MIT Center for Energy and Environmental Policy Research

Working Paper Series

Temperature and Cognitive Performance

Benjamin Krebs

DECEMBER 2022

CEEPR WP 2022-019

Working Paper Series.

Since 1977, the Center for Energy and Environmental Policy Research (CEEPR) has been a focal point for research on energy and environmental policy at MIT. CEEPR promotes rigorous, objective research for improved decision making in government and the private sector, and secures the relevance of its work through close cooperation with industry partners from around the globe. Drawing on the unparalleled resources available at MIT, affiliated faculty and research staff as well as international research associates contribute to the empirical study of a wide range of policy issues related to energy supply, energy demand, and the environment.

An important dissemination channel for these research efforts is the MIT CEEPR Working Paper series. CEEPR releases Working Papers written by researchers from MIT and other academic institutions in order to enable timely consideration and reaction to energy and environmental policy research, but does not conduct a selection process or peer review prior to posting. CEEPR's posting of a Working Paper, therefore, does not constitute an endorsement of the accuracy or merit of the Working Paper. If you have questions about a particular Working Paper, please contact the authors or their home institutions.

Temperature and Cognitive Performance

Benjamin Krebs*

January 16, 2023

Abstract

I estimate the effect of temperature on cognitive performance in online brain training games. As this setting represents everyday cognitive tasks, the results are indicative of how temperature affects people on a daily basis. With rising average temperatures and more frequent extreme heat, a thorough understanding of this relationship is central. I find that, above a threshold, a 1°C increase in ambient air temperature leads to a performance reduction of 0.13%. The effect is mostly driven by individuals living in relatively cold areas, who are less adapted to hot temperatures.

Keywords: Air temperature, cognitive performance, climate change
JEL codes: Q54, J24

*Center for Energy and Environmental Policy Research, Massachusetts Institute of Technology, Cambridge, USA.
Contact: bkrebs@mit.edu
For comments and suggestions I am grateful to Benedikt Janzen, Michael Keller, Christopher Knittel, Benjamin Koch, Patrick Leisibach, Simon Luechinger, Margarida Matos, Matthew Neidell, Xibo Wan, and Sylvain Zueger.

1 Introduction

Climate change entails vast economic and social consequences. A recent body of research has identified adverse effects of extreme temperatures on economic growth and production (Dell et al. 2012; Burke et al. 2015), agricultural output (Schlenker and Roberts 2009; Lobell et al. 2011), labor productivity (Graff Zivin and Neidell 2014; Somanathan et al. 2021), mortality (Deschênes and Greenstone 2011; Barreca 2012; Mullins and White 2020; Heutel et al. 2021), conflict (Hsiang et al. 2013), migration (Missirian and Schlenker 2017), and others.[1] Under the prospect of rising average temperatures and more frequent extreme weather events in most places of the world, this literature naturally predicts an intensification of these effects.

One central aspect of any human activity is cognitive performance. It determines labor productivity and serves as a prerequisite for human capital accumulation. As extreme temperatures are a detrimental factor to the human body, it is essential to learn how they affect cognitive performance, and ultimately, what climate change means for this relationship. Not surprisingly, a recent body of literature (discussed below) investigates the link between temperature and human capital, learning, and cognitive performance.

In this paper, I estimate the effect of temperature on cognitive performance. I use data from an online mental arithmetic training game, called Raindrops, with more than 31,000 individuals and 1.15 million games played in 748 3-digit U.S. ZIP Codes between 2015 and 2019. The data include only paying subscribers, many of whom regularly train their mental arithmetic skills. This context provides a unique opportunity to investigate how temperature affects people in a familiar environment, representative of everyday situations. It is a setting that is currently not yet covered in the literature.

[1] See Carleton and Hsiang (2016) for an overview.

Previous studies have investigated performance either in college admission tests (or other academic exams) or in cognitive tests from surveys.[2] Test-taking students find themselves in a particularly stressful, non-everyday situation with potentially long-lasting effects on their career path. While understanding the role of environmental factors in this context is highly relevant, individuals are potentially much more sensitive to these factors when writing a test than they usually are. This limits the external validity for less stressful situations people encounter every day. Meanwhile, most individuals rarely participate in surveys. They are confronted with an unfamiliar setting, in which they play a more passive role. Again, susceptibility to temperature might be quite different compared to when performing more frequent tasks.

The individuals in this study engage in a non-stressful, but cognitively challenging task, and many do so very frequently. Thus, this paper fills a knowledge gap by analyzing a setting that represents everyday-life tasks. Adding this context to the literature is important, because most tasks people perform on their job and in everyday life are frequent tasks that are not as decisive as admission tests, and not as rare as cognitive tests in surveys. Therefore, the results I estimate are potentially more representative of how temperature affects cognitive performance in a broad range of tasks, especially in economies with a large service sector.

Two other data characteristics support the external validity of this study. First, my analysis includes all ages from 18 to 80 from all over the contiguous United States and, thus, covers a very broad societal range. The literature concerned with academic exams includes mostly adolescents. The average effects I estimate in this paper are, therefore, arguably more indicative of how temperature affects a general population. Moreover, I can look into effect heterogeneity between younger and older

[2] See a summary thereof below.

people. Second, as I observe a very broad temperature range[3], I can investigate both hot and cold temperatures in the same context. Previous studies analyzed only one end of temperature extremes, mostly heat.

Finally, as I observe the individuals' number of false entries per game, I add to our understanding of the temperature-cognition relationship by differentiating between problem-solving speed and the error rate. This is an important distinction as speed decline and error proneness have distinct ramifications. In many settings, errors are arguably more costly than speed, e.g., in medical procedures, or the assembly of consumer goods. If a higher error rate is the main driver of a lower cognitive performance, protecting people from heat exposure in settings where mistakes are costly seems particularly important. This paper is the first to make this distinction.

I find that hot temperatures significantly reduce mental arithmetic performance. I run both piecewise-linear regressions that allow for different slopes in two different temperature ranges and regressions with 3°C temperature bin indicators. Using the linear regression model, I find that, above a defined threshold, a 1°C increase in the average air temperature during the 24 hours preceding a play lowers the number of correct answers by 0.084, or 0.13%. The threshold value is 16.5°C and represents the arithmetic mean of the best-performance bin (15-18°C) from the bin regressions. Below the threshold, the corresponding coefficient is insignificant and close to zero (0.010).

The coefficients from the temperature bin-indicator regressions tend to confirm the assumed linearity of the temperature effect. They indicate that people attain significantly lower scores when playing in the bins above 21°C, compared to the bin with the highest average performance (15-18°C). The scores decrease by 0.484 (21-24°C), 0.588 (24-27°C), and 0.954 ($\geq$27°C) points. These figures amount to a

[3] The lowest percentile of temperature is -11.1°C, and the highest percentile is 31.5°C.

drop of 0.73%, 0.90%, and 1.46%, respectively. Similar to the linear regressions, I do not find consistent evidence for adverse effects of cold temperatures.

The results exhibit an important heterogeneity: Individuals from relatively cold ZIP Codes (below-median 2015-2019 average temperatures) experience a larger performance drop than individuals from relatively hot ZIP Codes (above-median 2015-2019 average temperatures). The estimate from the piecewise-linear regressions for the above-threshold range is -0.142 for cold ZIP Codes (-0.21%), but only -0.042 for hot ZIP Codes (-0.07%) and not statistically significant.

The baseline result implies that, with rising temperatures and no adaptation, people will perform below their capacity more often. Running separate analyses for cold and hot ZIP Codes gives insight into how adaptation to climate change might mitigate these adverse effects (Dell et al., 2014; Auffhammer, 2018). As the cited literature above shows, hotter regions are generally better adapted to heat. As temperatures rise, adaptation investments in colder regions will potentially close this gap. People living in colder regions might therefore react less to hot temperatures in the future, similar to people currently living in hotter regions. However, as climate change will also result in more extreme temperatures, ranges that are currently rare will occur more frequently. Even the better-adapted, hotter regions can be expected to experience larger performance drops. These two effects run in opposite directions. Predictions about how climate change affects our cognitive performance thus hinge on central assumptions about the degree of potential adaptation.

In an additional result, I provide evidence for effect accumulation: While a single hot day or two consecutive hot days – defined as the average temperature being greater than 21°C – only lead to minor reductions in the Raindrops score, more than two consecutive hot days cause consistent performance reductions, even in the hot ZIP Code sample. If temperature is above 21°C for a week, the average performance drop for the whole sample is -0.789, or -1.21%. This finding is of particular importance as climate models not only predict higher average temperatures but also more frequent heat waves (Keellings and Moradkhani, 2020).

Individuals using the brain training software chose when to do so. This raises concerns about bias from two types of potential selection issues, at the extensive and the intensive margin. First, people might be less likely to play when temperatures are high. This will only be an issue if more temperature-sensitive individuals (those who experience a larger performance drop due to extreme temperatures) are less likely to play than others when temperatures are extreme. In this case, the coefficients will be biased toward zero. While I cannot directly test this, I show that, on average, individuals do not have a lower probability to play on hot days. However, they do use the software more often when it is cold. Due to the potential selection issue at the lower end of temperatures, I only cautiously interpret the performance results for cold temperatures.

Second, people might play fewer times on particularly hot or cold days. As they improve their performance with the number of times they play on a day, the average score from a very hot (or very cold) day will be lower than from a mild day if they play fewer times due to the heat (or cold). This would be an issue independent of the selection of people who are affected. In this case, the coefficients would be biased downwards – my results would overestimate the true adverse effect. I show that neither hot nor cold temperatures seem to substantially affect the intensive margin.

This paper is closest related to recent literature on the effect of temperature on performance in either academic exams or survey tests. The seminal paper in this literature is Graff Zivin et al. (2018). They estimate the effect of temperature on assessments of cognitive ability from the National Longitudinal Survey of Youth, and find a decline in math performance for contemporaneous temperatures above 21°C (significant from 26°C), but not for reading. The authors also analyze long-term effects, namely by looking at temperature realizations between tests, which take place annually or biannually, and find smaller, imprecisely estimated effects. Garg et al. (2020) use data from two Indian children surveys. They show that the number of hot days in the year before the survey-based test negatively affects school-age children's math and reading scores and that reduced agricultural income is a main channel. Yi et al. (2021) employ survey data from the China Health and Retirement Longitudinal Study to find impaired math and verbal skills from short-term heat stress.

Studies analyzing college admission or other academic exams are Cho (2017) for college entrance exams in Korea, Graff Zivin et al. (2020) for China's National College Entrance Examination, Cook and Heyes (2020) for exams taken at the University of Ottawa, Park et al. (2020) for the PSATs in the U.S., Park (2022) for the Regents Exams in New York City, Park et al. (2021) for the PISA test in 58 countries and annual math and English tests in 12,000 U.S. school districts, Melo and Suzuki (2021) for the Exame Nacional do Enismo Mèdio in Brazil, and Roach and Whitney (2021) for standardized tests in the U.S.

With the exception of Cook and Heyes (2020), who focus on cold temperatures, all of these studies find hot temperatures to negatively affect students' test scores. The tested subjects include math (all papers), a first language (Graff Zivin et al., 2020; Park, 2022; Melo and Suzuki, 2021; Roach and Whitney, 2021), a foreign language (Cho, 2017; Graff Zivin et al., 2020), and other subjects, though most of the papers use data on overall scores only, not individual subjects. Cook and Heyes (2020) find a reduction in university exam performance with negative temperatures (in °C). To my knowledge, the only paper that does not use survey or academic test data is Bao and Fan (2020), who use Chinese data from an online role-playing game. As they only include data from March 2011 they focus mainly on cold temperatures.

2 Data

2.1 Cognitive performance

To measure cognitive performance, I use data from Lumosity's online brain training game Raindrops.[4] Individuals have to solve mental arithmetic problems that fall down in raindrops before they hit the water at the bottom of the screen (see Figure A1). The problems disappear when solved correctly. After the third raindrop has hit the water the game is over. The dataset includes 31,029 players and 1,151,059 plays from the years 2015 to 2019.[5] It only covers web plays (no mobile apps). The spatial resolution is 3-digit ZIP Codes of which the dataset contains 748.

The main outcome variable is the number of correct answers. In an extension, I also use the error rate, measured as the number of correct entries divided by the total number of entries. As people's performance heavily depends on their play behavior (e.g., how many times they have played, how long ago they last played), I include the following three control variables: the log number of plays an individual has played so far, the log number of plays an individual has played since taking a break of at least one hour, and the log number of days since the last play. Further, I interact each of the three variables with three age range indicators (50–64, 65–74, and 75 and older).[6]

[4] Krebs and Luechinger (2021) use the same data in a different context. As they provide an extensive data description, I only provide a short overview here. I follow their data-cleaning process.

[5] Note that these are not the same numbers as in Krebs and Luechinger (2021). The reason is that I do not lose observations from missing pollution values.

[6] Krebs and Luechinger (2021) use the same control variables.

2.2 Temperature and weather covariates

To generate the temperature and weather covariates I use data from the NOAA National Centers for Environmental Information (NCEI) Global Hourly Integrated Surface Database for the years 2015–2019.[7] I only include weather monitoring stations that were operational throughout the whole period. I include the variables air temperature, air dew point temperature, wind speed, atmospheric pressure, and precipitation. To generate ZIP Code averages, I exert the following three steps: First, for each station, I drop all variables that have more than 25% missing observations. Second, I interpolate missing values of all available variables at the station level with an inverse distance-weighted average of all stations within a radius of 50 kilometers and a power parameter of 2. And third, I generate the ZIP Code average from all stations within a ZIP Code. In case there is no station within a ZIP Code, I use the inverse distance-weighted average of all stations within 50 kilometers of the ZIP Code centroid, again with a power parameter of 2, to attribute the hourly weather variables. I drop variables of ZIP Codes that have more than 25% missing values.

The main explanatory variable is the average air temperature during the 24 hours preceding a play. In a robustness check, I use the average air temperature during the 48 hours preceding a play and the average heat index during the 24 hours preceding a play as the exogenous variable. The covariates are the 24-hour average of relative air humidity, wind speed, atmospheric pressure, and precipitation, and the quadratic function of each of those variables. The 24-hour averages include the current hour when the game was played. At least 18 hours (75%) have to be non-missing values, otherwise, the observation is coded as missing. I calculate the relative humidity from a function of air temperature and dew point temperature[8],

[7] https://www.ncei.noaa.gov/data/global-hourly/archive/

[8] https://bmcnoldy.rsmas.miami.edu/Humidity.html

and the average heat index according to the National Weather Service[9].

2.3 Summary statistics

Table 1 shows the means and standard deviations of the main two dependent variables (number of correct answers and the error rate), and the main independent variable (average air temperature during the 24 hours preceding a play). It also includes two alternatives for the main explanatory variable used as robustness checks (average air temperature during the 48 hours preceding a play and average heat index during the 24 hours preceding a play). The table has separate columns for the whole sample, the cold-ZIP Codes sample (below-median 2015-2019 average temperatures), and the hot-ZIP Codes sample (above-median 2015-2019 average temperatures). Figure A2 displays a map of the contiguous United States with different colors for cold and hot ZIP Codes.

The overall average number of correct answers is 65.4. People living in cold ZIP Codes, on average, score roughly 5 units higher than people living in hot ZIP Codes, which amounts to 8.3% more correct answers. The difference in the error rate is 5.4% (9.7% erroneous entries in hot ZIP Codes vs. 9.2% in cold ZIP Codes). Both these differences are statistically significant at the 99% level. The average temperature is 14.1°C. The difference between hot and cold ZIP Codes is 7.5°C. The heat index variable is slightly lower, mainly due to the cold ZIP Codes sample, which indicates a higher relative humidity in hot ZIP Codes.

[9] https://www.wpc.ncep.noaa.gov/html/heatindex_equation.shtml

Table 1 Summary statistics

Variable	All ZIP Codes	Cold ZIP Codes	Hot ZIP Codes
No. of correct answers	65.417	68.446	63.228
	(37.441)	(39.804)	(35.476)
Error rate	0.095	0.092	0.097
	(0.071)	(0.070)	(0.072)
Temperature past 24h in °C	14.076	9.748	17.204
	(9.940)	(10.109)	(8.545)
Temperature past 48h in °C	14.083	9.757	17.209
	(9.812)	(9.956)	(8.422)
Heat index past 24h in °C	13.429	8.614	16.907
	(11.105)	(11.155)	(9.679)
Observations	1,151,059	482,812	668,247

Notes: Means and standard deviations (in parentheses) of the two main cognitive outcome variables, and the temperature variables for the full sample, the cold-ZIP Codes sample (below-median 2015-2019 average temperatures), and the hot-ZIP Codes sample (above-median 2015-2019 average temperatures). "Temperature past 24h in °C", "temperature past 48h in °C", and "heat index past 24h in °C" are the average air temperature during the 24 hours preceding a play, the average air temperature during the 48 hours preceding a play, and the average heat index during the 24 hours preceding a play, respectively, all in °C.

3 Identification

I use two different approaches to estimate the effect of temperature on cognitive performance. Equation 1 represents the piecewise-linear regression model.

$$C_{izt} = \beta T_{jt} + \gamma A_{jt} + \delta(T_{jt} \times A_{jt}) + \boldsymbol{W'_{zt}}\eta + \boldsymbol{G'_{izt}}\theta + \boldsymbol{\iota}_i + \boldsymbol{\tau}_t + \varepsilon_{izt} \tag{1}$$

C_{izt} is cognitive performance of individual i in ZIP Code z at time t. T_{jt} is the average air temperature during the 24 hours preceding the play in °C. A_{jt} is an indicator equal to 1 if the air temperature during the past 24 hours was above a certain threshold value. $\boldsymbol{W'_{zt}}$ is a vector of weather variables described in Section 2.2. $\boldsymbol{G'_{izt}}$ is a vector of play covariates described in Section 2.1. $\boldsymbol{\iota}_i$ absorbs the individual effects to control for any time-invariant differences between individuals, such as the innate ability. $\boldsymbol{\tau}_t$ absorbs year, month-of-year, day-of-week, hour-of-day effects. These fixed effects flexibly control for time trends, seasonal patterns, and differences across the timing of a play between weekdays and between the hours of a day. Finally, ε_{izt} is the error term.

The two coefficients of interest are β and δ. β represents the effect of temperature on cognitive performance below the threshold value. To investigate the effect of temperature on cognitive performance above the threshold value, I run a linear combination test for $\beta + \delta \neq 0$. I report the results in the regression tables. The threshold value is 16.5°C and is determined as the arithmetic mean of the temperature bin with the highest performance from the indicator regression (Equation 2).

A major advantage of Equation 1 is that the coefficients report average marginal effects over the full temperature range above and below the threshold value. However, the model hinges on the linearity assumption. Therefore, I also estimate the temperature bin-indicator regression model represented in Equation 2.

$$C_{izt} = \boldsymbol{T'_{jt}}\theta + \boldsymbol{W'_{zt}}\eta + \boldsymbol{G'_{izt}}\theta + \boldsymbol{\iota}_i + \boldsymbol{\tau}_t + \varepsilon_{izt} \tag{2}$$

$\boldsymbol{T'_{jt}}$ is a vector of temperature bin indicators. The temperature is, as in Equation 1, the average air temperature during the 24 hours preceding the play in °C. The bins correspond to 3°C temperature bins from 0 to 27, plus one bin for temperatures below 0, and one for temperatures above 27. I chose < 0 as the bottom bin as 0 is the (rounded) 10^{th} percentile of the temperature distribution (with the exact value being 0.47). Likewise, 27 is the closest (rounded) number dividable by 3 to the 90^{th} percentile (with the exact value being 26.1). I define the bin with the highest average performance in the regression as the reference.

The model in Equation 2 flexibly allows for non-linearity in the effect of temperature on cognitive performance. It is therefore a suited complement to Equation 1 to test the linearity assumption. A drawback of this approach is that the highest bin might be statistically significantly different from other bins by chance. This would make it seem like people perform worse in all other bins even if this was not the case. As the linear decline with higher temperatures in Section 4.1 does not support this concern, I refer to Equation 1 as my main results.

4 Results

4.1 Performance

Table 2 shows the coefficients from the piecewise-linear regression model outlined in Equation 1. The coefficients in the first row ("Air temperature") correspond to β.[10] The linear combination test in the second row ("Air temp. + air temp. × above threshold") reports the effect of air temperature on performance when temperature is above the threshold. This corresponds to a test of $\beta + \delta \neq 0$ from Equation 1. I do not report the coefficient from the air temperature × above threshold interaction, δ, as it is simply equal to the coefficient reported in the linear combination test minus β.

When focusing on all ZIP Codes included in my dataset (column 1), cold temperatures do not seem to affect people's performance in the Raindrops game. The coefficient of air temperature below the threshold value is positive but very close to zero (0.010) and statistically insignificant. Above the threshold, temperatures negatively affect performance. An increase of 1°C decreases the number of correct answers by 0.084, which amounts to 0.13%.

To evaluate potential adaptation effects, I run separate regressions for relatively

[10] Note that air temperature is measured in °C and, if not indicated otherwise, is calculated as the average air temperature during the 24 hours preceding the play. The threshold temperature is 16.5°C.

cold and relatively hot ZIP Codes.[11] Columns 2 and 3 of Table 2 report the respective results. They show that it is the cold ZIP Codes that drive the overall result. In the cold-ZIP Codes sample, a temperature increase of 1°C lowers the number of correct answers by 0.142, which is 0.21%. The effect for hot ZIP Codes is -0.042 (0.066%) and statistically insignificant.

Table 2 Air temperature and cognitive performance: piecewise-linear regressions

Dependent variable: No. of correct answers	All ZIP Codes	Cold ZIP Codes	Hot ZIP Codes
Air temperature	0.010	0.010	0.012
	(0.021)	(0.017)	(0.017)
Linear combination test:	-0.084**	-0.142***	-0.042
Air temp. + air temp. × above theshold	(0.038)	(0.048)	(0.030)
Observations	1,151,059	482,812	668,247
Individuals	31,029	12,708	18,321
ZIPs	748	374	374

Notes: Results from piecewise-linear regressions of the number of correct answers on the average air temperature during the 24 hours preceding the play (in °C), and an interaction with the above-threshold indicator. The standard errors (in parentheses) are clustered on ZIP Codes. "Air temperature" denotes the coefficient below the threshold, and the linear combination test amounts to the slope above the threshold. The threshold is 16.5°C. The regression in column 1 includes all observations. Columns 2 and 3 show the results from separate regressions for the cold-ZIP Codes sample (below-median 2015-2019 average temperatures), and the hot-ZIP Codes sample (above-median 2015-2019 average temperatures). The control variables include individual effects, time effects, play controls, weather controls, and an indicator for above threshold temperature (see Section 3).

This result suggests that hot ZIP Codes are at present better equipped to cope with high temperatures. However, as colder regions' temperature distributions shifts towards what hotter regions experience today, individuals living in colder regions might adapt and, therefore, become less susceptible to hotter temperatures than they currently are. Concurrently, temperatures that have been rare in the past

[11] I differentiate based on the ZIP Code mean temperature over the sample period of 2015 to 2019. I define cold (hot) ZIP Codes as the ones that had a below-median (above-median) average temperature during these years.

will become more common. As discussed in the introduction, these opposing effects hinder a thorough analysis of future developments.

The temperature bin-indicator regressions broadly confirm the results from the piecewise-linear regressions. Panel A of Figure 1 (equivalent to column 1 of Table A1) shows that the coefficients for the temperature bins below the reference bin of 15-18°C are mostly small and insignificant, with two of them being significant at the 90% level. The coefficient from the 18-21 bin is almost zero. Above that bin, there seems to be a linear trend with the coefficients and t-values becoming larger in absolute terms. These results tend to support the linearity assumption made in Equation 1.

Panel B of Figure 1 (equivalent to columns 2 and 3 of Table A1) confirms the findings from the cold-hot differential piecewise-linear regressions. There are no significant effects for below-reference temperatures, except for one outlier (6-9°C) for cold ZIP Codes. This outlier does not withstand the robustness checks I present in Section 4.3. All bins above 21°C are statistically significant at the 95% level for cold ZIP Codes. They are roughly twice the size (in absolute terms) as the coefficients from all ZIP Codes and seem to follow a linear trend as well. The results from the hot-ZIP Codes regressions are not significant, except for the top bin, which is significant at the 95% level. While there is a slight downward trend above 19.5°C, it seems likely that the 18-21°C bin is an upward outlier.

How do these results compare to the results from the literature? In Table A2 I summarize the main results from the relevant papers. I include all papers discussed in Section 1 that investigate short-term effects and use a linear approach. The table shows the temperature variable, the outcome variable, the temperature cutoff, the mean of the outcome variable, and the estimated coefficient. From the latter two, I calculate the percentage change of the outcome variable in percent for a 1°C increase

Figure 1 Air temperature and cognitive performance: 3°C-bins regressions

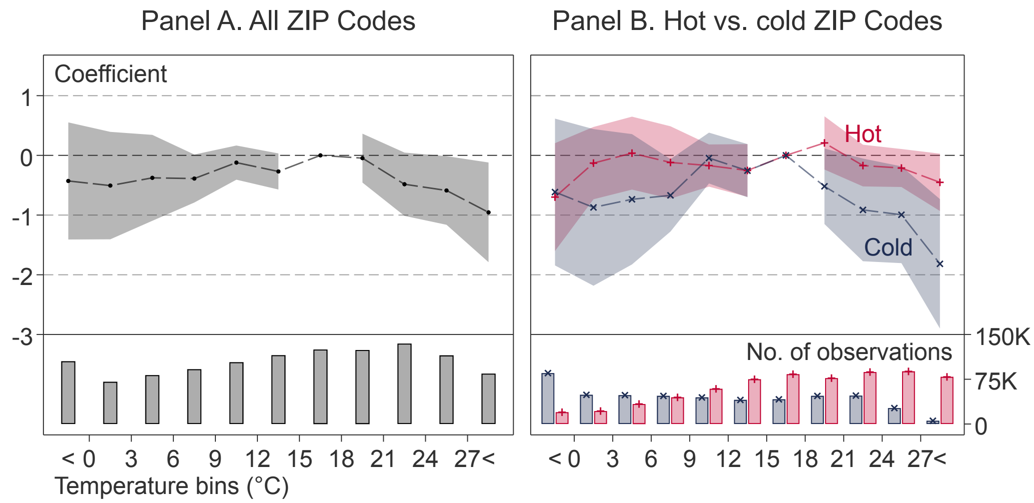

Notes: Coefficients with 95% confidence intervals (left y-axis), and number of observations (right y-axis) from regressions of the number of correct answers on 3°C-bin indicators of the average air temperature during the 24 hours preceding the play (x-axis) as presented in Table A1. The standard errors are clustered on ZIP Codes. The reference bin is 15-18°C. The regression in Panel A includes all observations. Panel B shows the results from separate regressions for the cold-ZIP Codes sample (below-median 2015-2019 average temperatures), and the hot-ZIP Codes sample (above-median 2015-2019 average temperatures). The control variables include individual effects, time effects, play controls, and weather controls (see Section 3).

in temperature (semi-elasticity) and display them in the last column. This makes comparisons across papers more convenient.

The five papers that meet the mentioned requirements all estimate larger effects than my main result, by a magnitude of 3 to 13. What might be a possible explanation for this? In contrast to previous studies using assessment tests, the stakes in brain training games are low. People might be less affected by heat in non-stressful, everyday settings. Künn et al. (2019) report similar findings for particulate matter pollution, another environmental stressor, in the context of cognitive performance.

Another potential explanation is age. Most of the discussed studies focus on adolescents, while the median age in this study is 56 years and ranges from 18 to 80 years. As there are only relatively few adolescents in my dataset, I lack the power to test for differential effects between such small age groups. Yet, in an extension in

Section 4.3 I estimate effect heterogeneity across two age groups (below and above the median age of 56 years) and find no evidence for differences between those groups.

I see two main limitations to interpretations of why I estimate smaller effects than the previous literature. First, all studies differ in a multitude of dimensions. E.g., the choice of temperature cutoff might affect the estimated slope. And second, some of the previous studies focus on a relatively constrained location. These locations might be better or worse adapted to heat, compared to the average effect across the contiguous United States I estimate in this paper.

4.2 Selection

One central difference to the previous literature is that individuals using the brain training software chose when to play. This is in contrast to college admission exams that have a predetermined date, or telephone surveys that come unexpected to study participants. In this section, I provide two selections test mentioned in Section 1.

The first test concerns the extensive margin. I construct a dataset that includes an observation for every hour from each individual's first to last play in the data. The observations contain the respective ZIP Code's weather variables and the individual's play covariates. The main outcome variable is an indicator equal to 1 if an individual played during a specific hour, and 0 otherwise. Table A3 shows the summary statistics for this dataset. As mentioned, I am not able to test whether susceptible individuals are more or less likely to refrain from playing than others on a particularly hot or cold day. Instead, I test whether people, on average, are more (or less) likely to play depending on the temperature.

Figure A3 (equivalent to Table A5) shows the results from the bin regressions similar to the main result. Note that I multiply the outcome variable by 100 to make the coefficients more readable. As the bin regressions imply a different underlying functional form compared to the main results, I do not estimate a linear model. The results in Panel A of Figure A3 for the whole sample suggest that temperatures do not affect the probability of playing, except for the coldest bins. The coefficient of 0.0107 for temperatures below 0°C means that people have a 3.06% higher probability of playing (0.0107 / 100 / 0.0035), compared to the reference bin (15-18°C). Separating between cold and hot ZIP Codes, I find no effects for hot temperatures either. The positive effect of cold temperatures seems to kick in earlier in cold ZIP Codes.

These findings suggest that there is no extensive margin selection for hot temperatures. However, people are more likely to play when it is cold. This could potentially bias the results if, e.g., it is mostly non-susceptible people who are more likely to play during cold weather. In that case, a potential negative effect of cold temperatures would be biased toward zero. Therefore, I interpret the null result from Section 4.1 with caution.

To address intensive-margin selection, I construct a dataset that has an observation for every hour during which an individual used the software. As the dependent variable, I use the log number of plays for each of these observations.[12] Thus, a coefficient of 0.01 corresponds to a 1% increase in the number of plays. This allows me to test whether people play more (or less) often, conditional on playing at all, when temperatures are extreme. The summary statistics for this dataset are in Table A4.

Figure A4 (equivalent to Table A6) shows the 3°C-bins regression results. As for the extensive margin I do not estimate the linear model, as the bin regressions do not support a corresponding functional form. The results do not provide consistent evidence for a change in the number of plays for hot temperatures. While the

[12]Note that the number of plays is at least 1.

coefficient for the 18-21°C bin is significantly negative, purely driven by the hot ZIP Codes, there is no trend and the pattern indicates that this is likely an outlier. People in cold ZIP Codes seem to play somewhat more during cold weather, but the extent is very small. All in all, I find little evidence for selection at the intensive margin.

4.3 Robustness

To test the robustness of my estimates, I run three alternations of my baseline results: First, instead of averaging temperature over 24 hours, I take the average of air temperature during the 48 hours preceding each play. The degree to which past temperatures affect people's cognitive performance has an underlying function that is ex-ante unclear. A 24-hour cutoff is obviously random and neglects that temperatures from longer ago might still affect people's performance. If temperatures in the 24-to-48-hours range do not affect cognitive performance, the coefficients should be closer to zero than in the main results regression, due to attenuation bias.

The results are robust to using the 48-hours average. Table A7 shows the linear regression results, and Figure A5 (equivalent to Table A8) the coefficients from the bin regressions. The coefficients for the linear combination test are somewhat larger in absolute terms for the whole sample (-0.093 vs. -0.084) as well as for the cold-ZIP Codes sample (-0.177 vs. -0.142). This suggests that temperatures in the 24-to-48-hour range are indeed relevant.

In the second robustness check, I use a heat index instead of air temperature. The effect of hot temperatures on the human body varies with relative humidity[13]. The heat index, a function of temperature and relative humidity, takes that into account directly, rather than solely controlling for relative humidity.

[13] https://www.weather.gov/ama/heatindex

Again, the results stay qualitatively unchanged. I present the piecewise-linear regression results in Table A9, together with the bin regression results in Figure A6 (equivalent to Table A10). The coefficient from the cold-ZIP Code sample is somewhat smaller (-0.117 vs. -0.142), while the coefficient from the hot-ZIP Code sample is slightly larger (-0.050 vs. -0.042), both in absolute terms.

Finally, I use the natural logarithm of 1 plus the number of correct answers as the dependent variable. As temperature effects on the number of correct answers might depend on individuals' baseline performance or average score, this specification would be a more suited functional form.

Again, there are no qualitative changes to the results. The coefficients for the linear combination test in Table A11 can be interpreted as percentage changes for a 1°C change in temperature above the threshold of 16.5°C. The figures are smaller than in Section 4.1 (-0.08% vs. 0.13% for the whole sample, -0.11% vs. -0.21% for the cold-ZIP Codes sample, and -0.06% vs. 0.07% for the hot-ZIP Codes sample). While the cold-ZIP Codes sample coefficient is less precisely estimated, the precision of the hot-ZIP Codes sample coefficient increases. The bin regression results in Figure A7 (equivalent to Table A12) look fairly similar as well.

4.4 Extensions

The main finding of this paper raises the question about the underlying mechanism. Do people simply solve the arithmetic problems more slowly and, thus, achieve fewer points, or do they make more mistakes and, thereby, lose time to find the right answer?

One way to investigate these mechanisms is to look at the error rate, the number of erroneous entries per total number of entries. Figure 2 (equivalent to Table A14) reports the bin regressions results. None of the estimated bin indicators returns a significant coefficient and there seems to be no trend, neither for cold nor for hot temperatures. This is confirmed in the liner regression results in Table A13. These results provide suggestive evidence that hot temperatures affect the solving speed, rather than the error rate.

Figure 2 Air temperature and error rate: 3°C-bins regressions

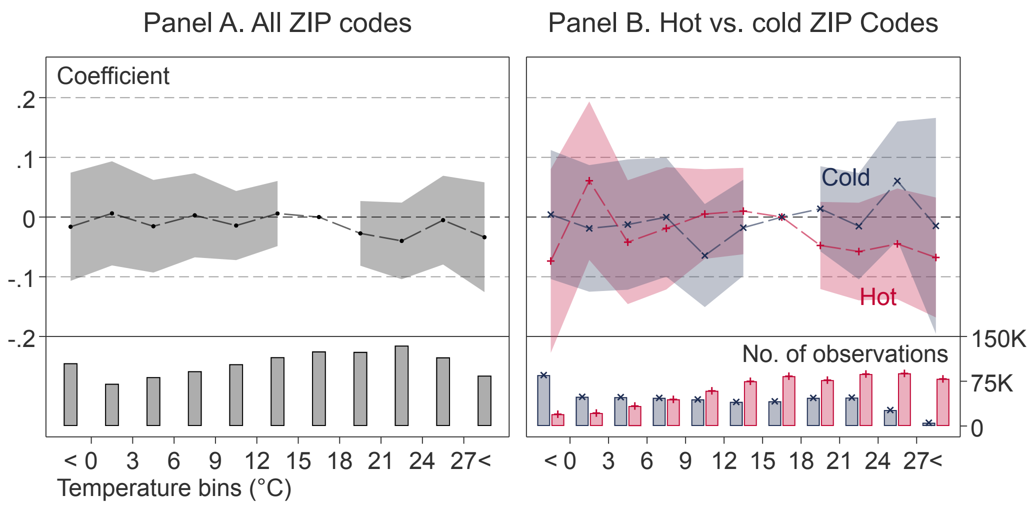

Notes: Coefficients with 95% confidence intervals (left y-axis), and number of observations (right y-axis) from regressions of the error rate (number of incorrect answers / total answers) on 3°C-bin indicators of the average air temperature during the 24 hours preceding the play (x-axis) as presented in Table A14. The standard errors are clustered on ZIP Codes. The reference bin is 15-18°C. The regression in Panel A includes all observations. Panel B shows the results from separate regressions for the cold-ZIP Codes sample (below-median 2015-2019 average temperatures), and the hot-ZIP Codes sample (above-median 2015-2019 average temperatures). The control variables include individual effects, time effects, play controls, and weather controls (see Section 3).

Another central aspect is how cumulative heat exposure affects cognitive performance. Specifically, is ongoing heat for multiple days worse than a single heat day alone?

Figure 3 (equivalent to Table A15) summarize the results of different heat period lengths. I run regressions similar to Equation 2, except that I use only one dummy variable equal to 1 if the average air temperature during the 24 hours preceding a play is greater than 21°C (baseline). This is the lower end of the bin with a negative coefficient at the 90% level in the main results. The coefficients for this dummy are -0.489 (all ZIP Codes), -0.666 (cold ZIP Codes), and -0.263, (hot ZIP Codes), all significant at the 95% level. These numbers correspond to a performance reduction of 0.75%, 0.97%, and 0.42%, respectively.

To disentangle, I calculate the average temperature of seven different time periods: period 1 is hours 0 to 23 preceding a play (the same as above and throughout the paper), period 2 is hours 24 to 47 preceding a play, etc. Accordingly, period 7 is hours 144 to 167. I then generate three indicator variables based on the average temperature of these seven periods. The first indicator is equal to 1 if the average temperature in period 1, but not 2, or in periods 1 and 2, but not 3 was greater than 21°C. This indicator refers to hot temperatures for one or two days, but not longer ("1-2 days..." in Figure 3). The second indicator is equal to 1 if the average temperature in periods 1, to 3, but not 4, or in periods 1 to 4, but not 5, or in periods 1 to 6, but not 7 was greater than 21°C. This indicator refers to hot temperatures for at least three but not more than six days ("3-6 days..." in Figure 3). Finally, the third indicator is equal to 1 if the average temperature in all seven periods was greater than 21°C. This indicator refers to hot temperatures for at least seven days in a row. I then run similar regressions to the baseline, but with these three heat period length indicators instead of just one indicator. The reference is a day with an average temperature below or exactly 21°C.

The estimated coefficient strictly increases with the length of the heat period, from -0.220 (-0.34%) for 1 or 2 days, to -0.568 (-0.87%) for 3 to 6 days, and -0.789 (-1.21%) for 7 or more days with average temperatures above 21°C. This pattern is consistent for both cold and hot ZIP Codes. The performance drop is -0.413 (-0.60%), -0.857 (-1.25%), and -1.075 (-1.57%) for cold ZIP Codes, and -0.014 (-0.02%), -0.27 (-0.43%), and -0.479 (-0.76%) for hot ZIP Codes, respectively.

While these coefficients are not statistically different from each other, all coefficients increase with the length of the heat period.

Figure 3 Effect accumulation

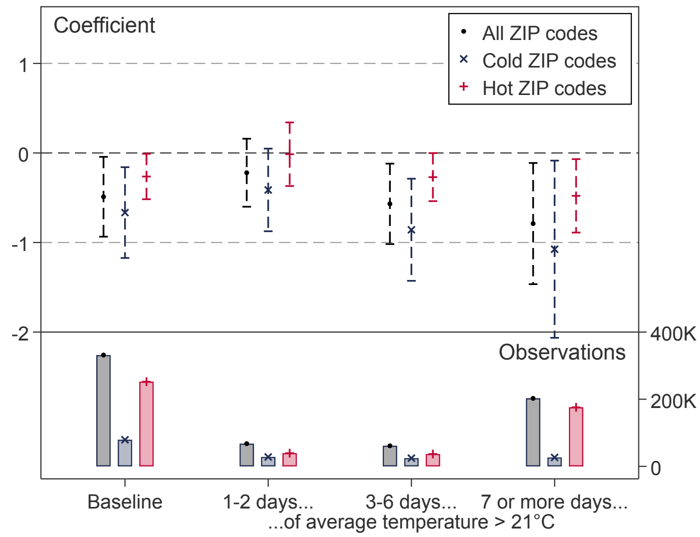

Notes: Coefficients with 95% confidence intervals (left y-axis), and number of observations (right y-axis) from regressions of the number of correct answers on and indicator = 1 if the average temperature was above 21°C during different temporal periods before a play (x-axis) as presented in Table A15. The standard errors are clustered on ZIP Codes. The baseline is 24 hours preceding a play. The reference is a day with an average temperature below or exactly 21°C. I run separate regressions for all observations, the cold-ZIP Codes sample (below-median 2015-2019 average temperatures), and the hot-ZIP Codes sample (above-median 2015-2019 average temperatures). The control variables include individual effects, time effects, play controls, and weather controls (see Section 3).

Analyzing the channels for this potential increase in the effect of heat on cognitive performance is beyond the scope of this paper. At least two mechanisms could be at play: First, a recent paper shows that people sleep less when temperatures are high (Minor et al., 2022). The observed pattern would be what to expect if the lack of sleep is a principal reason for reduced cognitive performance, and continuing sleep deprivation worsens cognitive performance. The second channel is more mechanical: Buildings take time to heat up. Without air conditioning, homes will usually be hotter after more days of heat. Hence, the effect might worsen with consecutive hot days.

The final extension I provide is on effect heterogeneity across ages. I estimate separate regressions for people who are below, or exactly, 56 years old, and people who are above 56 years old. 56 years is the median age. Table A16 shows the piecewise-linear results, Figure A8 and Table A17 show the temperature bins results, respectively. While the coefficients from the piecewise-linear regressions reveal a slightly larger effect for older people, the bins regressions coefficients are somewhat higher (in absolute terms) for younger people. The differences are generally very small and the coefficients are close to the baseline results for both age groups. They are, however, less precisely estimated. These results do not provide any evidence for effect heterogeneity.

5 Conclusion

I estimate the effect of temperature on cognitive performance for a broad population in the United States using a rich dataset on mental arithmetic from a brain training software. I find that hot temperatures reduce people's performance, with larger effects in colder, less adapted regions. The driver for the lower performance seems to be slower problem-solving, rather than higher error proneness. I do not find any significant effects of cold temperatures. The results for the cold-temperatures range should be taken with a grain of salt, as there are potential selection issues.

This paper fills a gap in the literature by focusing on a setting that is more representative of everyday-live situations and studying a broader population than the previous literature. The estimated coefficients are small, compared to studies that apply a similar approach in an educational or survey setting. Understanding the drivers behind temperature sensitivity in cognitively demanding settings seems like great potential for future research.

References

Auffhammer, Maximilian (2018). Quantifying economic damages from climate change. *Journal of Economic Perspectives 32*(4), 33–52.

Bao, Xiaojia and Qingliang Fan (2020). The impact of temperature on gaming productivity: Evidence from online games. *Empirical Economics 58*(2), 835–867.

Barreca, Alan I. (2012). Climate change, humidity, and mortality in the United States. *Journal of Environmental Economics and Management 63*(1), 19–34.

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

Carleton, Tamma A. and Solomon M. Hsiang (2016). Social and economic impacts of climate. *Science 353*(6304).

Cho, Hyunkuk (2017). The effects of summer heat on academic achievement: A cohort analysis. *Journal of Environmental Economics and Management 83*, 185–196.

Cook, Nikolai and Anthony Heyes (2020). Brain freeze: Outdoor cold and indoor cognitive performance. *Journal of Environmental Economics and Management 101*, 102318.

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

Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken (2014). What do we learn from the weather? The new climate-economy literature. *Journal of Economic Literature 52*(3), 740–98.

Deschênes, Olivier and Michael Greenstone (2011). Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the US. *American Economic Journal: Applied Economics 3*(4), 152–85.

Garg, Teevrat, Maulik Jagnani, and Vis Taraz (2020). Temperature and human capital in India. *Journal of the Association of Environmental and Resource Economists 7*(6), 1113–1150.

Graff Zivin, Joshua, Solomon M. Hsiang, and Matthew Neidell (2018). Temperature and human capital in the short and long run. *Journal of the Association of Environmental and Resource Economists 5*(1), 77–105.

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

Graff Zivin, Joshua, Yingquan Song, Qu Tang, and Peng Zhang (2020). Temperature and high-stakes cognitive performance: Evidence from the national college entrance examination in China. *Journal of Environmental Economics and Management 104*, 102365.

Heutel, Garth, Nolan H. Miller, and David Molitor (2021). Adaptation and the mortality effects of temperature across us climate regions. *Review of Economics and Statistics 103*(4), 740–753.

Hsiang, Solomon M., Marshall Burke, and Edward Miguel (2013). Quantifying the influence of climate on human conflict. *Science 341*(6151), 1235367.

Keellings, David and Hamid Moradkhani (2020). Spatiotemporal evolution of heat wave severity and coverage across the United States. *Geophysical Research Letters 47*(9), e2020GL087097.

Krebs, Benjamin (2022). Temperature and cognitive performance. *Unpublished manuscript*.

Krebs, Benjamin and Simon Luechinger (2021). Air pollution, cognitive performance, and the role of task proficiency. *SSRN Working Paper No. 3947149*.

Künn, Steffen, Juan Palacios, and Nico Pestel (2019). Indoor air quality and cognitive performance.

Lobell, David B., Wolfram Schlenker, and Justin Costa-Roberts (2011). Climate trends and global crop production since 1980. *Science 333*(6042), 616–620.

Melo, Ana Paula and Mizuhiro Suzuki (2021). Temperature, effort, and achievement: Evidence from a large-scale standardized exam in Brazil. *Unpublished manuscript*.

Minor, Kelton, Andreas Bjerre-Nielsen, Sigga Svala Jonasdottir, Sune Lehmann, and Nick Obradovich (2022). Rising temperatures erode human sleep globally. *One Earth 5*(5), 534–549.

Missirian, Anouch and Wolfram Schlenker (2017). Asylum applications respond to temperature fluctuations. *Science 358*(6370), 1610–1614.

Mullins, Jamie T. and Corey White (2020). Can access to health care mitigate the effects of temperature on mortality? *Journal of Public Economics 191*, 104259.

Park, R. Jisung (2022). Hot temperature and high-stakes performance. *Journal of Human Resources 57*(2), 400–434.

Park, R. Jisung, A. Patrick Behrer, and Joshua Goodman (2021). Learning is inhibited by heat exposure, both internationally and within the United States. *Nature Human Behaviour 5*(1), 19–27.

Park, R. Jisung, Joshua Goodman, Michael Hurwitz, and Jonathan Smith (2020). Heat and learning. *American Economic Journal: Economic Policy 12*(2), 306–339.

Roach, Travis and Jacob Whitney (2021). Heat and learning in elementary and middle school. *Education Economics*, 1–18.

Schlenker, Wolfram and Michael J. Roberts (2009). Nonlinear temperature effects indicate severe damages to us crop yields under climate change. *Proceedings of the National Academy of sciences 106*(37), 15594–15598.

Somanathan, Eswaran, Rohini Somanathan, Anant Sudarshan, and Meenu Tewari (2021). The impact of temperature on productivity and labor supply: Evidence from Indian manufacturing. *Journal of Political Economy 129*(6), 1797–1827.

Yi, Fujin, Tianhao Zhou, Luyun Yu, Bruce McCarl, Yanjun Wang, Fei Jiang, and Yafeng Wang (2021). Outdoor heat stress and cognition: Effects on those over 40 years old in China. *Weather and Climate Extremes 32*, 100308.

Appendix. Figures

Figure A1 The Raindrops game

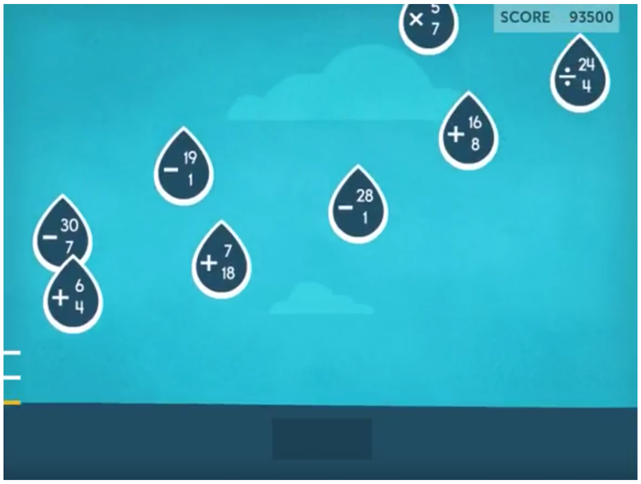

Source: Krebs and Luechinger (2021)

Figure A2 Hot and cold 3-digit ZIP Codes

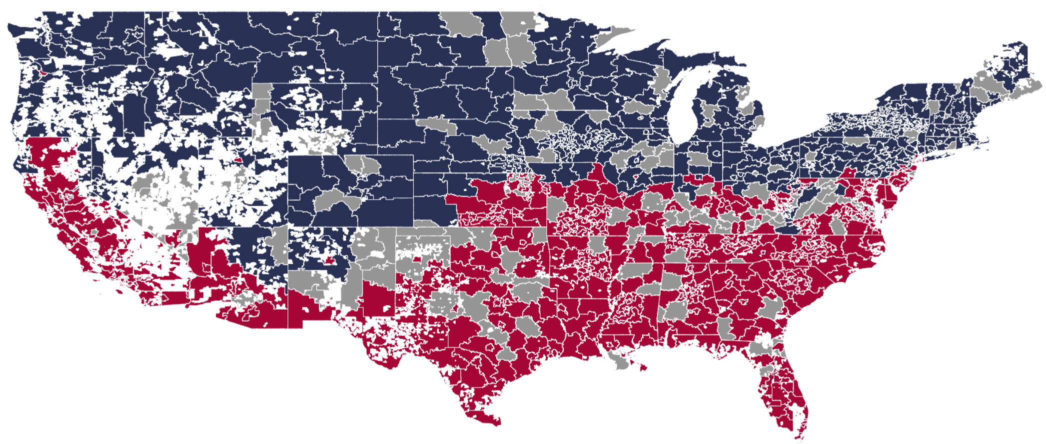

Notes: Map of 3-digit ZIP Codes of the contiguous United States with the hot-ZIP Codes sample (above-median 2015-2019 average temperatures) in red and the cold-ZIP Codes sample (below-median 2015-2019 average temperatures) in blue. Gray areas are ZIP Codes without a play observation. The average temperature of each ZIP Code is based on the NOAA data described in Section 2.2.

Figure A3 Air temperature and probability of playing: 3°C-bins regressions

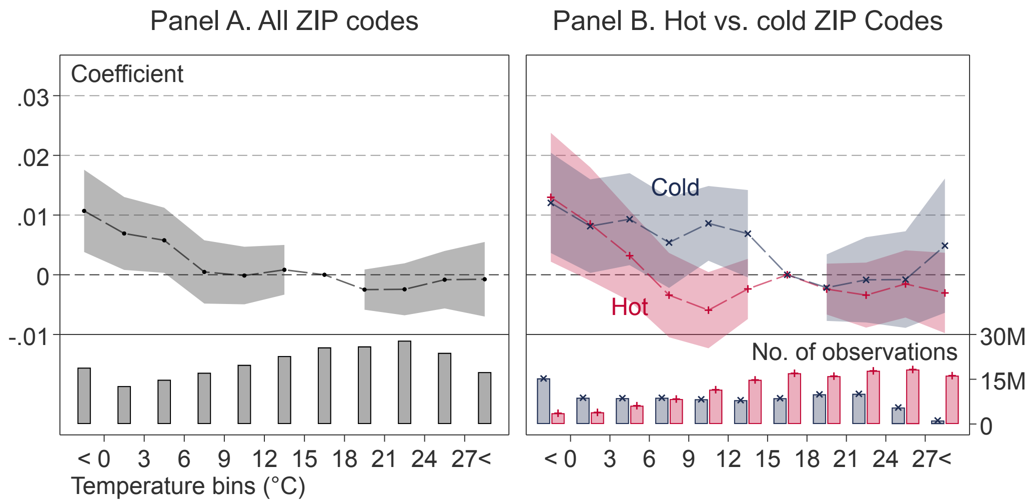

Notes: Coefficients with 95% confidence intervals (left y-axis), and number of observations (right y-axis) from regressions of play (indicator variable = 100 if an individual played during a specific hour) on 3°C-bin indicators of the average air temperature during the 24 hours preceding the play (x-axis) as presented in Table A5. The standard errors are clustered on ZIP Codes. The reference bin is 15-18°C. The regression in Panel A includes all observations. Panel B shows the results from separate regressions for the cold-ZIP Codes sample (below-median 2015-2019 average temperatures), and the hot-ZIP Codes sample (above-median 2015-2019 average temperatures). The control variables include individual effects, time effects, play controls, and weather controls (see Section 3).

Figure A4 Air temperature and frequency of playing: 3°C-bins regressions

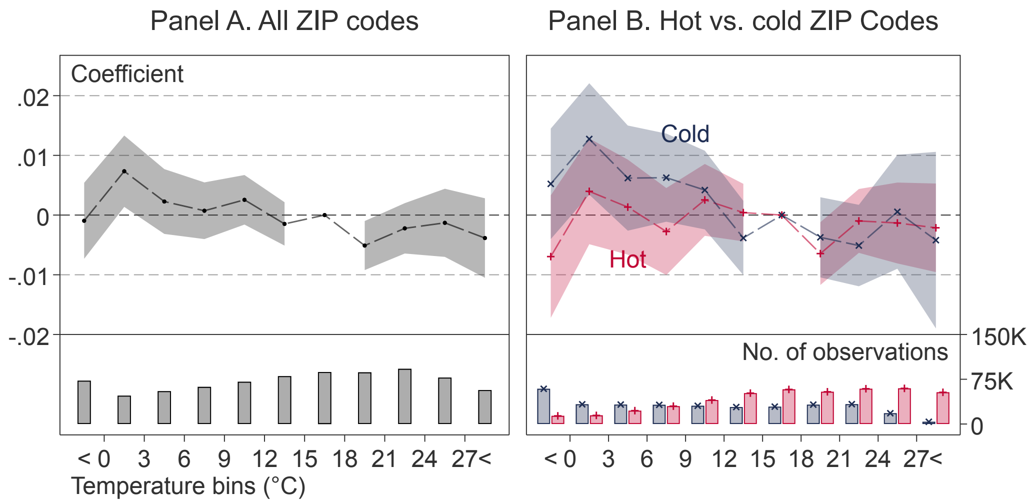

Notes: Coefficients with 95% confidence intervals (left y-axis), and number of observations (right y-axis) from regressions of the log number of plays (the log number of plays an individual engaged in during the current hour) on 3°C-bin indicators of the average air temperature during the 24 hours preceding the play (x-axis) as presented in Table A6. The standard errors are clustered on ZIP Codes. The reference bin is 15-18°C. The regression in Panel A includes all observations. Panel B shows the results from separate regressions for the cold-ZIP Codes sample (below-median 2015-2019 average temperatures), and the hot-ZIP Codes sample (above-median 2015-2019 average temperatures). The control variables include individual effects, time effects, play controls, and weather controls (see Section 3).

Figure A5 Air temperature past 48 hours and cognitive performance: 3°C-bins regressions

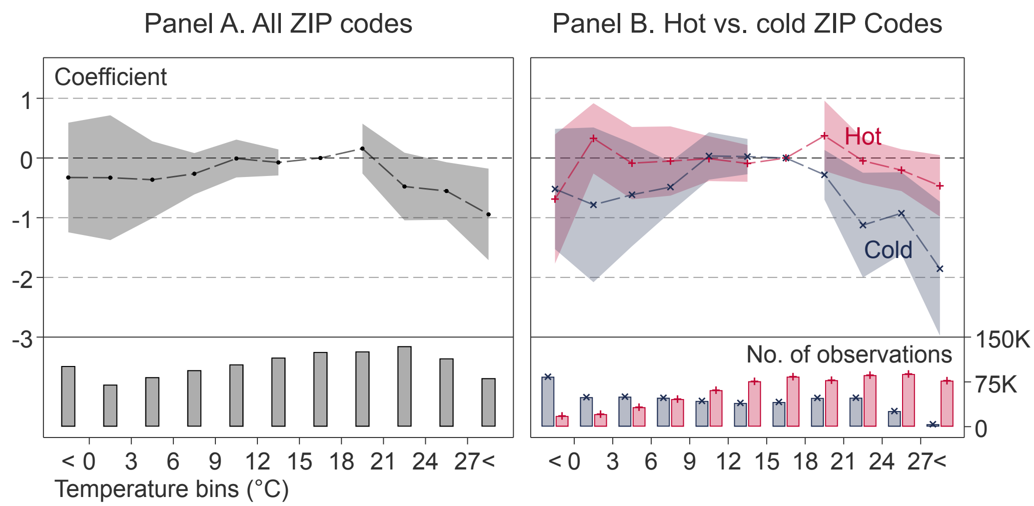

Notes: Coefficients with 95% confidence intervals (left y-axis), and number of observations (right y-axis) from regressions of the number of correct answers on 3°C-bin indicators of the average air temperature during the 48 hours preceding the play (x-axis) as presented in Table A8. The standard errors are clustered on ZIP Codes. The reference bin is 15-18°C. The regression in Panel A includes all observations. Panel B shows the results from separate regressions for the cold-ZIP Codes sample (below-median 2015-2019 average temperatures), and the hot-ZIP Codes sample (above-median 2015-2019 average temperatures). The control variables include individual effects, time effects, play controls, and weather controls (see Section 3).

Figure A6 Heat index and cognitive performance: 3°C-bins regressions

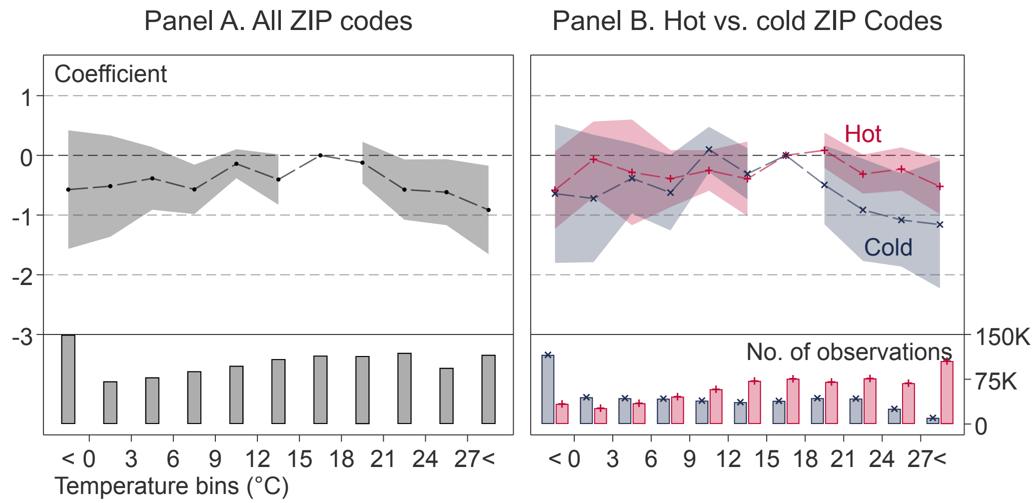

Notes: Coefficients with 95% confidence intervals (left y-axis), and number of observations (right y-axis) from regressions of the number of correct answers on 3°C-bin indicators of the average heat index during the 24 hours preceding the play (x-axis) as presented in Table A10. The standard errors are clustered on ZIP Codes. The reference bin is 15-18°C. The regression in Panel A includes all observations. Panel B shows the results from separate regressions for the cold-ZIP Codes sample (below-median 2015-2019 average temperatures), and the hot-ZIP Codes sample (above-median 2015-2019 average temperatures). The control variables include individual effects, time effects, play controls, and weather controls (see Section 3).

Figure A7 Air temperature and log cognitive performance: 3°C-bins regressions

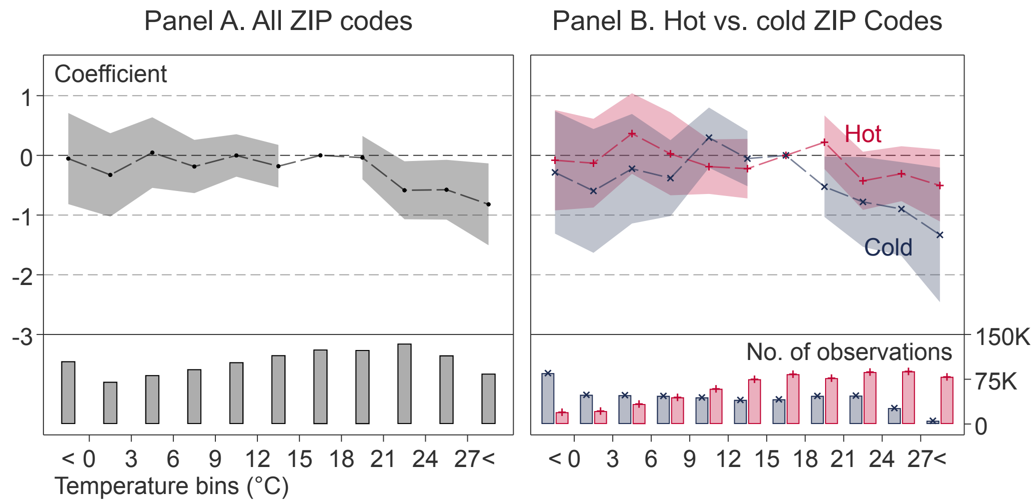

Notes: Coefficients with 95% confidence intervals (left y-axis), and number of observations (right y-axis) from regressions of log(number of correct answers + 1) on 3°C-bin indicators of the average air temperature during the 24 hours preceding the play (x-axis) as presented in Table A12. The standard errors are clustered on ZIP Codes. The reference bin is 15-18°C. The regression in Panel A includes all observations. Panel B shows the results from separate regressions for the cold-ZIP Codes sample (below-median 2015-2019 average temperatures), and the hot-ZIP Codes sample (above-median 2015-2019 average temperatures). The control variables include individual effects, time effects, play controls, and weather controls (see Section 3).

Figure A8 Air temperature and cognitive performance by age: 3°C-bins regressions

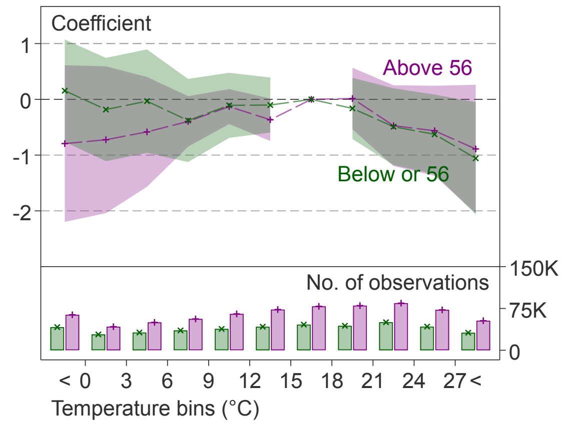

Notes: Coefficients with 95% confidence intervals (left y-axis), and number of observations (right y-axis) from regressions of the number of correct answers on 3°C-bin indicators of the average air temperature during the 24 hours preceding the play (x-axis) as presented in Table A17. The standard errors are clustered on ZIP Codes. The reference bin is 15-18°C. The figure shows the results from separate regressions for individuals below, or exactly, 56 years old, and individuals above 56 years old. The control variables include individual effects, time effects, play controls, and weather controls (see Section 3).

Appendix. Tables

Table A1 Air temperature and cognitive performance: 3°C-bins regressions

Dependent variable: No. of correct answers	All ZIP Codes	Cold ZIP Codes	Hot ZIP Codes
Air temp. < 0°C	-0.428	-0.614	-0.699
	(0.500)	(0.625)	(0.459)
Air temp. 0-3°C	-0.506	-0.872	-0.130
	(0.458)	(0.667)	(0.308)
Air temp. 3-6°C	-0.376	-0.736	0.037
	(0.367)	(0.555)	(0.311)
Air temp. 6-9°C	-0.388*	-0.670**	-0.119
	(0.205)	(0.309)	(0.308)
Air temp. 9-12°C	-0.121	-0.045	-0.171
	(0.146)	(0.216)	(0.180)
Air temp. 12-15°C	-0.270*	-0.256	-0.253
	(0.154)	(0.226)	(0.224)
Air temp. 18-21°C	-0.046	-0.518	0.208
	(0.209)	(0.321)	(0.227)
Air temp. 21-24°C	-0.484*	-0.913**	-0.171
	(0.270)	(0.437)	(0.178)
Air temp. 24-27°C	-0.588**	-0.995**	-0.212
	(0.293)	(0.412)	(0.162)
Air temp. ≥ 27°C	-0.954**	-1.817***	-0.452*
	(0.425)	(0.551)	(0.244)
Observations	1,151,059	482,812	668,247
Individuals	31,029	12,708	18,321
ZIPs	748	374	374

Notes: Results from regressions of the number of correct answers on 3°C-bin indicators of the average air temperature during the 24 hours preceding the play as depicted in Figure 1. The standard errors (in parentheses) are clustered on ZIP Codes. The reference bin is 15-18°C. The regression in column 1 includes all observations. Columns 2 and 3 show the results from separate regressions for the cold-ZIP Codes sample (below-median 2015-2019 average temperatures), and the hot-ZIP Codes sample (above-median 2015-2019 average temperatures). The control variables include individual effects, time effects, play controls, and weather controls (see Section 3).

Table A2 Results from previous literature

Paper	Temperature variable	Outcome variable	Temperature cutoff	Mean of outcome var.	Coefficient (Std. err.)	Δ outcome var. in %
Krebs (2022)	Average temp. during 24h preceding play	Mental arithmetic performance	16.5°C	65.4	-0.084 (0.038)	-0.13
Graff Zivin et al. (2018)	Same-day number of degree days	National Longitudinal Survey of Youth	21°C	49.7	–0.211 (0.090)	-0.42
Graff Zivin et al. (2020)	Same-day number of degree days	China's National College Entrance Examination	14°C	519	-0.0034[i] (0.0004)	-0.34
Park (2022)	Average temp. around test time	Regents Exams in New York City	none[ii]	64.9	-0.290[iii]	-0.45
Melo and Suzuki (2021)	Temp. during exam	Exame Nacional do Enismo Mèdio in Brazil	none[ii]	506.5	-0.879[iii]	-1.73
Yi et al. (2021)	Same-day heat stress degree days	China Health and Retirement Longitudinal Study	25°C	2.8	-0.025 (0.013)	-0.89

[i] Coefficient and standard error from log-lin regression. [ii] Tests taken during relatively warm weather. [iii] No standard error is provided because the coefficient is translated from a regression on the standardized test score.
Notes: Comparison of the main result of this paper to the main result of five other papers that focus on short-term effects and apply a linear approach. The table shows the temperature variable, the outcome variable, the temperature cutoff, the mean of the outcome variable, and the estimated coefficient. I calculate the percentage change of the outcome variable in percent for a 1°C increase in temperature, displayed in the last column.

Table A3 Summary statistics of extensive-margin selection

Variable	All ZIP Codes	Cold ZIP Codes	Hot ZIP Codes
Played	0.0035	0.0036	0.0034
	(0.0589)	(0.0596)	(0.0583)
Temperature past 24h in °C	14.5942	10.3918	17.5256
	(9.7932)	(10.1184)	(8.3951)
Observations	226,926,601	93,249,007	133,677,594

Notes: Means and standard deviations (in parentheses) of the extensive margin selection variable (indicator variable = 1 if an individual played during a specific hour), and the average air temperature during the 24 hours preceding a play in °C. See Section 4.2 for a sample construction description.

Table A4 Summary statistics of intensive-margin selection

Variable	All ZIP Codes	Cold ZIP Codes	Hot ZIP Codes
Number of plays	1.4541	1.4475	1.4590
	(1.5234)	(1.5149)	(1.5295)
Temperature past 24h in °C	14.0711	9.7705	17.2051
	(9.9291)	(10.1390)	(8.4982)
Observations	789,305	332,720	456,585

Notes: Means and standard deviations (in parentheses) of the intensive margin selection variable (the number of plays an individual engaged in during an hour in which the individual played), and the average air temperature during the 24 hours preceding a play in °C. See Section 4.2 for a sample construction description.

Table A5 Air temperature and probability of playing: 3°C-bins regressions

Dependent variable: Play	All ZIP Codes	Cold ZIP Codes	Hot ZIP Codes
Air temp. < 0°C	0.0107***	0.0120***	0.0130**
	(0.0035)	(0.0043)	(0.0055)
Air temp. 0-3°C	0.0069**	0.0081**	0.0085*
	(0.0031)	(0.0040)	(0.0048)
Air temp. 3-6°C	0.0058**	0.0093**	0.0032
	(0.0028)	(0.0039)	(0.0039)
Air temp. 6-9°C	0.0005	0.0054	-0.0034
	(0.0027)	(0.0039)	(0.0036)
Air temp. 9-12°C	-0.0001	0.0086***	-0.0059*
	(0.0025)	(0.0032)	(0.0032)
Air temp. 12-15°C	0.0008	0.0069*	-0.0024
	(0.0021)	(0.0037)	(0.0026)
Air temp. 18-21°C	-0.0025	-0.0022	-0.0024
	(0.0017)	(0.0028)	(0.0022)
Air temp. 21-24°C	-0.0024	-0.0008	-0.0034
	(0.0022)	(0.0036)	(0.0028)
Air temp. 24-27°C	-0.0008	-0.0008	-0.0015
	(0.0025)	(0.0041)	(0.0029)
Air temp. ≥ 27°C	-0.0007	0.0049	-0.0030
	(0.0032)	(0.0057)	(0.0034)
Observations	226,926,601	93,249,007	133,677,594
Individuals	30,081	12,344	17,737
ZIPs	745	373	372

Notes: Results from regressions of play (indicator variable = 100 if an individual played during a specific hour) on 3°C-bin indicators of the average air temperature during the 24 hours preceding the play as depicted in Figure A3. The standard errors (in parentheses) are clustered on ZIP Codes. The standard errors (in parentheses) are clustered on ZIP Codes. The reference bin is 15-18°C. The regression in column 1 includes all observations. Columns 2 and 3 show the results from separate regressions for the cold-ZIP Codes sample (below-median 2015-2019 average temperatures), and the hot-ZIP Codes sample (above-median 2015-2019 average temperatures). The control variables include individual effects, time effects, play controls, and weather controls (see Section 3).

Table A6 Air temperature and frequency of playing: 3°C-bins regressions

Dependent variable: Log no. of plays	All ZIP Codes	Cold ZIP Codes	Hot ZIP Codes
Air temp. < 0°C	-0.0010	0.0052	-0.0070
	(0.0032)	(0.0047)	(0.0052)
Air temp. 0-3°C	0.0073**	0.0127***	0.0040
	(0.0030)	(0.0047)	(0.0045)
Air temp. 3-6°C	0.0023	0.0062	0.0013
	(0.0028)	(0.0045)	(0.0040)
Air temp. 6-9°C	0.0007	0.0063*	-0.0028
	(0.0024)	(0.0038)	(0.0037)
Air temp. 9-12°C	0.0026	0.0042	0.0025
	(0.0021)	(0.0034)	(0.0031)
Air temp. 12-15°C	-0.0015	-0.0038	0.0004
	(0.0019)	(0.0032)	(0.0025)
Air temp. 18-21°C	-0.0051**	-0.0037	-0.0065**
	(0.0021)	(0.0034)	(0.0027)
Air temp. 21-24°C	-0.0022	-0.0051	-0.0010
	(0.0021)	(0.0035)	(0.0027)
Air temp. 24-27°C	-0.0013	0.0006	-0.0013
	(0.0029)	(0.0049)	(0.0034)
Air temp. ≥ 27°C	-0.0038	-0.0042	-0.0021
	(0.0034)	(0.0075)	(0.0038)
Observations	789,305	332,720	456,585
Individuals	30,081	12,344	17,737
ZIPs	745	373	372

Notes: Results from regressions of the log number of plays (the log number of plays an individual engaged in during the current hour) on 3°C-bin indicators of the average air temperature during the 24 hours preceding the play as depicted in Figure A4. The standard errors (in parentheses) are clustered on ZIP Codes. The reference bin is 15-18°C. The regression in column 1 includes all observations. Columns 2 and 3 show the results from separate regressions for the cold-ZIP Codes sample (below-median 2015-2019 average temperatures), and the hot-ZIP Codes sample (above-median 2015-2019 average temperatures). The control variables include individual effects, time effects, play controls, and weather controls (see Section 3).

Table A7 Air temperature past 48 hours and cognitive performance: piecewise-linear regressions

Dependent variable: No. of correct answers	All ZIP Codes	Cold ZIP Codes	Hot ZIP Codes
Air temperature	0.013	0.016	0.013
	(0.023)	(0.019)	(0.018)
Linear combination test:	-0.093**	-0.177***	-0.045
Air temp. + air temp. × above theshold	(0.041)	(0.054)	(0.030)
Observations	1,150,343	482,500	667,843
Individuals	31,027	12,706	18,321
ZIPs	748	374	374

Notes: Results from piecewise-linear regressions of the number of correct answers on the average air temperature during the 48 hours preceding the play (in °C), and an interaction with the above-threshold indicator. The standard errors (in parentheses) are clustered on ZIP Codes. "Air temperature" denotes the coefficient below the threshold, and the linear combination test amounts to the slope above the threshold. The threshold is 16.5°C. The regression in column 1 includes all observations. Columns 2 and 3 show the results from separate regressions for the cold-ZIP Codes sample (below-median 2015-2019 average temperatures), and the hot-ZIP Codes sample (above-median 2015-2019 average temperatures). The control variables include individual effects, time effects, play controls, weather controls, and an indicator for above threshold temperature (see Section 3).

Table A8 Air temperature past 48 hours and cognitive performance: 3°C-bins regressions

Dependent variable: No. of correct answers	All ZIP Codes	Cold ZIP Codes	Hot ZIP Codes
Air temp. < 0°C	-0.327	-0.519	-0.688
	(0.468)	(0.512)	(0.550)
Air temp. 0-3°C	-0.330	-0.784	0.328
	(0.533)	(0.659)	(0.299)
Air temp. 3-6°C	-0.364	-0.616	-0.087
	(0.328)	(0.439)	(0.308)
Air temp. 6-9°C	-0.264	-0.487**	-0.051
	(0.176)	(0.215)	(0.295)
Air temp. 9-12°C	-0.009	0.034	-0.011
	(0.161)	(0.202)	(0.191)
Air temp. 12-15°C	-0.074	0.024	-0.091
	(0.111)	(0.150)	(0.157)
Air temp. 18-21°C	0.158	-0.282	0.372
	(0.213)	(0.212)	(0.299)
Air temp. 21-24°C	-0.477*	-1.124**	-0.048
	(0.288)	(0.446)	(0.190)
Air temp. 24-27°C	-0.553**	-0.925***	-0.204
	(0.246)	(0.349)	(0.178)
Air temp. ≥ 27°C	-0.945**	-1.855***	-0.466*
	(0.390)	(0.571)	(0.262)
Observations	1,150,343	482,500	667,843
Individuals	31,027	12,706	18,321
ZIPs	748	374	374

Notes: Results from regressions of the number of correct answers on 3°C-bin indicators of the average air temperature during the 48 hours preceding the play as depicted in Figure A5. The standard errors (in parentheses) are clustered on ZIP Codes. The reference bin is 15-18°C. The regression in column 1 includes all observations. Columns 2 and 3 show the results from separate regressions for the cold-ZIP Codes sample (below-median 2015-2019 average temperatures), and the hot-ZIP Codes sample (above-median 2015-2019 average temperatures). The control variables include individual effects, time effects, play controls, and weather controls (see Section 3).

Table A9 Heat index and cognitive performance: piecewise-linear regressions

Dependent variable: No. of correct answers	All ZIP Codes	Cold ZIP Codes	Hot ZIP Codes
Heat index	0.011	0.012	0.012
	(0.020)	(0.016)	(0.015)
Linear combination test:	-0.081***	-0.117***	-0.050*
Air temp. + air temp. × above theshold	(0.031)	(0.043)	(0.026)
Observations	1,151,059	482,812	668,247
Individuals	31,029	12,708	18,321
ZIPs	748	374	374

Notes: Results from piecewise-linear regressions of the number of correct answers on the average heat index during the 24 hours preceding the play (in °C), and an interaction with the above-threshold indicator. The standard errors (in parentheses) are clustered on ZIP Codes. "Air temperature" denotes the coefficient below the threshold, and the linear combination test amounts to the slope above the threshold. The threshold is 16.5°C. The regression in column 1 includes all observations. Columns 2 and 3 show the results from separate regressions for the cold-ZIP Codes sample (below-median 2015-2019 average temperatures), and the hot-ZIP Codes sample (above-median 2015-2019 average temperatures). The control variables include individual effects, time effects, play controls, weather controls, and an indicator for above threshold temperature (see Section 3).

Table A10 Heat index and cognitive performance: 3°C-bins regressions

Dependent variable: No. of correct answers	All ZIP Codes	Cold ZIP Codes	Hot ZIP Codes
Heat index < 0°C	-0.573	-0.641	-0.582*
	(0.506)	(0.590)	(0.330)
Heat index 0-3°C	-0.516	-0.722	-0.066
	(0.432)	(0.543)	(0.322)
Heat index 3-6°C	-0.386	-0.384	-0.284
	(0.267)	(0.300)	(0.450)
Heat index 6-9°C	-0.570***	-0.624*	-0.389
	(0.211)	(0.323)	(0.243)
Heat index 9-12°C	-0.141	0.099	-0.253
	(0.124)	(0.193)	(0.172)
Heat index 12-15°C	-0.404*	-0.308	-0.393
	(0.215)	(0.218)	(0.318)
Heat index 18-21°C	-0.122	-0.495	0.085
	(0.178)	(0.335)	(0.151)
Heat index 21-24°C	-0.574**	-0.913**	-0.314*
	(0.257)	(0.435)	(0.166)
Heat index 24-27°C	-0.617**	-1.082***	-0.229
	(0.281)	(0.397)	(0.184)
Heat index ≥ 27°C	-0.914**	-1.159**	-0.520**
	(0.378)	(0.544)	(0.237)
Observations	1,151,059	482,812	668,247
Individuals	31,029	12,708	18,321
ZIPs	748	374	374

Notes: Results from regressions of the number of correct answers on 3°C-bin indicators of the average heat index during the 24 hours preceding the play as depicted in Figure A6. The standard errors (in parentheses) are clustered on ZIP Codes. The reference bin is 15-18°C. The regression in column 1 includes all observations. Columns 2 and 3 show the results from separate regressions for the cold-ZIP Codes sample (below-median 2015-2019 average temperatures), and the hot-ZIP Codes sample (above-median 2015-2019 average temperatures). The control variables include individual effects, time effects, play controls, and weather controls (see Section 3).

Table A11 Air temperature and log cognitive performance: piecewise-linear regressions

Dependent variable: No. of correct answers	All ZIP Codes	Cold ZIP Codes	Hot ZIP Codes
Air temperature	-0.003	0.008	-0.010
	(0.016)	(0.016)	(0.018)
Linear combination test:	-0.079**	-0.113**	-0.057*
Air temp. + air temp. × above theshold	(0.031)	(0.049)	(0.030)
Observations	1,151,059	482,812	668,247
Individuals	31,029	12,708	18,321
ZIPs	748	374	374

Notes: Results from piecewise-linear regressions of log(number of correct answers + 1) on the average air temperature during the 24 hours preceding the play (in °C), and an interaction with the above-threshold indicator. The standard errors (in parentheses) are clustered on ZIP Codes. "Air temperature" denotes the coefficient below the threshold, and the linear combination test amounts to the slope above the threshold. The threshold is 16.5°C. The regression in column 1 includes all observations. Columns 2 and 3 show the results from separate regressions for the cold-ZIP Codes sample (below-median 2015-2019 average temperatures), and the hot-ZIP Codes sample (above-median 2015-2019 average temperatures). The control variables include individual effects, time effects, play controls, weather controls, and an indicator for above threshold temperature (see Section 3).

Table A12 Air temperature and log cognitive performance: 3°C-bins regressions

Dependent variable: Log (no. of correct answers + 1)	All ZIP Codes	Cold ZIP Codes	Hot ZIP Codes
Air temp. < 0°C	-0.054	-0.285	-0.082
	(0.389)	(0.522)	(0.427)
Air temp. 0-3°C	-0.327	-0.596	-0.131
	(0.356)	(0.527)	(0.378)
Air temp. 3-6°C	0.046	-0.224	0.366
	(0.301)	(0.467)	(0.343)
Air temp. 6-9°C	-0.187	-0.381	0.026
	(0.227)	(0.323)	(0.353)
Air temp. 9-12°C	-0.002	0.297	-0.189
	(0.181)	(0.256)	(0.232)
Air temp. 12-15°C	-0.182	-0.054	-0.223
	(0.182)	(0.235)	(0.253)
Air temp. 18-21°C	-0.038	-0.526**	0.220
	(0.185)	(0.257)	(0.228)
Air temp. 21-24°C	-0.585**	-0.780**	-0.427*
	(0.247)	(0.383)	(0.248)
Air temp. 24-27°C	-0.575**	-0.898**	-0.307
	(0.255)	(0.399)	(0.234)
Air temp. $\geq$ 27°C	-0.820**	-1.332**	-0.504
	(0.349)	(0.573)	(0.306)
Observations	1,151,059	482,812	668,247
Individuals	31,029	12,708	18,321
ZIPs	748	374	374

Notes: Results from regressions of log(number of correct answers + 1) on 3°C-bin indicators of the average air temperature during the 24 hours preceding the play as depicted in Figure 1. The standard errors (in parentheses) are clustered on ZIP Codes. The reference bin is 15-18°C. The regression in column 1 includes all observations. Columns 2 and 3 show the results from separate regressions for the cold-ZIP Codes sample (below-median 2015-2019 average temperatures), and the hot-ZIP Codes sample (above-median 2015-2019 average temperatures). The control variables include individual effects, time effects, play controls, and weather controls (see Section 3).

Table A13 Air temperature and error rate: piecewise-linear regressions

Dependent variable: No. of correct answers	All ZIP Codes	Cold ZIP Codes	Hot ZIP Codes
Air temperature	0.000	-0.001	0.002
	(0.002)	(0.002)	(0.004)
Linear combination test:	0.001	0.004	-0.002
Air temp. + air temp. × above theshold	(0.004)	(0.006)	(0.004)
Observations	1,151,059	482,812	668,247
Individuals	31,029	12,708	18,321
ZIPs	748	374	374

Notes: Results from piecewise-linear regressions of the error rate (number of incorrect answers / total answers) on the average air temperature during the 24 hours preceding the play (in °C), and an interaction with the above-threshold indicator. The standard errors (in parentheses) are clustered on ZIP Codes. "Air temperature" denotes the coefficient below the threshold, the linear combination test amounts to the slope above the threshold. The threshold is 16.5°C. The regression in column 1 includes all observations. Columns 2 and 3 show the results from separate regressions for the cold-ZIP Codes sample (below-median 2015-2019 average temperatures), and the hot-ZIP Codes sample (above-median 2015-2019 average temperatures). The control variables include individual effects, time effects, play controls, weather controls, and an indicator for above threshold temperature (see Section 3).

Table A14 Air temperature and error rate: 3°C-bins regressions

Dependent variable: Error rate	All ZIP Codes	Cold ZIP Codes	Hot ZIP Codes
Air temp. < 0°C	-0.016	0.004	-0.074
	(0.046)	(0.055)	(0.078)
Air temp. 0-3°C	0.006	-0.019	0.061
	(0.044)	(0.054)	(0.067)
Air temp. 3-6°C	-0.015	-0.013	-0.042
	(0.040)	(0.055)	(0.053)
Air temp. 6-9°C	0.003	-0.000	-0.019
	(0.036)	(0.051)	(0.052)
Air temp. 9-12°C	-0.014	-0.065	0.005
	(0.029)	(0.044)	(0.038)
Air temp. 12-15°C	0.006	-0.018	0.010
	(0.028)	(0.041)	(0.037)
Air temp. 18-21°C	-0.027	0.014	-0.048
	(0.028)	(0.036)	(0.037)
Air temp. 21-24°C	-0.040	-0.015	-0.058
	(0.033)	(0.045)	(0.042)
Air temp. 24-27°C	-0.005	0.060	-0.045
	(0.038)	(0.051)	(0.047)
Air temp. $\geq$ 27°C	-0.034	-0.015	-0.068
	(0.047)	(0.092)	(0.051)
Observations	1,151,059	482,812	668,247
Individuals	31,029	12,708	18,321
ZIPs	748	374	374

Notes: Results from regressions of the error rate (number of incorrect answers / total answers) on 3°C-bin indicators of the average air temperature during the 24 hours preceding the play as depicted in Figure 2. The standard errors (in parentheses) are clustered on ZIP Codes. The reference bin is 15-18°C. The regression in column 1 includes all observations. Columns 2 and 3 show the results from separate regressions for the cold-ZIP Codes sample (below-median 2015-2019 average temperatures), and the hot-ZIP Codes sample (above-median 2015-2019 average temperatures). The control variables include individual effects, time effects, play controls, and weather controls (see Section 3).

Table A15 Effect accumulation

Dep. var.: No. of correct answers	All ZIP Codes		Cold ZIP Codes		Hot ZIP Codes	
Av. air temp. past 24h > 21°C (baseline)	-0.4894**		-0.6657**		-0.2630**	
	(0.2270)		(0.2579)		(0.1290)	
Av. air temp. > 21°C for 1 or 2 days		-0.2206		-0.4130*		-0.0139
		(0.1934)		(0.2345)		(0.1806)
Av. air temp. > 21°C for 3 to 6 days		-0.5682**		-0.8578***		-0.2696**
		(0.2288)		(0.2899)		(0.1364)
Av. air temp. > 21°C for 7 days or more		-0.7886**		-1.0750**		-0.4790**
		(0.3446)		(0.5032)		(0.2086)
Observations	1,142,089	1,142,089	478,616	478,616	663,473	663,473
Individuals	31,006	31,006	12,702	12,702	18,304	18,304
ZIPs	747	747	374	374	373	373

Notes: Results from regressions of the number of correct answers on and indicator = 1 if the average temperature was above 21°C during different temporal periods before a play as presented in Figure 3. The standard errors are clustered on ZIP Codes. The reference is a day with an average temperature below or exactly 21°C. I run separate regressions for all observations, the cold-ZIP Codes sample (below-median 2015-2019 average temperatures), and the hot-ZIP Codes sample (above-median 2015-2019 average temperatures). The control variables include individual effects, time effects, play controls, and weather controls (see Section 3).

Table A16 Air temperature and cognitive performance by age: piecewise-linear regressions

Dependent variable: No. of correct answers	Below or 56 years old	Above 56 years old
Air temperature	-0.012	0.024
	(0.021)	(0.030)
Linear combination test:	-0.068	-0.091*
Air temp. + air temp. × above theshold	(0.049)	(0.049)
Observations	431,871	719,188
Individuals	15,758	15,271
ZIPs	722	728

Notes: Results from piecewise-linear regressions of the number of correct answers on the average air temperature during the 24 hours preceding the play (in °C), and an interaction with the above-threshold indicator. The standard errors (in parentheses) are clustered on ZIP Codes. "Air temperature" denotes the coefficient below the threshold, and the linear combination test amounts to the slope above the threshold. The threshold is 16.5°C. Columns 1 and 2 show the results from separate regressions for individuals below or 56 years old, and individuals above 56 years old. The control variables include individual effects, time effects, play controls, weather controls, and an indicator for above threshold temperature (see Section 3).

Table A17 Air temperature and cognitive performance by age: 3°C-bins regressions

Dependent variable: No. of correct answers	Below or 56 years old	Above 56 years old
Air temp. < 0°C	0.155	-0.794
	(0.468)	(0.715)
Air temp. 0-3°C	-0.183	-0.724
	(0.472)	(0.671)
Air temp. 3-6°C	-0.032	-0.584
	(0.472)	(0.501)
Air temp. 6-9°C	-0.379	-0.396*
	(0.380)	(0.231)
Air temp. 9-12°C	-0.107	-0.130
	(0.297)	(0.159)
Air temp. 12-15°C	-0.102	-0.366*
	(0.251)	(0.194)
Air temp. 18-21°C	-0.162	0.015
	(0.280)	(0.281)
Air temp. 21-24°C	-0.491	-0.472
	(0.351)	(0.369)
Air temp. 24-27°C	-0.627*	-0.564
	(0.363)	(0.412)
Air temp. ≥ 27°C	-1.057**	-0.890
	(0.513)	(0.586)
Observations	431,871	719,188
Individuals	15,758	15,271
ZIPs	722	728

Notes: Results from regressions of the number of correct answers on 3°C-bin indicators of the average air temperature during the 24 hours preceding the play as depicted in Figure A8. The standard errors (in parentheses) are clustered on ZIP Codes. The reference bin is 15-18°C. Columns 1 and 2 show the results from separate regressions for individuals below or 56 years old, and individuals above 56 years old. The control variables include individual effects, time effects, play controls, and weather controls (see Section 3).

Contact.

MIT CEEPR Working Paper Series
is published by the MIT Center for Energy and Environmental Policy Research from submissions by affiliated researchers.
For inquiries and/or for permission to reproduce material in this working paper, please contact:

General inquiries: ceepr@mit.edu
Media inquiries: ceepr-media@mit.edu

MIT Center for Energy and
Environmental Policy Research

MIT Center for Energy and
Environmental Policy Research
Massachusetts Institute of Technology
77 Massachusetts Avenue, E19-411
Cambridge, MA 02139-4307
USA

ceepr.mit.edu

MASSACHUSETTS INSTITUTE OF TECHNOLOGY