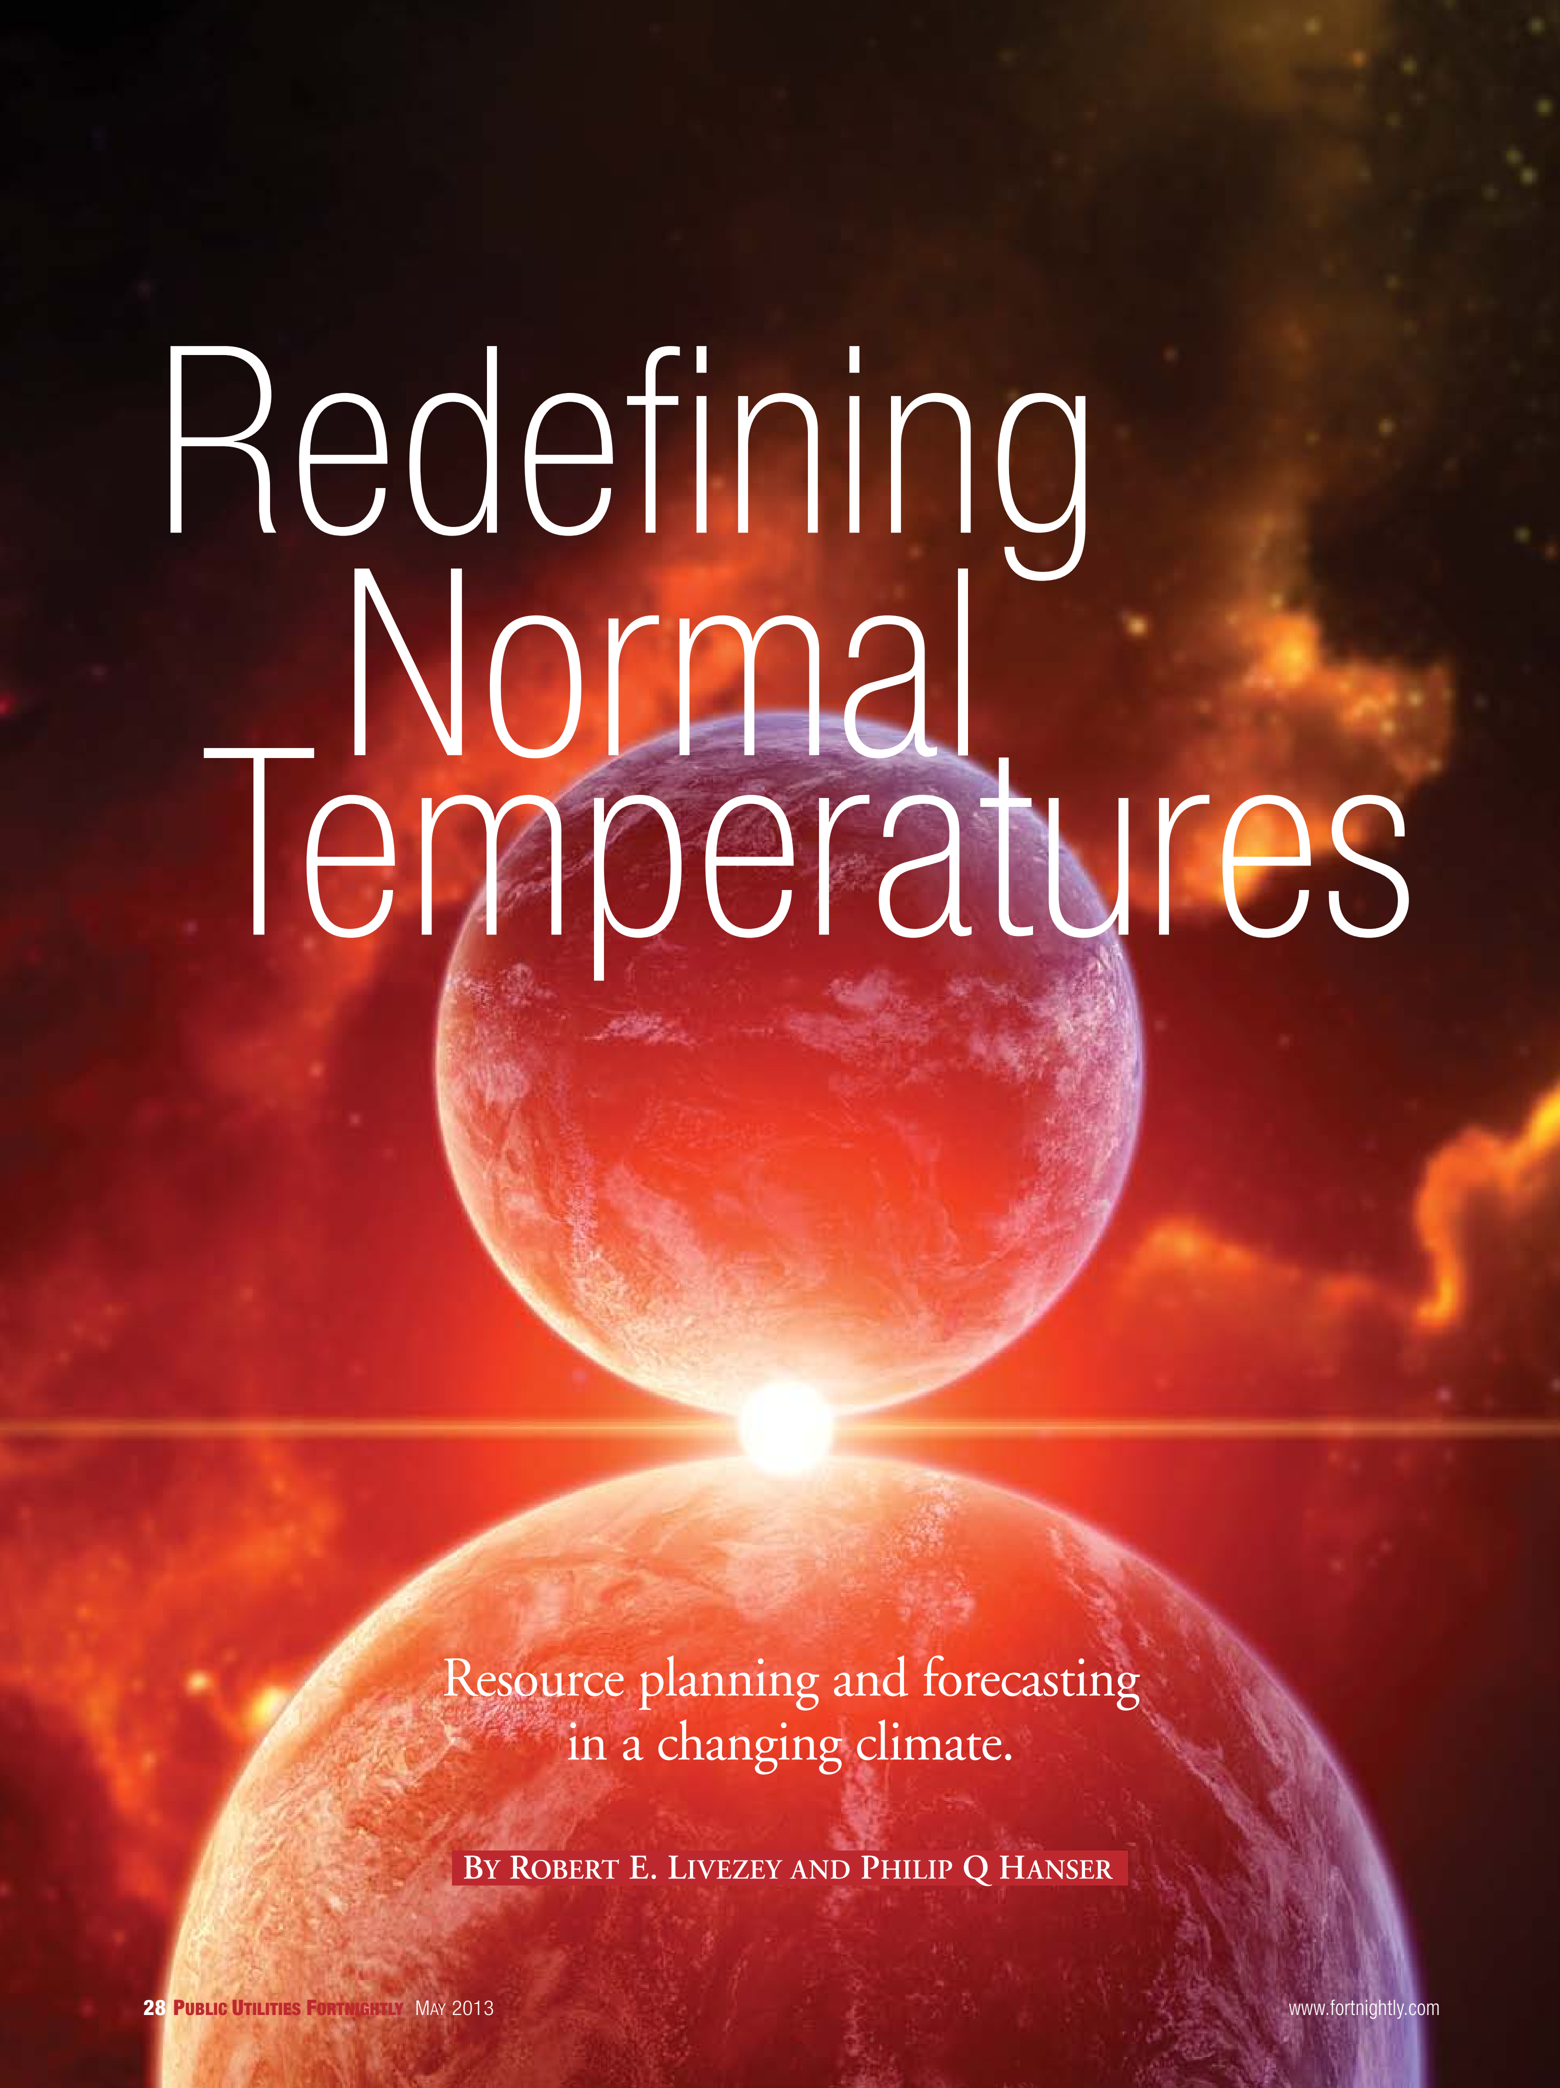

Redefining Normal Temperatures

Resource planning and forecasting in a changing climate.

BY ROBERT E. LIVEZEY AND PHILIP Q HANSER

The utility industry has long wrestled with the effects of temperature on energy use, whether it's summer heat waves driving peak electric demand to new heights and the electric system to record stress, or frigid winters straining gas pipeline capacities.

Extreme temperatures are likely to pose continual problems and perplex utilities' operators and planners. Defining what's "normal" for weather and temperature creates its own challenges. An empirically objective approach, based on historical data and climate science, can help utilities avoid the effects of inappropriately defining normal weather.

The Uses of 'Normal Weather'

"Normal weather" is defined as the statistical expectation of the temperature and precipitation for a location. As a measure of the climate or average weather, it's a fundamental input to many facets of utilities' operations, planning, and finances.

The electric grid's capability to carry power during the summer is often limited by how much transmission lines sag in the heat, which can result in the potential to fault to ground. The operating efficiency of power plants and transformers decays during the summer because of their reduced ability to rid themselves of waste heat. Gas in pipelines expands as ambient temperatures rise, raising the pressure within them. Utility resource planners need normal temperatures to base their estimates of future needs. Most demand forecasts are based on normal temperature with extremes treated as a means of bounding the limits of forecasts. Thus, there are myriad ways in which temperature affects utilities' operations and planning.

On the financial side, besides a utility's use of normal temperatures to forecast future loads for assessing and planning for resource needs, that same load forecast is translated into a sales forecast for estimating future revenues and will form part of the billing units basis for meeting its revenue requirement and, thus, rates. That revenue forecast will likely be shared with the financial executives who help secure the utility's finances. The finance community often asks, "What would the utility's revenues have been if the weather wasn't so cold (or hot) or had been colder (or hotter)?" Again, normal temperatures serve as the basis for assessing these scenarios.

30-Year Averages

The traditional and most commonly used approach for estimating normal temperatures—frequently called "normals"—is the use of a 30-year average of surface temperatures observed at geographically appropriate weather stations.[1] Normals can be estimated for specific days and weeks, as well as months and seasons, depending upon whether daily observations are archived. Such 30-year averages have been international standards for measuring climate for more than 70 years.

Failing to account for the trends in climate can have large and material implications.

As long as the climate is stationary (not changing) averaging over 30 years produces normal temperature estimates that are sufficiently accurate—*i.e.,* they're representative of the period but also reasonable predictors of next year's values. Moreover, given an unchanging climate, the 30-year normal will be unbiased (on average neither too cold nor too warm) and quite stable (changing little) when updated annually. These attributes, of course, are the reason for widespread acceptance of the 30-year normals as a standard. What error there is in the 30-year standard's representation of a particular year, and any change that occurs when it's updated, is entirely a result of the random and unpredictable year-to-year fluctuations for a particular week, month, or season, which climatologists refer to as "climate noise."

However, if the climate is changing rapidly, a retrospective average as long as 30 years will be biased one way or another depending on whether the climate is cooling or warming. In these instances, 30-year averages can no longer be considered representative of the current climate, and will produce less-accurate estimates of expected temperatures because the error is now a combination of the random error from the climate noise and the bias error from the climate change.

U.S. climates have been warming widely at moderate to rapid rates for three to four decades. This is illustrated for the nation as a whole for both the coldest and warmest times of the year by temperature histories in Figure 1,[2] and geographically for January through March in Figure 2.[3] Consequently, traditional

1. Official" normals are 30-year averages that are updated at the end of every decade, e.g., 1971-2000, 1981-2010, etc.

Robert E. Livezey retired in 2007 as chief of the National Weather Service/NOAA climate services division. **Philip Q Hanser** is a principal with The Brattle Group. The views in this article are the authors' and don't reflect those of The Brattle Group or its clients.

2. See Wilks, D. S., and R. E. Livezey. "Performance of Alternative 'Normals' for Tracking Climate Changes, Using Homogenized and Non-homogenized Seasonal U.S. Surface Temperatures," *Journal of Applied Meteorology and Climatology*, in press.
3. See Livezey, R.E., K.Y. Vinnikov, M.M. Timofeyeva, R. Tinker, and H.M. van den Dool. "Estimation and extrapolation of climate normals and climatic

FIG. 1 — TRACKING CLIMATE CHANGE

Time histories of: a) cold; and b) warm season temperatures, averaged over 1218 locations across the United States, showing the warming climate. Double-headed arrows are 30- and 15-year averages for December through April 2010 and 2011 and June through October 2011, respectively; the dashed lines are straight-line trends; and the solid lines are hinges (see article text) fitted to all but the last data year (x's).

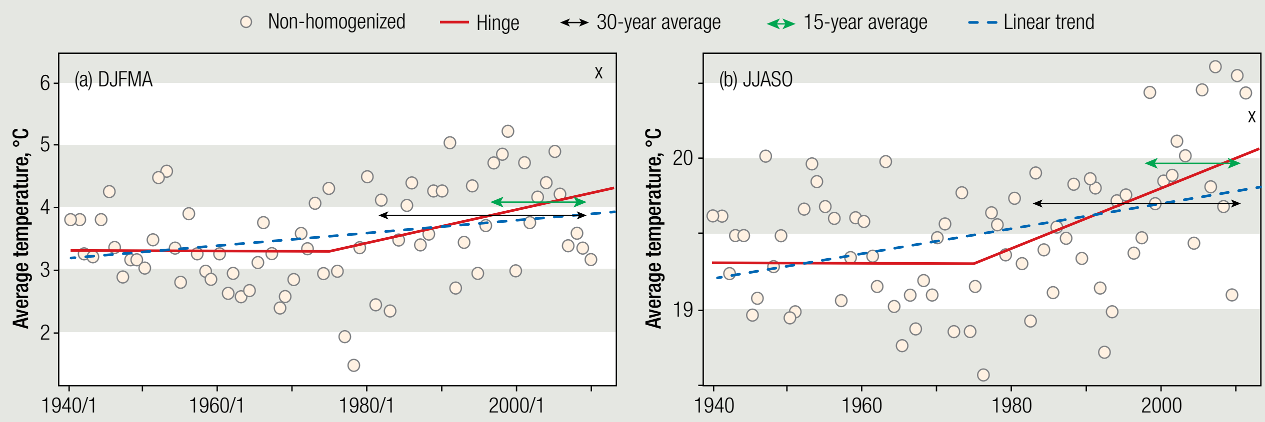

Source: "Performance of Alternative 'Normals' for Tracking Climate Changes, Using Homogenized and Non-homogenized Seasonal U.S. Surface Temperatures," *Journal of Applied Meteorology and Climatology* (forthcoming).

30-year normals generally will be expected to be cold-biased, often substantially. These biases for 2012 are evident in Figure 1, where the long double arrows denote 30-year averages ending in 2011 and the x's denote 2012 values.

Thus, a first consideration for determining normal temperatures is finding and using alternatives to 30-year normals that reduce both bias and total error, hopefully without a significant sacrifice in year-to-year stability. Three well-studied alternative approaches illustrate the pros and cons of each alternative. Before introducing these alternatives, there's a second important consideration, specifically the quality and representativeness of the observational temperature records to which these methods will be applied.

Weather Station Records

Temperature records from weather stations are collected and maintained by the National Oceanic and Atmospheric Administration's (NOAA) National Climate Data Center (NCDC). All of these records are available with some degree of minimal quality control, but a subset of these stations with relatively long histories have been homogenized by NCDC.[4]

The purpose of the homogenization is to ensure that the temperature record has no missing values and is only representative of meteorological and climatological changes that have taken place at the station's location.[5] Weather stations can naturally relocate, the environment around a station can drastically change, equipment might require recalibration, and observing protocols can change. All of these changes can mask or distort real, underlying climate changes at the station, so it's essential to remove them with the best science available before normals can be determined. This is the case regardless of the methods used, because all normals rely to some degree or another on history, some much more than others. In fact, all official normals from NCDC are produced from fully homogenized records, because a 30-year retrospective view will be particularly sensitive to inhomogeneities in the history.

There's little climate change from 1940 until sometime in the 1970s and mostly steady warming thereafter.

The degree to which non-climatological trends or abrupt changes compromise the climate history at stations varies widely, but most stations will be noticeably affected. Figure 3 shows an egregious rate case example for Broken Bow, Neb. The figure graphs the heating season (October through April average) temperature histories, both before and after homogenization is performed on the data. Note that the trend in the normal in the non-homogenized history is negative, suggesting a cooling climate. After homogenization, the trend is reversed, indicating warming. Even though the most recent 30-year averages are about the same for both records, inferences about their respective biases and suitability for representing normal temperature are quite different. Data inhomogeneities have completely distorted the climate history at this location. This and many other examples demonstrate the necessity to determine temperature normals only from homogenized data.

trends," *Journal of Applied Meteorology and Climatology*, 46 (2007), 1759-1776.

4. The set of 1218 station records referred to as the U.S. Historical Climate Network (USHCN) is available in a fully homogenized version.
5. See Menne, M.J., and C.N. Williams. "Homogenization of temperature series via pairwise comparisons," *Journal of Climate*, 22 (2009), 1700-1717. The website http://variable-variability.blogspot.com/2012/01/homogenization-of-monthly-and-annual.html has a less-technical description of the process.

When Normal is Moving

Even with limited knowledge of warming trends, two simple alternatives to a 30-year average can reduce cold biases and potentially reduce errors in representing the current climate or predicting next year's value: 1) use of a straight-line trend fit to the temperature history with the normal defined as the recent end-point of the fit; and 2) a shorter period average—*e.g.,* 15 years. Both of these alternatives are shown for the warming temperature histories in both panels of Figure 1 and the second panel of Figure 3; trend fits are the dashed lines and 15-year averages are the shorter double-headed arrows.

Using a straight-line trend fit to a very long record (more than four decades) turns out not to be an effective method, even though it produces normals that will change little when updated annually—*i.e.,* they are very stable. The reason is apparent in the three examples in Figures 1 and 3b: the biases of the trend-fit normals—the last points on the dashed trend lines—are still substantially cold, at best only slightly warmer than for the 30-year normal. The trends are underestimated because the warming has only been taking place since the mid-1970s. A better approach would be to restrict the trend fit to just the warming period, ensuring a steeper slope and a warmer, less-biased current value. However, the reduction in bias comes with a substantial reduction in stability, so simple trend fits aren't recommended. Instead, an approach called the "hinge fit," which doesn't require this trade-off, can be used.

The other way to reduce the cold bias in an alternative normal for a warming climate is to use a shorter-than 30-year average. This is easy to visualize if one considers a steadily warming climate: a most recent 30-year average would be most representative of the climate 15 years ago, a 20-year average of the climate 10 years ago, etc. Note that in all three examples in Figures 1 and 3b, the 15-year average is warmer than the 30-year average by a meaningful amount.

However, this reduction in bias with shorter-period averages comes with a price. The shorter-period averages have higher statistical sampling error, the so-called "standard error" of estimation, a consequence of the climate noise referred to earlier. Thus, there's no advantage gained unless the increase in standard error for a shorter-period average is less than the decrease in bias error. The averaging period with the smallest sum of the bias error and standard error is what climatologists call the "optimum climate normal" (OCN).[6]

FIG. 2 REGIONAL TEMPERATURE TRENDS

January through March warming trend estimates from 1975 to 2005 based on hinge fits (see article text). Note the large contrast from northwest to southeast.

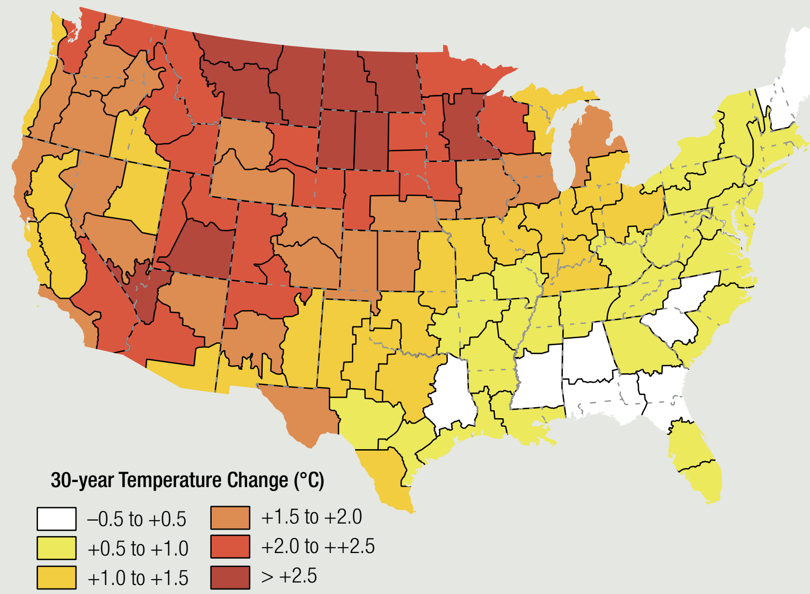

See Livezey, R.E., K.Y. Vinnikov, M.M. Timofeyeva, R. Tinker, and H.M. van den Dool. "Estimation and extrapolation of climate normals and climatic trends," *Journal of Applied Meteorology and Climatology*, 46 (2007), 1759-1776.

Numerous studies have demonstrated that OCNs for U.S. locations at all times of the year are overwhelmingly averages over much less than 30 years. There are several approaches available for tailoring OCN to a particular location and season,[7] but a number of studies in the last two decades have suggested that, overall, the best averaging period is between 10 and 15 years.[8] NOAA's Climate Prediction Center has used 10-year OCNs since the mid-1990s to make seasonal temperature forecasts based on the earliest of these modern studies.[9] Recently, however, convincing evidence has been published showing that 15-year averages have actually been the

30-year averages can no longer be considered representative of the current climate, and will produce less-accurate estimates of expected temperatures.

6. See Livezey, *et al.*, *op. cit.*
7. See Wilks and Livezey, *op. cit.* and Wilks, D.S. "Projecting 'normals' in a non-stationary climate," *Journal of Applied Meteorology and Climatology*, 52 (2013), 289-302.
8. See Wilks and Livezey, *op. cit.*, Wilks, *op. cit.*, and Huang, J., H.M. van den Dool, and A.G. Barnston. "Long-lead seasonal temperature prediction using optimal climate normals," *Journal of Climate*, 9 (1996), 809-817.
9. See Huang, *et al.*, *ibid.*

FIG. 3 HOMOGENIZATION: BEFORE AND AFTER

Heating season (October through April) average temperatures for Broken Bow NE (a) before and (b) after homogenization of the history. "1940" denotes 1940 and '41, etc.; lines are the same as in Figure 1.

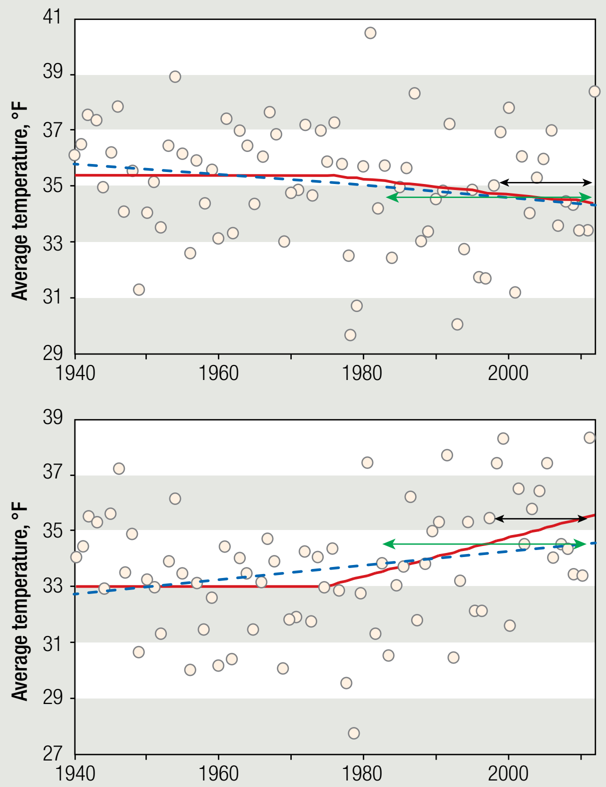

Source: Author's analysis, NOAA National Climate Data Center data

FIG. 4 MIDWEST WARMING

Winter (December through February) average homogenized temperatures for Minneapolis-St. Paul Airport. "1939" denotes 1939 and '40, etc.; lines are the same as in Figure 1.

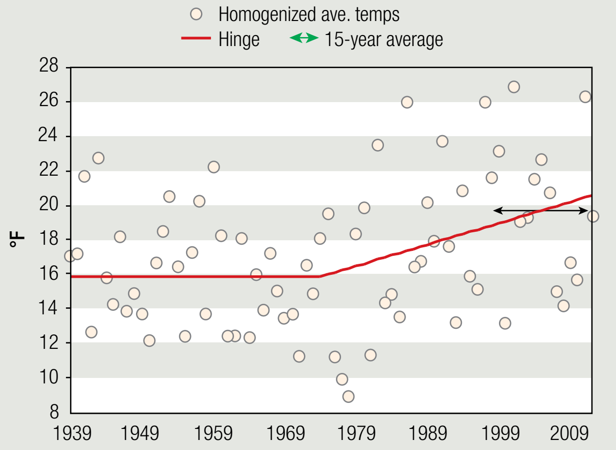

Source: Author's analysis, NOAA National Climate Data Center data

best performing alternative normals overall during this period, especially under demanding circumstances like the major interruption of warming in the cold-season United States (especially in the West) from 2007 and '08 to 2010 and '11.[10]

In summary, the OCN, whether it consists of a fixed 15-year average or is tailored to a location and season, is the superior viable alternative to 30-year averages in the context of both bias and total error. Obviously, some stability is sacrificed to achieve these gains.

Hinge-Fit Alternative

Under some circumstances, a better alternative to OCN is the hinge fit, with even smaller biases and errors without any compromise in stability whatsoever. A good way to think of the hinge fit is as a simple statistical model of climate trends. Examining both panels of Figure 1, along with Figure 3b, a similar pattern is evident in all three histories; there's little climate change from 1940 until sometime in the 1970s and mostly steady warming thereafter, with climate noise superimposed throughout. It turns out that this pattern is evident in temperature histories for the globe as a whole, the globe's land masses and oceans separately, and its two hemispheres and the different continents, whether for the warm or cold halves of the year. The pattern is also replicated by computer models of the global climate.[11]

Based on these observations, the first author proposed that changing normals be represented by a hinge form, specifically no change until 1975, then steady, straight-line warming thereafter. Thus, the current normal would be represented by the last point on the fit, in the same manner as the simple trend fit. In mathematical jargon, the hinge is piecewise-continuous and -linear[12] with zero slope to 1975.[13]

Least-squares hinge fits are represented in all of the panels in Figures 1 and 3 by solid lines. In Figure 1 the hinge fits are extrapolated one year to compare to the x data points and a bit beyond, even though the fits were based on only the observations excluding these last points. To avoid making any assumptions about future warming, the hinge estimates of current normals in Figure 1 should be considered the points on the hinge corresponding to the last "o" data points. With this in mind, the hinge current normal estimates in the three warming cases are all slightly warmer than the 15-year OCNs, so their cold biases should be slightly smaller.

Thus, for the examples in Figures 1 and 3, the hinge fits have just a small advantage regarding bias. However, the hinge fits are far more stable than the OCNs; after all, the hinges are fitted to

10. See Wilks and Livezey, *op. cit.* and Wilks, *op. cit.*

11. See Solomon, S., D. Qin, M. Manning, Z. Chen, M. Marquis, K. B. Averyt, M. Tignor, and H. L. Miller, Eds., *Climate Change 2007: The Physical Science Basis*. Cambridge, UK: Cambridge University Press, 2007.

12. A data plot is termed "piecewise-continuous" if a function is defined throughout the interval, its constituent functions are continuous on that interval, and there's no discontinuity at each endpoint of the subdomains within that interval. The plot is "piecewise-linear" if the function is comprised of straight-line sections.

13. See Livezey, *et al.*, *op. cit.*

over 60 years of data. And for roughly 25 percent of all location-seasons since the early 1990s—those cases with the strongest warming trends—the hinge fit will substantially improve on the OCN's bias with hardly any or no sacrifice in total error.[14] With very strong warming, like that for winter-time Minneapolis-St. Paul shown in Figure 4, the hinge fit is expected to outperform the OCN in every respect.[15] Here, the hinge current normal is almost a degree warmer than a 15-year average.[16]

An important and legitimate question has been raised about the validity of the hinge model. The assumption that modern climate change began in 1975 is well-supported and corroborated by independent and modeled data, yet it's still arbitrary. The recent studies, referred to above, that examined the performance of alternative normals since the mid-1990s also considered three other variant hinge models with: 1) fitted non-zero slopes prior to 1975 (allowing warming or cooling up to the change point year); 2) fitted change points (other than 1975); or 3) both (fitted change points with warming or cooling allowed up to the change point years). The use of fitted change points rather than fixing them at 1975 in variants 2 and 3 degraded performance substantially. Just allowing warming or cooling up to 1975 rather than keeping the normal unchanged in variant 1 made little difference in performance. These results validate the original design of the hinge-fit normal.[17]

Thus, two viable alternatives to 30-year averages for normals estimation—OCNs (15-year or less averages) and hinge fits—are available to reduce both biases and total errors in representing current climates. Some stability is sacrificed with use of the OCN, but is gained by use of the hinge. The choice of which to use depends on how fast warming is taking place for a particular location-season. Some guidance is available in the literature for making this choice.[18]

Utility Forecasting

A variety of aspects of utility planning and forecasting would benefit from better estimations of normalized weather. The first is sales and revenue projections. Failure to properly normalize the weather is particularly problematic for retail gas distribution rates, which are essentially entirely volumetric and more dependent on weather-related uses than are electric rates. Rates are typically set based on a forecast of expected sales and are trued-up based on the difference between the observed and the normal weather. If, for example, the winter normal weather forecast is too low, then projected sales volumes will be overstated and the utility will suffer from reduced rates of return, because the volumetric charge will be set too low. The gas utility is then placed in the position of having to go back to its public utility commission to true-up its rates to achieve its revenue requirements, not a comfortable position for either the utility or its commission.

Electric utilities suffer similarly and for the same reason, although with differing implications. Wall Street requires quarterly updating of weather-normalized sales. As with gas utilities, a biased basis for normal weather will yield biased estimates of normalized sales. The typical normalizing sales model looks something like: "Sales = b + m1 x HDD + m2 x CDD," where m1 = usage/heating degree-day; m2 = usage/cooling degree-day; and b = base usage (non-weather sensitive). The electric utility will then calculate its normalized sales as: "Normalized sales = Actual sales + m1 x ("normal" HDD – actual HDD) + m2 x ("normal" CDD – actual CDD)."

The assumption that modern climate change began in 1975 is well-supported and corroborated by independent and modeled data, yet it's still arbitrary.

For electric utilities, the problem becomes the variability of the forecast; not just a levels problem, as it is for gas utilities. If the utility relies upon traditional 30-year averages, then year-to-date forecasts completed after the winter season likely will show sales below those forecast because the normalized winter weather will, on average, be based on an over-forecast of heating-degree days. Forecasts following the summer will be the reverse; normal cooling-degree days will, on average, be lower than actuals. These swings will increase the variability of the financials, a situation that Wall Street usually dislikes, putting the utility—and its forecaster—in an uncomfortable position.

For those utilities, including regional transmission organizations setting future energy and capacity requirements, the use of incorrect weather normals will lead to possibly understated or misstated needs. If the weather, on average, is warmer than expected during the summer, then several resource capabilities will be reduced below expectations. For example, in warm weather, generator efficiency is lower, transmission line loadings might need to be reduced because of increased line sag, and the efficiency of transformers is reduced, all of which could imply higher resource requirements than projected as a result of the reduced capabilities of these resources. In addition, although peak temperatures might not be any higher, the duration of temperatures above projected normals could increase resource requirements. On the other hand, during

14. See Wilks and Livezey, *op. cit.*
15. See Wilks and Livezey, *op. cit.* and Livezey, *et al.*, *op. cit.*
16. The OCN optimally tailored to this case (DJF at Minneapolis/St. Paul AP) is a 12-year average.
17. See Wilks and Livezey, *op. cit.* and Wilks, *op. cit.*
18. See Wilks and Livezey, *op. cit.* and Livezey, *et al.*, *op. cit.*

the winter, the reverse might be true, although the effects aren't generally as significant for resource planning during the summer.

In addition, nonstationarity also applies to other climatological phenomena, certainly for rainfall,[19] an issue of very large importance for generation. Changes in rainfall patterns clearly affect hydroelectric generation, but in many areas the use of water for cooling electric generation stations represents the largest single use of water beyond that for human consumption; in some areas, even greater than that for human consumption. Thus, shortfalls in water have enormous implications for generation availability and use. However, research is still needed to identify the best alternatives to 30-year rainfall averages. Regardless, normals that mislead as a result of an incorrect approach to their forecast, whether for temperature or precipitation, put utilities in a position where they don't properly account for the resource risks they face. Thus, failing to account for the trends in climate can have large and material implications.

Beyond temperature and precipitation—for example trends in high-impact storms including wind, snow, ice, and other short-term but highly damaging weather events—there's little specific to recommend. At this time there's no expert consensus or convincing evidence that the risks of these hazards are generally increasing. Arguments for their increase to date are largely anecdotal or supported by inadequate models. Continued climate warming will certainly lead to shifts in weather patterns, but there is a large uncertainty as to what these shifts will be and where. In the case of extreme weather, the best strategy for utilities is to continue to objectively monitor the sequence of year-to-year weather and the developing peer-reviewed science.

Reduction in bias with shorter-period averages comes with a price: a higher statistical sampling error.

Adapting to the current trends in climate requires accounting for weather in a way that's consistent with the data. For utilities, one aspect of this is to properly account for what normal weather will be going forward. Two alternatives—OCNs and hinge-fits—both have a basis in climate science and empirical verification, and thus provide better estimates of normal weather than traditional 30-year averages do. The risks of improper accounting for changes in normal weather are real, and ignoring them is potentially costly.

19. See Livezey, *et al.*, *op. cit.*