2 is smaller by six municipalities, due to discomfort index data limitation. Standard errors in parentheses are clustered at the municipal level. Driscoll-Kray standard errors in brackets. *** significance at the 1% level, ** at the 5% level, * at the 10% level. They refer to cluster robust standard errors.

Table 4
Effect of weather on search.

Dep. Var.	(1) N. Viewers	(2) View	(3) Daily N. Items visualized	(4) Daily Time on Page (sec)	(5) Daily N. Revisits	(6) N. days of search in the period	(7) N. items visualized in the period
Sample	Municipalities	All users	AC Viewers	AC Viewers	AC Viewers	Buyers	Buyers
Panel A: AC							
Temp	0.0013	0.0002	−0.0036	1.3123	0.0258***	−0.2867***	−0.1150*
	(0.0010)	(0.0002)	(0.0044)	(1.6177)	(0.009)	(0.0964)	(0.0661)
	[0.0009]	[0.0002]	[0.0034]	[1.2886]	[0.0089]	[0.1612]	[0.0897]
Obs.	538,062	193,582	39,772	39,772	39,772	2250	2250
Numb. of id	3899	48,076	7098	7098	7098	2250	2250
Mean Dep	0.0328	0.0917	0.916	172.3	1.454	3.400	5.609
Panel B: Dryer							
Temp	−0.0014	−0.0002	−0.0037	−2.4041	−0.0204	1.1906***	0.2021***
	(0.0014)	(0.0003)	(0.0048)	(1.8881)	(0.013)	(0.1257)	(0.0488)
	[0.0010]	[0.0002]	[0.0036]	[1.9581]	[0.0137]	[0.3003]	[0.0717]
Wind sp	−0.0008	−0.0007	0.0053	3.9383	−0.0235	−0.3203	−0.2544**
	(0.0019)	(0.0007)	(0.0082)	(4.4335)	(0.027)	(0.2883)	(0.1207)
	[0.0016]	[0.0004]	[0.0103]	[3.7472]	[0.0192]	[0.4797]	[0.2127]
Obs.	538,062	193,582	38,822	38,822	38,822	3424	3424
Numb. of id	3899	48,076	6921	6921	6921	3424	3424
Mean Dep	0.0328	0.0924	0.912	239.2	1.703	5.055	4.893
Individual FE	No	Yes	Yes	Yes	Yes	No	No
Weekly FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Municipality FE	Yes	No	No	No	No	No	No
Promo Dummy	No	Yes	Yes	Yes	Yes	No	No

Notes: The dependent variable in Column 1 is the number of persons viewing an AC (Panel A) or dryer (Panel B) in municipality m and day t; in Column 2, it is a dummy variable equal to 1 if users viewed an AC (Panel A) or dryer (Panel B) page at least once in the day t and 0 otherwise; in Column 3, it is the number of distinct ACs (Panel A) or dryers (Panel B) viewed in the day; in Column 4, it is the seconds spent viewing AC (Panel A) or dryer (Panel B) pages in the day; in Column 5, it is the number of times the same AC (Panel A) or dryer (Panel B) is revisited in the day; in Column 6, it is the number of days that passed between the day the user first entered the retailer's website and the day of purchase of an AC (Panel A) or dryer (Panel B); in Column 7, it is the number of distinct ACs (Panel A) or dryers (Panel B) viewed in the period between entry and purchase. The sample in Columns 1 and 2 includes all users of the retailer's website who viewed at least one appliance page (AC, dryer, washing machine, dishwasher, refrigerator, freezer) in the period; in Columns 3-5, it includes AC (Panel A) or dryer (Panel B) viewers; in Columns 6 and 7, it includes only AC (Panel A) or dryer (Panel B) buyers. Standard errors in parentheses are clustered at the municipality level. Driscoll-Kray standard errors in brackets. *** significance at the 1% level, ** at the 5% level, * at the 10% level. They refer to cluster robust standard errors.

pushes the search process toward its final stage, which takes up a larger share of the overall search. The opposite holds for dryers, whose purchases decrease as temperature increases.[10]

Second, we analyze the duration and intensity of the search as a whole for buyers. In particular, we test whether the average temperature and wind speed buyers experience over their navigation period affect total search duration and the total number of distinct items viewed. We estimate the following specification:

$$y_{im} = \beta_1 Temperature_m + h_w + \varepsilon_{im} \tag{5}$$

where y_{im} measures the number of days elapsed between the day the user first entered the retailer's website and the day of purchase of the AC/dryer or the number of distinct ACs/dryers viewed in this period. The temperature is the average temperature over the period between the day the user first entered the retailer's website and the day of purchase of the AC/dryer. The sample only includes buyers.

Higher temperature leads to shorter search periods (Column 6) and fewer products viewed overall (Column 7) for ACs but has the opposite effect for dryers. These results confirm our hypotheses that higher temperatures make the search process faster and more superficial for ACs and longer and more intense for dryers.

4.4. Energy efficiency of sales and views

Next, we break down the impact of weather on sales and views by energy class. A first inspection of the data reveals that ACs purchased on the website belong, on average, to lower classes than those viewed. For example, let us take the sample of all ACs viewed at least once and split them into energy-efficiency classes. A quarter of ACs viewed belong to class A or less.[11] However, within the sample of ACs purchased, 32 percent belong to group A or less (Fig. 4). We detect the opposite relationship for dryers, with purchased dryers having higher energy efficiency than those viewed. The figure indicates that dryers belonging to the most efficient class, A+++, represent 19 percent of viewed dryers and 26 percent of sales.

We turn to regression analysis to study the effect of weather on the class of the products sold. Evaluating the effect of weather on buying an appliance in a specific energy class from the website requires that we control for other features of the product. For instance, as demonstrated in Appendix Table A.1, energy-efficiency class and prices are positively correlated, meaning that more energy-efficient products tend to be more expensive. We build a dataset where each observation is a product code, or model, p viewed by individual i on day t. Products viewed multiple times on the same day appear only once per day in the dataset. For each product, the data include some characteristics, including the 2022 price. We then run the following specification at the individual-day and product code level.

$$y^c_{ipmt} = \beta_1 Temperature_{mt} + \beta_2 Price_p + \beta_3 Promo_{pt} + \beta_4 X_p + h_w + g_i + \varepsilon_{ipmt} \tag{6}$$

y^c_{ipmt} is an indicator variable equal to 1 if individual i located in municipality m purchased on date t an AC/dryer with product code p that belongs to efficiency class c, where c are the classes A or less, A+, A++ and A+++; $Price_p$ is the price of product code p, retrieved in 2022, as described in Section 2; and $Promo_{pt}$ is an indicator for whether a promotion was applied to product p on date t. In the AC specification, X_p is a vector of product characteristics, such as the number of external and internal units and whether the AC is portable. The sample includes users who viewed an AC/dryer.[12]

Panel A of Table 5 displays the effect of temperature on the efficiency class of ACs sold on the website. Higher temperature shifts the distribution of sales toward products in the lowest class (A or less).[13] Panel B presents results for dryers. We find no strong evidence of an effect of temperature or wind speed on the likelihood of buying dryers in any energy class.

To check the validity of these last results and ensure that the different samples and units of observation are not affecting them, we use this product-individual-day level dataset to replicate our main specification that employed an individual-day level dataset. The positive and significant effect of temperature on the probability of AC sales in any energy class is confirmed (Appendix Table A.5, Column 1, Panel A). Similarly, the effect of temperature and wind speed on dryers sold in Panel B is consistent with the findings in Table 2.

We also estimate product views by class, using Eq. (6), where y^c_{ipmt} is a dummy variable equal to 1 if the AC or dryer viewed by individual i in day t belongs to one of the four energy classes. Panel A of Table 6 shows the results of this analysis for ACs. Higher temperatures take viewers' attention away from high-efficiency products. In particular, higher temperatures increase the probability of viewing items in low-efficiency classes, A or less and A+, and decrease the probability of viewing ACs in the A++ class. Panel B displays the effects of temperature and wind speed on dryers' energy class. We find no effect of weather on the different classes, except for A or less.

---

[10] That the increase in the number of revisits does not entail an increase in search time can be explained by the fact that revisits take less than first product views on average: 61 and 68 s for revisits versus 85 and 99 s for first views for ACs and dryers, respectively (both p-values = 0.000).

[11] In our sample, very few appliances fall in energy classes lower than A. Therefore, we group products in classes below A together with the A class.

[12] The number of observations in this individual-day-product dataset is smaller than in the individual-day one for three main reasons: users view few products on average, as shown in Table 1; the former dataset considers only days when users viewed an AC/dryer, but the latter includes all navigation days; and prices are not available for all products.

[13] He et al. (2022) analyzes the effect of temperature on Energy Star AC purchase compared to no Energy Star AC purchase. They find that the probability of such a choice increases with a temperature above 22 °C. We think our results are not at odds with their findings in the US context for two main reasons. First, 75 percent of their transactions concern Energy Star ACs. Second, they study sales at physical shops, which may be affected differently by psychological mechanisms. We also explore whether the differences between our results and those of He et al. (2022) may be due to the different settings of our studies. Penetration rates of ACs in the Italian markets are lower than in the US markets. In Appendix Table A.4, we repeat the analysis focusing on Italian regions with the highest AC penetration rates (top quartile), and find results that are consistent with our main specification.

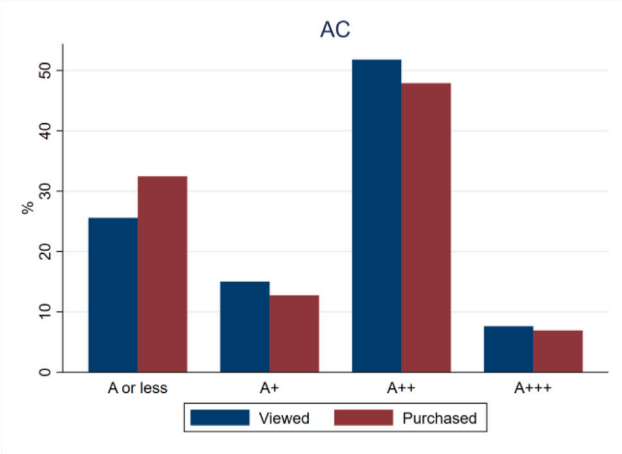

(a)

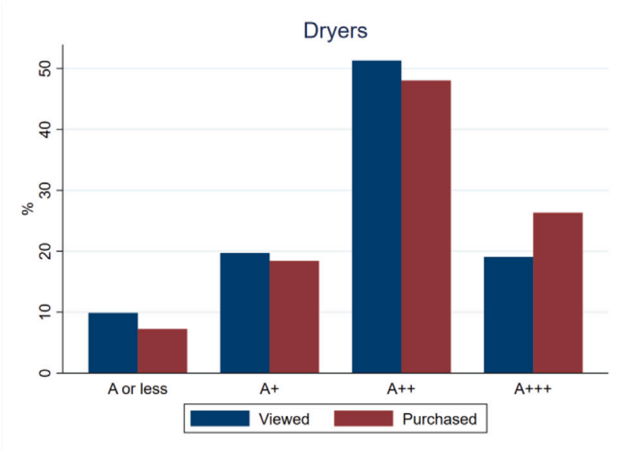

(b)

Fig. 4. Share of purchases and views by energy class.
Notes: The figure provides the number of ACs (panel a) and dryers (panel b) viewed and purchased by energy class as a share of the total number of ACs and dryers viewed and purchased over the study period, respectively.

Overall, search and purchase decisions for ACs are consistent in that we observe a qualitatively similar shift away from high-efficiency classes and toward low-efficiency ones in response to a higher temperature. These findings align with the theoretical framework described in Section 3. ACs' attributes related to being an effective remedy against heat become salient when the temperature is high. This leads to other attributes, such as their energy efficiency, receiving relatively less attention during hot days. Inattention to operating costs is considered one of the behavioral biases that constrain the adoption of energy-efficient appliances and vehicles (Allcott and Taubinsky, 2015; Allcott and Knittel, 2019; d'Adda et al., 2022). During hot days, neglect of energy costs increases because of the salience-induced effect of temperature. This behavioral bias affects the entire search process leading to a sale.

For dryers, the temperature appears to have a weaker influence on the salience of the energy classes. This result is consistent with the model's prediction that when the temperature is high, the ability to dry clothes may become less salient for individuals. However, it is not clear which other attributes become more salient as a result.

Studies have attempted to distinguish salience from projection bias. Given an intertemporal choice framework, whereby an agent receives utility from consuming a good from time t until period T, projection bias is defined as a situation where the agent's future utility is influenced by the state at time t (Loewenstein et al., 2003). A consumer's decision is not optimal because she employs current tastes to predict future tastes. BGS explain projection bias as resulting from salience-induced attention and thus include it in their framework, making the debate between salience and projection bias redundant.

However, our results shed light on a more nuanced distinction within the behavioral biases literature: the difference between utility and states in the context of projection bias. The former refers to a situation where agents correctly anticipate the future states

Table 5
Effect of weather on purchases by energy efficiency classes.

Dep. Var.	(1)	(2)	(3)	(4)
	Purchase			
	A or less	A+	A++	A+++
Sample	Products viewed			
Panel A: ACs				
Temperature	0.0019***	0.0007	0.0011	−0.0002
	(0.0006)	(0.0005)	(0.0007)	(0.0001)
	[0.0006]	[0.0005]	[0.0007]	[0.0001]
Observations	32,613	32,613	32,613	32,613
Number of ind	6311	6311	6311	6311
Mean Dep	0.015	0.010	0.020	0.001
Panel B: Dryers				
Temperature	0.0001	0.0003	−0.0011*	−0.0011
	(0.0002)	(0.0004)	(0.0006)	(0.0007)
	[0.0001]	[0.0003]	[0.0006]	[0.0008]
Wind speed	0.0004	−0.0008	−0.0021	−0.0008
	(0.0003)	(0.0011)	(0.0015)	(0.0016)
	[0.0002]	[0.0007]	[0.0015]	[0.0010]
Observations	36,561	36,561	36,561	36,561
Number of ind	6706	6706	6706	6706
Mean Dep	0.002	0.009	0.029	0.025
Individual FE	Yes	Yes	Yes	Yes
Weekly FE	Yes	Yes	Yes	Yes
Price	Yes	Yes	Yes	Yes
Promo Dummy	Yes	Yes	Yes	Yes

Notes: The dependent variable is a dummy variable equal to 1 if individual i purchased an AC (Panel A) or a dryer (Panel B) in the energy class indicated at the top of the column in day t and 0 otherwise. The sample includes all different models (product codes) of ACs (Panel A) or dryers (Panel B) viewed by individual i in day t. Standard errors in parentheses are clustered at the municipal level. Driscoll-Kray standard errors in brackets. *** significance at the 1% level, ** at the 5% level, * at the 10% level. They refer to cluster robust standard errors.

Table 6
Effect of weather on search by energy efficiency classes.

Dep. Var.	(1)	(2)	(3)	(4)
	View			
	A or less	A+	A++	A+++
Sample	Products viewed			
Panel A: ACs				
Temperature	0.0035**	0.0040*	−0.0072**	−0.0013
	(0.0015)	(0.0022)	(0.0028)	(0.0011)
	[0.0013]	[0.0015]	[0.0015]	[0.0008]
Observations	32,613	32,613	32,613	32,613
Number of ind	6311	6311	6311	6311
Mean Dep	0.338	0.216	0.400	0.0338
Panel B: Dryers				
Temperature	0.0013**	−0.0002	−0.0011	−0.0000
	(0.0006)	(0.0012)	(0.0020)	(0.0018)
	[0.0006]	[0.0012]	[0.0011]	[0.0014]
Wind speed	0.0003	−0.0021	0.0012	0.0005
	(0.0014)	(0.0027)	(0.0042)	(0.0033)
	[0.0013]	[0.0041]	[0.0051]	[0.0028]
Observations	36,561	36,561	36,561	36,561
Number of ind	6706	6706	6706	6706
Mean Dep	0.0335	0.175	0.447	0.345
Individual FE	Yes	Yes	Yes	Yes
Weekly FE	Yes	Yes	Yes	Yes
Price	Yes	Yes	Yes	Yes
Promo Dummy	Yes	Yes	Yes	Yes

Notes: The dependent variable is a dummy variable equal to 1 if the AC (Panel A) or the dryer (Panel B) viewed by individual i in day t belongs to the energy class indicated at the top of the column and 0 otherwise. The sample includes all different models (product codes) of ACs (Panel A) or dryers (Panel B) viewed by individual i in day t. Standard errors in parentheses are clustered at the municipal level. Driscoll-Kray standard errors in brackets. *** significance at the 1% level, ** at the 5% level, * at the 10% level. They refer to cluster robust standard errors.

but erroneously predict the utility they will receive by consuming the good in that state. The latter captures mistaken beliefs about how the world will look in the future, with agents overestimating the likelihood that future states resemble the current one.

We argue that projection bias of states implies that higher temperatures on a day should lead users to believe that they will use an AC more and a dryer less because hot days will be more frequent. For ACs, increased predicted usage comes with increased predicted running costs. Therefore, we would expect more, rather than less, attention to energy efficiency if the primary mechanism were projection bias of states. Our results that temperature shifts AC views and purchases away from high-efficiency classes are instead in line with projection bias of utility, whereby consumers are affected by mistaken beliefs of the utility derived from using the AC.

4.5. Robustness

We check the robustness of our results along different dimensions.

Deviations from long-term averages. We use a specification where temperature is measured as deviations from long-term averages. In particular, we generate deviations as the difference between temperature on a day and the average temperature in the same week over the previous 10 years. This specification is adopted by studies using several years of data, to control for seasonality in temperatures (Lamp, 2023), and to better control for the effect of long-term trends in temperatures. Our results are confirmed. High temperatures relative to long-term averages increase sales of ACs and decrease those of dryers (Appendix Table A.6).

Nonlinearity of temperature effects. We estimate the effect of temperature using a more flexible specification that allows heterogeneous effects across temperature bins. We replace daily temperature with 2-degree bins in our main specification on appliance purchases. Appendix Figure A.4 shows the regression results graphically. Consistent with the main findings, higher temperature increases the likelihood of AC purchases and decreases that of dryers. Conversely, if we add squared temperature to our main specification, the coefficients of the temperature terms turn nonsignificant (Appendix Table A.7). We thus reject the hypothesis of a nonlinear relationship between temperature and AC/dryer purchases.

Event-history analysis of the search process. We adopt a more flexible way to model the search process by performing an event-history analysis.[14] We focus on buyers and ask whether the likelihood that a purchase occurs on the first, second, or nth day of the search is affected by temperature on that day. This analysis confirms that higher temperatures lead to faster purchases for ACs and slower ones for dryers both when temperature enters the analysis linearly and when we use temperature bins (Appendix Table A.8). The estimates in Columns (1) and (3) indicate that a temperature increase by one degree Celsius is associated with a 3.5 percent increase and an 8.4 percent decrease in the likelihood of buying an AC or dryer, respectively, on a day among buyers, conditional on not having made a purchase yet. Appendix Figure A.5 displays the results of this analysis for a selected subset of temperature bins. It shows how the share of buyers who still have not made a purchase changes throughout the search. It drops on the first day, consistent with the large share of buyers who purchase on the same day as they start their search, and then decreases at a slower rate as the number of days of search increases. Higher temperatures are associated with higher drops in the early portion of the survival curve among AC buyers and lower drops among buyers of dryers. These results are consistent with our regression analysis of purchases and search.

Examining the role of prices. The impact of temperature on purchases might be due to the retailer changing its prices as the weather gets hot. Or, more generally, the seasonal trend in temperature may be correlated with seasonality in prices, and our temperature effects may, therefore, be capturing the effect of prices. Our empirical strategy addresses this concern in two ways. First, we identify temperature effects by exploiting local variations, while prices are set at the national level. Second, the inclusion of the "Promo" dummy in Eq. (4) controls for possible promotions that the online retailer may strategically launch during hot days. We also test whether different types of promotions are correlated with temperature, and find no evidence to support this claim (Appendix Figure A.6).[15]

We also conduct additional analysis to show that our results are robust when we control for prices (under the limitations of price data availability already discussed). We begin by exploiting the data on the daily price of products that we retrieved for 2018 on a subset of 22 ACs in our sample to examine the relationship between daily temperature and price (Appendix Figure A.7). An exploratory analysis indicates a positive relationship between temperature and price. Considering the small sample size, we find suggestive evidence that prices increase on average by 4.2 Euro for each additional degree of temperature.[16] Such a positive correlation is present for products in all energy efficiency classes. Investigating the reasons for this would be beyond the scope of this paper. However, the price increase when the temperature rises implies that users who purchase an AC on a hot day do so at greater expense than on a cooler day.

Next, we check whether our results are robust to controlling for price.[17] The results, presented in Appendix Table A.5, Column 2, suggest that the effects of temperature and wind speed on the probability of purchases are robust to including current prices. We

[14] We use a logit-hazard model where the time indicators and the covariates are associated with the logistic transformation of the hazard (Singer and Willett, 1993).

[15] The retailer at that time did not have advertising campaigns on the press, TV, or online, so we can rule out that other promotional activities were correlated with weather.

[16] The coefficient of temperature is statistically significant at the 1 percent level. Results are available upon request.

[17] As described in the analysis, we use the information on current product prices retrieved in 2022. However, they are shown to be highly correlated with prices at the study time.

also further explore the role of promotions. Rather than including an indicator for whether a promotion was offered on any of the ACs/dryers during the day in our user-level sample, we include a set of dummies for the specific promotions targeting the different products in the product-level sample. Among these are so-called *countdown offers*, i.e., promotions that try to generate a sense of urgency in the users and that the online retailer may strategically launch during hot days. The coefficients of temperature and wind speed are robust to the inclusion of a complete set of promo dummies (Column 3) or to controlling only for the countdown offer dummy.

We may worry that other retailers strategically place offers on AC/dryers depending on daily temperature, affecting the behavior of users in our dataset. We do not observe offers on ACs or dryers of other retailers, so we have no data to directly test the relevance of this concern. However, the fact that our results are consistent when we include (Column 3) or exclude (Column 1) the promotions applied by our partner retailer suggests that unobserved promotions by other retailers may similarly not exert a significant influence on the behaviors that we observe in our data.

Robustness to alternative specifications. We test alternative specifications by changing the set of controls included in the regression and the analysis sample.

First, we make more stringent assumptions on the temporal and spatial pattern of purchase variation. Specifically, we replace week-fixed effects with province-by-week fixed effects to allow the weekly variation in purchases to differ across provinces, which corresponds to NUTS3 geographical units (Appendix Table A.9, Columns 1 and 2).[18] In an alternative specification, we also allow for within-week variation in purchase patterns. If individuals have more time for shopping, including online shopping, during weekends, using week fixed effects would not capture this pattern. Therefore we add weekend fixed effects to our main specification (Appendix Table A.9, Columns 3 and 4). The results remain stable to the alternative specifications.

Second, we test the sensitivity of our results to different definitions of the sample. In our main specification (4) with individual fixed effects, the sample includes only users who visit the website on multiple days. However, 57% of viewers appear only for a single day in the database (Column 1 of Table 1). As a robustness check, we run that specification without individual fixed effects, on the sample that includes users from both single and multiple days (Appendix Table A.10, Columns 1 and 2). Our results are robust to changes in the analysis sample and the specification used. The coefficient of temperature on purchases of ACs in Column 1 is positive and statistically significant, albeit smaller than our main specification. Columns 3 and 4 report the estimates of a specification that does not employ individual fixed effects on our preferred sample of users with multiple observations and find similar results.[19] This implies that the impact of weather on users buying on the same day as we observe them entering the website is similar to that of users navigating on multiple days. Finally, the sample of buyers in our main specification also includes individuals who purchased more than one appliance of the same type in the period.[20] As a robustness check, we drop from the sample multiple transactions and keep only the first purchase. The results are robust (Table A.11).

Energy class of products viewed and purchased. Finally, we test whether our results on the energy class of viewed products are robust to different ways to define them. We analyze the effect of temperature on the class of the first product viewed in the day or the class of the least or most efficient products viewed during the day. We find results consistent with previous findings (Appendix Table A.12). Temperature tends to lower the class of ACs but not dryers. During hot days, the class of the first AC, least efficient AC, and most efficient AC viewed are lower.

5. Alternative mechanisms

Salience is consistent with all the evidence that we collect using our data: the effect of temperature on the likelihood of purchase of AC and dryers, the additional effect of wind speed on dryer purchases, and the impact of temperature on the energy class of purchased ACs. In addition, we show how salience works through the search process in generating decision outcomes. Neglecting nonsalient and complex attributes helps explain why the search process is quicker and more superficial for ACs, yet the opposite holds for dryers. The reduced attention paid to the energy efficiency of ACs at high temperatures is also reflected in the time spent viewing products in different classes during the search process.

We consider three possible alternative explanations for our results and discuss the evidence related to each.

5.1. Rational behavior

Users' behavior may be rational. They may already consider buying an appliance, but they may not be perfectly aware of the value they place on the appliance. If so, it would be rational for them to wait to actually buy until the value of the appliance to them is perfectly revealed, thanks to daily weather. In the case of ACs, a hot day may make people realize that no other adaptation behavior – using a fan, letting air circulate in the house, etc. – is sufficient to make the heat bearable; or they may already own an AC and realize that it is broken when turning it on a hot day; or they may be monitoring daily variations in prices and exploit discounts present on hot days. For dryers, cooler and damper weather may have the same effects. In all these cases, weather affects

[18] We include province-by-week fixed effects, rather than municipality-by-week fixed effects, because our main weather variable is municipality-time specific. The latter approach would capture most of the variability we exploit in the analysis.

[19] This should alleviate concerns related to our potential limited ability to track users as they navigate on the website on multiple days.

[20] Multiple purchases are primarily due to issues with the initial order, such as problems with the payment method (e.g., prepaid card without enough credit) or delivery. See d'Adda et al. (2022) for further details.

the timing of a purchase that would ultimately take place. In other words, weather induces users to "pull the trigger" on a purchase that they were going to make anyway.

We conduct a set of exercises to exclude any of the above instances as the explanation of our results. First, we check for the presence of intertemporal substitution. If weather only affects the timing, but not the likelihood of a purchase, we should observe that temperature induces substitution across days in the number of sales. For instance, if warm weather two days ago encourages someone to purchase an AC on that day, then they would not purchase it today even if it is hot. Similarly, if they bought an AC today because it is hot, then they will not purchase it in two days when it is hot again. A similar reasoning holds for dryers. We test for intertemporal substitution in sales by including 15-day temperature lags in our main specification.[21] If an intertemporal substitution occurs, we would expect the coefficient at time 0 to be equal in magnitude but opposite in sign to the sum of the lagged coefficients. This is not what we find: the lagged coefficients are both positive and negative for ACs and dryers (Appendix Figure A.8). Similarly, the effect of the current temperature on purchases is not accompanied by a corresponding opposite effect of temperature in subsequent days. Based on this evidence, we exclude that our results simply capture the effect of temperature on the timing of purchases by people already considering buying an AC or a dryer. Temperature changes the likelihood of purchases rather than only shifting their timing.

Second, we explore the possibility that users already own ACs, but realize that their systems do not work only when a hot day comes along and they turn the AC on. Our results would thus be explained by the probability of replacing an AC being higher on hot days because that is when users discover a malfunction. We believe that this is not the case for two reasons. First, as mentioned, the penetration rate of AC in Italy is still limited. Recall that the penetration rate of AC was 48.8% in 2021 and 29.4% in 2013. Moreover, the average lifespan of an AC is between 10 and 30 years (Litardo et al., 2023). It is, therefore, unlikely that the main driver of purchases in 2018 is replacement. Second, the probability of repairing, rather than immediately replacing it with a new one, is positively related to the item replacement cost (Jaeger-Erben et al., 2021), which is high for large appliances such as ACs. This reasoning holds even more for dryers, given the lower penetration rate of dryers relative to ACs in Italy.

To more formally test the role of replacement, we split the sample by quartiles of the regional penetration rate of ACs and run separate regressions (Appendix Table A.13).[22] If replacement were relevant, we would observe larger effects of temperature on sales in regions with high AC penetration. This analysis can also assess the validity of another rational explanation for our results, namely that users do not know the weight π_k that they attach to the cooling power of ACs, as described in Eq. (1). People living in areas with low rates of penetration of ACs might be particularly unaware of their usefulness. If this were the case, we would expect the effect of temperature on sales to be stronger in regions with low penetration rates. We do not find support for either explanation. First, contrary to the replacement explanation, temperature also affects sales in regions in the second quartile of AC penetration, and the magnitude of the coefficients on temperature are smaller in the regressions run on the sub-samples of regions in the top two quartiles than in the two bottom ones. Second, temperature significantly affects sales in the top quartile of AC penetration, against the learning explanation.

Third, we discuss the possibility that users who are set on buying an appliance may monitor the website and purchase the item during days in which prices are low. If daily variations in prices are negatively correlated with temperature for ACs, our results may be explained by low prices rather than high temperatures. Our exploration of the correlation between prices and temperature revealed that, contrary to this explanation, AC prices and temperature are positively correlated (Appendix Figure A.7).[23]

Fourth, in our sample, about 40% of AC or dryer buyers appear only for a single day in the database (Columns 2 and 3 of Table 1), implying that they purchase on the same day as they enter the website. Recall that our results are robust to including these buyers and this result is confirmed by event-history analysis (Appendix Table A.10 and Appendix Figure A.5). Available statistics indicate that buyers of large appliances start their search for an appliance online rather than in physical shops (Flavián et al., 2020). Therefore, we exclude that single-day buyers were set on buying the appliance before we observed them on the website.

These additional tests and considerations reasonably discard the possibility that temperature only affects the timing of a purchase that would occur anyway. Below, we consider other rational explanations whereby temperature may induce people to purchase an appliance under correct beliefs on the future usefulness of the appliance.

First, users may buy the appliance to address an urgent need generated by the daily weather. In the case of AC, this would be the need for cooling on a hot day. The utility derived from satisfying such needs during hot days may outweigh any cost consideration, regardless of any correct beliefs on the utility from future usage. In our setting, it is unlikely that any urgent need can be satisfied by purchasing an AC. Most ACs purchased online require delivery and installation, which may take weeks.[24] This creates a gap between the purchase time and when buyers will perceive the benefits from it. Such a gap is likely to increase further at peak demand. We can test whether purchases aim to satisfy an urgent need by exploiting the presence in our data of portable ACs, which require no installation and thus allow users to quickly satisfy an urgent need, as usage is possible immediately upon delivery. We test whether our main result is driven by the purchases of portable ACs by testing the effect of temperature on ACs that require installation and portable ACs, separately. The effect of temperature on the likelihood of purchase is indistinguishable between the two groups of ACs (Appendix Table A.14).[25] This result suggests that the desire to satisfy an urgent need is not behind the effect of temperature

[21] This strategy follows Busse et al. (2015). We use fewer lags, 15 versus 60, because our sample period (3 months) is shorter than theirs (7 years).

[22] We cannot conduct the same analysis for dryers because regional breakdowns of dryers' penetration rates are not available.

[23] As noted above, this correlation is positive both for high and low-efficiency ACs.

[24] This consideration generalizes to the ACs bought in physical stores.

[25] The t-test on the equality of the coefficients in the two regressions is 0.74.

on our observed AC purchases. In the case of dryers, this explanation is even less compelling, as we struggle to consider the need to dry clothes as one that must be urgently met.

Another possibility is that current temperature may affect purchases because of rational learning about future temperatures. In a warming climate, hot days may be associated with acquiring knowledge about the increasing frequency of heat waves and rising global temperatures. AC purchases would be a rational response to this knowledge. While we do not rule out the relevance of this mechanism, our empirical specification should isolate the impact of short-term weather variations. Moreover, our results on the impact of wind speed on dryer purchases and of temperature on the energy class of AC purchased confirm that this is not the only mechanism at work in our setting.

Finally, the temperature effect on ACs could result from a general tendency to stay home on hot days to avoid the heat. This could lead to more browsing and buying online. If that were the case, we would observe a similar effect of temperature on both AC and dryer purchases. This is not what we find, as the temperature effect on dryers is of the opposite sign than that on ACs. Moreover, if this were the mechanism, we would find an effect of temperature on purchases of other appliances. We test this hypothesis by running our main specification on the likelihood of purchase of other appliance types (refrigerators, dishwashers, and washing machines) and find no effect (Appendix Table A.15).

5.2. Selection

Our results could be due to selection. A specific type of buyer, with a higher willingness to pay and a lower preference for energy efficiency, could be more likely to browse and buy on hot days. Our existing analysis offers three arguments against this possibility. First, the use of individual fixed effects in our main specification alleviates the potential relevance of this explanation, as it allows us to isolate the effect of temperature variations across days for the same individual. Second, the null effect of temperature and wind speed on the likelihood of search at individual and municipality levels suggests that temperature does not induce selection. Third, the robustness of our results to the multiple specifications and samples we adopt is also reassuring.

We can also conduct heterogeneity analysis to assess the relevance of selection for our results on the energy efficiency of AC sales. The higher probability of buying low energy-efficient ACs during hot days may be driven by selection, in that more liquidity-constrained users may be more likely to go on the website to buy an AC on hot days. To test this possibility, we separately analyze the effect of temperature on purchases of products in the lowest class (A or less), by quartiles of the income of the municipality where users live (Appendix Table A.16). An explanation based on liquidity constraints would imply a stronger effect of temperature in lower-income municipalities. Transactions of ACs of any class during hot days are concentrated among users living in the richest municipalities (Column 7) and these users tend to purchase low-efficiency appliances (Column 8). The coefficient is in line with that reported in Table 5, thus discarding the liquidity constraint hypothesis.

A second type of selection refers to the representativeness of our sample. Our sample comprises users of an online retailer's website: as discussed above, internet users tend to be wealthier and more educated than the average citizen. Our results may thus not be generalizable to the entire population. To assess the relevance of such a lack of representativeness, we conduct heterogeneity analysis by internet access and education. Specifically, we split the regions in our sample into quartiles based on the average share of internet users in the region, and run our main specification on the sub-sample of users living in each quartile (Appendix Table A.17). We do the same for education (Appendix Table A.18). Our results show that the effect of temperature on AC sales is of the same sign and similar magnitude across sub-samples. These results suggest that our findings could hold beyond the selected sample of internet users.

5.3. Cognitive ability

A growing literature documents the effect of temperature on cognitive ability and preferences. Evidence shows that heat, specifically temperature above 26 degrees C, makes it harder for people to focus and reduces cognitive performance in high-stakes decision settings (Heyes and Saberian, 2019; Graff Zivin et al., 2020). If so, the effect of heat on the decision process could result from cognitive impediment rather than salience. Two of our results are not consistent with this direct effect. First, we find opposite effects of temperature on AC and dryer sales and search and no effect of temperature on sales of other types of appliances. Second, the impact of wind speed on dryer purchases suggests that our findings are not specific to temperature.

6. Conclusion

We find that salience bias explains the effect of daily weather on the decision to purchase ACs and dryers, the search process leading up to it, and the energy efficiency of viewed and purchased ACs. Methodologically, our results offer novel evidence of how salience works through the search process in generating decision outcomes. Moreover, our analysis adds to existing literature, documenting a behavioral bias in a policy domain that is particularly relevant in the face of the current climate and energy crises. Transitory conditions significantly influence decisions with long-lasting and large economic and environmental consequences.

Our results suggest that policymakers should focus on incentives or information interventions to boost the energy efficiency of AC purchases during hot days when many of these decisions are likely to occur. If individuals purchase ACs on those days, then encouraging efficient choices at that time is likely to have long-term repercussions on aggregate energy consumption.

Our work has limitations, primarily concerning the nature of our setting and the data available. As our data come from an online retailer, we have no information on what its users do when they are not on its website, which may include searching for or purchasing products elsewhere. Users of online retailers may also not be a representative sample of the population—although online shopping is increasingly widespread—and our lack of individual-level characteristics does not allow us to assess the extent of any selection issues. Further research should address these limitations.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jebo.2024.106703.

References

Acland, D., Levy, M.R., 2015. Naiveté, projection bias, and habit formation in Gym attendance. Manage. Sci. 61 (1), 146–160.

Allcott, H., Knittel, C., 2019. Are consumers poorly informed about fuel economy? Evidence from two experiments. Am. Econ. J. Econ. Policy 11 (1), 1–37.

Allcott, H., Taubinsky, D., 2015. Evaluating behaviorally motivated policy: Experimental evidence from the lightbulb market. Amer. Econ. Rev. 105 (8), 2501–2538.

Augenblick, N., Rabin, M., 2019. An experiment on time preference and misprediction in unpleasant tasks. Rev. Econ. Stud. 86 (3), 941–975.

Bendt, P., 2010. Are We Missing Energy Savings in Clothes Dryers?. Technical report, ACEE Summer Study on Energy Efficiency in Buildings, Ecos.

Bordalo, P., Gennaioli, N., Shleifer, A., 2012. Salience theory of choice under risk. Q. J. Econ. 127 (3), 1243–1285.

Bordalo, P., Gennaioli, N., Shleifer, A., 2013. Salience and consumer choice. J. Polit. Econ. 121 (5), 803–843.

Bordalo, P., Gennaioli, N., Shleifer, A., 2022. Salience. Ann. Rev. Econ. 2022 14, 521–544.

Buchheim, L., Kolaska, T., 2017. Weather and the psychology of purchasing outdoor movie tickets. Manage. Sci. 63 (11), 3718–3738.

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

Chang, T.Y., Huang, W., Wang, Y., 2018. Something in the air: Pollution and the demand for health insurance. Rev. Econ. Stud. 85 (3), 1609–1634.

Cohen, F., Glachant, M., Söderberg, M., 2017. The cost of adapting to climate change: evidence from the US residential sector. GRI Working Papers.

Colelli, F.P., Wing, I.S., Cian, E.D., 2023. Air-conditioning adoption and electricity demand highlight climate change mitigation–adaptation tradeoffs. Sci. Rep. 13 (1), 4413, Number: 1 Publisher: Nature Publishing Group.

Conlin, M., O'Donoghue, T., Vogelsang, T.J., 2007. Projection bias in catalog orders. Am. Econ. Rev. 97 (4), 1217–1249, Publisher: American Economic Association.

Cornes, R.C., van der Schrier, G., van den Besselaar, E.J.M., Jones, P.D., 2018. An ensemble version of the E-OBS temperature and precipitation data sets. J. Geophys. Res.: Atmos. 123 (17), 9391–9409.

d'Adda, G., Gao, Y., Tavoni, M., 2022. A randomized trial of energy cost information provision alongside energy-efficiency classes for refrigerator purchases. Nat. Energy 7 (4), 360–368, Number: 4 Publisher: Nature Publishing Group.

Davis, L., Gertler, P., Jarvis, S., Wolfram, C., 2021. Air conditioning and global inequality. Global Environ. Change 69, 102299.

Flavián, C., Gurrea, R., Orús, C., 2020. Combining channels to make smart purchases: The role of webrooming and showrooming. J. Retail. Consumer Services 52, 101923.

Graff Zivin, J., Song, Y., Tang, Q., Zhang, P., 2020. Temperature and high-stakes cognitive performance: Evidence from the national college entrance examination in China. J. Environ. Econ. Manag. 104, 102365.

He, P., Liu, P., Qiu, Y.L., Liu, L., 2022. The weather affects air conditioner purchases to fill the energy efficiency gap. Nature Commun. 13 (1), 5772.

Heyes, A., Saberian, S., 2019. Temperature and decisions: Evidence from 207,000 court cases. Am. Econ. J. Appl. Econ. 11 (2), 238–265.

Houde, S., 2018. How consumers respond to product certification and the value of energy information. Rand J. Econ. 49 (2), 453–477.

IEA, 2018. The Future of Cooling, Opportunities for energy- efficient air conditioning. Technical report, International Energy Agency, Paris.

ISTAT, 2014. I consumi energetici delle famiglie, anno 2013. Istituto Nazionale di Statistica, Rome, Italy.

ISTAT, 2018. Rilevazione sulle forze di lavoro. Istituto Nazionale di Statistica, Rome, Italy.

ISTAT, 2022a. I consumi energetici delle famiglie, anno 2021. Istituto Nazionale di Statistica, Rome, Italy.

ISTAT, 2022b. Multiscopo sulle famiglie: aspetti della vita quotidiana. Istituto Nazionale di Statistica, Rome, Italy.

Jaeger-Erben, M., Frick, V., Hipp, T., 2021. Why do users (not) repair their devices? A study of the predictors of repair practices. J. Clean. Prod. 286, 125382.

Kahneman, D., Tversky, A., 1979. Prospect theory: An analysis of decision under risk. Econometrica 47 (2), 263–291, Publisher: [Wiley, Econometric Society].

Kőszegi, B., Rabin, M., 2006. A model of reference-dependent preferences. Q. J. Econ. 121 (4), 1133–1165, Publisher: Oxford University Press.

Lamp, S., 2023. Sunspots that matter: The effect of weather on solar technology adoption. Environ. Resour. Econ. 84 (4), 1179–1219.

Liao, Y., 2020. Weather and the decision to go solar: Evidence on costly cancellations. J. Assoc. Environ. Resour. Econ. 7 (1), 1–33.

Litardo, J., Gomez, D., Boero, A., Hidalgo-Leon, R., Soriano, G., Ramirez, A.D., 2023. Air-conditioning life cycle assessment research: A review of the methodology, environmental impacts, and areas of future improvement. Energy Build. 296, 113415.

Loewenstein, G., O'Donoghue, T., Rabin, M., 2003. Projection bias in predicting future utility*. Q. J. Econ. 118 (4), 1209–1248.

Ministero dell'Economia e delle Finanze, 2019. Redditi delle persone fisiche per l'anno 2018.

Mistry, M.N., 2020. A high spatiotemporal resolution global gridded dataset of historical human discomfort indices. Atmosphere 11 (8), 835.

Qin, Y., Wu, J., Yan, J., 2019. Negotiating housing deal on a polluted day: Consequences and possible explanations. J. Environ. Econ. Manag. 94, 161–187.

Rapson, D., 2014. Durable goods and long-run electricity demand: Evidence from air conditioner purchase behavior. J. Environ. Econ. Manag. 68 (1), 141–160.

Simonsohn, U., 2010. Weather to go to college. Econ. J. 120 (543), 270–280.

Singer, J.D., Willett, J.B., 1993. It's about time: Using discrete-time survival analysis to study duration and the timing of events. J. Edu. Statist. 18, 155–195, Place: US Publisher: American Educational Research Assn.

Stango, V., Zinman, J., 2023. We are all behavioural, more, or less: A taxonomy of consumer decision-making. Rev. Econ. Stud. 90 (3), 1470–1498.