

Some Initiatives in a Business Forecasting Course

Singfat Chu
National University of Singapore

Journal of Statistics Education Volume 15, Number 2 (2007),
<http://www.amstat.org/publications/jse/v15n2/chu.html>

Copyright © 2007 by Singfat Chu all rights reserved. This text may be freely shared among individuals, but it may not be republished in any medium without express written consent from the author and advance notification of the editor.

Key Words: ARIMA, Logistic Regression, Pedagogy

Abstract

The paper reports some initiatives to freshen up the typical undergraduate business forecasting course. These include (1) students doing research and presentations on contemporary tools and industry practices such as neural networks and collaborative forecasting (2) insertion of Logistic Regression in the curriculum (3) productive use of applets available on the Internet to convey abstract concepts underlying ARIMA models and (4) showcasing forecasting tools in timely or familiar applications. These initiatives align with the best practices framed across the “Making Statistics More Effective in Schools of Business” (MSMESB) conferences. Course experiences and student feedback are also discussed.

1. Introduction

Surveys consistently indicate that business forecasting is mostly done in a subjective fashion, relying much on experience, intuition and field intelligence. For example, Sanders and Manrodt (2003a, 2003b, 1994) Winklhofer and Diamantopoulos (2002), Klassen and Flores (2001), Mady (2000) and Duran and Flores (1998) all report personal judgment, jury of executive opinion and sales force composite (a method of developing a sales forecast that uses how much each member of the sales force expects to sell) as the dominant forecasting practices in their surveys carried out with business executives in the USA, UK, Canada, Egypt and Mexico. In addition to convenience, these judgmental methods may be viewed as more reliable than statistical methods, which assume that the cause system underlying the process remains the same. Sanders and Manrodt (2003b) report that 61% of respondents in their survey routinely used their judgment to adjust

software-produced forecasts. However, the empirical evidence does not vouch for such practices. For example, Sanders and Manrodt (2003a, 2003b) find smaller errors on average in forecasts generated by quantitative methods compared to judgmental methods.

Recently, Lovallo and (2002 Nobel laureate in Economics) Kahneman (2003) highlight the preponderance of over-optimistic forecasts made by executives due to cognitive biases and to organizational pressures. They recommend that more realistic forecasts could and should be produced by supplementing “inside view” forecasts with “outside view” or statistical forecasts.

The marketplace demand for graduates with analytical skills provides further impetus for exposing students to the salient features of statistical forecasting models that are of pertinent interest to the business world, namely the calibration of trends, seasonal indices and extraneous or “causal” effects. Many user-friendly software packages are available nowadays to facilitate such analyses. Trade publications, such as the MS/OR magazine, regularly report the state of the art in forecasting packages e.g. Yurkiewicz (2003).

These market opportunities and software availability motivate us to equip our students, who will be future managers and entrepreneurs, with sound business forecasting skills. The purpose of this paper is to report on a dedicated business forecasting course that was added to my institution’s undergraduate curriculum in August 2002 and offered to some 60 students each semester.

The paper is organized as follows. Section 2 surveys the literature on forecasting education and the recommended pedagogy framed throughout the “Making Statistics More Effective in Schools of Business” (MSMESB) conferences. This paves the way for Section 3, which drafts the philosophy underlying the course design. Section 4 focuses on the initiatives that were implemented in the business forecasting course. The paper concludes with a discussion on the experiences of the instructor and the students in this course.

2. State of Forecasting and Statistical Education

Hanke (1984, 1989) and Hanke and Weigand (1994) surveyed the state of forecasting education in institutions accredited by the Association to Advance Collegiate Schools of Business (AACSB). The percentage of institutions with a formal forecasting course at either the undergraduate or graduate level ranged between 54% and 62% in the 3 surveys. Few changes were noted over time in a curriculum concentrating on quantitative forecasting techniques i.e. regression, smoothing, decomposition and ARIMA models (see e.g. Hays (2003) for an illustration of some of these techniques towards the forecast of computer usage in her MBA Decision Analysis course). By 1993, software such as SAS and Minitab were commonly used by students. However, based on their own experience and the survey results, Hanke and Weigand (1994) highlighted that the exposure of students to real world situations and the communication of analyses and managerial implications were deemed as lacking.

Such pedagogical shortcomings were among several which prompted the organization of several “Making Statistics More Effective in Schools of Business” (MSMESB) conferences beginning in 1986. Love and Hildebrand (2002) succinctly summarize the discussions at these conferences in a set of 10 best pedagogical practices:

1. *Practice What We Preach: Remember Our Customers.*
2. *Focus on Useful Tools.*
3. *Use Projects.*
4. *Students Need to Use Technology Well.*
5. *Lecture Less.*
6. *Focus on Statistical Thinking.*
7. *Encourage Collaboration.*
8. *Use Case Studies.*
9. *Use Real Data.*
10. *Presentation of Ideas Matter.*

3. Course Design

I paid attention to these recommendations while designing the elective course “Forecasting for Managerial Decisions” offered by the Department of Decision Sciences. First and foremost, I was aware that my “customers” would consist of mostly Supply Chain Management (SCM) majors and a few from other Business concentrations. Their future jobs required knowledge beyond the typical 1-week overview of forecasting models in their core Operations Management course. They needed exposure to a wider collection of forecasting tools. The presentation of these tools needed to be focused on their applicability rather than their theoretical or computational intricacies. In that respect, the instructor and students would use software extensively throughout the course. I opted for SPSS and MS Excel add-in StatTools .

In line with the first MSMESB recommendation “*Practice What We Preach: Remember Our Customers*”, I made it a point to focus applications, case studies and projects pertinent to the business world. For instance, I once provided a consultancy dataset (transformed for confidentiality purposes) as part of a case study to forecast the spot price of a commodity using historical shipment and inventory statistics. The efficacy of cases and projects towards learning is well documented in the literature (e.g. Love (1998, 2000) and Roback (2003)). In a business environment, they integrate:

- (a) Real data challenges e.g. filtering messy data and deciding which data to analyze to meet the objectives of a case or project,
- (b) Investigating alternative yet sensible models for separating effects from noise and
- (c) Translating statistical results into value-added actions and strategies in the form of a non-technical Memo to Management.

A thorough treatment of regression (including indicator variables, interaction, model selection, residual analysis, data transformation and interpretation of parameters) and time series models (e.g. moving average, exponential smoothing, ARIMA) constitutes the core of a forecasting course. My students had basic exposure to regression and smoothing concepts in their Statistics and Operations Management courses. Nevertheless, I started with a refresher on simple linear regression and simple exponential smoothing before extending the ideas to multiple regression and Winters' exponential smoothing. ARIMA models would be new to them. In the next section, I provide details on how I motivated ARIMA models as well as a new addition to the syllabus, namely Logistic Regression. The latter would be a valuable forecasting tool for Marketing majors as it would help them investigate factors influencing what leads to consumer choices.

Teamwork was the device used to implement the ensemble of MSMESB recommendations 3 to 10. Each semester, I split each of my 2 sections into 8 teams each comprising 3 – 4 students. Each team was responsible for (a) one topic presentation (b) 3 case studies assigned as homework and (c) a capstone project. These activities would address the business realism and communication shortcomings noted by Hanke (1994). The topic presentation was an innovation. It is described in detail in the next section.

4. Course Initiatives

This section highlights some initiatives implemented in the course. These include the expansion of standard teamwork to include the research and presentation of examinable material, the presentation of ARIMA and Logistic Regression concepts to the students and the discussion of forecasting tools using timely applications reported in the media.

4.1 Diversified Teamwork

The course promoted teamwork with 40% of the course marks allocated to team activities. Each team had to submit the same 3 case studies assigned as homework, a capstone project and a topic presentation. The case studies, which differed each semester, highlighted the 3 pillars of the course, namely regression, time series models and Logistic regression. For regression, I favored a Customer Relationship Management (CRM) angle. An example (Albright et al. (2002), pg 107) involves the Hytex Company, which has to decide how many catalogs to send to different customer segments. An interesting case study requiring time-series tools (to compute daily seasonalities) is Marriott Rooms Forecasting (do a search on "Marriott" at <http://www.dardenbusinesspublishing.com>) where the hotel chain has to decide whether to accept a booking order for a block of its rooms. I posted the best-case submissions and my general comments on the Forum page of the course website to promote discussion and collaborative learning among the students.

The final 2 class meetings were dedicated to the presentation of the capstone projects. Each team was allocated 40 minutes for the presentation followed by a Q&A session. In each meeting, 3 teams were constrained to highlight a different pillar of the course

namely, regression, time series and logistic regression models while the fourth and last team was free to showcase any forecasting tool. Teams selected their areas on a first-come-first-served basis 2 weeks before presentation. Constraining the projects to the 3 pillars of the course allowed the students to refresh ideas, clear out doubts in the Q&A sessions and thereby prepare effectively for the final exam. Each team either collected primary data or downloaded data from the Internet, organized and analyzed the business related data, interpreted the findings and submitted recommendations to Management in simple, non-technical language.

Examples of project presentations include regression models to explain salaries in the Logistics industry, the prices of hotel rooms and airplane e-tickets; time series models to investigate trends and seasonalities in container throughput at the world's busiest port (Singapore), water consumption, tourist arrivals, and postal deliveries (interestingly, the students have inferred that email has hardly dented the volume of snail mail); logistic regression models to explain the adoption of a service provider (e.g. credit card) or the choice of a particular brand (e.g. cellular phone, beverage etc.). The quality of the projects has risen from semester to semester as teams tried to do better than past projects that I posted on the course website.

The novelty in the team activities resided in the topic presentation. For this, each team searched for materials over the Internet and pertinent journals available in our digital library (Harvard Business Review, Interfaces, Journal of Business Forecasting, MS/OR etc.) and presented them for about 30 minutes at the beginning of each class. I chose the topics, allocated them randomly to the teams on Day 1 of the course and told them to discover and present the pertinent information without any assistance on my part. The topics covered subjective and quantitative forecasting tools and the latest industry practices e.g. Delphi Forecasting Method, Bass Diffusion of Innovations Model, Technical Analysis, Neural Networks, Collaborative Forecasting, etc.

This active learning exercise generated 3 benefits: sharing in class leadership, students being held accountable for their learning and diversity in the activities during the weekly 3 hour long class. The teams knew that their presentations had to be original, as I had archived past presentations on my computer. Also, I declared the topics as examinable on the final exam. Accordingly, the students researched their topics thoroughly. The non-presenting students also had an incentive to ask questions to clear their doubts. Teams instinctively prepare well when their grades depend substantially (I awarded 30% of the grade to Q&A performance) on how they answer a holistic range of questions from their fellow peers and the instructor.

I insisted that all the team members share equally in the presentation with each speaking for about 5 minutes. This was ample time to gauge the level of participation of each speaker in the preparation and coordination of the presentation. In the better-organized teams, I noticed that the members tended to anticipate questions and would handle them collectively rather than on an individual basis.

As I had 2 sections and hence 2 presentations on the same topic, I posted both on the course webpage to further promote collaborative learning across the sections. Personally, the presentations allowed me to infer how students wanted materials to be presented to them. For example, I noticed that teams put in substantial effort to find local, timely or familiar applications to motivate the utility of a forecasting tool prior to detailing its concepts. They rehearsed their presentations and managed to fit them into the allocated time, something that lecturers do not always succeed at.

4.2 Making ARIMA models less daunting

To motivate new ideas especially those that appear daunting, it helps to start with insightful applications or analogies. In 2002-3, I exploited timely media reports on the usage of mass spectrometry in the investigation of suspected anthrax-tainted letters in the USA (e.g. <http://www.chemguide.co.uk/analysis/masspecmenu.html>) to motivate a connection between chemical fingerprinting and ARMA signature ACFs and PACFs. I used applets available at <http://www.aranya.com/resources/java/> to simulate these signatures. Students clearly preferred this hands-on and dynamic learning experience to static ACF and PACF signatures illustrated in textbooks. The partial award of the 2003 Nobel Prize in Economics to Clive Granger for his work in co-integration provided an exciting background for driving the concept of differencing to stationarity (see <http://nobelprize.org/economics/laureates/2003/public.html>)

Demonstrating that models already discussed in the course were actually special cases of the ARIMA model also helped in its motivation. For instance, Hanke and Wichern (2005, p 431, Q7) supply a dataset and a guided exercise to help the students infer that a linear trend model is actually a special case of the ARIMA (0,1,1) model. This follows from,

$$y_t = a + bt + e_t \Rightarrow y_t - y_{t-1} = b + e_t - e_{t-1}$$

Hanke and Wichern (2005, p 424) further demonstrate that another special case of the ARIMA (0,1,1) model is the simple exponential smoothing model. Students may be further tasked to establish how the quadratic trend model or Winters' Exponential Smoothing Model relates to higher order ARIMA models.

Going through the above helped the introduction of a challenging forecasting tool and made the students realize the utility of ARIMA models to model situations beyond first order serial dependence. I was particularly delighted when two teams confidently and satisfactorily used ARIMA models for their capstone project in the first semester of academic year 2003-4.

4.3 Adding Logistic Regression to the Syllabus

I included binary logistic regression in the course. This forecasting tool has wide applicability in business, for example, explaining factors underlying consumer choices and the success of ventures and products. As Logistic Regression is not covered in my adopted (or most) forecasting texts, I provided notes and invited the students to search for

additional exposition either in texts available in our library or in lecture materials (from other universities) uploaded on the Internet.

I intentionally carried out a linear regression on binary response data to discuss its shortcomings and thereby motivate the need for an S-curve logistic model. I mentioned the concepts underlying maximum likelihood estimation without getting into the details. More importantly, I stressed the interpretation of odds (i.e. ratio of success to failure probabilities) and log odds output by computer software. Students were reminded to select their explanatory variables carefully because logistic regression models could also be afflicted by such problems as multicollinearity, and serial correlation.

4.4 Impact Applications and Datasets

Students were exposed to the challenges of real data in their case studies and project. I also made it a point to illustrate applications that were familiar to them. GDP data were used to illustrate the concepts of Trend, Seasonality, Cycles and Errors underlying Winters' Exponential Smoothing Model, the Time Series Decomposition Model and the seasonal ARIMA model. Financial data extracted from annual reports were used to illustrate linear and exponential growth. The relationship between price and supply was demonstrated using timely auction data on the Certificate of Entitlement (COE), a vehicle ownership management scheme implemented in space-challenged Singapore.

I found many timely forecasting applications in the media and on the Internet. For example, my students appreciated newspaper reports on the purchase of weather forecasts towards such activities as inventory planning at Wal-Mart and the trading of commodity futures. Prior to the Athens 2004 Olympics, several forecasts of medal tallies appeared on the Internet (e.g. search "Regression Olympic Medal" on Google). The "fun" regression models posted there provided interesting material for class discussion. Another prediction that interested my students concerned the [Nov 2004 US presidential elections](#).

When the SARS (Severe Acute Respiratory Syndrome) pandemic affected many Asian countries in 2003, the developers of a credit-card size frontal digital thermometer which could flash frontal temperatures within 15 seconds (compared to 1 to 2 minutes for analogue thermometers) used linear regression to benchmark its measurements against the standards which were based on oral temperatures.

Consultancy reports on [3G mobile phone adoption](#) and [real estate valuation](#) offered the students interesting windows into professional forecasting practices.

5. Take-Aways

My initiatives in the course consisted mainly of making the students more involved in their learning. They experienced the utility of forecasting tools in interesting applications and case studies. I used current events as background to communicate concepts. For specialized courses like Forecasting and possibly Quality Management, my experience

indicates that a diversity of learning activities comprising analytics, case studies, projects and student research and presentation of the latest industry practices and pertinent qualitative topics benefits all stakeholders: students, faculty and marketplace.

My relatively small sections facilitated the pedagogy. Larger classes present harder-to-resolve challenges with ever limited time, people and resources. For instance, students expect timely feedback, say 1 week after case submission. As the teams approached the problems from different angles or used different portions of the dataset, I had to understand their rationale, provide constructive feedback on where they went right and wrong and benchmark their efforts and their results transparently relative to their peers.

Since its first offering in 2002, the course has attracted an average rating between 4.2 and 4.4 out of 5 on the students' evaluation of 7 equally weighted aspects of pedagogy. The ratings pertain to the enhancement in their thinking skills, the timeliness and usefulness of feedback, the approachability of the instructor, the impetus on their research, the increase in subject interest, the degree of business realism and the integration of the subject within the entire curriculum. Tellingly, the average course ratings have been about 10%-15% higher than the averages for the other business courses offered at the same level. Furthermore, 40% or more of the students nominated the course for innovative teaching each semester. These contributed in part to the 2004 Faculty Outstanding Educator Award to the author.

Appendix

In this section, I describe a dataset that I have used to implement some of the ideas described above. The dataset pertains to the Certificate of Entitlement (COE) scheme, which is an auction mechanism used uniquely in land-scarce Singapore to manage vehicular ownership. My students have always been interested in this dataset as (1) the auction results are widely publicized and (2) I update the data after each auction.

The COE scheme started in August 1990. It uses a uniform price mechanism to allocate rights to put new vehicles on the road. It works as follows: suppose the policy makers decide that 1000 new vehicles can be put on the roads each month. The public is invited to bid for these rights. At the end of the auction, the 1000th highest bid is deemed as the clearing price. This means that the top 1000 bidders will be allocated a right to vehicular ownership and they will all pay a uniform price equal to the last (or 1000th) successful bid.

From August 1990 to June 2001, the COE auction was held once a month and it employed a closed-bid format (i.e. bids were secret and only disclosed at the close of the auction). Between July 2001 and March 2002, the auction was held twice a month alternating between closed-bid and open-bid formats. Under the open-bid format, a bidder could check the current clearing price and change his/her own bid on a dedicated website during the auction window which spanned 3 days. From April 2003 onwards, the biweekly auction has followed the open-bid format exclusively.

In the dataset (COE.xls), I provide data on the current and lagged clearing prices, the number of bidders and the quota of rights for each auction. I also provide a dummy variable to indicate closed or open bid formats. The dataset can be employed to illustrate simple as well as more intricate forecasting models. For example, students can consider the suitability of the Winters Exponential Smoothing model or the ARIMA model for the exclusively closed or open bid sub-periods (i.e. August 1990-June 2001 or from April 2003 onwards).

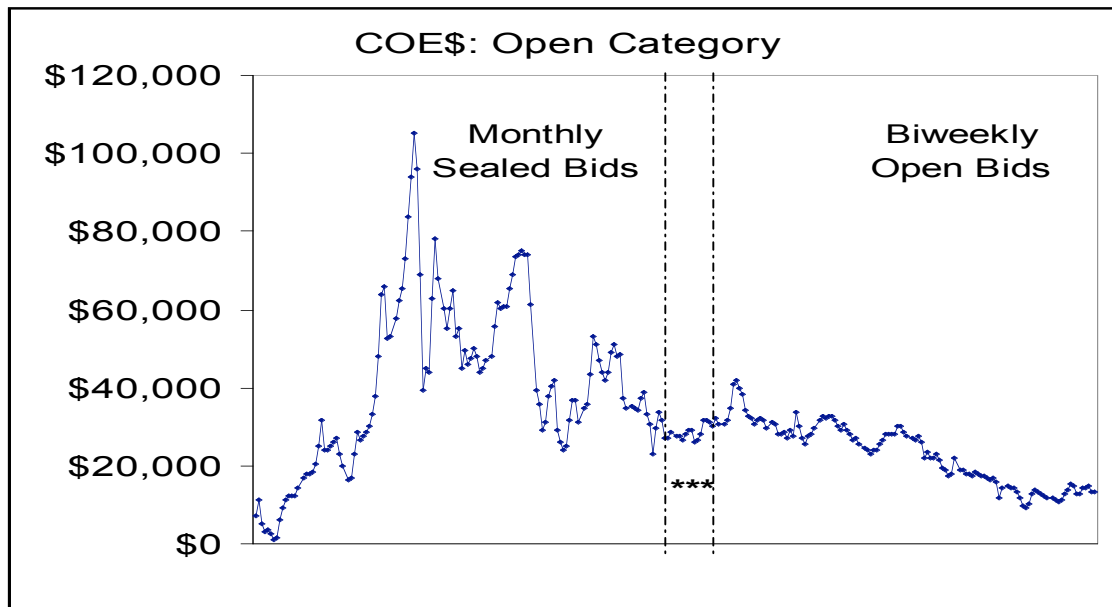


Figure 1: COE clearing prices from August 1990 to Jan 2007

***: June 2001 to Mar 2002: Alternating biweekly closed and open bid formats

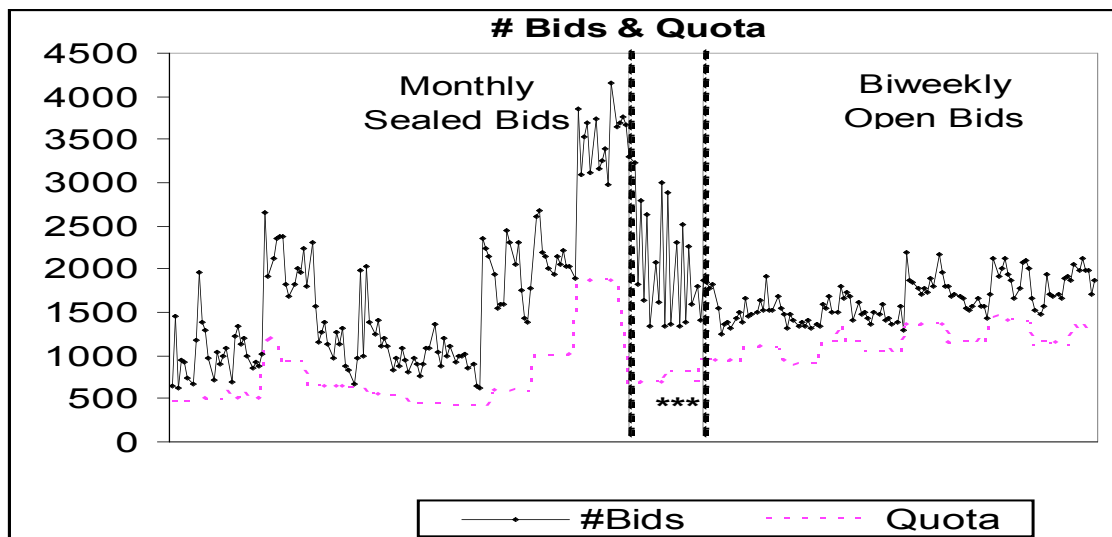


Figure 2: # Bids and Quota for each COE auction

Analyzing the entire time series is more intricate due to the change in auction format from closed to open bid. A suitable approach is to employ a regression model with explanatory variables such as preceding price, number of bidders, quota available, an auction format dummy and interaction terms to infer changes in bidding behaviour (i.e. are the regression coefficients for number of bidders and quota the same for closed and open bid auction formats?).

Residual analysis reveals 2 problems namely, (1) fanning-out and (2) first-order autocorrelation. Students will have to be guided through remedial actions such as log-transformations and the Cochran-Orcutt model. Interestingly, the conclusion from this analysis is that the change in auction format from closed to open bid did not affect the sensitivity of the COE price to either the preceding price, the number of bids and the quota. These findings provide rich discussion materials for the students as they integrate for example concepts such as price, supply and demand and differences in auction mechanism they may have come across in Economics or Strategy courses.

REFERENCES

Albright, S. C., Winston, W. L. and Zappe, C. J. (2002), "Managerial Statistics", Duxbury, California, USA.

Duran, J. A., and Flores, B. E. (1998), "Forecasting Practices in Mexican Companies," *Interfaces*, 28, 56-62.

Hanke, J. E. (1984), "Forecasting in Business Schools: a Survey," *Journal of Forecasting*, 3, 229-234.

Hanke, J. E. (1989), "Forecasting in Business Schools: a Follow-Up Survey," *International Journal of Forecasting*, 5, 259-262.

Hanke, J. E., and Weigand, P. (1994), "What are Business Schools Doing to Educate forecasters," *Journal of Business Forecasting Methods & Systems*, 13(3), 10-12.

Hanke, J. E., and Wichern, D. W. (2005), "Business Forecasting", 8th edition, Pearson Prentice-Hall, New Jersey

Hays, J. M. (2003), "Forecasting Computer Usage," *Journal of Statistical Education* [Online], 11(1).

www.amstat.org/publications/jse/v11n1/datasets.hays.html

Klassen, R. D., and Flores, B. E. (2001), "Forecasting Practices of Canadian Firms: Survey Results and Comparisons," *International Journal of Production Economics*, 70, 163-174.

Lovall, D., and Kahneman, D. (2003), "Delusions of Success: How Optimism Undermines Executives' Decisions," *Harvard Business Review*, 81(7), 56-63.

Love, T. E. (1998), "A Project-Driven Second Course," *Journal of Statistics Education* [Online], 6(1).
www.amstat.org/publications/jse/v6n1/love.html

Love, T. E. (2000), "A Different Approach to Project Assessment," *Journal of Statistics Education* [Online], 8(1).
www.amstat.org/publications/jse/secure/v8n1/love.cfm

Love, T. E., and Hildebrand, D. K. (2002), "Statistics Education and the Making of Statistics More Effective in Schools of Business Conferences", *The American Statistician*, 56(2), 107-112.

Mady, M. T. (2000), "Sales Forecasting Practices of Egyptian Public Enterprises: Survey Evidence," *International Journal of Forecasting*, 16, 359-368.

Roback, P. J. (2003), "Teaching an Advanced Methods Course to a Mixed Audience," *Journal of Statistics Education* [Online], 11(2).
<http://www.amstat.org/publications/jse/v11n2/roback.html>

Sanders, N. R., and Manrodt, K. B. (2003a), "The Efficacy of using Judgmental versus Quantitative Forecasting Methods in Practice Forecasting," *OMEGA*, 31, 511-522.

Sanders, N. R., and Manrodt, K. B. (2003b), "Forecasting Software in Practice: Use, Satisfaction and Performance," *Interfaces*, 33 (5), 90-93.

Sanders, N. R., and Manrodt, K. B. (1994), "Forecasting Practices in United States Corporations – Survey Results," *Interfaces*, 24, 92-100.

Winklhofer, H., and Diamantopoulos, A. (2002), "A Comparison of Export Sales Forecasting Practices among UK firms," *Industrial Marketing Management*, 31, 479-490.

Yurkiewicz, J. (2003) "Forecasting Software Survey: Predicting Which Product is Right for You," *MS/OR Today*, <http://www.lionhrtpub.com/orms/orms-2-03/frsurvey.html>

Singfat Chu
NUS Business School
National University of Singapore
bizchucl@nus.edu.sg
