Weather Risk Management

At The "Frozen Falls Fuel Company"[1]

By
Bob Dischel, Ph.D., CCM

The following paper discusses methods of hedging with over-the-counter weather derivatives--specifically, Heating Degree Day (HDD) swaps and options. The examples used in this paper cite hypothetical over-the-counter derivative strategies used by the management of the fictitious firm, the "Frozen Falls Fuel Company."

It is May of 1999 and the Frozen Falls Fuel Company faces the possibility of failing. The company is a family business founded and run by "Flash" Burns. Flash built a solid company and had always provided what he thought were adequate reserves against winter's "normal" variations. But he had not anticipated the revenue shortfalls of the last three winters. The most recent and infamous winter of 1998-99, the warmest in the history of recorded temperature, emptied the company's coffers. Another warm winter would be ruinous.

Flash runs the company from its headquarters in the northern town of Frozen Falls. The town is about 60 miles (100 kilometers) west of Frozen Lake airport, the nearest government weather measurement site. The company's rural service region is large, sits on the American-Canadian border and includes both the town and the airport.

After last winter's losses, Flash vowed not to gamble on the weather again. Over the last half-year he went to conferences and seminars that focused on the emerging weather risk market. He learned how to hedge the company's natural exposure to the weather. Looking beyond the market makers' offers of "Have I got a deal for you", he recognizes the opportunities the market provides for weather risk management at his company. He sees that the market has a strong quantitative base where all discussions of prices begin, and that the advantage in the market seems to go to the meteorologically informed and quantitatively skilled participants.

Flash recognizes that he needs help, and engages a consultant, "Storm" Wharning. Storm is a meteorologist experienced in capital markets who specializes in weather and weather derivatives. He can advise Flash on the prospects for the coming winter, and he says he can also evaluate weather financial options.

As his first task, Storm derives the time series of the 53 winters measured at Frozen Lake Airport from data he downloads from the National Climatic Data Center website (figure 1 below). Based on the data, he makes some obvious statements:

- Frozen Lake is a cold place. Average monthly temperatures in winter can fall to 14°F (-5°C), with frequent temperatures below zero.

- There is almost perfect correspondence between average winter temperature and November through March heating degree-days (correlation coefficient is 0.98).

- There is no clear and convincing trend in these 53 winters, and

- Three warmer-than-average winters in a row have happened twice before in this last half-century, but four in a row have not.

Figure 1

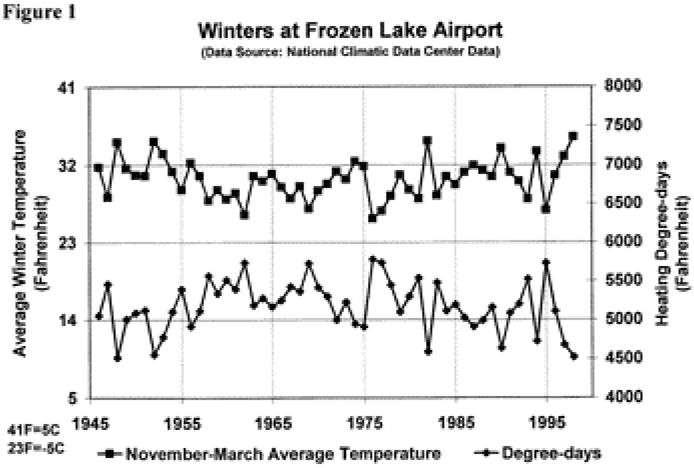

Flash is not very impressed with the significance of Storm's first effort. It told him much of what he already knew. He admits, however, that it did put numbers on the patterns of the last half-century, where he previously had only anecdotal experience.

He asks Storm "Will it be warm again this winter and what should risk management cost?" Storm equivocates. The official National Weather Service outlook for next winter, made just this month, six months in advance of winter, is for yet another warmer-than-normal season. This causes panic in the company headquarters. Not again!

Storm says that given the information in hand, the official forecast is the best forecast possible. He cannot do any better at this time and agrees with the forecast. Nonetheless, he tries to comfort Flash, reminding him that forecasts often change as more information becomes available, and that forecasts this far in advance do not deserve the highest confidence. As the time to winter shortens, confidence in the updated forecast will increase. In other words, forecasts made six months in advance of a season are sometimes shaky. "It depends," he says, but he is unable to explain in simple language on what it depends. Flash and Storm are caught up in a whirlpool of El Niño, La Niña and Jet Streams.

They do agree that Storm will watch the flow of weather information for anything that may alter the forecast. Meanwhile, he begins the next effort for the company - the assessment of the basis risk between heating in the company's distribution region and the weather measurements at Frozen Lake Airport. He plans to use all the weather data in the region, official as well as unofficial, to quantify any lack of correlation between temperature at the airport and the regional weather. What is the regional weather that drives consumers in the serviced area to raise the settings on their thermostats, and is the airport temperature a good measure of this? He will be in touch.

Flash searches through the company records and builds a temperature based revenue model for the company. As with Storm's analysis, everyone already knew the qualitative shape of the outcome, but with Flash's business model they are again able to put numbers to it (figure 2 below). Flash has quantified the company's natural weather exposure. If they could only estimate weather probabilities, they would know the risk of the company's implicit weather bet.

Figure 2

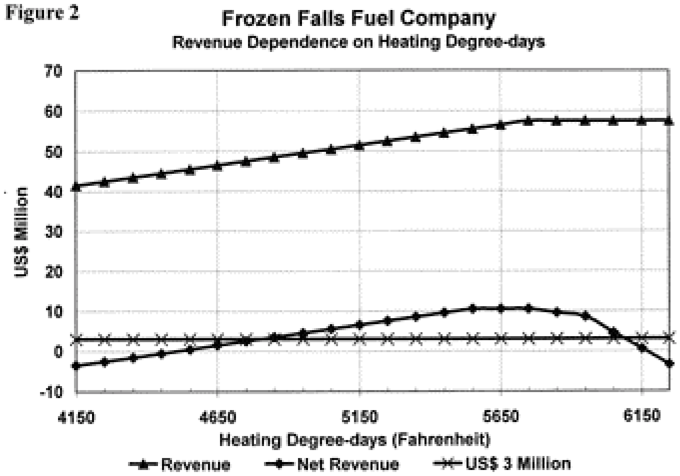

He shows in the chart above that gross revenue levels off in extremely cold winters (the upper curve). It is exactly these the times when the extremely low temperatures exceed the limitations of the consumers' systems and they cannot generate more heat. Sadly, Flash notes that it is at these times that the demand for fuel exceeds his available supplies. If this happens, he would have to purchase fuel in the open market: Purchase prices would soar as his sales price remains level. Net revenue in these extreme events could go negative (the lower curve in figure 2 above). It is a frightening prospect.

Flash discovers in his business model that the critical net revenue for the company is not zero, but USD 3 million. If revenue sinks below that level, the company's future is threatened. When he adds a level line at USD 3 million to figure 2, he finds that the critical level is crossed twice. Once if winter is so warm as to yield less than 4,750 heating degree-days, and again if it is so cold that it totals more than 6,050 heating degree-days. The company simply cannot tolerate a winter yielding less than 4,750 degree-days, or more than 6,050 degree-days.

He understands the company's implied bet on the weather: Average and near-average winters generate positive net revenue, but warmer-than-average winters and extremely colder-than-average winters result in reduced, even negative, revenue. How likely this is to occur, he cannot yet say. He makes a note to ask Storm about probabilities.

Historically, the weather bet tended to work in the company's favor, and normally Flash would let it ride. But with the company in a fragile state, he wants to take the weather bet off the table. If he does not, he is gambling more than just revenue this year. However, what he should pay for downside protection, and how much protection he needs is unclear. Yet this is no time to be incautious - he must hedge.

Now that he knows which extreme winters must be hedged, he telephones Storm and asks about the likelihood of such winters occurring. Storm uses a published temperature simulation model to estimate the probabilities of winter's severity and calculate a weather option's fair value. Storm had planned to implement the model on his laptop using instructions from a handbook published by the option model developer Robert Dischel[2]. It was not necessary to do that because the developer put the simulation model on the Internet. Storm sends a copy of the handbook to Flash by e-mail and offers to explain the model on his next visit.

When Storm next arrives he captures the attention of Flash - he has done his homework. He has built a histogram of historical winters, and has used the simulation model and daily temperatures of 30 past winters at the airport to simulate 6,000 possibilities for next winter (figure 3 below). The histogram of a few decades of history is uneven, to say the least (as depicted by the bars). Reducing it to a smooth curve would be difficult. However, the probability distribution of 6,000 simulations produces a smooth and credible alternative (the curve). At first glance, the modelled probabilities look like a normal or binomial distribution, but the discerning eye picks up the asymmetry of the distribution and the tilt to the right.

Figure 3

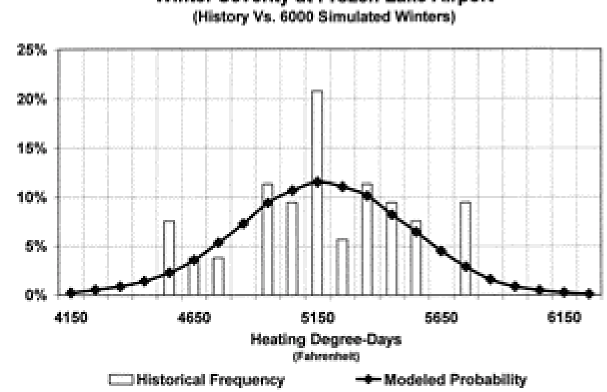

The unbiased model (without factoring in the forecast) indicates winters that are warmer than 4,750 degree-days - one of the critical values calculated by Flash - can be expected to occur about 14% of the time, or about one out of seven years. (The last two winters were warmer than 4,750.) Winters colder than 6,050 degree-days can be expected about 2% of the time. Together they make a plan. Flash will buy an out-of-the-money HDD call that will pay the company when winter is extremely cold, and they will sell a near-the-money swap that will pay if winter is warm. The company will have to pay cash for the call and give up some potential company revenue in the swap, but only if winter is cold and revenue is good.

Storm estimates that a no-cost swap (which he calls a no-cash swap) should be near the 30 year average of 5,175 degree-days. Flash estimates that he needs USD 10,000 per degree-day to stay above USD 3 million for most possible warm winters. Using both Flash's and Storm's models, they are able to look at the combination of revenues and probabilities.

Storm points out that if they sell a swap and winter is extremely cold, the call must cover both the potential revenue loss and the amount they will have to pay out if they are short a swap. To compensate for this, they hope to buy a call struck below 6,050 degree-days (perhaps at 5,850), to help offset the payout on the swap. Storm logs onto the internet and uses Robert Dischel's "D1" option pricing model for weather derivatives to estimate the fair value of a call, excluding bid-ask spread, risk premium and broker fee[3]. For a call struck at 5,850 that pays USD 20,000 per degree-day, but is capped at USD 6 million, he calculates the fair value at USD 75,000.

They construct figure 4 with the net revenue curve from figure 2 (the rightmost curve in figure 4 with values read off the left axis) and the modeled probability curve from figure 3 (the lower curve in Figure 4 with values read off the right axis). When they multiply revenue at each degree-day occurrence by the modeled probability of that occurrence, they get the probable revenue curve (the middle curve in figure 4 with values read off the right axis). They note that in both wings, warm and very cold winters, there is a non-zero probability of negative revenue.

Figure 4

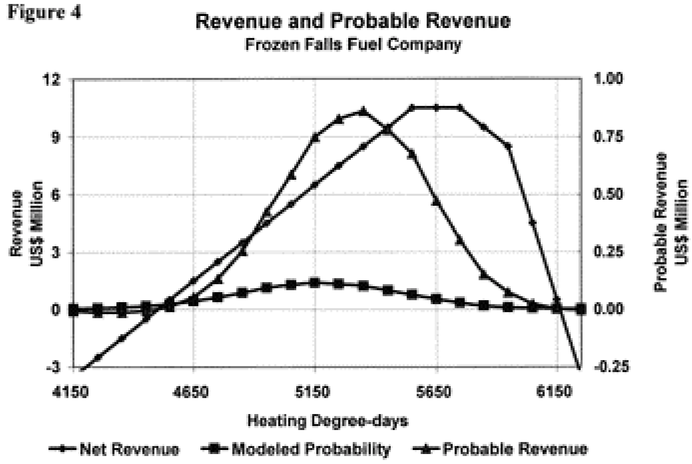

Storm explains that the integral under the probable revenue curve is the company's expected revenue of USD 6.3 million - sort of the fair value of the revenue stream. With this probable revenue curve they can calculate the dollar value of what they must give up to get the downside protection the company needs. Flash is puzzled by this assertion, but Storm says it will soon be clear enough.
By this time, it is June. They bring the call and swap proposals to a broker.

They are lucky with the option. They are able to buy the call struck at 5,850 degree-days from an insurance company through the broker. The cost is not too bad - USD 100,000. Clearly, the forecast for a warmer-than-average winter has worked to their benefit, as buyers of a call, by reducing the insurance company's demand for risk premium.

They are not, however, so lucky with the swap. They find that the market is valuing the swaps differently to them - Storm is using 30 years as a reference for his valuations, but the market makers are using 10 or 20 year averages. Storm insists that 30 years is more representative of the climate, but it does not change the market price. The forecast and the shorter averaging periods have made swaps cheap and the company is the seller. The market is buying swaps at 4,800 degree-days. Together they construct figure 5, in which the two lower curves describe the payout by degree-days of each of the two swaps: one at 4,800 and one at 5,175 degree-days. They multiply these two curves by Storm's modeled probabilities to get the two probable payout curves in the same figure.

Figure 5

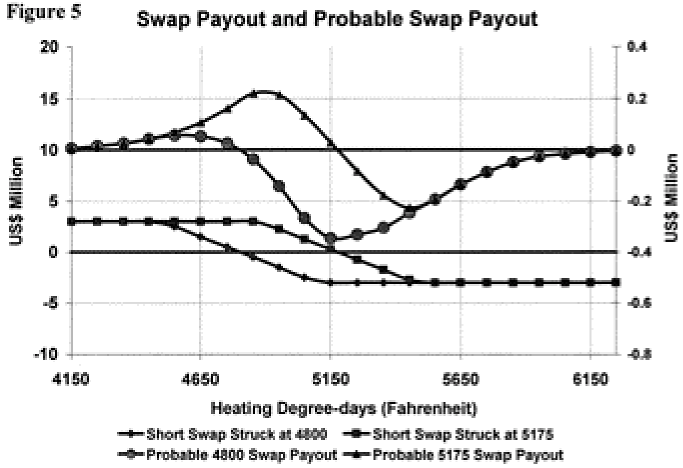

Storm explains what was unclear earlier.

For the swap struck at 5,175 degree-days (the uppermost curve in figure 5), the probable payout is read from the scale at the right, and he shows Flash that there is a cost to no-cost swaps. "For example," he explains, "A swap struck at the modeled mean would have equal areas above and below the zero-line in figure 5. The two areas will sum to zero - they will cancel each other if it is a fair-valued swap." He seems to be arguing with no one in particular when he goes on:
"These are not really no-cost swaps, but rather no-cash up-front swaps. Perhaps because this one struck at the mean of the modeled degree-days we should call it an 'equal-cost' swap, as both parties give up an equal amount of potential cash. If the two areas don't sum to zero, then someone is implicitly paying in advance in the amount of the non-zero sum." He smiles smugly as if he has just won a debate. Flash does not seem to notice.

That is, not until they look at the swap struck at 4,800 (second curve from the top in figure 5). They are shocked by the difference in areas above and below the zero-line. Flash is very unhappy about having to give up so much potential upside gain for so little downside protection. But his back is to the wall – his company's existence is at stake. Yet, paying too much for a hedge, he reasons, can be as unproductive and hazardous as not hedging at all. He reluctantly agrees to the swap, but reduces the amount to half what he would have wanted. He sells the swap at USD 5,000 per degree-day, not the USD 10,000 he estimated he needed. At least he has cut the severity of revenue shortfall by hedging.

By the time Flash and Storm finish their analyses and complete these first two transactions it is mid-June. In mid-July, Storm telephones with the good news and the bad news. The bad news is that he has been following the seasonal weather information and he sees no reason to deviate from the earlier forecast - he still expects a warmer-than-average winter. The good news is that the government forecast center running a "*coupled atmosphere-ocean model*" says it is going to get colder soon. Although he, Storm, does not agree with the new forecast, the market has reacted. Swaps just got more expensive! It might be a good time to sell.

Flash tries to decide what to do. If he believes in the forecast, then the downside protection he already has is unnecessary and good times are coming. "Yet," he grumbles to himself, "Forecasts can be wrong. After all, don't I have two opposite forecasts right now?"

Unwilling to risk the company, he makes a prudent business decision: He will double the partial hedge. He will sell another swap in the same amount as the first at USD 5,000 per degree-day, but only if the terms are more favorable than before. He calls the broker and offers the swap. After a few days, he gets the good news - he can sell the swap, but because the new forecast calls for a colder-than-normal winter, he can sell it near his originally desired level of 5,175 degree-days.

He and Storm go back to the earlier data and construct figure 6. First, they carry over the revenue curve from figure 4 as the upper curve in figure 6. Then they add to the revenue, the payout of the three hedges: the long call at 5,850, the swap at 4,800, and the available swap at 5,175. The result is the second curve from the top in figure 6, the hedged revenue. (The values for both the revenue curves are read from the left hand scale.)

Then they multiply each of these curves by the modeled probabilities resulting in the two lower curves in figure 6, the probable unhedged revenue and the probable hedged revenue. (Values for these two are read from the scale on the right.)

Figure 6

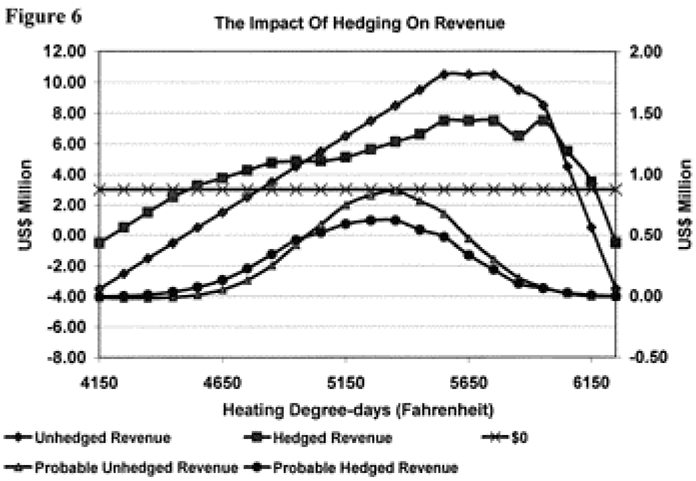

They note that the wings on both the unhedged revenue curve and the probable hedged revenue curve, where revenue would be less than USD 3 million are almost completely missing on the two hedged revenue curves.

They calculate that the probability of the company receiving revenue less than USD 3 million this coming winter is too small to hedge any further, and there will probably not be a market in the small expected value of these shortfalls. Negative revenue is possible only at the far extremes of winter at almost zero-probabilities. They have removed probable negative revenues and have raised the wings of the revenue curve by lowering its crown.

When he sees exactly what he did to protect the company - how he exchanged the upper curve in each pair in figure 6 for the lower curve in the pair, plus the USD 100,000 premium he paid for the degree-day call, Flash breaks out the champagne.

The company will survive this winter - whatever the weather.

Bob Dischel, now a consultant to the weather market, led quantitative activities at two of the US' largest insurers. There he developed measures of risk to portfolios from options and from asset / liability mismatches. Bob was a University Professor and is certified by the American Meteorological Society as a Consulting Meteorologist. He may be contacted at +1-212-799-4256, or by e-mail to bobdischel@worldnet.att.net. Bob also has a web site at www.wxpx.com.

1 - This paper was originally published by Applied Derivative Trading in their April 1999 edition. See www.adtrading.com .
2 - See Dischel, Robert, 1999, "THe D1 Stochastic Temperature Model" at www.adtrading.com. Dischel, Robert 1998a, "Option Pricing - Black-Scholes Won't Do", Weather Risk, A Special Report in the October 1998 Issues of RISK and Energy and Power Risk Management, Dischel, Robert, 1999a, "At Last: A Model for Weather Risk", March 1999 Issue of Energy and Power Risk Management.
3 - Ibid.