Quantifying the contributions of climate change and adaptation to mortality from unprecedented extreme heat events

Christopher W. Callahan[a,1], Jared Trok[a], Andrew J. Wilson[b], Carlos F. Gould[c], Sam Heft-Neal[b], Marshall Burke[a,b,d], and Noah S. Diffenbaugh[a]

This manuscript was compiled on August 26, 2025

Understanding the mortality effects of the most extreme heat events is central to climate change risk analysis and adaptation decision-making. Accurate representation of these impacts requires accounting for the effects of prolonged sequences of hot days on mortality, the change in that mortality due to anthropogenic forcing, and the potential compensating effects of adaptation to heat. Here, we revisit the August 2003 heat wave in France, a canonical event in a region with rich climate and mortality data, to understand these influences. We find that standard heat mortality exposure-response functions underpredict excess deaths in August 2003 by 55%, but that accounting for the temporally compounding effects of hot days better matches observed mortality. After accounting for compounding effects and applying a machine learning approach to single-event climate attribution, we attribute 6,079 deaths in August 2003 to climate change, ten times higher than previous estimates. Finally, we show that recent adaptation to heat in France has reduced the projected death tolls of future 2003-like events by more than 75%.

Heat-related mortality | Impact attribution | Climate adaptation

U nderstanding the contribution of climate change to human mortality from unprecedented extreme heat events is a critical priority (1). Unprecedented extreme climate events can stress infrastructure or adaptation measures that have been benchmarked to recent experience, posing challenges for societal resilience (2). These record-breaking events are increasing due to anthropogenic forcing (3–5), with future climate change likely to generate even more extreme events than those that have been recently witnessed (6). With respect to historical events, attributing the observed health impacts of extreme heat to climate change (7–10) has the potential to inform ongoing climate litigation (11, 12) and loss and damage compensation (13). With respect to the future, evaluating the impacts of previously unseen events made newly possible by global warming is essential to understanding the health risks of future global temperature levels (14, 15) and informing cost-benefit analysis tools such as the social cost of greenhouse gases (16, 17).

At the same time, people have a well-documented ability to adapt to extreme weather, often leading to reductions over time in the effect of heat exposure on mortality (18–23). If the conditions that generated historical extreme heat events recur at present or future levels of warming (15), they may occur not only in a different climate context but also against the backdrop of an evolving temperature-mortality relationship. As a result, accurately quantifying the past and future health risks of extreme heat requires evaluating the competing influences of climate warming and adaptation on mortality.

However, characterizing the impacts of unprecedented extreme events poses specific analytical challenges. Empirically derived exposure-response functions are a standard tool to quantify the health impacts of climate change, but are estimated using data that by definition do not include unprecedented future events. Further, these models tend to treat hot days as additively separable predictors of mortality, neglecting potential compounding effects of multiple days in sequence. Multiple hot days may result in heat accumulation in both the built environment and human bodies (24), and mortality during previous unprecedented events may have been driven by sequences of warm nights that prevented people from cooling

Significance Statement

We revisit the August 2003 heat wave in France to better understand the influences of climate change and human adaptation on mortality from unprecedented extreme heat events. We reconfirm previous findings of approximately 16,000 excess deaths during the event, but show that standard heat-mortality exposure-response functions underestimate this death toll by 55% in an out-of-sample test. Using an exposure-response function accounting for consecutive hot days that more effectively predicts mortality in 2003, we attribute more than 6,000 deaths in August 2003 to climate change. Additionally, we show that the heat-mortality relationship has moderated since 2003. As a result, we project future 2003-like events to cause 77% fewer deaths in France compared to a world without that moderation.
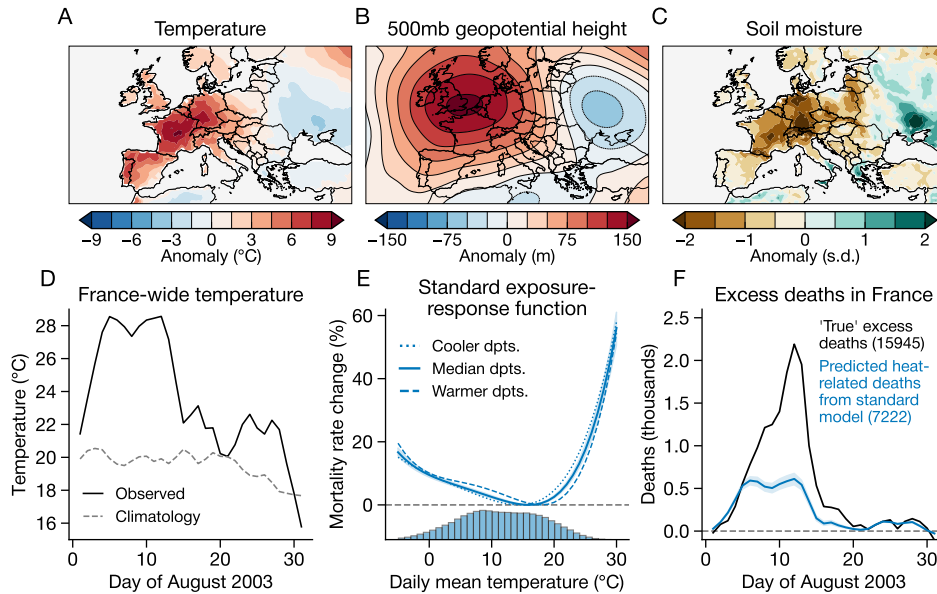
Fig. 1. Physical and epidemiological characteristics of the August 2003 heat wave in Europe. a, b, c) Anomalies of temperature (a), 500-mb geopotential height (b), and soil moisture (c) averaged over 1-14 August 2003. d) Population-weighted average temperature across French departments in August 2003 (black), along with the 1980-2002 mean (gray, dashed). e) Exposure-response function relating daily mean temperature to mortality rates over 1980-2002, using a fourth-order polynomial, 5 lags of temperature, and a continuous interaction with department mean temperature (Methods). Different lines show the response functions at the three terciles of department mean climate. Lower histogram shows distribution of daily temperatures in the sample. f) Excess deaths (black) and heat-related deaths (blue) across France. Excess deaths are based on deviations relative to averages and heat-related deaths are based on the exposure-response function in panel (e). In (e) and (f), shading shows 95% confidence intervals.

themselves after hot days (25). On the other hand, many statistical studies find little additional effect of the sequencing of hot days above their independent effects (26–28). As a result, it remains unclear whether standard statistical models estimated on the full historical distribution of temperatures are suitable for quantifying the effects of unprecedented sequences of very hot temperatures.

Here we revisit the August 2003 heat wave in France, an event which offers insight into the contributions of climate change and adaptation to extreme heat mortality for multiple reasons. First, when it occurred, 2003 was the hottest summer in Europe in at least the previous 500 years (29, 30) and the level of heat that was reached was partly due to climate change (31), making 2003 a useful test case for events that are out-of-sample relative to recent experience. Second, mortality in France appeared to be uniquely sensitive to heat prior to 2003 (19), yielding severe mortality during this event (32–34). For example, many of the victims lived in small, poorly ventilated apartments directly under zinc roofs that were extremely effective at trapping heat, potentially exposing residents to the compounding effects of multiple days of heat accumulation (35, 36). Third, France collects detailed daily mortality data, enabling robust statistical analysis. Finally, France adopted a series of adaptation measures immediately following the 2003 event, including the expansion of air conditioning in vulnerable locations such as nursing homes (18) and heat action plans that include educational messaging and proactive visits to isolated people during hot periods (37). Comparing the temperature-mortality relationship before and after 2003 thus offers a simple way to assess the effectiveness of these adaptation measures and other societal changes (18–22).

We take five key steps in our analysis (Fig. S1). First, we evaluate the skill of standard exposure-response functions when applied out-of-sample to the 2003 event, to determine whether empirically derived functions can skillfully represent the impacts of unseen events. Second, we develop response functions that explicitly incorporate temporally compounding heat. Third, w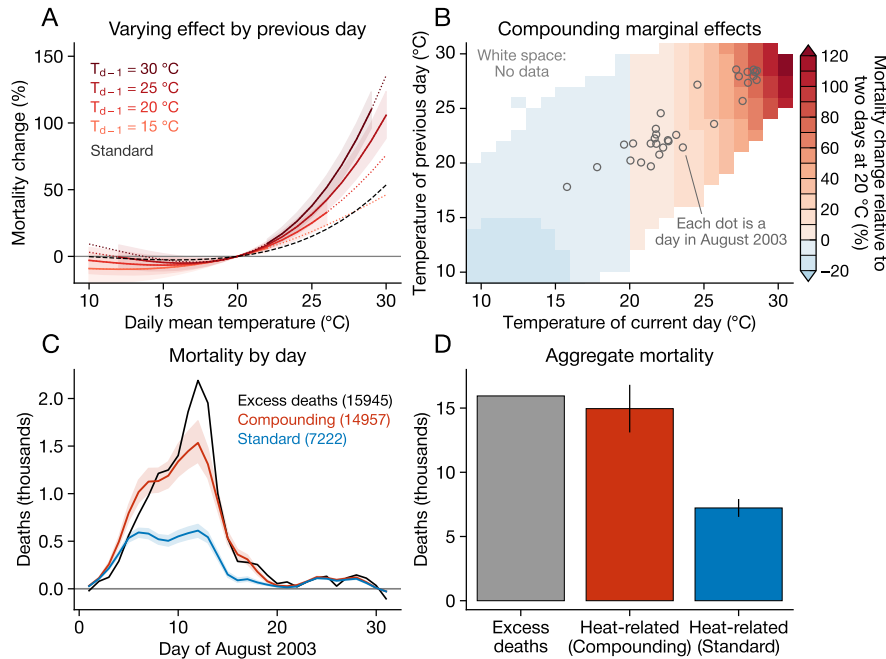e combine these exposure-response functions with a machine learning-based approach to extreme climate event attribution (38) to quantify the contribution of climate change to mortality in August 2003. Fourth, we evaluate the change in exposure-response functions before and after 2003 and quantify the effect of adaptation on heat-related mortality in the recent period. Finally, using the same machine learning-based approach to project the intensity of 2003-like events if such events were to recur in a warmer climate, we compare the extent to which recent adaptation can offset the mortality impacts of increasingly intense future heat events.

Mortality during August 2003

The first two weeks of August 2003 were characterized by extreme temperatures centered on France, Germany, and Spain (Fig. 1a), a high-pressure system centered north of France (Fig. 1b), and dry soils across much of the continent (Fig. 1c). Temperatures across France peaked at the end of the first week of August and through the second week, with a peak of 28.6 °C (daily mean) on 12 August (Fig. 1d).

To understand the death toll of this event, we first derive standard exposure-response functions that relate daily temperatures in French administrative regions (*"départements"* or "departments") to mortality in those regions, accounting for spatial heterogeneity as a function of baseline climate (Methods). We use a fourth-order polynomial in daily mean temperature across each day and the 5 days following it to account for lagged responses to temperature (Methods). We find a strong nonlinear relationship between temperature and mortality, where mortality rates increase at both cold and hot temperatures (Fig. 1e). We also observe spatial heterogeneity in this relationship that may suggest some baseline adaptation to temperature, with warmer departments exhibiting a slightly weaker response to high temperatures (Fig. 1e). Our findings are similar to the responses found in many previous studies, including those using two-stage pooled time series models (39). They are also similar when we use different numbers of lags, polynomial orders, fixed effects, or temperature variables (Fig. S2).

249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272
273

311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335

Heat wave mortality accounting for temporal compounding

Fig. 2. Temporal compounding of heat mortality. a) Mortality exposure-response functions varying by the temperature of the previous day. Confidence intervals are only shaded in the range of current-day temperatures that occur given the previous day's temperature, and lines are dotted when extrapolating outside of that range. Black dashed curve indicates the "standard" exposure-response function from Fig. 1e. Each model also includes a continuous interaction with department mean temperature and is shown at the population-weighted average department temperature for simplicity. b) Two-dimensional surface of marginal effects, illustrating the effect of combinations of current days (x-axis) and previous days (y-axis), compared to two consecutive days at 20 °C. Cells are not colored if that combination of current and previous days does not occur in the pre-2003 data. Gray dots show population-weighted France-wide temperatures on each day of August 2003. c) Predicted mortality from the heat wave by day from both the compounding and standard models. Solid line shows mean and shading shows 95% range. d) Total heat-related and excess deaths in August 2003. Bar height shows average prediction and black line shows 95% range.

To assess the "true" death toll in August 2003, we also calculate excess deaths relative to region- and time-specific baselines (Methods). Excess deaths are a standard epidemiological approach to quantify elevated mortality without specifying a cause of death or parametric exposure-response function. We estimate ~15,900 excess deaths in France across all of August (Fig. 1f), which aligns well with other estimates (32). Note that hereafter we use "excess deaths" to refer to estimates of total elevated mortality relative to averages and "heat-related deaths" to refer to mortality predicted by a temperature exposure-response function.

The standard exposure-response model shown in Fig. 1e is estimated using data from 1980 through 2002, but not including 2003, so we can perform an out-of-sample prediction of heat-related mortality during August 2003 and compare it to our estimate of excess deaths. Using this standard temperature-mortality association (Fig. 1e) to predict the August 2003 death toll underestimates total mortality by 55%: 7,222 heat-related deaths (95% confidence interval [CI]: 6,494 - 7,917) compared to 15,945 excess deaths (Fig. 1f). This underestimate is not unique to our specification; alternative polynomials, fixed effects, lag lengths, and temperature exposures yield similar results (Fig. S3).

Because excess deaths do not specify a cause or exposure-response function, an alternative interpretation of this result is that there were only ~7,200 heat-related deaths and some other cause explains the remaining ~8,700. However, there is no other known cause concurrent with the heat wave that would explain such a large number of excess deaths (35), and the magnitude of the mortality increase far exceeded typical variation from other causes (32). We thus interpret the gap between excess deaths and heat-related deaths as an underestimate from the exposure-response function rather than an overestimate from the excess deaths calculation.

Temporally compounding heat-related mortality

We hypothesize that neglecting the unique effects of multiple hot days in sequence may contribute to the underestimate from the standard model, which does not consider the order of hot days within our 5 days of lags. To incorporate temporal compounding, we modify our regression to distinguish between hot days that occur in sequence and hot days that occur after typical or cool days. Specifically, again using the pre-2003 data, we estimate response functions where the current day's temperature is interacted with the previous day's temperature. This specification allows the effect of current-day temperatures on mortality to vary according to the previous day (Methods). We include a separate interaction with department mean climate to allow spatial heterogeneity, as in the standard model (Fig. 1e); this interaction is included in all calculations, but we present results evaluated at the population-weighted average department temperature in the text and figures for clarity.

We find that, before 2003, sequencing a hot day after another hot day can nearly double the mortality effect of that hot day relative to if it occurred in isolation (Fig. 2a, 2b). For the average department, a 30 °C day increases mortality by 76% if it follows a 20 °C day, relative to two days at 20 °C (Fig. 2b). However, when a 30 °C day follows another 30 °C day, the mortality increase rises to 136% (Fig. 2b).

The strong differentiation between the responses conditional on previous days suggests an important role for temporal compounding in shaping heat wave mortality. Indeed, predicting August 2003 mortality using the compounding model yields a death toll that is closer to the total excess deaths than the standard model (Fig. 2c). The mean prediction from the compounding model is 14,957 heat-related deaths (CI: 13,350 - 16,810) compared to the mean of 7,222 from the standard model, and the confidence intervals from the compounding model include the total excess deaths value of 15,945 (Fig. 2d).
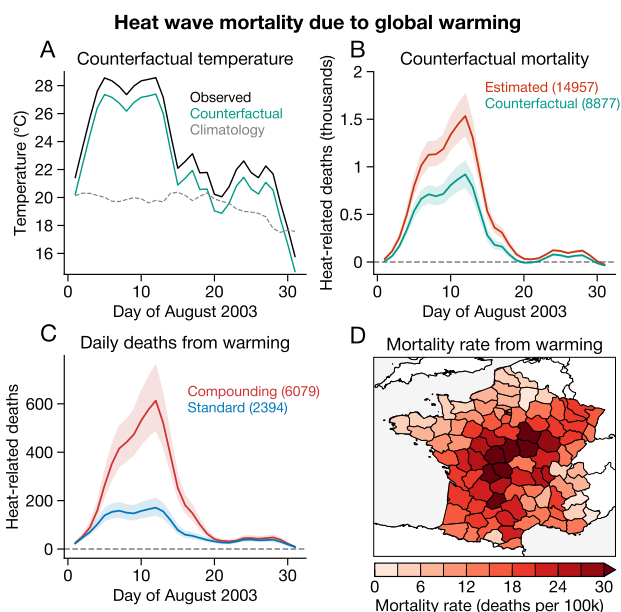
Fig. 3. Mortality attributable to global warming in August 2003. a) Observed (black), counterfactual (green), and climatological (gray) temperatures across France in August 2003. b) Mortality in August 2003 under observed conditions (red, as in Fig. 2d) and under counterfactual conditions (green). c) Mortality attributable to climate change, calculated as the predictions from observed conditions minus the predictions from counterfactual conditions. Blue curve shows the analogous calculation using the standard model. In (b) and (c), solid line shows the mean prediction and shading shows 95% range. d) Mortality rate due to climate change in each French department, defined as deaths per 100,000 population.

While the skill of our predictions is increased relative to the standard model, our predictions still underestimate peak daily mortality, especially on 12 and 13 August (Fig. 2c), yielding an underestimate of total mortality throughout the month (Fig. 2d).

To evaluate the robustness of the effect of temporal compounding, we test two alternative approaches for the interaction between current and previous days, and find that both yield very similar results (Methods). One factor that slightly affects the results is the number of lags in the regression. We use 5 lags because the effects of heat appear to accrue only in the first several days after exposure (Fig. S4, S5), though the effects of cold days take longer (Fig. S4) (40). Using 10 lags in the compounding model yields generally similar results, though with a notable overestimate of mortality in the third week of August after the peak of the heat (Fig. S6), yielding an overestimate of aggregate event mortality (17,998 heat-related deaths vs. 15,945 excess deaths).

Global warming contributions to heat-related mortality

An improved representation of the underlying heat-mortality response allows us to return to the question of how many heat-related deaths in August 2003 were due to increasing global average temperature. We use convolutional neural networks trained on global climate models to simulate counterfactual August 2003 temperatures at 0 °C of global mean temperature change, rather than the ~0.8 °C observed at the time of the event (Methods), as the historical increase in GMT has been driven by human emissions of greenhouse gases (41). We find

that climate change increased temperatures across France by an average of 1.2 °C in the first two weeks of August 2003 (Fig. 3a).

Applying our new exposure-response functions to these observed and counterfactual temperatures, we find that the counterfactual event in the absence of global warming would have caused 8,877 excess deaths, rather than the 14,957 we estimate occurred with observed temperatures (Fig. 3b). As a result, we attribute 6,079 heat-related deaths to climate change (CI: 5,043 - 7,373), 41% of the mortality from the event and more than double the analogous result from the standard model that does not account for temporal compounding (Fig. 3c). Climate change-driven mortality was substantial across all of France, but concentrated in the center of the country (Fig. 3d), with >30 deaths per 100,000 people contributed by global warming in some departments.

The contribution of climate change to mortality witnessed in 2003 raises the question of the mortality from a similar event if it occurred in the near future. To answer this question, it is important to consider that France may have adapted to heat extremes following 2003 by adopting measures such as heat action plans (37), potentially altering the future death toll of a physically similar event (18). To test this question, we re-estimate the exposure-response function using data from 2004-2019, under the assumption that shifts in the response over time indicate adaptation.

The response of mortality to temperature in 2004-2019 is milder than in 1980-2002 (Fig. 4a); in the later period using the standard model, a 30 °C day vs. a 20 °C day increases mortality by 23% for the average department, compared to 54% in the earlier period. And while there remains differentiation between the effects of hot days following previous hot or mild days, the compounding model yields results that are similarly muted compared to before 2003. Consecutive 30 °C days increase mortality by 46% compared to consecutive 20 °C days, a 66% reduction compared to the analogous figure of 136% before 2003 (Fig. 4b)

We observe small apparent protective effects of high temperatures for some combinations of current and previous days (Fig. 4a, 4b), likely a result of extrapolating a polynomial to regions where there is little or no data. The confidence intervals on these responses contain zero, and the standard model still shows mortality increases at high temperatures (Fig. 4a), so in the aggregate high temperatures remain dangerous even after 2003.

Our counterfactual simulations are based on projecting the same meteorological conditions (Fig. 1a, b, c), but at higher levels of annual global mean temperature (GMT; Methods). We use annual GMT anomaly values of 1.5 °C and 2 °C for present and near-future conditions, respectively, emphasizing that these values refer to *annual* GMT rather than long-term global warming levels (42, 43). (For reference, the annual GMT anomaly was ~0.8 °C in 2003 and ~1.5 °C in 2024.)

If the 2003 meteorological conditions were to occur at global temperature levels of 1.5 °C or 2 °C but with the pre-2003 exposure-response function, we estimate heat-related mortality of 26,842 and 39,707 deaths, respectively (Fig. 4c). That is, if the temperature-mortality relationship had not changed following 2003, we estimate that near-term warming could approximately double the death toll witnessed in 2003. On the other hand, incorporating post-2003 changes in the
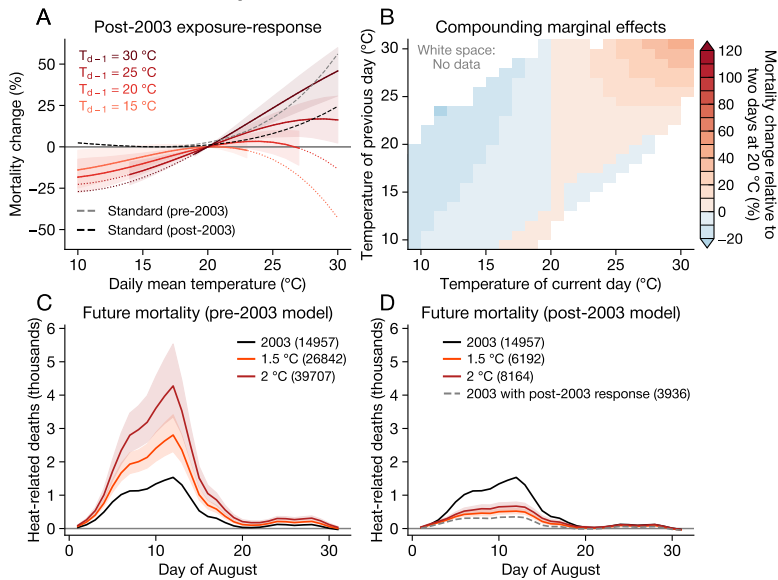
Fig. 4. Projected mortality from 2003-like events in the near future. a) Exposure-response functions as in Fig. 2 using the standard approach (black dashed line) and accounting for temporal compounding (red lines), fitted to data from 2004-2019 instead of 1980-2002. The gray dashed line shows the standard model from the 1980-2002 sample (as in Fig. 2a). Confidence intervals are only shaded in the range of current-day temperatures that occur given the previous day's temperature, and lines are dotted when extrapolating outside of that range. b) Two-dimensional marginal effects as in Fig. 2b, but for the 2004-2019 sample. Colorbar is the same as Fig. 2b. c) Predicted heat-related deaths for a 2003-like event at annual global mean temperature anomalies of 1.5 °C and 2 °C using the 1980-2002 (i.e., pre-2003-event) exposure-response functions. Black line shows the prediction for 2003 at its original temperature. All predictions use the model incorporating temporal compounding. d) As in (c), but using the 2004-2019 exposure-response functions. Gray dashed line shows predictions using 2003 temperatures but with the post-2003 exposure-response functions. In (c) and (d), solid line indicates mean projection and shading shows 95% range.

response function reduces the projected death toll by nearly an order of magnitude (Fig. 4d), with excess deaths of 6,192 and 8,164 at 1.5 °C and 2 °C. In other words, France's lower post-2003 sensitivity to extreme heat has reduced the death toll of a future 2003-like event by 77%. Similarly, if these adaptations had been undertaken prior to 2003, we calculate that the August 2003 death toll would have been 3,936, a ~74% reduction (Fig. 4d).

However, we emphasize that even with this adaptation, a 2003-like event at 1.5 °C annual GMT would, at its peak, increase daily mortality by 38% above its average daily rate. This is substantially lower than the 111% increase estimated during the peak of the 2003 event, but nevertheless highlights the magnitude of the health impacts of extreme heat, even after substantial adaptation.

Discussion

Our results demonstrate that during at least one widely studied extreme event, the sequencing of multiple extreme hot days played a key role in driving heat-related mortality. This finding differs from other work arguing that hot days have an additively separable effect on health (26–28). However, given that temporal compounding appears to play a much smaller role following 2003, our results may not generalize to all events, especially since pre-2003 France appeared to be an outlier in its heat sensitivity even relative to the rest of Europe (19). Still, given projected future increases in record-shattering heat events (6), understanding the impacts of sequences of hot days in other regions remains a research priority.

Our estimate of >6,000 deaths attributable to climate change in August 2003 in France is more than ten times larger than the previous estimate by Mitchell et al. (7) of 506. There are several potential reasons for this. We calculate mortality across all of France, rather than just in Paris. Mitchell et al. also used the fraction of attributable risk to quantify human influence on mortality, a metric which is not necessarily appropriate for impact attribution (10). Additionally, however, the exposure-response function

used in Mitchell et al. may underestimate heat-related mortality. Specifically, they used an exposure-response function that estimates a peak of 5 heat-related deaths per 100,000 population in Paris. By contrast, our excess deaths calculation (Fig. 1f) shows a peak of nearly 14 deaths per 100,000 in Paris (Fig. S7). While we previously noted that the timing of peak mortality may be uncertain, Mitchell et al. also reported 34 cumulative heat-related deaths per 100,000 population in Paris across June-August 2003, whereas we find 51 excess deaths per 100,000 population in Paris across the same time period, approximately 50% larger. Together, these results suggest that if Mitchell et al.'s exposure-response curve were extended to the rest of France, it may underestimate true excess mortality.

This discussion illustrates that a key contribution of our work is to derive and use an exposure-response function that is more appropriate for the event in question than a standard function. Future work on the impacts of extreme climate events should take care to ensure that exposure-response functions estimated from a full distribution of climate variables are skillful at representing particular extreme events of interest.

On the other hand, we find that the post-2003 response function is very different from the pre-2003 response (Fig. 4a), which may reflect adaptations undertaken in response to the death toll in 2003. Indeed, when we simulate the same heat event with two different response functions (Fig. 4b, c), we find that the milder response function generates a ten-fold reduction in mortality, suggesting large health benefits from these adaptations. Unfortunately, many other countries have not adopted the same measures as France (44), and adaptation to extreme heat appears limited on a global scale (19). So, progress in France may not indicate widespread adaptation elsewhere, though it does help to quantify the potential benefits of adaptation in the context of intensifying heat extremes.

It is notable that our model of temporal compounding underestimates daily deaths at the peak of the event (Fig. 2c). There are a number of social factors that may have con-

tributed to the extreme mortality peak, such as overstretched health infrastructure resulting both from the August vacation in France and from enormous demand for health care during the event (35, 45). Given that our models are estimated using data only prior to 2003, which may not include similar pressures on the health system, they may not be entirely effective at capturing the precise temporal dynamics of the most extreme heat events.

An additional factor that may have contributed to the elevated peak of mortality is air pollution, which we do not consider explicitly. Ozone concentrations were relatively high during the 2003 heat wave, potentially contributing to mortality (46, 47). However, spatially explicit data on daily ozone concentrations are limited over our time period of interest, making it difficult to directly model the effect of pollution on mortality. Further, high temperatures can contribute to increases in air pollution (48), so controlling for pollution concentrations in our statistical models may inappropriately exclude an important pathway by which temperature affects mortality. Understanding the dynamic interactions between temperature, air pollution, and mortality is an important area of active research.

Ours is an analysis of opportunity in some respects. France makes local daily mortality publicly available, but many other governments either do not collect or do not share comparable data, posing a challenge for researchers. Further, our machine-learning-based attribution approach has been shown to skillfully simulate temperature during the August 2003 event, but other heat waves such as the 2021 Pacific Northwest event have proven difficult to simulate both by our method (38) and others (49, 50). Alongside more sophisticated exposure-response functions, additional advances in physical event attribution may be necessary to understand the climate change contribution to mortality during more recent unprecedented events.

Conclusion

Our analysis uses econometric and machine learning tools to reveal several important insights about the influences of climate change and adaptation on mortality during the canonical 2003 heat wave in France. After accounting for the effects of multiple hot days in sequence, climate change can be linked to around 40% of the mortality of this event, even though the GMT anomaly was only approximately half of its current value (0.8 °C in 2003 vs. 1.5 °C in 2024). Even if global temperatures are stabilized near their current levels, this additional warmth may contribute more than half of mortality during future similar extreme heat events (15). However, we also reveal significant potential to reduce these harms if strong adaptation actions are taken. Widespread adoption of policies similar to those undertaken in France following 2003 may be necessary to avert mass mortality from future extreme heat.

Materials and Methods

Data. Our primary climate dataset is from the E-OBS station-based data product (51), which we use in the regression models and mortality prediction. E-OBS data are aggregated to the level of French departments, weighting by the population of each grid grid cell within the department. To plot the maps in Fig. 1, we use ERA5 reanalysis data (52) averaged over 1-14 August, with anomalies defined relative to grid cell and day of year.

Daily mortality data spanning 1980-2019 on the universe of deaths in France are made available by INSEE, the French statistical agency (https://www.insee.fr/fr/information/4769950). These data are provided at the commune level, a relatively fine geographic resolution, but we aggregate them to departments for comparison with climate data and to merge them with population data. We drop overseas territories from this analysis and focus only on the 94 departments in continental France.

Excess mortality. Excess mortality is a standard calculation to assess deviations in mortality from expected conditions. The procedure is twofold: (1) model mortality as a function of spatial and temporal baseline factors; and (2) subtract these baseline values from observed mortality during some time period of interest.

We model all-age, all-cause mortality over 1980-2019 as a function of department-specific day-of-year averages and department-specific annual averages, meaning we allow each department to have its own seasonal cycle and long-term trend. Specifically, we estimate an Ordinary Least Squares model for the log of the mortality rate (M) in department i as a function of day-of-year d and year y:

$$log(M_{idy}) = \mu_{iy} + \delta_{id} + \epsilon_{idy} \quad [1]$$

For each department-day in August 2003, we then subtract the predicted values using this equation from the observed mortality rate to calculate excess mortality.

Standard temperature-mortality exposure-response function. The goal of an exposure-response function is to describe a relationship between an exposure (e.g., temperature) and an outcome of interest (e.g., mortality rates). A standard approach is to regress mortality rates on a function of temperature, usually nonlinear, as well as non-parametric controls for all region- and time-specific average factors that could confound this relationship ("fixed effects"). Our approach models log mortality rates as a function of local daily temperature, department-by-day-of-year fixed effects, and department-by-year fixed effects:

$$log(M_{idy}) = \sum_{j=0}^{L} \left[f(T_{i(d-j)y} + f(T_{i(d-j)y}) \times \overline{T}_i \right] \quad [2]$$

$$+ \sum_{j=0}^{L} \left[\vec{\lambda}_j \vec{X}_{i(d-j)y} \right] + \mu_{iy} + \delta_{id} + \epsilon_{idy}$$

We note several features of this equation. First, the fixed effects ($\mu_{iy} + \delta_{id}$) are the same as those used in the excess mortality estimation (Eqn. 1). In effect, then, this approach seeks to isolate the temperature-driven component of excess mortality.

Second, we include L lags of daily temperature to account for the delayed effects of heat and cold. In our main analysis we use 5 lags, since the effect of heat appears to occur primarily in the first several days and to decay well before the fifth day later (Fig. S4). (We test the sensitivity of this choice by extending the lags to 10 and 30 or reducing them to 3; Fig. S2, S3.) We estimate independent coefficients for each daily lag of temperature, rather than specifying a particular parametric fit across the lags. This simple distributed lag approach is similar to that used in other papers on temperature and mortality (40), but does differ from other work that applies a smooth function or spline across the lag dimension (39). Because the effect of extreme heat on mortality only occurs within the first 3-4 days after exposure (Fig. S4), a more sophisticated specification across the lag dimension appears unnecessary in our setting.

Third, we interact daily temperatures with each department's long-term mean temperature ($\overline{T}_i$) to allow the response function to vary over space. Different locations might make different investments in adaptation resources based on their climate; for example, Mediterranean regions of France may be more likely to invest in air conditioning since they are warmer on average (22). Interacting a nonlinear function of temperature with long-term climate is a standard approach to account for spatial heterogeneity (17, 53, 54).

Finally, $\vec{X}$ refers to a vector of controls that we include alongside temperature, specifically daily mean relative humidity and daily accumulated precipitation.

In our main analysis, we use a fourth-order polynomial for $f(\cdot)$ and daily mean temperature for the exposure variable, both following other work (17, 19). Daily mean temperature, relative to daily maximum or minimum, has the advantage of balancing the effects of both hot days and warm nights, which aligns with our interest in the effects of heat accumulation. In Fig. S2, we show results using a cubic model or daily maximum or minimum temperature, which are qualitatively similar to our preferred specification.

Incorporating temporal compounding. As with most other work in this domain, the exposure-response function in Eqn. 2 treats temperatures on different days as linearly additive. Two hot days have the same effect on mortality if they occur on $d-5$ and d or if they occur on $d-1$ and d. Our approach to temporal compounding is to relax this assumption by interacting the temperature on each day with the temperature of the previous day. This approach asks the question: Does a hot day have the same effect when it occurs after another hot day as when it occurs after an average or cold day?

We estimate several variations of the following model:

$$log(M_{idy}) = \sum_{j=0}^{L} \left[f(T_{i(d-j)y} + f(T_{i(d-j)y}) \times g(T_{i(d-j-1)y}) \right]$$

$$[3]$$

$$+ \sum_{j=0}^{L} \left[f(T_{i(d-j)y}) \times \overline{T}_i + \vec{\lambda}_j \vec{X}_{i(d-j)y} \right] + \mu_{iy} + \delta_{id} + \epsilon_{idy}$$

This model adds the interaction with temperatures on day d with temperature on day $d-1$. For lagged days $d-j$, we interact that day's temperature with the temperature on day $d-j-1$. For example, temperature on day $d-2$ is interacted with temperature on day $d-3$.

The other components of this regression equation are the same as the standard model, including the fixed effects, the interaction with mean temperature, and the controls for humidity and rainfall.

The temperatures of a day and the day preceding it are not independent, statistically or physically. Fig. S8 shows a strong correlation between the temperature on day $d-1$ and the temperature on day d. While this correlation does not prevent us from identifying the effect of temporal compounding, it does mean there are many combinations of days that do not appear in the data (e.g., a 10 °C day followed by a 30 °C day does not occur in the data; see white space in Fig. 2b). An alternative approach is to interact temperature on day d with temperature *anomalies* on day $d-1$, which are less strongly correlated and may therefore provide increased statistical power (Fig. S8). Reproducing our predictions for August 2003 using this approach yields very similar results, with aggregate mortality of ~14,000 deaths (Fig. S6).

An additional concern is that our main model specifies a linear interaction between days, when the true relationship may be nonlinear. To test this assumption, we re-estimate the compounding model using natural cubic spline in the previous day's temperature, with a knot at the sample-wide median daily temperature of ~11.7 °C. This approach allows separate nonlinear interactions for cold previous days and hot previous days and again yields extremely similar results to our main analysis (Fig. S6).

In all regressions, we sample uncertainty in exposure-responses by generating a multivariate normal distribution of each vector of coefficients ($n = 500$) based on the variance-covariance matrix of each model (17). We cluster standard errors by department to account for autocorrelation in mortality. In the figures and text, we present the mean and 95% confidence interval (2.5th - 97.5th percentiles) of these 500 samples.

Machine learning predictions. We use the machine-learning-based extreme event attribution method developed by Trok et al. (38). This approach trains convolutional neural networks (CNNs) on an ensemble of climate model simulations to predict regional daily temperature from the daily meteorological conditions (including geopotential height, surface pressure, and soil moisture), the calendar day, and the annual global mean temperature (GMT) anomaly in the year leading up to the event. The trained networks are then applied to ERA5 reanalysis in an out-of-sample prediction to predict regional temperatures as a function of observed meteorological conditions. Trok et al. (38) showed that these predictions closely reproduce observed daily temperatures during the 2003 event.

Because this approach takes meteorology, day of year, and GMT as separate predictors, we can use the trained models to create new predictions which maintain the same meteorology and day of year, but prescribe an alternative GMT. We prescribe annual GMT anomalies of 0 °C for the counterfactual predictions without human influence (Fig. 3), given the strong evidence that the historical increase in GMT has been driven by human influence on the climate system (41). For the near-future events (Fig. 4), we prescribe annual GMT anomalies of 1.5 °C and 2 °C.

The results of this prediction procedure are "counterfactual" events, quantifying the regional temperature that would have resulted from the same meteorological conditions if they had occurred during a pre-industrial or future climate instead of the actual climate of 2003. While our experimental setup uses machine learning to create predictions of historical events under different global temperature conditions, our approach is a version of previous "storyline" event attribution studies that distinguish between the meteorological drivers of extreme events and the human contribution to increasing the intensity of those events (55, 56).

The target for the predictions is the average temperature over a region in southern and central Europe that encompasses France, exactly as in Trok et al. (38). We train the CNNs on three realizations each of five climate models from the sixth phase of the Coupled Model Intercomparison Project (CanESM5, UKESM1-0-LL, HadGEM3-GC31-LL, MIROC6, and MPI-ESM1-2-LR). Because this machine learning methodology requires relatively extensive data—daily data from three-dimensional meteorological variables for multiple realizations from a given climate model—these five models are the only ones currently available for training. The CNN is trained across the pooled sample of climate models, but includes an indicator variable for the model, allowing us to make separate predictions for each model. In training, we also vary the random seed 10 times to account for randomness in the training process. One final change we make to the method developed in Trok et al. is to use daily mean temperature as the predictand instead of daily maximum temperature.

We use a simple delta-method bias correction before applying these predictions to our observational data. First, for each day in August 2003, we take the difference between the predictions at the counterfactual global temperatures (0, 1.5, and 2 °C) and the predictions at the observed global temperature (0.8 °C in 2003). We then apply the region-wide "delta" for each day uniformly to each department's temperature for that day to create counterfactual temperature time series for each department.

Calculating heat-related mortality. We calculate heat-related mortality by applying exposure-response functions to observed or counterfactual time series and climatological baselines. Because the dependent variable in the regressions is log mortality, comparing the function of two different temperature values yields percent changes in mortality. We then multiply these percent changes in mortality by baseline numbers of deaths to calculate additional deaths due to heat.

When calculating mortality from the 2003 event, we compare observed temperatures in 2003 to the 1980-2002 average for each corresponding calendar day, and multiply the resulting percent differences by the 1980-2002 average number of deaths for each calendar day.

For calculations of counterfactual heat-related mortality, we combine the 500 samples of the regression coefficients and the 50 different machine learning predictions (the CNN makes separate predictions for each of five climate models across each of ten random seeds), yielding 25,000 total estimates. In our main analysis, we simply pool all CNNs trained on the different different climate models. The different climate models yield slightly different estimates, but all are qualitatively similar (Fig. S9).

1. KL Ebi, et al., Hot weather and heat extremes: health risks. *The Lancet* **398**, 698–708 (2021).
2. V Thompson, et al., The most at-risk regions in the world for high-impact heatwaves. *Nat. Commun.* **14**, 2152 (2023).
3. NS Diffenbaugh, et al., Quantifying the influence of global warming on unprecedented extreme climate events. *Proc. Natl. Acad. Sci.* **114**, 4881–4886 (2017).
4. N Christidis, GS Jones, PA Stott, Dramatically increasing chance of extremely hot summers since the 2003 European heatwave. *Nat. Clim. Chang.* **5**, 46–50 (2015).
5. V Thompson, et al., The 2021 western North America heat wave among the most extreme events ever recorded globally. *Sci. Adv.* **8**, eabm6860 (2022).
6. E Fischer, S Sippel, R Knutti, Increasing probability of record-shattering climate extremes. *Nat. Clim. Chang.* **11**, 689–695 (2021).
7. D Mitchell, et al., Attributing human mortality during extreme heat waves to anthropogenic climate change. *Environ. Res. Lett.* **11**, 074006 (2016).
8. AM Vicedo-Cabrera, et al., The footprint of human-induced climate change on heat-related deaths in the summer of 2022 in Switzerland. *Environ. Res. Lett.* **18**, 074037 (2023).
9. TM Beck, et al., Mortality burden attributed to anthropogenic warming during Europe's 2022 record-breaking summer. *npj Clim. Atmospheric Sci.* **7**, 245 (2024).
10. SE Perkins-Kirkpatrick, et al., On the attribution of the impacts of extreme weather events to anthropogenic climate change. *Environ. Res. Lett.* **17**, 024009 (2022).
11. M Allen, Liability for climate change. *Nature* **421**, 891–892 (2003).
12. CW Callahan, JS Mankin, Carbon majors and the scientific case for climate liability. *Nature* **640**, 893–901 (2025).
13. I Noy, et al., Event attribution is ready to inform loss and damage negotiations. *Nat. Clim. Chang.* **13**, 1279–1281 (2023).
14. NS Diffenbaugh, D Singh, JS Mankin, Unprecedented climate events: Historical changes, aspirational targets, and national commitments. *Sci. Adv.* **4**, eaao3354 (2018).
15. C Callahan, et al., Intensifying risk of mass human heat mortality if historical weather patterns recur. *EarthArXiv preprint* (2025).
16. K Rennert, et al., Comprehensive evidence implies a higher social cost of CO_2. *Nature* **610**, 687–692 (2022).
17. T Carleton, et al., Valuing the global mortality consequences of climate change accounting for adaptation costs and benefits. *The Q. J. Econ.* **137**, 2037–2105 (2022).
18. A Fouillet, et al., Has the impact of heat waves on mortality changed in France since the European heat wave of summer 2003? A study of the 2006 heat wave. *Int. J. Epidemiol.* **37**, 309–317 (2008).
19. M Burke, et al., Are we adapting to climate change? *Natl. Bureau Econ. Res. Work. Pap.* (2024).
20. A Gasparrini, et al., Temporal variation in heat–mortality associations: A multicountry study. *Environ. Heal. Perspectives* **123**, 1200–1207 (2015).
21. A Barreca, K Clay, O Deschenes, M Greenstone, JS Shapiro, Adapting to climate change: The remarkable decline in the US temperature-mortality relationship over the twentieth century. *J. Polit. Econ.* **124**, 105–159 (2016).
22. C Salesse, Heat waves, mortality and adaptation in France. *working paper* (2024).
23. CR Lay, et al., City-level vulnerability to temperature-related mortality in the USA and future projections: A geographically clustered meta-regression. *The Lancet Planet. Heal.* **5**, e338–e346 (2021).
24. K Laaidi, et al., The impact of heat islands on mortality in Paris during the August 2003 heat wave. *Environ. Heal. Perspectives* **120**, 254–259 (2012).
25. B Dousset, et al., Satellite monitoring of summer heat waves in the Paris metropolitan area. *Int. J. Climatol.* **31**, 313–323 (2011).
26. A Gasparrini, B Armstrong, The impact of heat waves on mortality. *Epidemiology* **22**, 68–73 (2011).
27. Y Guo, et al., Heat wave and mortality: a multicountry, multicommunity study. *Environ. Heal. Perspectives* **125**, 087006 (2017).
28. S Hajat, et al., Impact of high temperatures on mortality: is there an added heat wave effect? *Epidemiology* **17**, 632–638 (2006).
29. J Luterbacher, D Dietrich, E Xoplaki, M Grosjean, H Wanner, European seasonal and annual temperature variability, trends, and extremes since 1500. *Science* **303**, 1499–1503 (2004).
30. D Barriopedro, EM Fischer, J Luterbacher, RM Trigo, R García-Herrera, The hot summer of 2010: Redrawing the temperature record map of Europe. *Science* **332**, 220–224 (2011).
31. PA Stott, DA Stone, MR Allen, Human contribution to the European heatwave of 2003. *Nature* **432**, 610–614 (2004).
32. A Fouillet, et al., Excess mortality related to the August 2003 heat wave in France. *Int. Arch. Occup. Environ. Heal.* **80**, 16–24 (2006).
33. A Le Tertre, et al., Impact of the 2003 heatwave on all-cause mortality in 9 French cities. *Epidemiology* **17**, 75–79 (2006).
34. JM Robine, et al., Death toll exceeded 70,000 in Europe during the summer of 2003. *Comptes Rendus. Biol.* **331**, 171–178 (2008).
35. RC Keller, *Fatal isolation: The devastating Paris heat wave of 2003.* (University of Chicago Press), (2015).
36. S Vandentorren, et al., August 2003 heat wave in france: risk factors for death of elderly people living at home. *The Eur. J. Public Heal.* **16**, 583–591 (2006).
37. M Pascal, et al., France's heat health watch warning system. *Int. J. Biometeorol.* **50**, 144–153 (2006).
38. JT Trok, EA Barnes, FV Davenport, NS Diffenbaugh, Machine learning–based extreme event attribution. *Sci. Adv.* **10**, eadl3242 (2024).
39. A Gasparrini, et al., Mortality risk attributable to high and low ambient temperature: a multicountry observational study. *The Lancet* **386**, 369–375 (2015).
40. O Deschenes, E Moretti, Extreme weather events, mortality, and migration. *The Rev. Econ. Stat.* **91**, 659–681 (2009).
41. IPCC, Summary for policymakers in *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, eds. V Masson-Delmotte, et al. (Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA), (2021).
42. RA Betts, et al., Approaching 1.5°C: How will we know we've reached this crucial warming mark? *Nature* **624**, 33–35 (2023).
43. NS Diffenbaugh, EA Barnes, Data-driven predictions of peak warming under rapid decarbonization. *Geophys. Res. Lett.* **51**, e2024GL111832 (2024).
44. MI Bittner, EF Matthies, D Dalbokova, B Menne, Are European countries prepared for the next big heat-wave? *The Eur. J. Public Heal.* **24**, 615–619 (2014).
45. JF Dhainaut, YE Claessens, C Ginsburg, B Riou, Unprecedented heat-related deaths during the 2003 heat wave in Paris: consequences on emergency departments. *Critical care* **8**, 1 (2003).
46. K Dear, G Ranmuthugala, T Kjellström, C Skinner, I Hanigan, Effects of temperature and ozone on daily mortality during the August 2003 heat wave in France. *Arch. Environ. & Occup. Heal.* **60**, 205–212 (2005).
47. L Filleul, et al., The relation between temperature, ozone, and mortality in nine French cities during the heat wave of 2003. *Environ. Heal. Perspectives* **114**, 1344–1347 (2006).
48. D Jacob, D Winner, Effect of climate change on air quality. *Atmospheric Environ.* **43**, 51–63 (2009).
49. S Bartusek, K Kornhuber, M Ting, 2021 North American heatwave amplified by climate change-driven nonlinear interactions. *Nat. Clim. Chang.* **12**, 1143–1150 (2022).
50. E Bercos-Hickey, et al., Anthropogenic contributions to the 2021 Pacific Northwest heatwave. *Geophys. Res. Lett.* **49**, e2022GL099396 (2022).
51. RC Cornes, G van der Schrier, EJ van den Besselaar, PD Jones, An ensemble version of the E-OBS temperature and precipitation data sets. *J. Geophys. Res. Atmospheres* **123**, 9391–9409 (2018).
52. C3S, ERA5: Fifth generation of ECMWF atmospheric reanalysis of the global climate. *Copernicus Clim. Chang. Serv. Clim. Data Store* (2017).
53. Y Guo, et al., Global variation in the effects of ambient temperature on mortality: a systematic evaluation. *Epidemiology* **25**, 781–789 (2014).
54. S Hajat, T Kosatky, Heat-related mortality: a review and exploration of heterogeneity. *J. Epidemiol. & Community Heal.* **64**, 753–760 (2010).
55. KE Trenberth, JT Fasullo, TG Shepherd, Attribution of climate extreme events. *Nat. Clim. Chang.* **5**, 725–730 (2015).
56. TG Shepherd, A common framework for approaches to extreme event attribution. *Curr. Clim. Chang. Reports* **2**, 28–38 (2016).