Forecasting: Trying to Reason from Hurricane Season

B. Jay Coleman,

Forecasting: Trying to Reason from Hurricane Season

B. Jay Coleman

Dept of Management, Marketing, & Logistics
Coggin College of Business
University of North Florida
Jacksonville
Florida 32224-2645
jcoleman@unf.edu

Abstract

Hurricane season presents an excellent illustration of a wide range of issues encountered in developing and responding to business forecasts. The high profile of hurricane forecasts – particularly during the 2004 season in which four hurricanes made landfall in Florida, and during the 2005 season when hurricane Katrina devastated New Orleans and the surrounding region – makes them a convenient, visible, and ideal illustrative classroom example. This paper outlines 25 business forecasting principles and their direct parallels in hurricane forecasting. This compilation has been successfully used by author as a basis for an in-class review of forecasting in an undergraduate operations management introductory course. A report of the positive student feedback and comparative test results is provided.

Editor's note: This is a pdf copy of an html document which resides at http://ite.pubs.informs.org/Vol6No2/Coleman/

1. Introduction

For those living and teaching in Florida and along the Gulf Coast, the 2004 and 2005 hurricane seasons were devastating ordeals. Virtually every day from mid-August through late September – well, in 2005, even until late November's Hurricane Gamma – featured incessant attention to forecast tracks, preparing for a possible hit, actual evacuation for many, and/or dealing with landfall aftermath. In short, positives were rare.

However, in looking for at least a shred of a silver lining (and shreds were about all that was left in some cases), and with apologies to Jimmy Buffett's 1974 ditty, "Trying to Reason with Hurricane Season," hurricane season presents an excellent illustration of a wide range of issues encountered in developing and responding to business forecasts. The high profile of hurricane forecasts in the media, and the magnitude of the decisions based on those forecasts, makes them a convenient and visible example to use in the classroom to illustrate a large number of forecasting principles. Given the significant familiarity with this partic-

ular brand of forecasting, particularly for students from (or attending school in) locations along the Atlantic and Gulf states, this application area is one with which a large number of students can directly identify – a statement which cannot be made for many other applications that could be used instead. This makes it an ideal foundation for a class discussion.

In the following section, 25 principles of business forecasting are outlined, along with their parallels readily seen in (and illustrated by) hurricane forecasts. In several cases, the four Florida hurricanes of 2004 – Charley, Frances, Ivan, and Jeanne – as well as 2005's Hurricane Katrina, are used as illustrative examples. These principles and their hurricane parallels have been successfully used by author as a basis for an in-class review of forecasting in an undergraduate operations management introductory course. A report of the positive student feedback and comparative test results from the exercise is provided in the final section. (For readers interested in further discussion of forecasting principles, Armstrong (2001) provides an excellent review of 139 of them.)

2. Forecasting Principles and Parallels from Hurricane Season

Lack of forecasts, lack of good data, or a lack of attention to either, can be disastrous. When a major hurricane hit Galveston, Texas in 1900, hurricane forecasting and data collection were sketchy at best. Moreover, even the short-term warnings that were issued were largely ignored, since there was little visible evidence of the approaching storm. With scant warning and a *laissez-faire* attitude, at least 6000 lives (out of 38,000 residents) were lost (Handbook of Texas Online, 2002). By contrast, due to sophisticated forecasting, data collection, and an attentive population, the four 2004 storms that hit Florida's 17 million residents took about 70 lives (Drye, 2004). Tragically however, despite "startlingly accurate" (Strohm, 2005) National Weather Service (NWS) forecasts, one of the reasons for Katrina's horrific impact and its more than 1300 deaths (Hamilton, 2006) was "complacency" among individuals and multiple levels of government regarding the NWS's "uniquely detailed and strongly worded" warnings, including the projection of "human suffering incredible by modern standards" (Whittell, 2005). The moral: investing in good forecasting – and heeding it – can be *well* worth the effort. The same is true in business. Although it pales in comparison to a loss of life discussion, cost savings and customer service improvements associated with effective forecasting practices can easily reach into millions of dollars each year in benefits for large operations.

Never assume the point estimate will be perfect. Many in the Punta Gorda area of Charley's 2004 landfall relied on projections just a few hours prior that predicted a direct hit on Tampa to the north, as a Category 2 storm. Instead, Charley quickly strengthened to a Category 4 and shifted just slightly southward. Due to the angle of the track relative to the Florida coast, this moved the landfall point considerably, and those who assumed the forecast track would be perfect paid dramatically for their error. Despite this lesson just a year earlier, Katrina's extensive effect on south Florida in 2005 was partly due to residents focusing too heavily on the exact forecast track as the storm neared the state prior to its initial U.S. landfall (Northwest Hemisphere Hurricane Center, 2005). Businesses often make the same mistake. As noted in Armstrong (2001, p. 34), most firms only use a point forecast, which – as evidenced by the hurricane stories just described – does not allow the user to assess risks. While the point estimate is useful, due diligence in business calls for things like safety stock, safety time, and extra staffing to cover the risk of demand exceeding this number.

Long-term forecasts are far more subject to error than are short-term forecasts. By now, anyone following the media coverage of projected forecast tracks is very familiar with the so-called forecast "cone," which widens considerably as the forecast extends into the future. For example, the average 24-hour error in the forecast track is about 85 miles, whereas the typical 48-hour error is 160 miles (Pain, 2005). There's simply not a better widely publicized example of the behavior of prediction error than that. Very similar error patterns exist in business forecasts, which are much more accurate for tomorrow or next week that they are for six months out.

The best forecasts are rarely linear – particularly in the longer term. Nearly all actual and forecasted tracks curved significantly at some point – and the more long-term the forecast, the more apt that a curve was included. The damped trend in product demand for periods further into the future behaves similarly.

Linear approximations can be quite good <u>in the short term</u>. When forecasting tracks a few hours out, a straight line projection based on the short-term trend was generally quite good. Many basic time series models for product demand forecasting, such as Holt's exponential smoothing model or simple regression using time as the lone predictor variable, are based on this principle, meaning that such linear models can be effective if kept to near-term forecasts.

Recent historical data are much more useful when developing forecasts than are old data. Relying on the movement of storms when they were in the middle of the Atlantic would have been virtually useless when making forecasts as the storms neared landfall. Exponential smoothing forecasting models are built on this assumption, as weights on past data drop off dramatically the older those data get.

Consider seasonality. The frequency of hurricane occurrence follows a classical normal distribution from June through November, with a peak in September. Moreover, the strongest hurricanes typically come during the Cape Verde season in August and September. Likewise, demand for most (if not nearly all) products

and services exhibits seasonal patterns of spikes and dips that can last for any length of time, including weeks or months at a stretch (Shugan and Radas, 1999).

Outliers aren't necessarily a portent of permanent shifts, so don't overly respond to them – but watch the following data points. Before 2004, four hurricanes hitting one state in one season hadn't happened since 1886 – and that was in another state (Texas). Although occasional reports have noted that some Florida residents might consider moving elsewhere as a result, doing so would likely be over-responding to an outlier, as experts suggest the 2004 season was (Drye, 2004). However, the tremendous intensity of the 2005 season has led some to conclude that 2004 may not necessarily have been an outlier year (Berger, 2005). In business, unusual circumstances will often result in similarly unusual demand, such as the demand for generators in an area directly hit by a storm. As a result, forecasts for generator sales in that area the following season might very well need to be increased over pre-strike estimates. However, they shouldn't necessarily be increased to whatever the peak demand was in the wake of the hit. Such would probably be too responsive, as a direct hit in successive years is typically unlikely.

Forecasts should be updated frequently to reflect new information, and they do often change considerably as a result. Relying on a forecast track that is just a few days old would have had Jeanne nearly missing Florida to the east, and would have had Ivan staying well to the east of the Alabama coast. Neither of these ultimately proved true. Although the 48-hour forecast for Katrina was dead-on, just 24 hours prior (i.e., 72 hours out) the forecast cone did not even include New Orleans (Perry, 2006). In dynamic business settings, such as with new products, new businesses, and in highly competitive environments where the moves of other firms have great impact, demand often changes rapidly, and thus forecasts should too.

Conditions can change dramatically very quickly, so be diligent. Jeanne went from nearly dissipating north of Haiti to a Category 3 storm before making Florida landfall. Charley's eleventh hour strengthening to major hurricane status surprised many. Early reports from New Orleans after Katrina's strike suggested that the city had survived the storm well – and then the levees broke. Demand conditions in today's e-business environment can and do change dramatically as well, given the ready availability of consumer, industry, and firm information and its speed of flow; the easy implementation of price changes and sales promotions (both for you and the competition); and the shortening of product life cycles (Gung et al., 2002).

Mathematical models can improve forecasting dramatically. Much of the relative accuracy of modern hurricane forecasting is due to the dozen or so "computer" models that are used by the National Hurricane Center, or NHC (DeMaria, 1997). Forecast tracks from five of these are commonly reported on popular weather web sites such as the Weather Underground website[1] . In the last decade, these models have improved hurricane forecasts by a full day, meaning they can predict 48 hours in advance as accurately as they predicted 24 hours in advance 10 years ago (Perry, 2006). Mathematical approaches such as time series analysis and multiple regression can also greatly improve business forecasting. As noted in Armstrong (2001, p. 12), there is ample evidence that quantitative methods are superior to qualitative (subjective judgment) approaches for business forecasting, as they are less prone to bias and are better at analyzing the available data.

Causal models that consider other factors and interactions with competitors are often better than time-series models. Many of the computer models use other factors besides the historical track of the storm (i.e., the equivalent of historical demand data). One of the main factors included is the interaction with other "competing" weather: pressure troughs or ridges, upper level winds, dry air, and even other hurricanes (e.g., Jeanne's track was influenced to some extent by the effects of Ivan just before it). Armstrong (2001, p. 12) emphasizes the value of causal models in business forecasting, particularly when there's good knowledge of what the expected causes are (e.g., advertising's impact on sales), past data is available on those causes, future values of those causes can be reasonably forecasted (see below), and the cost of obtaining the forecast is superseded by its expected benefits.

When using causal models, a significant problem is that you often must also predict the behavior of the independent variables. Much of the discussion of forecast tracks in the media focused on the speed at which a particular

---

[1] http://www.wunderground.com/

pressure ridge or trough would strengthen, dissipate, or move to a point where it would affect the hurricane. Often *these* forecasts were in error, which then significantly affected the forecasts of hurricane movement and intensity. This is also a problem with business forecasting. For example, even if a strong relationship is identified between interest rates and demand for new home construction, in order to use such a model for forecasting, future interest rates would have to be predictedbefore new home demand could be predicted.

Different math models give different forecasts, depending on the assumptions and factors considered. The forecast tracks of the five reported computer models often differed greatly for all 2004 storms; however, when they largely agreed, they were seldom in significant error. When forecasts are divergent in business settings, one should always question the assumptions made in their computation. For example, are they based on the same set of data? Is the data accurate, and has it been accurately entered? Is there information – such as an upcoming sales promotion – that was not necessarily known or accounted for by all persons or methods used in developing the forecast? Are there components in the product's historical demand that were not addressed (or addressable) by the methods employed? For example, if simple exponential smoothing was used on data containing trend and seasonality (which it's not designed to do), the resulting forecasts will likely differ significantly from models that are designed to handle both of these (such as time series decomposition or Winter's exponential smoothing).

Although math models can help considerably, the human factor should never be ignored (corollary: validate math model results). For example, one independent hurricane forecasting site[2] sometimes observed that the computer models at times would make unrealistic assumptions; e.g., that a storm would plow right through a high pressure ridge. That model's projected storm track was thus completely discounted for the current forecast. Also, the official NHC forecast track rarely (if ever) specifically matched any one particular computer model, meaning that human decision-making still came heavily into play. This is particularly impor-

tant when using time series models in business settings, as such approaches are "naive" in that they aren't designed to address *why* demand has fluctuated as it has. Employee knowledge of market conditions such as new product introductions or changes in the competition can help determine when time series forecasts should be modified or ignored.

Focus forecasting is a pretty good idea – as is considering several different viewpoints. The same web site[3] also observed that some model forecasts would work well for some storms, while others would work well with other storms. Thus, constantly relying on the same method for all forecasts wouldn't have been best. This is essentially the construct behind focus forecasting, which calls for trying a variety of methods (e.g., multiple types of moving averages) on past data, and then using the one that would have worked best recently to make the next forecast. However, a different approach might be best for another product or for the next time a forecast is needed (Krajewski and Ritzman, 2005, p. 565).

Other people rely heavily on your forecasts, so don't take them lightly. As hurricanes approach landfall, the NHC web site[4] can receive over 100 million hits *per day* (Edwards, 2003). Although numbers like that are obviously unlikely in business settings, it is rare that a forecast will be used only by its developer. For example, forecasts generated by marketing personnel often drive decisions by production planners.

Some forecasts are far more important than others. Hurricane forecasts are far more important than those for minor issues like an afternoon rain shower. Spend your resources on the more important forecasts (e.g., your "A" items) and don't sweat the smaller stuff nearly as much.

Confidence interval forecasts are better than point estimates. Without the aforementioned forecast cone – the confidence (or prediction) interval for the forecast track – interests anywhere near the single forecast track could not have gauged their level of risk, and thus could not have made the most appropriate decisions. With just a single forecast track (and no cone), most interests would not have even recognized that they had *any*

---

[2] http://www.nwhhc.com/
[3] http://www.nwhhc.com/
[4] http://www.nhc.noaa.gov/

significant level of risk. As Armstrong (2001, p. 35) points out, "prediction intervals are needed when decisions are affected by uncertainty, which means nearly always." In business settings, at a minimum such intervals give best - and worst-case scenarios for planning purposes.

The point estimate – the most likely occurrence – is generally in the middle of the confidence interval. Forecast tracks were always in the center of the forecast cone. This is also typical of their business counterparts, particularly under the common assumption of normally distributed forecast errors (see below).

The probability of the actual event is not uniform across the confidence interval. As was widely reported in discussions of the forecast cone, the probability of a direct hit from the hurricane dropped dramatically as the deviation from the single predicted track increased. In other words, the distribution of the prediction error was more similar to a normal distribution than a uniform distribution. The assumption of normally distributed forecast errors is often used when developing confidence (i.e., prediction) intervals for product demand, meaning that demand is more likely to be nearer the forecast than far away from it.

Measure historical accuracy to determine the amount of forecast error (or width of the confidence interval). As was also widely reported in the media, the width of the cone reflected how far off past actual hurricane tracks had been from what had been forecasted. This underscores the necessity to retrospectively examine past performance.

Being on one side of the actual event can be worse than being on the other side. To those in hurricane alley, it's widely known that in relation to the hurricane's forward direction, being on the "right" side of the eye is far more devastating than being on the "left" side. The same can be true of demand forecasting. For example, if shortage costs are high relative to holding costs, forecasts that are too low can be far more damaging than forecasts that are too high.

Forecasts should reflect the applicable stage of the product life cycle. The birth, growth (strengthening), maturity, and death (post-landfall) stages of a hurricane closely mimic those of the product life cycle. Forecasts of hurricane intensity directly reflected the applicable stage.

Finally, continuously scanning the horizon for threats is strongly encouraged, and don't assume a threat will simply go away. After everything Floridians and Gulf Coast residents endured in 2004 and 2005, that's simply become a way of life.

3. Use as an In-Class Review Exercise

Early in June of the summer term of 2005 – or in other words, just after the June 1 start of the 2005 hurricane season – I implemented the above list of hurricane and business forecasting parallels as the basis for a classroom review exercise in an undergraduate introductory course in operations management. The institution involved is located in Florida, and most of the 30+ students in the course had experienced the infamous 2004 hurricane season first-hand. The significant across-the-board familiarity with hurricane forecasting made the application one with which nearly all of the students involved could directly identify.

The exercise was performed as a review of demand forecasting principles, at the conclusion of the forecasting section of the course. All forecasting coverage had otherwise already taken place, including discussion of basic principles, qualitative versus quantitative methods, using causal models versus time series models, principles of time series analysis, computing forecasts using various time series models, forecast error analysis, and confidence intervals. The exercise was conducted as a simple brainstorming session with the class as a whole. Without initially revealing virtually any of the 25 principles highlighted in the previous section (only the first one was noted, to get them started), the class, working as a whole, was asked to identify as many parallels between demand forecasting and hurricane forecasting as they could. As an incentive, all class members were promised a small amount of extra credit for each parallel the class identified beyond the first five, with this credit being applied towards their short pop quiz grade, an evaluation component that comprised 5% of their total course grade.

Over a time period of approximately 25 minutes, the students identified nearly all of the 25 parallels listed in the previous section – a very surprising but certainly encouraging result. Participation was heavy and lively – far more so than normal(!) – due in part undoubtedly to the extra credit opportunity, but also likely due to the attention and importance given to the topic in their

locale. As each item was mentioned, it was recorded on the projection screen overhead for each class member to note, a brief elaboration of the point was provided, and in the process an effort was made to tie it to other points already mentioned. At the conclusion of the discussion, all those parallels previously identified and listed in the previous section were briefly reviewed.

As an assessment of the student feedback to the exercise, a simple four-question anonymous survey about the exercise was administered at the beginning of the very next class meeting. Using a 5-point scale from "strongly disagree" to "strongly agree," students were asked to assess whether the exercise (1) was a good use of class time, (2) improved their understanding of the topic of forecasting, (3) was a valuable addition to the class, over and above its value as an opportunity to earn extra credit, and (4) should be used in future terms. The results of the survey are shown in Table 1. Approximately 84% of the class agreed or strongly agreed that the exercise was a worthwhile use of time, it improved their understanding, and it was a valuable addition to the course. A slightly higher number believed that it should be used in the future.

Table 1: The results of the short survey of students regarding the in-class review exercise were quite positive (31 of the 36 students in the course responded).

Good use of class time	Improved understanding	Valuable addition	Use in future
Mean	**4.26**	**4.06**	**4.35**	**4.26**
Strongly Agree	41.9%	25.8%	54.8%	45.2%
Agree	41.9%	58.1%	29.0%	41.9%
Neutral	16.1%	12.9%	12.9%	6.5%
Disagree	0.0%	3.2%	3.2%	6.5%
Strongly Disagree	0.0%	0.0%	0.0%	0.0%

Obviously, the feedback could have been significantly skewed by the extra credit component to the exercise.

Given this consideration, and given that the question regarding improved understanding was the only one for which "strongly agree" was not the modal response, an additional assessment was made using results from the first major exam in the course. This exam took place approximately two weeks after the exercise, and included questions on both forecasting and aggregate production planning. The student performances on the forecasting questions were compared to the forecasting performances of 135 students in four sections of the same course taught by the author during the

previous three terms, in which this exercise had not been used. The average forecasting grade for the 36 students in the class that participated in the exercise was 8.78 percentage points (or nearly a letter grade) higher than the performance of their counterparts in the previous terms. (The average grade was also the highest average grade among all four terms compared.) Although this difference was significant only at a level of 0.137, it nevertheless provided some degree of evidence that the exercise was useful and cognitively productive.

References

Armstrong, J. M. (Ed.) (2001), "Standards and Practices for Forecasting", from *Principles of Forecasting*, Boston: Kluwer Academic Publishers, http://www-marketing.wharton.upenn.edu/forecast/standardshort.pdf (last accessed on January 29, 2006).

Berger, E. (2005), "The Coast is Clear (For Now)," *Houston Chronicle* (November 30), p. A1, http://www.chron.com/CDA/archives/archive.mpl?id=2005_4017825 (last accessed on February 27, 2006).

DeMaria, M. (1997), "Summary of the NHC/TPC Tropical Cyclone Track and Intensity Guidance Models," *National Oceanic and Atmospheric Administration / National Weather Service*, http://www.nhc.noaa.gov/aboutmodels.shtml (last accessed on July 7, 2005).

Drye, W. (2004), "2004 Hurricane Season May Be Costliest on Record," *National Geographic News* (September 27), http://news.nationalgeographic.com/news/2004/09/0927_040927_jeanne.html#main (last accessed on July 7, 2005).

Edwards, R. (2003), "Keeping His Eye on the Hurricane," http://www.fcw.com/fcw/articles/2003/1006/mgt-noaa-10-06-03.asp (last accessed on July 7, 2005).

Gung, R. R., Y. T. Leung, G. Y. Lin, and R. Y. Tsai (2002), "Demand Forecasting TODAY," *OR/MS Today*, Vol. 29. No. 6 (December), http://www.lionhrtpub.com/orms/orms-12-02/frdemandforecasting.html (last accessed on February 27, 2006).

Hamilton, J. (2006), "Busiest Hurricane Season Ends with New Storm," http://www.npr.org/templates/story/story.php?storyId=5031529 (last accessed on January 29, 2006).

Handbook of Texas Online (2002), "Galveston Hurricane of 1900," http://www.tsha.utexas.edu/handbook/online/articles/view/GG/ydg2.html (last accessed on July 7, 2005).

Krajewski, L. J. and L. P. Ritzman (2005), *Operations Management: Processes and Value Chains* (7th edition), Upper Saddle River, New Jersey: Pearson Education, Inc.

Northwest Hemisphere Hurricane Center (2005), "NWHHC Statement on Hurricane Katrina Destruction," http://www.nwhhc.com/Katrina.html (last accessed on January 28, 2006).

Pain, J. (2005), "Dire Katrina Predictions Were on Track," http://www.livescience.com/forcesofnature/ap_050916_hurricane_forecasting.html (last accessed on January 27, 2006).

Perry, T. S. (2006), "Hurricane Watchers Hit Their Mark," http://www.spectrum.ieee.org/print/2353 (last accessed on February 27, 2006).

Shugan, S. M. and S. Radas (1999), "Services and Seasonal Demand," in *Handbook of Services Marketing and Management,* T. A. Swartz and D. Iacobucci (Eds.), Sage Publications, pp. 147-170, http://bear.cba.ufl.edu/shugan/profile/SSD.pdf (last accessed on February 27, 2006).

Strohm, C. (2005), "Weather Service Officials Gave Dire, Accurate Warnings Before Katrina Hit," http://www.govexec.com/dailyfed/0905/092205c1.htm (last accessed on January 27, 2006).

Whittell, G. (2005), "Warnings Were Loud and Clear - But Still City Drowned," *The Times* (September 28), http://www.timesonline.co.uk/article/0,,23889-1770245,00.html (last accessed on January 27, 2006).