With or Without U? Biased Effects of Temperature Extremes in the Presence of Global Warming*

Benjamin F. Jones† Northwestern Kellogg
Jacob Moscona‡ MIT
Benjamin A. Olken§ MIT
Cristine von Dessauer¶ MIT

March 26, 2026

Abstract

Estimates of climate impacts often produce U-shaped results, where extremely high and low temperatures have the largest effects. We show that these results can be spurious. Specifically, warming induces trends in extreme temperature exposure that are mechanically correlated with a location's baseline temperature. U-shaped bias can then emerge if an outcome variable also has trends that are related to baseline temperature, for any reason. We demonstrate this problem in simulations and with real outcomes. We then develop solutions. In applications using US data, some results in the literature are unaffected by these corrections, while other results change substantially.

*We thank Isaiah Andrews, Olivier Deschenes, Raffaele Ferrari, Peter Hull, Peter Huybers, Jesse Shapiro, Joe Shapiro, Jonathan Roth, Wolfram Schlenker, Andy Solow, Jim Stock, and Mengze Wang for helpful comments and suggestions. We thank Harufumi Nakazawa for outstanding research assistance.

†Northwestern Kellogg School of Management. Email: bjones@kellogg.northwestern.edu.
‡MIT, Department of Economics. Email: moscona@mit.edu.
§MIT, Department of Economics. Email: bolken@mit.edu.
¶MIT, Department of Economics. Email: cristine@mit.edu.

1 Introduction

Rising attention to climate change and concerns about the potential economic consequences of global warming have spurred extensive research on temperature's economic impacts. Since temperature variation across places can be correlated with other factors, economists have turned to panel models to identify temperature's effects (e.g., Deschênes and Greenstone 2007; Dell et al. 2012; see Dell et al. 2014 for a review). The panel approach leverages fluctuations in temperature over time, within a given place, to isolate idiosyncratic variation in temperature and provide causal inference.

When implementing these approaches, researchers often consider potential nonlinearities, where exposure to temperature extremes could be especially important. One common approach, for example, counts the number of days a location experiences each year in distinct temperature bins (e.g., the number of days $> 90°$F, the number of days $80 - 90°$F, and so on) and estimates a panel regression with these day counts as regressors. A recurrent finding across a broad range of outcomes is that temperature impacts are driven by extreme hot days and/or cold days, often generating "U-shaped" or "inverse-U-shaped" effects (e.g., Ranson, 2014; Graff Zivin and Neidell, 2014; Barreca et al., 2016; Hsiang et al., 2017; Park et al., 2020; Cohen and Dechezleprêtre, 2022; Carleton et al., 2022). For example, exposure to both very hot and very cold days have been shown to substantially lower crop yields (e.g., Schlenker and Roberts, 2009) and increase mortality (e.g., Deschênes and Greenstone, 2011). These non-linear estimates have become influential; for example, they have been used to assess the social cost of carbon (EPA, 2023).

In this paper, we consider a serious bias that can emerge in these regressions that has the curious result of producing spurious,"U shaped" patterns of the type frequently observed in the literature. Specifically, under general warming, exposure to extreme hot and cold temperatures will mechanically exhibit place-specific trends tied to a location's initial temperature. Warmer locations will accumulate more hot days over time than colder locations, and colder locations will lose more cold days than warmer ones. As a result, compared to cold areas, warm areas will exhibit relative upward trends in exposure to both hot and cold days.[1] When the outcome variable also has place-specific trends that are correlated with initial temperature, for any reason, estimates of the effects of temperature extremes become biased. Larger outcome trends in initially hot places thus

[1]The relative "upward" trend in cold days is subtler to see; it is because cold places see a faster downward trend than hot places in the number of cold days.

generate a spurious U-shaped relationship, with both hot days and cold days appearing to be positively associated with the outcome. Smaller outcome trends in initially hot places instead produce an inverted U shape. We refer to this phenomenon as *U Bias*.

We study this bias in three steps. First, we illustrate the bias and its broad scope. Using real world temperature data from U.S. counties, we show substantive, place-specific trends in exposure to temperature extremes that are mechanically related to a location's baseline temperature. We then study panel models with temperature bins, and show using simulated outcome data that we indeed obtain spurious U-shaped effects of temperature extremes, the same pattern that is common in the literature, even when the there is no actual relationship between contemporaneous temperature and the outcome. While we focus on binned temperature regressions in our main analysis, given their prominence in the literature, we also show that the bias extends to other non-linear methods for measuring temperature exposure, including polynomials of daily temperatures (e.g., Carleton et al., 2022) and so-called 'killing degree days' (e.g., Butler and Huybers 2013). We further document that trends in several outcome variables studied in the literature are sufficiently associated with local temperatures to generate substantial bias. We also show that U bias can generate spurious findings of "adaptation," which is often inferred from differential effects of temperature shocks in hot versus cold locations (e.g., Hultgren et al., 2025).

Second, we explore potential solutions. One approach focuses on counterfactual controls. By controlling directly for the expected exposure to each part of the temperature distribution, one can isolate random variation in temperature exposures. Implementing this procedure requires constructing the counterfactual set of possible temperature realizations in each place and time. This is not trivial, since each location has a different temperature distribution and trend; indeed, constructing such distributions is an active and unresolved area of research in climate science (e.g., Wang et al., 2024). Nevertheless, using location-by-month temperature distribution data, we develop a strategy that can capture changing exposure to extreme temperatures and account for U bias in simulations.

An alternative approach leverages short-panel variation by including location-by-time-window fixed effects. This approach isolates the identifying variation through temperature fluctuations across nearby years, where variation in the idiosyncratic component of temperature is more likely to dominate any trend component. This approach also performs well, though it requires higher-frequency data, making its breadth of application more limited. Finally, we explore the potential benefits of other controls that are variously

included in other temperature damage studies, including lags of the dependent and independent variables, broader location-specific time fixed effects (e.g., state-by-year fixed effects in a county-year panel), or location-specific linear trends. By and large, these sets of additional controls do not successfully address the bias we uncover.

Third, we revisit several results in the literature and investigate how correcting for this bias affects the estimates. We examine four outcome variables at the level of U.S. counties: population counts in each decade from the U.S. Census; annual mortality rates from the Centers for Disease Control; crop yields from the U.S. Department of Agriculture National Agricultural Statistics Service; and crime data from the FBI compiled by Ranson (2014). Correcting for bias has the potential to affect substantive conclusions, but results vary depending on the application. For example, the U-shaped relationship between temperature and mortality falls substantially when we account for bias. Extreme heat no longer raises all-cause mortality and extreme cold has a positive but smaller effect. On the other hand, the negative effects of extreme heat on crop yields is unaffected. The point of these exercises is not to establish the final word on any of these outcomes. Rather, these analyses illustrate how U bias can meaningfully affect empirical estimates by masking real effects or generating spurious ones.

2 The Core Idea: U Bias

In this section, we first describe the baseline approach used to identify non-linear effects of temperature. We then explain how general warming leads to differential trends across locations in exposure to temperature extremes. Finally, we show how these mechanical trends can generate "U-shaped" bias in temperature damage estimates.

Consider a panel regression that estimates the non-linear effects of temperature by counting days in different temperature "bins" (i.e., days below 10°F, days from 10°-20°F, etc.) and including these counts as regressors. This approach engages the natural possibility that hot days, temperate days, and cold days may have quite different effects.

Specifically, consider the following equation:

$$Y_{ct} = \alpha_c + \lambda_t + \sum_k \beta_k \times T_{kct} + u_{ct} \tag{1}$$

where Y_{ct} is the outcome variable for place c at time t, α_c are place fixed effects, λ_t are

time fixed effects, and T_{kct} represent counts of days in temperature bin k in a given time and place. If the number of days in each bin, location, and time period are as-good-as random conditional on fixed effects α_c and λ_t, then the β_k provide causal estimates of how the number of days in any given bin affects the outcome.

Existing work often finds large effects of temperature extremes, typically with U-shaped patterns. That is, models like (1) often estimate large positive or negative values of β_k for the highest and lowest bins. However, as we now show, these U-shaped patterns can be spurious. While we focus on the temperature-bin specification for simplicity and given its prevalence, all our conclusions are similar for other non-linear approaches (see Section 3.3). We next explain the intuition.

First, consider Figure 1a. This figure illustrates how warming shifts temperature distributions. We show a hot place and a cold place, with hot and cold temperature bins demarcated by vertical lines. We see that, even with uniform warming in both locations, there is a larger increase in exposure to very hot days in the hot location. The hot bin intersects with the dense part of the hot place's temperature distribution, while very hot days remain rare in the cold location. As temperatures shift to the right, this creates a mechanical, positive relationship between initial temperature and trends in exposure to extreme heat. The same argument holds, in reverse, for exposure to extreme cold, whereby the cold location sees a larger decline in exposure to cold days over time because the cold temperature bin intersects the thick part of the cold place's temperature distribution.

Figure 1b demonstrates such differential trends studying all U.S. counties, using real-world temperature data.[2] For each county, we consider the change in the number of very cold days (below 10°F) and very hot days (above 90°F), plotted against the county's mean temperature during the 1970s. Taking a cold county in the 1970s, we see a large decline in the number of very cold days over time but little or no increase in the count of very hot days. Taking a hot county in the 1970s, we see little or no decrease in very cold days but a large increase in very hot days. The result is differential trends in exposure to extreme temperatures that are correlated with a place's baseline temperature.

The less obvious but key observation here is that, compared to a colder place, a hotter place will exhibit relatively more hot days *and* relatively more cold days with time. The latter point is less intuitive, but follows (see Figure 1b) because the cold place's number

[2]We use the ERA5-Land database (Muñoz-Sabater et al., 2021) as the main source of temperature data throughout the paper. We also replicate all results using daily U.S. county-level data from Schlenker (2024) and raw station level data from the Global Historical Climatology Network (GHCN) (Menne et al., 2012).

of cold days are declining while the hot place's cold days are not. As a result, initially hot places appear relatively more exposed to both hot and cold temperature extremes with time.

Now return to the panel model in equation (1). We have just seen that warming induces a mechanical, positive relationship between local temperature and trends in T_{kct} for the extreme hot and cold bins. This will lead to biased estimates of β_k whenever outcomes trend in ways that are correlated with local temperature, for whatever reason. For example, let's say that outcomes are rising more in places that were initially hot. Then estimating (1) would yield a U shape relationship, even when in fact there is no relationship between contemporaneous temperature and the outcome. Conversely, if outcomes are rising more in places that were initially cold, we would end up with an inverted U shape. In both cases, we seem to have non-linear findings, but instead we have a potentially powerful form of non-linear bias that can produce entirely spurious results. We demonstrate this problem empirically in the following section.[3]

3 Simulations

3.1 Demonstrating U Bias with simulated outcomes and real temperature data

To illustrate the potential for spurious results, we consider what happens when we estimate equation (1) using real temperature data in U.S. counties (1970-2020) and simulated outcome data. We use the actual ERA-5-Land temperature data and count the days in each county-year in each temperature bin. For the outcome, we create data where, by construction, there is no relationship between annual temperature realizations and outcomes. The outcome variable is:

$$Y_{ct} = t \cdot Temp_c^{Base} + \epsilon_{ct} \tag{2}$$

where ϵ_{ct} is drawn from a normal distribution centered at zero and with variance σ_ϵ^2.[4] We estimate the baseline panel model (1) for 1,000 draws of the outcome and plot the

[3] We further demonstrate the existence of U-bias using a simple econometric theory (Appendix Section A).

[4] We set σ_ϵ^2 to be equal to the variance of $t \cdot Temp_c^{Base}$ over the entire sample, so that the error term in the outcome is of similar variance to trends, but this is not essential.

resulting distributions of the β_k.

The estimates imply positive effects of exposure to temperature extremes—both hot and cold days—despite there being no true relationship between between temperature bins and the outcome (Figure 2a). Reversing the sign of the trend in the outcome leads to inverse-U-shaped results (Figure 2b). In both cases, these "results" are entirely spurious.

A key ingredient to the bias is that the outcome has a trend that is correlated with baseline temperature: removing the differential trend in the outcome eliminates the bias (Figure 2c). Thus, the presence of U-bias will depend on the outcome being studied. On the other hand, the differential trends in the temperature variables arise mechanically. These differential trends are not due to some idiosyncratic features of warming patterns across space—even *uniform warming* with would lead to U bias (Appendix B.1).

More generally, the magnitude of the spurious effects depends on the magnitude of the outcome trends. To illustrate this phenomenon in the data, we simulate alternative outcomes using equation (2) in which we vary the absolute value of the coefficient in front of $Temp_c^{Base}$.[5] Spurious effects of temperature extremes are more pronounced for larger outcome trends. This motivates a potential diagnostic test for bias: bias is likely to be stronger (weaker) if the correlation between outcome trends and location-level temperature is stronger (weaker).

To gauge whether this bias could be quantitatively meaningful in realistic settings, we run additional simulations where we calibrate the magnitude of differential trends in our simulated outcomes to those of variables that have been studied in the context of climate damage estimates (Figure A.4), including mortality, crime rates, and crop yields.[6] We find that these calibrated outcome trends are sufficient to generate large, spurious findings. For example, the magnitude of the bias for the outcome calibrated to mortality trends is of similar size to effects reported in the literature (Deschênes and Greenstone, 2011) (see Figure A.4a). These results (thus far) do not necessarily imply that existing estimates are biased; instead, they suggest that the bias we identify above could be large enough to explain effects in prior work. While we find spurious U-shaped effects in all cases, the magnitudes vary across outcomes. For example, spurious effects for trends calibrated to crop yields are by far the smallest. These differences suggest that bias could affect estimates differently depending on the differential trend in the outcome variable.

[5] We add a coefficient β in front of t, which we we let range from 0.0001σ to 0.01σ, where $\epsilon_{ct} \sim \mathcal{N}(0, \sigma)$

[6] Details of how we do this calibration can be found in Appendix B.2.

3.2 Heterogeneous effects and "adaptation"

There is growing interest in potential mitigating effects of human adaptation to climate change. A common strategy for studying adaptation estimates the effects of extreme heat separately for already-hot locations, which have plausibly learned to adapt, and for colder locations, which have not (e.g., Butler and Huybers, 2013; Carleton et al., 2022; Hultgren et al., 2025). Weaker marginal effects of hot days in hotter locations are interpreted as adaptation. However, the bias we document here—and the temperature bins in which the bias is most severe—may differ systematically between hotter and colder places (or more generally, between any two sub-samples). This can generate spurious results that look like "adaptation" but instead reflect differences in trend parameters.

To investigate this possibility, we estimate county-specific trends for cold days (days below 10°F) and hot days (days above 90°F) and relate these trends to baseline temperature separately for counties with below and above median baseline temperature. Figure 3a plots these relationships for trends in cold days. Cold places (shown in blue) have a steeper slope in baseline temperature compared to the hot places (shown in red). Figure 3b shows that the converse is true for hot days. Thus, the key trends vary systematically between hot and cold locations, and for the hot and cold temperature bins.

We next return to the simulations above and report estimates of equation (1) separately for counties with above versus below median values of $Temp_c^{Base}$. The results are reported in Figure 3c. We find a larger spurious "effect" of extreme cold in hotter counties and, to a lesser extent, a larger "effect" of extreme heat in colder counties. This would be interpreted as evidence of adaptation but is driven entirely by differences in the extent of bias across sub-samples. Alternative specifications show even more obvious (but spurious) evidence of adaptation (Figure A.5a and A.5b) in both hot and cold places.

3.3 Sensitivity and Extensions

We present a series of sensitivity checks of our main simulation analysis. First, we show that our results extend to other non-linear parameterizations of temperature. This includes "killing degree days" (Figure A.6c), which count the number of degree days in excess of some temperature threshold and are commonly used when studying the exposure of crops to extreme heat (e.g., Butler and Huybers, 2013); using polynomials in daily temperature realizations (Figure A.6d), which some papers use instead of binning to reduce the number of estimated parameters (e.g., Carleton et al., 2022); and using either

finer or coarser temperature bins (Figures A.6a and A.6b);

We additionally consider other units of aggregation beyond the U.S. county-year level. In some cases researchers aggregate either spatially (e.g., to states) or temporally (e.g., to decades). We find similar levels of bias when we run the simulations at the state-year level (Figure A.7a) or county-decade level (Figure A.7b). Finally, we find similar results using alternative sources of temperature data, including Schlenker (2024)'s balanced panel of weather stations from PRISM (Figure A.7c) and raw station level data from GHCN (Menne et al., 2012) (Figure A.7d).[7]

4 Solutions

In this section, we investigate potential solutions to U bias. In Section 4.1, we consider commonly-used control strategies that are prevalent in existing work and investigate if they remove the bias. In Section 4.2, we develop new approaches that are designed to directly remove U bias. For each potential solution, we return to the simulation methods in Section 3. For a correction to be successful, the coefficient estimates should be centered at zero; if unsuccessful, the spurious U-shaped relationship persists.

4.1 Existing Control Strategies

Existing papers often include additional controls beyond the place and year fixed effects in equation (1). To the extent that these controls successfully absorb trends in the right- or left-hand-side variables that drive bias, they could present a potential solution.

State-by-year fixed effects. Some studies include state-by-year fixed effects when analyzing the county-level effects of temperature.[8]. While the bias is somewhat attenuated, state-by-year fixed effects do not eliminate spurious effects of temperature in simulations (see Figure A.8a). The bottom and top bins falsely reject zero (p <0.05) in 14% and 26% of the simulation runs, respectively (Figure A.12). Furthermore, the reduction in false-rejections comes primarily from the fact that state-year fixed effects increase noise rather than reducing bias in the coefficients.

[7]We aggregate these station data to the county level following Deschênes and Greenstone (2011).
[8]There is debate over whether including such fixed effects is best practice (see, e.g., Fisher et al., 2012)

Lags of dependent and independent variables. Some studies include lags of the dependent and independent variables as controls (e.g., Nath et al., 2024). However, including additional lags as controls does not address U bias (Figure A.8b). The coefficients on the top and bottom temperature bins both falsely reject zero over 90% of the time (Figure A.12).[9]

Place-specific trends. Some studies include location-specific trends as controls. For example, in a county-year panel, one could include a separate linear trend for each county. These are more common in estimates of the *linear* effect of temperature or studies focusing on more aggregated geographic units (e.g., Burke et al., 2015). The inclusion of county-specific linear trends solves the problem in our baseline simulation (Figures A.8c and A.12); however, this is precisely because the differential trend in the outcome is itself linear. If one does not know the functional form of the outcome trend, it is not clear that place-specific *linear* trends will work. For example, if we conduct modified simulations where the outcome trends in (2) are quadratic (Figure A.11) or cubic (Figure A.14), the inclusion of place-specific linear trends does not perform as well.

While one could always include higher-order trends, there are two key challenges. First, this only works if the researcher is confident that they know the functional form of the trend in the outcome variable. In real-world applications, this functional form is unknown. Second, it comes with a substantial degrees-of-freedom cost. In a county-by-year level regression, for example, adding county-specific linear trends requires adding over 3,000 controls; place-specific higher order trends would add commensurately more.

4.2 Alternative Approaches

4.2.1 Controlling for non-random temperature exposure

Our first alternative strategy seeks to control explicitly for the non-random component of temperature exposure. To do so, we build a statistical model of the counterfactual temperature distribution in each time and place. Using this counterfactual distribution, the researcher can calculate the expected value of each temperature bin variable and then add these as controls to the panel regression. The identifying variation that is left is the random component of the realization—the idiosyncratic shocks that motivated the panel approach in the first place. Importantly, this approach does not require knowing

[9] Intuitively, lags can help account for *common* dynamics but not *place-specific* dynamics or trends.

anything about the functional form of the dependent variable; is more parsimonious than the control strategies presented above; and could in theory be applied in any context. The main challenge is correctly estimating counterfactual temperature distributions that can be used to construct the control variables.[10]

In particular, we can estimate the augmented version of equation (1) that includes non-random exposure to each temperature bin, $\widehat{T}_{kct}$, as controls:

$$Y_{ct} = \alpha_c + \lambda_t + \sum_k \beta_k T_{kct} + \sum_k \xi_k \widehat{T}_{kct} + u_{ct} \tag{3}$$

The stochastically realized number of days in each bin, T_{kct}, is random conditional the expected number of days in each bin, $\widehat{T}_{kct}$, which captures the fact that rising temperatures cause differential trends in different temperature bins. Thus, as shown by Borusyak and Hull (2023) for example, estimates of the β_k from equation (3) are unbiased.[11] We verify in simulations where we know the DGP (i.e. where we simulate both outcome variables and temperatures) that indeed, this approach works if one can control for the correct $\widehat{T}_{kct}$ (Appendix B.1.2). The challenge, then, is constructing $\widehat{T}_{kct}$ empirically.

Constructing counterfactual temperature distributions. To apply the method above with real-world temperature data, we build a statistical model of changes in expected exposure to each temperature bin, exploiting seasonal difference in both temperature distributions and warming trends across space.[12,13]

Our procedure is as follows. The first goal is to capture a baseline temperature distribution (independent from warming) for each county using all the data from our sample

[10] An alternative approach may have been to use an off-the-shelf climate model built to simulate the distribution of possible temperatures in each place as a function of the global mean temperature. However, these "emulators" are typically designed to forecast possible future climate, rather than being purpose-built to back-cast possible counterfactual realizations, and tend to not perform as well for extreme temperatures (e.g., Wang et al., 2024). Consistent with this, state-of-the-art emulator results from Wang et al. (2024) perform poorly in our context when used to construct $\widehat{T}_{kct}$. These results are available from the authors upon request.

[11] Borusyak and Hull (2020) note weather as a potential application of their re-centering approach to non-random right-hand side variables (see Appendix D.8 of the working paper version of their paper) when weather is stationary. However, a key point of our analysis is that the non-stationarity in temperature (i.e., global warming) is what generates differential trends and U bias. Therefore, the key challenge is to determine an appropriate counterfactual for weather, i.e. to construct $\widehat{T}_{kct}$, which we develop here.

[12] We thank climate scientist Peter Huybers for suggesting this approach.

[13] We provide code to create these counterfactuals using any daily temperature dataset using the *cftemp* Stata package, available on our websites.

period. To do so, first define x_{cdmt} as the observed temperature realization for county c, day d, month m and year t. We then construct the mean observed temperature, $\bar{x}_{cmt}$, for each county-month-year, and estimate the linear time trend in $\bar{x}_{cmt}$ separately for each county-month across all years. Specifically, we use the following regression:

$$\bar{x}_{cmt} = \mu_{cm} + \gamma_{cm} \cdot t + \epsilon_{cmt} \tag{4}$$

where $\hat{\gamma}_{cm}$ is the estimated linear trend coefficient for month m in county c.[14,15] Then, we subtract off the mean shift that has taken place as of t from the daily realizations x_{cdmt}:

$$x^{detrend}_{cdmt} = x_{cdmt} - \hat{\gamma}_{cm} \cdot t \tag{5}$$

Combining all de-trended temperature realizations generates our county-month level temperature distributions, $F_{cm}(x)$.

The second goal is to combine these distributions with county-month warming trends in order to generate the counterfactual temperature distributions for each year. To do this, we project each county-month distribution forward using its time trend $\hat{\gamma}_{cm}$. This generates an empirical distribution of potential temperature realizations for each county, month, and year, $F_{cmt}(x)$, which can be aggregated across months to generate the counterfactual distribution at the county-year level, $F_{ct}(x)$. We can then calculate the share of the distribution in each county-year that falls into each temperature bin. After normalizing by the number of days in the year, this generates $\widehat{T}_{kct}$.

Correcting for bias. We use the simulations from Section 3 to see whether controlling for $\widehat{T}_{kct}$ eliminates U bias. Adding $\widehat{T}_{kct}$ derived from the procedure above as controls, we find that the coefficient estimates are now all centered at zero (Figure 4a, shown in red squares). This indicates that controlling for the non-random component of exposure to each bin removed the bias. The extreme temperature coefficient estimates are centered around zero and reject the null hypothesis at the appropriate rate (Figure A.10).[16]

[14] In our main analysis, we use empirical Bayesian methods to shrink our estimates of the γ_{cm} toward the average estimate, though results are similar without this step.

[15] As an alternative with fewer parameters, rather than estimate a separate trend for each county-month in equation (4), we can use a county-specific fourth-order Chebychev polynomial to approximate seasonal variation in temperature (as in Schlenker and Roberts, 2006) and estimate the place-specific time trend for each of the four Chebychev polynomial terms. In practice, this approach performs similarly.

[16] The results are also similar if we exclude the trend shrinkage step or use Chebychev polynomials to capture seasonality (Figure A.10).

Conversely, when alternatively simulate an outcome where we do add causal effects of contemporaneous temperature on the outcome (of 5 for the bottom and top bin), this approach accurately recovers the true effect (Figure 4b, again shown in red squares).

This approach reduces bias *not* by reducing the residual trend in the outcome variable, which could change across applications, but rather by fixing the binning bias on the "right hand side" of the panel regression—by correcting for differential trends in exposure to extreme temperature. This suggests that this approach could be applicable across a broad range of contexts.

4.2.2 Isolating short-panel comparisons

A alternative approach isolates comparisons among nearby years. Specifically, one can include interactions between place and longer-time-period fixed effects; e.g. including county-by-five-year-interval fixed effects. This strategy exploits the fact that in a small time period, idiosyncratic temperature variation is likely to be large compared to the trend, so the bias becomes smaller in practice.[17]

We find that this approach entirely eliminates the bias in our baseline simulation (Figure A.8d, shown with green Xs). The simulated coefficient estimates on the top and bottom bin are centered around zero and reject the null at the appropriate rate (Figure A.12), including with more complicated—quadratic and cubic—trends in the outcome (Figures A.11 and A.14).

While effective, there are also potential limitations to this approach. First, the choice of time window remains *ad hoc*. Second, it is only relevant in applications that use higher-frequency (e.g., yearly) temperature variation since it requires exploiting variation *within* aggregated time windows. This would not be applicable in situations when the unit of observation is a location-by-decade (for example) or when using long-difference specifications. Some outcome data do not exist at an annual frequency (e.g., the U.S. Population Census, as we explore below, or the U.S. Census of Agriculture data, as in Moscona and Sastry, 2023). Moreover, aggregating across time is a common strategy to understand adaptation (e.g., Burke and Emerick, 2016). The short-panel approach would not help in such contexts.

[17] Longer panels allow trends to dominate the variation; this can also be seen in the theoretical result that longer panels make U-bias worse (Appendix A, Result 2).

5 Applications

The final step of our analysis is to investigate how correcting for U bias affects estimates of the relationship between temperature bins and a variety of real-world outcomes. We focus on four sets of outcomes that are relevant for understanding the economic consequences of climate change: county-level population, mortality, crop yields, and crime.

For consistency, we use the same specifications, temperature data (ERA5-land) and variable construction throughout.[18] We use data on each outcome from standard sources. We compile county-level population data from each round of the U.S. Census from 1970 to 2010; mortality data by county and age group from the Centers for Disease Control (CDC) compressed mortality files;[19] yields for corn, wheat, and soybeans at the county-year level from the U.S. Department of Agriculture National Agricultural Statistics Service; and crime at the count-year level using data compiled from FBI reports by Ranson (2014).[20]

For each outcome, we first estimate a standard, un-corrected specification, which includes counts of days in each temperature bin as well as county and time fixed effects (equation (1)). We then estimate augmented versions that include our suggested counterfactual controls for the expected number of days in each bin (equation (3), shown in red squares). Where possible, we also report the "short-panel" correction approach by including five-year bin indicators interacted with county fixed effects, shown in green Xs. The results are shown in Figure 5.

5.1 Population

We first examine total county-level population. The reallocation of population, including through migration, has long been considered an important potential mechanism of adaptation to environmental change (e.g., McLeman and Smit, 2006; Cruz and Rossi-Hansberg, 2024). We report the effect of temperature extremes on log population, both with and without our proposed solution to binning bias (Figure 5a).[21] Uncorrected esti-

[18] Replicating each paper's exact specification and exact temperature dataset for the hundreds of papers in this literature is beyond the scope of this paper. Rather, our goal here is to highlight the importance of ensuring binning bias is addressed moving forward.

[19] The public version of these data omits death counts for county-by-year-by-age group triplets when there are few deaths. We focus on over-65 mortality, for which only 0.6% of the observations is missing.

[20] We count the number of violent crimes—including murder, rape, and aggravated assault—and non-violent or property crimes—including robbery and larceny—in each U.S. county, month, and year.

[21] Our "short panel" approach is not possible in this context since the unit of observation for population is a county-decade, making it impossible to exploit variation within shorter periods.

mates show a U-shape, and suggest that exposure to both extreme cold and extreme heat is *positively* associated with increased population. However, when we correct for U bias, the estimates suggest a much flatter relationship between temperature and population, with the exception of a significant, negative effect of the top temperature bin. Thus, naive estimates suggest that population moves *toward* parts of the country with greater exposure to temperature extremes, but in fact, extreme positive heat in a decade actually reduce population, consistent with out-migration as an adaptation mechanism.

5.2 Mortality

The effect of temperature extremes on mortality has been studied extensively (e.g., Deschênes and Greenstone, 2011; Barreca et al., 2016; Carleton et al., 2022) and the estimated positive effect of extreme heat on mortality represents a large share of predicted economic damages from climate change (Hsiang et al., 2017). While naive estimates of the effect of temperature on mortality (i.e., the mortality rate for people age 65 and over) return the U-shaped relationship that is familiar from the literature, we find much weaker evidence of this pattern when we account for U bias (Figure 5b). The effects of extreme cold are substantially smaller in magnitude, and the effect of extreme heat disappears. The pattern is similar if we pool effects on mortality across all age groups (see Figure A.16a).[22] These results suggest that estimates of the effect of U.S. temperature extremes on mortality may be biased.

5.3 Crop Yields

The agricultural sector is perhaps most exposed to environmental change, and a large literature has established a stark, negative effect of extreme heat on crop yields (e.g., Schlenker and Roberts, 2009; Hultgren et al., 2025). We estimate the effect of temperature on corn yields and find a clear, inverted-U shaped relationship both in the native specification and corrected specifications (Figure 5c). The conclusion is similar for wheat and soybean yields (Figures A.16b and A.16c). Thus, estimates of the relationship between crop yields and temperature do not seem biased (in either direction) by spurious trends. This is consistent with our prior result that yield trends across counties are not

[22]The all-age group specification takes into account differing mortality trends and impacts across age groups (e.g., Deschênes and Greenstone, 2011). The pooled regression is the same specification as equation (1) but with separate observations for each age group, weighted by group-county-year population.

strongly correlated with temperature (see Section 3).

5.4 Crime

Weather has been shown to affect both violent and non-violent crime (e.g., Ranson, 2014; Cohen and Gonzalez, 2024). Following the convention in the literature, we estimate an augmented version of equation (1) where crime observations are at the county-year-month level and we include county, month, and year fixed effects. We construct analogous counterfactual measures of non-random exposure to each temperature bin. Naive estimates suggest that the violent crime rate is higher at high temperatures, with a weaker, negative effect of cold days. After correcting for bias, the positive effect of heat on crime is smaller while the negative effect of cold is substantially larger (Figure 5d). Thus, correcting for bias moderates one conclusion from the naive estimates (i.e., that heat leads increases crime) and strengthens another (i.e., that cold reduces crime). Turning to non-violent crime, naive estimates suggest that cold days reduce crime rates while hot days do the opposite. Once we correct for bias, the negative effect of cold days on non-violent crime is even larger while the positive effect of hot days disappears (Figure A.16d).

6 Conclusion

A large literature uses panel models to estimate temperature's non-linear effects on various economic outcomes. Common findings are U-shaped or inverted U-shaped relationships, where more extreme temperatures have larger effects. These estimates have become important for informing the impacts of temperature, quantifying potential damages from global climate change, and guiding policy intervention.

In this paper, we show that warming itself can bias these estimates – what we call "U bias" - by introducing non-random, place-specific trends in extreme temperature exposure. We show how generalized warming can lead to this non-linear form of bias and argue that it may be relevant across important applications. We then develop solutions, including approaches that isolate short-panel comparisons or control directly for the non-random component of temperature fluctuations. We apply our proposed corrections to a range of real-world outcomes. In some cases, results become substantially weaker when the correction is applied. In other cases, bias appears to have masked important results. And in other cases, results seem unaffected. Overall, addressing and correcting for U bias

appears important for obtaining reliable estimates of temperature effects.

References

Barreca, A., Clay, K., Deschenes, O., Greenstone, M., and Shapiro, J. S. (2016). Adapting to climate change: The remarkable decline in the us temperature-mortality relationship over the twentieth century. *Journal of Political Economy*, 124(1):105–159.

Borusyak, K. and Hull, P. (2020). Non-random exposure to exogenous shocks: Theory and applications. Working Paper 27845, National Bureau of Economic Research.

Borusyak, K. and Hull, P. (2023). Nonrandom exposure to exogenous shocks. *Econometrica*, 91(6):2155–2185.

Burke, M. and Emerick, K. (2016). Adaptation to climate change: Evidence from us agriculture. *American Economic Journal: Economic Policy*, 8(3):106–140.

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

Butler, E. E. and Huybers, P. (2013). Adaptation of us maize to temperature variations. *Nature Climate Change*, 3(1):68–72.

Carleton, T., Jina, A., Delgado, M., Greenstone, M., Houser, T., Hsiang, S., Hultgren, A., Kopp, R. E., McCusker, K. E., Nath, I., et al. (2022). Valuing the global mortality consequences of climate change accounting for adaptation costs and benefits. *The Quarterly Journal of Economics*, 137(4):2037–2105.

Cohen, F. and Dechezleprêtre, A. (2022). Mortality, temperature, and public health provision: evidence from mexico. *American Economic Journal: Economic Policy*, 14(2):161–192.

Cohen, F. and Gonzalez, F. (2024). Understanding the link between temperature and crime. *American Economic Journal: Economic Policy*, 16(2):480–514.

Cruz, J.-L. and Rossi-Hansberg, E. (2024). The economic geography of global warming. *Review of Economic Studies*, 91(2):899–939.

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

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

Deschênes, O. and Greenstone, M. (2007). The economic impacts of climate change: evidence from agricultural output and random fluctuations in weather. *American economic review*, 97(1):354–385.

Deschênes, O. and Greenstone, M. (2011). Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the us. *American Economic Journal: Applied Economics*, 3(4):152–185.

EPA (2023). Report on the social cost of greenhouse gases: Estimates incorporating recent scientific advances. Technical report, U.S. Environmental Protection Agency. Accessed: 2024-03-18.

Fisher, A. C., Hanemann, W. M., Roberts, M. J., and Schlenker, W. (2012). The economic impacts of climate change: evidence from agricultural output and random fluctuations in weather: comment. *American Economic Review*, 102(7):3749–3760.

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

Hsiang, S., Kopp, R., Jina, A., Rising, J., Delgado, M., Mohan, S., Rasmussen, D. J., Muir-Wood, R., Wilson, P., Oppenheimer, M., et al. (2017). Estimating economic damage from climate change in the united states. *Science*, 356(6345):1362–1369.

Hultgren, A., Carleton, T., Delgado, M., Gergel, D. R., Greenstone, M., Houser, T., Hsiang, S., Jina, A., Kopp, R. E., Malevich, S. B., et al. (2025). Impacts of climate change on global agriculture accounting for adaptation. *Nature*, 642(8068):644–652.

McLeman, R. and Smit, B. (2006). Migration as an adaptation to climate change. *Climatic change*, 76(1):31–53.

Menne, M. J., Durre, I., Korzeniewski, B., McNeill, S., Thomas, K., Yin, X., Anthony, S., Ray, R., Vose, R. S., Gleason, B. E., and Houston, T. G. (2012). Global historical climatology network - daily (ghcn-daily), version 3.32. [Accessed 2025-04-12].

Moscona, J. and Sastry, K. A. (2023). Does directed innovation mitigate climate damage? evidence from us agriculture. *The Quarterly Journal of Economics*, 138(2):637–701.

Muñoz-Sabater, J., Dutra, E., Agustí-Panareda, A., Albergel, C., Arduini, G., Balsamo, G., Boussetta, S., Choulga, M., Harrigan, S., Hersbach, H., et al. (2021). Era5-land: A state-of-the-art global reanalysis dataset for land applications. *Earth system science data*, 13(9):4349–4383.

Nath, I. B., Ramey, V. A., and Klenow, P. J. (2024). How much will global warming cool global growth? Technical report, National Bureau of Economic Research.

Park, R. J., Goodman, J., Hurwitz, M., and Smith, J. (2020). Heat and learning. *American Economic Journal: Economic Policy*, 12(2):306–39.

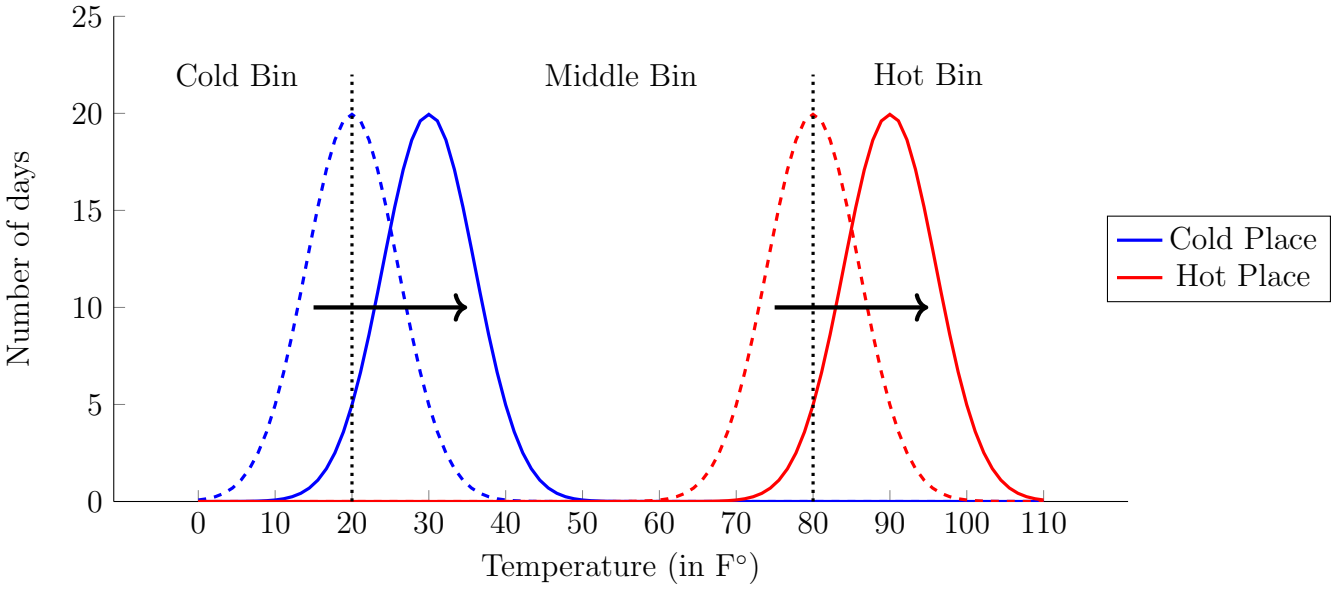
Ranson, M. (2014). Crime, weather, and climate change. *Journal of environmental economics and management*, 67(3):274–302.

Schlenker, W. (2024). Daily weather data - continental USA (version 2018).

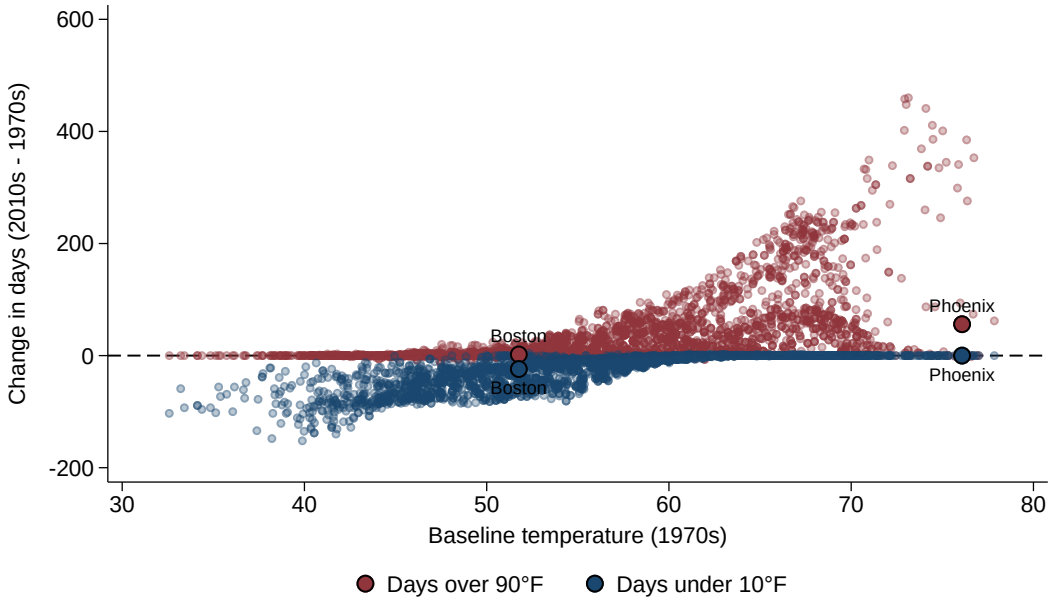
Schlenker, W. and Roberts, M. J. (2006). Nonlinear effects of weather on corn yields. *Review of agricultural economics*, 28(3):391–398.

Schlenker, W. and Roberts, M. J. (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.

Wang, M., Souza, A. N., Ferrari, R., and Sapsis, T. (2024). Stochastic emulators of spatially resolved extreme temperatures of earth system models. *Authorea Preprints*.



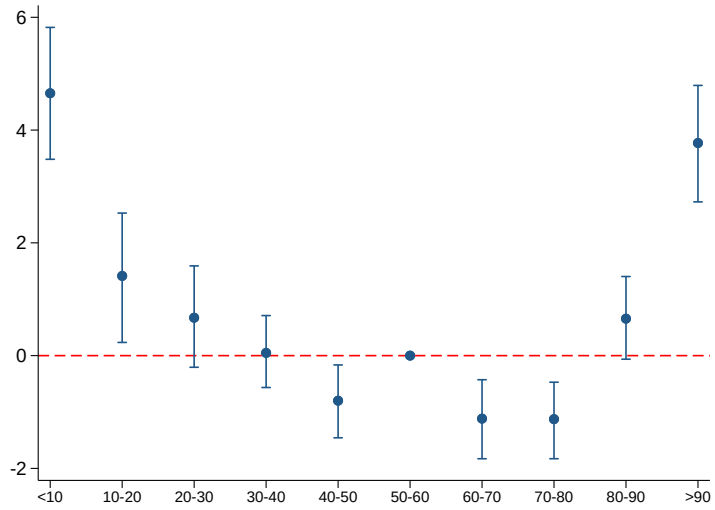
(a) Differential Trends in Temperature Extremes: Uniform Warming Diagram



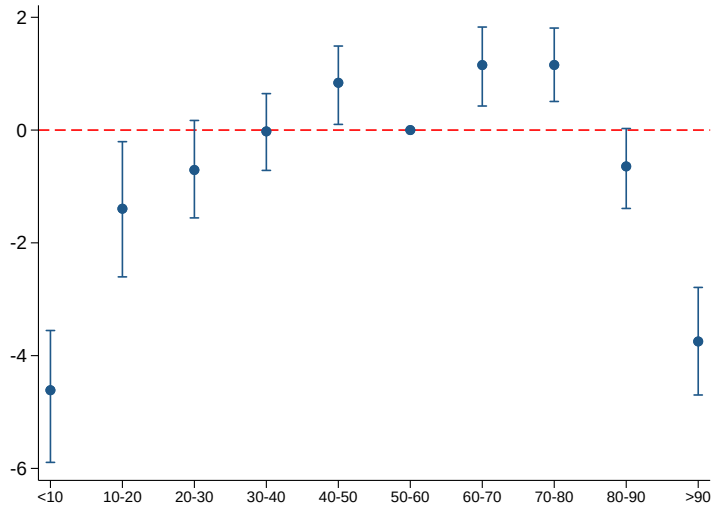
(b) Differential Trends in Temperature Extremes: U.S. County-Level Data

Figure 1: Differential Trends in Temperature Extremes

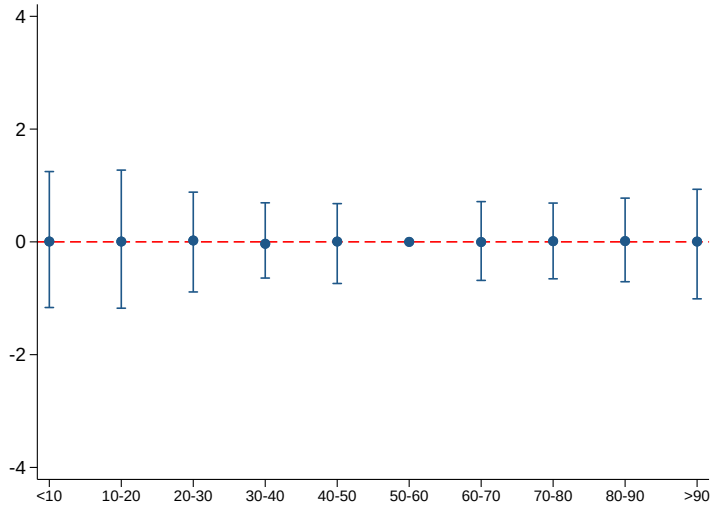
Notes:



(a) Positively Correlated Outcome Trend



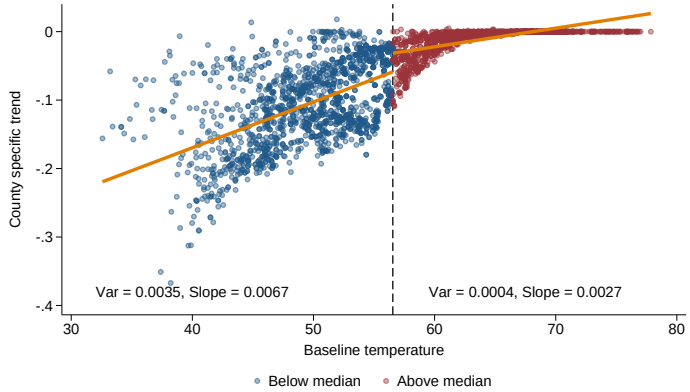
(b) Negatively Correlated Outcome Trend



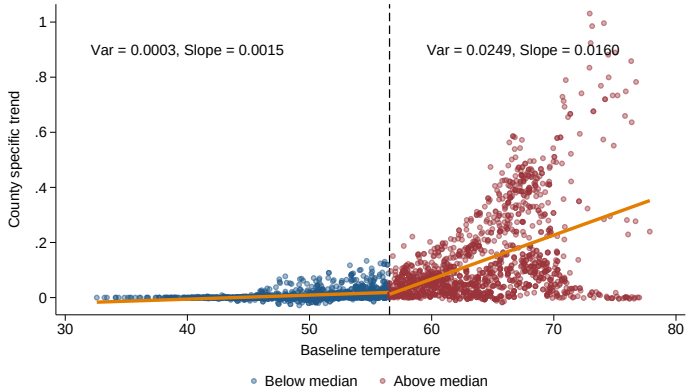
(c) Uncorrelated Outcome Trend

Figure 2: Binned Regression Estimates with Real Temperature Data and Dynamics

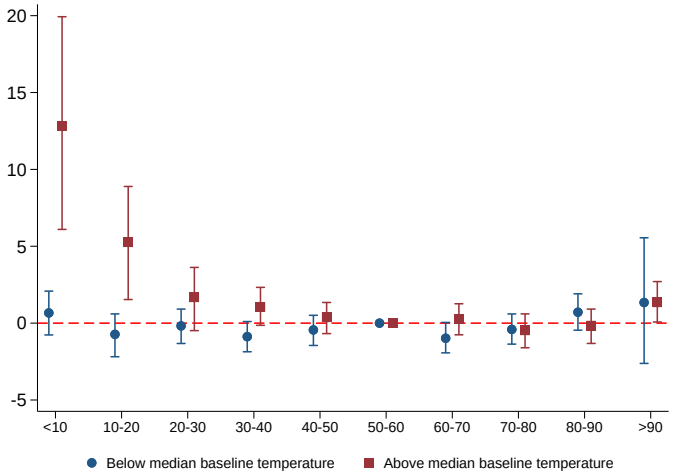
Notes: Each sub-figure reports estimates from equation (1), in which the outcome is simulated using the following equation: $Y_{ct} = aTemp_c^{base} + \epsilon_{ct}$. The error term is drawn from a normal distribution centered at zero with variance equal to the variance of $aTemp_c^{base}$ over the full sample. The temperature bins are constructed from historical, daily temperature realizations using the ERA-5 Land dataset. In each sub-figure, we report the mean and 95% confidence intervals of each coefficient from (1) taking 1,000 separate draws of the outcome variable. In Panel (a), the outcome is simulated to have a positive differential trend in baseline temperature (i.e., $a = t$); in Panel (b), the outcome is simulated to have a negative differential trend in baseline temperature (i.e., $a = -t$); and in Panel (c) the outcome has no relationship with baseline temperature (i.e., $a = 0$).



(a) County-specific Trends in Cold Days



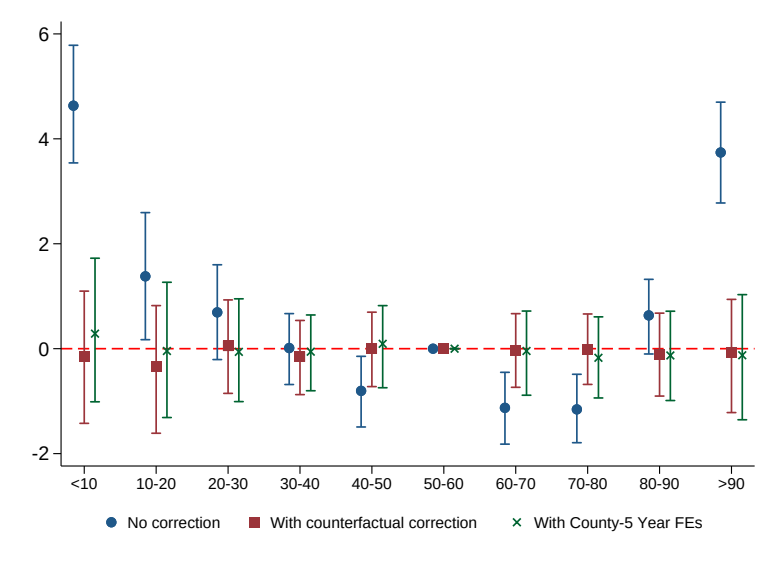
(b) County-specific Trends in Hot Days



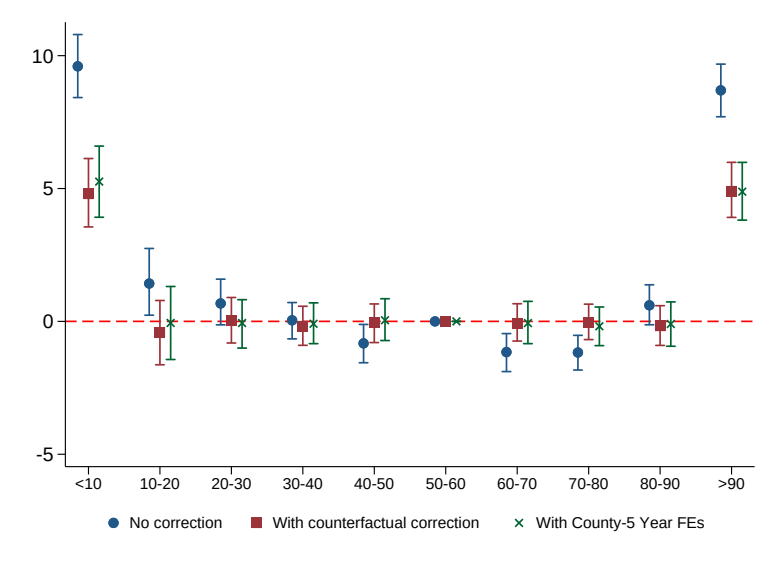
(c) Binned Regression Estimates

Figure 3: Binned Regression Estimates and "Adaptation"

Notes: Panel (a) reports the relationship between baseline temperature and county-specific trends (1970s-2010s) in the number of days below 10°F. We report this relationship separately for countries with above and below median average temperatures, and report the slope and trend variance for both sets of counties. Panel (b) reports the relationship between baseline temperature and county-specific trends (1970s-2010s) in the number of days above 90°F. We report this relationship separately for countries with above and below median average temperatures, and report the slope and trend variance for both sets of counties. Panel (c) reports estimates of equation (1), separately for counties with above versus below median baseline temperature. The outcome is simulated to have a positive differential trend in baseline temperature ($Y_{ct} = tTemp_c^{base} + \epsilon_{ct}$) and daily temperature realizations are from the ERA-5 Land dataset. We report the mean and 95% confidence intervals for each coefficient estimate, taking 1,000 draws of ϵ_{ct} from a normal distribution centered at zero with variance equivalent to the in-sample variance of $tTemp_c^{base}$.



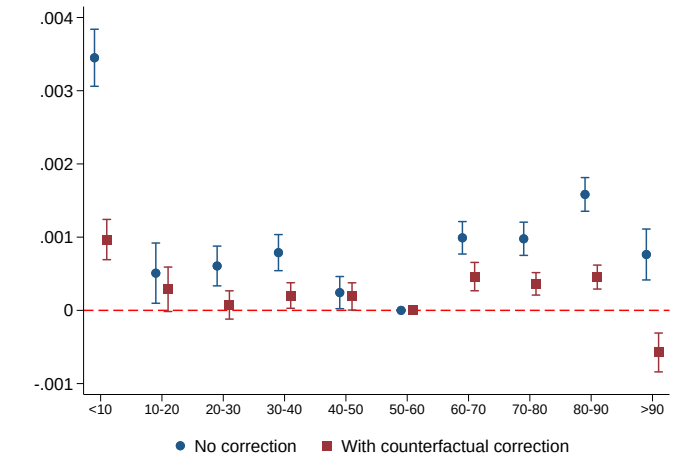
(a) Counterfactual Control Correction, No True Effect



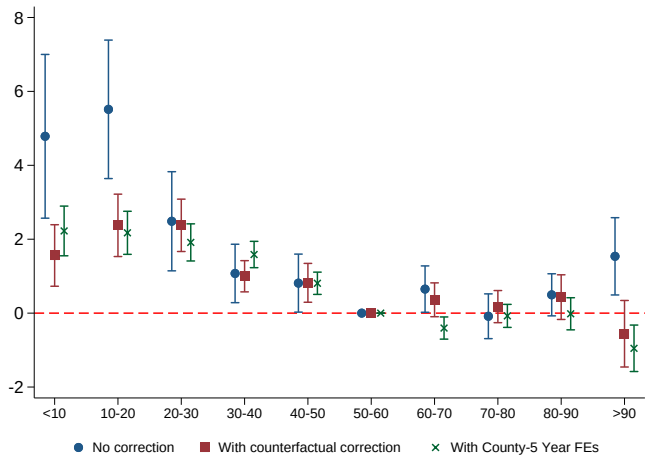
(b) Counterfactual Control Correction, True Effect of Extreme Temperatures

Figure 4: Correcting for Bias with Counterfactual Temperature Controls

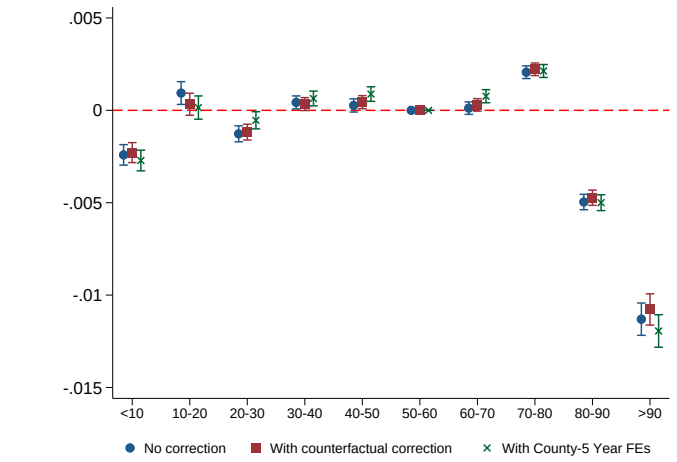
Notes: Both sub-figures report estimates of equation (1) alongside estimates of equation (3) that also include the i) expected counterfactual exposure to each temperature bin or ii) 5-year times county fixed effects as controls. In Panel (a), the outcome is simulated to have a positive differential trend in baseline temperature ($Y_{ct} = tTemp_c^{base} + \epsilon_{ct}$) but no causal effect of exposure to any temperature bin. In Panel (b), the outcome has the same differential trend but there is also a positive effect of 5 of exposure to the top and bottom temperature bins. Daily temperature realizations used to construct each temperature bin are from the ERA-5 Land dataset. For each regression specification, we report the mean and 95% confidence intervals for each coefficient estimate, taking 1,000 draws of $\epsilon_{c,t}$ from a normal distribution centered at zero with variance equivalent to the in-sample variance of $tTemp_c^{base}$.



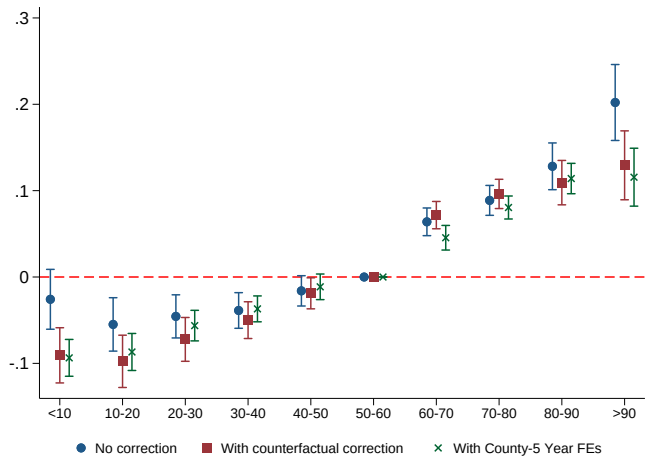
(a) log Population



(b) Over-65 Mortality Rate



(c) log Corn Yield



(d) Violent Crime Rate

Figure 5: Real Outcome Data (Counterfactual Control Correction)

Notes: This figure reports estimates of (1) for a series of outcome variables, alongside estimate of (3) that also include the i) expected counterfactual exposure to each temperature bin or ii) 5-year times county fixed effects as controls. Daily temperature realizations in each county used to construct each temperature bins are from the ERA-5 Land dataset. In Panel (a), the unit of observation is a county-decade and the outcome is log of population measured from the U.S. Census (1970-2010). In Panel (b), the unit of observation is a county-year (1970-2002) and the outcome variable is the mortality rate for individuals over the age of 65 (per 100,000 inhabitants), measured using the CDC mortality files. In Panel (c), the unit of observation is a county-year (1985-2006) and the outcome variable is log of corn yield, measured using NASS Quickstats. In Panel (d), the unit of observation is a county-month (1970-2009) and the outcome variable is violent crime rate (i.e., murder, rape, and aggravated assault) measured using data from the FBI. 95% confidence intervals are reported and standard errors are clustered by county.

Appendix

A Theoretical Framework

In this appendix, we provide further analysis of U-bias using a formal model. We develop conditions for when this bias will occur and analyze cases where U-shapes, inverted U-shapes, or other forms of non-linear bias will emerge.

To begin, consider a standard panel regression featuring temperature bins. The idea is that hot days, temperate days, and cold days may have quite different effects on outcome variables of interest. This possibility is often analyzed in models of the following form:

$$Y_{it} = \alpha_i + \lambda_t + \sum_k \beta_k \times T_{kit} + u_{it} \tag{6}$$

where Y_{it} is the outcome variable for place i at time t, α_i are place-specific fixed effects, λ_t are time fixed effects, and T_{kit} represent counts of days in temperature bin k in a given time and place. The bins consider different ranges of temperature, often 10°F or 5°C wide. The β_k are intended to provide the treatment effects of interest, estimating how more days in any given bin affect the outcome.

A.1 The Outcome Variable

First consider the data generating process for the outcome variable. We will consider the case where researchers run a panel model as in (6) but the true data generating process for the outcome variable is

$$Y_{it} = \alpha_i + \lambda_t + \sum_k \beta_k \times T_{kit} + \gamma_i t + \varepsilon_{it} \tag{7}$$

where the error is *iid* and $\mathbb{E}[\varepsilon | t, T_{kit}] = 0$. The key here is the introduction of the $\gamma_i t$, which bring in potential place-specific trends in the outcome variable (and are omitted in the regression specification (6)). These trends may covary with other place-specific characteristics and are independent of the error terms, $\mathbb{E}[\gamma_i \epsilon_{it}] = 0$.

A.2 The Temperature Bins

Now consider the data generation process for the temperature bins. For simplicity, consider $k = 3$ temperature bins, where C_{it} is a count of cold days, M_{it} is a count of moderate temperature days, and H_{it} is a count of hot days. Let periods be one year. The sum of the day counts thus add up to 365 days in any given place and year; given this collinearity,

we drop the count of moderate days, M_{it}, from the regression. Estimates of treatment effects of C_{it} and H_{it} will be interpreted as differences from the effect of moderate days.

Let the true data generating process for counts of cold days and hot days be

$$C_{it} = \mu_{C,i} + \delta_{C,i}t + v_{C,it} \tag{8}$$

$$H_{it} = \mu_{H,i} + \delta_{H,i}t + v_{H,it} \tag{9}$$

where the error terms are *iid* with $\mathbb{E}[v_{.,it}|t] = 0$, $\mathbb{E}[\delta_{.,i}v_{.,it}] = 0$, and variances σ_C^2 and σ_H^2 respectively.[23]

As with the dependent variable (see equation (7)), we have introduced the potential for place-specific differential trends into these temperature variables, represented here by $\delta_{C,i}t$ and $\delta_{H,i}t$. As described in the introduction, these differential trends are a natural consequence of global warming. In particular, if we think of warming as shifting the entire distribution of temperature in a given place to the right, we expect fewer cold days with time ($\delta_{C,i} < 0$). But since initially hot places experience few if any cold days to start with, this decline will inevitably be much less negative in initially hot places. Similarly, with warming, we expect more hot days with time ($\delta_{H,i} > 0$) but this increase will inevitably be much smaller in initially cold places, which may still see very few very hot days (see Figure 1a). Thus these day counts will tend to have differential trends depending on the baseline temperature of the place.

A.3 Forms of Bias

We can now consider biases that can result when econometricians implement a traditional panel model. Namely, consider the panel regression

$$Y_{it} = \alpha_i + \lambda_t + \beta_C C_{it} + \beta_H H_{it} + u_{it} \tag{10}$$

but where the true data generating processes are given by (7), (8), and (9).

Given the outcome variable's data generating process, the error term in (10) is

$$u_{it} = \gamma_i t + \epsilon_{it}. \tag{11}$$

Bias can then result when estimating (10) should the trending terms in either C_{it} or H_{it} have any covariance with these error trends. Intuitively, time fixed effects cannot solve this problem because, while they account for common dynamics across all places, they do not account for these differential trends by place. The resulting biases are derived as

[23] Note that, since total days in each year add up to 365, the bin-level error terms cannot be fully independent across all bins. Here we let the errors for cold and hot days be *iid*, with the moderate days being determined from these other independent processes. Extensions to the baseline model might consider non-independent errors.

follows.

Result 1. *For large T, the biases in the coefficient estimates are*

$$Bias(\hat{\beta}_C) = \frac{Var(\delta_{H,i})\,Cov(\delta_{C,i},\gamma_i) - Cov(\delta_{C,i},\delta_{H,i})\,Cov(\delta_{H,i},\gamma_i)}{Var(\delta_{C,i})\,Var(\delta_{H,i}) - \left(Cov(\delta_{C,i},\delta_{H,i})\right)^2} \tag{12}$$

$$Bias(\hat{\beta}_H) = \frac{Var(\delta_{C,i})\,Cov(\delta_{H,i},\gamma_i) - Cov(\delta_{C,i},\delta_{H,i})\,Cov(\delta_{C,i},\gamma_i)}{Var(\delta_{C,i})\,Var(\delta_{H,i}) - \left(Cov(\delta_{C,i},\delta_{H,i})\right)^2}. \tag{13}$$

These biases admit U-shape patterns (both biases are positive), inverted U-shape patterns (both biases are negative), or non-U-shape patterns, according to conditions summarized in Figure A.1.

Proof See Online Appendix.

Note that, from this result, the existence of place-specific trends in the outcome and temperature bin variables does not necessarily imply bias. In particular, there will be no bias if the place-specific trends in the outcome variable are not covariant with the place-specific trends in the temperature variables (i.e., if $\text{Cov}(\delta_{C,i},\gamma_i) = 0$ and $\text{Cov}(\delta_{H,i},\gamma_i) = 0$). However, a sufficient condition for bias is simply that at least one of the temperature bin variables shows differential trends across place that are correlated with differential outcome trends.

The bias formulas in equations (12) and (13), and the result in Figure A.1, allow for arbitrary correlations between trends in the temperature bins and trends in the outcomes. But as discussed above, we are motivated by a particular case where place-specific trends are correlated with initial temperature—i.e., where hot and cold day counts will have trends that are correlated by construction with a place's initial temperature (see discussion above). In this case, we can derive more specific results.

Specifically, suppose there is a relationship between trends in the temperature bins and initial temperature of a given place i, which we denote $T_{0,i}$. Write:

$$\delta_{C,i} = \psi_C + \omega_C T_{0,i} + e_{C,i} \tag{14}$$

$$\delta_{H,i} = \psi_H + \omega_H T_{0,i} + e_{H,i} \tag{15}$$

$$\gamma_i = \psi_Y + \omega_Y T_{0,i} + e_{Y,i} \tag{16}$$

where the error terms $e_{.,i}$ are all *iid* with variances $\sigma^2_{e_.}$. Let the variance of $T_{0,i}$ be $\sigma^2_{T_0}$. With warming, we expect $\psi_C < 0$, so that cold days are on average trending downwards, and $\omega_C > 0$, so that places with initially low temperature will have the most rapid decline in the number of cold days. Similarly, with warming, we expect that $\psi_H > 0$ so that hot days are on average trending upwards, and $\omega_H > 0$ so that places with initially low temperature will have smaller increases in the number of hot days. We then have the following result:

Result 2. *With the trends defined in (14), (15), and (16), the biases in the coefficient estimates are*

$$Bias(\hat{\beta}_C) = \frac{\omega_C \omega_Y}{\omega_C^2 + \frac{A_C}{A_H}\omega_H^2 + A_H} \tag{17}$$

$$Bias(\hat{\beta}_H) = \frac{\omega_H \omega_Y}{\omega_H^2 + \frac{A_H}{A_C}\omega_H^2 + A_C} \tag{18}$$

where $A_C = \frac{\sigma_{e_C}^2}{\sigma_{T_0}^2} + \frac{T-1}{S_t}\frac{\sigma_C^2}{\sigma_{T_0}^2}$, $A_H = \frac{\sigma_{e_H}^2}{\sigma_{T_0}^2} + \frac{T-1}{S_t}\frac{\sigma_H^2}{\sigma_{T_0}^2}$ *and* $S_t = \sum_{t=1}^{T}(t-\bar{t})^2$. *The bias is U-shaped if* $\omega_Y > 0$, *the bias is inverted U-shaped if* $\omega_Y < 0$, *and there is no bias if* $\omega_Y = 0$.

Proof See Online Appendix.

Result 2 elucidates several features of binning bias that we will explore in the next sections. As a general intuition, the bias problem will naturally be worse when differential trends on both sides of the regression are more tightly related. These trend relationships depend on the trends' correlations with baseline temperature (the ω parameters) and the degree of noise (the σ^2 parameters).

Consider four key features, focusing on the case where $\omega_Y > 0$ (generating a U-shape). First, bias increases in ω_Y. Intuitively, the larger the correlation between place-specific outcome trends and initial temperature, the greater the bias will be. Second, the bias in $\hat{\beta}_C$ declines as $\sigma_{e_C}^2$ rises. Intuitively, as a temperature bin's trends become driven more by random noise than a systematic relationship to initial temperature, the bias for the bin is attenuated. Third, and perhaps counter-intuitively, longer panels tend to make binning bias *worse*. The reason is that idiosyncratic inter-annual variation (via σ_C^2 and σ_H^2), which provides the shocks to each bin's day counts and which are intended to identify temperature effects, becomes dominated by any trend the longer those trends proceed.[24] Thus longer panels tend to amplify the binning bias problem.

Finally, the effect of ω_C on the bias in $\hat{\beta}_C$ is more subtle. On the one hand, a smaller slope with initial temperature (smaller ω_C) causes the signal to noise ratio to go down, which works to reduce the bias (the same way that increasing the noise, $\sigma_{e_C}^2$, reduces bias). On the other hand, a smaller slope with initial temperature (smaller ω_C) means that smaller variations in day counts can now predict the given variation in the outcome, resulting in larger bias. The latter force will actually dominate when noise is relatively low (e.g., in longer panels or with small $\sigma_{e_C}^2$). Overall, small but precise relationships between temperature bin trends and initial temperature are especially pernicious, driving large bias. We will see these features in practice in the empirical analyses to come.

[24] In Result 2, as T increases, the A_C and A_H terms decline, since T/S_T is decreasing in T.

B Additional Simulation Details

B.1 Uniform Warming with Simulated Temperatures

To model the effects of U bias under uniform warming, we consider synthetic temperature data and synthetic outcome data in order to home in on the mechanism of interest. We generate temperature data with uniform warming across U.S. counties, in which the average temperature in each county increases by 1.8°F (i.e., 1°C) from its baseline in 1970.[25] For each county and year, we draw 365 daily temperatures from a normal distribution centered at the average, with a common fixed variance. We then count the number of days in each county-year in each temperature bin, defining the bins in 10°F intervals. We generate the outcome using equation (2).

B.1.1 U Bias with Uniform Warming

With these data, we estimate a version of equation (1) in which i indexes counties and t indexes years for 1,000 draws of the simulated outcome and temperature realizations. Estimates of β_k are presented in Figure A.2a, where each dot represents the median coefficient estimate and the bars contain 95% of the coefficient estimates across all draws. Strikingly, the coefficient estimates in Figure A.2a imply large positive effects of exposure to both hot and cold days despite there being *no* relationship between contemporaneous temperature extremes and the outcome.

Thus, a rightward shift in the temperature distribution—combined with place-specific trends in the outcome that are correlated with baseline temperature—generates a spurious U-shaped relationship even with completely uniform, trends in temperature. As in our main results, when the trend in the outcome is *negative* (i.e., t is replaced with $-t$ in equation (2)) the panel model generates a spurious inverted-U-shaped relationship between temperature exposure and the outcome (Figure A.2b), and when there are no differential trends in the outcome the spurious effects are absent (Figure A.2c).

B.1.2 Validating the Non-Random Temperature Exposure Controls

To validate the approach in Section 4.2.1, we exploit the fact that in this uniform-warming simulation we know the exact data generating process (DGP) for temperature. Since the independent variables are constructed from simulated data we can compute the expected number of days that fall into each temperature bin. We then include the expected number of days in each bin as controls in the panel model, as in equation (3). We find indeed that controlling for the expected number of days in each bin removes the bias (Figure A.9). The method is successful both when the trend in the outcome is a linear function of baseline temperature (Figure A.9a) or when we use more complicated, non-linear functional forms

[25] This 1.8°F increase is roughly the change in average U.S. temperature over the sample period.

(Figure A.9b). Moreover, when we add in *actual* time-specific effects (e.g., where extra days in the hottest and coldest bin do affect outcomes), we recover those effects accurately (Figure A.9c). Thus, this approach correctly purges bias from the estimates.

B.2 Calibrated Outcome Trends

To calibrate differential outcome trends in our simulation analysis to those of real-world outcomes, we first estimate:

$$Y_{ct} = \delta_c + \eta_t + \alpha_c \cdot t + \epsilon_{ct} \tag{19}$$

where Y_{ct} is actual outcome data (e.g., the mortality rate in a county c and year t) and $\hat{\alpha}_c$ captures any county-specific time trend for that outcome.

We then run the following model:

$$\hat{\alpha}_c = \omega + \zeta \cdot Temp_c^{Base} + \epsilon_c \tag{20}$$

where $\hat{\zeta}$ captures any correlation between baseline temperatures and county-specific time trends.

Finally, we simulate Y_{ct} so that its trend in baseline temperature matches that of each real-world outcome of interest:

$$\widetilde{Y}_{ct} = \hat{\delta}_c + [\hat{\omega} + \hat{\zeta} \cdot Temp_c^{Base}] \cdot t + \xi_{ct} \tag{21}$$

where $\hat{\delta}_c$ is estimated from equation (19) and $\hat{\omega}$ and $\hat{\zeta}$ are estimated from equation (20). Finally, using these calibrated outcomes and actual temperature data, we again estimate equation (1)). Again, by construction, there should be no actual relationship between temperature realizations and outcomes.

Online Appendix

Proof of Result 1

We first remove place and time fixed effects by applying the two-way fixed effects transformation. We can then proceed to the estimated treatment coefficients from the residualized regression specification, following the Frisch–Waugh–Lovell theorem. For the true model, the adjusted (demeaned) outcome equation is:

$$\tilde{Y}_{it} = \beta_C \tilde{C}_{it} + \beta_H \tilde{H}_{it} + \tilde{\gamma}_i \tilde{t} + \tilde{\varepsilon}_{it} \tag{22}$$

Residual cold days are

$$\tilde{C}_{it} = \tilde{\delta}_{C,i} \tilde{t} + \tilde{v}_{C,it} \tag{23}$$

where $\tilde{\delta}_{C,i} = \delta_{C,i} - \delta_{C,\text{avg}}$ (and $\delta_{C,\text{avg}} = \frac{1}{N} \sum_{i=1}^{N} \delta_{C,i}$) and where $\tilde{t} = t - \bar{t}$ (and $\bar{t} = \frac{1}{T} \sum_{s=1}^{T} s$).

Residual hot days are

$$\tilde{H}_{it} = \tilde{\delta}_{H,i} \tilde{t} + \tilde{v}_{H,it} \tag{24}$$

where $\tilde{\delta}_{H,i} = \delta_{H,i} - \delta_{H,\text{avg}}$ (and $\delta_{H,\text{avg}} = \frac{1}{N} \sum_{i=1}^{N} \delta_{H,i}$).

Residual place-specific trends in the outcome variable are

$$\tilde{\gamma}_i = \gamma_i - \gamma_{\text{avg}} \tag{25}$$

where γ_{avg} is the average place-specific trend across all places, $\gamma_{\text{avg}} = \frac{1}{N} \sum_{i=1}^{N} \gamma_i$.

Thus, the two-way fixed effects eliminate unit intercepts and the common time pattern but leave unit-specific deviations.

Estimating the Model without Individual-Specific Trends

When we estimate the model without accounting for the residual individual-specific trend $\tilde{\gamma}_i \tilde{t}$, we have:

$$\tilde{Y}_{it} = \beta_C \tilde{C}_{it} + \beta_H \tilde{H}_{it} + \tilde{u}_{it} \tag{26}$$

where the error term is:

$$\tilde{u}_{it} = \tilde{\gamma}_i \tilde{t} + \tilde{\varepsilon}_{it} \tag{27}$$

raising potential for bias if the residual trending terms in either $\tilde{C}_{it}$ or $\tilde{H}_{it}$ have any covariance with the residual trend in the error term $\tilde{u}_{it}$.

OLS Estimators

The OLS estimators for β_C and β_H are:

$$\begin{pmatrix} \hat{\beta}_C \\ \hat{\beta}_H \end{pmatrix} = \left(\tilde{\mathbf{D}}'\tilde{\mathbf{D}}\right)^{-1} \tilde{\mathbf{D}}'\tilde{Y} \tag{28}$$

where $\tilde{\mathbf{D}}$ is the matrix of adjusted treatment variables:

$$\tilde{\mathbf{D}} = \begin{pmatrix} \tilde{C}_{11} & \tilde{H}_{11} \\ \tilde{C}_{12} & \tilde{H}_{12} \\ \vdots & \vdots \\ \tilde{C}_{NT} & \tilde{H}_{NT} \end{pmatrix} \tag{29}$$

Bias in the Estimators

The bias in the estimated coefficients arises due to the correlation between $\tilde{u}_{it}$ and the treatment variables $\tilde{C}_{it}$ and $\tilde{H}_{it}$.

$$\begin{pmatrix} \text{Bias}(\hat{\beta}_C) \\ \text{Bias}(\hat{\beta}_H) \end{pmatrix} = \left(\tilde{\mathbf{D}}'\tilde{\mathbf{D}}\right)^{-1} \tilde{\mathbf{D}}'\tilde{u} \tag{30}$$

Define $\mathbf{V} = \tilde{\mathbf{D}}'\tilde{\mathbf{D}}$ and $\mathbf{Z} = \tilde{\mathbf{D}}'\tilde{u}$. In calculating the bias, we will focus on results for large samples, where we deploy the weak law of large numbers and the continuous mapping theorem.

Computing V

Define the elements of the matrix $\mathbf{V}$ as

$$\mathbf{V} = \begin{pmatrix} V_{CC} & V_{CH} \\ V_{CH} & V_{HH} \end{pmatrix} \tag{31}$$

The term V_{CC} relates to the variance of $\tilde{C}_{it}$.

$$\begin{aligned} V_{CC} &= \sum_{i=1}^{N}\sum_{t=1}^{T} \tilde{C}_{it}^2 \\ &= \sum_{i=1}^{N}\sum_{t=1}^{T} (\tilde{\delta}_{C,i}\tilde{t} + \tilde{v}_{C,it})^2 \\ &= \sum_{i=1}^{N}\sum_{t=1}^{T} \tilde{\delta}_{C,i}^2\tilde{t}^2 + \sum_{i=1}^{N}\sum_{t=1}^{T} 2\tilde{\delta}_{C,i}\tilde{t}\tilde{v}_{C,it} + \sum_{i=1}^{N}\sum_{t=1}^{T} \tilde{v}_{C,it}^2 \end{aligned}$$

Take the probability limits for large N, so that each summation over i converges to its expectation. Define $S_t = \sum_{t=1}^{T} \tilde{t}^2$ where time is deterministic. Note that $\mathbb{E}[\tilde{\delta}_{C,i}^2] = \text{Var}(\delta_{C,i})$, and $\mathbb{E}\left[\tilde{v}_{C,it}^2\right] = (1 - 1/T)\sigma_C^2$.[26] The (true) error terms, $v_{C,it}$ and $v_{H,it}$, are *iid*, so that $\mathbb{E}\left[\tilde{\delta}_{C,i}\tilde{v}_{C,it}\right] = 0$. Then it follows that

$$V_{CC} \to NS_t\text{Var}(\delta_{C,i}) + N(T-1)\sigma_C^2 \tag{32}$$

The term V_{HH} relates to the variance of $\tilde{H}_{it}$. Following similar steps we have

$$V_{HH} \to NS_t\text{Var}(\delta_{H,i}) + N(T-1)\sigma_H^2 \tag{33}$$

The term V_{CH} relates to the covariance between $\tilde{C}_{it}$ and $\tilde{H}_{it}$, which is computed as:

$$V_{CH} = \sum_{i=1}^{N}\sum_{t=1}^{T} \tilde{C}_{it}\tilde{H}_{it} \tag{34}$$

$$= \sum_{i=1}^{N}\sum_{t=1}^{T} (\tilde{\delta}_{C,i}\tilde{t} + \tilde{v}_{C,it})(\tilde{\delta}_{H,i}\tilde{t} + \tilde{v}_{H,it}) \tag{35}$$

$$= \sum_{i=1}^{N}\sum_{t=1}^{T} \tilde{\delta}_{C,i}\tilde{\delta}_{H,i}\tilde{t}^2 + \sum_{i=1}^{N}\sum_{t=1}^{T} \tilde{t}\tilde{\delta}_{C,i}\tilde{v}_{H,it} + \sum_{i=1}^{N}\sum_{t=1}^{T} \tilde{t}\tilde{\delta}_{H,i}\tilde{v}_{C,it} + \sum_{i=1}^{N}\sum_{t=1}^{T} \tilde{v}_{C,it}\tilde{v}_{H,it} \tag{36}$$

Again, taking probability limits, and noting that the true errors (v) are *iid* and mean zero conditional on regressors, it follows that all summations converge to zero except the first, which is the covariance of $\tilde{\delta}_{C,i}$ and $\tilde{\delta}_{H,i}$. Thus

$$V_{CH} \to NS_t\text{Cov}(\delta_{C,i}, \delta_{H,i}) \tag{37}$$

Inverting the Matrix V

The inverse of $\mathbf{V}$ is:

$$\mathbf{V}^{-1} = \frac{1}{\det(\mathbf{V})}\begin{pmatrix} V_{HH} & -V_{CH} \\ -V_{CH} & V_{CC} \end{pmatrix} \tag{38}$$

where the determinant of $\mathbf{V}$ is:

$$\det(\mathbf{V}) = V_{CC}V_{HH} - V_{CH}^2 \tag{39}$$

[26] We are not (yet) assuming large T, so the empirical variance of the idiosyncratic disturbances that depend on both i and t are not identical to the true variance in short T panels.

Taking probability limits and using the above results, we have

$$\det(\mathbf{V}) \to \left(NS_t\text{Var}(\delta_{C,i}) + N(T-1)\sigma_C^2\right)\left(NS_t\text{Var}(\delta_{H,i}) + N(T-1)\sigma_H^2\right) - \left(NS_t\text{Cov}(\delta_{C,i},\delta_{H,i})\right)^2 \tag{40}$$

$$= N^2S_t^2\left(\left(\text{Var}(\delta_{C,i}) + \frac{(T-1)\sigma_C^2}{S_t}\right)\left(\text{Var}(\delta_{H,i}) + \frac{(T-1)\sigma_H^2}{S_t}\right) - \left(\text{Cov}(\delta_{C,i},\delta_{H,i})\right)^2\right) \tag{41}$$

Computing Z

The covariance vector $Z = \tilde{\mathbf{D}}'\tilde{u}$ has elements

$$\mathbf{Z} = \begin{pmatrix} Z_C \\ Z_H \end{pmatrix} \tag{42}$$

which we compute as follows.

$$Z_C = \sum_{i=1}^{N}\sum_{t=1}^{T}\tilde{C}_{it}\tilde{u}_{it} \tag{43}$$

$$= \sum_{i=1}^{N}\sum_{t=1}^{T}(\tilde{\delta}_{C,i}\tilde{t} + \tilde{v}_{C,it})(\tilde{\gamma}_i\tilde{t} + \tilde{\varepsilon}_{it}) \tag{44}$$

$$= \sum_{i=1}^{N}\sum_{t=1}^{T}\tilde{t}^2\tilde{\delta}_{C,i}\tilde{\gamma}_i + \sum_{i=1}^{N}\sum_{t=1}^{T}\tilde{\delta}_{C,i}\tilde{\varepsilon}_{it}\tilde{t} + \sum_{i=1}^{N}\sum_{t=1}^{T}\tilde{\gamma}_i\tilde{v}_{C,it}\tilde{t} + \sum_{i=1}^{N}\sum_{t=1}^{T}\tilde{v}_{C,it}\tilde{\varepsilon}_{it} \tag{45}$$

Taking probability limits, and noting that the true errors (v and ϵ) are *iid* and mean zero conditional on regressors, it follows that all summations converge to zero except the first, which incorporates the covariance of $\tilde{\delta}_{C,i}$ and $\tilde{\gamma}_i$. Thus

$$Z_C \to NS_t\text{Cov}(\delta_{C,i},\gamma_i) \tag{46}$$

Similarly for Z_H.

$$Z_H \to NS_t\text{Cov}(\delta_{H,i},\gamma_i) \tag{47}$$

Computing Bias

Recall that the bias is:

$$\begin{pmatrix} \text{Bias}(\hat{\beta}_C) \\ \text{Bias}(\hat{\beta}_H) \end{pmatrix} = \mathbf{V}^{-1}\mathbf{Z} \tag{48}$$

Using the matrix definitions, we have

$$\text{Bias}(\hat{\beta}_C) = \frac{1}{\det(\mathbf{V})}\left(V_{HH}Z_C - V_{CH}Z_H\right) \tag{49}$$

Using the probability limits calculated for each term above, and assuming a non-zero determinant (full rank), this bias converges to

$$\text{Bias}(\hat{\beta}_C) \to \frac{\left(NS_t\text{Var}(\delta_{H,i}) + N(T-1)\sigma_H^2\right)\left(NS_t\text{Cov}(\delta_{C,i},\gamma_i)\right) - \left(NS_t\text{Cov}(\delta_{C,i},\delta_{H,i})\right)\left(NS_t\text{Cov}(\delta_{H,i},\gamma_i)\right)}{N^2S_t^2\left(\left(\text{Var}(\delta_{C,i}) + \frac{(T-1)\sigma_C^2}{S_t}\right)\left(\text{Var}(\delta_{H,i}) + \frac{(T-1)\sigma_H^2}{S_t}\right) - \left(\text{Cov}(\delta_{C,i},\delta_{H,i})\right)^2\right)}$$

Or, simplifying, we can write this in terms of variances and covariances as

$$\text{Bias}(\hat{\beta}_C) \to \frac{\left(\text{Var}(\delta_{H,i}) + \frac{T-1}{S_t}\sigma_H^2\right)\left(\text{Cov}(\delta_{C,i},\gamma_i)\right) - \left(\text{Cov}(\delta_{C,i},\delta_{H,i})\right)\left(\text{Cov}(\delta_{H,i},\gamma_i)\right)}{\left(\text{Var}(\tilde{\delta}_{C,i}) + \frac{T-1}{S_t}\sigma_C^2\right)\left(\text{Var}(\delta_{H,i}) + \frac{T-1}{S_t}\sigma_H^2\right) - \left(\text{Cov}(\delta_{C,i},\delta_{H,i})\right)^2} \tag{50}$$

Similarly, we calculate $\text{Bias}(\hat{\beta}_H)$ as

$$\text{Bias}(\hat{\beta}_H) = \frac{1}{\det(\mathbf{V})}\left(V_{CC}Z_H - V_{CH}Z_C\right) \tag{51}$$

which converges for large N as

$$\text{Bias}(\hat{\beta}_H) \to \frac{\left(\text{Var}(\delta_{C,i}) + \frac{T-1}{S_t}\sigma_C^2\right)\left(\text{Cov}(\delta_{H,i},\gamma_i)\right) - \left(\text{Cov}(\delta_{C,i},\delta_{H,i})\right)\left(\text{Cov}(\delta_{C,i},\gamma_i)\right)}{\left(\text{Var}(\delta_{C,i}) + \frac{T-1}{S_t}\sigma_C^2\right)\left(\text{Var}(\delta_{H,i}) + \frac{T-1}{S_t}\sigma_H^2\right) - \left(\text{Cov}(\delta_{C,i},\delta_{H,i})\right)^2} \tag{52}$$

Bias with Large T

For large T, S_t grows on the order of T^3. Therefore, the terms involving T/S_t become negligible compared to other terms. Simplifying on these grounds, we then have

$$\text{Bias}(\hat{\beta}_C) = \frac{\text{Var}(\delta_{H,i})\text{Cov}(\delta_{C,i},\gamma_i) - \text{Cov}(\delta_{C,i},\delta_{H,i})\text{Cov}(\delta_{H,i},\gamma_i)}{\text{Var}(\delta_{C,i})\text{Var}(\delta_{H,i}) - \left(\text{Cov}(\delta_{C,i},\delta_{H,i})\right)^2} \tag{12}$$

$$\text{Bias}(\hat{\beta}_H) = \frac{\text{Var}(\delta_{C,i})\text{Cov}(\delta_{H,i}, \gamma_i) - \text{Cov}(\delta_{C,i}, \delta_{H,i})\text{Cov}(\delta_{C,i}, \gamma_i)}{\text{Var}(\delta_{C,i})\text{Var}(\delta_{H,i}) - (\text{Cov}(\delta_{C,i}, \delta_{H,i}))^2} \tag{13}$$

as was to be shown.

Bias Map

With these results, we can consider the direction of the bias. Rearranging the numerators of the bias equations (12) and (13), define the ratios

$$R_A = \frac{\text{Var}(\delta_{H,i})\text{Cov}(\delta_{C,i}, \gamma_i)}{\text{Cov}(\delta_{C,i}, \delta_{H,i})\text{Cov}(\delta_{H,i}, \gamma_i)} \tag{53}$$

$$R_B = \frac{\text{Var}(\delta_{C,i})\text{Cov}(\delta_{H,i}, \gamma_i)}{\text{Cov}(\delta_{C,i}, \delta_{H,i})\text{Cov}(\delta_{C,i}, \gamma_i)} \tag{54}$$

Further note that that a positive determinant of the matrix $\boldsymbol{V}$ implies that

$$R_A R_B > 1 \tag{55}$$

The potential bias cases are summarized in Figure A.1. ■

Proof of Result 2

Consider the trend processes given by (14), (15), and (16), as repeated here:

$$\begin{aligned}
\delta_{C,i} &= \psi_C + \omega_C T_{0,i} + e_{C,i} \\
\delta_{H,i} &= \psi_H + \omega_H T_{0,i} + e_{H,i} \\
\gamma_i &= \psi_Y + \omega_Y T_{0,i} + e_{Y,i}
\end{aligned}$$

It then follows that the covariances are

$$\begin{aligned}
\text{Cov}(\delta_{C,i}, \delta_{H,i}) &= \omega_C \omega_H \sigma^2_{T_0} \\
\text{Cov}(\delta_{C,i}, \gamma_i) &= \omega_C \omega_Y \sigma^2_{T_0} \\
\text{Cov}(\delta_{H,i}, \gamma_i) &= \omega_H \omega_Y \sigma^2_{T_0}
\end{aligned}$$

and the variances are

$$\text{Var}(\delta_{C,i}) = \omega_C^2 \sigma_{T_0}^2 + \sigma_{e_C}^2$$
$$\text{Var}(\delta_{H,i}) = \omega_H^2 \sigma_{T_0}^2 + \sigma_{e_H}^2$$

Plug these variance and covariance results in to the finite T results above (equations (50) and (52)). Collecting terms and simplifying leads to the forms:

$$\text{Bias}(\hat{\beta}_C) = \frac{\omega_C \omega_Y}{\omega_C^2 + \frac{A_C}{A_H}\omega_H^2 + A_H} \tag{56}$$

$$\text{Bias}(\hat{\beta}_H) = \frac{\omega_H \omega_Y}{\omega_H^2 + \frac{A_H}{A_C}\omega_H^2 + A_C} \tag{57}$$

where we define $A_C = \frac{\sigma_{e_C}^2}{\sigma_{T_0}^2} + \frac{T-1}{S_t}\frac{\sigma_C^2}{\sigma_{T_0}^2}$ and $A_H = \frac{\sigma_{e_H}^2}{\sigma_{T_0}^2} + \frac{T-1}{S_t}\frac{\sigma_H^2}{\sigma_{T_0}^2}$.

Since $\omega_C > 0$ and $\omega_H > 0$ and all terms in the denominator are positive, it follows by inspection that the biases are both positive (a U-shape) if $\omega_Y > 0$ and the biases are both negative (an inverted U-shape) if $\omega_Y < 0$. Further, by inspection, there is no bias if $\omega_Y = 0$. ■

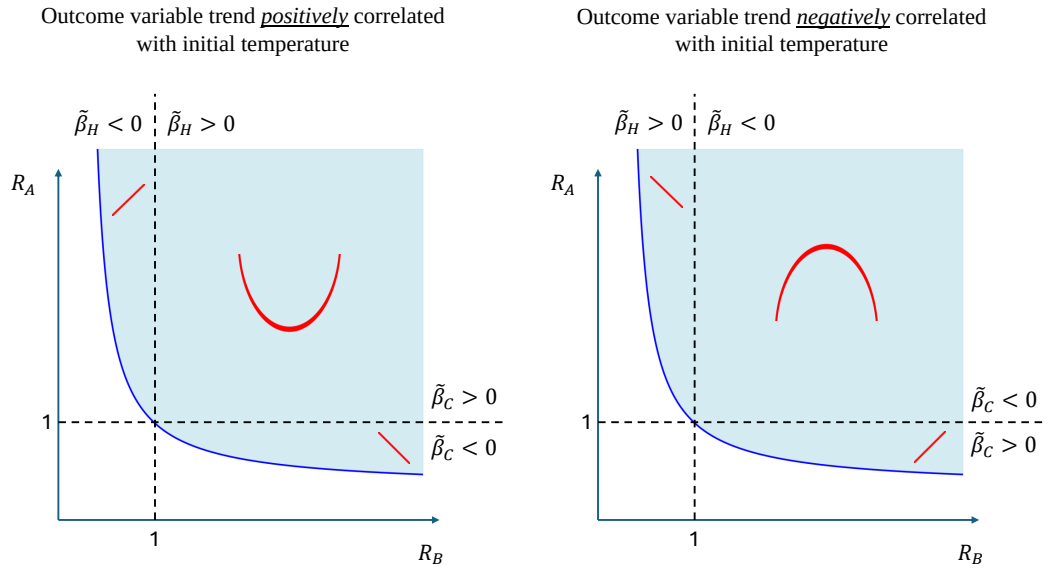
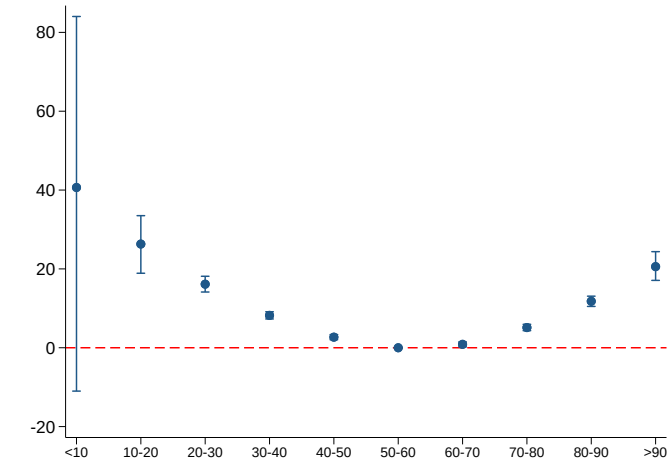
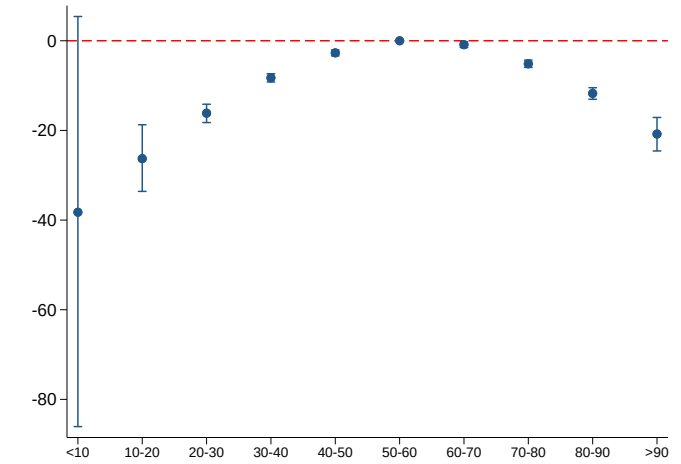


Figure A.1: The Direction of Bias

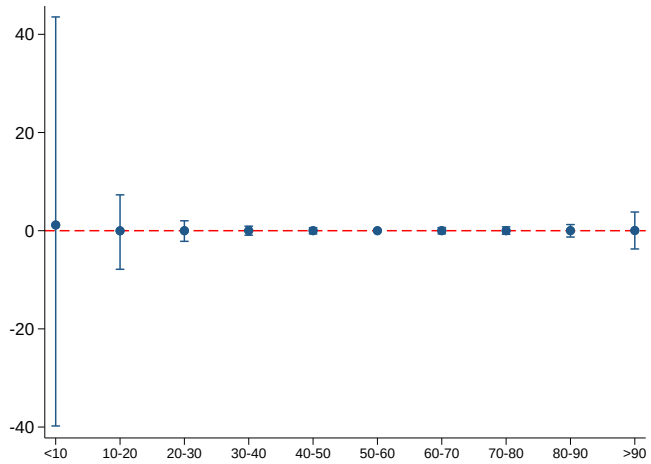
Notes: This figure considers the sign of the biases on the coefficients $\hat{\beta}_C$ and $\hat{\beta}_H$ using Result 1. U-shapes and inverted U-shapes are common, but upward and downward sloping biases are also possible, with the bias directions indicated in red. The vertical and horizontal axes are ratios formed from the numerators of the bias equations. The ratio $R_A = \text{Var}(\delta_{H,i})\text{Cov}(\delta_{C,i}, \gamma_i)/[\text{Cov}(\delta_{C,i}, \delta_{H,i})\text{Cov}(\delta_{H,i}, \gamma_i)]$ is drawn from the bias equation (12). The ratio $R_B = \text{Var}(\delta_{C,i})\text{Cov}(\delta_{H,i}, \gamma_i)/[\text{Cov}(\delta_{C,i}, \delta_{H,i})\text{Cov}(\delta_{C,i}, \gamma_i)]$ is drawn from the bias equation (13). The condition $R_A R_B > 1$ is necessary and guarantees a non-singular regression matrix. In the left panel, which features U shapes, we assume that the outcome trends are positively covariant with the trends in both hot and cold days. In the right panel, which features inverted U shapes, we assume that the outcome trends are negatively covariant with the trends in both hot and cold days.



(a) Positively Correlated Outcome Trend



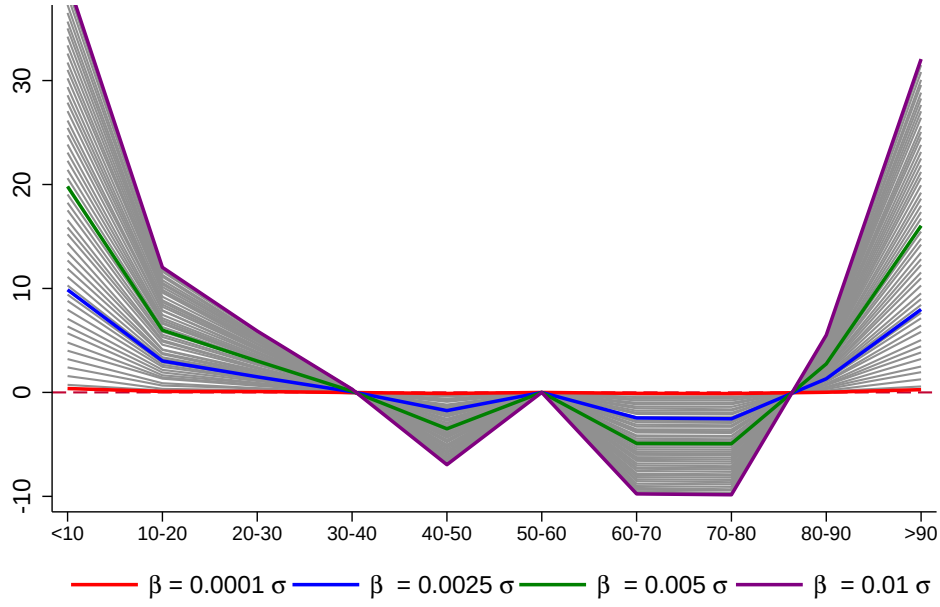
(b) Negatively Correlated Outcome Trend



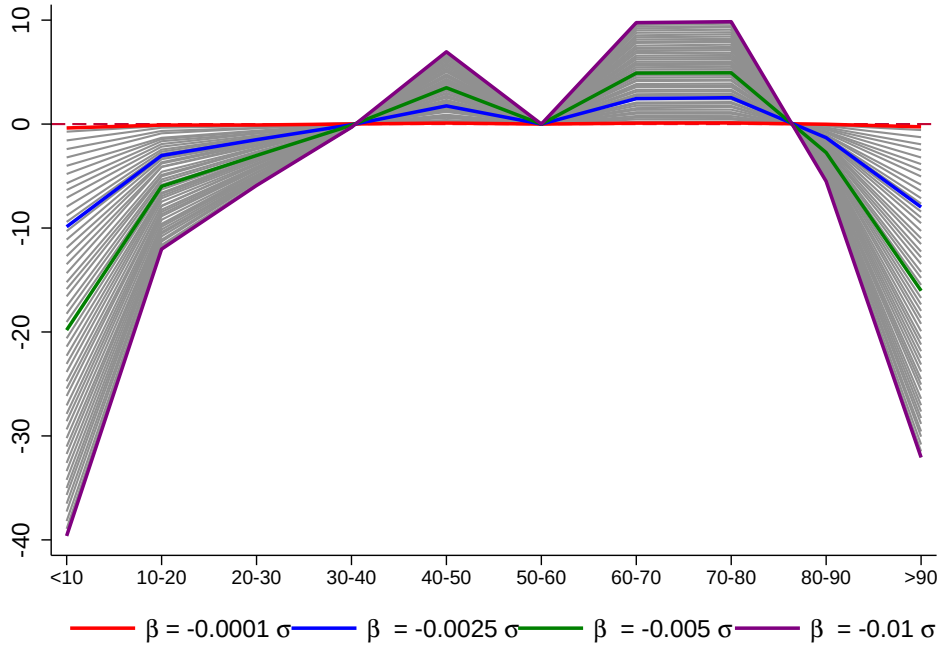
(c) Uncorrelated Outcome Trend

Figure A.2: Binned Regression Estimates with Simulated Outcome and Uniform Temperature Change

Notes: Each sub-figure reports estimates from the panel model (1), in which the outcome is simulated using the following equation: $Y_{ct} = aTemp_c^{base} + \epsilon_{ct}$. The error term is drawn from a normal distribution centered at zero with variance equal to the variance of $aTemp_c^{base}$ over the full sample. The temperature data are drawn so that all counties have the same warming trend (1.8°F over the sample period), starting with a normal distribution centered at the county-level baseline temperature. In each sub-figure, we report the mean and 95% confidence intervals of each coefficient from (1) taking 1,000 separate draws of the outcome and temperature variables. In Panel (a), the outcome is simulated to have a positive differential trend in baseline temperature (i.e., $a = t$); in Panel (b), the outcome is simulated to have a negative differential trend in baseline temperature (i.e., $a = -t$); and in Panel (c) the outcome has no relationship with baseline temperature (i.e., $a = 0$).



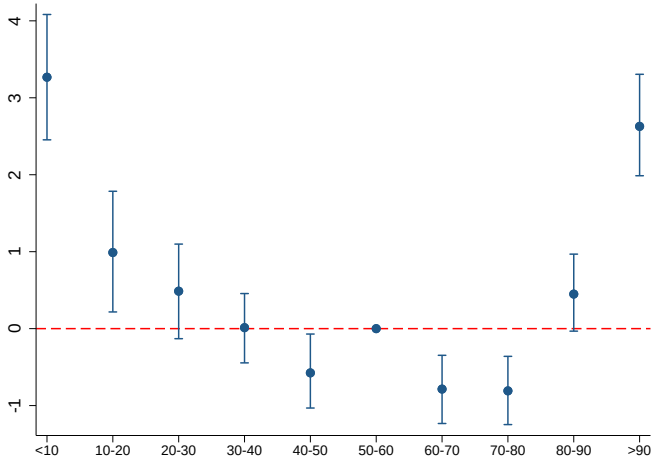
(a) Spurious Effects by Magnitude of Positive Outcome Trend



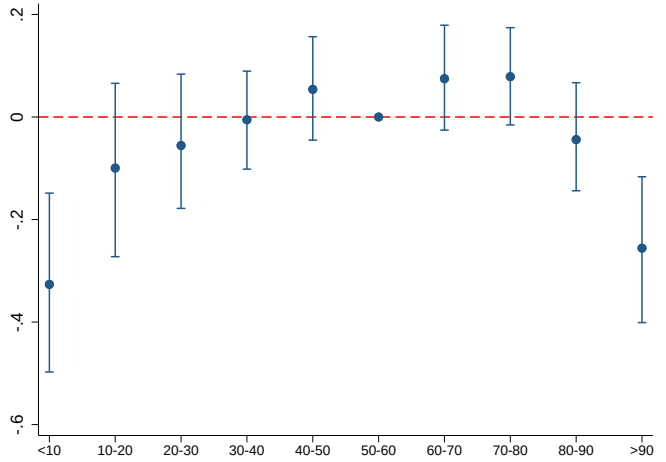
(b) Spurious Effects by Magnitude of Negative Outcome Trend

Figure A.3: Magnitudes of Outcome Trend and Spurious Effects

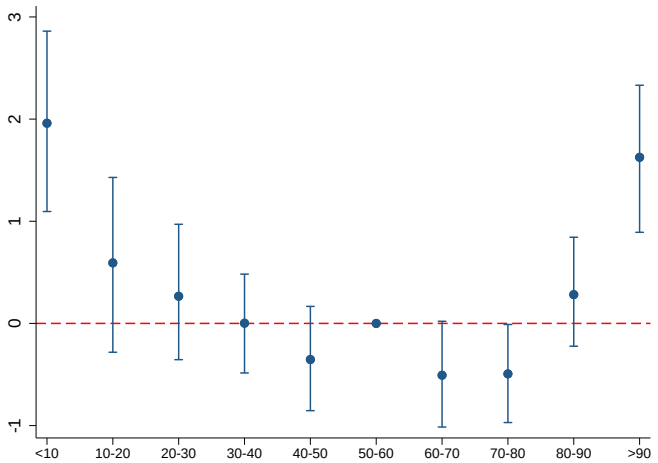
Notes: Each sub-figure reports estimates from equation (1), in which the outcome is simulated using the following equation: $Y_{ct} = \beta \cdot t \cdot Temp_c^{base} + \epsilon_{ct}$, $\epsilon_{ct} \sim \mathcal{N}(0, \sigma)$, and σ equals the standard deviation of $tTemp_c^{base}$ over the full sample. The temperature bins are constructed from historical, daily temperature realizations using the ERA-5 Land dataset. In Panel (a), the outcome is simulated to have a positive differential trend in baseline temperature (i.e., $\beta \in [0.0001\sigma, 0.01\sigma]$); in Panel (b), the outcome is simulated to have a negative differential trend in baseline temperature (i.e., $\beta \in [-0.01\sigma, -0.0001\sigma]$). In each sub-figure, we report the mean of each coefficient from (1) taking 1,000 separate draws of the outcome variable for each value of β and highlight four examples in red ($\beta = \pm 0.0001\sigma$), blue ($\beta = \pm 0.0025\sigma$), green ($\beta = \pm 0.005\sigma$) and purple ($\beta = \pm 0.01\sigma$).



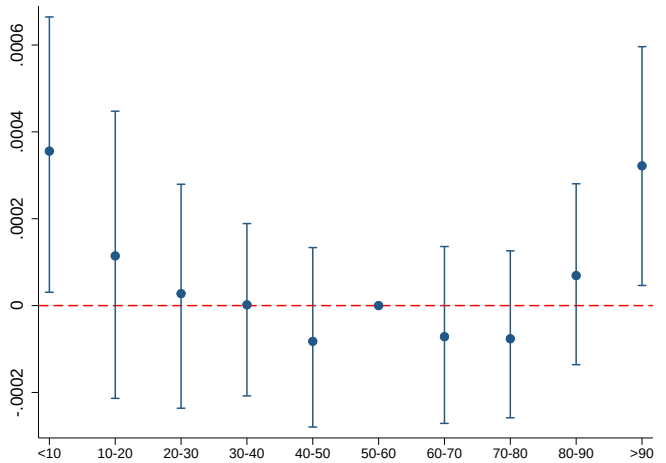
(a) Outcome Calibrated to Mortality for Ages 65+



(b) Outcome Calibrated to Violent Crime Rates



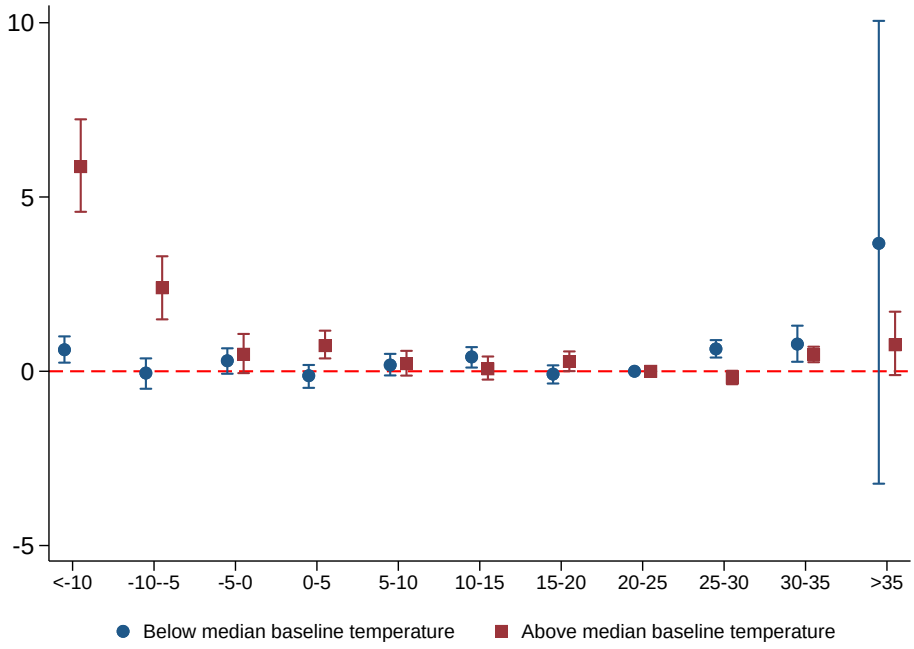
(c) Outcome Calibrated to Nonviolent Crime Rates



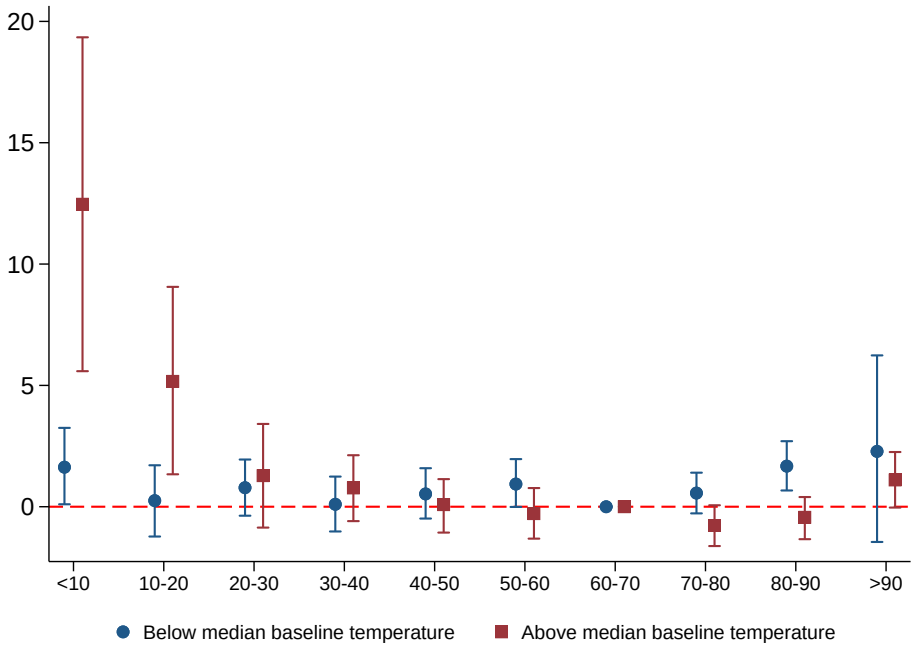
(d) Outcome Calibrated to Corn Yields

Figure A.4: Simulation with Calibrated Outcomes and Real Temperature

Notes: Each sub-figure reports estimates from equation (1), in which the outcome is calibrated to real data using equation (21). The temperature bins are constructed from historical, daily temperature realizations using the ERA-5 Land dataset. In each sub-figure, we report the mean and 95% confidence intervals of each coefficient from (1) taking 1,000 separate draws of the outcome variable. In Panel (a), the outcome is calibrated to the mortality rate for people age 65 and over using CDC mortality files. In Panels (b) and (c), the outcome is calibrated to violent/non violent crime rate respectively (i.e., murder, rape, and aggravated assault for Panel (b) and larceny and robbery for Panel (c)) measured using data from the FBI. In Panel (d), the outcome variable is calibrated to log of corn yield, measured using NASS Quickstats.



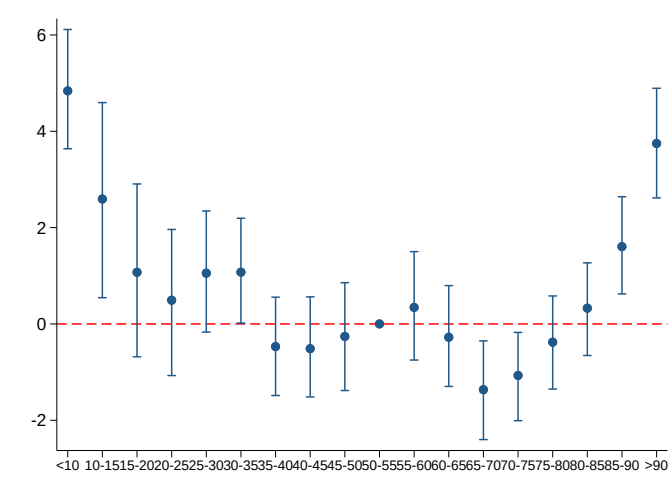
(a) Binned Regression Estimates: Celsius



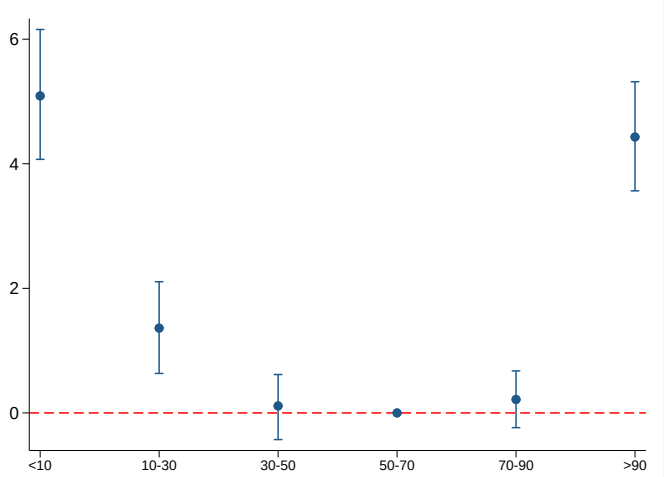
(b) Binned Regression Estimates: Omitting the 60-70°F Bin

Figure A.5: Binned Regression Estimates and "Adaptation": Alternative Binning

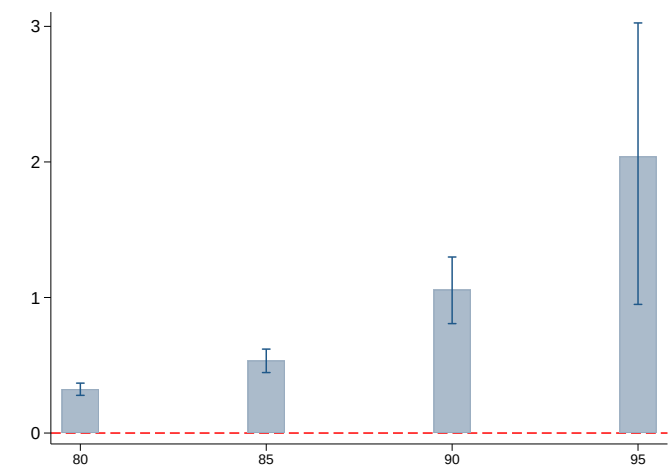
Notes: We report estimates of equation (1), separately for counties with above versus below median baseline temperature. The outcome is simulated to have a positive differential trend in baseline temperature ($Y_{ct} = tTemp_c^{base} + \epsilon_{ct}$) and daily temperature realizations are from the ERA-5 Land dataset. In Panel (a), the dataset is constructed in Celsius, and the 20-25°C (68-77°F) bin is omitted for normalization. In Panel (b), the dataset is in Fahrenheit but we omit the 60-70°F bin instead of the 50-60°F bin as in Figure 3c. We report the mean and 95% confidence intervals for each coefficient estimate, taking 1,000 draws of ϵ_{ct} from a normal distribution centered at zero with variance equivalent to the in-sample variance of $tTemp_c^{base}$.



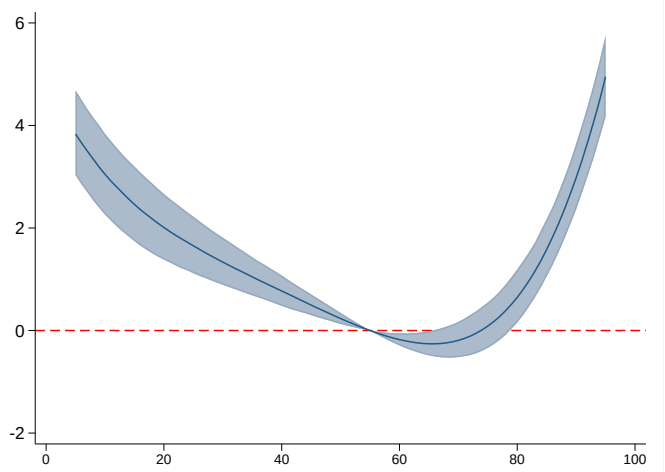
(a) Narrower Bins (5°F)



(b) Wider Bins (20°F)



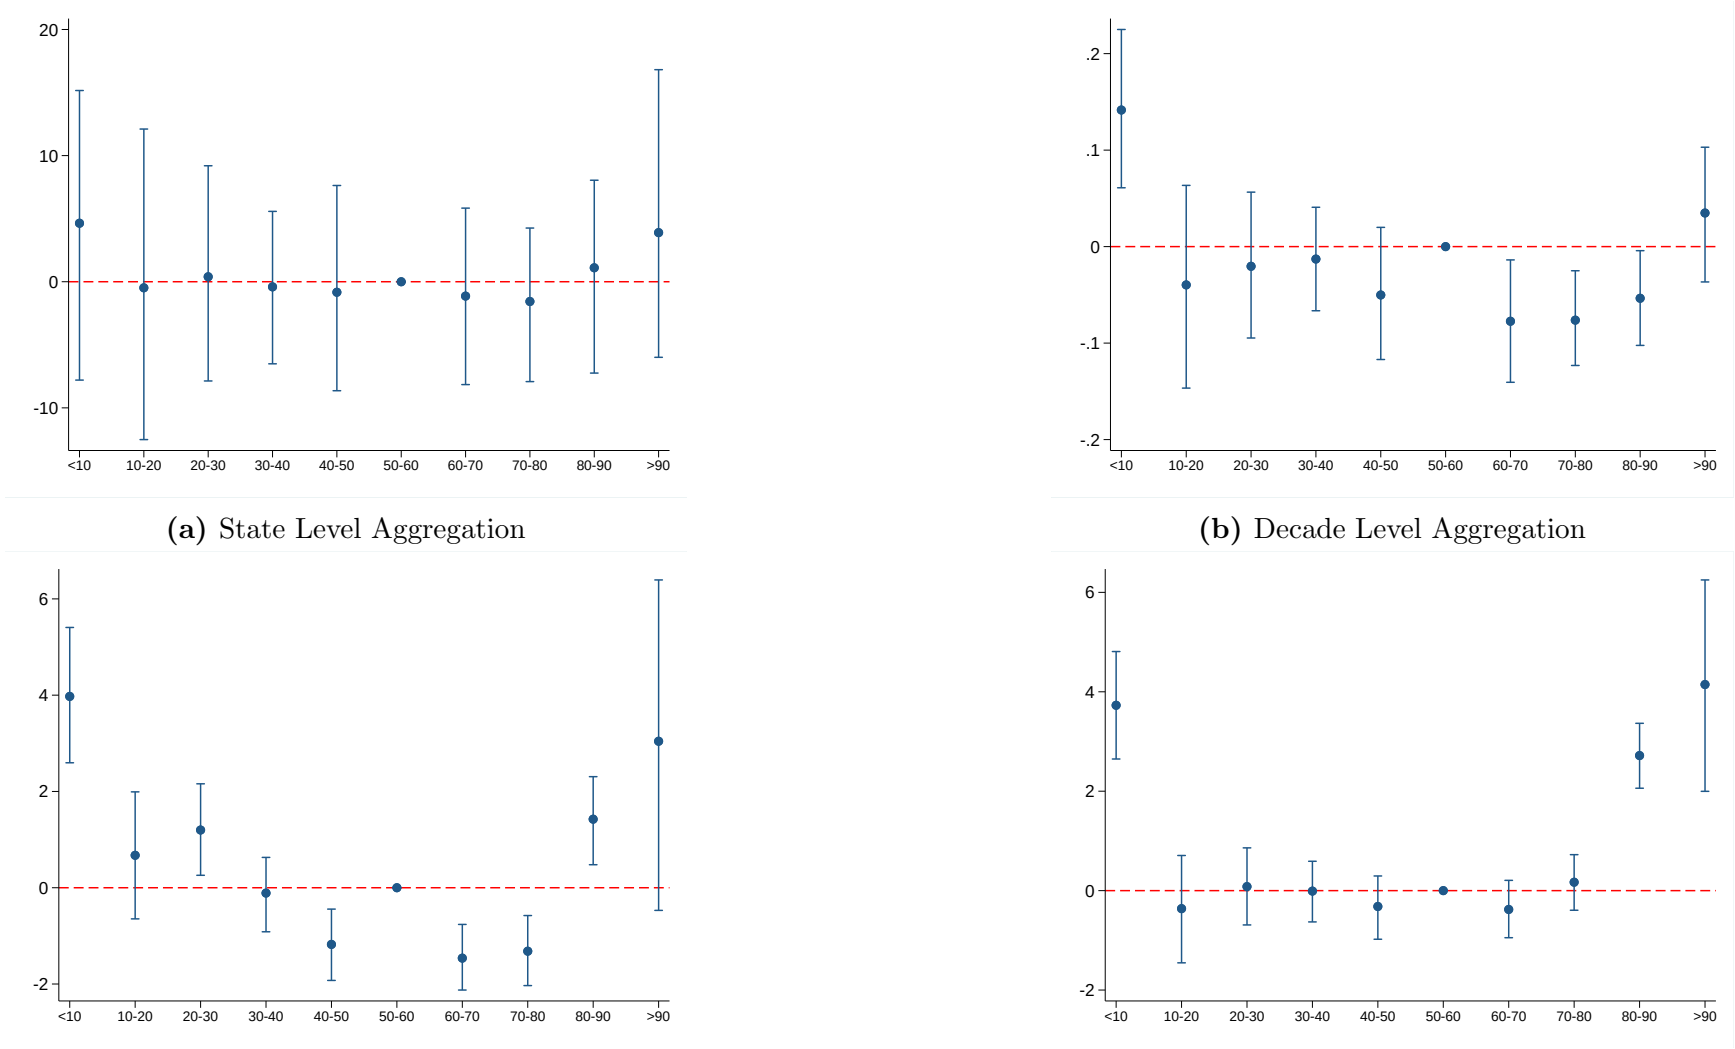
(c) Killing Degree Days



(d) Polynomials in Daily Temperature

Figure A.6: Binning Bias: Alternative Nonlinear Specifications

Notes: We conduct the same exercise as Figure 2, which uses simulated outcome data and real temperature data, under a variety of alternative non-linear parameterizations of temperature shocks. Panels (a) and (b) present results when daily temperatures are binned at 5°F and 20°F intervals, respectively. Panel (c) implements a regression of the simulated outcome on killing degree days (KDDs). For a given threshold T^*, KDD is defined as $\sum_{t=T^*}\sum_{d=1}\mathbb{1}(t < Temp_d \leq t+1)](t+0.5-T^*)$ (i.e., the number of degree-days in a given year above a threshold T^*). We report the mean and 95% confidence intervals for the coefficient estimates on KDD from setting 80, 85, 90, and 95°F as the threshold. Panel (d) present results from the polynomial specification used in Carleton et al. (2022). This specification regresses the outcome variable on a fourth-order polynomial of daily temperatures with county fixed effects and year fixed effects. Each polynomial term is constructed by raising the daily temperatures to that power and summing them within each county-year. This figure plots the predicted values of the fourth-order polynomials at each °F, where we normalize by subtracting the predicted value of 55°F. In all specifications, the outcome is simulated to have a positive differential trend in baseline temperature ($Y_{ct} = tTemp_c^{base} + \epsilon_{ct}$) and daily temperature realizations are from the ERA-5 Land dataset. We report the mean and 95% confidence intervals for each estimate, taking 1,000 draws of $\epsilon_{c,t}$ from a normal distribution centered at zero with variance equivalent to the in-sample variance of $tTemp_c^{base}$.



(a) State Level Aggregation

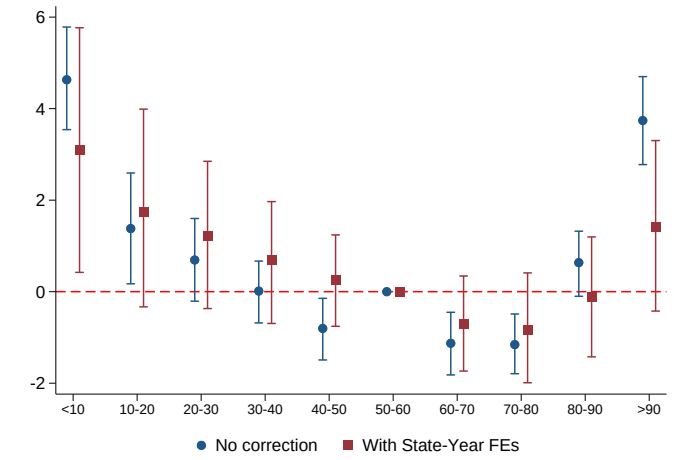
(b) Decade Level Aggregation

(c) PRISM Temperature Data

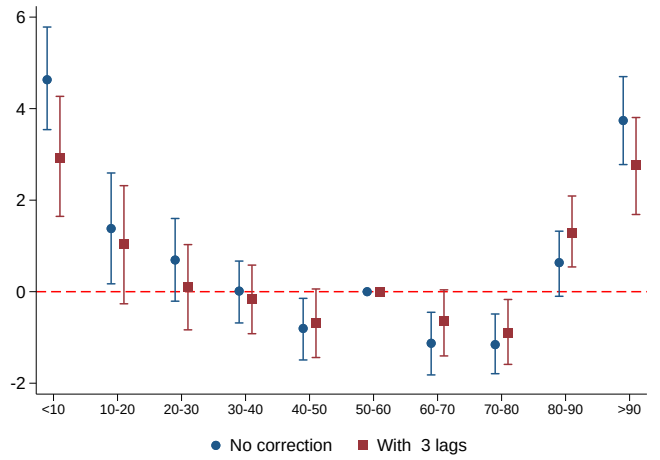
(d) GHCN Temperature Data

Figure A.7: Binning Bias: Sensitivity and Extensions to Alternative Levels of Aggregation and Temperature Datasets

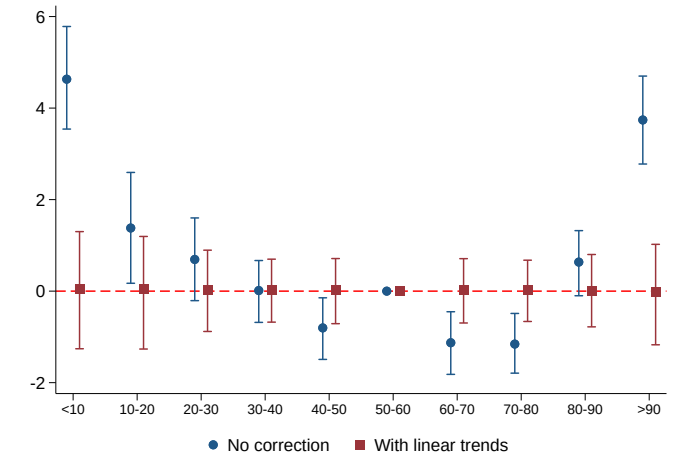
Notes: We conduct the same exercise as Figure 2, which uses simulated outcome data and real temperature data, under a variety of aggregation levels and temperature datasets. Each sub-figure reports estimates of equation (1). The outcome is simulated to have a positive differential trend in baseline temperature ($Y_{ct} = tTemp_c^{base} + \epsilon_{ct}$) and daily temperature realizations are from the ERA-5 Land dataset for Panels (a) and (b). We report the mean and 95% confidence intervals for each coefficient estimate, taking 1,000 draws of $\epsilon_{c,t}$ from a normal distribution centered at zero with variance equivalent to the in-sample variance of $tTemp_c^{base}$. Panel (a) presents results when the outcome and temperature variables are aggregated at the state level; Panel (b) presents results when the outcome and temperature variables are aggregated at the decade level. Panels (c) and (d) present results when daily temperature realizations are from PRISM and GHCN datasets, respectively.



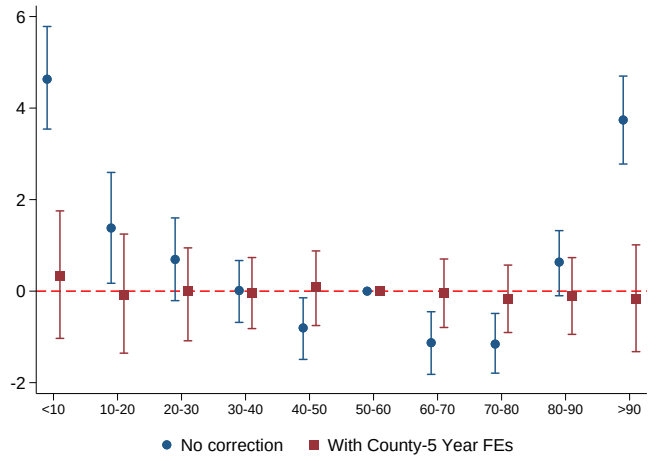
(a) State-Year Fixed Effects



(b) Lags of Dependent/Independent Variables



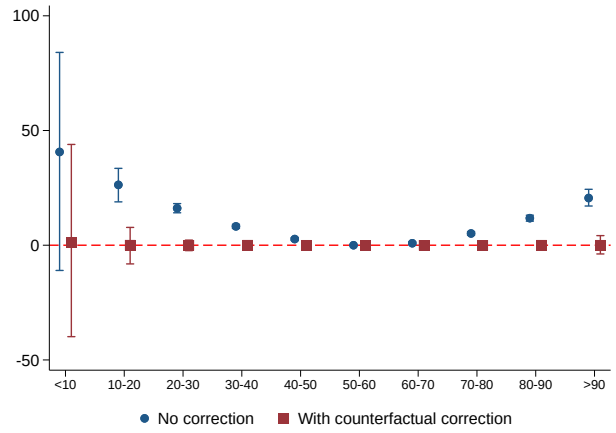
(c) County-Specific Linear Trends



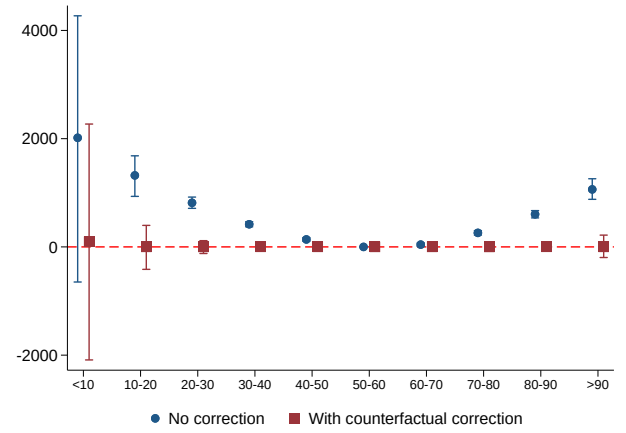
(d) Five-Year Indicators × County Fixed Effects

Figure A.8: Correcting for Bias with Reduced-Form Approaches

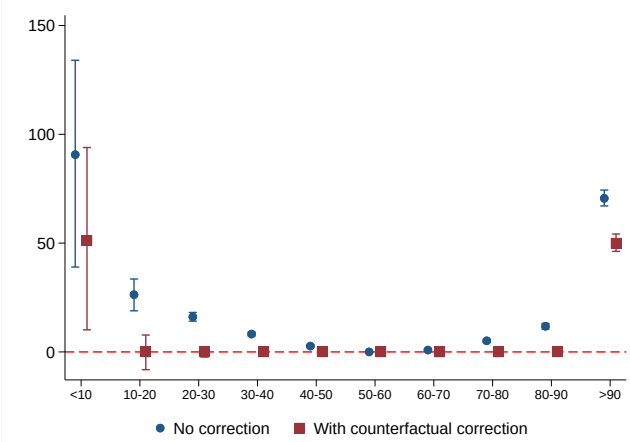
Notes: Simulations are the same as in Figure 4, but here we report estimates of equation (1) alongside estimates that employ alternative, reduced-form corrective strategies as indicated in the title for each panel. Daily temperature realizations used to construct each temperature bin are from the ERA-5 Land dataset. The outcome is simulated to have a positive differential trend in baseline temperature ($Y_{ct} = tTemp_c^{base} + \epsilon_{ct}$). For each regression specification, we report the mean and 95% confidence intervals for each coefficient estimate, taking 1,000 draws of $\epsilon_{c,t}$ from a normal distribution centered at zero with variance equivalent to the in-sample variance of $tTemp_c^{base}$.



(a) Counterfactual Correction: Linear Trend



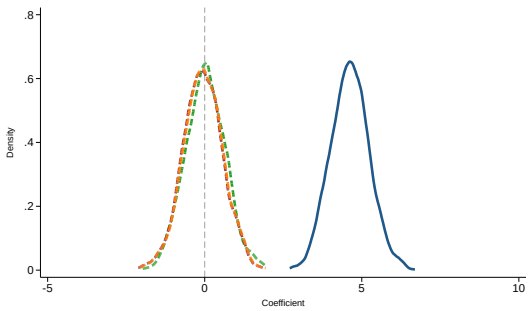
(b) Counterfactual Correction: Quadratic Trend



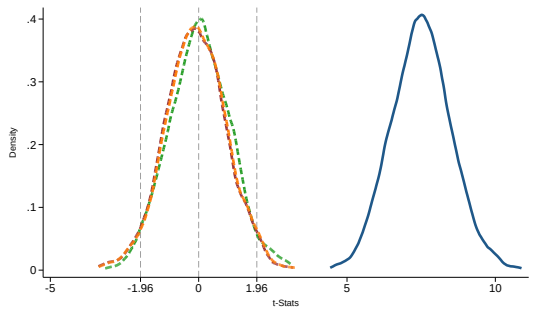
(c) Counterfactual Correction: True Effects

Figure A.9: Counterfactual Temperatures as a Solution to Binning Bias

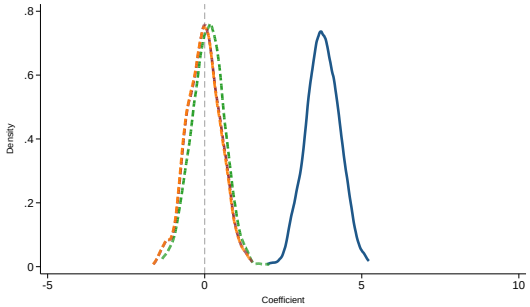
Notes: Each sub-figure reports estimates of equation (1), alongside of estimates of the same equation that include the expected counterfactual exposure to each temperature bin as controls. In Panels (a) and (c), the outcome is simulated to have a differential linear trend in baseline temperature, $Y_{ct} = t \cdot Temp_c^{base} + \epsilon_{ct}$. The error term is drawn from a normal distribution centered at zero with variance equal to the variance of $tTemp_c^{base}$ over the full sample. In Panel (b), the outcome is simulated to have a differential quadratic trend in baseline temperature $Y_{ct} = t^2 \cdot T_c^{Base} + \epsilon_{ct}$. The error term is drawn from a normal distribution centered at zero with variance equal to the variance of $t^2 Temp_c^{base}$ over the full sample. The temperature data are drawn so that all counties have the same warming trend (1°C over the sample period), starting with a normal distribution centered at the county-level baseline temperature. In Panels (a) and (c) the 1°C increase follows a linear trend over the period, while in Panel (b) the 1°C increase follows a quadratic trend. In Panel (c), the outcome has the same differential trend in the outcome variable as Panel (a) but there is also a positive effect of 50 of exposure to the top and bottom temperature bins. We report the mean and 95% confidence intervals for each coefficient estimate from (1), taking 1,000 draws of the outcome variable.



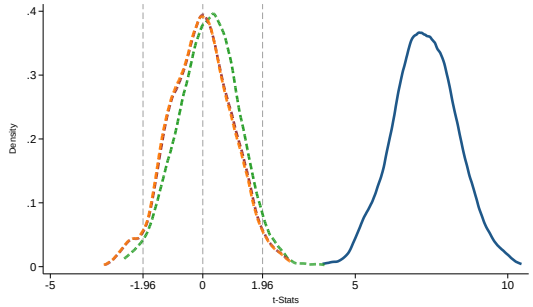
(a) Under 10 °F Bin: Coefficient



(b) Under 10 °F Bin: t-Stat



(c) Over 90 °F Bin: Coefficient



(d) Over 90 °F Bin: t-Stat

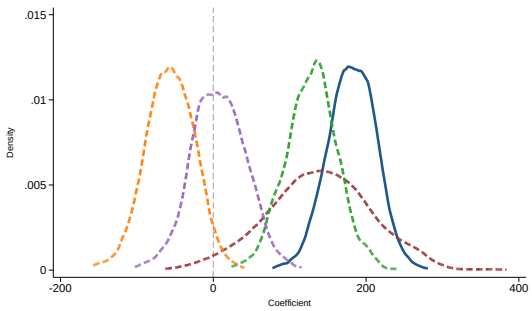
Figure A.10: Comparison of Counterfactual Temperature Solutions with Simulated Outcome Variable (Linear Trends)

—— No correction (Under 10 Bin: 100.0%, Over 90 Bin: 100.0%)

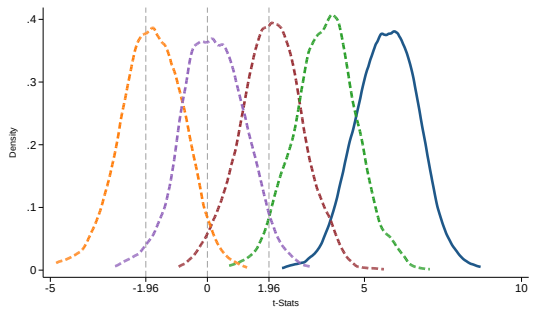
- - - Counterfactual Temperature Control (5.4%, 5.5%)

- - - Counterfactual Temperature Control without Bayes (8.8%, 8.2%) - - - Chebyshev (8.4%, 4.9%)

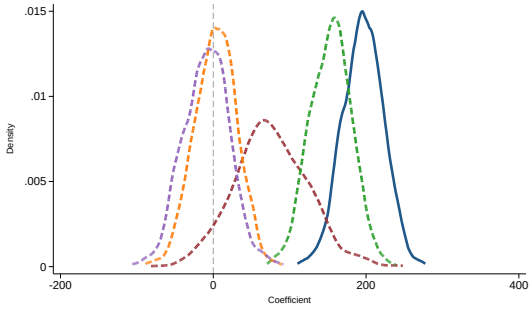
Notes: We conduct the same exercise as Figure 2, which uses simulated outcome data and real temperature data. Here, we plot the density of coefficients and t-statistics from estimating equation (10) 1,000 times, which only uses the two extreme temperature bins, alongside densities from specifications that also include the expected counterfactual exposure to each temperature bin as controls. Each color and line plots densities that use different methods to construct the counterfactuals, as indicated in the legend. The outcome is simulated to have a positive differential trend in baseline temperature ($Y_{ct} = tTemp_c^{base} + \epsilon_{ct}$) and daily temperature realizations are from the ERA-5 Land dataset. Each time, we take a draw of $\epsilon_{c,t}$ from a normal distribution centered at zero with variance equivalent to two-times the in-sample standard deviation of $tTemp_c^{base}$. Numbers inside the parentheses in the legend denote the number of runs per 1,000 simulations that the coefficient is significant at the 5% level, for the under 10°F bin and the over 90°F bin respectively.



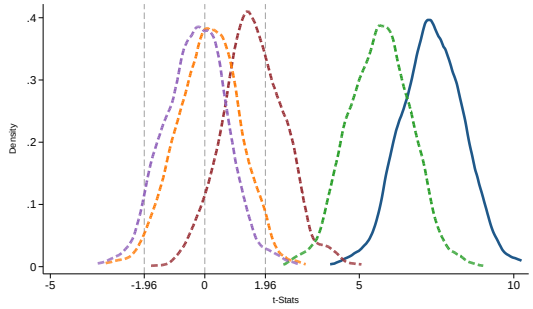
(a) Under 10 °F Bin: Coefficient



(b) Under 10 °F Bin: t-Stat



(c) Over 90 °F Bin: Coefficient



(d) Over 90 °F Bin: t-Stat

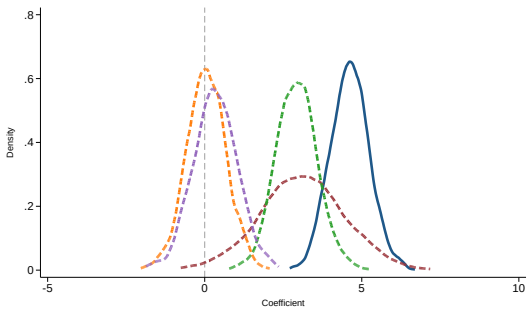
Figure A.11: Comparison of Reduced Form Solutions with Simulated Outcome Variable (Quadratic Trends)

—— No correction (Under 10 Bin: 100.0%, Over 90 Bin: 100.0%) - - - State-Year Fixed Effects (31.2%, 71.5%)

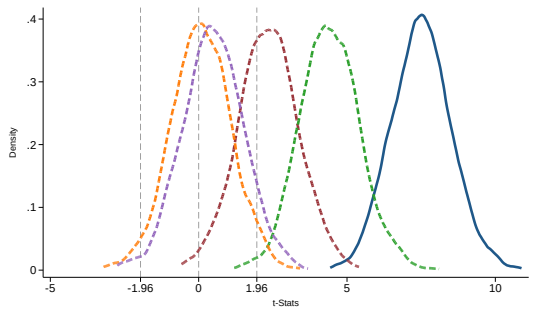
- - - With 3 Lags (100.0%, 100.0%) - - - County-Specific Linear Trends (40.5%, 4.5%)

- - - County-5 Year Fixed Effects (14.3%, 5.6%)

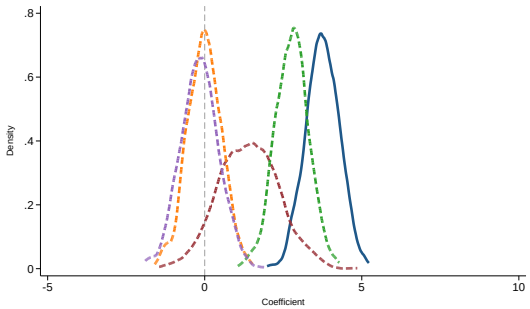
Notes: We conduct the same exercise as Figure 2, which uses simulated outcome data and real temperature data. Here, we plot the density of coefficients and t-statistics from estimating equation (10) 1,000 times, which only uses the two extreme temperature bins, alongside densities from specifications that use alternative reduced-form corrective strategies as indicated in the legend. The outcome is simulated to have a positive quadratic differential trend in baseline temperature ($Y_{ct} = t^2 Temp_c^{base} + \epsilon_{ct}$) and daily temperature realizations are from the ERA-5 Land dataset. Each time, we take a draw of $\epsilon_{c,t}$ from a normal distribution centered at zero with variance equivalent to two times the in-sample standard deviation of $t^2 Temp_c^{base}$. Numbers inside the parentheses in the legend denote the number of runs per 1,000 simulations that the coefficient is significant at the 5% level, for the under 10°F bin and the over 90°F bin respectively.



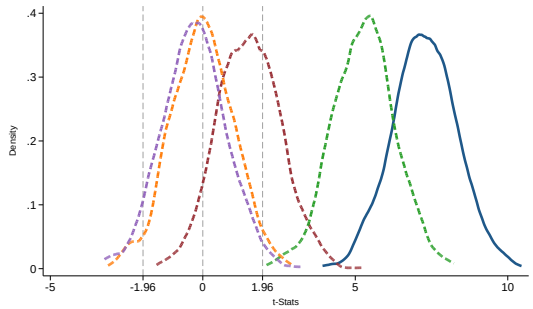
(a) Under 10 °F Bin: Coefficient



(b) Under 10 °F Bin: t-Stat



(c) Over 90 °F Bin: Coefficient



(d) Over 90 °F Bin: t-Stat

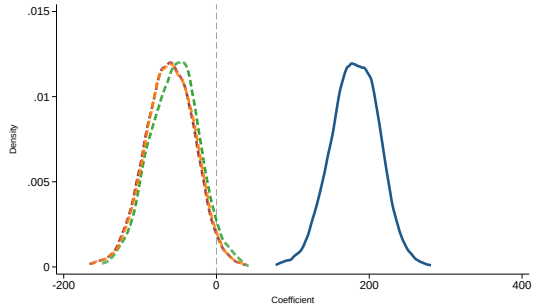
Figure A.12: Comparison of Reduced Form Solutions with Simulated Outcome Variable (Linear Trends)

—— No correction (Under 10 Bin: 100.0%, Over 90 Bin: 100.0%) - - - State-Year Fixed Effects (41.4%, 72.8%)

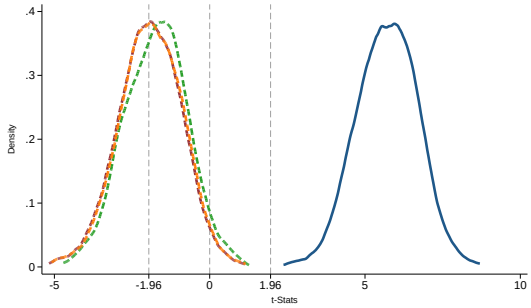
- - - With 3 Lags (100.0%, 100.0%) - - - County-Specific Linear Trends (5.4%, 4.5%)

- - - County-5 Year Fixed Effects (10.8%, 5.2%)

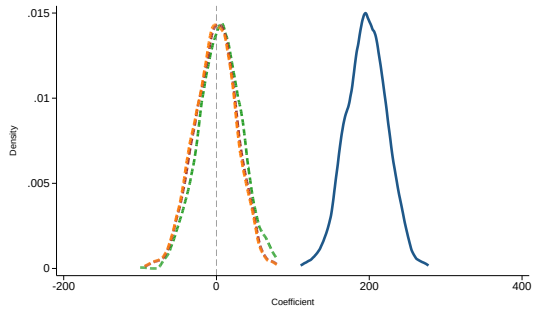
Notes: We conduct the same exercise as Figure 2, which uses simulated outcome data and real temperature data. Here, we plot the density of coefficients and t-statistics from estimating equation (10) 1,000 times, which only uses the two extreme temperature bins, alongside densities from specifications that use alternative reduced-form corrective strategies as indicated in the legend. The outcome is simulated to have a positive differential trend in baseline temperature ($Y_{ct} = tTemp_c^{base} + \epsilon_{ct}$) and daily temperature realizations are from the ERA-5 Land dataset. Each time, we take a draw of $\epsilon_{c,t}$ from a normal distribution centered at zero with variance equivalent to two times the in-sample standard deviation of $tTemp_c^{base}$. Numbers inside the parentheses in the legend denote the number of runs per 1,000 simulations that the coefficient is significant at the 5% level, for the under 10°F bin and the over 90°F bin respectively.



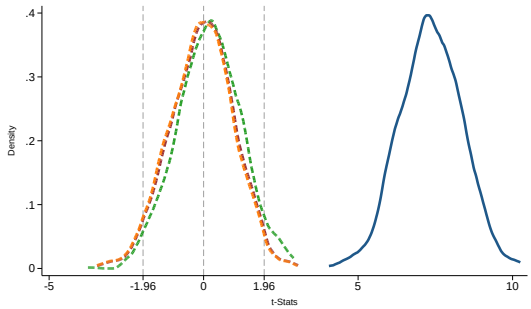
(a) Under 10 Bin: Coefficient



(b) Under 10 Bin: t-Stat



(c) Over 90 Bin: Coefficient



(d) Over 90 Bin: t-Stat

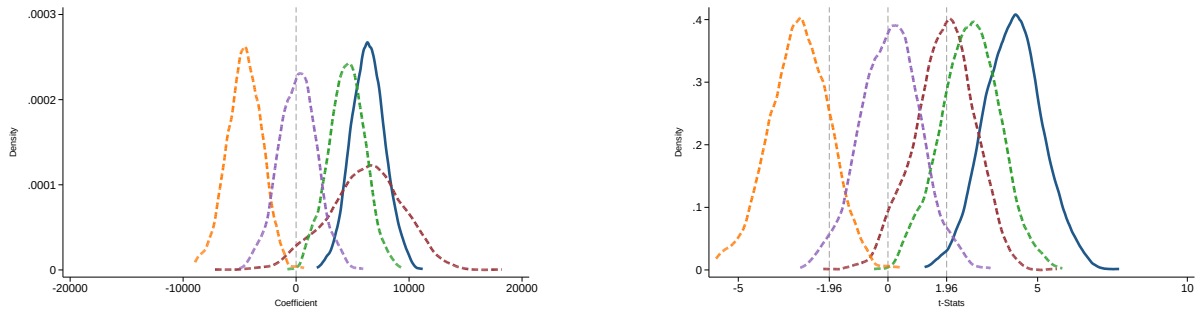
Figure A.13: Comparison of `cftemp` Solutions with Simulated Outcome Variable (Quadratic Trends)

—— No correction (Under 10 Bin: 100.0%, Over 90 Bin: 100.0%)

- - - Counterfactual Temperature Control (33.6%, 5.2%)

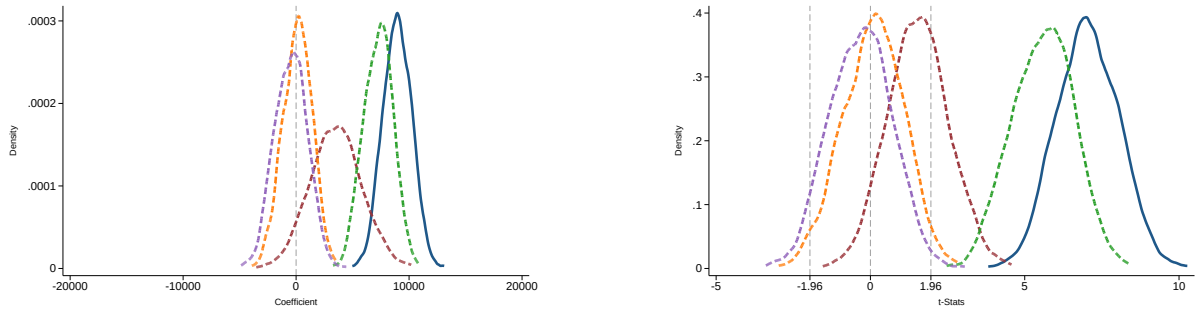
- - - Counterfactual Temperature Control without Bayes (19.3%, 5.9%) - - - Chebyshev (22.0%, 5.2%)

Notes: We conduct the same exercise as Figure 2, which uses simulated outcome data and real temperature data. Here, we plot the density of coefficients and t-statistics from estimating equation (10) 1,000 times, which only uses the two extreme temperature bins, alongside densities from specifications that also include the expected counterfactual exposure to each temperature bin as controls. Each color and line plots densities that use different methods to construct the counterfactuals, as indicated in the legend. The outcome is simulated to have a positive quadratic differential trend in baseline temperature ($Y_{ct} = t^2 Temp_c^{base} + \epsilon_{ct}$) and daily temperature realizations are from the ERA-5 Land dataset. Each time, we take a draw of $\epsilon_{c,t}$ from a normal distribution centered at zero with variance equivalent to two times the in-sample standard deviation of $t^2 Temp_c^{base}$. Numbers inside the parentheses in the legend denote the number of runs per 1,000 simulations that the coefficient is significant at the 5% level, for the under 10°F bin and the over 90°F bin respectively.



(a) Under 10 °F Bin: Coefficient

(b) Under 10 °F Bin: t-Stat



(c) Over 90 °F Bin: Coefficient

(d) Over 90 °F Bin: t-Stat

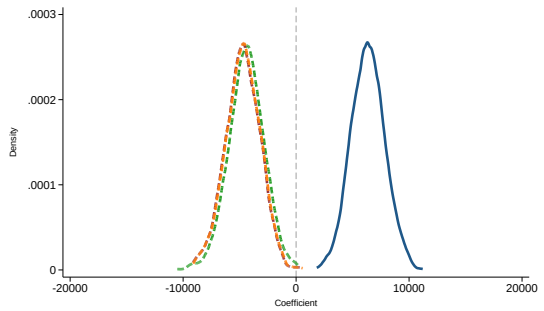
Figure A.14: Comparison of Reduced Form Solutions with Simulated Outcome Variable (Cubic Trends)

—— No correction (Under 10 Bin: 100.0%, Over 90 Bin: 100.0%) - - - State-Year Fixed Effects (21.3%, 68.3%)

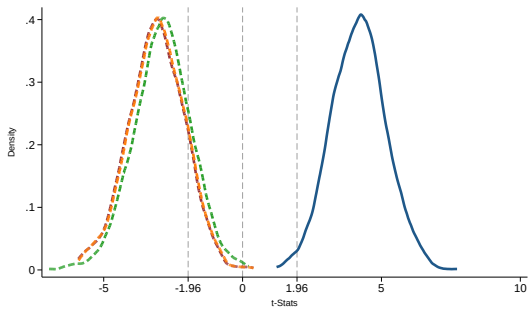
- - - With 3 Lags (99.0%, 100.0%) - - - County-Specific Linear Trends (81.2%, 6.0%)

- - - County-5 Year Fixed Effects (13.5%, 5.4%)

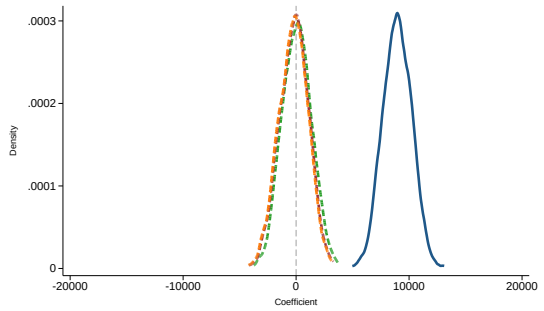
Notes: We conduct the same exercise as Figure 2, which uses simulated outcome data and real temperature data. Here, we plot the density of coefficients and t-statistics from estimating equation (10) 1,000 times, which only uses the two extreme temperature bins, alongside densities from specifications that use alternative reduced-form corrective strategies as indicated in the legend. The outcome is simulated to have a positive cubic differential trend in baseline temperature ($Y_{ct} = t^3 Temp_c^{base} + \epsilon_{ct}$) and daily temperature realizations are from the ERA-5 Land dataset. Each time, we take a draw of $\epsilon_{c,t}$ from a normal distribution centered at zero with variance equivalent to two times the in-sample standard deviation of $t^3 Temp_c^{base}$. Numbers inside the parentheses in the legend denote the number of runs per 1,000 simulations that the coefficient is significant at the 5% level, for the under 10°F bin and the over 90°F bin respectively.



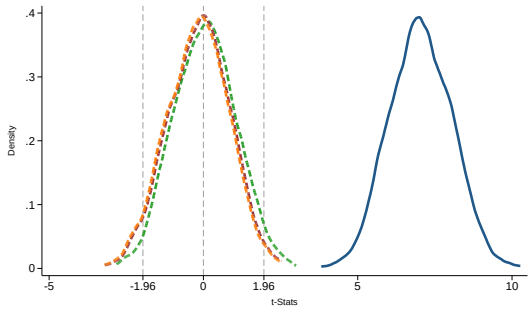
(a) Under 10 °F Bin: Coefficient



(b) Under 10 °F Bin: t-Stat



(c) Over 90 °F Bin: Coefficient



(d) Over 90 °F Bin: t-Stat

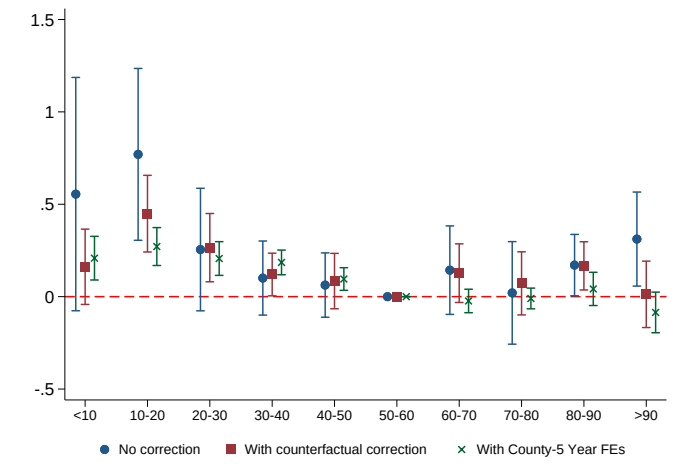
Figure A.15: Comparison of `cftemp` Solutions with Simulated Outcome Variable (Cubic Trends)

—— No correction (Under 10 Bin: 100.0%, Over 90 Bin: 100.0%)

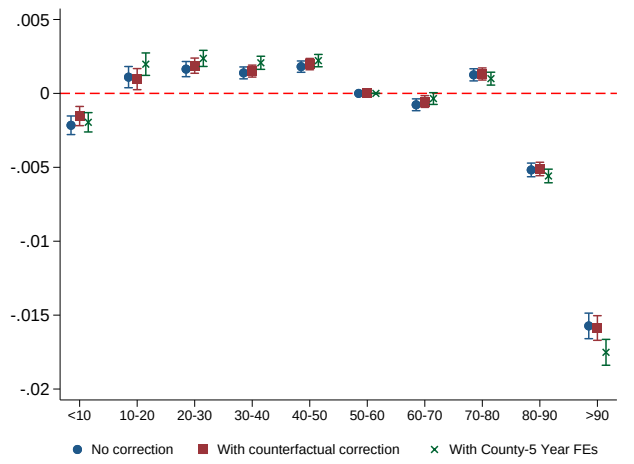
- - - Counterfactual Temperature Control (76.2%, 6.8%)

- - - Counterfactual Temperature Control without Bayes (61.3%, 6.2%) - - - Chebyshev (64.7%, 7.9%)

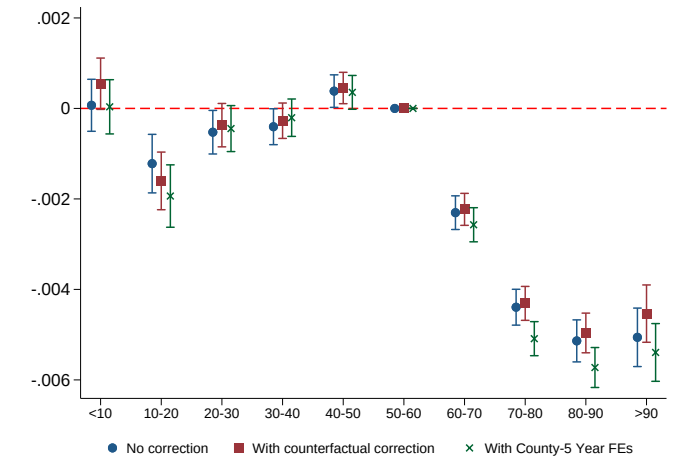
Notes: We conduct the same exercise as Figure 2, which uses simulated outcome data and real temperature data. Here, we plot the density of coefficients and t-statistics from estimating equation (10) 1,000 times, which only uses the two extreme temperature bins, alongside densities from specifications that also include the expected counterfactual exposure to each temperature bin as controls. Each color and line plots densities that use different methods to construct the counterfactuals, as indicated in the legend. The outcome is simulated to have a positive cubic differential trend in baseline temperature ($Y_{ct} = t^3 Temp_c^{base} + \epsilon_{ct}$) and daily temperature realizations are from the ERA-5 Land dataset. Each time, we take a draw of $\epsilon_{c,t}$ from a normal distribution centered at zero with variance equivalent to two times the in-sample standard deviation of $t^3 Temp_c^{base}$. Numbers inside the parentheses in the legend denote the number of runs per 1,000 simulations that the coefficient is significant at the 5% level, for the under 10°F bin and the over 90°F bin respectively.



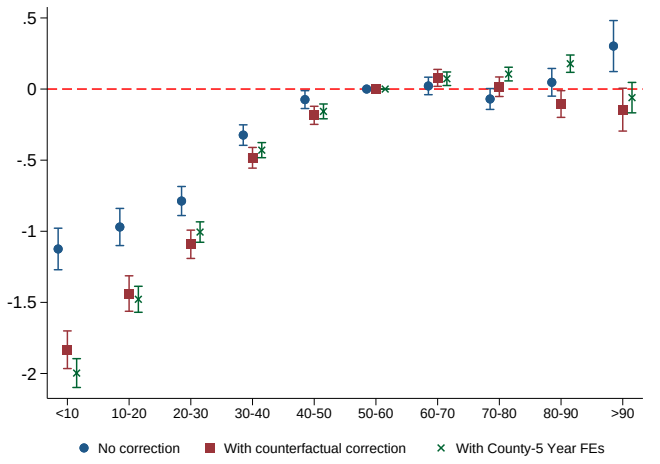
(a) Mortality Rate (All Ages, Stacked)



(b) log Soy Yield



(c) log Wheat Yield



(d) Nonviolent Crime

Figure A.16: Additional Real Outcomes (Counterfactual Controls with Empirical Bayes and County-5 Year Fixed Effects)

Notes: This figure reports estimates of (1) for outcome variables not reported in Figure 5, alongside estimates of (3) that also include the i) expected counterfactual exposure to each temperature bin or ii) 5-year times county fixed effects as controls. Daily temperature realizations in each county used to construct each temperature bins are from the ERA-5 Land dataset. In Panel (a), the unit of observation is a age group-county-year (1970-2002) and the outcome variable is the mortality rate for each age group, measured using the CDC mortality files. The age groups are constructed as ages < 1, 1-44, 45-64, and > 65 following Deschênes and Greenstone (2011). In Panels (b) and (c), the unit of observation is a county-year (1985-2006) and the outcome variable is log of soybean yield (b) and log of wheat yield (c), measured using NASS Quickstats. In Panel (d), the unit of observation is a county-month (1970-2009) and the outcome variable is nonviolent crime (i.e., robbery, larceny) measured using data from the FBI. 95% confidence intervals are reported and standard errors are clustered by county.