J2.4 SKILLFUL SEASONAL DEGREE-DAY FORECASTS AND THEIR UTILITY IN THE WEATHER DERIVATIVES MARKET

Jeffrey A. Shorter, Todd M. Crawford, Robert J. Boucher, James P. Burbridge
WSI Corporation, Billerica, MA

1. INTRODUCTION

The majority of weather derivatives are temperature-based, using degree-days, and cover periods ranging from one week to multiple months. They are usually priced using climatological data, since there is a perceived lack of skill for forecasts beyond 10 days. The prevailing view of seasonal forecasting was succinctly summarized by a weather trader at the June 2001 Weather for Risk Management Association meeting, who quipped, “Seasonal forecasts are garbage.”

However, this view is beginning to change due to two recent developments: 1) the demonstration of skillful operational seasonal forecasts and 2) enabling weather traders to construct statistical pricing models by optimally blending forecasts and climatology. This presentation will address both of these topics.

The first development, the achievement of seasonal forecast skill, has taken significant amounts of time and capital. WSI’s Energycast Trader group has been delivering national seasonal forecasts to its clients since April 2000. The forecast skill in the intervening months has been assessed using both mean absolute errors with respect to standard climatology and directional correctness. In November 2000, WSI started providing specific degree-day forecasts for selected cities on seasonal timescales.

The second development discussed involves the construction of statistical pricing models for use by weather traders using current forecasts and historical weather data. We will discuss (1) traditional pricing models currently in use, (2) blending of seasonal forecasts into the traditional pricing models, and (3) utilization of seasonal forecasts to guide buying and selling decisions. The pricing models will require a forecast probability distribution, rather than a deterministic forecast. This distribution is derived based on both the forecast accuracy and consistency of a 20-year dynamical climate model run.

*Corresponding author address: Jeffrey A. Shorter, WSI Corporation, 4 Federal Street, Billerica, MA 01821-3569; e-mail: jshorter@wsi.com

2. SEASONAL FORECAST SKILL

Starting in November 2000, WSI has issued twice-monthly seasonal degree-day forecasts for each of 11 cities (ATL, BOS, ORD, CVG, IAH, LIT, LGA, PHL, DCA, SAC, PHX). The 11 cities were chosen to match those on the Enron Online trading platform.

The first (primary) forecast is issued two days before the Climate Prediction Center (CPC) seasonal forecast. The second (update) forecast is issued 1-4 days before the termination of trading in a given month (“bid day”). Forecasts are provided for both the entire subsequent three-month “season” and each of the following three months individually (hereafter referred to as Month 1, 2, and 3, respectively).

Forecasts were made for either HDDs or CDDs, depending upon the time of year. The 1990-1999 climatology was used to determine which type of degree-day was more prevalent in a given month – e.g., if BOS averages 200 HDDs and 100 CDDs in a given month, then we forecasted only HDDs.

The primary metric we have used to evaluate forecast accuracy is “directional correctness.” Directional correctness simply means that our forecast was “on the right side of the fence,” either predicting above normal and verifying above normal or predicting below normal and verifying below-normal.

A secondary metric, accuracy, is simply the ratio of the forecast error to the observed value. However, accuracy is not a particularly robust metric since a fixed forecast error will produce significantly different accuracy values in December than in March. For example, in December an observed HDD value may be 1000, while in March 100 HDDs are more common. Assuming the same forecast error, e.g., 10 degree-days, for both months produced vastly different values of accuracy (1% versus 10%). However, the metric is still useful for comparing, for example, monthly forecasts at different ranges verifying on a given month (Month 1, 2, and 3 month forecasts verifying in September, for example).

2.1 WSI Seasonal Forecasts

To date, we have issued 8 primary and 8 updated seasonal forecasts (DJF through JAS). Our seasonal forecasts have been quite skillful, as the table of directional correctness shows (Table 1). On the whole, the seasonal forecasts have demonstrated a 77% rate of directional correctness to date, with no locations less than 50%.

CITY	DIRECTIONAL CORRECTNESS
SAC	100
PHX	77
IAH	100
LIT	88
ATL	94
ORD	69
CVG	81
DCA	50
PHL	63
LGA	69
BOS	69

Table 1. Percentage of directionally correct 3-month seasonal forecasts for each of the 11 cities

In fact, all 16 forecasts have been directionally correct at both Houston and Sacramento. At Houston (Table 2 and Figure 1), WSI correctly predicted colder-than-normal temperatures from DJF to FMA, slightly warmer-than-normal in MAM, and then cooler-than-normal from AMJ to JAS. At Sacramento (Figure 2), WSI correctly predicted colder-than-normal DJF to MAM, warmer-than-normal in AMJ and MJJ, and cooler-than-normal in JJA and JAS.

	Normal	Forecast	Observed
DJF	***924***	***1082***	***1238***
JFM	***750***	***967***	***968***
FMA	***425***	***498***	***497***
MAM	499	516	563
AMJ	1032	956	1025
MJJ	1474	1404	1377
JJA	1696	1584	1613
JAS	1612	1543	1536

Table 2. Degree-day normals, forecasts, and verification for IAH in 2001. Italicized and bolded values represent HDD forecasts; the remainder are CDD forecasts.

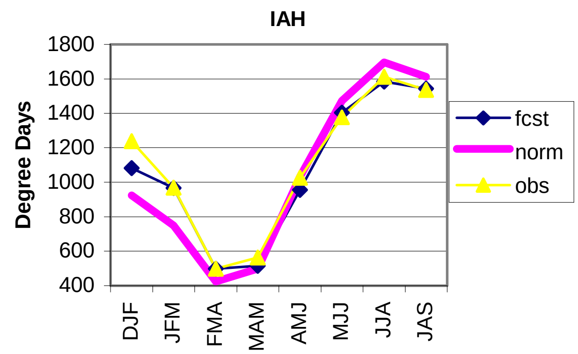

Figure 1. Primary seasonal degree-day forecasts, verification, and normals for Houston, TX. HDDs were forecasted for DJF through FMA, CDDs thereafter

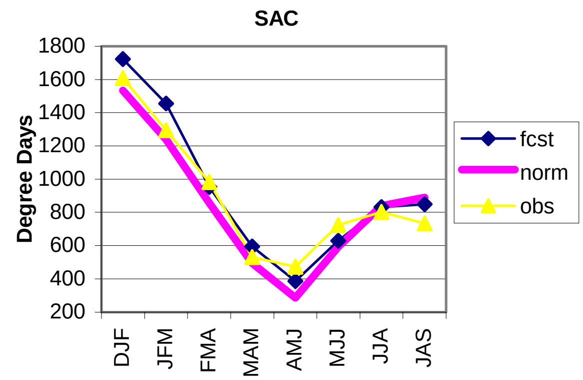

Figure 2. Primary seasonal degree-day forecasts, verification, and normals for Sacramento, CA. HDDs were forecasted for DJF through MAM, CDDs thereafter

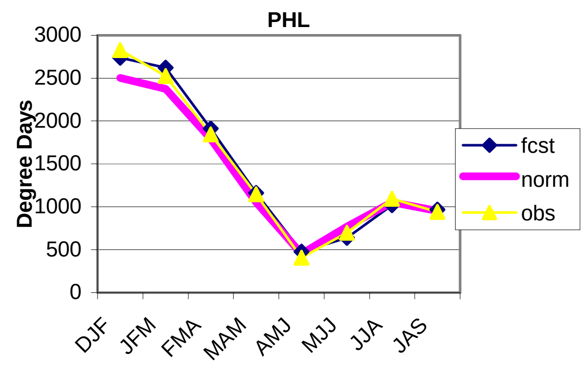

Figure 3. Primary seasonal degree-day forecasts, verification, and normals for Philadelphia, PA. HDDs were forecasted for DJF through AMJ, CDDs thereafter. Note that the most anomalous periods (DJF, JFM, and MJJ) were forecasted well.

The WSI forecasts were successful to a lesser extent in the Northeast, with an average of 63% directionally correct forecasts. A typical result in the Northeast is illustrated in Figure 3 for Philadelphia.

2.2 Primary vs. Updated Forecasts

As mentioned above, WSI issues seasonal forecasts twice-monthly for the benefit of our clients. The forecasts are based primarily on the processed output of our in-house dynamical climate model. Further guidance is gained from additional factors including soil moisture, snow cover, and month-to-month persistence. The update forecast can also be influenced by more traditional medium-range guidance if a strong signal is apparent early in Month 1.

A comparison of the primary and update forecasts is outlined in Table 3. There are only small differences between the directional correctness of the primary forecast and the update forecast. In both Month 3 and Season, the directional correctness values are similar for the primary and update forecasts. For Month 2, the value is slightly higher for the updated forecast. A closer examination of the numbers (not shown) reveals that most of this advantage can be attributed to our July 25th updated forecast for September. In this forecast, an appropriate statistical method was given more weight instead of our climate model forecast, with very successful results. Month 1 updated forecasts are also slightly better than the primary forecasts. This is primarily driven by human adjustment of the forecast due to medium-range guidance that extends well into the next month, and allows the forecaster to get a "headstart" on the month. It should be noted that the similarity between the scores of the primary and update forecasts indicates that the climate model is doing well and does not need to be adjusted most of the time.

	PRIMARY	UPDATE
Month 1	61	67
Month 2	63	68
Month 3	61	61
SEASON	76	77

Table 3. 11-city average forecast percentage directional correctness

3. FORECAST PROBABILITY DISTRIBUTIONS

The WSI operational seasonal forecasts are provided in probability distributions that are derived from the output of a 20-year dynamical model run. The previous section, which outlined the significant skill of the forecasts, used the center of the probability distribution as a proxy for a deterministic forecast for the analysis.

Representing the forecasts in a probability format is necessary since current pricing and risk management models evaluate risk based on the probability of the seasonal climate. The utilization of the probability distribution will become apparent in the description of the traditional pricing in Section 4.

3.1 Establishing Climatological PDF

In order to construct risk models to utilize with our degree-day forecasts, we have used an historical data set of observed seasonal HDD/CDD data for each of the 12 seasons (JFM to DJF). For each of the 11 locations and 12 seasons, the mean and standard deviation is calculated for use in constructing the climatological PDF. An example is shown in Figure 4 for Atlanta during JJA.

3.2 Establishing WSI's Forecast PDF

The method used in creating WSI's seasonal forecasts produces both a degree-day forecast and a measure of confidence in that forecast. These two measures are then used to create a forecast PDF to complement the climatological PDF. The confidence "factor" was derived from the skill of the historical climate model runs. This factor ranges from 0 to 100, where 100 represents a perfect forecast with no uncertainty, and 0 means the seasonal forecasting method has produced no useful information at a given location. In this case the forecast would simply be the same as the climatological PDF.

Two illustrative examples are shown in Figure 4. The tall, narrow distribution represents a high level of confidence, while the short, squat distribution represents a less confident forecast.

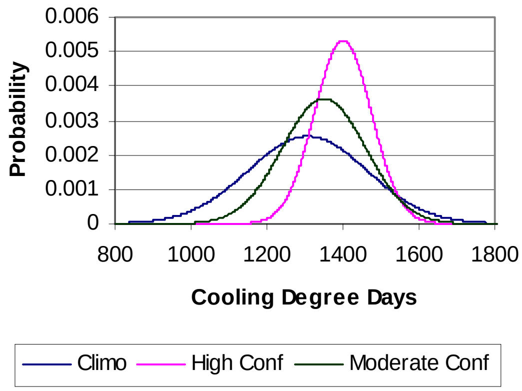

Figure 4. Examples of probability distributions

4. APPLICATION OF SEASONAL FORECASTING TO WEATHER DERIVATIVES

In this section, we will outline the traditional approach for pricing weather derivatives, an approach for blending seasonal forecasts into the existing methodology, and the use of seasonal forecasts in making derivative buying decisions.

4.1 Traditional Pricing Models

Pricing models for weather derivatives are typically based on observed distributions of recent (usually the last 10 years) temperature data. Some pricing models take into account previous trades that are still active and some even account for cross-commodity relationships, e.g., between weather and the price of electricity.

For the classic "swap" contracts, the mean of the distribution is used with the addition of a bid/ask spread. For example, if the 10-year average number of cooling degree-days was 1180 and the bid-ask spread was 20, then the short strike would be 1180-10 = 1170 and the long strike would be 1180+10 = 1190. For "puts", "calls", and "extreme weather" contracts, the trader uses the distribution to calculate the probability of return and sets the strike accordingly. It should be noted that each weather desk has their own proprietary pricing models.

4.2 Incorporating Seasonal forecasts into the Traditional Pricing Models

The WSI seasonal forecasts are delivered with probability distributions and as such could be directly substituted for the climate data in the pricing models. This may eventually be the case, but until traders embrace seasonal forecasting the more likely scenario is a blending of the forecast with climatology.

Applying a weight to each probability distribution and convolving them can help to achieve an optimal blend of climatology and the WSI seasonal forecast. The relative weights can vary depending on the trader's confidence in the seasonal forecast. Once the convolved distribution is computed, the 50% probability point would be derived. From this, a standard "swap" strike point can be computed.

The resulting probability distributions are shown in Figure 5. The shifting of the 50% probability point is illustrated in Figure 6 as a series of cumulative probability distributions going from pure climatology to pure forecast. The 50% points are also given in Table 4.

The blending approach has the net effect of biasing the strike points toward the forecasts. Thus this only adds the risk associate with the magnitude of the bias.

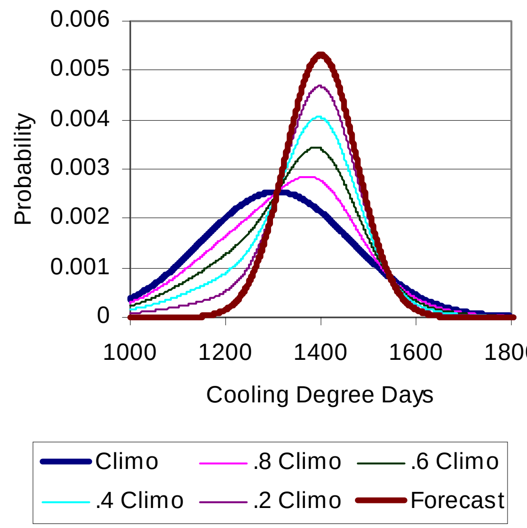

Figure 5. A series of probability distributions derived by weighting the climatological and forecast distributions.

4.3 Incorporating seasonal forecasts into the buying decisions

A second approach for utilizing seasonal forecasts is to use them for evaluating derivatives in the market. If a derivative is priced at the 10-year average with a narrow bid/ask spread and the seasonal forecast is showing a strong probability

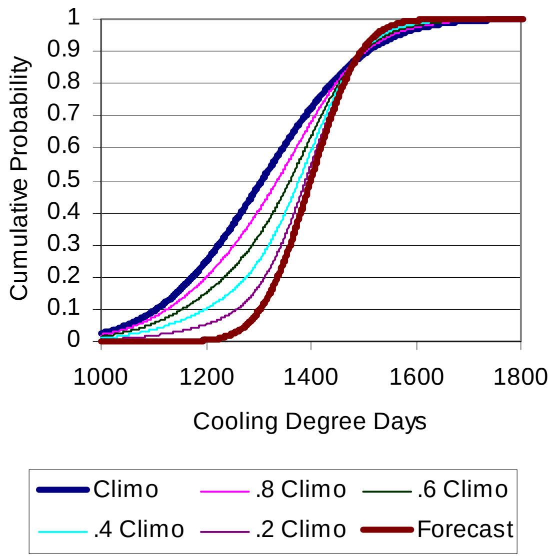

Figure 6. Blending climatology and seasonal forecasts in the form of cumulative probability distributions. The bold curve to the left is the pure climatology line. The narrow line to the right of it in a blend of 80% climatology and 20% forecast. The subsequent narrow lines are 60%, 40%, and 20% climatology. The bold line to the right is 100% forecast.

Climatology / Forecast Blend	50% point
Pure Climatology	1307
80% Climatology +20% Forecast	1339
60% Climatology + 40% Forecast	1359
40% Climatology + 60% Forecast	1378
20% Climatology + 80% Forecast	1390
Pure Forecast	1400

Table 4. Illustrates the shifting of the 50% probability point as the seasonal forecast is blended with the climatological distribution. Numbers are based off the 1990-1999 ATL cooling degree-day data and a fictional forecast.

for either warmer- or colder-than-normal, then the trader could buy the derivative. This approach lends itself more to speculating and has more risk associated with it.

This approach was used in evaluating the utility of the November 2000-September 2001 WSI seasonal forecasts. Assuming swaps could be bought at the 10-year average with a bid/ask spread of 40 degree days, and would pay $100 per degree day with a cap of $20,000 the return would have been $987,000. The payout per city is in Table 5. Note that all 11 cities provided a positive return.

CITY	Payout
SAC	114
PHX	37
IAH	131
LIT	127
ATL	177
ORD	85
CVG	77
DCA	13
PHL	77
LGA	77
BOS	72
Total	987

Table 5. Trading payout (in thousands of dollars) by city, assuming swaps were bought at the 10 year average with a 40 degree-day bid/ask spread. Note that every city provided a positive payout. The period was from November 2000 to September 2001.

5. SUMMARY

From November 2000 to September 2001, WSI validated seasonal forecasts (two per month, 16 total) covering the following three months for eleven cities. The forecasts were quite successful, with 77% directional correctness overall, and no location under 50%. At IAH and SAC, all 16 forecasts were directional correct.

Both the primary and update forecasts were skillful for all months, and especially for the whole season, where 76-77% of our forecasts were directionally correct.

The update forecasts for Month 1 and Month 2 were slightly better than the primary forecasts due to other forecasting factors that were used to improve on the skillful dynamical climate model output.

Statistical weather risk models have been established in order to use WSI's skillful seasonal forecasts in the weather derivatives market. Probability distribution functions were created for both climatological and forecast data at a given location in order to quantify the risk/reward relationship for any individual forecast.

We have shown that seasonal forecasts can be used by traders to improve profits in the weather derivatives market.