The UMAP Journal

Vol. 23, No. 3

Subscription Rates for 2002 Calendar Year: Volume 23

MEMBERSHIP PLUS FOR INDIVIDUAL SUBSCRIBERS

Individuals subscribe to *The UMAP Journal* through COMAP's Membership Plus. This subscription includes print copies of quarterly issues of *The UMAP Journal*, our annual collection *UMAP Modules: Tools for Teaching*, our organizational newsletter *Consortium*, on-line membership that allows members to search our on-line catalog, download COMAP print materials, and reproduce for use in their classes, and a 10% discount on all COMAP materials.

(Domestic)	**#2220**	**$75**
(Outside U.S.)	**#2221**	**$85**

INSTITUTIONAL PLUS MEMBERSHIP SUBSCRIBERS

Institutions can subscribe to the *Journal* through either Institutional Pus Membership, Regular Institutional Membership, or a Library Subscription. Institutional Plus Members receive two print copies of each of the quarterly issues of *The UMAP Journal*, our annual collection *UMAP Modules: Tools for Teaching*, our organizational newsletter *Consortium*, on-line membership that allows members to search our on-line catalog, download COMAP print materials, and reproduce for use in **any class taught in the institution, and** a 10% discount on all COMAP materials.

(Domestic)	**#2270**	**$395**
(Outside U.S.)	**#2271**	**$415**

INSTITUTIONAL MEMBERSHIP SUBSCRIBERS

Regular Institutional members receive only print copies of *The UMAP Journal*, our annual collection *UMAP Modules: Tools for Teaching*, our organizational newsletter *Consortium*, **and** a 10% discount on all COMAP materials.

(Domestic)	**#2240**	**$165**
(Outside U.S.)	**#2241**	**$185**

LIBRARY SUBSCRIPTIONS

The Library Subscription includes quarterly issues of *The UMAP Journal* and our annual collection *UMAP Modules: Tools for Teaching* and our organizational newsletter *Consortium*.

(Domestic)	**#2230**	**$140**
(Outside U.S.)	**#2231**	**$160**

To order, send a check or money order to COMAP, or call toll-free
1-800-77-COMAP (1-800-772-6627).

The UMAP Journal is published quarterly by the Consortium for Mathematics and Its Applications (COMAP), Inc., Suite 210, 57 Bedford Street, Lexington, MA, 02420, in cooperation with the American Mathematical Association of Two-Year Colleges (AMATYC), the Mathematical Association of America (MAA), the National Council of Teachers of Mathematics (NCTM), the American Statistical Association (ASA), the Society for Industrial and Applied Mathematics (SIAM), and The Institute for Operations Research and the Management Sciences (INFORMS). The Journal acquaints readers with a wide variety of professional applications of the mathematical sciences and provides a forum for the discussion of new directions in mathematical education (ISSN 0197-3622).

Second-class postage paid at Boston, MA
and at additional mailing offices.
Send address changes to:
The UMAP Journal
COMAP, Inc.
57 Bedford Street, Suite 210, Lexington, MA 02420

Vol. 23, No. 3 2002

Table of Contents

COMAP

Publisher's Editorial

Roots

Solomon A. Garfunkel
Executive Director
COMAP, Inc.
57 Bedford St., Suite 210
Lexington, MA 02420
s.garfunkel@mail.comap.com

In the beginning... Well, to be honest, I no longer remember the beginning very well. But when COMAP began in late 1980, we had a rather simple mission. We believed, then as now, in teaching mathematics through modeling and applications. We believed, and continue to believe, that students and teachers need to have persuasive answers to "What will we ever use this for?"

We set out to create a body of curriculum materials—in instructional modules, journals, newsletters and texts; in print, video, and electronic formats—that embodied this approach, presenting mathematics through its contemporary applications. And, over the years, we have worked at every educational level, from elementary school to graduate school, because we believe that students need to see mathematics as part of their personal experience from their first encounters with number and pattern.

But in the beginning... In the beginning, we worked at the undergraduate level. We—and here I include the first authors, field-testers, and users—were college mathematics faculty. Many of us received our degrees in the late 1960s—and we were legion. Remember, universities were granting approximately 1,400 doctorates per year in the mathematical sciences then (1,100 today), and over 90% were U.S. citizens (about 50% today)—all looking for jobs at the same 20 universities at the same time. Not surprisingly, we were flung far and wide; and many found themselves at colleges without a legacy of mathematical research and with consequently high teaching loads.

But we had the energy of youth and the idealism of the late 1960s. We would change the way mathematics was taught where we worked in college mathematics departments. And with a great deal of simplification, that is how the UMAP Module series and this *Journal* were born.

As I'm sure you know, in the last several years we have had a number of projects that focused on the high school level; readers of this *Journal* certainly

The UMAP Journal 23 (3) (2002) 185–186.

don't need a sermon on the importance of K–12 mathematics instruction. And, importantly for COMAP, we have greatly expanded our experiences, and I hope grown with these efforts. But in this editorial, I want to describe a new project we are undertaking that takes us back to our undergraduate roots.

`Mathmodels.org`—the Mathematical Modeling Forum, or just the Modeling Forum—is a new program recently funded by the National Science Foundation. Its purpose is to help teachers and students learn and participate in the modeling process. One of COMAP's most successful endeavors has been the Mathematical Contest in Modeling (MCM). But one failing of the contest format is that we cannot provide students with feedback on their papers. Were they on the right track? Did they not take into account a crucial variable? Also, contests are designed to select winning papers, not to promote cooperation between teams. `Mathmodels.org` is designed to do better.

The idea is simple. Have a Web site where modeling problems are posted and where students, individually and in teams, can present whole or partial solutions. The student work will be monitored by experienced faculty, who will give feedback at regular intervals as well as encourage cooperation with other students/teams. Students and faculty will communicate through threaded discussions. Prof. Pat Driscoll (U.S. Military Academy) is the project director, and I'll do some work on this as well. We see this project as a natural extension of MCM—the forum can clearly help students prepare for the contest. But more importantly, it can give them rich experiences with a wide variety of modeling problems without time constraints. For the faculty, the forum will provide an important source of problems and examples. We hope that you will join us as site mentors, problem sources, or just enthusiastic members of our modeling forum. We are very excited about the opportunity that this grant will afford us to strengthen our ability to meet our core mission as we strengthen our roots in the undergraduate mathematics community.

And speaking of the mathematics community, how many of each year's 1,100 new Ph.D.'s in mathematics know about COMAP or about this ?

College mathematics faculty are the main readers and subscribers to this *Journal*, which was founded for them and for their students. But COMAP cannot be sustained into the future by the *same* college mathematics faculty who contributed ideas, articles, and Modules 15, 20, or 25 years ago. We and they will pass; before then, however, we must pass the torch.

Hence we ask you to

- show this and other issues of this *Journal* and other COMAP materials to young members in your department,
- suggest that they consider joining COMAP (information is on the back side of the title page of this issue), and
- urge them to contact myself, editor Paul Campbell, any of the associate editors of this *Journal*, or any of the coaches of MCM teams listed in this issue about how to become active as an author, a reviewer of manuscripts, a book reviewer, or coach of an MCM or ICM team.

Modeling Forum

Results of the 2002 Mathematical Contest in Modeling

Frank Giordano, MCM Director
COMAP, Inc.
57 Bedford St., Suite 210
Lexington, MA 02420
f.giordano@mail.comap.com

Introduction

A total of 525 teams of undergraduates, from 282 institutions in 11 countries, spent the second weekend in February working on applied mathematics problems in the 18th Mathematical Contest in Modeling (MCM).

The 2002 MCM began at 8:00 P.M. EST on Thursday, Feb. 7 and officially ended at 8:00 P.M. EST on Monday, Feb. 11. During that time, teams of up to three undergraduates were to research and submit an optimal solution for one of two open-ended modeling problems. Students registered, obtained contest materials, downloaded the problems at the appropriate time, and entered completion data through COMAP'S MCM Web site.

Each team had to choose one of the two contest problems. After a weekend of hard work, solution papers were sent to COMAP on Monday. Ten of the top papers appear in this issue of *The UMAP Journal*.

Results and winning papers from the first sixteen contests were published in special issues of *Mathematical Modeling* (1985–1987) and *The UMAP Journal* (1985–2001). The 1994 volume of *Tools for Teaching*, commemorating the tenth anniversary of the contest, contains all of the 20 problems used in the first ten years of the contest and a winning paper for each. Limited quantities of that volume and of the special MCM issues of the *Journal* for the last few years are available from COMAP.

This year's Problem A was about controlling the amount of spray hitting passersby that is produced by wind acting on an ornamental fountain located in the midst of a plaza surrounded by buildings. The water flow is controlled by

The UMAP Journal 23 (3) (2002) 187–208.

a mechanism linked to an anemometer located on top of an adjacent building. Students were asked to design a control algorithm that would provide a balance between an attractive spectacle and a soaking.

Problem B focused on the challenge associated with airline practices overbooking of flight reservations. Students were asked to determine an optimal overbooking strategy in light of operational constraints evolving from the events of September 11, 2001.

In additional to the MCM, COMAP also sponsors the Interdisciplinary Contest in Modeling (ICM) and the High School Mathematical Contest in Modeling (HiMCM). The ICM, which runs concurrently with MCM, offers a modeling problem involving concepts in mathematics, environmental science, environmental engineering, and/or resource management. Results of this year's ICM are on the COMAP Web site at `http://www.comap.com/undergraduate/contests`; results and Outstanding papers appeared in Vol. 23 (2002), No. 1. The HiMCM offers high school students a modeling opportunity similar to the MCM. Further details about the HiMCM are at `http://www.comap.com/highschool/contests`.

Problem A: Wind and Waterspray

An ornamental fountain in a large open plaza surrounded by buildings squirts water high into the air. On gusty days, the wind blows spray from the fountain onto passersby. The water flow from the fountain is controlled by a mechanism linked to an anemometer (which measures wind speed and direction) located on top of an adjacent building. The objective of this control is to provide passersby with an acceptable balance between an attractive spectacle and a soaking: The harder the wind blows, the lower the water volume and the height to which the water is squirted, hence the less spray falls outside the pool area.

Your task is to devise an algorithm that uses data provided by the anemometer to adjust the water-flow from the fountain as the wind conditions change.

Problem B: Airline Overbooking

> *You're all packed and ready to go on a trip to visit your best friend in New York City. After you check in at the ticket counter, the airline clerk announces that your flight has been overbooked. Passengers need to check in immediately to determine if they still have a seat.*

Historically, airlines know that only a certain percentage of passengers who have made reservations on a particular flight will actually take that flight. Consequently, most airlines overbook—that is, they take more reservations than the capacity of the aircraft. Occasionally, more passengers will want to take a flight than the capacity of the plane, leading to one or more passengers being bumped and thus unable to take the flight for which they had reservations.

Airlines deal with bumped passengers in various ways. Some are given nothing, some are booked on later flights on other airlines, and some are given some kind of cash or airline ticket incentive.

Consider the overbooking issue in light of the current situation:

- fewer flights by airlines from point A to point B;
- heightened security at and around airports,
- passengers' fear, and
- loss of billions of dollars in revenue by airlines to date.

Build a mathematical model that examines the effects that different overbooking schemes have on the revenue received by an airline company, in order to find an optimal overbooking strategy—that is, the number of people by which an airline should overbook a particular flight so that the company's revenue is maximized. Ensure that your model reflects the issues above and consider alternatives for handling "bumped" passengers. Additionally, write a short memorandum to the airline's CEO summarizing your findings and analysis.

The Results

The solution papers were coded at COMAP headquarters so that names and affiliations of the authors would be unknown to the judges. Each paper was then read preliminarily by two "triage" judges at Southern Connecticut State University (Problem A) or at the U.S. Military Academy (Problem B). At the triage stage, the summary and overall organization are the basis for judging a paper. If the judges' scores diverged for a paper, the judges conferred; if they still did not agree on a score, a third judge evaluated the paper.

Final judging took place at Harvey Mudd College, Claremont, California. The judges classified the papers as follows:

	Outstanding	Meritorious	Honorable Mention	Successful Participation	Total
Wind and Waterspray	4	48	60	167	279
Airline Overbooking	6	38	61	138	246
	10	86	121	305	525

The ten papers that the judges designated as Outstanding appear in this special issue of *The UMAP Journal*, together with commentaries. We list those teams and the Meritorious teams (and advisors) below; the list of all participating schools, advisors, and results is in the **Appendix**.

Outstanding Teams

Institution and Advisor	Team Members
Wind and Waterspray Papers	
"Simulating a Fountain"	
Maggie L. Walker Governor's School	Lyric P. Doshi
Richmond, VA	Joseph E. Gonzalez
John A. Barnes	Philip B. Kidd
"The Fountain That Math Built"	
North Carolina School of Science	
and Mathematics	Alex McCauley
Durham, NC	Josh Michener
Daniel J. Teague	Jadrian Miles
"Wind and Waterspray"	
U.S. Military Academy	Tate Jarrow
West Point, NY	Colin Landon
David Sanders	Mike Powell
"A Foul-Weather Fountain"	
University of Washington	Ryan K. Card
Seattle, WA	Ernie E. Esser
James Allen Morrow	Jeffrey H. Giansiracusa
Airline Overbooking Papers	
"Things That Go Bump in the Flight"	
Bethel College	Krista M. Dowdey
St. Paul, MN	Nathan M. Gossett
William M. Kinney	Mark P. Leverentz
"Optimal Overbooking"	
Duke University	David Arthur
Durham, NC	Sam Malone
David P. Kraines	Oaz Nir

"Models for Evaluating Airline Overbooking"
Harvey Mudd College
Claremont, CA
Michael E. Moody

Michael B. Schubmehl
Wesley M. Turner
Daniel M. Boylan

"Probabilistically Optimized Airline Overbooking Strategies, or 'Anyone Willing to Take a Later Flight?'"
University of Colorado at Boulder
Boulder, CO
Anne M. Dougherty

Kevin Z. Leder
Saverio E. Spagnolie
Stefan M. Wild

"ACE Is High"
Wake Forest University (Team 69)
Winston-Salem, NC
Edward E. Allen

Anthony C. Pecorella
Elizabeth A. Perez
Crystal T. Taylor

"Bumping for Dollars: The Airline Overbooking Problem"
Wake Forest University (Team 273)
Winston-Salem, NC
Frederick H. Chen

John D. Bowman
Corey R. Houmard
Adam S. Dickey

Meritorious Teams

Wind and Waterspray Papers (48 teams)
Asbury College, Wilmore, KY, USA (Kenneth P. Rietz)
Beijing Institute of Technology, Beijing, P.R. China (Yao Cui Zhen)
Beijing University of Chemical Technology, Beijing, P.R. China (Yuan WenYan)
Beijing University of Posts and Telecommunication, Beijing, P.R. China (He Zuguo) (two teams)
Beijing University of Posts and Telecommunication, Beijing, P.R. China (Sun Hongxiang)
Bethel College, St. Paul, MN (William M. Kinney)
Boston University, Boston, MA (Glen R. Hall)
California Polytechnic State University, San Luis Obispo, CA (Thomas O'Neil)
Central South University, Changsha, Hunan, P.R. China (Xuanyun Qin)
The College of Wooster, Wooster, OH (Charles R. Hampton)
East China University of Science and Technology, Shanghai, P.R. China (Lu Yuanhong)
Goshen College, Goshen, IN (David Housman)
Hangzhou University of Commerce, Hangzhou, Zhejiang, P.R. China (Zhao Heng)
Hangzhou University of Commerce, Hangzhou, Zhejiang, P.R. China (Zhu Ling)
Humboldt State University, Arcata, CA (Roland H. Lamberson)
Jacksonville University, Jacksonville, FL (Robert A. Hollister)
James Madison University, Harrisonburg, VA (Caroline Smith)

Lafayette College, Easton, PA (Thomas Hill)
Lawrence Technological University, Southfield, MI (Scott D. Schneider)
Lawrence Technological University, Southfield, MI (Howard E. Whitston)
Luther College, Decorah, IA (Reginald, D. Laursen) (two teams)
Magdalen College, Oxford, Oxfordshire, United Kingdom (Byron W. Byrne)
Massachusetts Institute of Technology, Cambridge, MA (Daniel H. Rothman)
Nankai University, Tianjin, P.R. China (Huang Wuqun)
North China Electric Power University, Baoding, Hebei, P.R. China (Gu Gendai)
Northern Jiaotong University, Beijing, P.R. China (Wang Bingtuan)
Southern Oregon University, Ashland, OR (Kemble R. Yates)
State University of West Georgia, Carrollton, GA (Scott Gordon)
Trinity University, San Antonio, TX (Jeffrey K. Lawson)
Trinity University, San Antonio, TX (Hector C. Mireles)
University College Cork, Cork, Ireland (Donal J. Hurley)
University of Colorado at Boulder, Boulder, CO (Anne M. Dougherty)
University of Colorado at Boulder, Boulder, CO (Michael Ritzwoller) (two teams)
University of Elec. and Sci. Technology, Chengdu, Sichuan, P.R. China (Qin Siyi)
University of New South Wales, Sydney, NSW, Australia (James W. Franklin)
University of North Carolina, Chapel Hill, NC (Jon W. Tolle)
University of Washington, Seattle, WA (James Allen Morrow)
Wright State University, Dayton, OH (Thomas P. Svobodny)
Xavier University, Cincinnati, Ohio (Michael Goldweber)
Youngstown State University, Youngstown, OH (Angela Spalsbury)
Zhejiang University, Hangzhou, Zhejiang, P.R. China (Yang Qifan)

Airline Overbooking Papers (38 teams)
Albertson College of Idaho, Caldwell, ID (Mike P. Hitchman)
Asbury College, Wilmore, KY (Kenneth P. Rietz)
Beijing Institute of Technology, Beijing, Beijing, P.R. China (Zhang Bao Xue)
China University of Mining and Technology, Xuzhou, Jiangsu, P.R. China (Zhu Kaiyong)
Chongqing University, Chongqing, Shapingba, P.R. China (Yang Xiaofan)
Colgate University, Hamilton, NY (Warren Weckesser)
College of Sciences of Northeastern University, Shenyang, Liaoning, P.R. China (Han Tie-min)
Fudan University, Shanghai, P.R. China (Cai Zhijie)
Gettysburg College, Gettysburg, PA (James P. Fink)
Harbin Institute of Technology, Harbin, Heilongjiang, P.R. China (Wang Xuefeng)
Harvey Mudd College, Claremont, CA (Michael E. Moody)
Harvey Mudd College, Claremont, CA (Ran Libeskind-Hadas) (two teams)
Institut Teknologi Bandung, Bandung, Jabar, Indonesia (Edy Soewono)
Juniata College, Huntingdon, PA (John F. Bukowski)
Lipscomb University, Nashville, TN (Gary Clark Hall)
Maggie L. Walker Governor's School, Richmond, VA (John A. Barnes)
Maggie L. Walker Governor's School, Richmond, VA (Crista Hamilton)
Massachusetts Institute of Technology, Cambridge, MA (Martin Zdenek Bazant)

Nankai University, Tianjin, Tianjin, P.R. China (Ruan Jishou)
North Carolina State University, Raleigh, NC (Dorothy Doyle)
Northern Jiaotong University, Beijing, P.R. China (Wang Xiaoxia)
NUI Galway, Galway, Ireland (Niall Madden)
Pacific Lutheran University, Tacoma, WA (Zhu Mei)
School of Mathematics and Computer Science, Nanjing Normal University, Nanjing, Jiangsu, P.R. China (Zhu Qunsheng)
Shanghai Jiading No. 1 High School, Shanghai, P.R. China (Chen Li)
Shanghai Jiaotong University, Shanghai, P.R. China (Song Baorui)
South China University of Technology, Guangzhou, Guangdong, P.R. China (Lin Jian Liang)
South China University of Technology, Guangzhou, Guangdong, P.R. China (Zhuo Fu Hong)
Stetson University, DeLand, FL (Lisa O. Coulter)
Tianjin University, Tianjin, P.R. China (Rong Ximin)
Tsinghua University, Beijing, P.R. China (Hu Zhiming)
U.S. Military Academy, West Point, NY (Elizabeth Schott)
University of South Carolina, Columbia, SC (Ralph E. Howard)
University of Washington, Seattle, WA (Timothy P. Chartier)
Xidian University, Xi'an, Shaanxi, P.R. China (Zhang Zhuo-kui)
Youngstown State University, Youngstown, OH (Angela Spalsbury)
Youngstown State University, Youngstown, OH (Stephen Hanzely)

Awards and Contributions

Each participating MCM advisor and team member received a certificate signed by the Contest Director and the appropriate Head Judge.

INFORMS, the Institute for Operations Research and the Management Sciences, gave a cash prize and a three-year membership to each member of the teams from North Carolina School of Science and Mathematics (Wind and Waterspray Problem) and Wake Forest (Team 69) (Airline Overbooking Problem). Also, INFORMS gave free one-year memberships to all members of Meritorious and Honorable Mention teams.

The Society for Industrial and Applied Mathematics (SIAM) designated one Outstanding team from each problem as a SIAM Winner. The teams were from University of Washington (Wind and Waterspray Problem) and Duke University (Airline Overbooking Problem). Each of the team members was awarded a $300 cash prize and the teams received partial expenses to present their results at a special Minisymposium of the SIAM Annual Meeting in Philadelphia, PA in July. Their schools were given a framed, hand-lettered certificate in gold leaf.

The Mathematical Association of America (MAA) designated one Outstanding team from each problem as an MAA Winner. The teams were from U.S. Military Academy (Wind and Waterspray Problem) and Harvey Mudd College

(Airline Overbooking Problem). With partial travel support from the MAA, both teams presented their solutions at a special session of the MAA Mathfest in Burlington, VT in August. Each team member was presented a certificate by MAA President Ann E. Watkins.

Judging

Director
Frank R. Giordano, COMAP, Lexington, MA

Associate Directors
Robert L. Borrelli, Mathematics Dept., Harvey Mudd College, Claremont, CA
Patrick Driscoll, Dept. of Mathematical Sciences, U.S. Military Academy, West Point, NY

Contest Coordinator
Kevin Darcy, COMAP Inc., Lexington, MA

Wind and Waterspray Problem

Head Judge
Marvin S. Keener, Executive Vice-President, Oklahoma State University, Stillwater, OK

Associate Judges
William C. Bauldry, Appalachian State University, Boone, NC
Kelly Black, Mathematics Dept., University of New Hampshire, Durham, NH (SIAM)
Courtney Coleman, Mathematics Dept., Harvey Mudd College, Claremont, CA
Gordon Erlebacher, School of Computational Science and Information Technology, Florida State University, Tallahassee, FL
J. Douglas Faires, Youngstown State University, Youngstown, OH (MAA)
Ben Fusaro, Mathematics Dept., Florida State University, Tallahassee, FL
Mario Juncosa, RAND Corporation, Santa Monica, CA
John Kobza, Texas Tech University, Lubbock, TX (INFORMS)
Deborah Levinson, Compaq Computer Corp., Colorado Springs, CO
Veena Mendiratta, Lucent Technologies, Naperville, IL
Mark R. Parker, Mathematics Dept., Carroll College, Helena, MT
John L. Scharf, Carroll College, Helena, MT
Daniel Zwillinger, Newton, MA

Airline Overbooking Problem

Head Judge
Maynard Thompson, Mathematics Dept., University of Indiana, Bloomington, IN

Associate Judges
James Case, Baltimore, MD (SIAM)
Lisette De Pillis, Harvey Mudd College, Claremont, CA
William P. Fox, Francis Marion University, Florence, SC (MAA)
Jerry Griggs, University of South Carolina, Columbia, SC
Don Miller, Dept. of Mathematics, St. Mary's College, Notre Dame, IN (SIAM)
Lee Seitelman, Glastonbury, CT (SIAM)
Dan Solow, Mathematics Dept., Case Western Reserve University, Cleveland, OH (INFORMS)
Robert Tardiff, Salisbury State University, Salisbury, MD
Michael Tortorella, Lucent Technologies, Holmdel, NJ
Marie Vanisko, Carroll College, Helena, MT (MAA)
Larry Wargo, National Security Agency, Ft. Meade, MD (Triage)

Triage Sessions:

Wind and Waterspray Problem

Head Triage Judge
Patrick Driscoll, Dept. of Mathematical Sciences, U.S. Military Academy, West Point, NY

Associate Judges
Steve Horton, Michael Jaye, and Doug Matty, all of the U.S. Military Academy, West Point, NY

Airline Overbooking Problem

Head Triage Judge
Larry Wargo, National Security Agency, Ft. Meade, MD

Associate Judges
James Case, Baltimore, Maryland
Paul Boisen and 7 others from the National Security Agency, Ft. Meade, MD

Sources of the Problems

The Wind and Waterspray Problem was contributed by Tjalling Ypma, Mathematics Dept., Western Washington University, Bellingham, WA. The Airline Overbooking Problem was contributed by William P. Fox and Richard D. West, Mathematics Dept., Francis Marion University, Florence, SC.

Acknowledgments

Major funding for the MCM is provided by the National Security Agency and by COMAP. We thank Dr. Gene Berg of NSA for his coordinating efforts. Additional support is provided by the Institute for Operations Research and the Management Sciences (INFORMS), the Society for Industrial and Applied Mathematics (SIAM), and the Mathematical Association of America (MAA). We are indebted to these organizations for providing judges and prizes.

We thank the MCM judges and MCM Board members for their valuable and unflagging efforts. Harvey Mudd College, its Mathematics Dept. staff, and Prof. Borrelli were gracious hosts to the judges.

Cautions

To the reader of research journals:

Usually a published paper has been presented to an audience, shown to colleagues, rewritten, checked by referees, revised, and edited by a journal editor. Each of the student papers here is the result of undergraduates working on a problem over a weekend; allowing substantial revision by the authors could give a false impression of accomplishment. So these papers are essentially au naturel. Editing (and sometimes substantial cutting) has taken place: minor errors have been corrected, wording has been altered for clarity or economy, and style has been adjusted to that of *The UMAP Journal*. Please peruse these student efforts in that context.

To the potential MCM Advisor:

It might be overpowering to encounter such output from a weekend of work by a small team of undergraduates, but these solution papers are highly atypical. A team that prepares and participates will have an enriching learning experience, independent of what any other team does.

COMAP's Mathematical Contest in Modeling and Interdisciplinary Contest in Modeling are the only international modeling contests in which students work in teams. Centering its educational philosophy on mathematical modeling, COMAP uses mathematical tools to explore real-world problems. It serves the educational community as well as the world of work by preparing students to become better-informed and better-prepared citizens.

Appendix: Successful Participants

KEY:

P = Successful Participation
H = Honorable Mention
M = Meritorious
O = Outstanding (published in this special issue)

A = Bicycle Wheel Control Problem
B = Hurricane Evacuation Problem

INSTITUTION	CITY	ADVISOR	A	B
ALABAMA				
Huntingdon College	Montgomery	Vyacheslav V. Rykov	P,P	
ARKANSAS				
Hendrix College	Conway	Duff Gordon Campbell	H	H
ARIZONA				
McClintock High School	Tempe	Ivan Barkdoll	P	
CALIFORNIA				
Calif. Institute of Technology	Pasadena	Darryl H. Yong	H	
Calif. Polytechnic State Univ.	San Luis Obispo	Thomas O'Neil	M,P	
		Jennifer M. Switkes	P	
California State University at Monterey Bay	Seaside	Hongde Hu	P	
California State University	Bakersfield	Maureen E. Rush	P	
Claremont McKenna College	Claremont	Mario U. Martelli	P	
Harvey Mudd College	Claremont	Michael E. Moody		O,M
		Ran Libeskind-Hadas		M,M
Humboldt State University	Arcata	Roland H. Lamberson	M	
Pomona College	Claremont	Ami E. Radunskaya		H
Sonoma State University	Rohnert Park	Elaine T. McDonald	H	P
University of San Diego	San Diego	Jeffrey H. Wright		P
University of Southern Calif.	Los Angeles	Robert J. Sacker		H
		Geoffrey R. Spedding	H	
COLORADO				
Colorado College	Colorado Springs	Peter L. Staab	P	
Colorado State University	Fort Collins	Michael J. Kirby		P
Mesa State College	Grand Junction	Edward. K. Bonan-Hamada	P	
Regis University	Denver	Linda L. Duchrow	P	
U.S. Air Force Academy	USAF Academy	Gerald E. Sohan	P	
		James S. Rolf		H
		J. Gerken	H	
		James E. West	P	
University of Colorado	Boulder	Anne M. Dougherty	M	O
		Michael Ritzwoller	M,M	
Univ. of Southern Colorado	Pueblo	Bruce N. Lundberg	P	

INSTITUTION	CITY	ADVISOR	A	B
CONNECTICUT				
Sacred Heart University	Fairfield	Peter Loth	P,P	
Southern Connecticut State Univ.	New Haven	Ross B. Gingrich		P
		Therese L. Bennett		P
DELAWARE				
University of Delaware	Newark	Louis F. Rossi	P,P	
FLORIDA				
Embry-Riddle Aeronautical Univ.	Daytona Beach	Greg Scott Spradlin	P	
Florida Gulf Coast University	Fort Myers	Charles Lindsey	P	
Florida State University	Tallahassee	Mark M. Sussman	P	
Jacksonville University	Jacksonville	Robert A. Hollister	M,P	
Stetson University	DeLand	Lisa O. Coulter		M
University of Central Florida	Orlando	Heath M. Martin	P,P	
Florida Institute of Technology	Melbourne	Michael O. Gonsalves	P,P	
GEORGIA				
Georgia Southern University	Statesboro	Jacalyn M. Huband		P
		Laurene V. Fausett		H
State University of West Georgia	Carrollton	Scott Gordon	M,P	
IDAHO				
Albertson College of Idaho	Caldwell	Mike P. Hitchman		M
Boise State University	Boise	Jodi L. Mead	P	
Idaho State University	Pocatello	Rob Van Kirk	P	
ILLINOIS				
Greenville College	Greenville	Galen R. Peters	P	H
McKendree College	Lebanon	Raymond E. Robb		P
Monmouth College	Monmouth	Christopher G. Fasano	P	
Wheaton College	Wheaton	Paul Isihara	P	
INDIANA				
Earlham College	Richmond	Charlie Peck	H	
Goshen College	Goshen	David Housman	M	H
Indiana University	Bloomington	Michael S. Jolly	H	
Rose-Hulman Institute of Tech,	Terre Haute	David J. Rader	P	H
Saint Mary's College	Notre Dame	Joanne R. Snow	P,P	
IOWA				
Grinnell College	Grinnell	Mark Montgomery	P	H
		Royce Wolf	H.P	
Iowa State University	Ames	Stephen J. Willson		H
Luther College	Decorah	Reginald D. Laursen	M,M	
Mt. Mercy College	Cedar Rapids	K.R. Knopp	P	P
Simpson College	Indianola	Murphy Waggoner	P	P
		Werner S. Kolln	P	H
Wartburg College	Waverly	Mariah Birgen	P	

INSTITUTION	CITY	ADVISOR	A	B
KANSAS				
Kansas State University	Manhattan	Korten N. Auckly	P	
KENTUCKY				
Asbury College	Wilmore	Kenneth P. Rietz	M	M
Bellarmine University	Louisville	William J. Hardin	P	P
Northern Kentucky University	Highland Heights	Gail S. Mackin	P	
		Shamanthi Marie Fernando	P	
MAINE				
Colby College	Waterville	Jan Holly	P	
MARYLAND				
Salisbury University	Salisbury	Michael J. Bardzell		P
Goucher College	Baltimore	Robert Lewand	P	
Hood College	Frederick	Betty Mayfield	P	
Mt. Saint Mary's College	Emmitsburg	John J. Dropp		P
		John E. August	P	P
Towson University	Towson	Mike P. O'Leary	P	
Washington College	Chestertown	Eugene P. Hamilton	P	
MASSACHUSETTS				
Boston University	Boston	Glen R. Hall	M	
Greenfield Community College	Greenfield	Peter R. Letson	P	
Massachusetts Institute of Tech.	Cambridge	Daniel H. Rothman	M,P	
		Martin Zdenek Bazant		M
Olin College of Engineering	Needham	Burt S. Tilley	P	
Simon's Rock College	Great Barrington	Allen B. Altman	P	H
		Michael Bergman	H	H
University of Massachusets	Amherst	Edward A. Connors		P
University of Massachusetts	Lowell	James Graham-Eagle	H	
Western New England College	Springfield	Lorna B. Hanes	P	
Worcester	Arthur	C. Heinricher	P	
MICHIGAN				
Calvin College	Grand Rapids	Thomas L. Jager	P	
Eastern Michigan University	Ypsilanti	Christopher E. Hee	H	
Hillsdale College	Hillsdale	John P. Boardman		P
Hope College	Holland	Aaron C. Cinzori		P
Lake Superior State University	Sault Sainte Marie	John H. Jaroma	H	
Lawrence Technological Univ.	Southfield	Howard E. Whitston	M	
		Ruth G. Favro		H
		Scott D. Schneider	M	
		Valentina Tobos	P	
Siena Heights University	Adrian	Toni Carroll	P,P	

INSTITUTION	CITY	ADVISOR	A	B
MINNESOTA				
Augsburg College	Minneapolis	Rebekah N. Dupont	H	
Bemidji State University	Bemidji	Colleen G. Livingston		H
Bethel College	St. Paul	William M. Kinney	M	O
College of Saint Benedict, St. John's University	Collegeville	Robert J. Hesse	P,P	
Gustavus Adolphus College	St. Peter	Thomas P. LoFaro	P	H
Macalester College	St. Paul	Daniel T. Kaplan	P	P
		Elizabeth Shoop		P
Winona State University	Winona	Barry A. Peratt		P
MISSOURI				
Missouri Southern State College	Joplin	Patrick Cassens	P	
Northwest Missouri State University	Maryville	Russell N. Euler	P	
Southeast Missouri State University	Cape Girardeau	Robert W. Sheets	P	
Truman State University	Kirksville	Steve J. Smith		P
Washington University	St. Louis	Hiro Mukai		P
MONTANTA				
Carroll College	Helena	Holly S. Zullo	H	H
NEBRASKA				
University of Nebraska-Lincoln	Lincoln	Glenn W. Ledder		H
NEW JERSEY				
Rowan University	Glassboro	Hieu D. Nguyen		P
		Paul J. Laumakis	P	
William Paterson University	Wayne	Donna J. Cedio-Fengya	H	
NEW MEXICO				
New Mexico State University	Las Cruces	Caroline Sweezy		P
Western New Mexico University	Silver City	Thomas P. Gruszka		H
NEW YORK				
Colgate University	Hamilton	Warren Weckesser		M
Ithaca College	Ithaca	Jim Conklin	P	
Keuka College	Keuka Park	Catherine A Abbott		P
Manhattan College	Riverdale	Kathryn Weld		H
Rensselaer Polytechnic Institute	Troy	Peter R. Kramer	P	
Roberts Wesleyan College	Rochester	Gary L. Raduns		H
St. Bonaventure University	St. Bonaventure	Albert G. White	P	
U.S. Military Academy	West Point	David Sanders	O	
		Elizabeth Schott		M
		Gregory S. Parnell	H	
		Ray Eason	H	
Wells College	Aurora	Thomas A. Stiadle	P	
Westchester Community College	Valhalla	Sheela L. Whelan	P	

INSTITUTION	CITY	ADVISOR	A	B
NORTH CAROLINA				
Appalachian State University	Boone	Alan T. Arnholt	P	P
		Eric S. Marland	P,P	
Brevard College	Brevard	Clarke Wellborn	P	P
Duke University	Durham	David Kraines	P	O
Meredith College	Raleigh	Cammey E. Cole	P	
Mount Olive College	Mount Olive	Ollie J. Rose	P	P
North Carolina School of Science and Mathematics	Durham	Daniel J. Teague	O,P	
		John Kolena	H	
North Carolina State University	Raleigh	Thomas L. Honeycutt	H,P	
		Dorothy Doyle		M,P
University of North Carolina	Chapel Hill	Jon W. Tolle	M	
Wake Forest University	Winston-Salem	Frederick H. Chen		O
		Miaohua Jiang		P
		Edward E. Allen		O
OHIO				
Hiram College	Hiram	Brad Scott Gubser	P	
Malone College	Canton	David W. Hahn	P	P
Miami University	Oxford	Doug E. Ward		P
College of Wooster	Wooster	Charles R. Hampton	M	P
Wright State University	Dayton	Thomas P. Svobodny	M	
Youngstown State University	Youngstown	Angela Spalsbury	M	M
		Bob Kramer	P	
		Stephen Hanzely		M
University of Dayton	Dayton	Muhammad N. Islam	P	
Xavier University	Cincinnati	Michael Goldweber	M	
OKLAHOMA				
Southeastern Oklahoma State University	Durant	Christopher Moretti		P
OREGON				
Eastern Oregon University	La Grande	Anthony A. Tovar	P,P	
		Robert L. Brandon		P
Lewis and Clark College	Portland	Robert W. Owens	P	H
Portland State University	Portland	Gerardo A. Lafferriere	H	
Southern Oregon University	Ashland	Kemble R. Yates	M	H
University of Portland	Portland	Michael W. Akerman		P
PENNSYLVANIA				
Bloomsburg University	Bloomsburg	Kevin K. Ferland	P,P	
Bucknell University	Lewisburg	Karl Knaub		P
		Sally Koutsoliotas	H,P	
Clarion University	Clarion	Karen D. Bolinger	P	
		Richard M. Smaby	P	
		John W. Heard	P	

INSTITUTION	CITY	ADVISOR	A	B
Gettysburg College	Gettysburg	James P. Fink		M
Indiana University of Penna	Indiana	Frederick A. Adkins	P	
Juniata College	Huntingdon	John F. Bukowski		M
Lafayette College	Easton	Thomas Hill	M	P
Messiah College	Grantham	Lamarr C. Widmer	P	
Westminster College	New Wilmington	Barbara T. Faires	P,P	
SOUTH CAROLINA				
Charleston Southern University	Charleston	Stan Perrine	P	
University of South Carolina	Columbia	Ralph E. Howard		M
SOUTH DAKOTA				
SD School of Mines and Tech.	Rapid City	Kyle Riley	P	P
TENNESSEE				
Austin Peay State University	Clarksville	Nell K. Rayburn	P	P
Lipscomb University	Nashville	Gary Clark Hall		M
University of the South	Sewanee	Catherine E. Cavagnaro	P	P
TEXAS				
Abilene Christian University	Abilene	David Hendricks		P
Angelo State University	San Angelo	Robert L. Hamilton	P	
Baylor University	Waco	Frank H. Mathis	P	
Southwestern University	Georgetown	Therese N. Shelton	H,P	
Texas Christian University	Fort Worth	George T. Gilbert	P	
Trinity University	San Antonio	Hector C. Mireles	M	
		Jeffrey K. Lawson	M	
		Kenneth L. Nelson		H
VERMONT				
Johnson State College	Johnson	Christopher A. Aubuchon		P
VIRGINIA				
Chesterfield County Mathematics and Science High School	Midlothian	Diane Leighty	P	
Godwin High School Science and Mathematics Center	Richmond	Ann W. Sebrell		P
James Madison Univ	Harrisonburg	Joseph W. Rudmin	P	
		Caroline Smith	M	
		Dorn W. Peterson		P
		Paul G. Warne	P	
Maggie L. Walker Governor's School	Richmond	Crista Hamilton		M
		John A. Barnes	O	M
Roanoke College	Salem	Jeffrey L. Spielman	P	
University of Richmond	Richmond	Kathy W. Hoke	P	
University of Virginia's College at Wise	Wise	George W. Moss	P,P	

INSTITUTION	CITY	ADVISOR	A	B
Virginia Tech	Blacksburg	Catherine A. Stephens		P
		Laura J. Spielman	H	
Virginia Western Comm. College	Roanoke	Steven T. Hammer		P
		Ruth A. Sherman	P	
WASHINGTON				
Central Washington University	Ellensburg	Stuart F. Boersma	P,P	
Pacific Lutheran University	Tacoma	Mei Zhu	H	M
University of Puget Sound	Tacoma	John Riegsecker		P
		Michael Scott Casey		H,P
University of Washington	Seattle	Anne Greenbaum	H	
		James Allen Morrow	O,M	
		Timothy P. Chartier		M
Western Washington University	Bellingham	Tjalling J. Ypma	P	P
WISCONSIN				
Beloit College	Beloit College	Paul J. Campbell		P
Ripon College	Ripon	David W. Scott	P,P	
University of Wisconsin	River Falls	Kathy A. Tomlinson		P
AUSTRALIA				
University of New South Wales	Sydney	James W. Franklin	M,H	
CANADA				
Brandon University	Brandon, Manitoba	Doug A. Pickering		H
Memorial Univ. of Newfoundland	St. John's NF	Andy Foster		H
Dalhousie University	Halifax NS	Dorothea A. Pronk	H	
University of Western Ontario	London ON	Peter H. Poole	H	
University of Saskatchewan	Saskatoon SK	James A. Brooke		P
CHINA				
Anhui University	Hefei	Cai Qian		P
		Yang Shangjun	H	
Baoding Teachers' College	Baoding	Jing Shuangyan	P,P	
		Wang Xinzhe	P,P	
		Yuan Shaoqiang	P	P
Beijing Institute of Technology	Beijing	Chen Yi Hong	P	
		Cui Xiao Di		H
		Zhen Yao Cui	M	
		Zhang Bao Xue		M
Beijing Polytechnic University	Beijing	Xue Yi	P	
Beijing Union University	Beijing	Zeng Qingli	P	P
Beijing Univ. of Aero. and Astro.	Beijing	Peng Linping	H	
		Liu Hongying		P
		Wu Sanxing		P

INSTITUTION	CITY	ADVISOR	A	B
Beijing Univ. of Chemical Technology	Beijing	Jiang Guangfeng	H	
		Weiguo Lin		P
		Jiang Dongqing		P
		Liu Damin		P
		Yuan WenYan	M	
Beijing Univ. of Posts & Telecomm.	Beijing	He Zuguo	M,M	
		Luo Shoushan		H,H
		Sun Hongxiang	M	
Central South University	Chang Sha	Chen Xiaosong		H
		Qin Xuanyun	M	P
ChangChun University	Changchun	Yuan Shuai	P	
China University of Mining and Technology	Xuzhou	Wu Zongxiang		H
		Zhang Xingyong		H
		Zhou Shengwu		P
		Zhu Kaiyong		M
Chongqing University	Chongqing	He Zhongshi		H
		Li Chuandong	P	
		Wen Luosheng		P
		Yang Xiaofan		M
Dalian Nationalities University	Dalian	Wang Jinzhi		P,P
Dalian University	Dalian	Tan Xinxin	H	
Dalian University of Technology	Dalian	He Mingfeng		P,P
		Zhao Lizhong	P	P
		Wang Yi		P,P
Dong Hua University	Shanghai	Hu Liangjian		P
		Du Yugen		H
East China University of Science and Techn.	Shanghai	Su Chunjie		H,P
		Lu Yuanhong	M,H	
Educational Adminstration	Shenyang	Zhang Shujun		P,P
Fudan University	Shanghai	Hu Jinjin		P
		Cao Yuan		P
		Cai Zhijie		M
Guangdong Commercial College	Guangzhou	Xiang Zigui		P
Hangzhou University of Commerce	Hangzhou	Ding Zhengzhong		H
		Hua Jiu Kun		H
		Zhao Heng	M	
		Zhu Ling	M	
Harbin Engineering University	Harbin	Gao Zhenbin	P	
		Luo Yuesheng		P
		Shen Jihong		P
		Zhang Xiaowei		P
Harbin Institute of Technology	Harbin	Shao Jiqun		P
		Shang Shouting	P	H
		Wang Xuefeng	P	M

INSTITUTION	CITY	ADVISOR	A	B
Harbin Uinversity of Science and Technology	Harbin	Chen Dongyan	P	
		LiDongmei	P	
		Tian Guangyue		P
		Ni Xiaoying	P	
Hefei University of Technology	Hefei	Su Huaming	H	
		Du Xueqiao		H
		Zhou Yongwu		P
		Huang Youdu		P
Huazhong University of Science & Technology	Wuhan	Wang Yizhi		H
Hunan University	Changsha	Yi Kunnan		P
		Han Luo	P	H
Jiamusi University	Jiamusi	Wei Fan		P
		Zhi Gu Li		P
Jiamusi University College of Sciences	Jiamusi	Shan Bai Feng	P	
		Liv Tongfu		P
Jiangxi Normal University	NanChang	Wu Gen Xiu		P
Jilin Institute of Technology	Changchun	Chun Sun Chang	P	
		Yue Xi Ting		P
Jilin University	Changchun	Iv Xianrui	P	
		Fang Peichen		P
		Song Lixin	P	
		Zou Yongkui	H	
Jinan University	Guangzhou	Hu Daiqiang	H	
		Fan Suohai		P
		Ye Shiqi	P,P	
Nanjing Normal University	Nanjing	Fu Shitai		H,P,P
Nanjing Normal University School of Mathematics and Computer Science	Nanjing	Zhu Qunsheng		M,H
Nanjing University of Science and Technology	Nanjing	Zhang Zhengjun		P
		Xu Yuan	H	
		Xu Chungen	P	
		Huang Zhen You	P	
Nankai University	Tianjin	Huang Wuqun	M	
		Ruan Jishou		M,H
Nankai University School of Mathematics	Tianjin	Ruan Jishou		H
		Yang Qingzhi		P
National Univ. of Defence Tech.	Changsha	Wu Mengda		P
		Duan XiaoLong		P,P
		Wu Mengda		H
North China Electric Power University	Baoding	Gu Gendai	M	
		Xie Hong	H	
North China Inst. of Technology	Taiyuan	Xue Ya-kui	H	
		Lei Ying-jie	H	
		Yong Bi		P

INSTITUTION	CITY	ADVISOR	A	B
Northeastern University				
College of Sciences	Shenyang	Sun Ping	P	P
		Han Tie-min		M,P
Information Science and Engineering	Shenyang	Hao Peifeng	P	H
Inst. of Artificial Intelligence & Robotics	Shenyang	Cui Jianjiang	P	H
Northern Jiaotong University		Wang Bingtuan	M	
		Wu Faen	P	
		Wang Bingtuan		P
		Wang Xiaoxia		M
Northwest University	Xi'an	He Rui-chan	B,B	
Northwestern Polytechnical University	Xi'an	Zhang Li Ning	P	
		Hua Peng Guo	H	
		Sun Hao		P
		Xu Wei	H	
Peking University	Beijing	Lian Guijun	P	P
		Wang Ming	P	P
		Zhang Mingquan		P
		Shu Yousheng	P	
		Liu Yulong	P	P
Qufu Normal University	Qufu	Wang Feng	P	P
Shandong University	Jinan	Luan Junfeng		P
Shanghai Foreign Language School	Shanghai	Wan Baihe		P
		Pan LiQun		H,P
Shanghai Jiading No. 1 High School	Shanghai	Xie XiLin		H
		Chen Li		M
Shanghai Jiaotong University	Shanghai	Song Baorui		M,H
		Huang Jianguo		H,P
Shanghai Normal University	Shanghai	Zhu Detong		P
		Zhang Jizhou		P
		Guo Shenghuan		P
Shanghai University	Shanghai	Wang Yuandi	H	
Shanghai Univ. of Finance and Economics	Shanghai	Yang Xiao Bin	P	P
Shanghai XiangMing Senior High School	Shanghai	Wang Daren		P
Shanxi University	Taiyuan	Yang Aimin		P
		Li Jihong	P	
		Zhao Aimin	P	P
South China Normal University	Guangzhou	Wang Limin	H,P	
South China Univ. of Tech.	Guangzhou	Zhuo Fu Hong		M
		Lin Jian Liang		M
		Tao Zhi Sui		H
		Feng Zhu Feng		H
South-west Jiao Tong University	Chengdu	Zhao Lian Wen	H,H	
		Han Yang		P,P
Tianjin University	Tianjin	Rong Ximin		M,H
		Liu Zeyi	H	H
Tsinghua University	Shanghai	Hu Zhiming		M,H
		Ye Jun		P,P

INSTITUTION	CITY	ADVISOR	A	B
Univ. of Elec. Sci. Tech.	Chengdu	Du Hongfei	P	P
		Qin Siyi	M	P
Univ. of Sci. & Tech. of China	Hefei	Dou Dou		P
		Sun Guangzhong		H
		Li Yu		P
		Yang Zhouwang	H	
Wuhan University of Tech.	Wuhan	Peng Sijun	H	P
		Huang Zhangcan	H,P	
Xi'an Jiaotong University	Xi'an	Zhou Yicang	P	
		Wu Xiang Zhong	P	
		Dai Yonghong	H	
Xi'an University of Tech.	Xi'an	Cao Maosheng		P
Xidian University	Xi'an	Chen Hui-chan		H
		Liu Hong-wei	H	
		Ye Ji-min	H	
		Zhang Zhuo-kui		M
Zhejiang University	Hangzhou	Yang Qifan	M	H
		Yong He	P	P
Zhongshan University	Guangzhou	Chen Zepeng		P
		Tang Mengxi		P
		Yuan Zhuojian	H,H	
FINLAND				
Päivölä College	Tarttila	Merikki Lappi	H	
HONG KONG				
Hong Kong Baptist Univ.	Kowloon	Tong Chong-sze	H	
		Shiu Wai-chee	P	
INDONESIA				
Institut Teknologi Bandung	Bandungr	Edy Soewono		M
		Kuntjoro Adji Sidarto	H	
IRELAND				
National Univ. of Ireland	Galway	Niall Madden		M,P
Trinity College Dublin	Dublin	Timothy G. Murphy		P
University College Cork	Cork	Donal J. Hurley	M	
		James J. Grannell	H	
University College Cork	Cork	Supratik Roy		P
University College Dublin	Belfield	Ted Cox	H,P	
		Maria G. Meehan	P,P	
University College Dublin	Dublin	Peter Duffy	M	
		Maria G. Meehan		P,P
LITHUANIA				
Vilnius University	Vilnius	Ricardas Kudzma		P
SINGAPORE				
National Univ. of Singapore	Singapore	Victor Tan	B	

INSTITUTION	CITY	ADVISOR	A	B
SOUTH AFRICA				
University of Stellenbosch	Matieland	Jan H. Van Vuuren	H	P
UNITED KINGDOM				
Magdalen College	Oxford, England	Byron W. Byrne	M	

Editor's Note

For team advisors from China and Singapore, we have endeavored to list family name first, with the help of Susanna Chang, Beloit College '03.

Simulating a Fountain

Lyric P. Doshi
Joseph Edgar Gonzalez
Philip B. Kidd
Maggie L. Walker Governor's School
for Government and International Studies
Richmond, VA

Advisor: John A. Barnes

Introduction

We establish the mathematical behavior of water droplets emitted from a fountain and apply this behavior in a computer model to predict the amount of splash and spray produced by a fountain under given conditions. Our goal is a control system that creates the tallest fountain possible while limiting water spillage to a specified level.

We combine height and volume of the fountain spray, making both functions of the speed at which water exits the fountain nozzle. We simulate water droplets launched from the fountain, using basic physics to model the effects of drag, wind, and gravity. The simulation tracks the flight of droplets in the air and records their landing positions, for wind speeds from 0 to 15 m/s and water speeds from 5 to 30 m/s. It calculates the amount of water spilled outside of a pool around the fountain, for pool radii from 0 to 40 m.

We design an algorithm for a programmable logic controller, located inside an anemometer, to do a table search to find allowable water speeds for given pool radius, acceptable water spillage, and wind velocity. We test the control system with simulation, subjecting a fountain with a 4-m pool radius to wind speeds from 0 to 3 m/s with an allowable spillage of 5%. We also test the model for accuracy and for sensitivity to changes in the base variables.

The UMAP Journal 23 (3) (2002) 209–219.

Problem Analysis

Wind

The anemometer measures two main wind factors that affect the fountain: speed, which affects the force exerted on the water, and direction.

Fountain

The main components of the fountain are the pool and the nozzle. The factors associated with the pool are its radius, which remains constant within a trial, and the acceptable level of spillage, which describes the percentage of water that may acceptably fall outside of the fountain.

Nozzle

Major aspects of the nozzle are the radius of the opening, the angle relative to the vertical axis (normal axis), and the spread and speed of the water passing through it. The angle of the nozzle relative to the vertical axis determines the initial trajectory of the water. The spread, described in standard deviations from the angle of the nozzle, determines the extent to which the initial trajectory of droplets differs from the angle of the nozzle. For a given water speed and nozzle radius, the flow of water through the nozzle may be determined from

$$f = \pi r^2 v,$$

where f is flow, v is the water launch speed, and r is the radius of the nozzle. The radius is constant, so the flow and consequent volume are functions of the speed, the dominant controllable factor affecting the height of the stream.

Assumptions

... about Fountains

- The fountain is composed of a single nozzle located at the center of a circular pool.
- The ledge of the pool is sufficiently high to collect the splatter produced by particles impacting the surface of the water.
- Fountains with higher streams are more attractive than those with lower streams.

. . . about the Nozzle

- The nozzle has a fixed radius, but the speed of the water through it can be controlled.
- The nozzle is perpendicular to the ground.
- The nozzle responds rapidly to input from the anemometer.
- The nozzle produces a normally distributed spread of droplets with a low standard deviation.

. . . about Water Droplets

- Because the droplets are small and roughly spherical, they may be treated as spherical.
- The radii of droplets are normally distributed.
- The density of water is unaffected by conditions and therefore remains constant among and within droplets.
- The only outside forces exerted on a water droplet are gravity and the force exerted by the surrounding air, including drag and wind.
- Acceleration due to gravity is the same for all droplets.
- The effect of air perturbations produced by droplets on other droplets is insignificant.
- All droplets share the same constant drag coefficient.
- Droplet interactions and collisions do not increase the overall energy of the system nor increase significantly the distance traveled by droplets.

. . . about the Anemometer and Control System

- The anemometer and control system can rapidly evaluate the wind speed, apply a basic formula, and adjust the nozzle in changing wind conditions.

. . . about the Wind

- The wind speed is uniform regardless of altitude.
- Wind blows parallel to the ground without turbulence or irregularities.

Basic Description of Model

Water droplets are emitted from the nozzle and follow trajectories affected by wind and drag. The particles are tracked until they land, including recalculations of trajectories in case of changes in conditions, such as wind. The landing distance from the center of the fountain is recorded. Since the fountain pool is circular, only radial distance is important.

The model ignores wind direction (does not affect a circular fountain pool) and turbulence (insignificant and too complicated to model accurately).

We tested droplet collisions and found that they do not greatly affect the distance that droplets land from the center of the pool; so we ruled out incorporating complex interactions into the model. Further physical analysis supported that decision: Since energy and momentum are conserved, a droplet could not travel significantly farther after a collision.

Finally, we combined fountain height and volume into speed of the water out of the nozzle, because they are directly determined by the speed.

Our simulation tries all combinations of 11 different water speeds, from 5 to 30 m/s (at intervals of 2.5 m/s), with 16 wind speeds, from 0 to 15 m/s (at intervals of 1 m/s). Each combination is run for five trials of 10,000 droplets. Spillage is logged for radii from 0 to 40 m (at intervals of 0.1 m). The five trials are then averaged to construct an entry in a three-dimensional reference table with axes of radial distance from nozzle, wind speed, and water speed. The control system functions by referring in the table to the current wind speed, the fountain's radius, and the acceptable level of spillage to identify the corresponding maximum water speed.

The Underlying Mathematics

The simulation uses basic physics equations to model the flight of water droplets through the air.

Each droplet is acted on by three forces: gravity, drag, and wind. Drag is calculated from the following equation [Halliday et al. 1993]:

$$D = \tfrac{1}{2} C \rho A v^2,$$

where

D is the drag coefficient, an empirically-determined constant dependent mainly on the shape of an object;

ρ is the density of the fluid through which the object is traveling, in this case air;

A is the cross-sectional area of the object; and

$v = |\vec{v}|$ is the speed of the object relative to the wind.

The drag coefficient of a raindrop is 0.60 and the density of air is about 1.2 kg/m^3 [Halliday et al. 1993]. Drag acts directly against velocity, so the acceleration vector from drag can be found from Newton's law $\vec{F} = m\vec{a}$ as

$$\vec{a} = \frac{-D}{m}\frac{\vec{v}}{|v|} = \frac{\frac{1}{2}C\rho A|\vec{v}|^2}{m}\frac{\vec{v}}{|\vec{v}|} = \frac{\frac{1}{2}C\rho A|\vec{v}|}{m}\,\vec{v},$$

where $\vec{a}$ is the acceleration vector and m is mass.

We factor in gravity by subtracting the acceleration g of gravity at Earth's surface, 9.8 m/s^2, from the vertical component of the acceleration vector:

$$\vec{a}_z = -\frac{\frac{1}{2}C\rho A|\vec{v}|}{m}\,\vec{v}_z - g.$$

Next, we use the acceleration to find velocity, beginning with the expression

$$\frac{d\vec{v}}{dt} = -\frac{\frac{1}{2}C\rho A|\vec{v}|}{m}\,\vec{v} = \vec{a}.$$

To circumvent the difficulties of solving a differential equation for each component of the velocity vector, we use Euler's method to approximate the velocity at a series of discrete points in time:

$$\frac{d\vec{v}}{dt} = \vec{a}, \qquad \Delta\vec{v} \approx \Delta t\vec{a}, \qquad \vec{v}_1 \approx \vec{v}_0 + \Delta t\vec{a}_0.$$

We use a similar process to find the position of the droplet, resulting in

$$\vec{x}_1 \approx \vec{x}_0 + \Delta t\vec{v}_0.$$

With $\Delta t = 0.001$ s, error from the approximation is virtually zero.

Now that we have equations for describing the droplet in flight, we generate its initial position and velocity. First, we randomly select a value z from a standard Gaussian (normal) distribution (mean 0, standard deviation 1). We calculate the angle from a set mean μ and standard deviation σ of the distribution of possible angles as

$$\phi = z\sigma + \mu.$$

We randomly select another angle θ between 0 and 2π radians to be the angle between the velocity vector and the x-axis.

Thus, the initial velocity vector of the droplet in spherical coordinates is (ρ, θ, ϕ), where ρ is the magnitude of the velocity. Conversion to rectangular coordinates yields $(\rho\sin\phi\cos\theta, \rho\sin\phi\sin\theta, \rho\cos\phi)$.

We also randomly select a starting location within the nozzle (whose diameter is 1 cm) and create a radius for the droplet using a similar sampling from a normal distribution. The mass of the droplet is then

$$m = \tfrac{4}{3}\pi r^3\rho,$$

where ρ is the density of water, 998.2 kg/m^3 at 20° C [Lide 1995]. In the basic simulation, the ϕ distribution has a mean of 0 and a standard deviation of $\pi/60$ radians, and the radius distribution has a mean of 0.0015 m and a standard deviation of 0.0001 m.

In the basic simulation, the nozzle points straight up; however, we also test the effect of tilting the nozzle into the wind. The program first rotates the nozzle a set angle away from the z-axis ($\pi/18$, $\pi/9$, or $\pi/6$ radians). The initial position and velocity vectors are changed by the formula for rotating a point t radians about the x-axis, from z towards negative y [Dollins 2001]:

$$\begin{bmatrix} x' \\ y' \\ z' \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos t & -\sin t \\ 0 & \sin t & \cos t \end{bmatrix} \begin{bmatrix} x \\ y \\ z \end{bmatrix}.$$

Next, the program rotates the nozzle around the z-axis to point directly away from the wind (in spherical coordinates, the θ of the nozzle is equal to that of the wind vector). The formula to rotate a point t radians about the z-axis, from x towards y [Dollins 2001] is

$$\begin{bmatrix} x' \\ y' \\ z' \end{bmatrix} = \begin{bmatrix} \cos t & -\sin t & 0 \\ \sin t & \cos t & 0 \\ 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \end{bmatrix}.$$

Design of Program

We developed a program to simulate the fountain. The program component `Simulator.class` manages interactions among the other components of the program. `Particle.class` describes a water droplet in terms of position, velocity, radius, and mass. `Vector3D.class` creates and performs functions with vectors, including setting vector components, adding and subtracting vectors, multiplying vectors by scalars, finding the angle between vectors, and finding the magnitude of a vector.

`Emitter.class` creates a fountain by spraying droplets. It considers the nozzle radius, direction, and angle orientations and generates launch angle ϕ and launch location on the nozzle according to the prescribed distributions.

Launch speed is determined by `Anemometer.class`, which takes the wind-speed reading from the anemometer and sends that plus fountain radius and tolerable spillage percentage to `FindingVelocity.class`. This latter class does a table lookup and returns the maximum droplet speed for the spillage percentage. `Anemometer.class` then sets the droplet emission speed.

Once a droplet is emitted, its trajectory is updated every iteration using `Physics.class`, which checks `Wind.class` (which contains a vector of the current wind) in each iteration in calculating an updated trajectory. Then `Physics.class` iterates through the entire collection of particles and computes new velocities and positions based on the forces acting on them.

The `Analyzer.class` checks to see if any particles have hit the ground; their locations are recorded and they are removed from consideration. It then relays this information back to `Simulator.class`, where it is written to disk.

Results

A program run takes 5 min to model 2 sec of spray (10,000 droplets).

Scatterplots showing where droplets land appear uniform and radially symmetric (**Figure 1**); a side profile of the points appears uniformly distributed along a line and bilaterally symmetric (**Figure 2**).

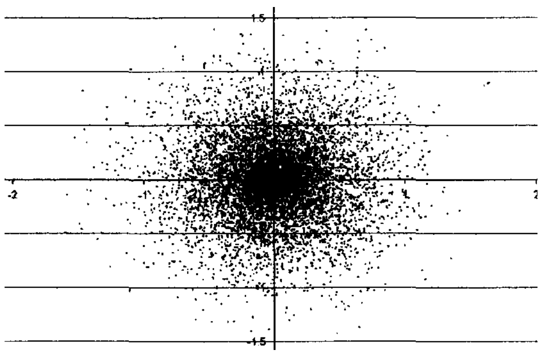

Figure 1. Fountain from overhead: launch speed 10 m/s, no wind.

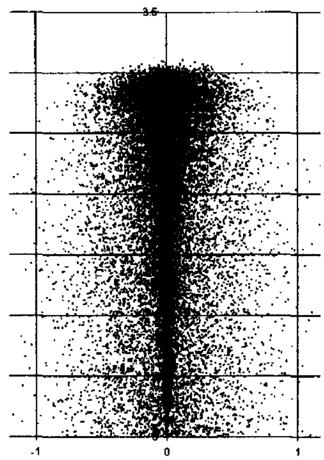

Figure 2. Fountain from the side: launch speed 10 m/s, no wind.

We then introduced wind in the positive x-direction. As expected, the landing plot and the side profile plot are skewed horizontally (**Figure 3**).

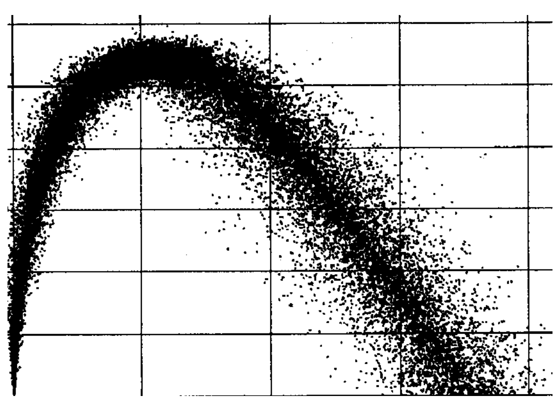

Figure 3. Fountain from the side: launch speed 10 m/s, wind of 5 m/s

Figures 1–3 conform very well to the actual appearance of fountains, indicating that our model creates an accurate portrait of a real fountain.

We used a pool radius of 4 m and an acceptable spillage of 5% to generate a table of water speeds. We then simulated control of the fountain by a theoretical anemometer using the table. The anemometer was subjected to sinusoidal wind ranging from 0 to 3 m/s. There was 7.6% spillage, a success since the extra loss above 5% is from droplets carried farther by an increase of wind after launch.

Analysis of Results

We tested the model for accuracy and sensitivity. We did some useful analysis of the physics of the model by creating a miniature version of the simulation on an Excel spreadsheet to track the trajectory of a single particle.

Our first test was of the accuracy of the Euler's method approximation. Continuous equations for the motion of a flying droplet can be easily developed if drag and wind are ignored, so we chose this scenario to test our approximation. We considered a particle with a speed of 10 m/s and a launch angle of $\pi/60$ radians. We calculated its trajectory using

$$x = (v_i \sin\phi)t, \qquad y = (v_i \cos\phi)t - \tfrac{1}{2}gt^2,$$

where

x is the position along the horizontal axis,

y is the position along the vertical axis,

v_i is the magnitude of the initial velocity,

t is time,

g is the acceleration of gravity, and

ϕ is the launch angle, measured from the vertical axis towards the horizontal.

We compared that trajectory with the one calculated Euler's method. The two were indistinguishable, showing that the Euler's method approximation results in virtually no error.

We also used the spreadsheet model to examine the effects of wind and drag on individual particle trajectories. **Figure 4** compares trajectories of particles with and without drag; and **Figure 5** compares the trajectories of two droplets, one with a 5 m/s wind and the other with no wind. Drag has a major effect and cannot be ignored.

Sensitivity

We tested the effect of changing some base factors in the model, using an initial water speed of 10 m/s. Fountain pool radii were chosen to highlight general trends in the data, either stability or sensitivity.

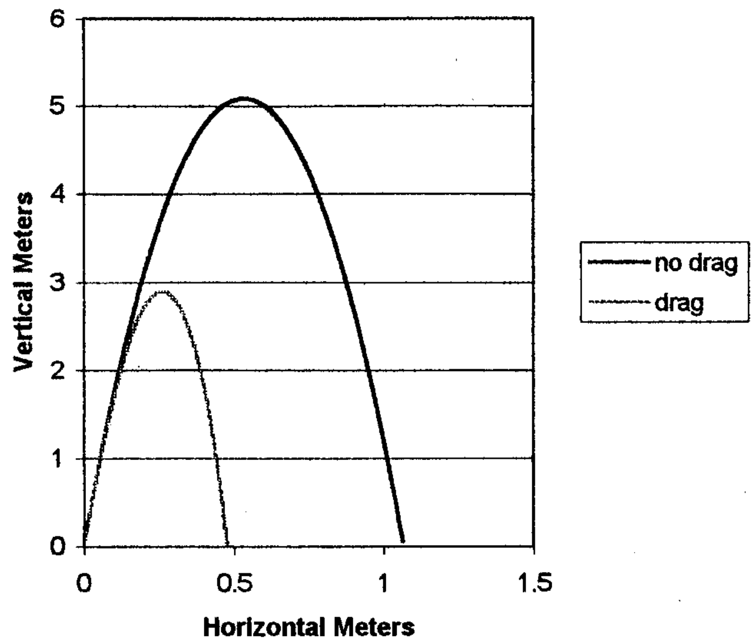

Figure 4. Droplet trajectories with and without drag.

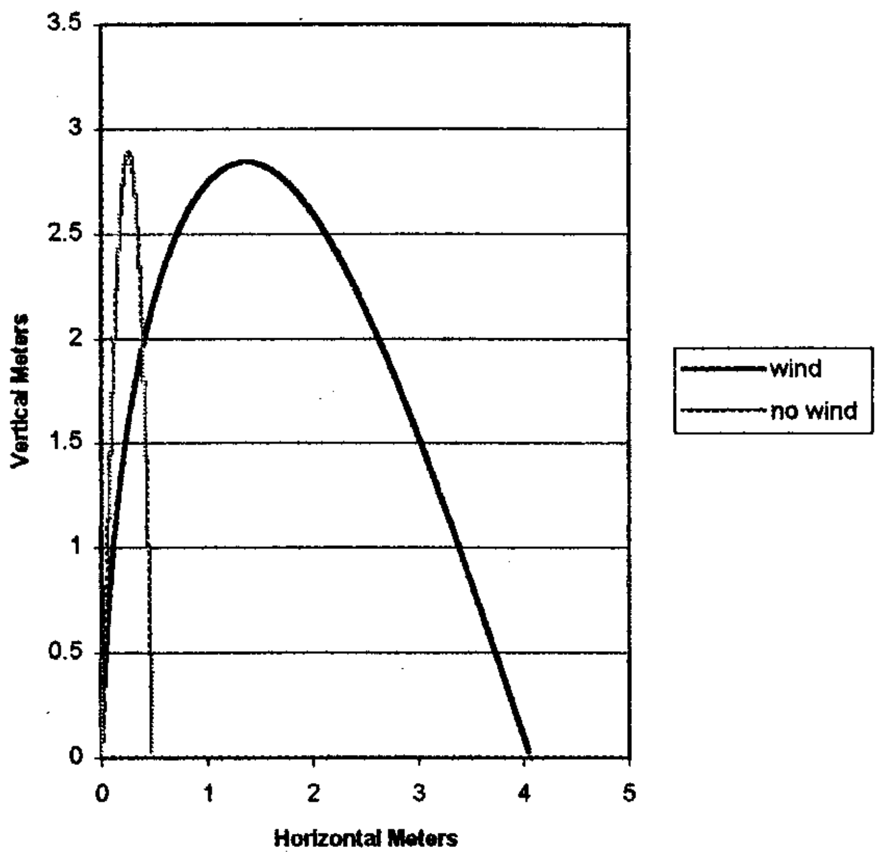

Figure 5. Droplet trajectories with and without wind.

Nozzle angle

We ran the simulation at a wind speed of 5 m/s with the nozzle tilted 0, $\pi/18$, $\pi/9$, or $\pi/6$ radians in the same direction as the wind vector. For a pool with a radius of 6 m, no water fell outside when the nozzle was pointed straight up and virtually none with a tilt of $\pi/18$ radians. With a tilt of $\pi/9$ radians, 47% of the water fell outside; for $\pi/6$ radians, 99.9% fell outside. The data suggest that tilting the nozzle into the wind could be used to prevent spillage.

Nozzle radius

With no wind and a pool radius of 2 m, virtually no water was spilled for nozzle radii of 0.25, 0.5, or 1 cm. With a 5 m/s wind, virtually all of the water was spilled at all three radii. The radius of the nozzle thus has virtually no effect on the percentage spilled, supporting our decision to use a percentage measure so as to allow the model to apply to fountains with different flow rates.

Water droplet size

In a fountain with a pool radius of 3.5 m, droplet radii of 0.75, 1.5, and 3 mm resulted in 94%, 53%, and 6% percent spillage. The sensitivity to droplet radius is a reflection of real-world behavior rather than a weakness of the model: Small particles, because of their low mass, are greatly affected by wind and drag.

Variability of launch angle

With a 3.5 m pool, a 5 m/s wind produced 15%, 45%, and 49% spillage for standard deviations of$\pi/180$, $\pi/20$, and $\pi/12$ radians. Thus, results are sensitive to the launch angles of the droplets, dictating that the angle be measured carefully before the model is used.

Strengths

As intended, the model controls the fountain height and volume according to conditions. It creates the tallest and therefore most interesting fountain possible while maintaining the set spillage level. At low spillage-level settings, no passersby get drenched nor is much water wasted.

The model is easy to adapt by changing parameters, including nozzle size, mean droplet size, mean launch angle and standard deviation, and mean droplet size and standard deviation.

Graphs of the droplets in midair show that the programmed fountain accurately depicts a real fountain.

Use of a table means that the radius or spill percentage can be changed without requiring recalculations. Since the control system does not do any calculation, it can respond almost instantaneously.

Weaknesses

A major problem occurs when wind speed increases quickly: Water droplets already emitted cannot be slowed down and will be carried away on the wind. However, any fountain system will suffer from this dilemma. To give the fountain a small buffer, the radius entered into the fountain control system can be set lower than the actual radius of the pool.

We model the wind as moving parallel to the ground with uniform speed. Real wind may vary with altitude and may blow from above or below the droplets. We also neglect wind turbulence.

We ignore droplet collisions. Some droplets may combine and then separate, causing slightly more splatter or mist; or the droplets' collisions may cause more of them to fall short of their expected trajectories, reducing spillage.

References

Dollins, Steven C. 2001. Handy mathematics facts for graphics. `http://www.cs.brown.edu/people/scd/facts` . Dated 6 November 2001; accessed 9 February 2002.

Halliday, David, Robert Resnick, and Jearl Walker. 1993. *Fundamentals of Physics*. 4th ed. New York: Wiley.

Lide, David R., ed. 1995. *CRC Handbook of Chemistry and Physics*. Boca Raton, FL: CRC Press.

Yates, Daniel, David Moore, and George McCabe. 1999. *The Practice of Statistics*. New York: W.H. Freeman.

The Fountain That Math Built

Alex McCauley
Josh Michener
Jadrian Miles
North Carolina School of Science and Mathematics
Durham, NC

Advisor: Daniel J. Teague

Introduction

We are presented with a fountain in the center of a large plaza, which we wish to be as attractive as possible but not to splash passersby on windy days. Our task is to design an algorithm that controls the flow rate of the fountain, given input from a nearby anemometer.

In calm weather, the fountain sprays out water at a steady rate. When the wind picks up, the flow should be attenuated so as to keep the water within the fountain's pool; in this way, we strike a balance between esthetics and comfort.

We consider the water stream from the fountain as a collection of different-sized droplets that initially leave the fountain nozzle in the shape of a perfect cylinder. This cylinder is broken into its component droplets by the wind, with smaller droplets carried farther. In the reference frame of the air, a droplet is moving through stationary air and experiencing a drag force as a result; since the air is moving with a constant velocity relative to the fountain, the force on the droplet is the same in either frame of reference.

Modeling this interaction as laminar flow, we arrive at equations for the drag forces. From these equations, we derive the acceleration of the droplet, which we integrate to find the equations of motion for the droplet. These allow us to find the time when the droplet hits the ground and—assuming that it lands at the very edge of the pool—the time when it reaches its maximum range from the horizontal position equation. Equating these and solving the initial flow rate, we arrive at an equation for the optimal flow rate at a given constant wind speed. Since the wind speeds are not constant, the algorithm must make its best prediction of wind speed and use current and previous wind speed measurements to damp out transient variations.

The UMAP Journal 23 (3) (2002) 221–234.

Our final solution is an algorithm that takes as its input a series of wind speed measurements and determines in real time the optimal flow rate to maximize the attractiveness of the fountain while avoiding splashing passersby excessively. Each iteration, it adds an inputted wind speed to a buffer of previous measurements. If the wind speed is increasing sufficiently, the last 0.5 s of the buffer are considered; otherwise, the last 1 s is. The algorithm computes a weighted average of these wind speeds, weighting the most recent value slightly more than the oldest value considered. It uses this weighted velocity average in the equation that predicts the optimal flow rate under constant wind. The result is the optimal flow rate under variable wind, knowing only current and previous wind speeds.

A list of relevant variables, constants, and parameters is in **Table 1**.

Table 1.

Relevant constants, variables, and parameters.

Physical constants	Description	Value
η_a	Viscosity of air	1.849×10^{-5} kg/m·s [Lide 1999]
ρ_w	Density of water	1000 kg/m^3
ρ_a	Density of air	1.2×10^{-6} kg/m^3
Situational constants		**Units**
A	Cross-sectional area of fountain nozzle	m^2
$f_{\max}$	Maximum flow rate of fountain's pump	m^3/s
R_p	Radius of fountain pool	m
r	Radius of smallest uncomfortable water droplet	m
dt	Sampling interval of anemometer	s
k	$k = 9\eta_a/2\rho_w r^2$	
Situational variables		
v_a	Instantaneous wind speed	m/s
f	Instantaneous flow rate of water from the fountain	m^3/s
n	$n = g/k + f/A$	m/s
Dynamic variables		
$x(t), y(t)$	Droplet's horizontal and vertical positions	m
$v_x(t), v_y(t)$	Droplet's horizontal and vertical speeds	m/s
$a_x(t), a_y(t)$	Droplet's horizontal and vertical accelerations	m/s^2
Situational parameters		
τ_d	Default sample wind velocity buffer time	s
τ_i	Buffer time for quickly increasing sample wind velocities	s
K	Weight constant	dimensionless

Assumptions

- Passersby find a higher spray more attractive.
- Avoiding discomfort is more important to passersby than the attractiveness of the fountain.
- The water stream can be considered a collection of spherical droplets, each of which has no initial horizontal component of velocity.
- Every possible size of sufficiently small water droplet is represented in the water stream in significant numbers.
- Water droplets remain spherical.
- The interaction between the water droplets and wind can be described as non-turbulent, or "laminar," flow.
- There exists a minimum uncomfortable water droplet size; passersby find it acceptable to be hit by any droplets below this size but by none above.
- When the wind enters the plaza, its velocity is entirely horizontal.
- The wind speed is the same throughout the plaza at any given time.
- The pool and the area around it are radially symmetric, so there is no preferred radial direction.
- We can neglect any buoyant force on the water due to the air, since the error introduced by this approximation is equal to the ratio of densities of the fluids involved, on the order of 10^{-3}, which is negligible.
- The anemometer reports wind speeds at discrete time intervals dt.

Analysis of the Problem

For a water stream viewed as a collection of small water droplets blown from a core stream, the interaction between the droplets and the air moving past them can best be described in the inertial reference frame of the moving air. In this frame, the air is stationary while the droplet moves horizontally through the air with a speed equal to the relative speed of the droplet and wind, $v_r = v_a - v_x$. In the vertical direction, $v_r = v_y$, since the wind blows horizontally.

In the air's frame of reference, the water droplet experiences a drag force opposing v_r. Assuming that the air moves at a constant velocity, this force is the same in both frames of reference. In the frame of the fountain, then, the droplet is being blown in the direction of the wind. The smaller water droplets are carried farther, so we need only consider the motion of the smallest uncomfortable water droplets, knowing that bigger droplets do not travel as far.

The water droplet initially has a vertical velocity $v_y(0)$ that is directly related to the flow rate of water through the nozzle of the fountain. This initial vertical velocity component can be controlled by changing the flow rate. The droplet's motion causes vertical air resistance, slowing the droplet and affecting how long (t_w) the droplet is in the air.

Since the vertical and horizontal components of a water droplet's motion are independent, t_w is determined solely by the vertical motion. Knowing this time allows us to find the horizontal distance traveled, which we wish to constrain to the radius of the pool.

When the wind is variable, however, we cannot determine exactly the ideal flow rate for any given time. We must instead act on the current reading but also rely on previous measurements of wind speed in order to restrain the model from reacting too severely to wind fluctuations. We need to react faster to increases in wind speed, since they result in splashing which is weighted more heavily.

Design of the Model

For our initial model, we assume that v_a is constant for time intervals on the order of t_w, so that any given droplet experiences a constant wind speed.

We model the water stream as a collection of droplets that are initially cohesive but are carried away at varying velocities by the wind. The distances that they travel depend on the wind speed v_a and the initial vertical velocity of the water stream through the nozzle, $v_y(0)$. Since the amount of water flowing through the nozzle per unit time is $f = v_y(0)A$, we have $v_y(0) = f/A$. The dynamics of the system, then, is fully determined by f and v_a. First, we find the equations of motion for the droplet.

Equations of Motion for a Droplet

For laminar flow, a spherical particle of radius r traveling with speed v through a fluid medium of viscosity η experiences a drag force F_D such that

$$F_D = (6\pi\eta r)v \qquad \text{[Winters 2002].}$$

Since a spherical water droplet has a mass given by

$$m = \rho_w \left(\tfrac{4}{3}\,\pi r^3\right),$$

the acceleration felt by the droplet is given by Newton's Second Law as the total force over mass. Since there are no other forces acting in the horizontal direction, the horizontal acceleration a_x is given by:

$$a_x(t) = \frac{d^2x}{dt^2} = \left(\frac{9\eta_a}{2\rho_w r^2}\right) v_r = k(v_a - v_x), \quad \textbf{(1)}$$

where $k = 9\eta_a/2\rho_w r^2$.

The droplet experiences both air drag and gravity in the vertical direction, so the vertical acceleration is

$$a_y(t) = -\left[\left(\frac{9\eta_a}{2\rho_w r^2}\right) v_y + g\right] = -k\left(v_y + \frac{g}{k}\right).$$

With constant v_a, we use separation of variables and integrate to find $v_x(t)$ and $v_y(t)$, using the facts that $v_x(0) = 0$ and $v_y(0) = f/A$. The results are

$$v_x(t) = v_a\left(1 - e^{-kt}\right), \qquad v_y(t) = ne^{-kt} - \frac{g}{k},$$

where $n = g/k + f/A$.

Integrating again, and using $x(0) = y(0) = 0$, we have

$$x(t) = \frac{v_a}{k}\left(kt + e^{-kt} - 1\right), \qquad y(t) = \frac{1}{k}\left(n\left(1 - e^{-kt}\right) - gt\right).$$

Determining the Flow Rate

Because f is the only parameter that the algorithm modifies, we wish to find the flow rate that would restrict the smallest uncomfortable water droplets to ranges within R_p, so that they would land in the fountain's pool.

After a time t_w, the droplet has fallen back to the ground. Thus, $y(t_w) = 0$. This equation is too difficult to solve exactly, so we use the series expansion for e^{-kt} and truncate after the quadratic term: $e^{-kt} \approx 1 - kt + (kt)^2/2$. Solving $y(t_w) = 0$, we find

$$t_w \approx \frac{2}{k}\left(1 - \frac{g}{nk}\right).$$

We know that the maximum horizontal distance $x(t_w)$ must be less than or equal to R_p, with equality holding for the smallest uncomfortable droplet. For that case, using the same expansion for e^{-kt} as above,

$$R_p = x(t_w) \approx \frac{v_a}{k}\left(kt_w - 1 + 1 - kt_w + \frac{(kt_w)^2}{2}\right) = \frac{v_a k}{2}\, t_w^2.$$

Solving for t_w and equating it to the earlier expression for t_w, we get

$$\sqrt{\frac{2R_p}{v_a k}} = t_w = \frac{2}{k}\left(1 - \frac{g}{nk}\right).$$

Recalling that in this equality only n is a function of f, we substitute for n and solve for f. The result is

$$f(v_a) = \frac{Ag}{\sqrt{\dfrac{2v_a k}{R_p} - k}}. \qquad \textbf{(2)}$$

As $v_a \to kR_p/2$, this equation becomes singular (see **Figure 2**). At lower values of v_a, it gives a negative flow rate. These wind speeds are very small; at such speeds, the droplets would not be deflected significantly by the wind. Since **(2)** assumes that the flow rate can be made arbitrarily high, it is unrealistic and invalid in application. To make the model more reasonable, we modify **(2)** to include the maximum flow rate achievable by the pump, $f_{\max}$:

$$F(v_a) = \begin{cases} \min\left(\dfrac{Ag}{\sqrt{\dfrac{2v_a k}{R_p} - k}}, f_{\max} \right), & v_a > kR_p/2; \\ f_{\max}, & v_a \le kR_p/2. \end{cases} \qquad \textbf{(3)}$$

An algorithm can use the given constants and a suitable minimal droplet size to determine the appropriate flow rate for a measured v_a. However, **(3)** assumes that the wind speed is constant over the time scale t_w for any given droplet. A more realistic model must take into account variable wind speed.

Variable Wind Speed

When wind speed varies with time, the physical reasoning used above becomes invalid, since the relative velocity of the reference frames is no longer constant. Mathematically, this is manifested in the equation for velocity-dependent horizontal acceleration; integrating is now not so simple, and we must resort to numerical means to find the equations of motion. Additionally, the algorithm can rely only on past and present wind data to find the appropriate flow rate. Our model needs to incorporate these wind data to make a reasonable prediction of the wind's velocity over the next t_w and determine an appropriate flow rate using **(3)**.

A *gust* is defined to be a sudden wind speed increase on the order 5 m/s that lasts for no more than 20 s; a *squall* is a similarly sudden wind speed increase that lasts longer [Weather Glossary 2002]. Our model should account for gusts and squalls, as well as for "reverse" gusts and squalls, in which the wind speed suddenly decreases. Since wind speeds can change drastically and unpredictably over the flight time of a droplet, our model will behave badly at times and there is no way to completely avoid this—only to minimize its effects.

The model's reaction to wind speed is not fully manifested until the droplet lands, after a time t_w (approximately 2 s). By the time our model has reacted to a gust or reverse gust, therefore, the wind speed has stopped changing. Without some type of buffer, in a gust our model would react by suddenly dropping flow rate as the wind peaked and then increasing it again as the wind decreased; the fountain would virtually cut off for the duration of any gust, which would release less water and thus seem very unattractive to passersby. Additionally, the water released just before the onset of the gust would be airborne as the wind speed picked up, splashing passersby regardless of any reaction by our model.

We exhibit an algorithm for analyzing wind data that makes use of **(3)**. Because velocity now varies within times on the order of t_w, we do not want to directly input the current wind speed but rather a buffered value, so that the model does not react too sharply to transient wind changes. The model should react more quickly to sudden increases in wind than to decreases, because increases cause splashing, which we weight more heavily than attractiveness.

The model, therefore, has two separate velocity buffer times: one, τ_d, the default, and another, τ_i, for when the wind increases drastically. We also weight more-recent values in the buffer more heavily, since we want the model to react promptly to wind speed changes but not to overreact. We weight each value in the velocity buffer with a constant value K plus a weight proportional to its age: Less-recent velocities are considered but given less weight than more recent ones. The weight of the oldest value in the buffer is K and that of the most recent is $K+1$, with a linear increase between the two. With the constraint that the weights are normalized (i.e., they sum to 1), the equation for the ith weight factor is

$$w_i = \frac{\left(K + \dfrac{i}{\tau - dt}\right) dt}{\left(K + \dfrac{1}{2}\right) \tau}.$$

The speeds are multiplied by their respective normalized weights and summed. This sum, v^*, is then used in **(3)** to find the appropriate flow rate for the fountain at a given time. We use τ_i rather than τ_d when the wind speed increases sufficiently over a recent interval, but not when it increases slightly or fluctuates rapidly. We switch from τ_d to τ_i whenever the wind speed increases over two successive 0.2 s intervals and by a total of at least 1 m/s over the entire 0.4 s interval.

Our algorithm follows the flowchart in **Figure 1** in computing the current flow rate. We wrote a C++ program to implement this algorithm, the code for which is included in an appendix. [EDITOR'S NOTE: We omit the code.]

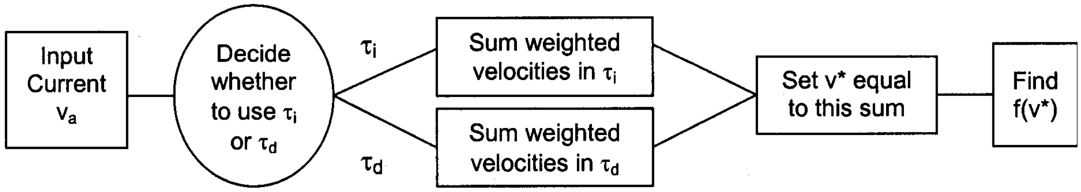

Figure 1. Flow chart for computing flow rate with variable wind speed.

Testing and Sensitivity Analysis

Sensitivity of Flow Equation

In our equation for flow rate, two variables can change: minimal droplet size and wind speed. While the minimal droplet size will not change dynamically, its value is a subjective choice that must be made by the owner of the fountain. The wind speed, however, will change dynamically throughout the problem, and the purpose of our model is to react to these changes.

We examined **(3)** for varying minimal drop sizes (**Figure 2**) and wind speeds (**Figure 3**). We used a fountain with nozzle radius 1 cm, maximum flow rate 7.5 L/s, and pool radius 1.2 m. (This maximum flow rate is chosen for illustrative purposes and is not reasonable for such a small fountain.)

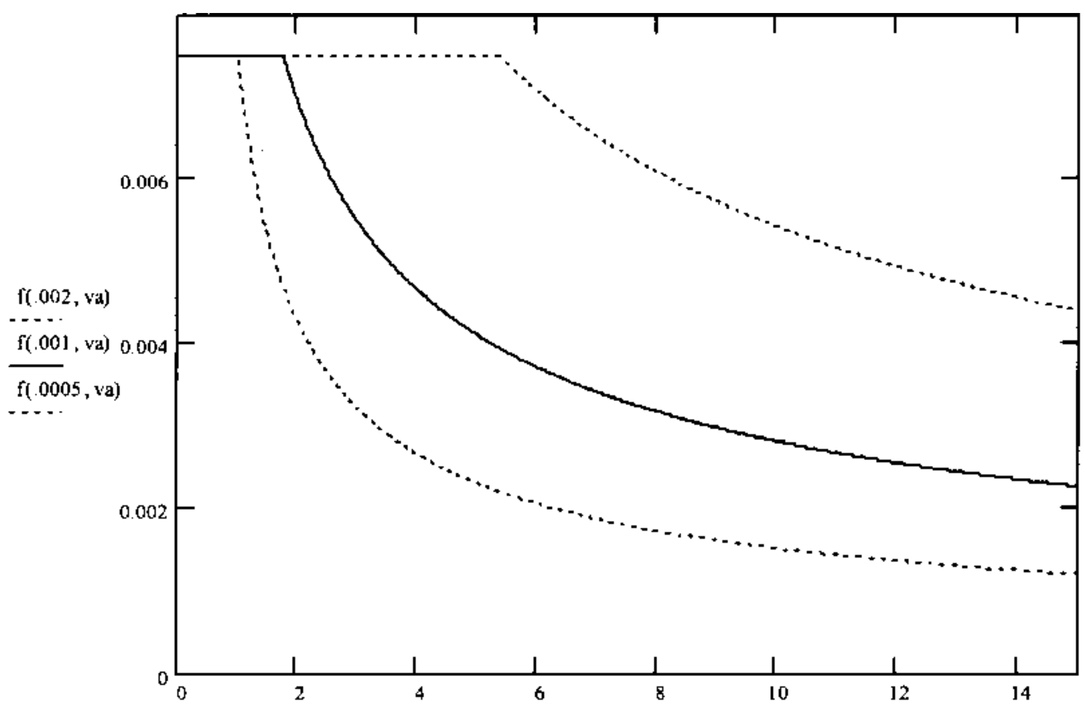

Figure 2. Graphs of flow rate f vs. wind speed v_a for several values of radius r of smallest uncomfortable droplet.

At any wind speed, as the acceptable droplet radius decreases, the flow rate decreases. At higher wind speeds, this difference is less pronounced; but at lower speeds, acceptable size has a significant impact on the flow rate. At very low wind speeds, the fountain cannot shoot the droplets high enough to allow the wind to carry them outside the pool, regardless of drop size. Our cutoff, $f_{\max}$, reflects that the fountain pump cannot generate the extreme flow needed to get the droplets to the edge of the pool in these conditions.

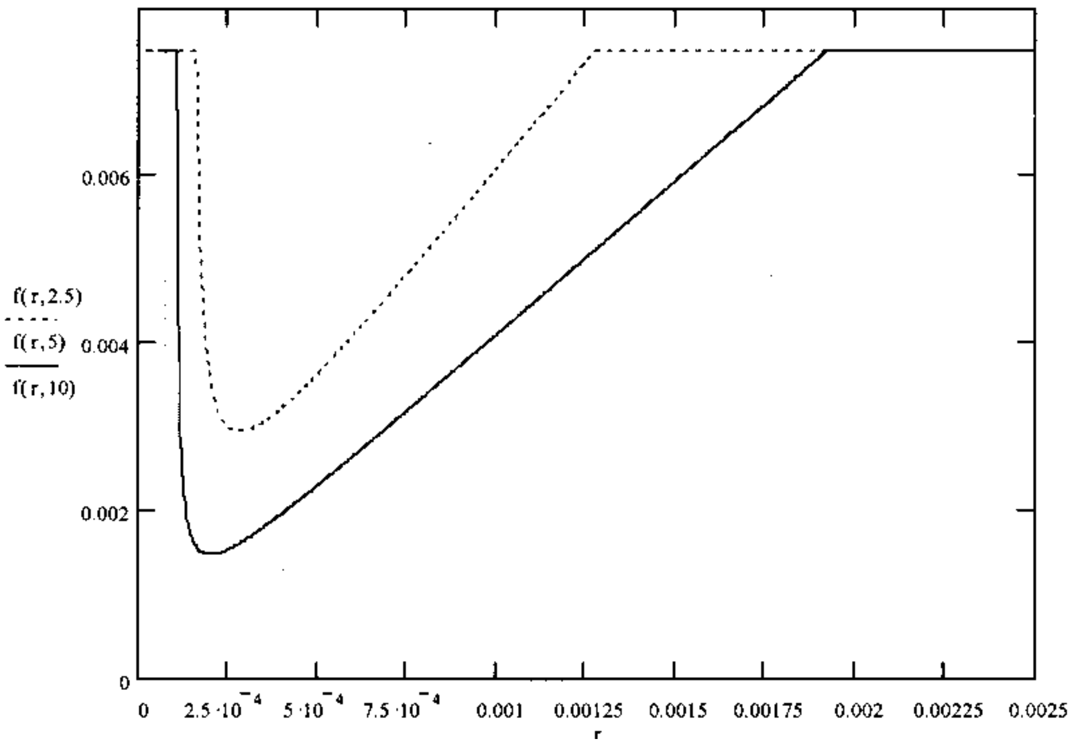

Figure 3. Graphs of flow rate f vs. radius r of smallest uncomfortable droplet for several values of wind speed v_a.

For any droplet size, as the wind speed increases, the flow rate must decrease to keep the droplets in the pool. For large r, a change in wind speed requires a greater absolute change in flow rate than for small r. For very small droplets, the drag force dominates the force of gravity, and an increase in flow also increases the drag force to such an extent that the particle spends no more time in the air. This behavior is readily apparent in **(1)** as r approaches zero. These extremely small values of r, though, describe droplets that are unlikely to discomfort passersby and thus are not significant to our model.

Sensitivity of Flow Algorithm

The results of the algorithm depend on the parameters τ_i, τ_d, and K, which determine the size of the buffer and weights of the velocities in the buffer. To test sensitivity to these parameters and to find reasonable values for them, we created the set of simulated wind speeds shown in **Figure 4**, including small random variations, on which to test our algorithm. This data set does not reflect typical wind patterns but includes a variety of extreme conditions.

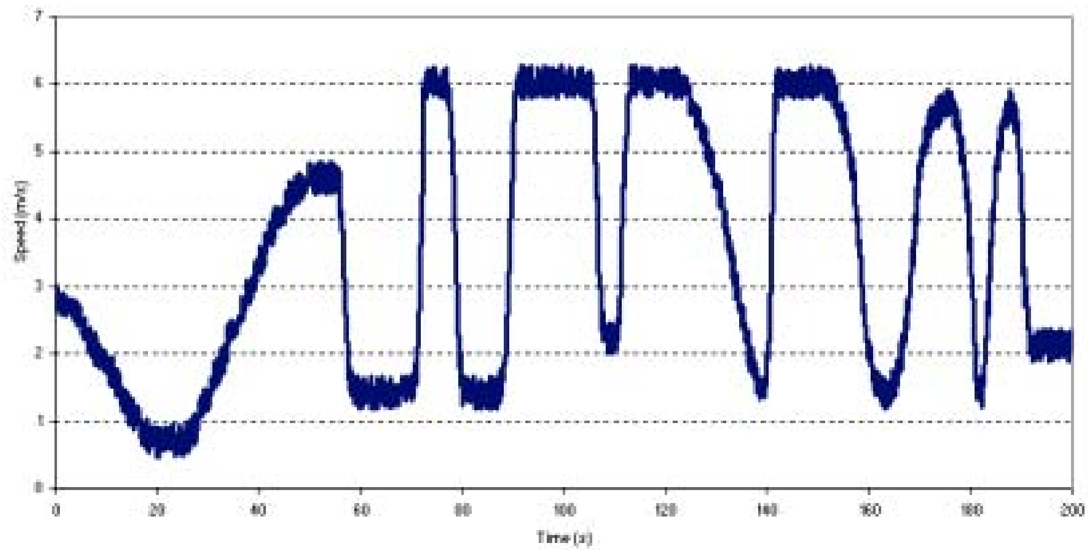

Figure 4. Simulation of wind speed for 3 min.

We wish to create a quantitative estimate of the deviation of our flow algorithm from ideal performance and then test the algorithm with different combinations of parameters to find the set that produces the smallest deviation under simulated wind conditions.

To measure how "bad" a set of flow choices is, we consider only the droplets that fall outside the pool. The "badness" is the sum over the run of the distances outside the pool at which droplets land.

To determine the distance, we need to know how droplets move through the air in varying wind speeds. Describing this motion in closed form is mathematically impossible without continuous wind data, so we approximate the equations of motion with an iterative process.

Since the time that a particle spends in the air, t_w, is not affected by the wind speed, we know t_w for each particle. We step through the time t_w in intervals of dt, computing the particle's acceleration, velocity, and position as

$$\begin{aligned} a_i &= k(v_{a,i} - v_i), & a_0 &= kv_{a,0}; \\ v_i &= v_{i-1} + a_{i-1}dt, & v_0 &= 0; \\ x_i &= x_{i-1} + v_i dt, & x_0 &= 0. \end{aligned}$$

When we reach t_w, the droplet has hit the ground, and we compare its horizontal position to the radius of the pool. We do this for each droplet, keeping track of both the largest absolute difference and the average difference.

To test the flow algorithm, we ran our program on the flow data with each combination of parameters. The parameter values that produced the least deviation were $\tau_i = 0.5$, $\tau_d = 1$, and $K = 10$. These values imply that only fairly recent wind speed measurements should be held in the buffer, with most recent velocity having a weight of $(K+1)/K = 1.1$ relative to the oldest. Lowering K beyond this value increases the deviation from the ideal, while increasing it further makes no difference. Similarly, increasing τ_i or τ_d increases the deviation, because the algorithm cannot respond quickly to changes in wind speed. Decreasing τ_i below 0.5 makes no difference, while decreasing τ_d would make the model too sensitive to short fluctuations in wind speed.

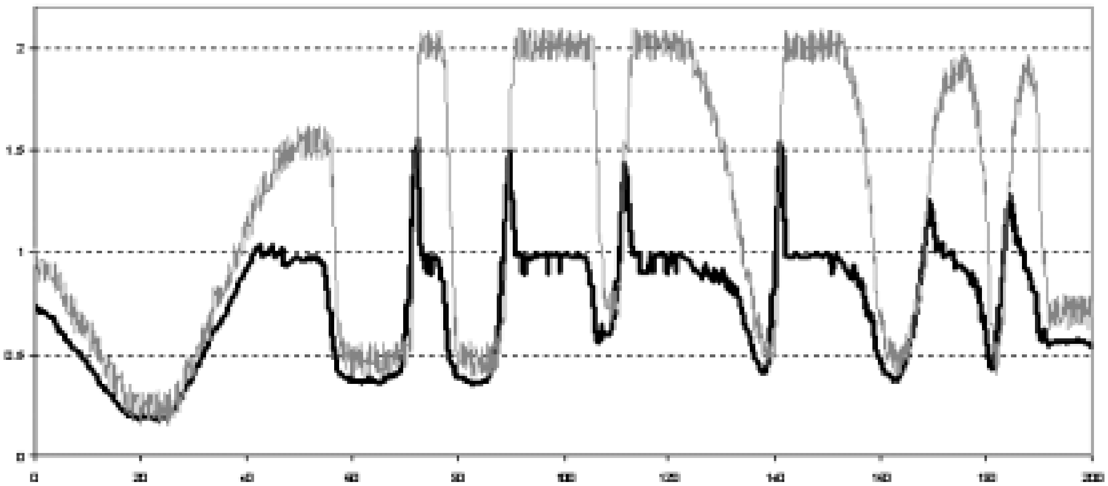

Figure 5. Range of droplets over the simulation overlaid with scaled wind speeds.

Justification

Validity of the Laminar Flow Assumption

Our model is based on a drag force proportional to v_r, which is not necessarily correct. For higher speeds or large droplet sizes, the drag becomes proportional to v_r^2. We thus need to determine whether reasonable physical scenarios allow us to model the drag force as proportional to v_r and not v_r^2.

For a sphere of radius r moving through the air with speed v_r, the Reynolds number R is defined to be

$$R = \frac{2\rho_a v_r}{\eta_a}\, r \qquad \text{[Winters 2002].}$$

When $R < 10^3$, there is little turbulence and laminar flow dominates, so air resistance is roughly proportional to v_r. If $R > 10^3$, the flow is turbulent and the drag force is proportional to v_r^2 [Winters 2002]. Using a physically reasonable relative speed of 4.5 m/s (corresponding to a wind speed of roughly 10 mph), we obtain $R = (5.8 \times 10^5)r$, which gives predominantly laminar flow when $r < 1.7$ mm. Because water droplets of diameter greater than 3 mm are uncomfortable, these provide an upper limit on the droplet sizes to consider. Because these smaller droplets bound the larger droplets in how far they go from the fountain (see below), all of our analysis is concerned with droplets whose sizes are within the allowed range for laminar flow.

Bounding the Droplet Range

For either laminar or turbulent flow, the acceleration due to drag scales with $F/m \propto r^{-n}$, where $1 \leq n \leq 2$. Larger droplets therefore experience a lower horizontal acceleration due to drag, while acceleration in the vertical direction is dominated by gravity ($k < 0.1g$); so the time that a particle spends in the air is roughly the same for droplets of varying radius. The heavier droplets have less horizontal acceleration, so they travel a shorter horizontal distance in the same amount of time than smaller droplets. The ranges are, therefore, shorter for larger droplets, so we can bound all uncomfortably-sized droplets by the range of the smallest such droplet.

Initial Shape of the Water Stream

We assume that the water coming out of the fountain nozzle has no initial horizontal velocity; that is, the stream is a perfect cylinder with the same radius as the nozzle. In fact, the stream is closer to the shape of a steep cone and the droplets have some horizontal velocity. In the absence of wind, this assumption has a significant impact on where the droplets land, since without wind the algorithm predicts a horizontal range of zero. However, in these cases, the flow rate is bounded by $f_{\max}$ regardless of initial velocity, so the natural spread of the fountain is irrelevant. In higher wind, the initial horizontal velocity is quickly dominated by the acceleration due to the wind and thus makes a negligible contribution to the total range.

Exclusively Horizontal Wind

We assume that the wind is exclusively horizontal. Since the anemometer measures only horizontal wind speed, that is the only component that we can consider in our model. Additionally, the buildings around the plaza would tend to act as a wind tunnel and channel the wind horizontally.

Quadratic Approximation of e^{-kt}

Because the series for e^{-kt} is alternating, the error from truncating after the second term is no greater than the third term, which is $(kt)^3/6$. The relative error is $(kt)^3/6e^{-kt} \approx 0.001$ for reasonable values of k and t, so our approximation introduces very little error.

Conclusions

Our final solution is an algorithm that takes as its input a series of wind speed measurements and determines in real-time the optimal flow rate to maximize the attractiveness of the fountain while avoiding splashing passersby excessively. It takes an inputted wind speed and adds it to a buffer of previous measurements. If the wind speed is increasing sufficiently, the last 0.5 s of the buffer are considered; otherwise, the last 1 s is. The algorithm computes a weighted average of these wind speeds, weighting the most recent value 10% more heavily than the oldest value considered. It then takes this weighted average and uses it in the equation that predicts the optimal flow rate under constant wind. The result is the optimal flow rate under variable wind, knowing only current and previous wind speeds.

Strengths and Weaknesses

Strengths

- Given reasonable values for the characteristics of the fountain and for wind behavior, our model returns values that satisfy the goal of maintaining an attractive fountain without excessively splashing passersby.
- The model can compute optimal flow rates in real time. Running one cycle of the algorithm takes a time on the order of 0.001 s, so the fountain's pump could be adjusted as fast as physically possible.
- The values for the parameters that determine the behavior of the algorithm, τ_d, τ_i, and K, are not arbitrary but instead are the values that perform best under simulation.
- Our algorithm is very robust; it works well under extreme conditions and can be readily modified for different situations or fountains.

Weaknesses

- A primary assumption is that the droplets coming from the fountain nozzle have no horizontal velocity. In reality, the nozzle sprays a cone of water, rather than a perfect cylinder; but this difference does not have a significant impact on the results.
- Another important assumption is laminar flow. The water droplets are of a size to experience a combination of laminar and turbulent flow, but describing such a combination of regimes is mathematically difficult and is known only through experimentation. A more rigorous representation of the drag force would increase the accuracy of our simulation, but doing so would

markedly increase the complexity of the algorithm and thus make real-time computation more difficult.

- We have ignored the abundances of droplet sizes in considering discomfort. If one droplet would spray passersby, we assume that enough droplets would spray passersby to make them uncomfortable. In fact, it is only significant numbers of droplets that discomfort passersby; but we do not know how many droplets would be released nor how many would be needed to be discomforting.

References

Goldstein, S., ed. 1965. *Modern Developments in Fluid Dynamics*. Vol. 2. New York: Dover.

Hughes, W.F., and J.A. Brighton. 1967. *Fluid Dynamics*. New York: McGraw-Hill.

Lide, David R., ed. 1999. *CRC Handbook of Chemistry and Physics*. 80th ed. Boca Raton, FL: CRC Press.

Winters, L. 2002. Theory of velocity dependent drag forces. `http://courses.ncssm.edu/ph220/labs/vlab1/theory.pdf` . Accessed February 2002.

Weather Glossary. 2002. `http://www.weather.com/glossary/` . Accessed February 2002.

Wind and Waterspray

Tate Jarrow
Colin Landon
Mike Powell
U.S. Military Academy
West Point, NY

Advisor: David Sanders

Introduction

Given anemometer readings from a nearby building, the task is to devise an algorithm that controls the height of a fountain in an open square. Our mission is to keep passersby dry and yet have the fountain look as impressive as possible. With ever-changing winds, we must devise a scheme to regulate the flow of water through the fountain to ensure that the bulk of the water shot into the air falls back to the ground within the fountain basin boundary.

Our model considers many factors and is divided into five basic parts:

- The conversion of wind speed on top of the building to wind speed at ground level based on height and the force of drag.
- The determination of initial velocity, maximum height, and time of flight from fountain nozzle characteristics, using Bernoulli's equation and the rate of flow equation of continuity.
- The assessment of the displacement effects of the wind on the water's ascent.
- The assessment of the displacement effects of the wind on the water's descent.
- The calculation of the optimal flow rate by comparing the water's total horizontal displacement to the radius of the fountain basin.

After creating this model in a MathCAD worksheet, we solved every function involved in this model as a function of the water flow rate. This worksheet takes the input from several variables such as the nozzle radius, the maximum flow rate the fountain can handle, the dimensions of the building on which the

The UMAP Journal 23 (3) (2002) 235–250.

anemometer is placed, and the dimensions of the fountain. From the inputs, the model finds the maximum flow rate that keeps the water in the fountain basin. As wind speed and direction vary, the model reacts to produce the optimal flow rate.

Testing the model shows that while the results are reasonable, the main source of error results from our drag calculations due to the interaction between wind and the buildings. To solve this error, measurements should be taken at both the building roof and the fountain itself. Although future work would resolve this issue and improve the model, our current model still provides realistic results.

We provide in **Table 1** a list of symbols used.

Problem Approach

We break the overall problem down into several smaller pieces, solve the pieces separately, and put the pieces together to find the overall solution.

- How the wind is affected as it flows around the buildings.
  - How the wind varies with height off the ground.
  - How the buildings slow the wind.
- How the wind affects the water from the fountain.
  - How the wind affects the water on the way up.
  - How the wind affects the water on the way down.
  - How to contain that total displacement within the basin.

Assumptions

Overall Assumptions

- **The plaza has a fountain in the center with four surrounding buildings.** Other arrangements can be handled with slight modifications.
- **The buildings are rectangular and have the same dimensions.** Most buildings are rectangular; for same-size buildings, we can use a single constant drag coefficient.
- **The distances from each building to the fountain are the same,** so each building has the same effect on the fountain water.
- **The acceptable splash area is the radius of the fountain basin.** A basin surrounds the water jet, and people walking outside the fountain do not want to get wet.

Table 1.
Table of symbols.

Symbol	Meaning (units)
R	rate of flow of the fountain (m^3/s)
R_e	Reynolds number
v	flow speed (m/s)
d	a relevant dimension (m)
ν	kinematic viscosity of the fluid
$F_{\vec{D}}$	force of drag (N)
ρ	density of the wind (kg/m^3)
v_{bh}	speed of the wind before the building at height h (m/s)
C_d	drag coefficient
A	surface area interacting with the wind (m^2)
v_z	wind speed measured by the anemometer at the height z (m)
h	height above ground (variable) (m)
z	height of the building (m)
α	terrain constant number = 0.105
$h_{\max}$	maximum height that the water reaches, a function of R (m)
K_i	kinetic energy of the wind-building system before the wind hits the building (J)
K_f	kinetic energy of the wind-building system after the wind passes the building (J)
W_{NC}	work done by nonconservative forces, drag of the building times the length over which it is applied (J)
$\vec{d}$	distance over which drag acts, length and width of the building (m)
b	width or half the length of one of the buildings (m)
v_h	speed of the wind after it passes the building at a height h (m/s)
m	mass of the air that interacts with the building in 1 s if the speed v_{bh} was constant over the face of the building (kg)
θ	angle at which the wind strikes the building (°)
A_p	cross-sectional area of the pipe at the nozzle tip (m^2)
v_f	speed of the water as it leaves the nozzle (m/s)
r_p	radius of the pipe at the nozzle tip (m)
g	acceleration due to gravity, 9.803 m/s^2
r_c	radius of the column of water at a time t after leaving the nozzle with a rate of flow R (m)
P	pressure on the water caused by the wind (N/s^2)
A_c	surface area of the column of ascending water(m^2)
ρ	density of air (kg/m^3)
m_T	total mass of the water in the air at a flow rate R (kg)
T_{Total}	total time that the water spends in the air with a flow rate R (s)
ρ_{water}	density of water (kg/m^3)
a_c	horizontal acceleration of the water in the column with a flow rate of R and a wind of speed v_h (m/s^2)
F_c	force on the column of water from the wind of speed v_h (N)
x_c	horizontal displacement of the ascending column of water with a flow rate R and wind speed v_h at a time t (m)
P_D	pressure on a drop of water from wind of speed v_h (N/m^2)
F_d	force on the drop from wind of speed v_h (N))
A_d	area of a drop (m^2)
m_d	mass of a drop of water (kg)
a_d	horizontal acceleration of the drop of water as a function of rate of flow R and time in air t (m/s^2)
a_{avg}	average horizontal acceleration of a drop during its descent at a rate of flow R and wind of speed v_h (m/s^2)

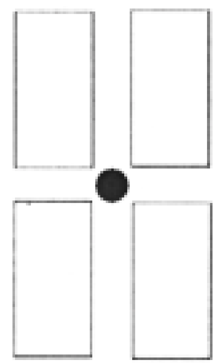

Figure 1. The fountain in the center of four buildings.

- **The fountain does not squirt water higher than the buildings,** although shooting water over the roofs would indeed be spectacular.
- **The fountain shoots water straight into the air.** This is important for our model so that we can predict how the water will flow up, how it will fall, and where it will fall.
- **The fountain nozzle creates a single sustained stream of water.** This assumption enables us to neglect drag as the water reaches its peak height. Furthermore, most fountains have a continuous flow of water.

Wind

- **The pertinent wind flow is around the sides of the buildings, not over them.** Since the fountain does not exceed the height of the buildings, it does not interact with wind that passes over the tops of the buildings. This assumption is important in calculating the drag caused by the buildings.
- **The flow of the wind continues in the same direction across the entire plaza.** The wind flows through the plaza in a constant direction, goes around obstacles, and resumes the same direction of motion. The wind does not get stuck in the plaza nor react to cars, people, doors, or windows in the plaza.
- **Wakes caused by buildings are not factors.** The wake that results when wind hits a building and goes around it does not change the velocity after the wake, so the wake force does not influence the wind's speed or direction.
- **The fountain is not in the wake of the buildings.** With this assumption, there is no need to worry about wake in our model. This is important because wake is too complex to be modeled.
- **The change in wind velocity is due solely to drag.** The reason that the wind decreases before and after hitting the building is because of drag. This assumption allows us to use the law of conservation of energy to predict the change in velocity.

- **The anemometer measures wind speed and direction at the top of the building before any effects of drag.** The anemometer must be at the top of the building on the windward side, elevated above the height of the building so as not to measure any of the effects of the building. To simplify, we assume that it is at the height of the building.
- **The wind pattern is the same across the entire plaza as measured at the anemometer.** If the pattern changed, the anemometer reading would be invalid.
- **The fountain is in a city or urban area.** This assumption allows us to determine the effect of the ground on wind speed at a given height.
- **The drag applied to wind at a certain height is equal to the average effect of drag,** that is, to the total drag caused by the building at the velocity at that height divided by the height of the building. This is slightly inaccurate but still produces a reasonable model.

Water Height

- **Water has laminar flow.** Water has a constant velocity at any fixed point, regardless of the time. A fluid may actually have various internal flows that complicate the model, but we consider the flow as the jet of water ascends to be constant so that we can model it as an ideal fluid.
- **Water has nonviscous flow.** The water experiences no viscous drag force in the pipe or in the air. The outer edge of the column of water actually interacts with the air and loses some energy to due to the viscosity of both fluids; but since air and water both have a low viscosity, this loss is negligible.
- **Water is incompressible.** The density of water is constant and does not change as the water moves up into the air and back down again.

Water Movement Sideways

- **The water jet upward flows as a cylinder.** Since the surface tension of the water holds it together unless it is acted upon by a force, the water should retain the dimensions of the nozzle from which it emerges.
- **The pressure of the wind is a force per area on the water column and on water drops.** Wind and water are both fluids, so the interaction between them is a complex relationship of their viscosities; however, we also know that wind creates a pressure difference that we can model. We model the force on the water as the pressure caused by a certain velocity of wind multiplied by the surface area of the body of water.

- **The largest particle of water that we want to contain is the size of average drop of water 0.05 mL.** The column of water breaks into smaller particles at the peak of its ascent, and they descend individually. We estimate that particles smaller than that size would be acceptable to bystanders hit by them. Any larger particle would have more mass, hence a higher mass-to-surface-area ratio, so the pressure could not push it as far.
- **Water drop behaves as a rigid body.** Since a drop is small, internal currents have very little effect. Additionally, the pressure acts over the entire surface area of the drop and should accelerate it as a single body.

Model Design

Effects of Buildings on Wind Velocity

Because buildings surround the fountain, the wind velocity at the anemometer on top of a building is different from that at fountain level. Buildings disrupt wind currents, slow the wind, and change its direction [Liu 1991, 62]. Buildings create areas of increased turbulence, as well as a wake—an area of decreased pressure—behind the building. Thus, the behavior of wind after it passes a building is so complex as to be almost impossible to model. Hence, we assume that the fountain is located outside of the wakes of the buildings.

Wind Speed Reduction

The wind inside a group of buildings is less than that outside of the group; the interaction between the wind and the buildings causes a decrease in speed. The drag between the building and the wind decreases the kinetic energy of the wind and hence its speed.

Since the fountain is squirting water into the air in a symmetrical shape, the wind affects where the water lands in the same way regardless of the wind's direction; so there is no need to find the wind direction after it hits the building.

Drag

Nevertheless, wind direction before the wind hits the building is an important factor. The angle at which the wind hits the building changes the surface area that the wind interacts with, and drag changes with area. The drag force $\vec{F}_d$ is given by

$$\vec{F}_d = \tfrac{1}{2}\rho v_{\text{bh}}^2 C_d A,$$

where ρ is the density of air, v_{bh} is the speed of wind at height h, C_d is the drag coefficient, and A is the surface area interacting with the wind. Therefore, we must know from which angle the wind approaches the building and how this affects the surface area perpendicular to the direction of the wind.

For a rectangular building with the narrow face to the wind, $C_d = 1.4$ [Macdonald 1975, 80].

Figure 2 diagrams the plaza and fountain. No matter which way the wind blows, it interacts with a narrow edge of a building. Wind from due east or west creates a problem for this model because of discontinuity in the the drag coefficient. Instead, we assume that the coefficient remains constant.

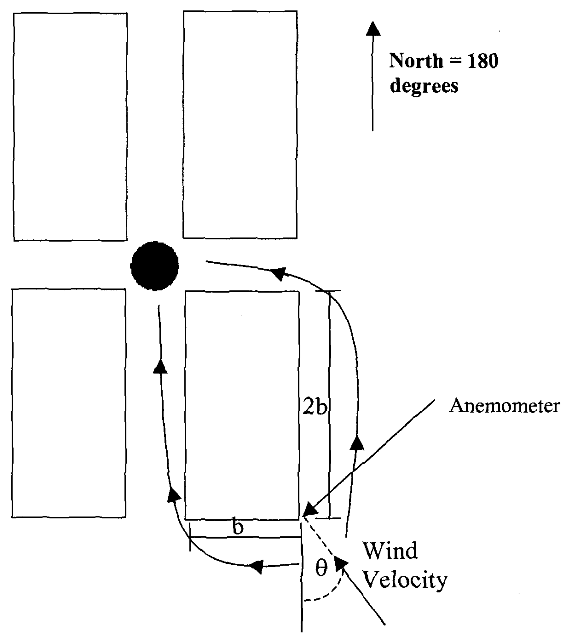

Figure 2. The plaza.

Wind Speed at Differing Heights

The speed of wind changes with the height from the ground because there is an additional force on the wind due to surface friction (dependent on the surface characteristics of the ground). The effect of this friction decreases as the wind speed is measured at a greater distance from the ground, creating faster speeds at greater heights.

Wind speed also varies because the temperature varies with height and location. However, if we assume that temperature and ground roughness are constant, a mean speed at a certain height can be modeled by

$$v_{\mathrm{bh}} = v_z \left(\frac{h}{z}\right)^{\alpha} \qquad \text{[Macdonald 1975, 47],} \qquad \textbf{(1)}$$

where v_{bh} is the speed of the wind before it hits the building, v_z is the wind speed measured by the anemometer at the height z of the building, h is the variable height of the water, and α is the terrain constant number. We use $\alpha = 0.105$, the value for ground roughness of a city center [Macdonald 1975, 48].

We assume that the greatest height of the water that the fountain hits, $h_{\max}$, does not exceed the height of the building, so we can neglect the drag from the building's roof (since the wind that goes over the building does not interact with or affect the water in the fountain).

Converting Drag to Work

We need to convert the drag force into a form that will enable us to determine the actual loss of speed. Since drag is a nonconservative force (energy is lost during its application), we can use conservation of energy in the form that says that the initial kinetic energy K_i minus the work W_{NC} done by the nonconservative force equals the final kinetic energy K_f, or

$$K_i = K_f + W_{\text{NC}}. \tag{2}$$

For the K terms, we use the kinetic energy equation $K = \frac{1}{2}mv^2$. For K_i, we have v_{bh}; for K_f, we have v_h.

Work is the dot product of the force and the distance that the force is in contact with the surface, or

$$W_{\text{NC}} = \vec{F}_d \cdot \vec{d}.$$

The work done is the drag force exerted by the building on the wind multiplied by the distance that the wind travels along the sides of the building.

With substitution, we find

$$W_{\text{NC}} = \tfrac{1}{2}\rho v_{\text{bh}}^2 C_d A d. \tag{3}$$

The drag coefficient C_d is for the entire building. However, we cannot have the entire building's drag force working on the speed at a specific height or we will overestimate the influence of the drag. Instead, we find the average drag per meter of the building. To do this, we divide **(3)** by the height z of the building, then substitute the result into **(2)**:

$$\tfrac{1}{2}mv_{\text{bh}}^2 = \tfrac{1}{2}mv_h^2 + \frac{\frac{1}{2}\rho v_{\text{bh}}^2 C_d A d}{z}.$$

Using **(1)**, we can find v_{bh} at any height h; but the equation still has several unknowns that stop us from solving for v_h: the mass m, the area A, and the distance d.

Mass of Air

The mass of wind that interacts with the building per second at height h is

$$m = v_{\text{bh}} A \rho t.$$

It is reasonable for convenience to use the average mass over 1 s.

Surface Area Interacting with Wind

As shown in **Figure 3**, the surface area as it relates to the drag due to wind is the cross section of the building perpendicular to the wind.

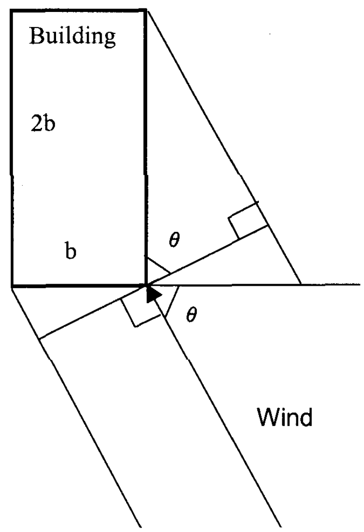

Figure 3. Orientation of wind to building.

Therefore, the surface area of the building based on the angle θ at which the wind strikes the building of width b is found using trigonometry and gives

$$A = (b|\cos\theta| + 2b|\sin\theta|)z,$$

where z is the height of the building. We take the absolute value of the cosine and sine because we use the direction of the wind measured by the anemometer in terms of a 360° compass.

Distance

The distance d that the wind goes over the building is $3b$, the length of one side plus the width of the building, because the wind will curve around the building.

Combining the Equations

Combining, solving for v_h, and using $\alpha = 0.105$ gives the speed v_h at height h. [EDITOR'S NOTE: We do not reproduce the complicated expression.]

Height of the Fountain

We find a function for the maximum height $h_{\max}(R)$ of the fountain in terms of the rate of flow R. We assume that the water acts as an ideal fluid and that the fountain shoots water straight into the air in a single sustained stream.

Volume Flow Rate and Bernoulli's Equation

We have from Halliday et al. [2001, 334]

$$R = A_p v_f, \quad \text{or} \quad v_f(R) = \frac{R}{A_p} = \frac{R}{\pi r_p^2},$$

where R is the rate of flow, v_f is its speed, A is the cross-sectional area of the pipe, and r_p is the radius of the pipe.

Based on the effect that we want the fountain to have, we make the water column (the radius of the pipe at the tip of the nozzle) have a 6-cm diameter, hence a radius of 0.03 m.

We use Bernoulli's equation [Halliday et al. 2001, 336], which relates forms of energy in a fluid, to calculate the maximum height of the water as it shoots into the air:

$$p_1 + \tfrac{1}{2}\rho\nu_1^2 + \rho g y_1 = p_2 + \tfrac{1}{2}\rho\nu_2^2 + \rho g y_2,$$

where p_1 and p_2 are the pressure of the water (both are zero since we are looking only at the water in the air) and g is the acceleration due to gravity. At the initial point, we consider the height of the nozzle as having zero gravitational potential energy, so the pressure head $\rho g y_1$ equals zero. Additionally, the speed v_1 is the speed from **(1)**. At the endpoint, the water has height $h_{\max}$ and the kinetic energy is zero. Substituting and simplifying gives

$$h_{\max}(R) = \frac{\left(\dfrac{R}{\pi r_p^2}\right)^2}{2g}.$$

With the radius r_p constant, the height of the top of the water stream varies directly with the square of the rate of flow R. **Figure 4** shows the heights for values of R between 0 and 0.04 m^2/s of water. Whatever mechanism pumps the water must be able to vary the flow rate by small amounts, particularly for large R, to maintain the maximum height allowable for the wind conditions.

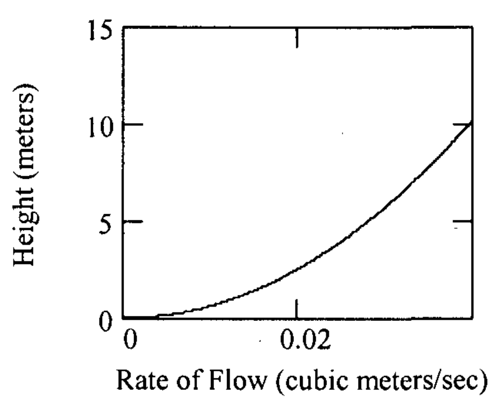

Figure 4. The effect of rate of flow on height of the fountain.

The Effect of Wind on the Water Ascent

Radius Change in Ascent

Photos of fountains show that the water ascends as a slowly widening column until it reaches its maximum height, then falls back on itself and scatters. We can derive an expression that shows the change in the radius as the cylinder of water ascends; but since the change is very small, on the order of 1 mm, we use the initial radius at the nozzle, r_p, in our calculations.

Wind Effects in Ascent

The other contributor to the water's horizontal movement is the wind, whose force can be determined from pressure. Pressure is force exerted over an area, so pressure multiplied by area gives the force:

$$P = \frac{F_c}{A_c},$$

where P is pressure, F_c is force, and A_c is area. The cylinder of water has height $h_{\max}$ and width twice the radius r_p, so $A_c = 2r_p h_{\max}$.

We find pressure in terms of wind speed using

$$P = \tfrac{1}{2}\rho v_h^2,$$

where ρ is the density of air and v_h is the speed of wind at height h. So we have

$$\tfrac{1}{2}\rho v_h^2 = \frac{F_c}{A_c} = \frac{F_c}{2r_p h_{\max}}.$$

Solving for F_c gives

$$F_c(R) = \rho r_p h_{\max} v_h^2 = \rho r_p \frac{\left(\dfrac{R^2}{\pi r_p^2}\right)^2}{2g} v_h^2.$$

Since we have the F from $F = ma$, finding mass should lead us to acceleration. Discovering the mass equation is pleasantly simple. If you take the flow rate and multiply it by the time that the water is in the air, then you know exactly how much water is in the air. Multiplying that volume by water's density gives the total mass, m_T:

$$m_T(R) = R\,\frac{t_{\text{total}}(R)}{2}\,\rho_{\text{water}},$$

where $\rho_{\text{water}} = 1000$ kg/m^3. We solve for the acceleration:

$$a_c(R) = \frac{F_c(R)}{m_T(R)}.$$

We use kinematics yet again to find how far the center of the water cylinder shifts, x_c, by the time it reaches the top of its ascent:

$$x_c(R) = \tfrac{1}{2} a_c(R) \left(\frac{t_{\text{total}}(R)}{2} \right)^2 .$$

The Effect of Wind on Water Descent

The water's surface tension holds it in a very cylinder-like column during its accent; but when the water reaches the top of its path, it runs out of kinetic energy and begins falling. Modeling the erratic behavior of the fall is somewhat difficult. At that point, turbulence caused by the competition between the gravitational force and the momentum of the ascending water overcomes the water's surface tension and smaller bodies of water descend individually.

Since we are concerned about the bystanders level of dryness, we want the fountain to shoot to a height that will keep within the fountain's basin all particles with potential to dampen the onlookers. Anything smaller than a drop from a common eyedropper, about 0.05 mL, will not considerably moisten a person. Therefore, we want to spray the fountain to a height that will not let the wind carry this size drop outside the fountain's basin.

Anything larger than such a drop has a greater mass-to-surface-area ratio, so it does not accelerate as much nor travel as far. So, we need model the flight of only such a drop.

We assume that the drop behaves as a rigid body. This assumption neglects the forces that act internally in the fluid and thereby overestimates the effect of the wind. Thus, this assumption may lower the maximum height of the water but will not result in any excessive water hitting bystanders.

Since the drop behaves as a rigid body, Newton's second law applies: The sum of all the external forces equals the mass of the drop of water, m_d, times the net acceleration:

$$\sum \vec{F}_d = m_d \vec{a}_d.$$

The only significant forces are the the wind (parallel to the ground) and gravity (vertical). The wind force F_d we know from $P = F_d/A_d$; we can calculate the surface area of the drop, and we know the pressure of the wind at height h. We can calculate the force on the drop as a function of height. Dividing that by the mass of the droplet gives its acceleration a_d parallel to the ground as a function of height h:

$$a_d(h) = \frac{P(h) A S_d}{m_d}.$$

Unfortunately this acceleration depends on height, which is a function of time t in the air and the rate of flow R; so we cannot use the constant-acceleration kinematics equations. Also, the nature of the equation for acceleration makes integration with respect to time an unwieldy task. We can, however, get the average acceleration by integrating the acceleration from the time at the peak to the total time in the air and dividing by half of the time in air:

$$a_{\text{avg}} = \frac{\int_{t_{\text{total}}/2}^{t_{\text{total}}} a_d\big(h(t,R)\big)\,dt}{t_{\text{total}}/2}.$$

Using a_{avg} as a constant, we can find the displacement x_d of the drop in the horizontal direction. We know that

$$x - x_0 = v_0 t + \tfrac{1}{2} a t^2.$$

Applying this equation to the motion of the drop, we see that

$$x_d(R,t) = \left[\frac{d}{dt} r\left(R, t_{\text{total}}/2\right) + \frac{d}{dt} x_c\left(R, t_{\text{total}}/2\right)\right] t + \frac{1}{2} a_{\text{avg}}(R) \left(\frac{t}{2}\right)^2 + x_c\left(R, t_{\text{total}}/2\right) + r_p.$$

Combined with the y position of the droplet, we get the flight path of **Figure 5**.

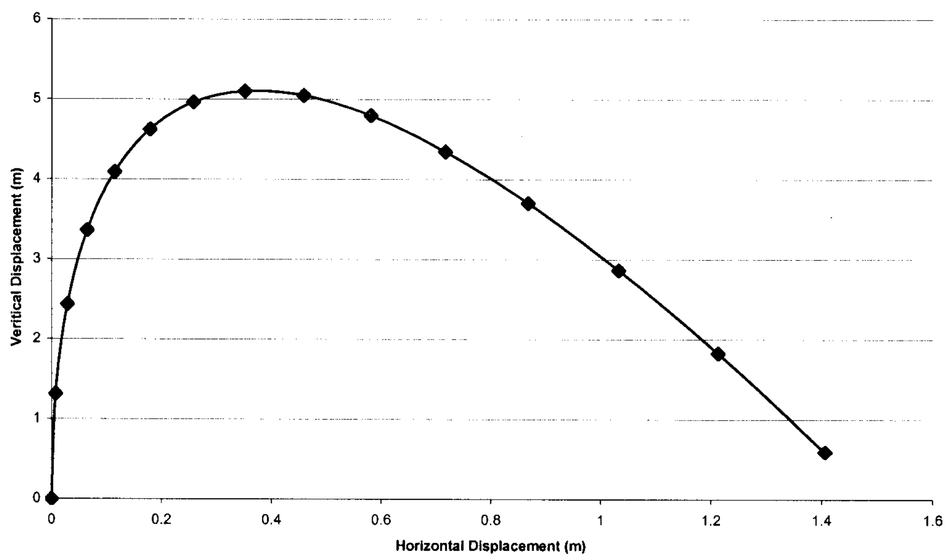

Figure 5. Water path from the fountain due to wind.

The initial speed is the derivative of two position functions, $r_c(R,t)$ (the radius of the column) and $x_c(R,t)$ (the displacement of the column due to the wind). Taking the rate that those distances change when the drop separates from the column gives the initial speed of the drop. In addition, the equation shows that the initial displacement, x_0, is the original width of the column plus the distance that the wind pushes the center of the cylinder, $x_c(R,t)$ for the time t that it takes for a particle of water to reach its peak, $t_{\text{total}}/2$.

The displacement depends on the rate of flow. This is useful, since we must moderate the rate of flow to control the amount of water that escapes the basin. We can now set the displacement $x_d(R,t)$ equal to the maximum allowable displacement—the radius of the basin—and solve for the rate of flow.

The Optimal Rate of Flow

Our computer algebra system choked on solving for R exactly in terms of the other parameters. Instead, we adapted an incremental approach with a simple program in MathCAD. The maximum value for R could be anything; but an available off-the-shelf industrial pump has a maximum value of 0.04 m^3/s [Fischer Process Industries 2002]. We set that as the upper limit for R. We set $R = 0.001$ m^3/s and increment it in steps of 0.001 m^3/s until the displacement is greater than the radius of the basin.

Results and Discussion

We discuss how well our model handles each of the six variables that affect the solution:

- the fountain nozzle radius,
- the height of the building,
- the wind speed,
- the angle between the wind and the building,
- the building width, and
- the fountain radius.

Wind Speed

As the wind speed increases, the flow rate must decrease inside the fountain basin. Since the flow rate decreases, the height should also decrease because less water is forced through the nozzle, causing a lower initial speed. Does our model reflect these phenomena? Yes, it does.

Angle Between Wind and Building

The angle has no apparent effect on the solution our model produces. How is this possible? It is possible because our fountain is surrounded by buildings. The way we calculated the buildings' effect on the wind created similar effects at any angle.

There were indeed variations present when we calculated how much the wind was slowed down by the building depending on the angle; however, these variations were too small to affect the fountain's setting.

Nozzle Radius

A smaller radius at a given flow rate means a higher speed. As the radius increases, so does the flow rate until the maximum rate is reached.

What does this do to the height of the fountain? Height increases as the radius increases (because the rate of flow increases as well) until the maximum rate is reached. If the radius still keeps increasing, the flow remains constant through a larger opening, causing a lower speed and therefore a lower height.

Height of Building

As the height of the building increases, the height of the fountain could increase as well. Our model doesn't accurately reflect this. The problem most likely lies in our drag calculations, the only place where building height shows up. Both the wind angle and the building height depend on the accuracy of our drag assumptions, and both have produced questionable results.

Building Width

We finally find some data that suggest that our drag equation is at least partially correct. As the width of the building increases, more surface area is created for wind/building interaction. The increased surface area leads to more drag, a lower wind speed by the time the wind reaches the fountain, and therefore a higher flow rate and higher height of the fountain.

Fountain Radius

Our goal is to keep the water contained in the fountain basin. If the radius increases, we can shoot the water higher up into the wind and still have it land in the fountain basin. Both the rate of flow and the height produce an increased basin radius.

Summary and Conclusions

Our task was to develop a model to take inputs of wind speed and direction measured on a rooftop and use them to regulate flow through a nearby fountain.

By breaking the problem down into parts, we developed a model that produces believable results; and we have shown how our model responds to different inputs.

The effects of wake formation and wind interaction against a building are the two biggest problems. We assumed that the wake has no effect, and we dealt with the building interaction—but our findings raise questions. Clearly, the best test would be a fountain in a wind tunnel.

It would make sense to install an anemometer into the fountain structure and use wind-speed readings from the fountain itself. Direction and surroundings would be insignificant; only the wind speed at the fountain would be important.

We assume that wind gusts do not occur. There must somehow be a warning for the fountain that a gust is coming. Perhaps the rooftop anemometer could gauge wind speed change and send a signal to the fountain to reduce flow until the wind speed returns to normal.

References

Bertin, John J. 1984. *Engineering Fluid Mechanics*. Englewood Cliffs, NJ: Prentice-Hall.

Fischer Process Industries. 2002. http://64.44.169.181:8080/examples/jsp/fpvcom/fpvresults.jsp . Accessed 11 February 2002.

Floating Fountains. 1998. http://www.atlanticfountains.com/floatingfountains.htm . Accessed 11 February 2002.

Halliday, David, Robert Resnick, and Jearl Walker. 2001. *Fundamentals of Physics*. 6th ed. New York: Wiley.

King, Steven. 1999. Air Movement. http://fridge.arch.uwa.edu.au/topics/thermal/airflow/airflow.html . Accessed 11 February 2002.

Liu, Henry. 1991. *Wind Engineering: A Handbook for Structural Engineers*. Englewood Cliffs, NJ: Prentice Hall.

Macdonald, Angus. 1975. *Wind Loading on Buildings*. New York: Wiley.

Mooney, Douglas, and Randall Swift. 1999. *A Course in Mathematical Modeling*. Washington, DC: Mathematical Association of American.

Tritton, D.J. 1977. *Physical Fluid Dynamics*. New York: Van Nostrand Reinhold.

Yuan, Jinchao. 2002. Wind Simulation Around Building Clusters by Numerical Methods. http://www.arche.psu.edu/courses/ae597J/Jinchao.pdf . Accessed 11 February 2002.

A Foul-Weather Fountain

Ryan K. Card
Ernie E. Esser
Jeffrey H. Giansiracusa
University of Washington
Seattle, WA

Advisor: James Allen Morrow

Introduction

We devise a fountain control algorithm to monitor wind conditions and ensure that a fountain at the center of a plaza fires water high enough to be dazzling while not drenching the pedestrian areas surrounding the fountain.

We construct a model of a fountain based on the physics of falling water droplets considered as a particle system. We examine the behavior of a fountain under various wind conditions through computer simulation. Using complex analytic techniques, we model the wind flow through the plaza and estimate how anemometer readings from a nearby rooftop relate to plaza conditions.

We construct four algorithms—two intelligent algorithms, a conservative approach, and an enthusiastic system—to control the fountain.

We devise a measure of unacceptable spray levels outside the fountain and use this criterion to compare performance. First, we examine the behavior of these algorithms under general abstract wind conditions. Then we construct a wind signal generator that simulates the conditions of several major cities from meteorological database data, and we compare the performance of our control systems in each city.

Simulations show that the Conservative and Enthusiastic algorithms both perform unacceptably in realistic conditions. The Weighted Average Algorithm works best in gusty cities such as Chicago, but the Averaging Algorithm is superior in calmer cities such as Los Angeles and Seattle.

The control algorithm cannot possibly respond to changes in conditions at anything below the 10 s scale, since wind is highly variable and the response of the anemometer is somewhat slow [Industrial Weather Products 2002]. The goal is therefore to design the algorithm to operate on a time scale of 10 s up to a couple of hours and adapt the height of the fountain to a maximum safe level.

The UMAP Journal 23 (3) (2002) 251–266.

Model of the Water Jet

We model the spray from the fountain as a particle system. As water droplets spew forth from the nozzle, they are subjected to forces (gravity, air drag, turbulence, etc.). We formulate a simplified differential equation governing the motion and then numerically integrate to find the trajectory for each droplet. This equation is based on a physically realistic model of small droplets (around 1 mm radius) and we scale it up to an effective model for larger clumps of water (up to 10 cm across) because the physics of turbulence and viscosity at the larger scale cannot be computed accurately.

We need the following assumptions:

- The drag force is proportional to the square of the speed and to the square of the radius [NASA 2002].
- Droplets break into smaller droplets when subjected to wind. Breakup rate is proportional to relative wind speed and surface area [Nobauer 1999].
- When a droplet breaks, turbulence causes the new droplet fragments to move slightly away from their initial trajectory.

Modeling a Single Droplet

We formulate the motion of a water droplet as

$$m\frac{d\vec{v}}{dt} = -mg\hat{z} + \eta|w|^2\hat{w}r^2,$$

where $\vec{v}$ is the velocity, $\vec{w}$ is the wind velocity relative to the motion of the droplet (wind vector minus velocity vector), m and r are the droplet's mass and radius, and η is a constant of proportionality. According to the Virtual Science Center Project Team [2002], a raindrop with radius 1 mm falls at a terminal velocity of 7 m/s; so we determine that $\eta = 0.855$ kg/m^3. Large drops fall quickly; very tiny drops fall very slowly, mimicking a fine mist that hangs in the air for a long time.

We assume droplet breakup is a modified Poisson process, with rate

$$\lambda_{\text{breakup}} = \lambda_0|w|r^2.$$

If the breakup rate did not depend on variable parameters $|w|$ and r^2, the process would be a standard Poisson process. We determine λ_0 by fitting the water stream of our fountain to the streams of two real fountains: the Jet D'Eau of Geneva, Switzerland, and the Five Rivers Fountain of Lights in Miami, Florida.

When a breakup occurs, we split the droplet into two new droplets and divide the mass randomly, using a uniform distribution. Air turbulence tends to impart to the two new droplets a small velocity component perpendicular to the relative wind direction $\vec{w}$. This effect causes a tight stream of water to spread out as it travels, even under zero-wind conditions. We let this velocity nudge have magnitude 2% of the particle's speed relative to the air and a random perpendicular direction. We give the two drops equal and opposite nudges.

Putting Water Drops Together to Make a Fountain

We define the water jet as a stream of large water drops. Their size is roughly the size of the nozzle, and they leave with an initial velocity equal to the nozzle's output velocity (**Figure 1**).

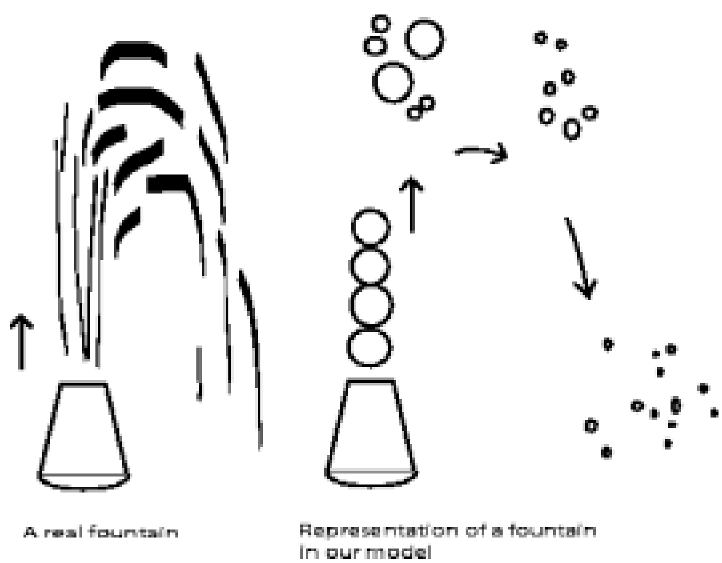

Figure 1. A continuous water jet is approximated by a discrete stream of water blobs.

The water blobs leave at a rate such that the flux of water is equal to the flux given by a nozzle-sized cylindrical stream moving at the same speed.

To model the turbulence in the jet as the water leaves the nozzle, we give each water blob a normal distribution of radius and initial speed:

- The standard deviation of blob radii is 10% of the nozzle size.
- The standard deviation of initial speeds is 5% of the initial speed.
- The blobs leave with an angular spread of 3°, consistent with industrial high-pressure nozzles [Spray Nozzles 2002].

Wind drag in particle streams is significantly reduced for particles following one another closely (NASCAR drivers and racing cyclists are intimately familiar with this phenomenon). These effects are already incorporated into the dynamics of large water blobs (which can be thought of as representing many small drops moving together). We therefore consider this to be an effective model for large drops rather than a realistic interaction model.

Fitting the Fountain

The Five Rivers Fountain of Lights in Daytona, Florida, is one of the largest fountains in the world. It consists of several water jets, and on low-wind days each propels a water stream 60 m high and 120 m out. The Jet D'Eau in Geneva, Switzerland, another impressive fountain, shoots a 30 cm-diameter stream of water at 60 m/s straight up. The water reaches a height of 140 m and on an average breezy day (wind speed 5 m/s) returns to earth approximately 35 m downwind from the nozzle [Micheloud & Cie 2002] (**Figure 2**).

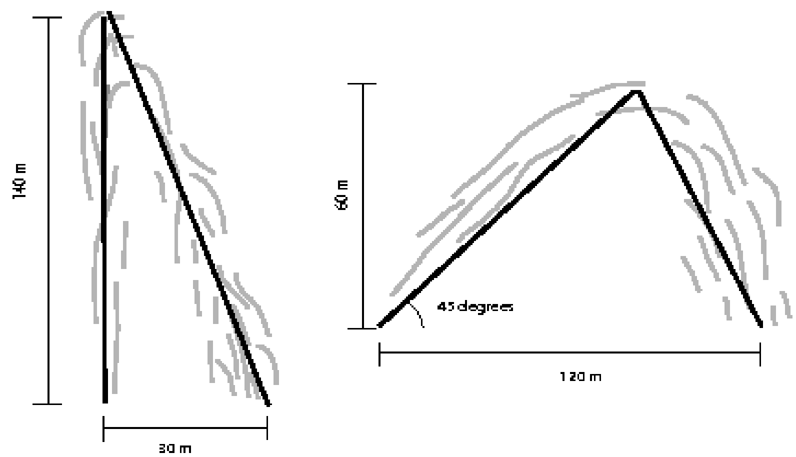

Figure 2. The Jet D'Eau and the Fountain of Lights.

To determine λ_0, we first match our geometry to the Five Rivers Fountain of Lights. We fix λ_0 so that with an initial velocity such that the stream reaches a height of 60 m, it returns to the ground at a distance of just over 100 m. Too large a λ_0 results in the water breaking up too quickly into tiny droplets, which have a much lower terminal velocity and thus fail to reach the desired distance; if the value is too small, then an unrealistically small amount of spray is produced and the water blob travels too far. The results are summarized in **Table 1**.

We set $\lambda_0 = 5000$. The results are highly insensitive to this parameter; varying λ_0 by a factor of 2 cause only a 15% changes in the distances. Therefore, even though our method for determining this parameter is fairly rough, the important behavior is much more strongly affected by other parameters.

Table 1.
Comparison between real fountains and our model.

	Jet D'Eau		Five Rivers Fountain	
	real	model	real	model
Height (m)	140	121	60	62
Distance (m)	35	30	120	100

We conclude from this comparison that our model reproduces the spray patterns of extreme fountains accurate to within about 15%. We expect that for a plaza-sized fountain, our model will be more accurate, since our formulas for breakup and drag force are derived under less extreme conditions.

Wind Flow Through the Plaza

Buildings and other structures in an urban environment can cause significant disturbances to wind flow patterns; rooftop and street-level conditions can often be quite different, so readings from a rooftop anemometer could be biased. To model the plaza wind, we assume:

- There are no significant structures between the buildings beside of the plaza.
- The plaza is large, so effects caused when wind flow leaves the plaza are negligible at the plaza center; the significant effects are entirely caused at the inward boundary passage.
- The air flow is smooth enough so that turbulent vortices are negligible.

Formulation

We approximate the geometry of the plaza as in **Figure 3** and use complex analytic flow techniques [Fisher 1990, 225].

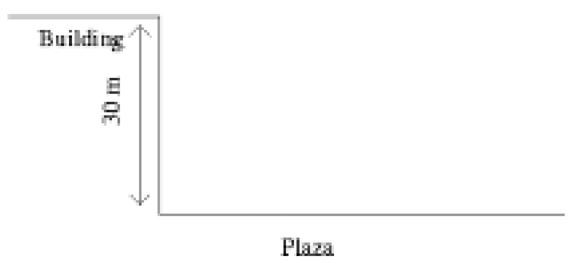

Figure 3. Schematic representation of the relevant features of the plaza.

With a Schwarz-Cristoffel mapping of a smooth horizontal flow from the upper half of the complex plane onto the region above the plaza, we obtain a flow function for the wind as it enters the plaza area:

$$\Gamma_c(t) = \frac{h_0}{\pi} \left\{ \left[(t+ic)^2 - 1 \right]^{1/2} + \log \left(t + ic + \left[(t+ic)^2 - 1 \right]^{1/2} \right) \right\},$$

where t parametrizes a streamline for each value of c. These streamlines are plotted in **Figure 4**, where the acceleration of the wind as it passes over the building edge and the decreased velocity in the plaza are both clearly visible.

The flow velocity $\vec{v}$ is inversely proportional to the streamline spacing, so the horizontal component of it is

$$v_x = \text{Im} \left[\frac{\partial \Gamma_c}{\partial c} \right].$$

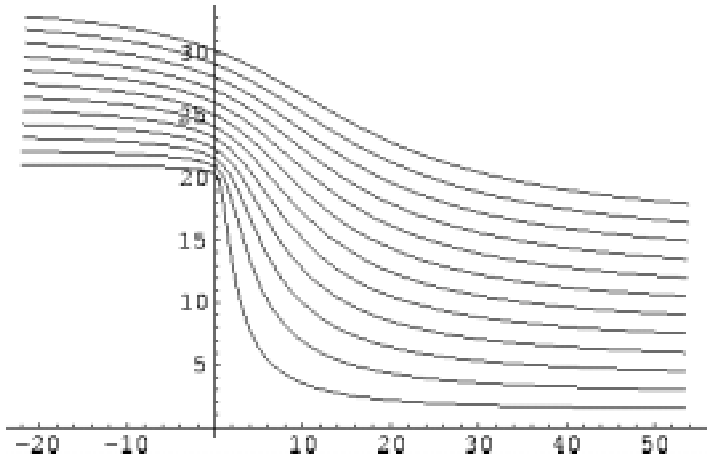

Figure 4. Streamlines for wind flow entering the plaza; decreased wind speed at the plaza level is apparent. Note the highly increased wind speed near the edge of the building.

The horizontal velocity profile for a streamline that passes about 3 m above the building roof (corresponding to $c = 0.6$) is plotted in **Figure 5**; 3 m is a reasonable height for an anemometer mounting. From these graphs, one can see that the wind speed through the plaza center (at a distance of 30 to 40 m from edge) is *approximately half* of the rooftop wind speed.

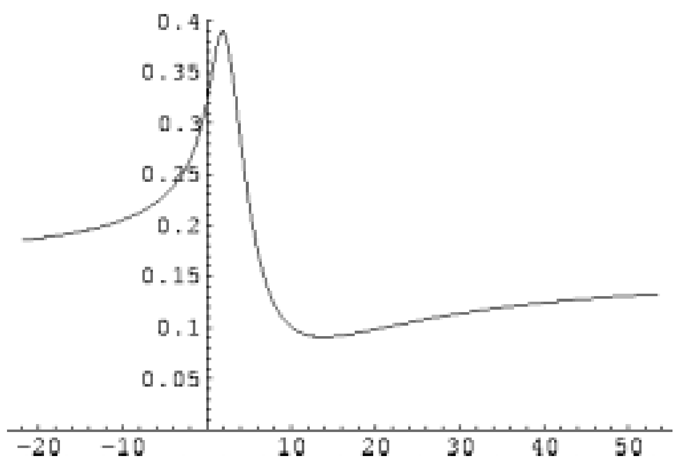

Figure 5. Horizontal velocity profile for the streamline corresponding to $c = 0.6$. This streamline passes above the building's roof at a height of 3 m, a reasonable anemometer mounting height.

This calculation is validated by its excellent agreement with the findings of Santamouris and Dascalaki [2000], who report that in flows perpendicular to a street the ground-level speeds are between zero and 55% of the free-stream speeds.

Results

We conclude from this flow model:

- **Placement of the anemometer is important!** It should be mounted near the center of the rooftop to minimize disturbances from the roof's edge.
- **The anemometer reports a wind speed that is highly biased!** Plaza-level wind moves approximately half as fast as the roof-level wind.
- **Wind speeds are spatially constant within the plaza airspace.** If the fountain is not significantly higher than the surrounding buildings, then spatial wind variation can be safely ignored.

Modeling Wind Variation Over Time

The control system must be able to handle a range of weather conditions, from calm up to strongly gusty. We abstract the wind patterns into three generalized types of increasing complexity:

- **Type 1: A low intensity constant breeze of a few m/s,** meant to test the algorithm's ability to judge the proper height for a given wind speed.
- **Type 2: A breeze varying smoothly over a timescale of a couple of minutes.** We use a sinusoidal oscillation in magnitude and direction, with a constant term to reflect the prevailing wind direction of the hour. This type tests the algorithm's ability to adapt to slowly changing conditions.
- **Type 3: Sudden unexpected wind gusts, with a few seconds duration and very high intensity.** We model the occurrence of a gust as a Poisson process and distribute the gust durations and intensities normally. The mean and variance are chosen to produce reasonable results. This is perhaps the most important test, since the gusty scenario can easily fool a naive algorithm.

Generating a Realistic Wind Signal

We parametrize the wind profile of a location by four numbers:

- The **mean steady wind** μ_{steady}.
- The **mean gust strength** μ_{gust}, where a gust is defined to be variation on the sub-15 s timescale.
- The **mean gust duration** t_{gust}.
- The **gust deviation** σ_{gust}.

From WebMET data [2001], we estimate these characteristic numbers for some major U.S. cities (**Table 2**).

We construct realistic wind signals from these characteristic numbers to correspond to our types:

Table 2.

Characteristic parametrization of several major U.S. cities. These parameters specify the plaza wind conditions, which are slightly milder than the free-stream conditions.

	μ_{steady} (m/s)	μ_{gust} (m/s)	t_{gust} (s)	σ_{gust} (m/s)
Seattle, WA	1.2	2.25	6.0	0.7
Chicago, IL	2.0	4.0	3.0	4.0
Boston, MA	2.3	4.2	4.0	2.2
Los Angeles, CA	1.7	2.0	3.0	0.7
Washington DC	1.3	3.4	3.0	1.0

- Type 1: constant wind of strength $\frac{2}{3}\mu_{\text{steady}}$,
- Type 2: sinusoidal oscillations of amplitude $\frac{2}{3}\mu_{\text{steady}}$, and
- Type 3: a gust signal with mean amplitude μ_{gust}, amplitude standard deviation σ_{gust}, duration mean t_{gust}, and deviation $\frac{1}{2}t_{\text{gust}}$.

Figure 6 shows a comparison of wind signals for Seattle and Chicago; the extreme gustiness of the "Windy City" is apparent.

We also create a "Hurricane Floyd" wind profile by multiplying a Chicago wind signal by a factor that damps the wind to zero early on (the calm before the storm) and then amplifies it to hurricane level over a period of 10 min.

Fountain Control Algorithms

The goal of the control algorithm is to respond to the anemometer data by maximizing the height of the fountain while minimizing the probability of the plaza area outside the fountain pool being drenched. The control algorithm has access to anemometer readings and direct control over the nozzle speed.

The control system must have some knowledge of how the water spray travel distance relates to nozzle speed and wind speed. We develop two complementary measures of spray distance and tabulate the relationship between them and nozzle/wind conditions. The algorithms that we develop combine this table with an estimate of possible future wind speed (based on the current wind and a stored recent history) to decide on a good nozzle speed.

Measures of Water Spray Distance

Our measures of spray distance are

- the radius within which 99% of sprayed water lands, and
- a threshold for acceptable water density outside the fountain that corresponds roughly to a light rain: 1 cm in 10 h (2.8×10^{-4} mm/s).

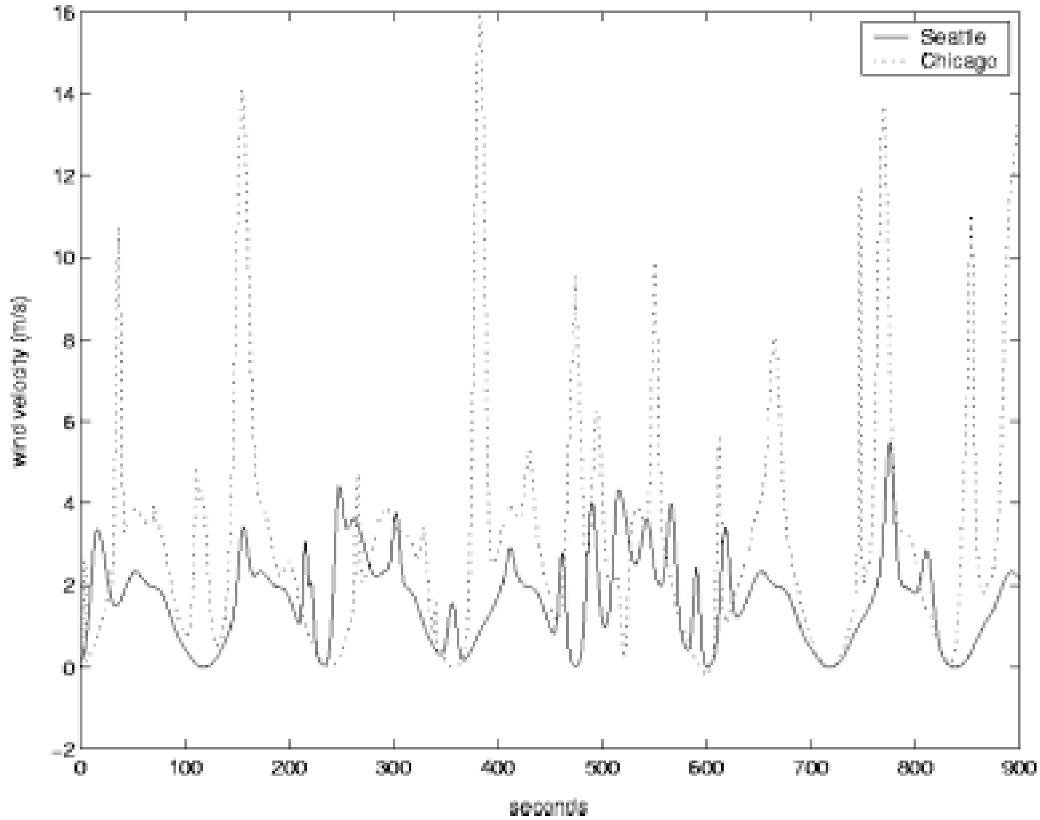

Figure 6. Horizontal velocity profile.

In simulations over a suitably long time period, we find that these two measures agree to within 1%.

We evaluate the performance of our control algorithm by measuring how the spray distance compares to the actual radius of the pool. If the spray radius goes beyond the pool radius, then people might become unacceptably wet. However, if this radius is significantly less than the pool radius, then we are not getting as much height out of the fountain as we could.

Constructing the Control System

We begin with a few useful assumptions:

- Variation in wind direction can be safely ignored. We use the triangle inequality: If the wind pushes a drop first in one direction and then in another, it will necessarily land nearer to the fountain than if it had been pushed in one direction continuously.
- The algorithm has access to real-time anemometer data averaged over 10-s intervals as well as at least a 10-min history of measurements. Even if the anemometer responds faster than 10 s, it is nonsensical to vary the fountain any faster than this, because the water requires approximately this much time to complete its flight.

- For concreteness, we focus on the plaza configuration of **Figure 7**. Most importantly, the fountain is at the center of a circular pool of radius 5 m.

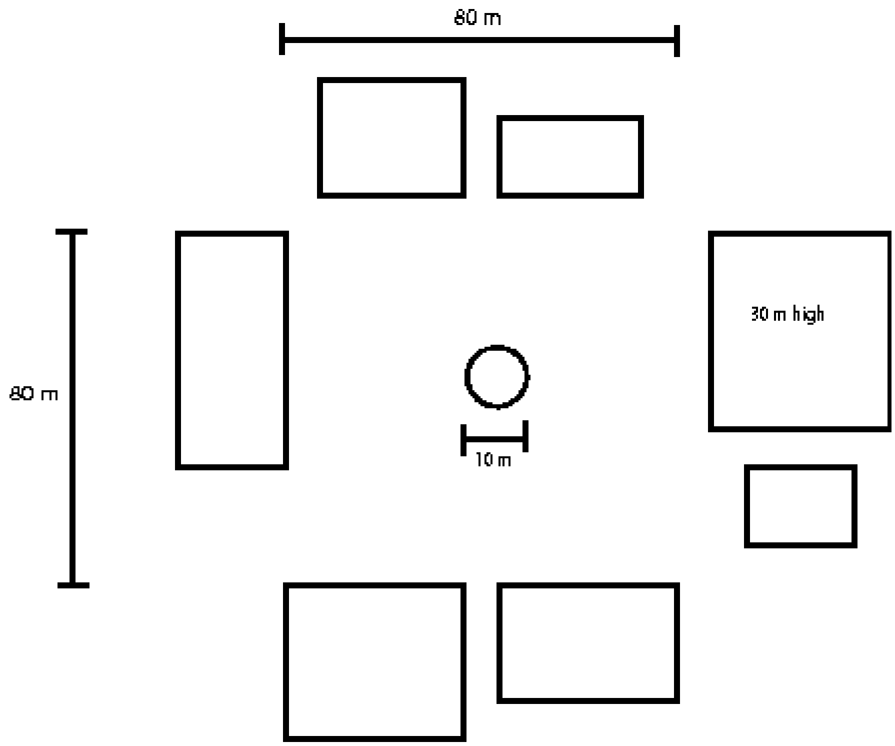

Figure 7. The layout of our hypothetical model plaza.

We display the spray distance as a function of wind speed and nozzle speed in **Figure 8**.

An estimate of how far a water droplet can travel starting at height z_0, falling at its terminal velocity v_t, and moving at horizontal wind speed w is

$$\text{distance} \approx \frac{z_0 w}{v_t}.$$

The smallest droplets that our simulations produce have radii of about 1 mm with corresponding terminal velocity 7 m/s. For specific heights and winds, we find that this rough estimate is usually within 30% of the corresponding minimum safe distance shown in **Figure 8**, a good indication that our simulations produce reasonable results.

The Control Algorithms

We formulate four control algorithms:

- **Averaging Algorithm:** This algorithm considers an average of the previous 10 min of wind data and the sample variance. The worst-case scenario is estimated to be a wind strength of one standard deviation above the average.

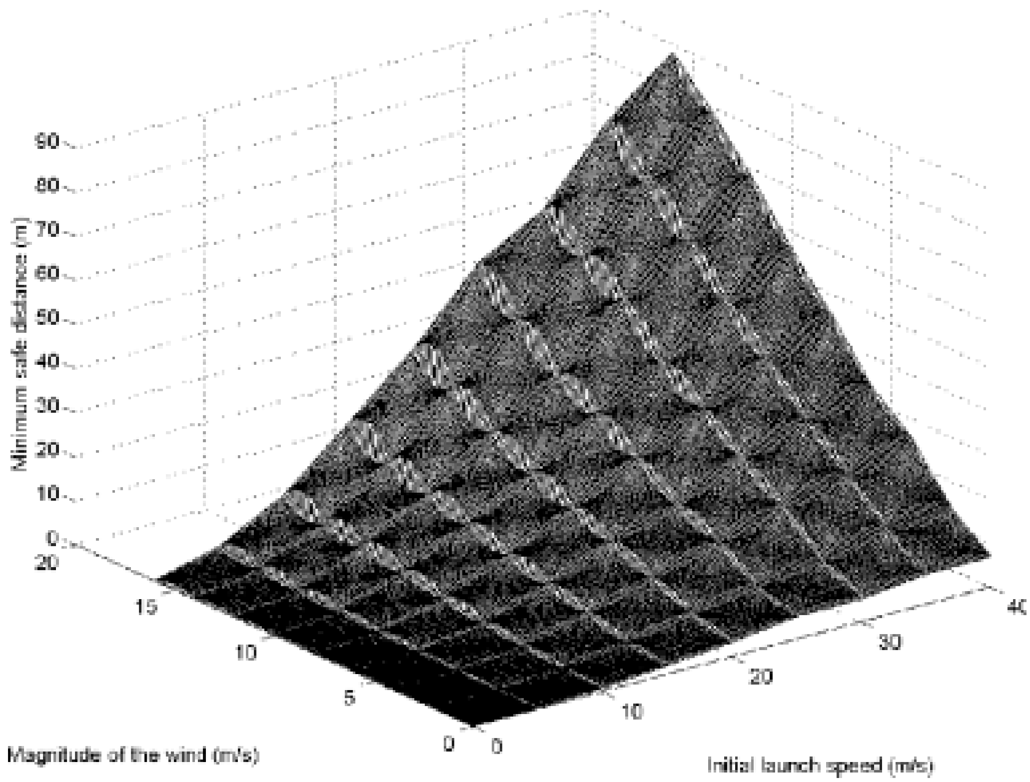

Figure 8. Linearly interpolated table of spray distance as a function of wind speed and nozzle speed. Each data point represents a burst of 5 nozzle-size water blobs.

- **Weighted Average Algorithm:** The key feature of this algorithm is that the data of the last 10 min are weighted linearly according to recentness. The current measurement gets the highest weight.
- **Conservative Algorithm:** This algorithm uses the maximum wind speed measured over the last 10 min to predict the worst-case wind. This is the most conservative approach—it will always err towards safety.
- **Enthusiastic Algorithm:** This algorithm ignores previous wind data history and puts the fountain to the maximum safe height given immediate conditions. No precaution is taken with regard to possible future wind behavior.

Results

Comparing the Algorithms

We test each algorithm against the following gamut of wind conditions:

- Type 1: constant wind
- Type 2: smoothly varying wind

- Type 3: highly variable gusty wind
- Real wind data from Seattle, Chicago, Boston, Los Angeles, and Washington DC.
- Hurricane Floyd-type winds!

We run several simulations of the fountain, each for 3 min, under the control of each algorithm—long enough to capture relevant wind features and give statistical significance to the results. For consistency, we run each algorithm under an identical wind signal (to remove random variation). We use the following criteria for comparing the performance of the algorithms:

- The average height of the water spout over the time of the simulation.
- The percentage of the total water contained within the pool.
- The ratio of the highest density of water landing outside the pool area to the maximum acceptable spray density (2.8×10^{-4} mm/s).

The results of our simulations (**Table 3**) indicates that the performance of the algorithms depends significantly on the wind data provided.

Strengths and Weaknesses

All of the algorithms perform equally well under constant wind conditions, but each has unique strengths and weaknesses.

- The *Enthusiastic Algorithm* consistently achieves the most spectacular fountain heights but at a cost. Since it considers only the current wind reading, it is always caught by surprise by sudden gusts or any increase in wind speed. Except in the constant-wind case, the algorithm systematically results in too much water being sprayed outside the fountain.
- The *Conservative Algorithm* always has the most paranoid estimate of how bad the wind could get, and all the water is usually contained in the fountain except in rare cases when sudden gusts greatly surpass the maximum recorded wind speed before the next measurement is made. However, the fountain height is often disappointingly low compared to the other algorithms, especially when a large gust of wind was recorded in the wind speed history.
- The *Weighted Average Algorithm* performs about as well as the *Averaging Algorithm*. Both contain most of the water but are often surprisingly conservative. In the Gusty Wind simulations, the Weighted Average Algorithm is even more conservative than the Conservative Algorithm; since both averaging algorithms consider the standard deviation of previous wind speed data, they become more conservative when recent wind speeds are highly variable. But if wind speeds change suddenly, as in the Hurricane Floyd case, the Weighted Algorithm reacts slightly faster than the Averaging Algorithm.

Table 3.

Comparisons of algorithm performance. When too much water spills out of the fountain, water densities become too computationally intensive to compute (denoted by $*$), and the fountain is operating well outside of acceptable parameters.

	Weighted Average	Average	Conservative	Enthusiastic
	Type 1: Constant Wind			
Average height	10.7 m	10.6 m	10.7 m	10.6 m
% contained	100%	100%	100%	100%
Density ratio	0	0	0	0
	Type 2: Smooth Wind			
Average height	12.0 m	12.4 m	12.1 m	20.2 m
% contained	100%	100%	100%	100%
Density ratio	0	.90	0	10321
	Type 3: Gusty Wind			
Average height	11.7 m	12.5 m	12.0 m	19.9 m
% contained	100%	100%	100%	99%
Density ratio	0	0	0	1357
	Hurricane Floyd-type wind!			
Average height	3.3 m	3.5 m	3.3 m	3.3 m
% contained	99%	98%	99%	98%
Density ratio	12	42	34	505
	Seattle			
Average height	10.4 m	10.6 m	5.0 m	20.7 m
% contained	99%	99%	100%	75.6%
Density ratio	61	25	0	$*$
	Chicago			
Average height	10.3 m	7.7 m	5.0 m	20.9 m
% contained	99%	99%	99%	62%
Density ratio	1357	2467	22	$*$
	Boston			
Average height	7.6 m	7.9 m	2.4 m	21.0 m
% contained	98%	97%	100%	95%
Density ratio	1964	11000	0	$*$
	Los Angeles			
Average height	7.6 m	10.5 m	5.9 m	10.2 m
% contained	99%	99%	100%	91%
Density ratio	2196	0.4	0	$*$
	Washington, DC			
Average height	8.7 m	10.2 m	7.7 m	20.8 m
% contained	99%	99%	100%	92%
Density ratio	3.36	18	0	$*$

Possible Extensions

Tiltable Nozzles

Water jets with directional control exist (firefighters use them extensively!). So, with a steady wind, aiming the fountain slightly into the wind may allow for a higher water stream without additional water spraying outside the pool.

For a range of constant wind speeds, we simulate the fountain at various tilt angles and find the angle that maximizes fountain height without unacceptable spray landing outside the pool (**Table 4**). For each run, we fire enough blobs (10) so that results are statistically significant.

Table 4.
Results of tilting the fountain into the wind.

Wind speed (m/s)	Maximum height (m) no tilt	tilt	angle
2	16.4	31.0	37.5°
5	10.8	22.5	8.5°
7	5.9	12.7	32.0°

The fountain can be made nearly twice as high by directing the nozzle into the wind. This would appear very encouraging indeed, were it not for two important points:

- **The spray distance is extremely sensitive to the tilt angle.** Variations of a single degree cause unacceptable amounts of water outside the pool area.
- **Real wind is rarely so constant.**

We therefore consider it infeasible to use tilting to increase the fountain height.

Multiple Nozzles

Our model can be extended to handle multiple nozzles by superimposing, provided that the stream-stream interaction is not significant.

Alternative Pool Geometries

We can handle fountains with noncircular pools, measuring the percentage of water that lands outside of the pool and requiring that no region gets too wet. If the fountain is in a city with wind predominantly in one direction, then an elliptical pool with major axis parallel to the wind direction may work better, though variation in wind direction can no longer be ignored by the model.

Other Considerations

- There are parameters that we did not incorporate in our model that may have effect in real life, such as temperature and barometric pressures.
- If a storm is approaching, the fountain should be turned off.
- At low temperature, we might set the algorithms to be more conservative, because it is very unpleasant to be wet in cold weather and ice formation can be dangerous.
- If the buildings around the plaza are significantly closer to the fountain than the 40 m considered in our simulations, then the dynamics of the wind near the fountain may be altered with the addition of eddies and other turbulence.
- For fountains that reach heights significantly higher than the nearby buildings, the magnitude of the wind will grow stronger farther above the plaza.
- A longer wind history could be incorporated into the algorithm.

Recommendations and Conclusions

If keeping the water spray contained in the pool is a much larger concern than shooting the fountain high into the air, then the Conservative Algorithm may be the best choice. Conversely, if water spray outside the fountain is not an overriding concern, than the Enthusiastic Algorithm may be best.

For a reasonable balance between safety and dazzle, the Conservative Algorithm and the Enthusiastic Algorithm are both totally inadequate:

Use either the Weighted Average Algorithm or the Averaging Algorithm.

The Weighted Average Algorithm responds faster to sharp changes in wind speed and performs better in places like Chicago where wind gusts are more frequent. However, if wind variations are fairly smooth, as in Los Angeles, then the Averaging Algorithm is the best choice.

References

Fisher, Stephen D. 1990. *Complex Variables*. 2nd ed. New York: Dover.

WebMET: The Meteorological Resource Center. 2001. `www.webmet.com` . Accessed 11 February 2002.

Santamouris, Matheos, and Elena Dascalaki. 2000. Wind speed in the urban environment. `www.brunel.ac.uk/research/solvent/pdf/report3.pdf` . Accessed 10 February 2002.

Industrial Weather Products Catalog. 2002. `www.scientificsales.com` . Accessed 10 February 2002.

Micheloud & Cie. 2002. `www.switzerland.isyours.com` . Accessed 10 February 2002.

NASA. 2002. Re-Living the Wright Way. `wright.wrc.nasa.gov` . Accessed 11 February 2002.

National Climatic Data Center. 2002. `www.ncdc.noaa.gov` . Accessed 10 February 2002.

Nobauer, Gerhard Thomas. 1999. *Droplet Motion and Break-up*. Oxford, England: University of Oxford Computing Laboratory-Mathematical Institute.

Spray Nozzles. 2002. `www.mrpressurewasher.com/spraynozzles.html` . Accessed 9 February 2002.

Virtual Science Center Project Team. 1997. ScienceNet. Accessed 9 February 2002.

Ross, Sheldon M. 2000. *Introduction to Probability Models*. 7th ed. New York: Academic.

Judge's Commentary: The Outstanding Wind and Waterspray Papers

Patrick J. Driscoll
Department of Systems Engineering
U.S. Military Academy
West Point, NY 10996
fp5543@exmail.usma.army.mil

Introduction

As so often is the case with events that iterate on an annual basis, many of the same lessons learned carry on from year to year, never losing their relevancy. Certainly, the MCM this year is no exception to the trend.

In an attempt to maintain some degree of economy in this commentary, I will resist the temptation to reiterate many of these again herein and point the interested reader to MCM commentaries previously appearing in this *Journal*.

However, there are several notable modeling issues that clearly surfaced in consideration of the Wind and Waterspray Problem that had an impact on the quality of the papers and are worth mentioning to assist teams in future competitions. In this vein, the following comments represent a compendium of observations during the final judging session and are taken in no particular order of preference or priority.

Style and Economy

As to style and clarity of the papers, it is probably sufficient to state that teams should bear in mind that they are writing to a population of modeling experts from both academia and industry who will spend a limited amount of time reading their paper. During this period, judges must assess the quality of a team's approach, the validity of their results, and the paper's completeness with regard to the modeling process. Contrast this with the hours and sometimes

The UMAP Journal 23 (3) (2002) 267–271.

days available for a professor to grade a similar project of this type, and it is apparent that teams must choose a writing style that maximizes clarity and gets across their modeling work in the most effective manner possible. Using concise and properly labeled tables and graphics to illustrate the trends and results of experimental trials that are commented on in the body of the report goes a long way towards achieving this goal.

The Specific Challenges of This Problem

The stated challenge of the Wind and Waterspray Problem was to develop an algorithm that uses data provided by an anemometer to adjust the water flow from a fountain as wind conditions change.

In a most general sense, an algorithm can be succinctly defined as a "method for the solution of all problems of a given class ... whose purpose is to describe a process in such a way that afterwards [it] can be imitated or governed by a machine" [Gellert et al. 1977, 340]. A basic characteristic of an algorithm is that it transforms given quantities (input) into other quantities (output) on the basis of a system of transformation rules. The input quantities (anemometer data) and output quantities (water flow characteristics) for the problem were clear from the problem description. The particular transformation rules for this problem were unspecified and left up to the individual teams to decide upon.

Formulating these transformation rules constituted the heart of each approach used to model the water flow and spray patterns associated with the fountain. The most predominant appeared to be Newton's Second Law of Motion, Bernoulli's formula, continuity equations, fuzzy membership sets, Poiseuille's equation, or Navier-Stokes equations, largely dependent on the assumptions that teams were willing to make.

The better papers walked the reader through the application of the approach chosen, clearly explaining exactly how each variable and parameter applied to the problem, and then used the known results of the specific approach directly.

How to Make Assumptions

Most technical report formats advise students to list and explain all their assumptions in one concise location, typically in the front portion of the report. While this advice is sound for constructing a technical report, it might be helpful to note that it is in contrast with the pattern of how assumptions occur chronologically during a modeling process. For the MCM, useful assumptions typically arise in one of two settings:

- either a team needs specific information concerning the problem that they do not have (and cannot get in the time allotted) and hence must make an assumption in order to carry on; or

- a team decides to make an assumption that simplifies some detail(s) of the problem in order to use the mathematics they are familiar with or risk not being able to complete their modeling effort in the time allotted.

Both of these situations arise naturally in the chronological flow of attacking a problem, and not during a single brainstorming effort at the onset.

When a paper contains a long list of assumptions, many of which neither get used nor justified in the modeling that follows, it is a clear indication that the team does not quite understand the roles that the assumptions play in the overall modeling effort. Such papers typically possess a very shallow or missing "Strengths and Weaknesses" section, which is supposed to constitute an analysis of one's model and results in consideration of the assumptions that were included by necessity. If a team does not know why they need a particular assumption, chances are that they will do a poor task of explaining why they made the assumption!

The lesson here is that teams should struggle mightily to make only the assumptions they need when they need them, thereby minimizing the diluting effect on model fidelity caused by an excessive number of assumptions.

The Importance of Model Validation

When all is said and done, a paper introducing a proposed algorithm must resolve the question, "Does the author provide me with sufficient evidence that it works?" While occasionally provided by way of convergence proofs, this type of evidence more commonly appears in MCM papers by way of computational testing. For the MCM, at least three categories of testing come to mind that support model *validation*:

- Once the team is convinced that their base model produces reasonable results, special cases of interest (e.g., no wind, no spread angle, etc.) should be tested.

- Recognizing that model parameters contain some amount of uncertainty, high, most likely, and low values of important parameters used in the base model should be examined by systematically altering these values and rerunning the model to see if the output results remain reasonable. For this MCM problem, these parameters might be drag coefficients, shapes of water droplets, wind speed and direction, and so on. This process essentially constitutes what is commonly referred to as *sensitivity analysis* of the parameters.

- The effects of relaxing a select number of simplifying assumptions made during the course of developing the model should be examined. However, it is fair to stress that this last category is safely performed only when time permits, because it generally requires substantial model modifications to examine the desired effects. A good example of this third category for the Wind and Waterspray problem would be adding the influence of surrounding buildings on wind speed and direction after they were previously assumed away. Such a change would be nontrivial and might consume more time than what is available.

Teams must link their computational results back to the problem that they are trying to solve. Tell the reader what to conclude from the results! This is what is referred to as *analyzing the results*. Never, ever, ever leave this task to the reader!

When the conclusions of these analyses remain the same despite changes in parameters such as those noted, it is appropriate to conclude that the model results are *robust*. These analyses also highlight any limitations of the model, which then provide a basis for recommending ways the model could be enhanced or improved in the future.

The Summary

The summary that the MCM asks for is a standalone object that should not be identical to the introduction to the paper. The summary should briefly

- state the problem,
- describe the approach taken to modeling the problem,
- state the most important results and conclusions the reader should remember, and
- mention any recommendations directly relevant to the problem.

The summary should not include a statement such as "read inside for results" or its equivalent. A good test a team can use to assess the quality of their summary is to ask, "If someone read only the summary without the rest of the report available, would it clearly tell the big picture story of what the problem was, what we did, what we concluded, and what we recommend?" As a note, most equations, code, and derivations belong somewhere else as well.

Advance Planning

With regard to time management, something that teams can do ahead of the contest is to decide

- what document-writing environment they intend to use;
- how equations will be entered and labeled;
- the outline format of the paper;

- how tables, figures, and graphics are going to look;
- how captions are going to be stated for all tables, figures, and graphics; and
- who will be responsible for what task in the final write-up.

Human nature being what it is, a sloppy or haphazard paper that looks as if it was put together 15 min before it had to be postmarked almost assuredly will be downgraded in the mind of a judge, independent of the specific results obtained.

Use of Sources

Finally, the observed trend continues that teams are becoming increasingly selective with regard to the Web sites that they will trust for credible information. I also encourage teams to maintain their effort to properly document sources used to support their work. This practice explicitly recognizes the intellectual property and work of others while strengthening the quality of their paper at the same time.

Reference

Gellert, W., H. Kustner, M. Hellwich, and H. Kastner. 1977. *The VNR Concise Encyclopedia of Mathematics*. New York: Van Nostrand Reinhold.

About the Author

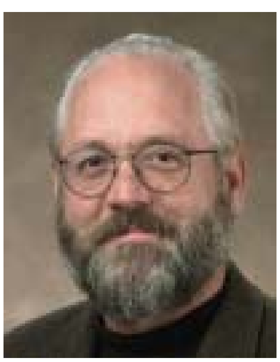

Pat Driscoll is Professor of Operations Research in the Department of Systems Engineering at the United States Military Academy. He holds an M.S. in both Operations Research and Engineering Economic Systems from Stanford University, and a Ph.D. in Industrial and Systems Engineering from Virginia Tech. His research focuses on mathematical programming, systems design for reliability, and information modeling. Pat is the INFORMS Head Judge for the MCM.

Things That Go Bump in the Flight

Krista M. Dowdey
Nathan M. Gossett
Mark P. Leverentz
Bethel College
St. Paul, MN

Advisor: William M. Kinney

Introduction

We develop a risk assessment model that allows an airline to specify certain parameters and receive recommendations for compensation policy for bumped passengers and for how much to overbook each flight. The basis is the potential cost of each bumped passenger compared to the potential revenue from booking an extra passenger. Our model allows an airline to compare quickly the likely results of different compensation and overbooking strategies.

To demonstrate how our model works, we apply it to Vanguard Airlines. Publicly available data provide all of the needed parameters for our model. Our software package reaches an overbooking policy by calculating and comparing the expected revenues for all possible situations and compensation policies.

Terms and Definitions

We set out terminology, taking much of it from Delta Airlines [2002].

- *Available seat miles* **(ASM)**: A measure of capacity which is calculated by multiplying the total number of seats available for transporting passengers by the total number of miles flown during a reporting period.
- *Revenue passenger mile* **(RPM)**: One revenue-paying passenger transported one mile. RPM is calculated by multiplying the number of revenue passengers by the number of miles they are flown for the reporting period.
- *Load factor* **(LF)**:A measure of aircraft utilization for a reporting period, calculated by dividing RPM by ASM.

The UMAP Journal 23 (3) (2002) 273–282.

- *Cost per available seat mile* **(CASM)**: Operating cost per available seat mile during a reporting period; also referred to as *unit cost*.
- *Revenue per available seat mile* **(RASM)**: Total revenue for a reporting period divided by available seat miles; also referred to as *unit revenue*.
- *"No-show"*: A person who purchased a ticket but does not attempt to board the intended flight.
- *Bumping*: The practice of denying boarding to a ticket holder due to lack of sufficient seating on the flight.
- *Voluntary bumping*: When passengers who purchased ticket for a flight give up their seats for some compensation offered by the airline.
- *Involuntary bumping*: When not enough passengers voluntarily give up their seats, the airline chooses whom to bump against their will.
- *Revenue*: Money gained by the airline from a flight minus penalties paid to bumped passengers. **This is *not* the standard definition of revenue** ("inflow of assets as result of sales of goods and/or services" [Porter 2001, 146]). We use this different definition to highlight the effect of bumping practices.)
- *Flight leg*: A direct flight from one airport to another with no stops.

Assumptions

- **Passenger airline traffic is returning to normal, so yearly industry statistics can be used.** Airline traffic trends are returning to the levels before the terrorist attacks on September 11 [Airline Transport Association 2001], so statistics from before that date are still valid.
- **We model U.S. flights only.** International flights have different policies.
- **The "no-show" rate is about 10%.** ["More airline passengers . . . " 1999].
- **Ticket prices may be represented by calculated averages.**
- **The number of passengers on the plane does not affect the cost of the flight to the airline.** The most significant part of the operating costs for a flight are fixed costs that are not be affected by the number of passengers.
- **The flight schedule is static.** The schedule of flights is outside of the scope of our problem statement. Thus, we make recommendations only about the overbooking strategy, not about changes to the schedule.
- **Airlines must follow the DOT "Fly-Rights" regulations.** These regulations outline the minimal compensation required to passengers when bumping occurs [U.S. Department of Transportation 1994].

- **Compounded overbooking takes care of itself (i.e., goes away naturally).** Consistent industry-wide statistics establish a 60% to 80% load factor [Airline Transport Association 2002], resulting in naturally combating the waterfall effect of one overbooked flight causing another to be even more overbooked.
- **There is sufficient demand for at least some flights to warrant overbooking.**
- **No-shows do not generate revenue.** No-shows are given a refund or (if original ticket was nonrefundable) a ticket voucher.
- **Taxes paid by a passenger are nonrefundable.**

Statement of Purpose

- Our first priority is to maximize revenue for the airline.
- Our second priority is to maximize customer service in the form of providing as much compensation to bumped passengers as is financially feasible.

Naive Model

The naive approach is to assume that since not all ticket buyers show up for the flight, we simply overbook the flight so that on average the plane fills to capacity. If on average 90% show up, we book to $100/90$ capacity.

However, the 90% is only an average; for some flights, more than 90% will show up, resulting in bumped passengers and a penalty for the airline paid to bumped passengers. Since the penalty is often more than the potential revenue for one more passenger, the airline could pay more in penalties than the extra revenue received. We need a way to factor the risk of penalties into our model.

Risk Assessment Model

We maximize revenue on each individual flight leg, which we regard as independent of other flight legs. Thus, optimizing the revenue of one flight does not adversely affect potential revenue from other flights.

Since an airline incurs an increased penalty the longer that a bumped passenger is delayed, an airline minimizes the penalty by transporting the passenger to their destination as quickly as possible. Therefore, bumped passengers are usually booked on the next flight or series of flights to their destination.

Expected Revenue of a Flight

Let a flight have capacity of c and we book b passengers. Let r be the potential revenue from a passenger and p the potential penalty cost of a passenger bumped . Finally, let x be the percentage of ticket holders who show up for the flight. The revenue generated by the flight is

$$\text{revenue}(x, b) = \begin{cases} xbr, & \text{if } xb \leq c; \\ cr - (xb - c)p, & \text{if } xb > c. \end{cases}$$

The percentage x of passengers who show up follows some probability distribution with density function $f(x)$ and an appropriate mean (in our case, 0.9). We find the value of b that maximizes the expected revenue for b passengers:

$$\text{expected\_revenue}(b) = \int_0^1 f(x) \cdot \text{revenue}(x, b)\, dx$$

Repeat this process for all flights and you have a complete recommendation for an overbooking policy.

Examining Compensation Policies

We can adjust our model even further by examining the effects of different compensation policies. Airlines have several forms of compensation at their disposal, from food to hotel stays to vouchers. The cost of the compensation policy is the penalty paid to a bumped passenger (p in our formulas above). By rerunning our expected revenue calculations for each compensation policy, we can see how each policy affects the maximum expected revenue of a flight.

Key Overbooking Flights

An airline can determine from historical data the "key" overbooking flights, the ones most likely to require overbooking. It can then use a compensation policy that concentrates on maximizing expected revenue for those flights.

From Theory to Reality: Vanguard Airlines

We illustrate our ideas by a case study of Vanguard Airlines, using publicly available information below [Vanguard Airlines 2001]. We assume that the January 2001 through September 2001 statistics provide an accurate picture of the airline:

- RASM = $ 0.073/seat-mile.
- RPM = 817,330 passenger-miles.

- ASM = 1,225,942 seat-miles.
- Operating expenses per ASM = $ 0.090/seat-mile.
- A full flight (Boeing 737-200 or MD-80 aircraft) holds $c = 130$ passengers.
- 95% of bumped passengers are volunteers [U.S. Department of Transportation 2001].

Applying the Model

We created a software package parameterized for adaptation to any airline.

Vanguard's Web site [2002] gives a list of flight legs, along with source cities, destination cities, departure times, and arrival times. All flight legs are flown daily, except for four; to keep our example simple, we ignore these exceptions and treat all flights as daily.

The potential revenue r per passenger is the average ticket price for the flight leg; we calculate it as flight-leg distance times revenue earned per passenger mile. The latter is total revenue (RASM $\times$ ASM) divided by passenger-miles flown (RPM). So we have

$$r = \frac{(\text{distance})(\text{RASM})(\text{ASM})}{\text{RPM}}.$$

We could not locate good data on the distribution of how many ticket buyers show up for the flight. In lieu of a real distribution, we use a truncated normal distribution with mean 0.9 and appropriately small standard deviation (0.05):

$$f(x) = \frac{1.023}{0.05\sqrt{2\pi}}\, e^{200(x-0.9)^2}.$$

Penalty costs depend on how long the passenger is delayed, so we search the flight schedule for the quickest alternative route for each flight leg. We require at least 30 min between connecting flights.

Compensation Policies

There are three main forms of compensating bumped passengers:

- **Cash Payment vs. Ticket Voucher**
  - Bumped passengers who arrive at their destination within one hour of their originally scheduled arrival receive no compensation.
  - Those who arrive between one and two hours after their originally scheduled arrival are eligible for compensation in the amount of their full ticket cost up to $200.

- A passenger who arrives two or more hours late is eligible for compensation in the amount of double their ticket cost up to $400.

Compensation is required only for passengers involuntarily bumped, but common practice is to offer similar amounts to attract volunteers for bumping. We assume that 95% of all "bumped" passengers are voluntary and we offer them vouchers in place of cash. We calculate that a $1.00 voucher costs the airline $0.82. Incorporating that 5% of bumped passengers receive cash, this plan costs (voucher_value) $\times$ 0.831 per bumped passenger.

- **Meal Compensation** In our software, a passenger sitting in an airport through particular intervals gets compensation for a meal: 6 A.M. to 9 A.M., breakfast ($10); 11 A.M. to 1 P.M., lunch ($10); 5 P.M. to 8 P.M., dinner ($15). This compensation is not mandated, so it serves only as customer service.
- **Providing Lodging** A quick survey of airport motels in Kansas City (the hub for Vanguard) showed that $50 is reasonable to cover a motel room along with transportation to and from the motel. Our plan offers overnight accommodation to a passenger stranded in an airport for at least 6 hrs including midnight who has a flight leaving after 4 A.M. This compensation is not mandated, so it serves only as customer service.

Choosing a Compensation Policy

We compare the impacts of the following policies:

- Meal compensation, hotel compensation, and cash
- Hotel compensation and cash
- Meal compensation, hotel compensation, and voucher
- Hotel compensation and voucher
- Meal compensation and cash
- Meal compensation and voucher

We tabulate penalties for each flight leg and each compensation policy and calculate an optimal number of passengers to book on each flight leg depending on the compensation policy. To ensure that bumping is no more likely than not needing to bump, we impose a maximum booking level of 10/9. We then calculate the expected revenue for each flight leg at the optimal booking level for each policy and rank the policies for each flight leg by expected revenue. [EDITOR'S NOTE: We omit the authors' extensive tables giving results for specific flights.]

An important consideration in choosing a compensation package is customer service. While there is little short-term impact on revenue from good or bad customer service, there can be significant long-term impact. We should give some preference to policies that offer greater customer satisfaction. When two or more policies produce the same revenue, our model chooses one that maximizes customer service.

The Best Compensation Policy for Vanguard

To determine its best compensation policy, Vanguard would need to examine historical data to determine flights most likely to require overbooking.

Responding to the Current Situation

We turn to issues currently facing the airline industry. Here we demonstrate how our model deals with unexpected circumstances.

Fewer Flights

The airline sets its schedule; our model adapts to it. In any case, flight traffic is increasing back to the level before September 11 (**Figure 1**).

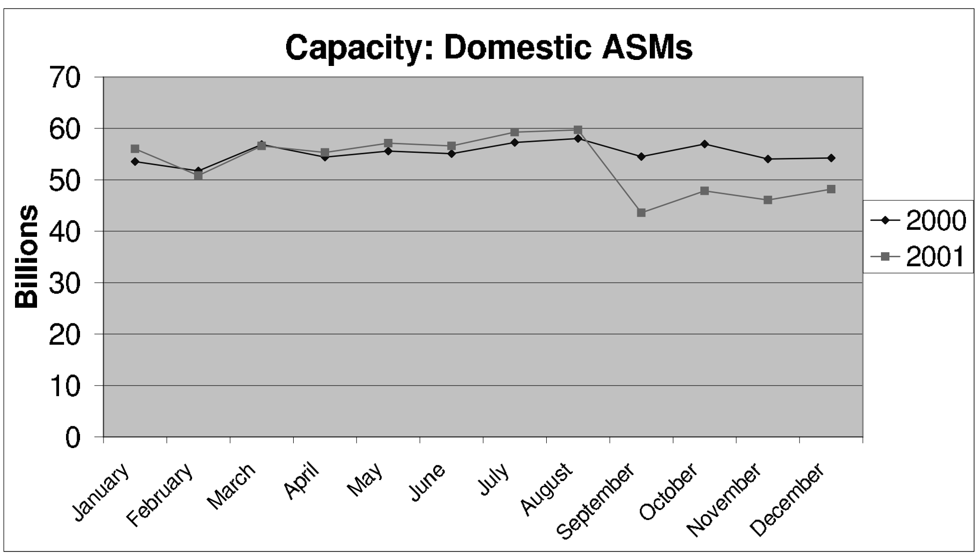

Figure 1. Domestic available seat miles (ASM) by month [Airline Transport Association 2002].

Heightened Security

Since the change in security policies at airports nationwide, both the checking in for a flight and layover gate changes could slow down passengers. Our model adjusts for that by factoring in 30 min for a layover.

Passengers' Fear

Passengers' fear could reduce no-shows (because those who purchase tickets are more serious about needing to fly) or increase them; there are no statistics to verify either effect. In either case, any effect of passengers' fear of flying seems to be declining [Airline Transport Association 2002].

Airlines' Losses

Revenue losses will likely make airlines cautious about taking on too much risk yet anxious to maximize revenue. Our model takes both goals into account, including enhancing revenue by dropping some customer service aspects.

Revenue loss also could cause an airline to schedule fewer flights to reduce costs. Our model gives an optimal recommendation adapted to the schedule.

Other Recommendations

- If two compensation packages have the same revenue benefits, choose the package that benefits the customer the most.
- Use vouchers instead of cash for compensation, because it costs less yet carries comparable perceived value for the customer.
- Give gate attendants some power to negotiate with angry customers, possibly including additional food vouchers.
- Upgrade bumped passengers to first class on their next flight when possible. This has no added cost in the case of an empty first class seat, yet has high value to the customer.
- Whenever possible, bump volunteers first, followed by passengers flying only one flight leg. This reduces the risk of further complicating a passenger's schedule.
- Ensure that the compensation policy is comparable to other airlines'.

Analysis of our Model

Strengths

The fundamental strengths of our model are its robustness and flexibility. All of the data is fully parameterized, so the model can be be applied to any airline. An airline can easily create probability distributions that accurately reflect not only average no-show percentages but also historical or per-flight trends. Although the industry may face constantly changing situations, our model adjusts to give the best recommendation possible.

Opportunities for Further Development

- The Vanguard implementation of our model ignores exceptions in flight schedules, assuming that all flights are daily. The incorporation of flight schedule exceptions into our implementation would be straightforward.
- Proprietary data would improve the accuracy of our Vanguard example, including the probability distribution of no-shows, average ticket prices, the current cost of various forms of compensation, and which flights are high-demand flights.

References

Airline Rules and Policies. 2002. http://www.onetravel.com/advisor/AR_rules_Menu.cfm . Accessed 9 February 2002.

Airline Transport Association. 2002. http://www.airlines.org/public/industry/ .

Delta Airlines. 2000. Glossary of Terms. http://www.delta.com/inside/investors/annual_reports/2000b_annual/06_glossary.html. Accessed 9 February 2002.

More airline passengers get "bumped". 2000. *USA Today* (6 October 2000). http://www.usatoday.com/life/travel/business/1999/bt002.htm .

Porter, Gary A., and Curtis L. Norton. 2001. *Financial Accounting*. New York: Harcourt College Publishers.

U.S. Department of Transportation. 1994. Fly-Rights Regulations. http://www.dot.gov/airconsumer/flyrights.htm#overbooking . September 1994. Accessed 9 February 2002.

________. 2001. Air Travel Consumer Report. http://www.dot.gov/airconsumer/0112atcr.doc . December 2001. Accessed 11 February 2002.

Vanguard Airlines. 2001. Vanguard Airlines 10-Q Filing. Securities and Exchange Commission. http://www.sec.gov/Archives/edgar/data/1000578/000100057801500016/r10q111401.htm. Filed 11 November 2001. Accessed 11 February 2002.

________. 2002. http://www.flyvanguard.com .

WebFlyer MileMarker. 2002. Air Distance Calculator. http://www.webflyer.com/milemarker/getmileage.cgi . Accessed 9 February 2002.

M e m o r a n d u m

February 11, 2002

To: Scott Dickson,

Chairman, CEO and President of Vanguard Airlines

From: Team 229

Airline Yield Management Research

Subject: **Policy Changes for Optimal Overbooking and Compensation**

After careful analysis of our company's current flight schedule, revenue per additional passenger, overbooking strategy, and several different compensation schemes for bumped passengers, we have the following recommendations to maximize revenue on high demand flights.

1) Since our historical records indicate that flights101b, 251a, 325b, 451a, 552a and 902a are the flights in highest demand, we recommend that our airline adopt a compensation package consisting of the following policies
   a) Provision of overnight accommodations for those stranded during applicable times
   b) Ticket vouchers for all sufficiently delayed passengers who will accept them in lieu of cash.
   c) In addition, we are recommending against the use of meal compensation.

   This policy ensures the least cost to the airline in the case of an overbooked passenger while maintaining the highest possible level of customer satisfaction at that cost.
2) While using this compensation package, we also advise you to overbook flights using the following numbers of allowable bookings for each high demand flight. A complete reference of allowable flight booking levels is available.

Flight	Allowable Bookings
101b	142
207a	140
325b	142
451a	142
552a	142
902a	142

Optimal Overbooking

David Arthur
Sam Malone
Oaz Nir
Duke University
Durham, NC

Advisor: David P. Kraines

Introduction

We construct several models to examine the effect of overbooking policies on airline revenue and costs in light of the current state of the industry, including fewer flights, increased security, passengers' fear, and billions in losses.

Using a plausible average ticket price, we model the waiting-time distribution for flights and estimate the average cost per involuntarily bumped passenger.

For ticketholders bumped voluntarily, the interaction between the airline and ticketholders takes the form of a least-bid auction in which winners receive compensation for foregoing their flights. We discuss the precedent for this type of auction and introduce a highly similar continuous auction model that allows us to calculate a novel formula for the expected compensation required.

Our One-Plane Model models expected revenue as a function of overbooking policy for a single plane. Using this framework, we examined the relationship between the optimal (revenue-maximizing) overbooking strategy and the arrival probability for ticketholders. We extend the model to consider multiple fare classes; doing so does not significantly alter optimal overbooking policy.

Our Interactive Simulation Model takes into account estimates for average compensation costs. It simulates the interaction between 10 major U.S. airlines with a market base of 10,000 people, factoring in passenger arrival probability, flight frequency, compensation for bumping, and the behavior of rival airlines. Thus, we estimate optimal booking policy in a competitive environment. Simulations of this model with likely parameter values before and after September 11 gives robust results that corroborate the conclusions of the One-Plane Model and the compensation-cost formula.

The UMAP Journal 23 (3) (2002) 283–300.

Overall, we conclude that airlines should maintain or decrease their current levels of overbooking.

Terms

- **Ticketholders**: People who purchased a ticket.
- **Contenders**: Ticketholders who arrive in time to board their flight.
- **Boarded passengers**: Contenders who board successfully.
- **Bumped passengers**: Contenders who are not given seating on their flight.
- **Voluntarily bumped passengers**: Bumped passengers who opt out of their seating in exchange for compensation.
- **Involuntarily bumped passengers**: Bumped passengers who are denied boarding against their will.
- **Compensation costs**: The total value of money and other incentives given to bumped passengers.
- **Flight Capacity**: The number of seats on a flight.
- **Overbooking**: The practice of selling more tickets that flight capacity.
- **Waiting time**: The time that a bumped passenger would have to wait for the next flight to the destination.
- **Load factor**: The ratio of the number of seats filled to the capacity.

Assumptions and Hypotheses

- Flights are domestic, direct, and one-way.
- The waiting time between flights is the amount of time until the scheduled departure time of the next available flight to a given destination.
- The ticket price is $140 [Airline Transport Association 2002], independent of when the ticket is bought, except when we consider multiple fares.
- Pre-September 11, the average probability of a ticketholder checking in for the flight (and thus becoming a contender) was 85% [Smith et al. 1992, 9].
- The pre-September 11 average load factor was 72% [Bureau of Transportation Statistics 2000].

Complicating Factors

Each of our models attempts to take into account the current situation facing airlines:

- **The Traffic Factor**
  On average, there are fewer flights by airlines between any given locations.
- **The Security Factor**
  Security in and around airports has been heightened.
- **The Fear Factor**
  Passengers are more wary of the dangers of air travel, such as possible terrorist attacks, plane crashes, and security breaches at airports.
- **The Financial Loss Factor**
  Airlines have lost billions of dollars in revenue due to decreased demand for air travel, increased security costs, and increased industry risks.

The Traffic Factor

Because there are fewer flights, it is likely that the demand for any given flight will increase. Flights are likely to be fuller; the average waiting time between flights to a destination is likely to increase, so bumped passengers will demand higher compensation.

The Security Factor

The increase in security will likely lead to an increase in the number of ticketholders who arrive at the airport but—due to security delays—do not arrive at their departure gates in time.

Successful implementation of security measures may lead to an improvement in the public perception of the airline industry and an increase in demand for air travel.

The Fear Factor

Increased fear of flying decreases demand for air travel, so security delays may not be as serious.

On the other hand, if a higher percentage of ticketholders are flying out of necessity, then the probability that a ticketholder becomes a contender may increase because of decreased cancellations and no-shows. However, fewer ticketholders are likely to agree to be bumped voluntarily at any price, so the percentage of involuntarily bumped passengers may increase.

The Financial Loss Factor

Because companies may seek to increase short-term profits in the face of recent losses, some airlines may implement more aggressive overbooking, which could induce an overbooking war between airlines [Suzuki 2002, 148]. The likely increase in the number of bumped passengers would lead to a rise in compensation costs that would partially offset increased revenue.

Decreasing the number of bumped passengers would improve the airlines' image and might spur demand, which would bolster future revenue.

One-Plane Model

Introduction and Motivation

We first consider the optimal overbooking strategy for a single flight, independent of all other flights. We will see later that its results are a good approximation to the results of the full-fledged Interaction Simulation Model.

Development

Let the plane have a capacity of C identical seats and let a ticket cost $T = \$140$ independent of when it is bought. Let the airline's overbooking strategy be to sell up to B tickets, if possible ($B > C$). We analyze this strategy in the case when all B tickets are sold.

We model the number of contenders for the flight with a binomial distribution, where a ticketholder becomes a contender with probability p. The average p for flights from the ten leading U.S. carriers is $p = 0.85$ [Smith et al. 1992]. The value of p for a particular flight depends on a host of factors—flight time, length, destination, whether it is a holiday season—so we carry out our analysis for a range of possible p values.

With our binomial model, the probability of exactly i contenders among the B ticket-holders is $\binom{B}{i}p^i(1-p)^{B-i}$.

We assume that each bumped passenger is paid compensation $(1+k)T = 140(1+k)$, for some constant k. Translated into everyday terms, this means that a bumped passenger receives compensation equal to the ticket price T plus some additional compensation $kT > 0$. Later, we relax the assumption that compensation cost is the same for each passenger, when we consider involuntary vs. voluntary bumping.

We define the compensation cost function $F(i, C)$ to be the total compensation the airline must pay if there are exactly i contenders for a flight with seating capacity C:

$$F(i,C) = \begin{cases} 0, & i \leq C; \\ (k+1)T(i-C), & i > C. \end{cases}$$

We calculate expected revenue R as a function of B:

$$\begin{aligned} R(B) &= \sum_{i=1}^{B} \binom{B}{i} p^i (1-p)^{B-i} (BT - F(i,C)) \\ &= 140B - 140(k+1) \sum_{i=C+1}^{B} \binom{B}{i} p^i (1-p)^{B-i} (i-C) \end{aligned}$$

We use a computer program to determine, for given C, p, and k, the overbooking strategy B_{opt} that maximizes $R(B)$. However, it is also possible to produce a close analytic approximation, which we now derive.

The revenue for a bumped passenger, $T-(k+1)T = -kT$, has magnitude k times that for a boarded passenger, T. Thus, the optimal overbooking strategy is such that the distribution of contenders is in some sense "balanced," with $1/(k+1)$ of its area corresponding to bumped passengers and the remaining $k/(k+1)$ corresponding to boarded passengers.

We approximate the binomial distribution of contenders with a normal distribution:

$$\frac{C - Bp}{\sqrt{Bp(1-p)}} \approx \Phi^{-1}\left(\frac{k}{k+1}\right),$$

where Φ is the cumulative distribution function of the standard normal distribution. Clearing denominators and solving the resulting quadratic in $\sqrt{B}$ gives

$$B'_{opt} = \left(\frac{-\Phi^{-1}\left(\frac{k}{k+1}\right)\sqrt{p(1-p)} + \sqrt{\Phi^{-1}(\frac{k}{k+1})^2 p(1-p) + 4pC}}{2p} \right)^2 \qquad \textbf{(1)}$$

as an analytic approximation to B_{opt}. For $k=1$, we get $B'_{\text{opt}} = C/p$.

This analytic approximation is always within 1 of the optimal overbooking strategy for $.80 \le p \le .90$ and $1 \le k \le 3$.

Results and Interpretation

The airline should be able to obtain good approximations to p and k empirically. Thus, it can take our computer program, insert its data for C, T, p, and k, and obtain the optimal overbooking strategy B_{opt}. **Figure 1** plots expected revenue $R(B)$ vs. B $C=150$, $k=1$, $p=0.85$, and $T=140$.

At $B=177$, the airline can expect revenue $R(177) = \$24,200$, which is more than 15% in excess of the expected revenue $R(150) = \$21,000$ from a policy of no overbooking.

Operating at a less-than-optimal overbooking strategy can have serious consequences. For example, American Airlines has an annual revenue of $20 billion [AMR Corporation 2000]. An overbooking policy B outside the range of $[173, 183]$ implies an expected loss of more than $1 billion over a 5-year period compared with the expected revenue at $B_{\text{opt}} = 177$.

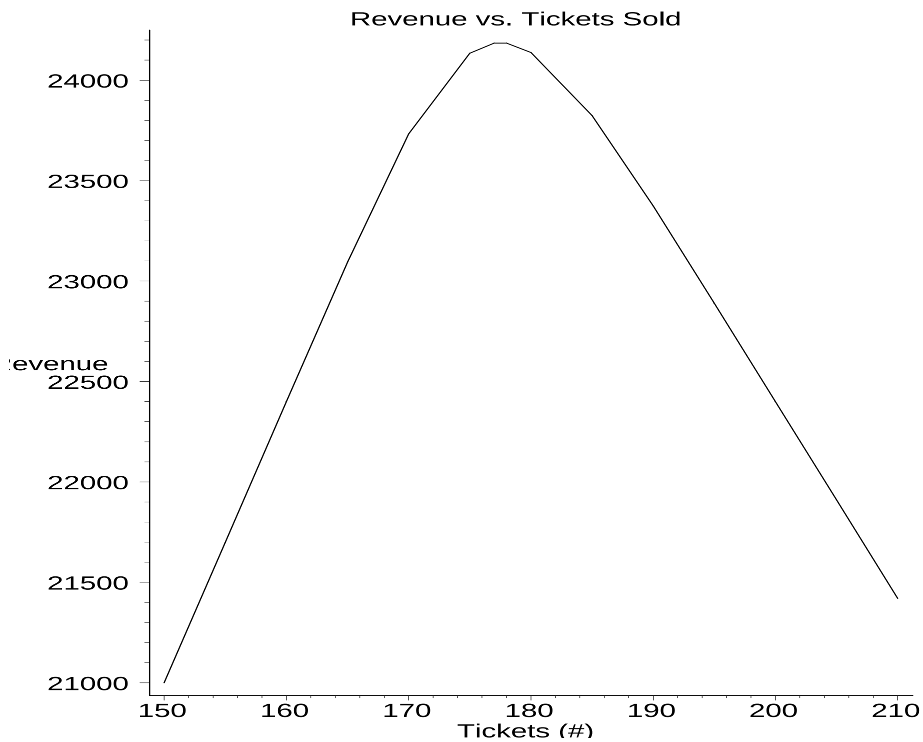

Figure 1. Revenue R vs. overbooking strategy B for $C = 150$, $k = 1$, $p = 0.85$, and $T = \$140$.

Limitations

The single-plane model

- fails to account for bumped passengers' general dissatisfaction and propensity to switch airlines;
- assumes a simple constant-cost compensation function for bumped passengers;
- ignores the distinction between voluntary and involuntary bumping;
- assumes that all tickets are identical—that is, everyone flies coach;
- assumes that all B tickets that the airline is willing to sell are actually sold.

Even so, the model successfully analyzes revenue as a function of overbooking strategy, plane capacity, the probability that ticket-holders become contenders, and compensation cost. Later, we develop a more complete model.

The Complicating Factors

First, though, we use the basic model to make preliminary predictions for the optimal overbooking strategy in light of market changes due to the complicating factors post-September 11.

Of the four complicating factors, only two are directly relevant to this model: the security factor and the fear factor. The primary effect of the security factor is to decrease the probability p of a ticketholder reaching the gate on time and becoming a contender. On the other hand, the primary effect of the fear factor is that a greater proportion of those who fly do so out of necessity; since such passengers are more likely to arrive for their flights than more casual flyers, the fear factor tends to increase p.

Figure 2 plots the optimal overbooking strategy B_{opt} vs. p for fixed $k = 1$ and $C = 150$.

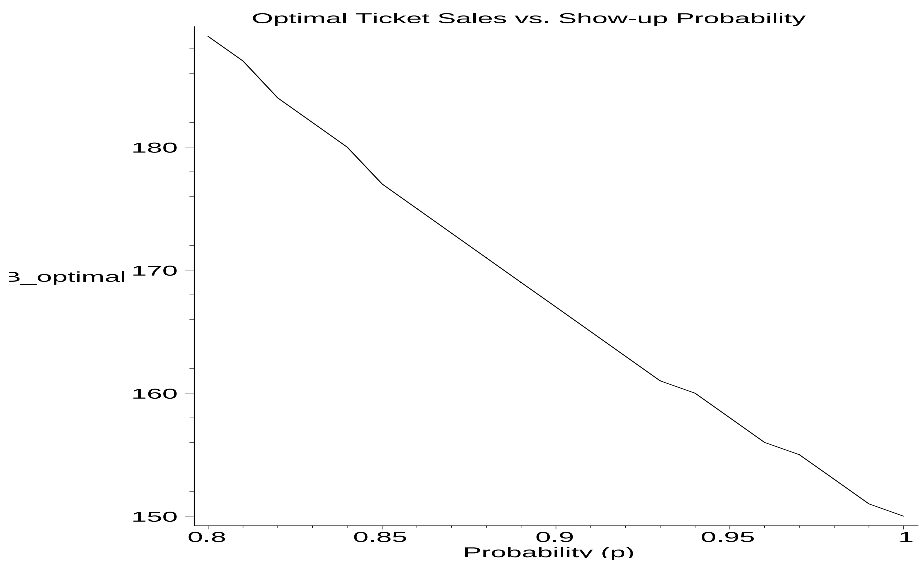

Figure 2. Optimal overbooking strategy vs. arrival probability p.

It is difficult to assess the precise change in p resulting from the security and fear factors. However, airlines can determine this empirically by gathering statistics on their flights, then use our graph or computer program to determine a new optimal overbooking strategy.

One-Plane Model: Multifare Extension

Introduction and Motivation

Most airlines sell tickets in different fare classes (most commonly first class and coach). We extend the basic One-Plane Model to account for multiple fare classes.

Development

For simplicity, we consider a two-fare system, with C_1 first-class seats and C_2 coach seats. We assume that a first-class ticket costs $T_1 = \$280$ and that a coach ticket costs $T_2 = \$140$. We consider an overbooking strategy of selling up to B_1 first class tickets and up to B_2 coach tickets, where the two types of sales are made independently of one another.

We assume that a first-class ticketholder becomes a first-class contender with probability p_1 and that a coach ticketholder becomes a coach contender with probability p_2. We use two independent binomial distributions as our model. First-class ticketholders are more likely to become contenders than coach passengers, since they have made a larger monetary investment in their tickets; that is, $p_1 > p_2$. Thus, the probabilities of exactly i first-class contenders and exactly j coach contenders are

$$\binom{B_1}{i} p_1^i (1-p_1)^{B_1-i}, \qquad \binom{B_2}{j} p_2^j (1-p_2)^{B_2-j}.$$

We model compensation costs as constant per bumped passenger but dependent on fare class, with $(k_1+1)T_1$ as compensation for a bumped first-class passenger and $(k_2+1)T_2$ for a bumped coach passenger. We define the compensation cost function:

$$F(i,j,C_1,C_2) = \begin{cases} 0, & i \le C_1,\ j \le C_2; \\ T_1(k_1+1)(i-C_1), & i > C_1,\ j \le C_2; \\ \max\{T_2(k_2+1)((j-C_2)-(i-C_1)),\ 0\}, & i \le C_1,\ j > C_2; \\ T_1(k_1+1)(i-C_1) + T_2(k_2+1)(j-C_2), & i > C_1,\ j > C_2. \end{cases}$$

The justification for the third case is that an excess of coach contenders is allowed to spill over into any available first-class seats. On the other hand, excess first-class contenders cannot be seated in any available coach seats; this fact is reflected in the second case.

We model expected revenue R as a function of the overbooking strategy (B_1, B_2):

$$R(B_1,B_2) = \sum_{i=1}^{B_1}\sum_{j=1}^{B_2} \binom{B_1}{i}\binom{B_2}{j} p_1^i (1-p_1)^{B_1-i} p_2^j (1-p_2)^{B_2-j} \cdot \left(B_1T_1 + B_2T_2 - F(i,j,C_1,C_2)\right)$$

Results and Interpretation

For fixed C_i, T_i, p_i, and k_i ($i = 1, 2$), we can find $(B_{1,\text{opt}}, B_{2,\text{opt}})$ for which $R(B_1, B_2)$ is maximal by adapting the computer program used to solve the one-fare case.

For example, for a plane with $C_1 = 20$ first class seats, $C_2 = 130$ coach seats, ticket costs of $T_1 = \$280$ and $T_2 = \$140$, and compensation constants $k_1 = k_2 = 1$, we obtain the optimal overbooking strategies listed in **Table 2**.

The optimal strategy involves relatively little overbooking of first-class passengers, since there is a much higher compensation cost. However, the total

Table 2.

Two-fare optimal overbooking strategies for selected arrival probabilities.

p_1	p_2	$B_{1,\text{opt}}$	$B_{2,\text{opt}}$
0.85	0.80	23	165
0.90	0.80	22	165
0.95	0.80	20	166
0.85	0.85	23	155
0.90	0.85	22	155
0.95	0.85	20	155
0.90	0.90	22	146
0.95	0.90	21	145

number of passengers (coach plus first-class) overbooked in an optimal two-fare situation is virtually the same as the total number overbooked in the one-fare situation. The upshot is that the effect of multiple fare classes on the optimal overbooking strategy is not very significant; so, when we construct our more general model, we do not take into account multiple fares.

Compensation Costs

The key element that separates different schemes for compensating bumped ticketholders is the degree of choice for the passenger. Airlines often hold auctions for contenders in which the lowest bids are first to be bought off of a flight.

We construct a model for involuntary bumping costs that is based on DOT regulations and takes into account the waiting time distribution for flights. Then we discuss auction methods for voluntary bumping and derive novel results for expected compensation cost for a continuous auction that matches actual ticket auctions fairly well.

Involuntary Bumping: DOT Regulations

The Department of Transportation (DOT) requires each airline to give all passengers who are bumped involuntarily a written statement describing their rights and explaining how the airline decides who gets on an overbooked flight and who does not [Department of Transportation 2002]. Travelers who do not get to fly are usually entitled to an "on-the-spot" payment of denied boarding compensation. The amount depends on the price of their ticket and the length of the delay:

- Passengers bumped involuntarily for whom the airline arranges substitute transportation scheduled to get to their final destination within one hour of their original scheduled arrival time receive no compensation.

- If the airline arranges substitute transportation scheduled to arrive at the destination between one and two hours after the original arrival time, the airline must pay bumped passengers an amount equal to their one-way fare, with a $200 maximum.

- If the substitute transportation is scheduled to get to the destination more than two hours later, or if the airline does not make any substitute travel arrangements for the bumped passenger, the airline must pay an amount equal to the lesser of 200% of the fare price and $400.

- Bumped passengers always get to keep their tickets and use them on another flight. If they choose to make their own arrangements, they are entitled to an "involuntary refund" for their original ticket.

These conditions apply only to domestic flights and not to planes that hold 60 or fewer passengers.

The function for the compensation cost for an involuntarily bumped passenger is

$$C(T,F) = \begin{cases} 0, & \text{if } 0 < T \leq 1; \\ \min(2F, F+200), & \text{if } 1 < T \leq 2; \\ \min(3F, F+400), & \text{if } 2 < T, \end{cases}$$

where T is waiting time and F is the fare price. We assume that all flights to a given location are direct and have the same flight duration. Thus, the waiting time between flights equals the difference in departure times, and the waiting time T is the time until the next flight to the destination departs. We assume that involuntarily bumped passengers always ask for a refund of their fare.

Involuntary Bumping: The Waiting Time Model

To use the compensation cost function to determine the average compensation (per involuntarily bumped passenger), we would need to know the joint distribution of fare prices and waiting times. Because this information would be extremely difficult to obtain, we opt instead for practical compromises:

- We restrict our attention to determining the expected compensation cost for the average ticket price, $140 [Airline Transport Association 2000].

- We specify a workable model for the distribution of waiting times that allows us to calculate this cost directly.

Our model for the distribution of waiting times is the exponential distribution, a common distribution for waiting times. Let T be a random variable representing waiting time between flights; then

$$Pr(T \leq t) = 1 - e^{-\lambda t}$$

and $E(T) = \tau = 1/\lambda$, where τ is the mean waiting time for the next available flight.

The expected cost of compensating an involuntarily bumped passenger who purchased a ticket of price P can be evaluated directly and is

$$\min(2P, P + 200)\left[e^{-\lambda} - e^{-2\lambda}\right] + \min(3P, P + 400)\left[e^{-2\lambda}\right].$$

From examining airline booking sites, we estimate the average daytime waiting time τ to be 2.6 h, not including the time between the last flight of the day and the first flight of the next day. If we include these night-next-day waiting times in our calculations, we obtain $\tau \approx 4.8$ h; this value corresponds to five flights per 24-hour period, which is fairly typical. Using the smaller, strictly daytime value $\tau = 2.6$ h, we obtain an expected compensation cost of \$255.

Voluntary Bumping: Auction Methods

In 1968, J.L. Simon proposed an auction among ticketed passengers. Each ticketed passenger contending for a seat on a flight would submit a sealed envelope bid of the smallest amount of money for which the contender would give up the seat and wait until the next available one. The airline would compensate the passengers who required the least money and require that they give up their seats. Passengers would never get bumped without suitable compensation, and airlines could raise their overbooking level much higher than they could otherwise. After Ralph Nader successfully sued Allegheny Airlines for bumping him, variants on this scheme have gradually become standard throughout the industry.

There are two reasonable ways to attempt an auction.

- Per Simon, force every contender to choose a priori a price for which they would give up their ticket. The airline could arrange all bumpings immediately.
- The actual practice by most airlines is to announce possible compensation prices in discrete time intervals. Customers can then accept any offer they wish to.

The first is attractive to the airlines because it is instant and minimizes compensation. The second, however, can be started well before a flight departs; and if intervals are increased gradually enough, the difference in cost is negligible. The methods should generate similar results, so for simplicity we concentrate on the second, though with continuous compensation offerings.

Voluntary Bumping: Continuous-Time Auction

In the literature, it is common to assume that if m passengers are compensated through an auction, the total cost for the airline should be linear in m, although some authors (such as Smith et al. [1992]) recognize that the function should be nonlinear and convex but do not analyze it further. In fact, we can say a great deal more with only a few basic assumptions. Indeed, suppose that

- n ticketholders check in for a flight with capacity C, with $n > C$.
- Each contender has a *limit price*, the smallest compensation to be willing to give up the seat.
- An airline can always rebook a ticketholder on one of its own later flights at no cost (i.e., it does not have to pay for a ticket on a rival airline).

In an ideal auction, the airline offers successively higher compensation prices; whenever the offer exceeds a contender's limit price, the contender gives up the ticket voluntarily. Suppose that ticketholders $(\Gamma_1, \Gamma_2, \ldots \Gamma_n)$ are ordered so that Γ_i's limit price is less than Γ_j's limit price for $i < j$. Define:

- $D(x)$ = the probability that a randomly selected ticketholder gives up the seat for a price x.
- Y_m = the compensation that the airline must pay Γ_m to give up the ticket.
- X_m = the total compensation that the airline must pay for m contenders give up their seats.

We have $X_m = \sum_{i=1}^{m} Y_i$. To determine $E[X_m]$, we determine $E[Y_i]$ for $i \leq m$. To do this, we need the following result:

$$E[Y_m] = \sum_{i=0}^{m-1} \binom{n}{i} \int_0^\infty \Big(D(x)\Big)^m \Big(1 - D(x)\Big)^{n-m} dx.$$

[EDITOR'S NOTE: We omit the authors' proof.]

Very little can be done beyond this point without further knowledge about the nature of $D(x)$. There is not much recent data on this; but when airlines were first considering moving to an auction-based system, K.V. Nagarajan [1978] polled airline passengers on their limit price. Although he performed little analysis, we find that the cumulative distribution function of this limit price fits very closely exponential curves of the form $1 - e^{-Ax}$ for a fixed A (**Figure 3**).

With $D(x) = 1 - e^{-Ax}$ for some constant A, then

$$E[X_m] = \frac{1}{A}\left[m - (n-m)\left(\frac{1}{n} + \frac{1}{n-1} + \frac{1}{n-2} + \ldots + \frac{1}{n-m+1}\right)\right].$$

[EDITOR'S NOTE: We omit the authors' proof.]

Using the approximation

$$\frac{1}{1} + \frac{1}{2} + \ldots + \frac{1}{n} \approx \ln n,$$

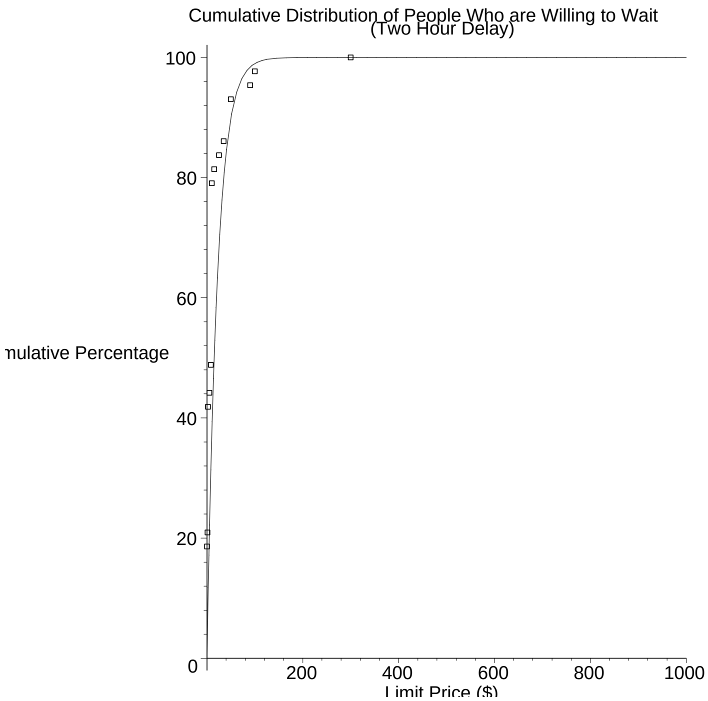

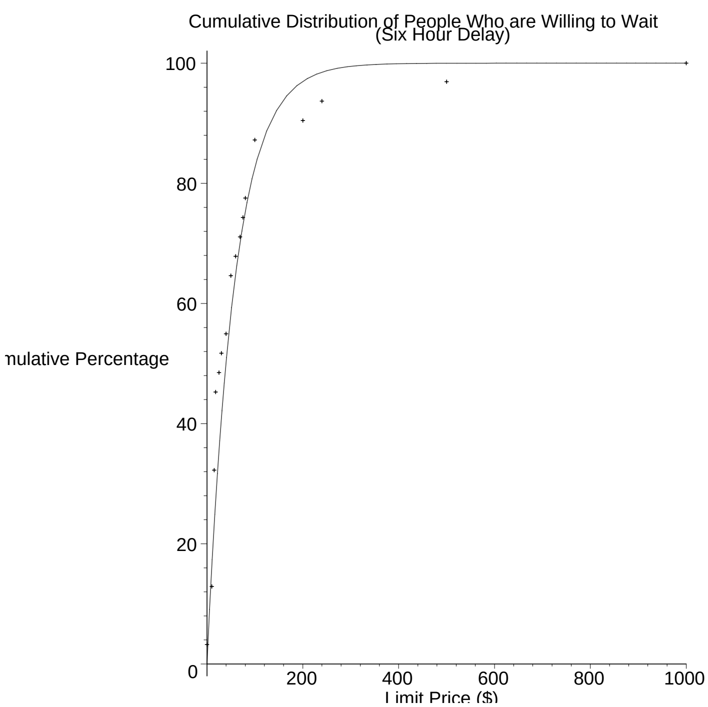

Figure 3. Polled distribution of ticketholder limit price, with best fit graphs $1 - e^{0.046x}$ for 2-hour wait and $1 - e^{0.0175x}$ for 6-hour wait (data from [Nagarajan 1978, 113]).

this becomes

$$E[X_m] \approx \frac{1}{A}\left[m - (n-m)\ln\left(\frac{n}{n-m}\right)\right].$$

There is no reason to believe that the value of A is constant across all scenarios. For example, contenders will certainly accept a smaller compensation if the next flight is departing soon. For our purposes, however, we assume that A is constant over all situations; and we estimate that on a flight with capacity $C = 150$ and only a small number of overbooked passengers, Γ_1 has a limit price of \$100. Then we have $\frac{1}{A} \cdot \frac{1}{150} \approx \$100$, so $A \approx \$1/15,000$.

Hence, the expected compensation required to bump m out of n ticketholders via auction is approximately

$$\frac{\$1}{15,000}\left[m - (n-m)\ln\left(\frac{n}{n-m}\right)\right],$$

compared to a cost of $\$255m$ (plus ill will) for involuntary bumping the same number of ticketholders.

Effects of Overbooking on Market Share

Constructing the Model

We focus on the 10 largest U.S. airlines (Alaska, America West, American, Continental, Delta, Northwest, Southwest, Trans World, United, US Air), which comprise 90% of the market. We use 1997–1998 statistics on their flight frequency and market share. [EDITOR'S NOTE: We omit the data table.]

Flights are modeled as identical in all respects except for market interest. The market is simulated as a group of initially 10,000 people, each loyal to one airline, who independently buy tickets on their airline with a fixed probability and meet reservations with a fixed probability. Each member of the market independently chooses to stay with an airline or change airline based on treatment regarding each flight.

Each company choose a number r, which specifies its overbooking strategy: On a flight of capacity C, the company will sell up to $B = Cr$ tickets.

In each time period, precisely one flight is offered. The chance that a given airline will offer that flight is proportional to the number of flights that it offers per year. We also determine a constant k that indicates the level of interest in this flight. Each flight has capacity $C = 150$ seats each sold at \$140.

The exact size of the market should have little effect on the result. We assume that the total market is initially made up of 10,000 independent people, each loyal to one carrier. The relative sizes of the company market shares are initialized according to 1997–98 industry data. We assume each person in the market flies on average the same number of times in a year.

We assume that each person in a company's market has probability k of wanting to buy a ticket for a flight by the company. We have k follow a normal distribution with mean fixed so that the average load factor on all flights is the industry average of 0.72 [Bureau of Transportation Statistics 2002].

Industry data prior to September 11 indicate a probability of .85 that a ticketholder will check in for the flight.

If necessary, each airline bumps some passengers voluntarily and some involuntarily, according to its strategy. The immediate cost of bumpings is set to the values that we derived in the previous section. We surmise that voluntarily bumped passengers are relatively happy and thus leave the airline with probability only .05, whereas involuntarily bumped passengers are furious and leave with probability .8.

A person who leaves an airline stays within the market with probability .9 (0.95 if bumped voluntarily) and simply switches to another airline; otherwise, the person leaves the market altogether. People trickle into the market fast enough to compensate for the loss of people due to dissatisfaction, thus allowing the market to grow slowly.

Simulation Results, Pre-September 11

We investigate the effect of different overbooking rates on profit. For each overbooking rate, we calculate net profit over 500 time periods (ensuring that the same random events occur regardless of the strategy tested). The strategy that maximizes profit for that time period is then determined and tabulated. We repeat this 40 times for each airline.

This leaves open the question of what strategies the companies not being tested should use. To determine this, we initially assume that each company would overbook by 1.17 (as computed in the single-plane model), run the program to get a first estimate of a good strategy, and use the optimal results from that preliminary run to set the default overbooking rates of each company in a final run. Finally, we use the industry figure that 5% of all bumped passengers are bumped involuntarily to set the company compensation strategies.

The optimal overbooking rate for all companies other than Alaska is between 1.165 and 1.176, close to but a little less than the results from the One-Plane Model. This is reasonable, since the most significant improvement that this simulation makes over the One-Plane Model is the consideration of lost customers, whose effect should slightly reduce the optimal overbooking rate.

The program generates very consistent answers on each run for every airline except Alaska. Alaska has far fewer passengers per flight than its competitors and rarely fills any plane entirely, so its overbooking policy has a negligible effect on its overall profit. Thus, the simulation is almost certainly too coarse to generate useful data on Alaska.

Adjusting the Model Due to September 11

We estimate the effects of the complicating factors after September 11 have on the simulation parameters:

- *Arrival probability p increases from* 0.85 *to* 0.90.
- *Flight frequency decreases by* 20% [Parker 2002].
- *Total arket size decreases by 15%*. Fourth quarter data from 2001 are not yet available, so we make an estimate. Our own experience is that flights are more crowded now, which suggests that the percentage of market size decrease is smaller than the percentage of flight frequency decrease. Thus, we estimate that market size has decreased by 15%.
- *Market return rate doubles*. The market size has decreased due to the fear factor, but Parker [2002] anticipates that demand will return to pre-September 11 levels by mid-2002. Moreover, public perception of airline safety is improving due to the security factor. Thus, the market return rate should be substantially higher than its pre-September 11 level.
- *Market exit rate decreases by 50%*. The market composition is now more heavily weighted towards those who fly only out of necessity; such fliers are much less likely than casual fliers to leave the market.
- *Percentage of bumps that are voluntary decreases from 95% to 90%*. There are fewer flights, hence the waiting time between flights is greater. Since passengers are more likely to be flying of necessity, they are much less interested in giving up a seat for compensation.
- *Compensation cost of voluntary bumping increases by 20%*.
- *Compensation cost of involuntary bumping increases by 20%*. Bumped passengers face longer waiting times; because of DOT regulations, average involuntary compensation costs must rise.
- *Competitors increase their overbooking levels from r to $r + 0.02$*. Due to financial losses, an airline can expect its competitors to focus more heavily on short-term profits than previously.

Simulation Results, Post-September 11

Using the parameter changes outlined, we ran the simulation again to estimate the effect of the events of September 11 on optimal overbooking strategies. The results are shown in **Table 3**.

There is again a strong correlation between the simulation results for these parameters and the corresponding results from the One-Plane Model.

From **Table 3**, it is clear that the events of September 11 have indeed had a significant effect on optimal overbooking rates. Indeed, for a company the

Table 3.
Optimal overbooking rates, from simulation results.

Airline	Pre-September 11	Post-September 11
Alaska	1.319	1.260
America West	1.169	1.094
American	1.171	1.094
Continental	1.170	1.096
Delta	1.173	1.095
Northwest	1.174	1.095
Southwest	1.173	1.095
Trans World	1.176	1.096
United	1.168	1.094
US Air	1.165	1.092

size of American Airlines, the 7% change in these rates could easily lead to a difference in profits on the order of $1 billion.

Thus, if our estimates of parameter changes due to September 11 are reasonable, *all major airlines should significantly decrease their overbooking rates*.

References

Air Transport Association. 2001. `http://www.air-transport.org/public/industry/display1.asp?nid=1138` .

AMR Corporation. 2000. `http://www.amrcorp.com/ar2000/fin_highlights.html` .

Bureau of Transportation Statistics. 2000. `http://www.bts.gov/oai/aviation_industry/BlueBook_Dec200.pdf` .

Department of Transportation. 2002. `http://www.dot.gov/airconsumer/flyrights.htm#overbooking` .

Infoplease. 2002. `http://www.infoplease.com/ipa/A077823.html` .

Nagarajan, K.V. 1979. On an auction solution to the problem of airline overbooking. *Transportation Research* A13: 111–114.

Parker, James D. 2001. Growth airline outlook even more favorable in new airline environment. `http://www.raymondjames.com/invbrf/airline.htm` .

Smith, Barry C., John F. Leimkuhler, and Ross M. Darrow. 1992. Yield management at American Airlines. *Interfaces* 22 (1) (January-February): 8–31.

Suzuki, Yoshinori. 2002. An empirical analysis of the optimal overbooking policies for U.S. major airlines. *Transportation Research* E38: 135–149.

Memorandum

Attn: Don Carty, CEO American Airlines
From: MCM Team 180
Subject: Overbooking Policy Assessment Results

We completed the preliminary assessment of overbooking policies that you requested. There is a great deal of money at stake here, both from ticket sales and also from compensation that must be given to bumped passengers. Moreover, if too many passengers are bumped, there will be a loss of good will and many regular customers could be lost to rival airlines. In fact, we found that the profit difference for American Airlines between a good policy and a bad policy could easily be on the order of $1 billion a year.

Using a combination of mathematical models and computer simulations, we considered a wide variety of possible strategies that could be tried to confront this problem. We naturally considered different levels of overbooking, but we also looked at different ways in which airlines could compensate bumped passengers. In terms of the second question, we find that the current scheme of auctioning off compensations for tickets, combined with certain calculated forced bumpings, is still ideal, regardless of changes to the market state.

Although we were forced to work without much recent data, we were also able to achieve reliable and consistent results for the optimal overbooking rate. In particular, we found that prior to September 11, American Airlines stood to maximize profits by selling approximately 1.171 times as many tickets as seats available.

We next considered how this number would likely be affected by the current state of the market. In particular, we focused on four consequences of the events on September 11: all airlines are offering fewer flights, there is heightened security in and around airports, passengers are afraid to fly, and the industry has already lost billions of dollars. Analyzing each of these in turn, we found that they did indeed have a significant effect on the market. In particular, American Airlines should lower its overbooking rate to 1.094 tickets per available seat.

In conclusion, we found that there is indeed a tremendous need to re-evaluate the current overbooking policy. According to our current data, we believe that the rate should be dropped significantly. It would be valuable, however, to supplement our calculations with some of the confidential data that American Airlines has access to, but that we do not.

Models for Evaluating Airline Overbooking

Michael P. Schubmehl
Wesley M. Turner
Daniel M. Boylan
Harvey Mudd College
Claremont, CA

Advisor: Michael E. Moody

Introduction

We develop two models to evaluate overbooking policies.

The first model predicts the effectiveness of a proposed overbooking scheme, using the concept of expected marginal seat revenue (EMSR). This model solves the discount seat allocation problem in the presence of overbooking factors for each fare class and evaluates an overbooking policy stochastically.

The second model takes in historical flight data and reconstructs what the optimal seat allocation should have been. The percentage of overbooking revenue obtained in practice serves as a measure of the policy's value.

Finally, we examine the overbooking problem in light of the recent drastic changes to airline industry and conclude that increased overbooking would bring short-term profits to most carriers. However, the long-term ill effects that have traditionally caused airlines to shun such a policy would be even more pronounced in a post-tragedy climate.

Literature Review

There are two major ways that airlines try to maximize revenues: overbooking (selling more seats than available on a given flight) and seat allocation (price discrimination). These measures are believed to save major airlines as much as half a billion dollars each year, in an industry with a typical yearly profit of about $1 billion dollars [Belobaba 1989].

The UMAP Journal 23 (3) (2002) 301–316.

Beckman [1958] models booking and no-shows in an effort to find an optimal overbooking strategy. He ignores advance cancellations, assuming that all cancellations are no-shows [Rothstein 1985]. A method that is easy to implement but sophisticated enough to allow for cancellations and group reservations was developed by Taylor [1962]. Versions of this model were implemented at Iberia Airlines [Shlifer and Vardi 1975], British Overseas Airways Corporation, and El Al Airlines [Rothstein 1985].

None of these approaches considers multiple fare classes. Littlewood [1972] offers a simple two-fare allocation rule: A discount fare should be sold only if the discount fare is greater than or equal to the *expected marginal return* from selling the seat at full-fare. This idea was generalized by Belobaba [1989] to include any number of fare classes and allow the integration of overbooking. We use expected marginal seat revenue in predicting about overbooking schemes.

There is a multitude of work on the subject [McGill 1999]—according to Weatherford and Bodily [1992], there are more than 124,416 classes of models for variations of the yield management problem, though research has settled into just a few of these. Several authors tried to better Belobaba's [1987] heuristic in the presence of three or more fare classes (for which it is demonstrably suboptimal) [Weatherford and Bodily 1992]; generally, these adaptive methods for obtaining optimal booking limits for single-leg flights are done by dynamic programming [Mcgill 1999].

After deregulation in 1978, airlines were no longer required to maintain a direct-route system to major cities. Many shifted to a hub-and-spoke system, and network effects began to grow more important. To maximize revenue, an airline may want to consider a passenger's full itinerary before accepting or denying their ticket request for a particular leg. For example, an airline might prefer to book a discount fare rather than one at full price if the passenger is continuing on to another destination (and thus paying an additional fare).

The first implementations of the origin-destination control problem considered segments of flights. The advantage to this was that a segment could be blacked out to a particular fare class, lowering the overall complexity of a booking scheme. Another method, *virtual nesting*, combines fare classes and flight schedules into distinct buckets [Mcgill 1999]. Inventory control on these buckets would then give revenue-increasing results. Finally, the bid-price method deterministically assigns a value to different seats on a flight leg. The legs in an itinerary are then summed to establish a bid-price for that itinerary; a ticket request is accepted only if the fare exceeds the bid-price [Mcgill 1999]

The most realistic yield management problem takes into account five price classes. The ticket demands for different fare classes are randomized and correlated with one other to allow for sell-ups and the recapture of rejected customers on later flights. Passengers can no-show or cancel at any time. Group reservations are treated separately from individuals—their cancellation probability distribution is likely different. Currently, most work assumes that passengers who pay full fare would not first check for availability of a lower-class ticket; a more realistic model would allow buyers of a higher-class ticket to be diverted by a lower fare. A full accounting of network effects would consider the relative value of what Weatherford and Bodily [1992] terms *displacement*—denying a discount passenger's ticket request to fly a multileg itinerary in favor of leaving one of the legs open to a full-fare passenger.

Unfortunately, while the algorithms for allocating seats and setting overbooking levels are highly developed, there has been little work done on the problem of evaluating how effective these measures actually are. Our solution applies industry-standard methods to find optimal booking levels, then examines the actual booking requests for a given flight to determine how close to an optimal revenue level the scheme actually comes.

Factors Affecting Overbooking Policy

General Concerns

The reason that overbooking is so important is because of multiple fare classes. With only one fare class, it would be easier for airlines to penalize customers for no-shows. However, while most airlines offer nonrefundable discount tickets, they prefer not to penalize those who pay full fare, like business travelers, because these passengers account for most of the profits.

The overbooking level of a plane is dictated by the likelihood of cancellations and of no-shows. An overbooking model compares the revenue generated by accepting additional reservations with the costs associated with the risk of overselling and decides whether additional sales are advisable. In addition, the "recapture" possibility can be considered, which is the probability that a passenger denied a ticket will simply buy a ticket for one of the airline's other flights. Since a passenger is more valuable to the airline buying a ticket on a flight that has empty seats to fill than on one that is already overbooked, a high recapture probability reduces the optimal overbooking level [Smith et al. 1992].

No major airline overbooks at profit-maximizing levels, because it could not realistically handle the problems associated with all the overloaded flights. This gives the overbooking optimization problem some important constraints. The total flight revenue is to be maximized, subject to the conditions that only a certain portion of flights have even one passenger denied boarding (one oversale), and that a bound is placed on the expected total number of oversales. Dealing with even one oversale is a hassle for the airline's staff, and they are not equipped to handle such problems on a large scale. Additionally, some research indicates that increased overbooking levels would most likely trigger an "overbooking war" [Suzuki 2002], which would increase short-term profits but would probably engender enough consumer resentment that the industry as a whole would lose business.

While the overbooking problem sets a limit for sales on a flight as a whole, proper seat allocation sets an optimal point at which to stop selling tickets for individual fare levels. For example, a perfectly overbooked plane, loaded exactly to capacity, could be flying at far below its optimal revenue level if too many discount tickets were sold. The more expensive tickets are not for first-class seats and involve no additional luxuries above the discount tickets, apart from more lenient cancellation policies and the ability to buy the tickets a shorter time before the flight's departure.

September 11, 2001

Since the September 11 terrorist attacks, there have been significant changes in the airline business. In addition to the forced cancellation of many flights in the immediate aftermath of the attacks and the extreme levels of cancellations and no-shows by passengers after that, passenger traffic has dropped sharply in general. The huge downturn in passenger levels has led to large reductions in service by most carriers.

In terms of the booking problem, there are fewer flights for passengers to spill over onto, which could increase the loss by an airline if it bumps a passenger from a flight. On the other hand, since passenger levels have fallen so far, it is less likely that an airline will overfill any given flight. The heightened security at airports will likely increase the passenger no-show rate somewhat, as passengers get delayed at security checkpoints. At the very least, it should almost completely remove the problem of "go-shows," passengers who show up for a flight but are not in the airline's records.

On the whole, optimal booking strategies have become even more vital as airlines have already lost billions of dollars, and some teeter on the brink of failure. Some overbooking tactics previously dismissed as too harmful in the long run might be worthwhile for companies in trouble. For example, an airline near failure might increase the overbooking rate to the level that maximizes revenue, without regard to the inconvenience and possible future resentment of its customers.

Constructing the Model

Objectives

A scheme for evaluating overbooking policies needs to answer two questions: how well should a *new* overbooking method perform, and how well is a *current* overbooking scheme already working? The first is best addressed by a simple model of an airline's booking procedures; given some setup for allocating seats to fare classes, candidate overbooking schemes can be laid on top and tested by simulation. This approach has the advantage that it provides insight into *why* an overbooking scheme is or is not effective and helps to illuminate the characteristics of an optimal overbooking approach.

The second question is, in some respects, a simpler one to answer. Given the actual (over)booking limits that were imposed on each fare class, and all available information on the actual requests for reservations, how much revenue might have been gained from overbooking, compared to how much actually was? This provides a very tangible measure of overbooking performance but very little insight into the reasons for results.

The enormous number of factors affecting the design and evaluation of an overbooking policy force us to make simplifying assumptions to construct models that meet both of these goals.

Assumptions

- **Fleet-wide revenues can be near-optimized one leg at a time.**

  Maximizing revenue involves complicated interactions between flights. For instance, a passenger purchasing a cheap ticket on a flight into a major hub might actually be worth more to the airline than a business-class passenger, on account of connecting flights. We assume that such effects can be compensated for by placing passengers into fare classes based on revenue potential rather than on the fare for any given leg. This assumption effectively reduces the network problem to a single-leg optimization problem.

- **Airlines set fares optimally.**

  Revenue maximization depends strongly on the prices of various classes of tickets. To avoid getting into the economics of price competition and supply/demand, we assume that airlines set prices optimally. This reduces revenue maximization to setting optimal fare-class (over)booking limits.

- **Historical demand data are available and applicable.**

  The model needs to estimate future demand for tickets on any given flight. We assume that historical data are available on the number of tickets sold any given number of days t before a flight's departure. In some respects, this assumption is unrealistic because of the problem of data *censorship*—that is, the failure of airlines to record requests beyond the booking limit for a fare class [Belobaba 1989]. On the other hand, statistical methods can be used to reconstruct this information [Boeing Commercial Airline Company 1982, 7–16; Swan 1990].

- **Low-fare passengers tend to book before high-fare ones.**

  Discount tickets are often sold under advance purchase restrictions, for the precise reason that it enables price discrimination. Because of restrictions like these, and because travelers who plan ahead search for cheap tickets, low-fare passengers tend to book before high-fare ones.

Predicting Overbooking Effectiveness

Disentangling the effects of overbooking, seat allocation, pricing schemes, and external factors on revenues of an airline is extremely complicated. To isolate the effects of overbooking as much as possible, we want a simple, well-understood seat allocation model that provides an easy way to incorporate various overbooking schemes. In light of this objective, we pass up several methods for finding optimal booking limits on single-leg flights detailed in, for example, Curry [1990] and Brumelle [1993], in favor of the simpler expected marginal seat revenue (EMSR) method [Belobaba 1989].

EMSR was developed as an extension of the well-known rule of thumb, popularized by Littlewood [1972], that revenues are maximized in a two-fare system by capping sales of the lower-class ticket when the revenue from selling an additional lower-class ticket is balanced by the *expected* revenue from selling the same seat as an upper-class ticket. In the EMSR formulation, any number of fare classes are permitted and the goal is "to determine how many seats *not to sell* in the lowest fare classes and to retain for *possible* sale in higher fare classes closer to departure day" [Belobaba 1989].

The only information required to calculate booking levels in the EMSR model is a probability density function for the number of requests that will arrive before the flight departs, in each fare class and as a function of time. For simplicity, this distribution can be assumed to be normal, with a mean and standard deviation that change as a function of the time remaining. Thus, the only information an airline would need is a historical average and standard deviation of demand in each class as a function of time. Ideally, the information would reflect previous instances of the particular flight in question. Let the mean and standard deviations in question be denoted by $\mu_i(t)$ and $\sigma_i(t)$ for each fare class $i = 1, 2, \ldots, k$. Then the probability that demand is greater than some specified level S_i is given by

$$\bar{P}_i(S_i, t) \equiv \frac{1}{\sqrt{2\pi}\,\sigma_i(t)} \int_{S_i}^{\infty} e^{(r-\mu_i(t))^2/\sigma_i(t)^2}\, dr.$$

This *spill probability* is the likelihood that the S_ith ticket would be sold if offered in the ith category. If we further allow $f_i(t)$ to denote the expected revenue resulting from a sale to class i at a time t days prior to departure, we can define

$$\mathrm{EMSR}_i(S_i, t) = f_i(t) \cdot \bar{P}_i(S_i, t),$$

or simply the revenue for a ticket in class i times the probability that the S_ith seat will be sold. The problem, however, is to find the number of tickets S_j^i that should be protected from the lower class j for sale to the upper class i (ignoring other classes for the moment). The optimal value for S_j^i satisfies

$$\mathrm{EMSR}_i(S_j^i, t) = f_j(t), \tag{1}$$

so that the expected marginal revenue from holding the S_j^ith seat for class i is exactly equal to (in practice, slightly greater than) the revenue from selling it immediately to someone in the lower class j. The booking limits that should be enforced can be derived easily from the optimal S_j^i values by letting the booking limit B_j for class j be

$$B_j(t) = C - S_j^{j+1} - \sum_{i<j} b_i(t), \tag{2}$$

that is, the capacity C of the plane, less the protection level of the class above j from class j and less the total number of seats already reserved. Sample EMSR curves, with booking limits calculated in this fashion, are shown in **Figure 1**.

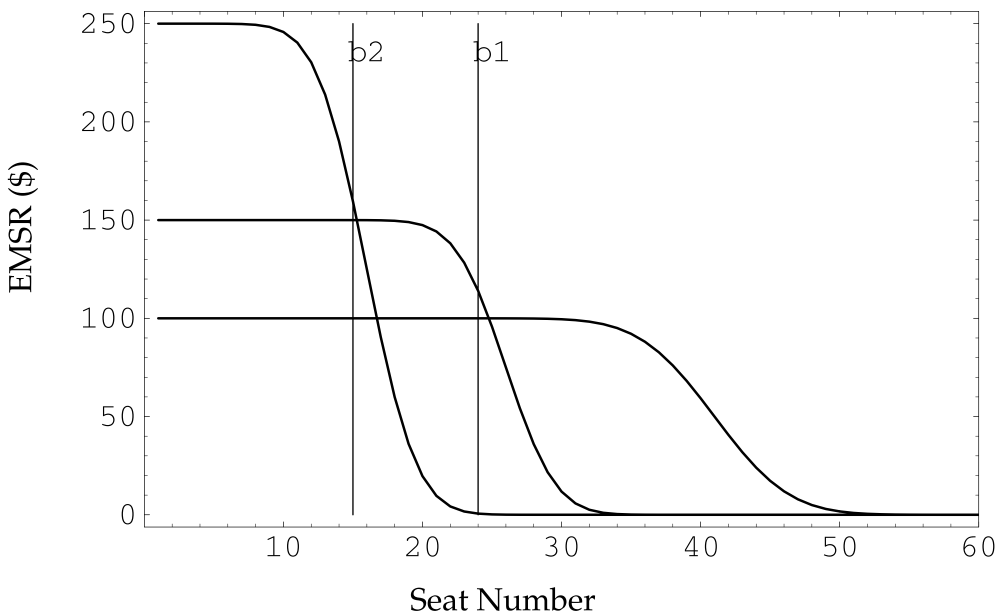

Figure 1. Expected marginal seat revenue (EMSR) curves for three class levels, with the highest-revenue class at the top. Each curve represents the revenue expected from protecting a particular seat to sell to that class. Also shown are the resulting booking limits for each of the lower classes—that is, the levels at which sales to the lower class should stop to save seats for higher fares.

This formulation does not account for overbooking; if we allow each fare class i to be overbooked by some factor OV_i, the optimality condition **(1)** becomes

$$\text{EMSR}_i(S_j^i, t) = f_j(t) \cdot \frac{OV_i}{OV_j}. \tag{3}$$

This can be understood in terms of an adjustment to f_i and f_j; the overbooking factors are essentially cancellation probabilities, so we use each OV_i to deflate the expected revenue from fare class i. Then

$$\bar{P}_i(S_j^i, t) \cdot \frac{f_i(t)}{OV_i} = f_j(t) \cdot \frac{f_j(t)}{OV_j},$$

which is equivalent to **(3)**. Note that the use of a single overbooking factor for the entire cabin (that is, $OV_i = OV$) causes the OV_i and OV_j in **(3)** to cancel. Nonetheless, the boarding limits for each class are affected, because the capacity of the plane C must be adjusted to account for the extra reservations, so now

$$C^* = C \cdot OV$$

and the booking limits in **(2)** are adjusted upward by replacing C with C^*.

The EMSR formalism gives us the power to evaluate an overbooking scheme theoretically by plugging its recommendations into a well-understood stable model and evaluating them. Given the EMSR boarding limits, which can be updated dynamically as booking progresses, and the prescribed overbooking factors, a simulated string of requests can be handled. Since the EMSR model involves only periodic updates to establish limits that are fixed over the course of a day or so, a set of n requests can be handled with two lookups each (booking limit and current booking level), one subtraction, and one comparison; so all n requests can be processed on $\mathcal{O}(n)$ time. An EMSR-based approach would thus be practical in a real-world real-time airline reservations system, which often handles as many as 5,000 requests per second. Indeed, systems derived from EMSR have been adopted by many airlines [Mcgill 1999].

Evaluating Past Overbookings

The problem of evaluating an overbooking scheme that has already been implemented is somewhat less well studied than the problem of theoretically evaluating an overbooking policy. One simple approach, developed by American Airlines in 1992, measures the optimality of overbooking and seat allocation separately [Smith et al. 1992]. Their overbooking evaluation process assumes optimal seat allocation and, conversely, their seat allocation evaluation scheme assumes optimal overbooking. Under this assumption, an overbooking scheme is evaluated by estimating the revenue under *optimal* overbooking in two ways:

- If a flight is fully loaded and no passenger is denied boarding, the flight is considered to be optimally overbooked and to have achieved maximum revenue.

- If n passengers are denied boarding, the money lost due to bumping these passengers is added back in and the n lowest fares paid by passengers for the flight are subtracted from revenue.

- On the other hand, if there are n empty seats on the plane, the n highest-fare tickets that were requested but not sold are added to create the maximum revenue figure.

Their seat-allocation model estimates the demand for each flight by calculating a theoretical demand for each fare class and then setting the minimum flight revenue (by filling the seats lowest-class first) and the maximum flight revenue (by filling the seats highest-class first). To estimate demand, we use the information on the flight's sales up to the point where each class closed. By assuming that demand is increasing for each class, we can project the number of requests that would have occurred had the booking limits been disregarded.

Given these projected additional requests and the actual requests received before closing, it is straightforward to compute the best- and worst-case overbooking scenarios. The worst-case revenue R_- is determined by using no booking limits and taking reservations as they come, and the best-case revenue R_+ is determined by accommodating high-fare passengers first, giving the leftovers to the lower classes. The difference between these two figures is the revenue to be gained by the use of booking limits. Thus, the performance of a booking scheme that generates revenue R is

$$p = \frac{R - R_-}{R_+ - R_-} \cdot 100\%, \qquad \textbf{(4)}$$

representing the percentage of the possible booking revenue actually achieved.

We select this method for evaluating booking schemes after the fact.

Analysis of the Models

Tests and Simulations

The EMSR method requires information on demand as a function of time. Although readily available to an airline, it is not widely published in a detailed form. Li [2001] provides enough data to construct a rough piecewise-linear picture of demand remaining as a function of time, shown in **Figure 2**.

This information can be inputted into the EMSR model to produce optimal booking limits that evolve in time. A typical situation near the beginning of ticket sales was shown in **Figure 1**, while the evolution of the limits themselves is plotted in **Figure 3**.

The demand information in **Figure 2** can also be used to simulate requests for reservations. By taking the difference between the demand remaining at day t and at day $(t-1)$ before departure, the expected demand on day t can be determined. The actual number of requests generated on that day is then given by a Poisson random variable with parameter λ equal to the expected number of sales [Rothstein 1971]. The requests generated in this manner can be passed to a request-handling simulation that looks at the most current booking limits and then accepts or denies ticket requests based on the number of reservations already confirmed and the reservations limit. An example of this booking process is illustrated in **Figure 4**.

The results of the booking process provide an ideal testbed for the revenue opportunity model employed to evaluate overbooking performance. The simulation conducted for **Figure 4** had demand values of $\{11, 41, 57\}$, for classes 1,

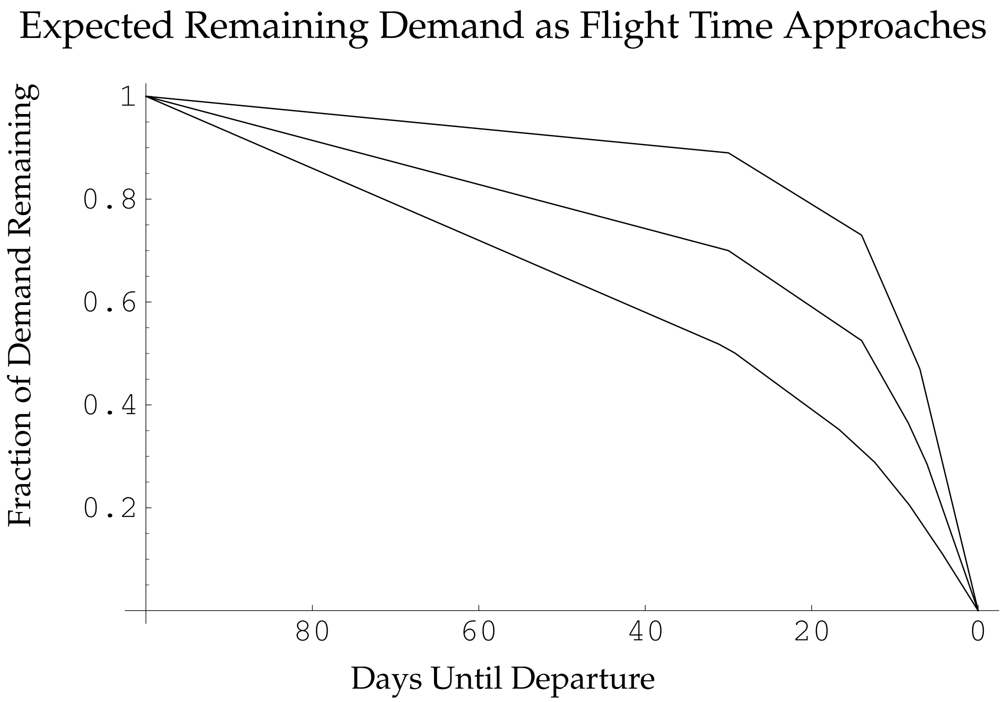

Figure 2. Demand remaining as a function of time for each of three fare classes, with the highest fare class on top. The curves represent the fraction of tickets that have yet to be purchased. Note that, for example, demand for high fare tickets kicks in much later than low-fare demand. (Data interpolated from Li [2001].)

2, and 3 respectively, before ticket sales were capped. A linear forward projection of these sales rates indicates that they would have reached $\{18, 49, 69\}$ had the classes remained open. Given fare classes $\{\$250, \$150, \$100\}$, the minimum revenue would be

$$R_- = \$100(69) + \$150(40) + \$250(0) = \$12,900$$

and the maximum revenue would be

$$R_+ = \$250(18) + \$150(49) + \$100(42) = \$16,050.$$

The actual revenue according to the EMSR formalism was

$$R = \$100(57) + \$150(41) + \$250(11) = \$14,600,$$

so the efficiency is

$$p = \frac{R - R_-}{R_+ - R_-} \cdot 100\% = \frac{\$14,600 - \$12,900}{\$16,050 - \$12,900} \cdot 100\% = 54\%,$$

without the use of a complicated overbooking scheme. This is not close to the efficiencies reported in Smith et al. [1992], which cluster around 92%. This relative inefficiency is to be expected, however, from a simplified booking scheme given incomplete booking request data.

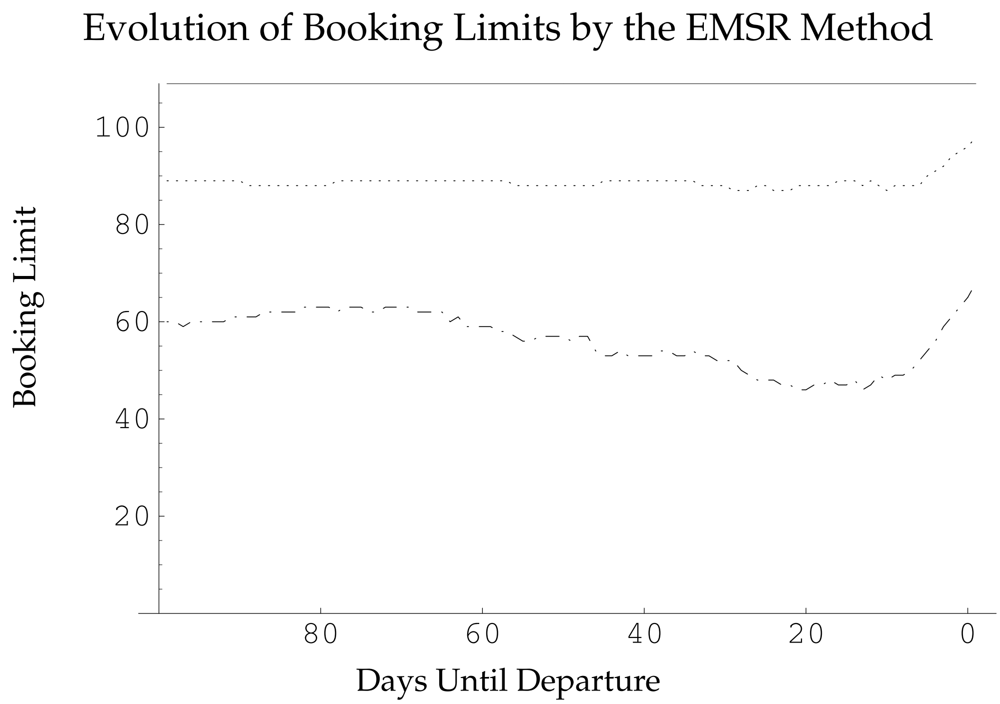

Figure 3. Booking limits for each class are dynamically adjusted to account for tickets already sold. For illustrative purposes, the number of tickets already sold is replaced here with the number of tickets that should have been sold according to expectations. In this case, the booking limits estimated at the beginning of the process are fairly accurate and require relatively little updating.

Figure 4. The EMSR-based booking limits are used to decide whether to accept or reject a sequence of ticket requests. These requests follow a Poisson distribution where the parameter λ varies with time to match the expected demand. Each fare class reaches its booking limit, as desired, so the flight is exactly full. Incorporating overbooking factors shifts the limits up accordingly.

Strengths and Weaknesses

Strengths

- **Applies widely accepted, industry-standard techniques.**

  Although more advanced (and optimal) algorithms are available and are used, EMSR and its descendants are still widely used in industry and can come close to optimality. The EMSR scheme, tested as-is on a real airline, caused revenue gains as much as 15% [Belobaba 1989].

  Our method for determining the optimality of a scheme after the fact is also based on tried and true methods developed by American Airlines [Smith et al. 1992].

- **Simplicity**

  Since it does not take into account as many factors as other booking models, EMSR is easier to deal with computationally. While a simple model may not be able to model a major airline with complete accuracy, an optimal pricing scheme can be made using only three fare classes [Li 2001].

Weaknesses

- **Neglects network effects**

  We treat the problem of optimizing each flight as if it were an independent problem although it is not.

- **Ignores sell-ups**

  In considering the discount seat allocation problem, we treat the demands for the fare classes as constants, independent of one other. This is not the case, because of the possibility of sell-ups. If the number of tickets sold in a lower fare class is restricted, then there is some probability that a customer requesting a ticket in that class will buy a ticket at a more expensive fare. This means it is possible to convert low-fare demand into high-fare demand, which would suggest protecting a higher number of seats for high fares than calculated by the model that we use. Sell-ups would be straightforward to incorporate into EMSR, but doing so would require additional information [Belobaba 1989].

- **Discounts possibility of recapture**

  Similar to sell-ups, the recapture probability is the probability that a passenger unable to buy a ticket at a certain price on a given flight will buy a ticket on a different flight. Depending on the recapture probability for each fare class, more or fewer seats might be allocated to discount fares.

Recommendations on Bumping Policy

In 1999, an average of only 0.88 passengers per 10,000 boardings were involuntarily bumped. Airlines are not required to keep records of the number of voluntary bumps, so it is impossible to determine a general bump rate.

Before bumping passengers involuntarily, the airline is required to ask for volunteers. Because there are no regulations on compensation for voluntary bumps, this is often a cheaper and more attractive method for airlines anyway. If too few people volunteer, the airline must pay those denied boarding 200% of the sum of the values of the passengers' remaining flight coupons, with a maximum of $400. This maximum is decreased to $200 if the airline arranges for a flight that will arrive less than 2 hours after the original flight. The airline may also substitute the offer of free or reduced fare transportation in the future, provided that the value of the offer is greater than the cash payment otherwise required. Alternatively, the airline may simply arrange alternative transportation if it is scheduled to arrive less than an hour after the original flight.

Auctions in which the airline offers progressively higher compensation for passengers who give up their seats are both the cheapest and the most common practice. As long as the airline does not engage in so much overbooking that it cannot find suitable reroutes for passengers bumped from their original itineraries, no alternatives to this policy need to be considered.

Conclusions

The two models presented in this paper work together to evaluate overbooking schemes by simulating their effects in advance and by quantifying their effects after implementation.

The expected marginal seat revenue (EMSR) model predicts overbooking scheme effectiveness. It determines the correct levels of protection for each fare class above the lowest—that is, how many seats should be reserved for possible sale at later dates and higher fares. Overbooking factors can be specified separately for each fare class, so the model effectively takes in overbooking factors and produces booking limits that can be used to handle ticket requests.

The revenue opportunity model attempts to estimate the maximum revenue from a flight under perfect overbooking and discount allocation. This is accomplished by estimating the actual demand for seats, then calculating the revenue that these seats would generate if sold to the highest-paying customers. Simple calculations produce the ideal overbooking cap and the optimal discount allocation for the flight. Thus, this model effectively represents how the airline would sell tickets if they had perfect advance knowledge of demand.

After the September terrorist attacks and their subsequent catastrophic effects on the airline industry, heightened airport security and fearful passengers will increase no-show and cancellation rates, seeming to dictate increasing overbooking levels to reclaim lost profits.

Airlines considering such action should be cautioned, however, that the negative effects of increased overbooking could outweigh the benefits. With reduced airline service, finding alternative transportation for displaced passengers could be more difficult. The effect of denying boarding to more passengers, along with the greater inconvenience of being bumped, could seriously shake consumers' already-diminished faith in the airline industry. With airlines already losing huge numbers of customers, it would be a mistake to risk permanently losing them to alternatives like rail and auto travel by alienating them with frequent overselling.

References

Beckman, J.M. 1958. Decision and team problems in airline reservations. *Econometrica* 26: 134–145.

Belobaba, Peter. 1987. Airline yield management: An overview of seat inventory control. *Transportation Science* 21: 63–73.

________. 1989. Application of a probabilistic decision model to airline seat inventory control. *Operations Research* 37: 183–197.

Biyalogorsky, Eyal, et al. 1999. Overselling with opportunistic cancellations. *Marketing Science* 18 (4): 605–610.

Boeing Commercial Airline Company. 1982. Boeing Promotional Fare Management System: Analysis and Research Findings. Seattle, WA: Boeing, Seattle.

Bowen, B.D., and D.E. Headley. 2000. Airline Quality Rating 2000. `http://www.uomaha.edu/~unoai/research/aqr00/` .

Brumelle, S.L., and J.I. McGill. 1993. Airline seat allocation with multiple nested fare classes. *Operations Research* 41: 127–137.

Chatwin, Richard E. 2000. Optional dynamic pricing of perishable products with stochastic demand and a finite set of prices. *European Journal of Operations Research* 125: 149–174.

Curry, R.E. 1990. Optimal airline seat allocation with fare classes nested by origins and destinations. *Transportation Science* 24: 193–204).

Davis, Paul. Airline ties profitability yield to management. 1994. *Siam News* 27 (5); `http://www.siam.org/siamnews/mtc/mtc694.htm` .

Li, M.Z.F. 2001. Pricing non-storable perishable goods by using a purchase restriction with an application to airline fare pricing. *European Journal of Operations Research* 134: 631–647.

Littlewood, K. 1972. Forecasting and control of passenger bookings. *AGIFORS Symposium Proceedings*: 12.

McGill, J.I., and G.Z. and Van Ryzin. 1999. Revenue management: Research overview and prospects. *Transportation Science* 33 (2): 233–256.

Rothstein, M. 1971. An airline overbooking model. *Transportation Science* 5: 180–192.

Rothstein, Marvin 1985. OR and the airline overbooking problem. *Operations Research* 33 (2): 237–247.

Ross, Robert. 1998. Embry-Riddle experts reveal airfare secrets. `http://comm.db.erau.edu/leader/fall98/priced.html` .

Shlifer, R., and Y. Vardi. 1975. An Airline Overbooking Policy. *Transportation Science* 9: 101-114.

Swan. 1990. Revenue management forecasting biases. Working paper. Seattle, WA: Boeing Commercial Aircraft.

Smith, Barry C., et al. 1992. Yield management at American Airlines. *Interfaces* 22 (1): 8–31.

Suzuki, Yoshinori. 2002. An empirical analysis of the optimal overbooking policies for U.S. major airlines. *Transportation Research* E38: 135–149.

Taylor, C.J. 1962. The determination of passenger booking levels. *AGIFORS Symposium Proceedings*, vol. 2. Fregene, Italy.

Weatherford, L.R., and S.E. Bodily. 1992. A taxonomy and research overview of perishable-asset revenue management: Yield management, overbooking, and pricing. *Operations Research* 40: 831–844.

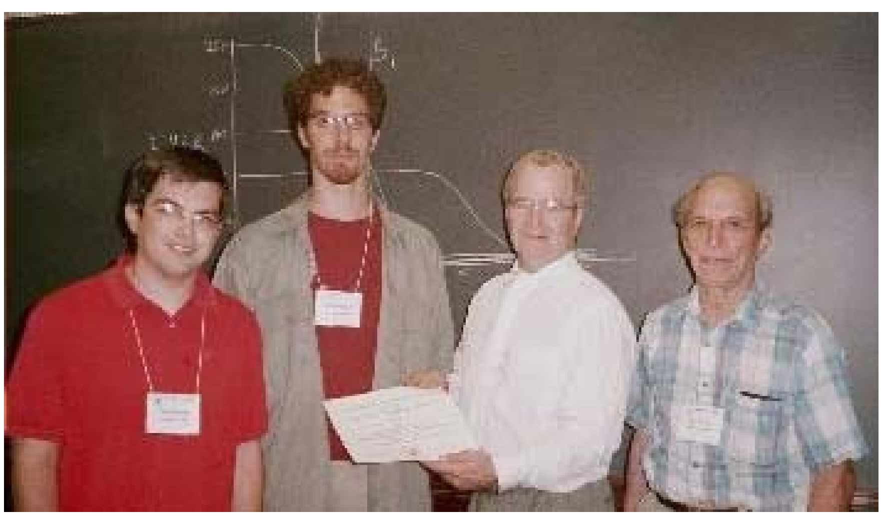

Presentation by Richard Neal (MAA Student Activities Committee Chair) of the MAA award to Daniel Boylan and Wesley Turner of the Harvey Mudd College team (Michael Schubmehl could not come), after their presentation at the MAA Mathfest in Burlington, VT, in August. On the right is Ben Fusaro, Founding Director of the MCM. Photo by Ruth Favro, Lawrence Tech University.

Letter to the CEO of a Major Airline

Airline overbooking is just one facet of a revenue management problem that has been studied extensively in operations research literature. Airlines have been practicing overbooking since the 1940's, but early models of overbooking considered only the most rudimentary cases. Most importantly, they did not take into account the revenue maximizing potential of price discrimination—charging different fares for identical seats. In order to maximize yield, it is particularly critical to price discriminate between business and leisure travelers. That is, when filling the plane, book as many full fare passengers and as few discount fare passengers as possible.

The implementation of a method of yield management can have dramatic effects on an airline's revenue. American Airlines managed its seat inventory to a $1.4 billion increase in revenue from 1989 to 1992—about 50% more than its net profit for the same period. Controlling the mix of fare products can translate into revenue increases of $200 million to $500 million for carriers with total revenues of $1 billion to $5 billion.

Though several decision models of airline booking have been developed over the years, comparing one scheme to another remains a difficult task. We have taken a two-pronged approach to this problem, both simulating and measuring a booking scheme's profitability.

In order to simulate a booking scheme's effect, we used the expected marginal seat revenue (EMSR) model proposed by Belobaba [1989] to generate near-optimal decisions on whether to accept or deny a ticket request in a given fare class. The EMSR model accepts as input overbooking levels for each of the fare classes that compose a flight, so different policies can be plugged in for testing.

Our approach to measuring a current scheme's profitability is similar to one used at American Airlines [Smith et al. 1992]. We compare the actual revenue generated by a flight with an ideal level calculated with the benefit of hindsight, as well as with a baseline level that would have been generated had no yield management been used. By calculating the percentage of this spread earned by a flight employing a particular scheme, we are able to gauge the effectiveness of different booking schemes.

It is our hope that these models will prove useful in evaluating your airline's overbooking policies. Simulations should provide insight into the properties of an effective scheme, and measurements after the fact will help to provide performance benchmarks. Finally, while it may be tempting to increase overbooking levels in order to compensate for lost revenues in the post-tragedy climate, our results indicate this will probably hurt long-term profits more than they will help.

Cordially,

Michael P. Schubmehl, Wesley M. Turner, and Daniel M. Boylan

Probabilistically Optimized Airline Overbooking Strategies, or "Anyone Willing to Take a Later Flight?!"

Kevin Z. Leder
Saverio E. Spagnolie
Stefan M. Wild
University of Colorado
Boulder, CO

Advisor: Anne M. Dougherty

Introduction

We develop a series of mathematical models to investigate relationships between overbooking strategies and revenue.

Our first models are static, in the sense that passenger behavior is predominantly time-independent; we use a binomial random variable to model consumer behavior. We construct an auction-style model for passenger compensation.

Our second set of models is more dynamic, employing Poisson processes for continuous time-dependence on ticket purchasing/cancelling information.

Finally, we consider the effects of the post-September 11 market on the industry. We consider a particular company and flight: Frontier Airlines Flight 502. Applying the models to revenue optimization leads to an optimal booking limit of 15% over flight capacity and potentially nets Frontier Airlines an additional $2.7 million/year on Flight 502, given sufficient ticket demand.

Frontier Airlines: Company Overview

Frontier Airlines, a discount airline and the second largest airline operating out of Denver International Airport (DIA), serves 25 cities in 18 states. Frontier offers two flights daily from DIA to LaGuardia Airport in New York. We focus on Flight 502.

The UMAP Journal 317–338.

Technical Considerations and Details

We discuss regulations for handling bumped passengers, airplane specifications, and financial interests.

Overbooking Regulations

When overbooking results in overflow, the Department of Transportation (DOT) requires airlines to ask for volunteers willing to be bumped in exchange for compensation. However, the DOT does not specify how much compensation the airlines must give to volunteers; in other words, negotiations and auctions may be held at the gate until the flight's departure. A passenger who is bumped involuntarily is entitled to the following compensation:

- If the airline arranges substitute transportation such that the passenger will reach his/her destination within one hour of the original flight's arrival time, there is no obligatory compensation.
- If the airline arranges substitute transportation such that the passenger will reach his/her destination between one and two hours after the original flight's arrival time, the airline must pay the passenger an amount equal to the one-way fare for flight to the final destination.
- If the substitute transportation is scheduled to arrive any later than two hours after the original flight's arrival time, or if the airline does not make any substitute travel arrangements, the airline must pay an amount equal to twice the cost of the fare to the final destination.

Aircraft Information

Frontier offers only one class of service to all passengers. Thus, we base our overbooking models on single-class aircraft.

Financial Considerations

Airline booking considerations are frequently based on the *break-even load-factor*, a percentage of airplane seat capacity that must be filled to acquire neither loss or profit on a particular flight. The break-even load-factor for Flight 502 in 2001 was 57.8%.

Assumptions

- **We need concern ourselves only with the sale of restricted tickets.** Frontier's are nonrefundable, save for the ability to transfer to another Frontier flight for $60 [Frontier 2001]. Restricted tickets represent all but a very small percentage of all tickets, and many ticket brokers, such as Priceline.com, sell only restricted tickets.

- **Ticketholders who don't show up at the gate spend $60 to transfer to another flight.**
- **Bumped passengers from morning Flight 502 are placed, at the *latest*, 4 h 35 min later on Frontier's afternoon Flight 513 to the same destination.** Frontier Airlines first attempts to place bumped passengers on other airlines' flights to the same destination. If it can't do so, Frontier bumps other passengers from the later Frontier flight to make room for the originally bumped passengers.
- **The annual effects/costs associated with bumping involuntary passengers is negligible in comparison to the annual effects/costs of bumping voluntary passengers.** According to statistics provided by the Department of Transportation, 4% of all airline passengers are bumped voluntarily, while only 1.06 passengers in 10,000 are bumped involuntarily. With a maximum delay for bumped passengers of 4 h 35 min, the average annual cost to Frontier of bumping involuntary passengers is on the order of $100,000—negligible compared to costs of bumping voluntary passengers.

The Static Model

Our first model for optimizing revenues is static, in the sense that passenger behavior is predominantly time-independent: All passengers (save no-shows) arrive at the departure gate independently. This model does not account for when passengers purchase their tickets. This system may be modeled by the following steps:

- Introduce a binomial random variable for the number of passengers who show up for the flight.
- Define a total profit function dependent upon this random variable.
- Apply this function to various consumer behavior patterns.
- Compute (for each behavioral pattern) an optimal number of passengers to overbook.

A Binomial Random Variable Approach

We let the binomial random variable X be the number of ticketholders who arrive at the gate after B tickets have been sold; thus, $X \sim \text{Binomial}(B, p)$. Numerous airlines consistently report that approximately 12% of all booked passengers do not show up to the gate (due to cancellations and no-shows) [Lufthansa 2000], so we take $p = .88$.

$$Pr\{i \text{ passengers arrive at the gate}\} = Pr\{X = i\} = \binom{B}{i} p^i (1-p)^{B-i}.$$

Modeling Revenue

We define the following per-flight total profit function and subsequently present a detailed explanation.

$$T_p(X) = (B - X)R + \begin{cases} \text{Airfare} \times X - \text{Cost}_{\text{Flight}}, & X \le C_{\bar{\$}}; \\ \text{Airfare} - \text{Cost}_{\text{Add}} \times (X - C_{\bar{\$}}), & C_{\bar{\$}} < X \le C; \\ \text{Airfare} - \text{Cost}_{\text{Add}} \times (X - C_{\bar{\$}}) - \text{Bump}(X - C), & X > C, \end{cases}$$

where

R = transfer fee for no-shows and cancellations,

B = total number of passengers booked,

Airfare = a constant

$\text{Cost}_{\text{Flight}}$ = total operating cost of flying the plane

Cost_{Add} = cost to place one passenger on the flight

Bump = the Bump function (to be defined)

$C_{\bar{\$}}$ = number of passengers required to break even on the flight

C = the full capacity of the plane (number of seats)

For Airfare, we use the average cost of restricted-ticket fare over a one-week period in 2002: \$316. $\text{Cost}_{\text{Flight}}$ is based on the break-even load-factor of 57.8%; for Flight 502, we take $\text{Cost}_{\text{Flight}} = \$24{,}648$ [Frontier Airlines 2001]. The average cost associated with placing one passenger on the plane is $\text{Cost}_{\text{Add}} \approx \$16$. The break-even occupancy is determined from the break-even load-factor; since Flight 502 uses an Airbus A319 with 134 seats, we take $C = 134$ and $C_{\bar{\$}} = 78$.

The Bump Function

We consider various overbooking strategies, the last three of which translate directly into various Bump functions.

- **No Overbooking**
- **Bump Threshold Model** We assign a "Bump Threshold" (BT) to each flight, a probability of having to bump one or more customers from a flight given B and p:

$$Pr\{X > \text{flight capacity}\} < \text{BT}.$$

We take BT = 5% of flight capacity. The probability that more than N ticket-holders arrive at the gate, given B tickets sold, is

$$Pr\{X > N\} = 1 - Pr\{X \leq N\} = 1 - \sum_{i=1}^{N} \binom{B}{i} p^i (1-p)^{B-i}.$$

This simplistic model is independent of revenue and produces (through simple iteration) an optimal number of ticket sales (B) for expecting bumping to occur on less than 5% of flights.

- **Linear Compensation Plan** This plan assumes that there is a fixed cost associated with bumping a passenger, the same for each passenger. The related Bump function is

  $$\text{Bump}(X - C) = B_{\$} \times (X - C),$$

  where $(X - C)$ is the number of bumped passengers and $B_{\$}$ is the cost of handling each.

- **Nonlinear Compensation Plan** Steeper penalties must be considered, since there is a chain reaction of expenses incurred when bumping passengers from one flight causes future bumps on later flights. Here we assume that the Bump function is exponential. Assuming that flight vouchers are still adequate compensation at an average cost of $2 * \text{Airfare} + \$100 = \$732$ when there are 20 bumped passengers, we apply the cost equation

  $$\text{Bump}_{\text{NL}}(X - C) = B_{\$}(X - C)e^{r(X-C)},$$

  where $B_{\$}$ is the ompensation constant and $r = r(B_{\$})$ is the exponential rate, chosen to fit the curve to the points $(0, 316)$ and $(20, 732)$.

- **Time-Dependent Compensation Plan (Auction)** The primary shortcoming of the nonlinear compensation plan is that it does not deal with flights with too few voluntarily bumped passengers, where the airline must increase its compensation offering. We now approximate the costs of an auction-type compensation plan.

  This plan assumes that the airline knows the number of no-shows and cancellations one-half hour prior to departure. The following auction system is employed. At 30 min before departure, the airline offers flight vouchers to volunteers willing to be bumped, equivalent in cost to the original airfare. This offer stands for 15 min, at which time the offer increases exponentially up to the equivalent of $948 by departure time. We chose this number as twice the original airfare (which is the maximum obligatory compensation for involuntary passengers if they are forced to wait more than 2 h), plus one more airfare costin the hope that treating the customers so favorably will result in future business from the same customers. These specifications are enough to determine the corresponding time-dependent Compensation function,plotted in **Figure 1**.

$$\text{Compensation}(t) = \begin{cases} 316, & 0 \le t \le 15 \text{ min;} \\ 105.33e^{0.07324\,t}, & 15 \text{ min} < t \le 30 \text{ min.} \end{cases}$$

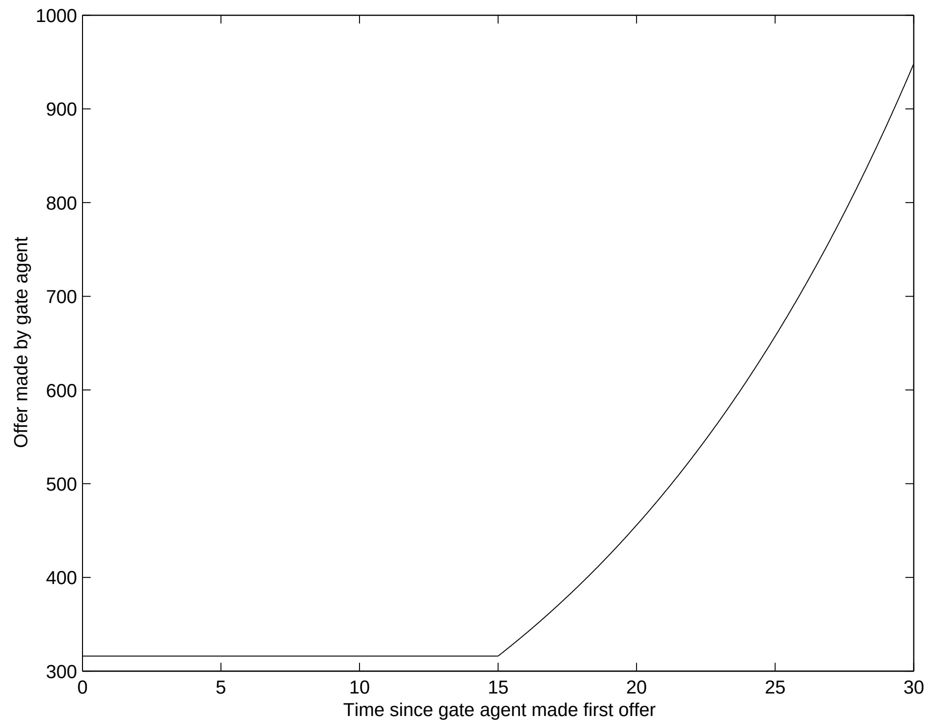

Figure 1. Auction offering (compensation)

Consideration of passenger behavior suggests that we use a Chebyshev weighting distribution for this effort (shown in **Figure 2**). A significant number of passengers will take flight vouchers as soon as they become available.

We simulate this random variable, which has probability density function

$$f(s) = \frac{1}{\pi\sqrt{1-s^2}}, \qquad s \in [-1, 1],$$

and cumulative distribution function

$$F(\tau) = \int_{-1}^{\tau} \frac{1}{\pi\sqrt{1-\eta^2}}\, d\eta = \frac{1}{2} + \sin^{-1}(\tau),$$

where η is a dummy variable. Inverting the cumulative distribution function results in a method for generating random variables with the Chebyshev distribution [Ross 1990]:

$$F^{-1}(\tau) = \sin\left[\pi(U - \tfrac{1}{2})\right],$$

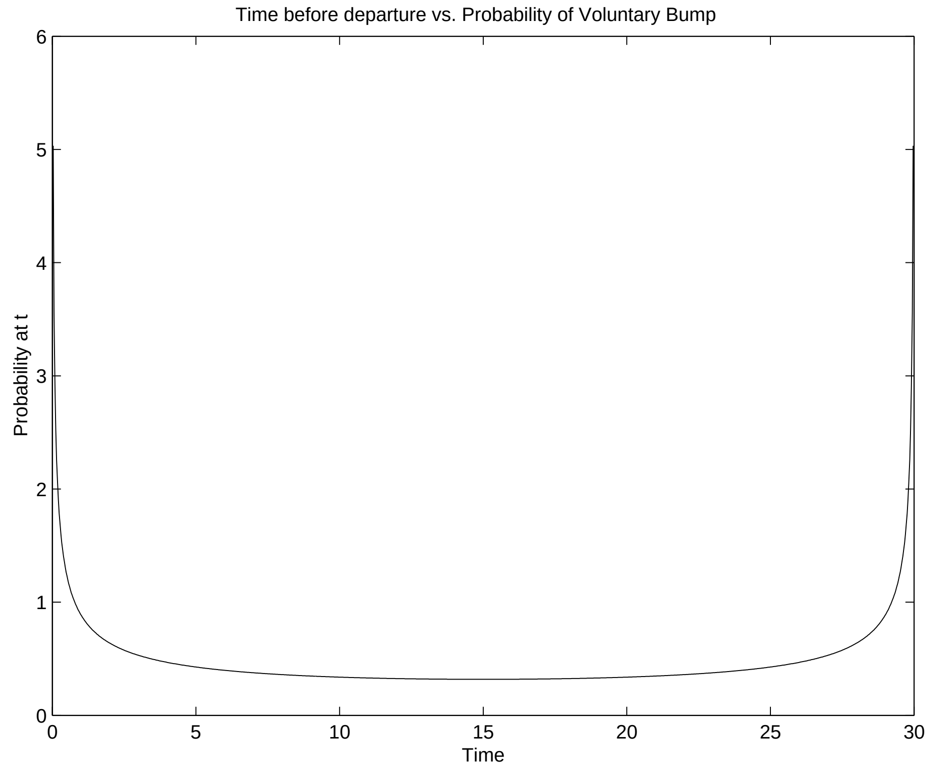

Figure 2. Chebyshev weighting function for offer acceptance

where U is a random uniform variable on $[0, 1]$.

With a linear transformation from the Chebyshev domain $[-1, 1]$ to the time interval $[0, 30]$ via $t = 15\tau + 15$, we find a random variable t that takes on values from 0 to 30 according to the density function $f(s)$. **Figure 3** shows the results of using this process to generate 100,000 time values. We use this random variable to assign times for compensation offer acceptance under the auction plan.

The total cost of bumping $(X-C)$ passengers is $\sum_{i=1}^{X-C} \text{Compensation}(t_i)$.

Optimizing Overbooking Strategies

Our goal is to maximize the expected value of the total profit function, $E[T_P(X)]$, given the variability of the bump function and the probabilistic passenger arrival model.

There are competing dynamic effects at work in the total profit function. Ticket sales are desirable, but there is a point at which the cost of bumping becomes too great. Also, the variability of the number of passengers who show up affects the dynamics. The expected value of the total profit function is

$$E[T_P(X)] = \sum_{i=1}^{B} T_P(i) \binom{B}{i} p^i (1-p)^{B-i}.$$

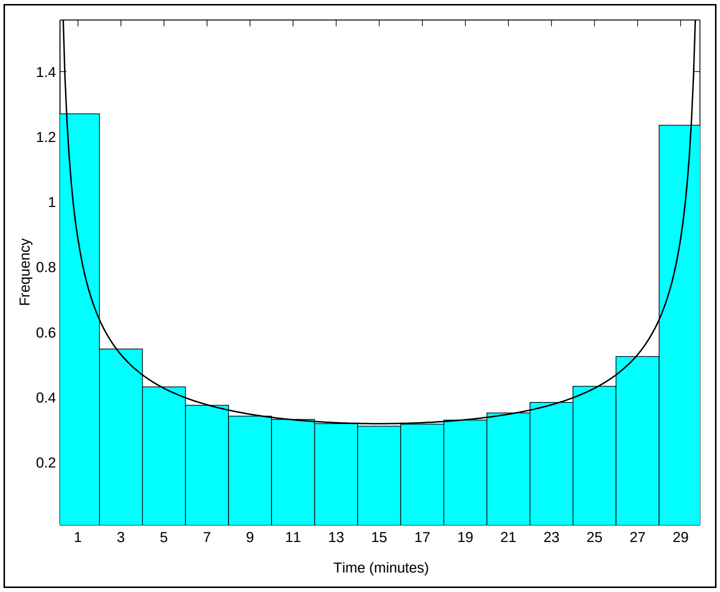

Figure 3. Histogram of 100,000 draws from the Chebyshev distribution.

We optimize the revenue by finding the most appropriate booking limit (B) for any bump function. Solving such a problem analytically is unrealistic; any solution would require the inversion of a sum of factorial functions. Therefore, we turn to computation for our results. We wrote and tested MatLAB programs that solve for B over a range of trivial bump Functions.

Results of Static Model Analysis

No Overbooking

If Frontier Airlines does not overbook its flights, it suffers a significant cost in terms of loss of opportunity. If the number of people that booked (B) equals plane capacity (C), the expected value of X (number of passengers who arrive at the gate) is $pB = pC = .88 \times 134 \approx 118$ passengers. Assuming (as in the total profit function) that each passenger beyond the 78th is worth \$300 in profit, the expected profit is nearly

$$(134 - 118) \times \$60 + \$300 \times (118 - 78) = \$12{,}960$$

per flight.This is only an estimate, since a smaller or larger proportion than 57.8% of ticket-holding passengers may arrive at the gate. The profit is sizeable but there are still (on average) 16 empty seats! The approximate lost opportunity cost is $\$300 \times 16 = \$4{,}800$! Thus, not overbooking sends Flight 502 on its way with only 63% of its potential profitability.

Bump Threshold Model

Using a 0.05 bump threshold, we compute an optimal number of passengers to book on Flight 502. Given the Airbus A319 capacity of 134 passengers and a passenger arrival probability of $p = .88$, the optimal number of tickets to sell to guarantee that bumping occurs less than 5% of the time is $B = 145$, or 107% of flight capacity.

Linear Compensation Plan

Table 1 shows the expected profit for various linear bump functions.

Table 1.
Linear bump functions compared.

Bump cost per passenger	Optimal # to book	Expected profit per flight
200	∞	∞
316	162	$17,817
400	156	$17,394
500	153	$17,121
600	152	$16,940
700	151	$16,799
800	151	$16,692
900	150	$16,601
1000	150	$16,526

If Frontier were to compensate bumped passengers less than the cost of airfare, bumping passengers would always cost less than revenue gained from ticket sales. Thus, assuming it could sell as many tickets as it wanted, Frontier would realize an unbounded profit on each flight! Obviously, the linear compensation plan is not realistic in this regime, and we must wait for subsequent models to see increased real-world applicability. These results agree with the result of using a simple bump threshold above and indicate an average profit of approximately $17,000. In comparison with using no overbooking strategy at all, Frontier gains additional profit of $4,000 per flight!

The actual dynamics of the problem may be seen in **Figure 4**, where competing effects form an optimal number of tickets to sell (B) when Frontier assumes a sizeable enough compensation average. We can also see the unbounded profit available in the unrealistic regime.

Nonlinear Compensation Plan

Numerical results for the more realistic nonlinear model paint a more reasonable picture.

Table 2 recommends booking limmits similar to (though slightly higher than) previous models. The dynamics may be seen in the **Figure 5**.

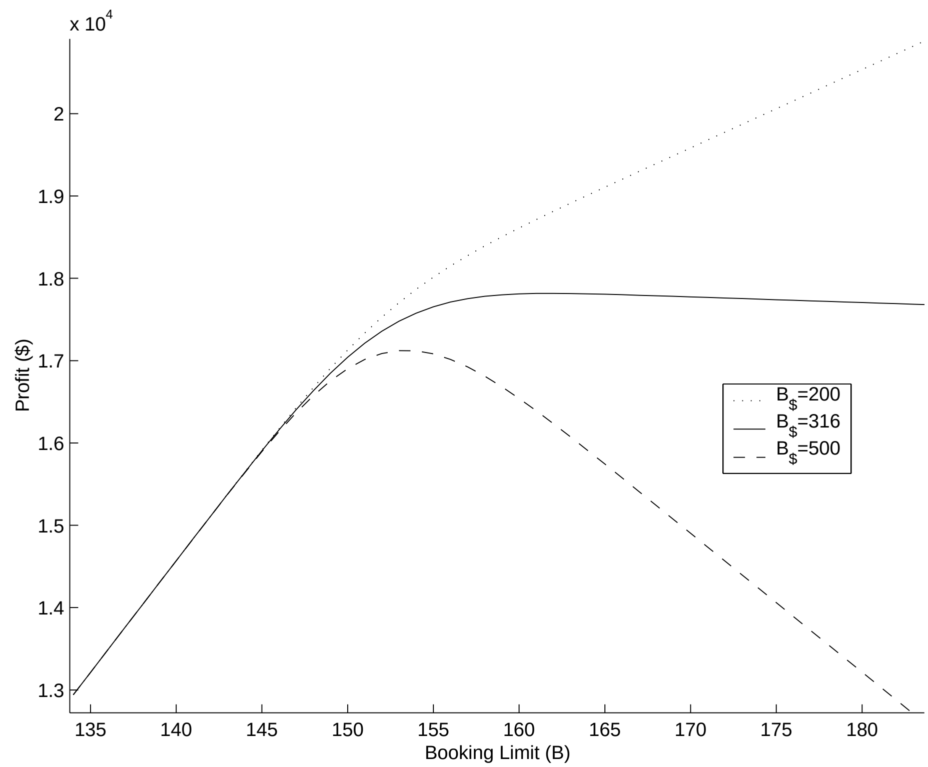

Figure 4. Per-flight profit vs. booking limit (B) for different bump costs (Linear Compensation Plan)

Table 2.
Nonlinear bump functions compared.

Bump function	Optimal number to book	Profit per flight
$50e^{0.134(X-C)}(X-C)$	160	\$18,700
$100e^{0.100(X-C)}(X-C)$	158	\$18,240
$200e^{0.065(X-C)}(X-C)$	156	\$17,722
$316e^{0.042(X-C)}(X-C)$	154	\$17,363

All nonlinear bump functions that we investigated result in a maximum realizable profit, as expected.

Time-Dependent Compensation Plan

The histogram of 1,000 runs using the time-dependent compensation plan in **Figure 6** shows that the optimal booking limit is most frequently $B = 154$. **Figure 7** is a graph of expected total profit versus the optimal booking limit for 15 trials, displaying the randomness due to the Chebyshev draws at higher values of B. If B is too low, then all models have the same profit behavior, because the randomness from the overbooking scheme is not introduced until customers are bumped. This graph also shows that regardless of random

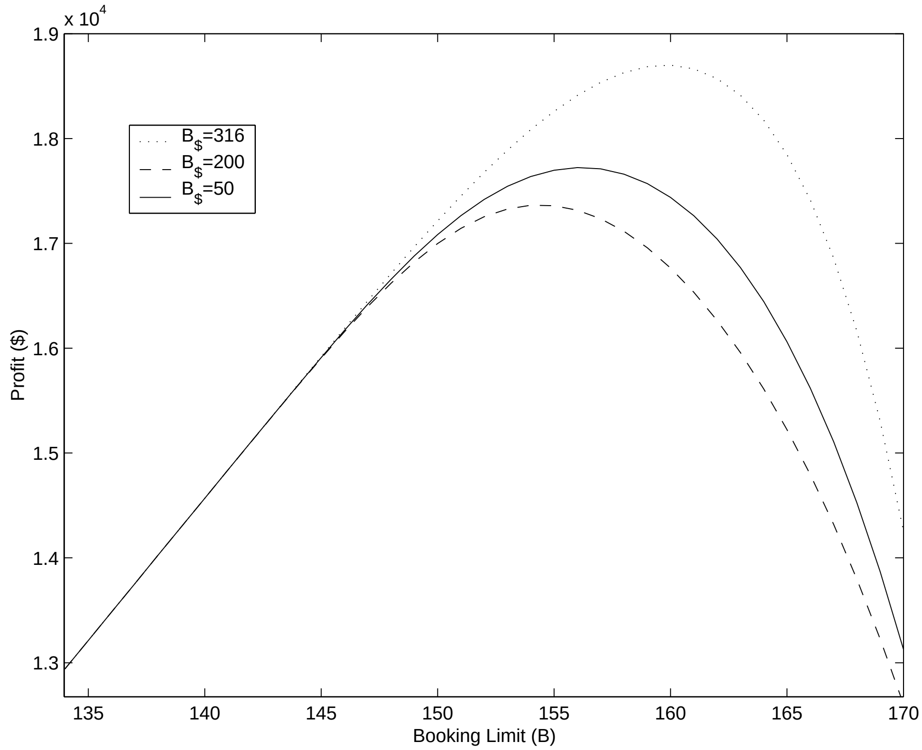

Figure 5. Per-flight profit vs. booking limit (B) for different bump functions (Nonlinear Compensation Plan).

effects, profitability is maximized around $B = 160$.

The Dynamic Model

Many of the assumptions in the binomial-based models are loosened in this dynamic setting. Continuous time allows for more detailed analysis of the order of events in the airline booking problem. Keeping track of the order of reservation requests, ticket bookings, and cancellations results in a model that attempts to recommend what ticketing agents should do at a certain time. In the "Firesale Model," we attempt to increase revenue by selling the tickets of cancellations to customers who would otherwise be denied tickets due to a fixed booking limit.

Reservation Process

We simulate the booking/reservations process, which often begins weeks before departure and continues right up until departure (due, for example, to other airlines booking their bumped customers into Frontiers' empty seats).

To model the stream of reservation requests, we employ a Poisson process $\{N(t), t \geq 0\}$—a counting process that begins at zero ($N(0) = 0$) and has

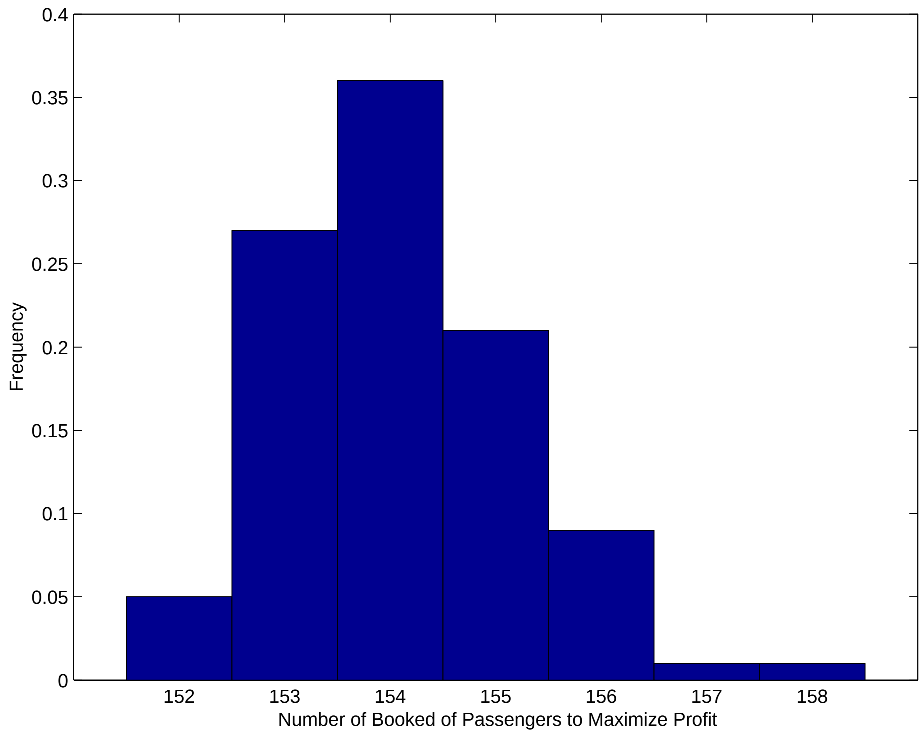

Figure 6. Time-dependent compensation plan simulated 1,000 times

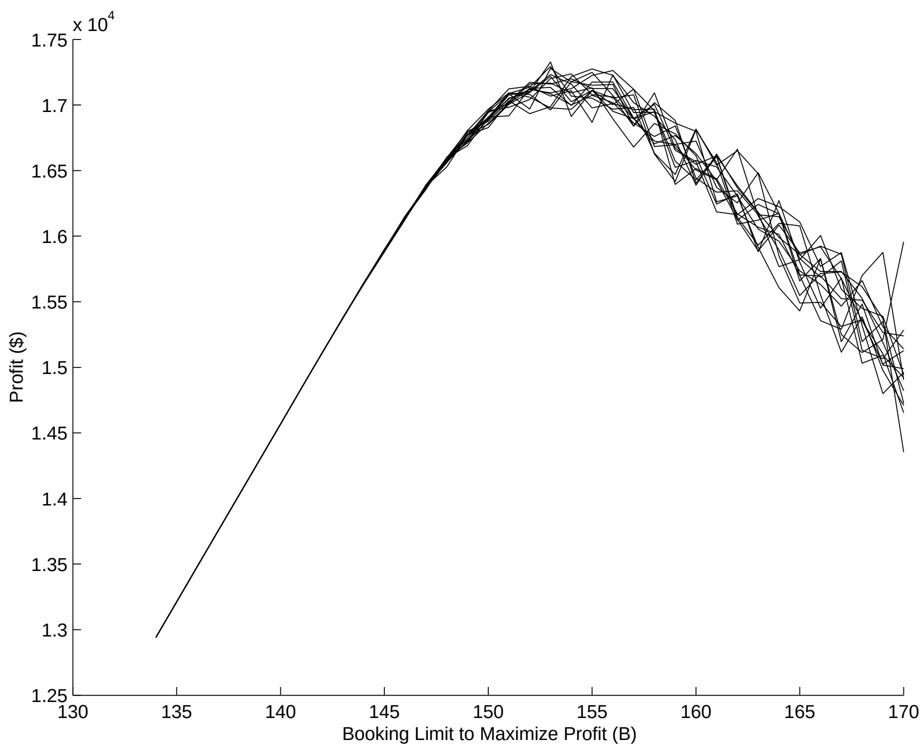

Figure 7. 15 time-dependent compensation plan simulations

independent increments, with the number of events in any interval of length t Poisson-distributed with mean λt [Ross 2000]. The interarrival times of a Poisson process are distributed according to an exponential distribution with rate parameter λ. Each reservation request comes with a variable number of tickets requested for that reservation. The number of tickets requested is generated from some specified batch distribution, BatchD, that we introduce later.

This arrangement results in a compound Poisson process (in this case, a"stuttering" process [McGill and Garrett 1999]), which provides a more reasonable fit to real-world reservation request data than simpler processes.

Simulating the first T time units of a Poisson process using the method in Ross [1990] results in a vector **at** of arrival times for the $A = \text{length}(\textbf{at})$ reservation requests received.

Another vector, **Bnum**, the number of tickets requested in each of the A reservations, is also generated according to the batch distribution. The density BatchD is shown in **Figure 9**; it states that callers reserve anywhere from 1 to 4 tickets at a time, with varying probabilities for each number. The total number of tickets (potential fares) requested is then $\sum_i(\textbf{Bnum}(i))$.

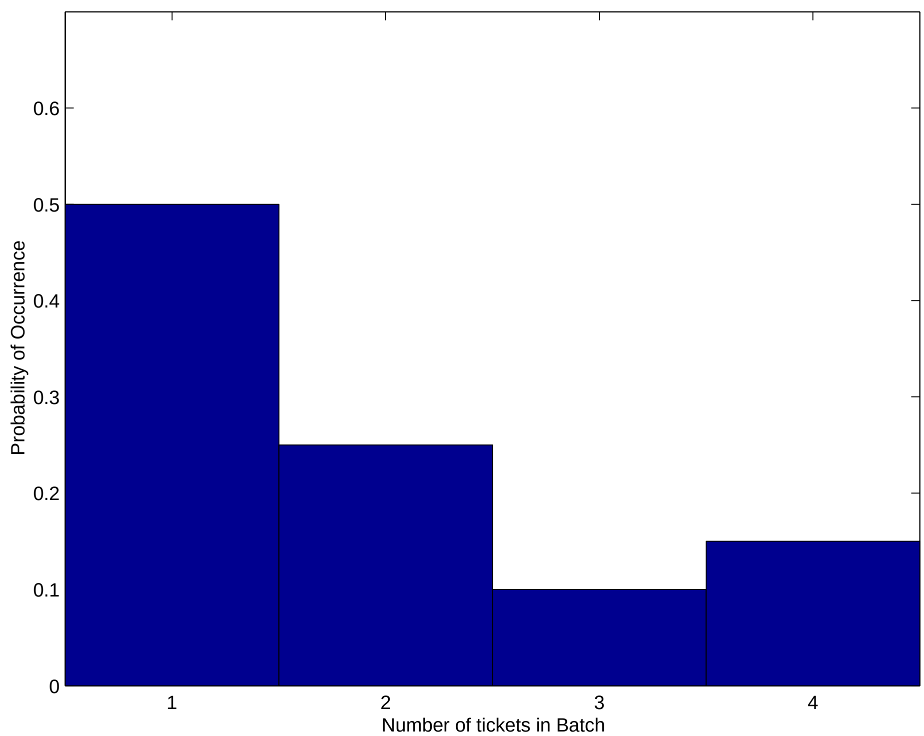

Figure 9. Density function for number of tickets in a batch of reservations.

The arrival rates for these reservation requests are derived by setting the expected value of the Poisson process over an interval of length T equal to the average ticket demand A_D that we expect. Then a rate of $\lambda = A_D/E_B T$, where E_B is the expected value of BatchD (1.9 in this case), will on average generate A_D tickets. The histogram in **Figure 10** shows the results of a simulation of 10,000

Poisson processes outputting the number of reservations requested when the average demand for tickets was 134 ($A_D = C$).

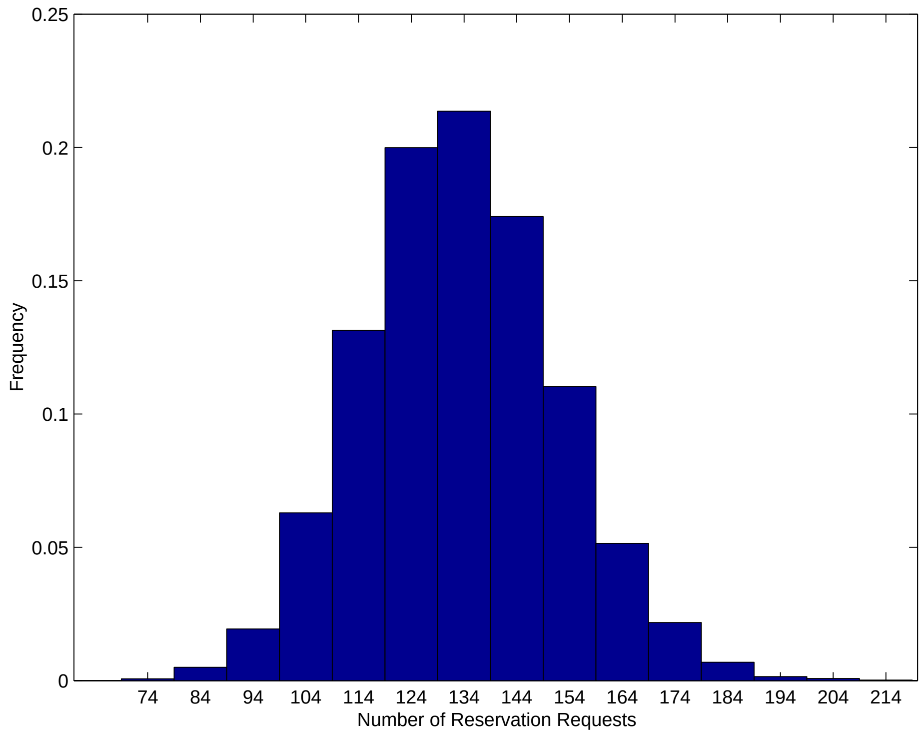

Figure 10. Histogram of number of reservation requests for 10,000 flights with an average demand of 134 tickets.

Cancellations and No-Shows

The binomial-based static models do not distinguish between cancellations (tickets voided before the flight departs) and no-shows (tickets not used or voided by flight departure); however, the dynamic model is well-suited for monitoring these events. We assume that 75% of unused tickets are cancellations and 25% are no-shows. Additionally, we assume that the time of cancellation for a set of tickets reserved together is uniformly distributed from the time that the tickets are granted to the time that the flight departs. This means that some cancellations occur almost immediately after the ticket(s) are granted (e.g., due to a typo on an online ticket service form), while some occur just before a plane is scheduled to depart (e.g., a last-minute change of plans). Lastly, we assume that multiple tickets in a single reservation behave equivalently (i.e., families act as unbreakable groups!).

To simulate this process, for each requested reservation a biased coin is flipped to determine with probability p if the group will keep their tickets. If not, another biased coin is flipped to determine whether the unused tickets are cancellations or no-shows. If a cancellation occurs, a cancellation time is drawn uniformly between that batch's arrival time and the flight departure time.

Dynamic Booking

Dynamic Test Model

We use the dynamic model to make the binomial-based models more realistic by eliminating some assumptions and introducing randomness. The Dynamic Test allows for "group tickets" (for both reservations and cancellations). The Dynamic Test requires that average ticket demand A_D be specified, so as to confirm the expected effects of less demand for tickets.

Firesale Model

The Firesale Model uses cancellation times to sell all possible tickets. If the number of tickets requested (at time t) for a particular reservation plus Tix (the number of tickets approved and still held at time t) is less than the predetermined booking limit (B), then a reservation request is approved. Conversely, if $Tix(t)$ is equal to the booking limit or if the sale of the multiple tickets requested in a reservation batch would push $Tix(t)$ over the booking limit, the request is rejected. Thus, for a process with no cancellations, reservation requests totaling less than the booking limit would be approved while subsequent requests would be rejected. The Fireside Model is highly dependent on the average demand (i.e., if demand is high enough, the airline would end up with an overwhelming majority of no-shows, as opposed to cancellations). The Firesale Model is the most realistic model developed in this paper.

Results of Dynamic Model Analysis

The Fireside Model attempts to capture a scenario where all tickets of cancellations are sold as long as there are customers willing to buy them. If demand for tickets is high enough, we expect to sell all tickets of cancellations, resulting in a large number of bumped passengers. However, because the airline profits $60 from each cancellation or no-show and because the numbers of both cancellations and no-shows continue to increase as more tickets are sold, reasonable results are expected for reasonable ticket demand.

Figure 11 plots expected profit as a function of bumping limit as determined from 1,000 Fireside Model simulations. An average demand twice that of capacity ($A_D = 268$) is used and a maximum profit is realized at a booking limit of 163. Most importantly, this figure displays how a small variation in booking limit could significantly alter profit. A change in either direction of 3 in corresponds to a loss of more than $1,000 profit.

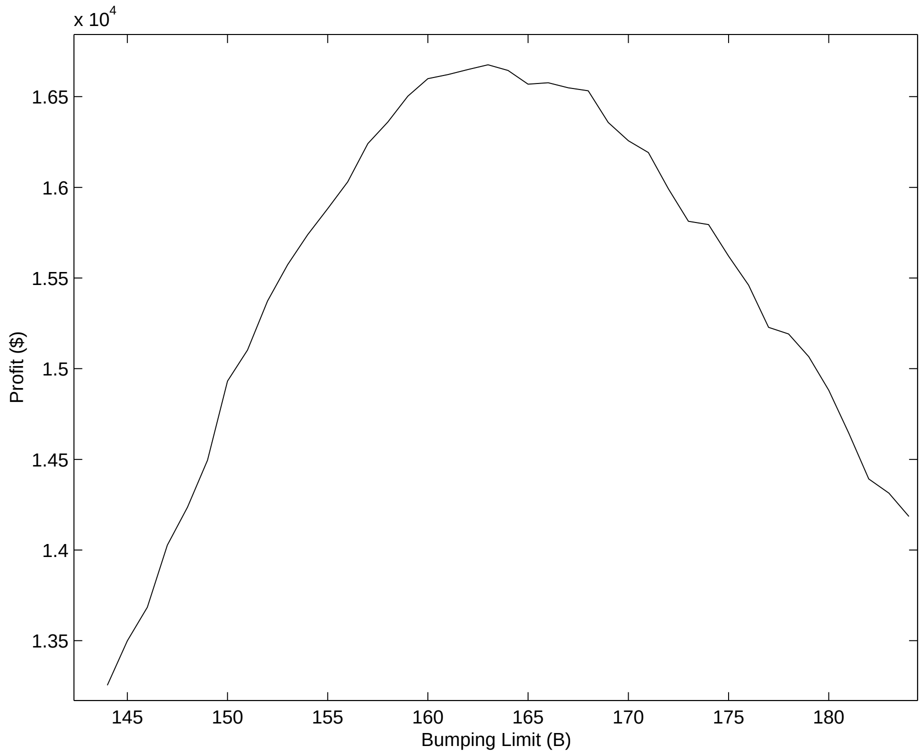

Figure 11. 1,000 simulations of the Fireside Model.

Dynamic Testing of the Static Model

The dynamic model allows us to test the results from the static (binomial-based) models in a more realistic setting. The Dynamic Test allows tickets to be reserved in batches and introduces the randomness experienced in real-world airline booking.

In all testing, 10,000 simulations are performed for each booking limit (B) and then expected profits are computed. Booking limit vs. Profit ($) is plotted for appropriate booking limit values. The average demand (A_D) used in this test is kept constant at twice the capacity of the airplane (so $A_D = 268$), to simulate a very large pool of customers so that the overbooking process could be tested.

Linear Compensation Plan

We tested two Bump costs ($B_\$ = \$316$ and $B_\$ = \$600$) with different behaviors (as predicted by the static model).

Figure 12 shows that for this compensation plan, an optimal booking limit is $B = 155$, an increase of 3 from the optimal value for the static model. However, profit drops off steeply for booking limits over 155, indicating that a more conservative strategy might be to lower the bumping limit to ensure that this steep decline is rarely reached.

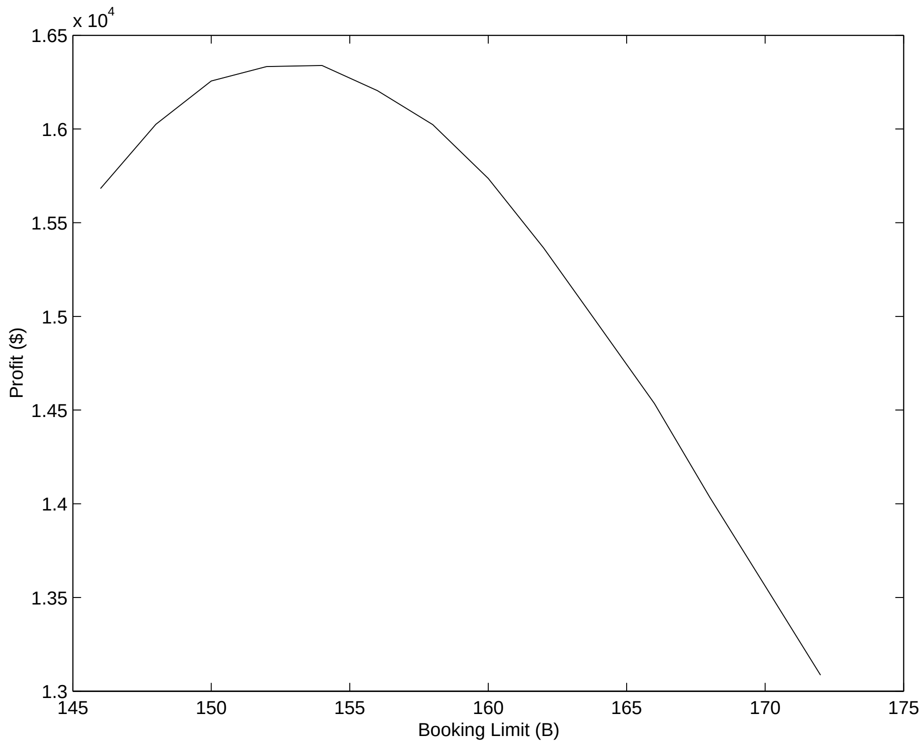

Figure 12. 10,000 simulations of the linear compensation plan with $B_{\$} = \$600$.

Figure 13 corresponds to a bump cost of 316; the optimal booking limit is now 166, again an increase (from 162).

Nonlinear Compensation Plan

We tested two nonlinear bump coefficients ($B_{\$} = 316$ and $B_{\$} = 100$) with different behaviors (as predicted by the static model).

Figures 14 and **15** demonstrate the negative effect of too high a booking limit. For nonlinear bump coefficients $B_{\$} = 316$ and $B_{\$} = 100$, optimal booking limits from the static model are 154 and 160, with Dynamic Test result values of 154 and 158.

Time-Dependent Compensation Plan

Figure 16 shows that the optimal booking limit for the time-dependent compensation plan is $B = 155$, an increase of 1 from the static model. Profit appears to rise relatively steeply until the optimal booking limit is reached and then falls steeply. Thus, in our most realistic static model, a careful overbooking plan matters the most! If the booking limit were altered by 3, the profit would shrink by more than $1,000, similar to the result detailed in the Fireside Model.

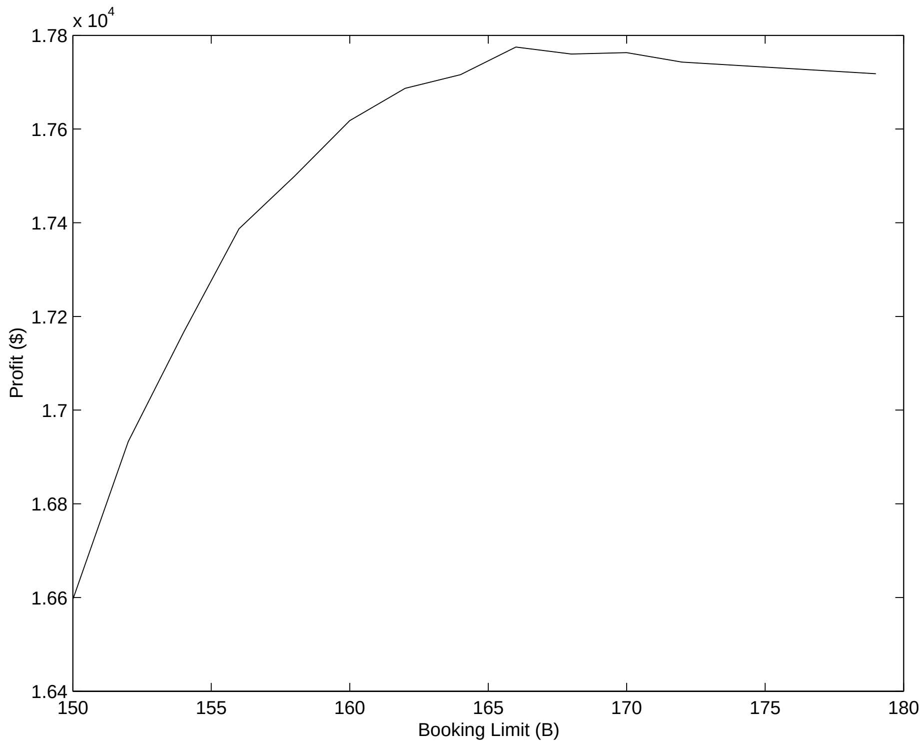

Figure 13. 10,000 simulations of the linear compensation plan with $B_\$ = \$316$.

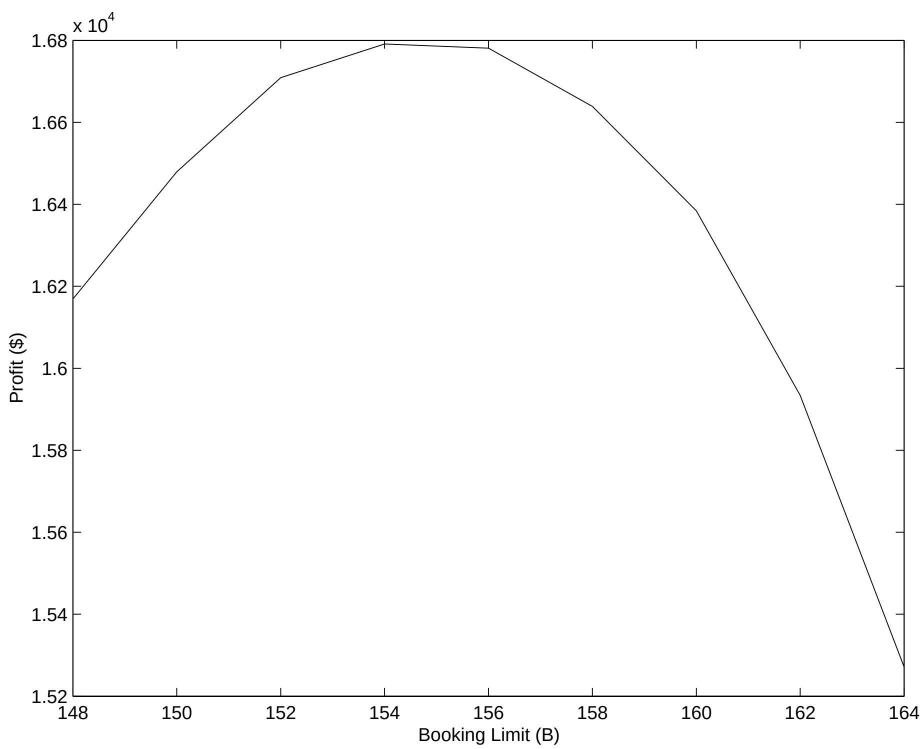

Figure 14. 10,000 simulations of the nonlinear compensation plan with $B_\$$=316.

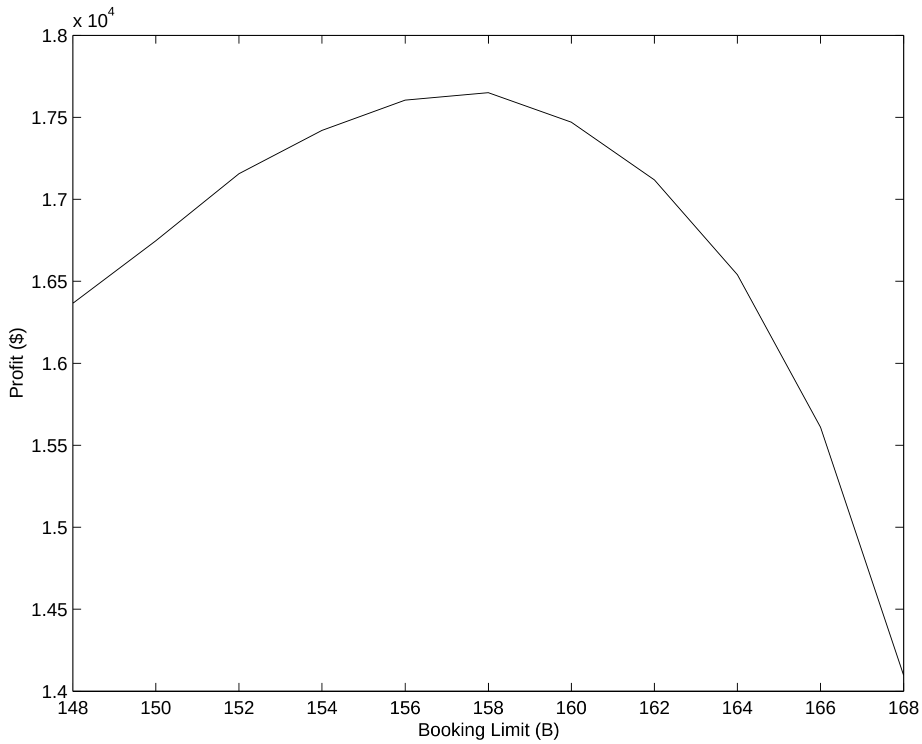

Figure 15. 10,000 simulations of the nonlinear compensation plan with $B_{\$}$=100.

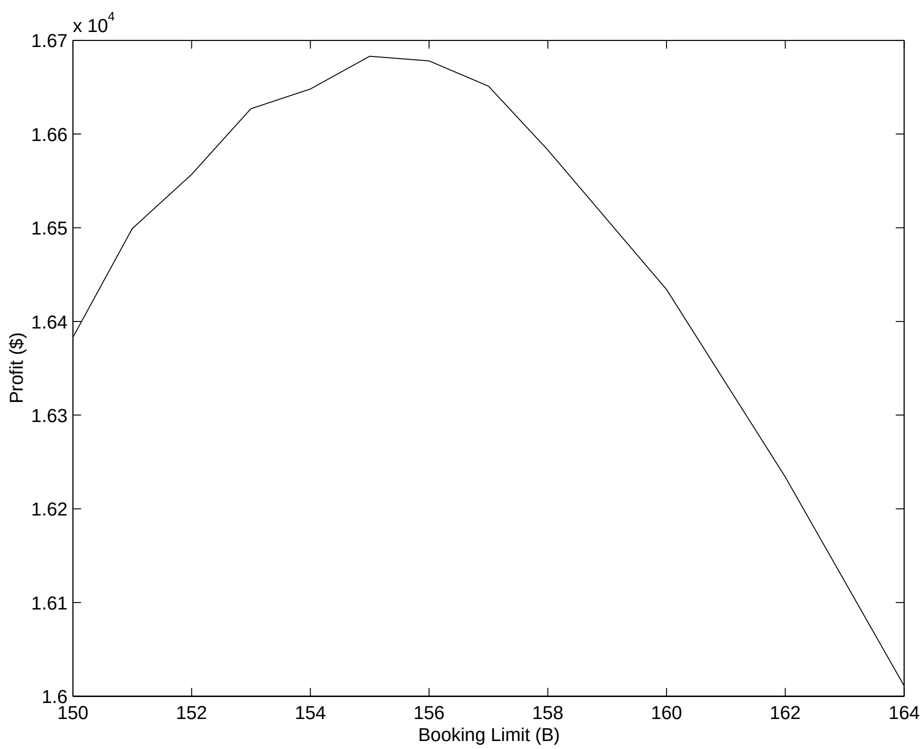

Figure 16. 10,000 simulations of the time-dependent compensation plan.

Post-September 11 Effects

Security checks (at Denver International Airport) add only 10 min to check-in ["Frontier operating at 80%" 2001], which may be considered negligible.

The most significant post-September 11 effect that the airlines must consider is the consumer fear. The individual probability of passenger arrival p should not change drastically, since ticket-purchasing customers after September 11 are fully aware of the risks involved. A consequence of September 11 that is difficult to model is the decrease in average demand for flight reservations.

Model Strengths and Weaknesses

Strengths

- Time-dependent auction model for pre-flight compensation: When Frontier begins to offer compensation to voluntarily bumped passengers one-half hour before departure, our model allows consumer behavior to influence the financial results.

- Time-dependent decision process in the dynamic model: The dynamic model allows ticketing agents to decide whether or not to accept reservation requests based on the number of tickets sold by then and based on time until departure.

- Multiple considerations of consumer behavior via bump functions: The implementation of multiple bump functions allow for testing alternative strategies for compensation. Profit and customer satisfaction may then be balanced depending upon the company's short-term or long-term interests.

- Varying degrees of model complexity: Our early models are simple, making sizeable simplifying assumptions to exhibit the most basic dynamics inherent in the problem. We take small steps of increasing complexity towards a more realistic model. The intuitive relationships between the results from each step lead to increased confidence in the stability and applicability of the most involved models.

Weaknesses

- Absence of a stability analysis: We lack an adequate mathematical understanding of the stability of our models. Varying parameters like p could potentially alter our results.

- Infinite customer pool in the static model: In our static model, we assume that for any booking limit we set, all tickets will be sold.

- Insufficient data: The only operational data that we could get from Frontier Airlines was its quarterly report, which contains general information on how many people flew, operating costs, revenues, number of flights flown, and occupancy rates. However, our model lacks information regarding cancellation rates, no-show rates, cost per flight, rates of reservation requests, and ratio of restricted tickets sold to unrestricted tickets sold. The lack of this information limits us because our parameters are not based on historical data, and therefore we cannot be confident in the accuracy of our rates.

Conclusion and Recommendations

Our models are quite consistent in recommending similar booking limits: 154 passengers on 134-seat Flight 502, 115% of capacity. This limit results in an average of $17,000 per flight; so this one flight alone, by employing one of our overbooking strategies, nets the company an extra $2.7 million profit per year, under the limiting assumption of an infinite demand pool.

References

U.S. Department of Transportation. 2001. Air Travel Consumer Report. `http//:www.dot.gov/airconsumer/` .

AlaskaAir. 2001. `http://www2.alaskaair.com` .

BBC News. 2001. Airlines receive $15bn aid boost. `http://news.bbc.co.uk/hi/english/business/newsid_1558000/1558854.stm` .

Boudreaux, Donald. 1998. Notes from FEE: Julian Simon Life Saver. `http://www.fee.org/freeman/98/9804/notes.html` .

Chatwin, Richard. 1999. Continuous-time airline overbooking with time-dependent fares and refunds. *Transportation Science* 33 (2) (May 1999).

Davis, Paul. 1994. Airline ties profitability yield to management. *SIAM News* 27 (5) (May/June 1994).

Delta Airlines. 2002. `http://www.delta.com/care/service_plan/index.jsp` .

Frontier Airlines reports fiscal third-quarter profit. 2002. *Business Wire* (2 February 2002).

Frontier Airlines Inc. 2001. Form 10-Q, Quarterly Report Pursuant to Section 13 or 15(d) of the Securities Exchange Act of 1934. Securities and Exchange Commission file number 0-24126, quarterly period ending September 30, 2001.

________. 2002. `http://www.frontierairlines.com` .

Frontier operating at 80%. 2001. *Denver Business Journal* (20 September 2001). http://denver.bizjournals.com/denver/stories/2001/09/17/daily38.html .

Graves, Gary. 2001. Airline decline. CBC News Online (7 November 2001). http://www.cbc.ca/news/indepth/background/airlinedecline.html .

Gupta, Diwikar, and William Cooper. 2001. When is yield improvement really an improvement? Working Paper, University of Minnesota, Department of Mechanical Engineering. September 4, 2001.

Inse, Amy. 2001. Airport's a lonely place these days. *Rocky Mountain News* (19 September 2001). http://www.insidedenver.com/drmn/america_under_attack/article/0,1299,DRMN_664_828402,00.html .

Karaesmen, Itir, and Garrett van Ryzin. 1998. Overbooking with substitutable inventory classes. (20 September 1998). Columbia University. http://www.columbia.edu/~gjv1/subob1.pdf .

Kesmodel, David. 2002. Frontier expects profits for quarter. *Rocky Mountain News* (7 February 2002).

Lufthansa. 2001. Overbooking. http://cms.lufthansa.com/de/dlh/en/nws/0,1774,0-0-77892,00.html .

Masselli, Jennifer. 2001. IT's critical role in airlines. *Information Week* (12 November 2001). http://www.informationweek.com/story/IWK20011112S0001 .

McGill, Jeffery, and Garrett van Ryzin. 1999. Revenue management: Research overview and prospects. *Transportation Science* 33 (2) (May 1999).

Ross, Sheldon M. 1990. *Simulation*. 2nd ed. San Diego, CA: Harcourt Academic Press.

________ . 2000. *Probability Models*. 7th ed. San Diego, CA: Harcourt Academic Press.

Subramanian, J., S. Stidham, and C.J. Lautenbacher. 1999. Airline yield management with overbooking, cancellations, and no-shows. *Transportation Science* 33 (2) (May 1999).

ACE is High

Anthony C. Pecorella
Crystal T. Taylor
Elizabeth A. Perez
Wake Forest University
Winston-Salem, NC

Advisor: Edward E. Allen

Introduction

We design a model that allows an airline to substitute its own values for ticket prices, no-show rates and fees, compensation for bumped passengers, and capacities to determine its optimal overbooking level.

Our model is based on an equation that combines the two cases involved in overbooking: The first sums all cases in which the airline doesn't fill all seats with passengers, and the second sums all cases in which there is an overflow of passengers due to overbooking. The model includes the possibility of upgrading passengers from coach to first-class when there is overflow in coach.

Furthermore, we use a binomial distribution of the probabilities of bumping passengers, given different overbooking percentages, to supply the airlines with useful information pertaining to customer relations.

We apply our model with different values of the parameters to determine optimal overbooking levels in different situations.

By using our model, an individual airline can find an optimal overbooking level that maximizes its revenue. A joint optimal overbooking strategy for all airlines is to agree to allow bumped passengers to fly at a discounted fare on a different airline.

Analysis of the Problem

From January to September 2001, 0.19% of passengers were bumped from flights due to overbooking. This seems like an inconsequential percentage, but it actually amounts to 730,000 people. Additionally, 4.4% of those bumped,

The UMAP Journal 339–349.

or 32,000 people, were denied their flights involuntarily [U.S. Department of Transportation 2002].

Since 10% to 15% of passengers who reserve a seat don't show up, airlines have little chance to fill their planes if they book only as many passengers as seats available. Overbooking by American Airlines helped save the airline $1.4 billion between 1989 and 1992.

We examine a fictional company to determine an optimal overbooking strategy that maximizes revenue. The goal is a model to increase revenue while maintaining favorable customer relations.

Our main model, the Expected Gain Model, provides a clear formula for what percentage of the seats to overbook. Based on sample no-show rates, ticket prices, and seat numbers, our Expected Gain Model shows that a 16% overbooking rate is the most effective choice.

Our other model, the Binomial Distribution Model, calculates, for various overbooking levels, the probability that a passenger will be bumped.

Assumptions

- There is no overbooking in first class (to maintain good relations with wealthy and influential passengers).
- Anyone bumped (voluntarily or involuntarily) is compensated with refund of ticket price plus an additional 100% of the ticket price.
- There are only two flight classes, coach and first-class.
- The fare is constant regardless of how far in advance the ticket is purchased. Overbooked passengers are given seats on a first-come-first-served basis, as is often the case. Therefore, ticket prices will average out for both those bumped and those seated.
- Each passenger's likelihood of showing up is independent of every other passenger.
- First-class ticket-holders have unrestricted tickets, which allow a full refund in case of no-show; coach passengers have restricted tickets, which allow only a 75% refund in case of no-show.
- There are no walk-ons.
- There are no flight delays or cancellations.
- The marginal cost of adding a passenger to the plane is negligible.

The Model

Equations

$$\text{prob}(x,y,r) = \binom{x}{y} r^y (1-r)^{x-y}$$

$$P_1(y) = \sum_{k=S_f+1-y}^{S_f} \text{prob}(S_f,k,R_f)\big[(-B_c)\big((y-(S_f-k)\big) + F_c(S_f-k)\big]$$

$$P_2(y) = \sum_{k=0}^{S_f-y} \text{prob}(S_f,k,R_f)F_c y$$

$$M_1(x) = \sum_{i=0}^{S_f-y} \text{prob}(x,i,R_c)\big(F_C i + N_c(x-i)\big)$$

$$M_2(x) = \sum_{j=S_c+1}^{x} \text{prob}(x,j,R_c)\big[S_c F_c + N_c(x-j) + P_1(j-S_c) + P_2(j-S_c)\big]$$

$$M(x) = M_1(x) + M_2(x)$$

Parameters

S_f = seating available for first-class

S_c = seating available for coach

R_f = show-up rate for first-class reservations

R_c = show-up rate for coach reservations

F_c = coach fare

N_c = no-show fee for coach

B_c = coach bump cost to airline

Variables

x = number of reservations

Functions

$M(x)$ = expected gain with x reservations

$\text{prob}(x,y,r)$ = probability of y events happening in x trials where r is the chance of a single event happening

$P_1[y], P_2[y]$: to be described later

Binomial Distribution Model

We create ACE Airlines, a fictional firm, to understand better how to handle overbooking. We examine binomial distributions of ticket sales, so we call this the Binomial Distribution Model.

ACE features planes with 20 first-class seats and 100 coach seats. The no-show rate is 10% for coach and 20% for first-class. **Figure 1** compares various overbooking levels with the chance that there will be enough available seats in first-class to accommodate the overflow. The functions are

$$y = \sum_{j=0}^{100+x} \binom{C}{j} (0.9)^j (0.1)^{C-j} \quad \text{(coach)},$$

$$y = \sum_{j=0}^{20-x} \binom{C}{j} (0.8)^j (0.2)^{20-j} \quad \text{(first class)},$$

where C reservations are made for coach and 20 are always made for first class.

Where the first-class line passes below the various overbooking lines indicates the probability at which we must start bumping passengers.

This simplistic model doesn't account for ticket prices, no-show fees, or refunds to bumped passengers and doesn't specifically deal with revenue either. Thus, it can act as a good reference for verifying the customer-relations aspect of any solution but can't give a good solution on its own. To be sure that ACE is receiving the most revenue it can, we must create a more in-depth model.

ACE coach fare is \$200. We refund \$150 on no-show coach tickets, thus gaining \$50 on each. To keep good customer relations, when we are forced to bump a passenger from a flight, we refund the ticket price with an additional bonus of 100% of the ticket price (thus, we suffer a \$200 loss).

We define $\text{prob}(x, y, r)$ as the binomial probability of y independent events happening in x trials, with a probability r of each event happening:

$$\text{prob}(x, y, r) = \binom{x}{y} r^y (1-r)^{x-y}.$$

Model for Coach

We first ignore first class and maximize profit based solely on overbooking the coach section, via the Simple Expected Gain Model. This model is defined in two parts. The first looks at the chances of the cabin not filling—$i < 100$ people showing up. ACE gets \$200 for each of the i passengers who arrive and fly and \$50 from each of the $(x - i)$ no-shows. We multiply the probability of each outcome (determined by the binomial distribution) by the resulting

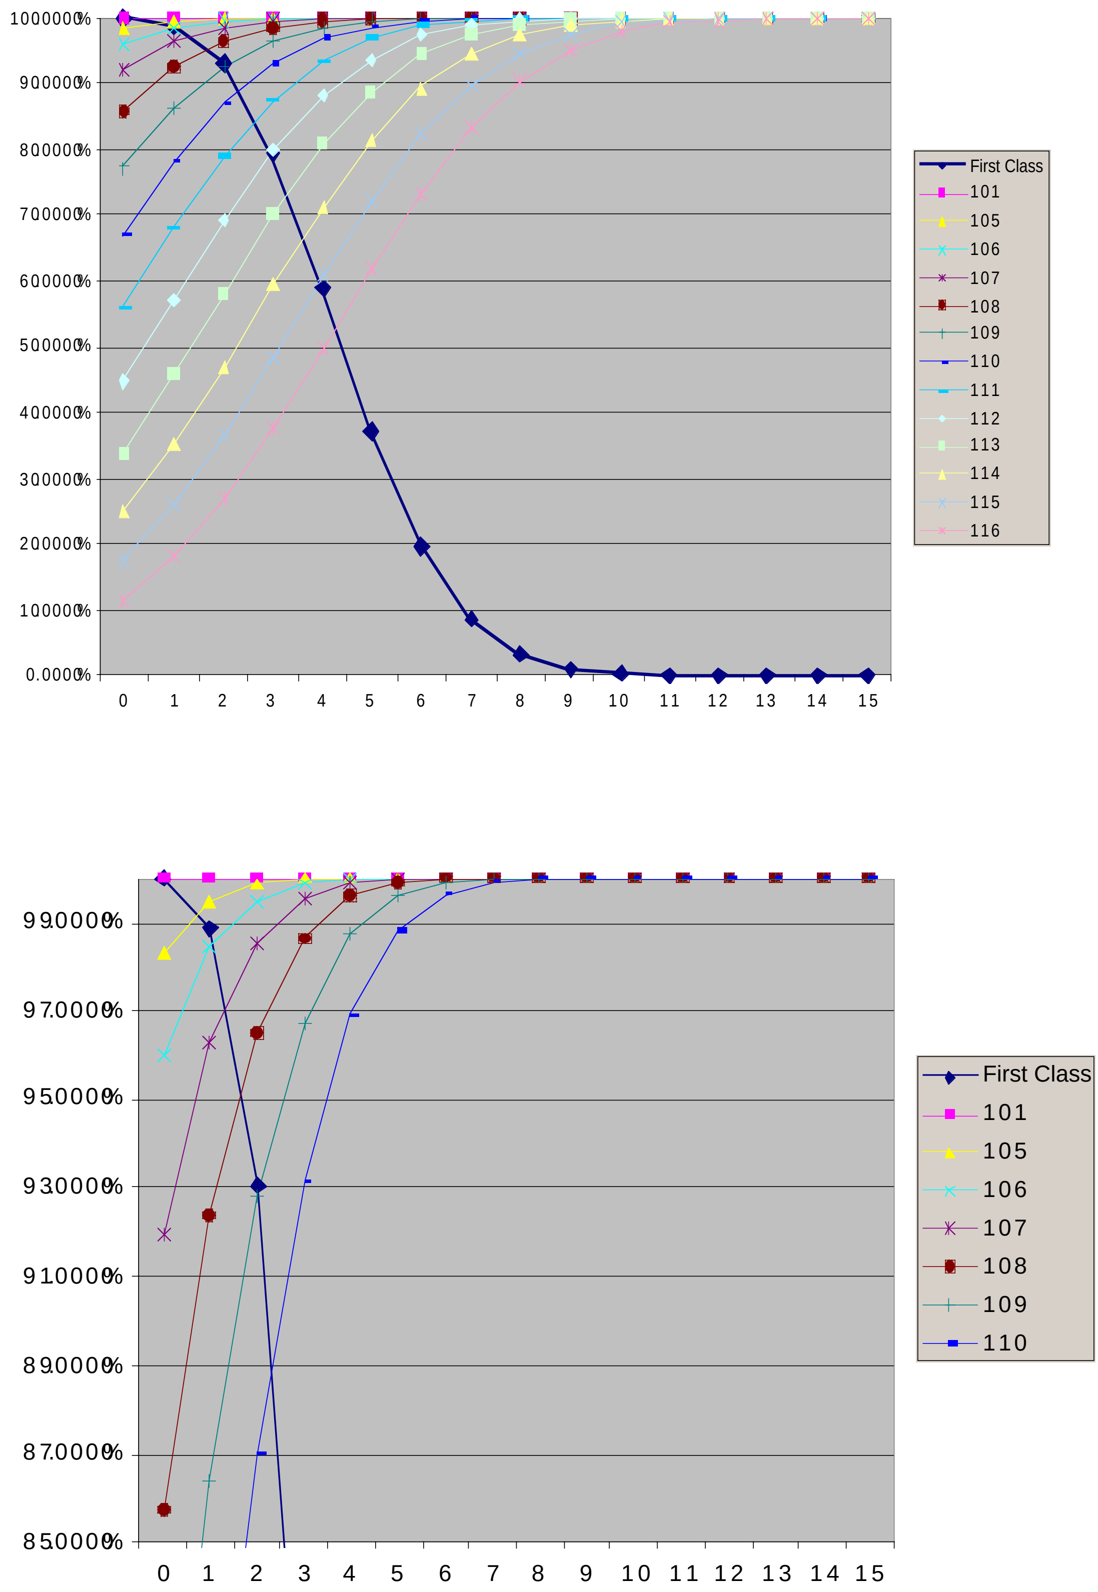

Figure 1. Probability of enough seats vs. overbooking level. The graph below is a close-up of the upper left corner of the graph above.

revenue and sum over all of these values of i to find the expected gain, M_1:

$$M_1(x) = \sum_{i=0}^{100} \text{prob}(x, i, 0.9)\big(200i + 50(x - i)\big).$$

The second part of the model focuses on overflow in the coach section, when $j > 100$. In this case, ACE is limited to \$200 fare revenue on 100 passengers, plus \$50 for each of the $(x-i)$ no-shows. However, for the $(j-100)$ passengers who arrive but have no seats, ACE bumps them and thus loses \$200 in compensation per passenger. We again multiply by the probability of each outcome and sum:

$$M_2(x) = \sum_{j=101}^{x} \text{prob}(x, j, 0.9)\big[200(100) + 50(x - j) - 200(j - 100)\big].$$

We add M_1 and M_2 to arrive at an expression M for revenue. From the graph for M, we discover that (independent of first class) for maximum revenue, ACE should overbook by about 11 people, expecting a net revenue from the coach section of \$20,055 (**Figure 2**).

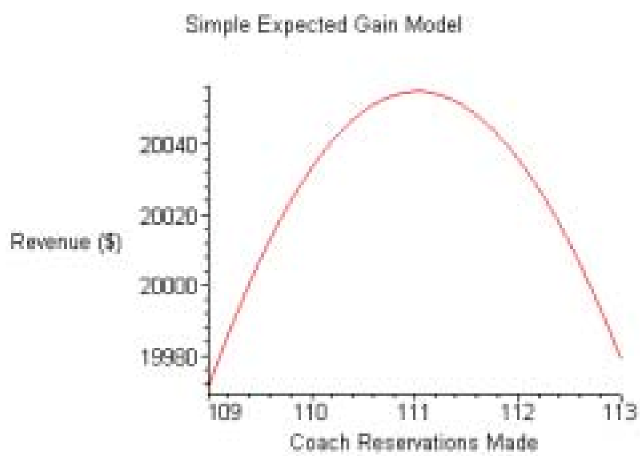

Figure 2. Simple Expected Gain Model: Revenue M vs. number of coach reservations.

Coach Plus First Class

When we add in consideration of first-class openings, ACE can overbook by even more while still increasing revenue, since it can upgrade coach overflow into first-class openings. The first part of the previous formula, M_1 can still be used, since it deals with the cases in which the coach section isn't filled anyway. Since ACE will not overbook first-class, ACE should always book it fully. Thus, fare for first-class is unimportant when considering how to maximize revenue. We further assume that ACE sells only unrestricted first-class tickets (there is no penalty to first-class no-shows).

The second part of the equation needs only a minor modification to adjust for seats made available by first-class no-shows. ACE still gets $200 for each of the 100 passengers who show up and gets $50 for each of the $(x - j)$ no-shows. The difference now is that instead of simply multiplying by $-$\$200 for each passenger over 100, we check for first-class openings and multiply $-$\$200 by just the number who end up bumped. Those upgraded to first-class still pay coach fare (thus, more than 100 coach passengers can pay that $200). This function, $P_1(y)$, with $y = j - 100$ the number of overflow coach passengers, gives the expected net revenue expected for that much overflow. Similarly, $P_2(y)$ gives the expected net revenue when ACE can seat all of the overflow. Thus, the new version of the second part of the formula reads:

$$M_2(x) = \sum_{j=101}^{x} \text{prob}(x, j, 0.9)\big[200(100) + 50(x - j) + P_1(j - 100) + P_2(j - 100)\big].$$

The form of the P_i functions is similar to the two parts of the model already discussed. The probability of there being few enough first-class passengers is multiplied by $200 (coach fare) times the number of extra coach passengers who can be seated $(j - 100)$. Recall that the show rate for first-class is 0.8:

$$P_2(y) = \sum_{k=0}^{20-y} \text{prob}(20, k, 0.8)(200)y.$$

The other case is when ACE can't seat all of the coach overflow. This time, we multiply by the loss of revenue from the coach spillover y, $200 for each of the $y - (20 - k)$ bumped customers, offset by $200 for each of the $(20 - k)$ passengers upgraded to first-class. The result is

$$P_1(y) = \sum_{k=21-y}^{20} \text{prob}(20, k, 0.8)\big[(-200)\big((y - (20 - k)\big) + 200(20 - k)\big]$$

At this point, we have all of the pieces of the expected gain model. We plot the equation $M = M_1 + M_2$ in **Figure 3** and find the maximum for $x \geq 100$.

The ideal overbooking level lies at 115 or 116 reservations, with a negligible difference in profit ($0.07) between them.

Applying the Model

We implemented our model in a computer program in which the parameters can be varied, including seating capacities, ticket price, and no-show fees.

In the case of our example, the optimum is very broad around 116. When deciding optimal overbooking levels, the airlines must balance revenue is against the chance of bumping. If ACE books 113 passengers instead, the revenue decreases by $105 per flight but the probability of no bumping rises to 73% from 53%. Similarly, if it books 114 passengers, it loses $35 per flight but there is a 67% probability that no one will be bumped.

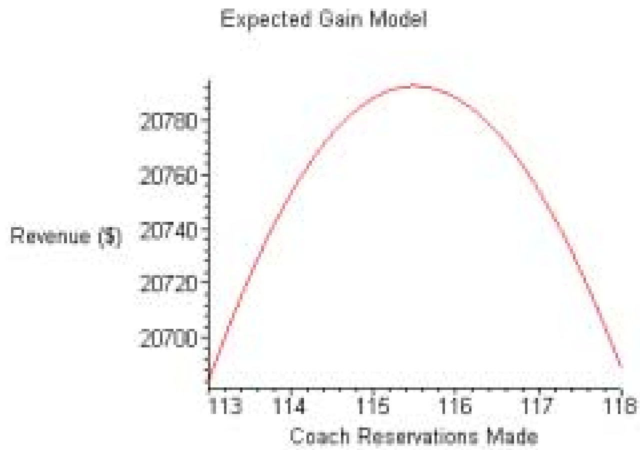

Figure 3. Expected Gain Model: Revenue M vs. number of coach reservations.

Fewer Flights

The decrease in air traffic by 20% since September 11 means fewer flights. Due to more-detailed security checks, it is necessary for planes to have longer turnaround times between flights. Adding 15 extra minutes at each turnaround would cause an airline such as Southwest to need almost 100 additional planes to maintain its previous air traffic flow. Therefore, there are fewer flights.

How does this circumstance affect our model?

Federal regulations do not require compensating a bumped passenger scheduled to reach the destination within an hour of the original arrival time. But now the probability of accomplishing that is much smaller than before September 11th; we do not consider it likely and do not include it in the model.

Can bumped passengers be put onto a later flight to arrive within two hours of their original scheduled time? If this happens, federal regulations require an airline to compensate them for a ticket, essentially flying them for free. There is no loss or gain from this transaction, which is certainly more desirable than paying every bumped passenger $200 on top of refunding ticket price.

If people are put onto later flights, ACE pays fewer passengers an extra $200. However, our Expected Gain Model attempts to maximize the number of people on the flight. Thus, the probability of a passenger being able to take a later flight is very low and the optimal overbooking level changes negligibly. For example, disregarding first class, our Expected Gain Model shows only a 7% chance of a coach seat available on the next flight. We conclude that the Expected Gain Model is just as effective and much simpler if we disregard the possibility of bumped passengers obtaining a seat on a later flight, so we assume that all bumped passengers are compensated with a refund of their ticket price and $200.

Heightened Security and Passengers' Fear

Demand for flying is down, despite funding for additional security, which—while well-justified—causes problems for airlines and passengers.

ACE is concerned about passengers who miss flights because of security checks. These, along with passengers' fear, can increase the no-show rate, which ACE must consider in its overbooking strategy. Passengers may reserve a seat but then decide that in light of events they are too frightened to get on the plane. With higher no-show rates, a higher overbooking rate may become optimal. In our expected gain model with no-show rates of 20% and 30% for coach and first-class, the optimal overbooking level jumps to 130 or 135 seats.

Dealing with Bumped Passengers

While most ways that airlines have dealt with bumping passengers are subtle and good business practice, some border on the absurd. For example, until 1978, United Airlines trained employees to bump people less likely to complain: the elderly and armed services personnel—two groups that perhaps instead should have priority in seating!

A strategy other than current compensations could be optimal for the airlines, but it depends on cooperation. If ACE could convince other airlines that they all should give bumped passengers discount tickets (usable on any of the airlines), then each airline would lose less money from compensating bumped passengers. This would create a mutually profitable situation for all airlines involved: The airline accepting bumped passengers would fill seats that would otherwise be empty; the airline bumping the passengers would cut the amount of compensation to the price of a discounted ticket.

Suppose that the amount of compensation is decreased by one-half. The optimal level of overbooking rises, as does revenue; but we cannot be sure that every bumped passenger can be placed on another flight.

Strengths and Weaknesses of the Model

Strengths

- Our model involves only basic combinatorics and elementary statistics.
- Because it is parametrized, the model can continue to be used as rates, seating capacities, and compensation amounts change.
- The model considers more than one class.
- An airline can attempt to find the balance between maximizing revenue and pleasing customers, depending on how much risk the airline chooses to take.

Weaknesses

- We do not consider business class; including it would have risked the model being too complicated. Business class should not have a large effect on revenue maximization, because no-show rates are lower and business people are more concerned with reaching their destination on time than surrendering their seats for compensation.
- Our model does not take into consideration how multiple flights affect each other. If putting passengers onto later flights were an option, revenue would increase slightly but doing so would also further complicate optimal overbooking levels on other flights.
- Because ACE is not bumping passengers to later flights, bumped passengers are left out in the cold with no flight and just a little bit of extra money—a resolution that does not provide positive customer relations.
- We allow no overbooking in first class. If ACE is willing to take the risk of downgrading or bumping first-class passengers, then revenue could increase slightly by overbooking first-class seats.
- In reality, anyone can buy a restricted or unrestricted ticket in either class. Therefore, a more complicated model would include the possibility of some coach no-shows receiving full-refund and some first-class no-shows paying a no-show fee.
- Our binomial distribution for showing up assumes independence among passengers. However, many people fly and show up in groups.

References

Airline Passengers to pay security fee next month. 2002. *USA Today* (3 January 2002). `http://www.usatoday.com/news/washdc/jan02/2002-01-03-airlinesecurity.htm` .

Alexander, Keith, and Frank Swoboda. 2002. US Airways posts record $1 billion quarterly loss. *Washington Post* (18 January 2002). `http://www.washingtonpost.com/ac2/wp-dyn?pagename=article&node=&contentId=A64400-2002Jan17` .

Bayles, Fred. 2001. New safety procedures likely to add costs to flying. *USA Today* (6 November 2001). `http://www.usatoday.com/money/biztravel/2001-11-06-security-costs.htm` .

Karaesmen, Itir, and Garrett van Ryzin. 1998. Overbooking with substitutable inventory classes. (20 September 1998). Columbia University. `http://www.columbia.edu/~gjv1/subob1.pdf` .

More airline passengers get bumped. 2000. *USA Today* (6 October 2000). `http://www.usatoday.com/life/travel/business/1999/bt002.htm` .

Schmeltzer, John. 2002. Weak economy, Sept. 11 continue to buffet carriers. *Chicago Tribune* (4 January 2002). `http://chicagotribune.com/business/chi-0201040415jan04.story?coll=chi-business-hed>` .

Spiegel, Murray R. 1988. *Theory and Problems of Statistics*. 2nd ed. New York: McGraw-Hill.

US airline chiefs seek assurances on finance and security. 2001. *AIRwise News* (18 September 2001). `http://news.airwise.com/stories/2001/09/1000839754.html` .

U.S. Department of Transportation. 2002. Air Travel Consumer Report. February 2002. `http://www.dot.gov/airconsumer/atcr02.htm` .

Memorandum

Date: 02/11/2002
To: CEO of ACE Airlines
From: Aviophobia University
RE: Your Troubles Solved

Today is your lucky day!! We know that airlines have been going through especially hard times recently and so we have come up with something that will solve your problems.

You and I both know that overbooking occurs not because you cannot count the number of seats on your plane, but rather because it is a brilliant business strategy that can increase revenue. We have created a model that allows you to find your optimal overbooking strategy.

Our model can consider your specific situation because it can account for different no-show rates and fees, seat capacities, ticket prices, and bumped passenger compensations. We have designed an easy-to-use computer program that allows you to quickly find your optimal overbooking strategy based on your figures. This program saves you time in a business where time is money. In addition, using our model will allow you to maximize your revenue without bringing in an expensive consultant.

When designing our model, we even used data concerning your airline, so half of the work is done for you! For your planes, fares, and policies, our model shows that 16% overbooking is optimal for maximizing revenue. However, we find that to reduce the probability of bumping too many passengers and still maintain a high revenue rate, 14% or 15% is ideal.

We hope that this information leads you to a profitable quarter and stock increase, which we would both find profitable.

Bumping for Dollars: The Airline Overbooking Problem

John D. Bowman
Corey R. Houmard
Adam S. Dickey
Wake Forest University
Winston-Salem, NC

Advisor: Frederick H. Chen

Introduction

We construct a model that expresses the expected revenue for a flight in terms of the number of reservations, the capacity of the plane, the price of a ticket, the value of a voucher, and the probability of a person showing up for the flight. When values are supplied for every variable but the first, the function can be maximized to yield an optimal booking that maximizes expected revenue.

We apply the model to three situations: a single flight, two flights in a chain of flights, and multiple flights in a chain of flights. We conclude that fewer flights will increase the value of the penalty or voucher and thus decrease the optimal number of reservations. Heightened security also lowers the optimal number of reservations. An increase in passengers' fear decreases the probability that a person will show up for a flight and thus increases the optimal number of reservations. Finally, the loss of billions of dollar in revenue has no effect on the optimal value of reservations.

We model the probability of a given number of people showing up as a binomial distribution. We express the average expected revenue of a flight in terms of the number of bookings made.

Starting with the Single-Flight case, we derive a model and revenue function for a flight unaffected by previous flights. From this situation, we expand the model to the Two-Flight case, in which the earlier flight affects the number of people who show up for the later flight. We generalize the model even further to the number of people showing up depending on many previous flights.

The UMAP Journal 351–364.

The Model

In each of the three situations modeled, we derive two formulas. The first, $P(k)$, describes the probability that k people show up for a flight. The second, Revenue(b, c, r, x, p), describes the expected revenue for a flight as a function of the number of reservations. We verified these theoretical equations by a Monte-Carlo simulation.

For the Single-Flight Model:

$$P(k) = \binom{n}{k} p^k (1-p)^{n-k},$$

$$\text{Revenue}(b, c, r, x, p) = \sum_{k=0}^{c+(b-c)} P(k)\big[r \min(k, c)\big) - x \max(k-c, 0)\big].$$

For the Two-Flight Model:

$$P_2(k) = P_1(k)\left[1 - \sum_{i=c+1}^{b} P_1(i)\right] + \sum_{j=1}^{b-c} P_1(k-j)P_1(c+j),$$

$$\text{Revenue}_2(b, c, r, x, p) = \sum_{k=0}^{c+2(b-c)} P(k)\big[r \min(k, c)\big) - x \max(k-c, 0)\big].$$

For the n-Flight Model:

$$P_n(k) = P_1(k)\left[1 - \sum_{i=c+1}^{c+(n-1)(b-c)} P_{n-1}(i)\right] + \sum_{j=1}^{(n-1)(b-c)} P_1(k-j)P_{n-1}(c+j),$$

$$\text{Revenue}_n(b, c, r, x, p) = \sum_{k=0}^{c+n(b-c)} P_n(k)\big[r \min(k, c)\big) - x \max(k-c, 0)\big].$$

The variables are:

b = number of reservations (or bookings) per flight

c = plane capacity

r = price of a ticket

x = value of a voucher

p = probability that a booked passenger shows up for a flight

Given p, c, r, and x, the method finds the b that maximizes revenue.

Derivation of the Single-Flight Model

The binomial distribution applies to calculating the probability that a number of passengers shows up for a flight:

- The probability involves repeated events (each trial calculates the probability of one person showing up) with only two possible outcomes (either the person is a show or no-show).
- We assume that people's actions do not influence one another; each person's chance of showing up is independent of another person's chance. This is not true in reality, as people often travel in groups; but this a necessary and appropriate simplification.
- We assume that the probability of a person arriving remains constant for each person.

We use the binomial distribution to calculate expected revenue. Airlines overbook their flights, knowing that some people will not take the flight. Given a certain overbooking strategy b (i.e., the maximum number of reservations taken for a particular flight, with $b > c$, the capacity of the plane), the expected revenue is

$$\text{Revenue}(b, r, p) = \sum_{k=0}^{c+(b-c)} P(k) r \min(k, c).$$

The function is incomplete, however, because it does not penalize the airline for the consequences of overbooking. The airline usually provides bumped passengers with either an airline ticket voucher or a cash reimbursement, valued at x per bumped person:

$$\text{Revenue}(b, r, p) = \sum_{k=0}^{c+(b-c)} P(k) \left[r \min(k, c) - x \max(k - c, 0) \right].$$

When $k \leq c$, the x term is zero; when $k > c$, the airlines is penalized for having to bump people.

The booking decision b and the capacity c of the plane are fixed before the model begins. This model considers just one flight in a complex network of flights; it does not allow for the possibility that passengers are bumped from a previous flight, since it assumes that the only passengers are those who made a reservation for this particular flight. The model also applies to just one flight: If the number of passengers who show up is greater than the capacity of the plane, those bumped passengers receive a voucher and—with a wave of the magic wand of assumption—disappear. Finally, regardless of the flight's destination (Hawaii or Death Valley), we assume that there is enough demand to fill the predetermined number of bookings.

Since p is constant throughout our model, the Revenue function is really dependent only on the number of bookings, the capacity of the plane, the cost of a ticket, and the cost of the penalty.

Application of the Single-Flight Model

We set $p = .9$. Since b must be an integer, the revenue function is not continuous. Thus, the analytic method of maximizing of the function (namely, differentiating and setting the derivative equal to zero) cannot be applied. Instead, we use Maple 6.

After setting values for the probability, plane capacity, and ticket and voucher values, we evaluate the function at $b = c$, then increment b until a maximum for Revenue is found.

We used three plane capacities: 10, 30, and 100. The values of the ticket price r, the voucher x, and the arrival probability p are held constant at $300, $300, and .9 for the examples in **Table 1**.

Table 1.

Results for the Single-Flight Model.

Capacity	Optimal overbooking	Revenue	Bump probability
10	11	$2,782	31%
30	33	$8,598	35%
100	111	$29,250	44%

The probabilities of bumping are larger than the industry frequency of about 20%. Worse, the model ignores all the problems created by these bumped passengers. The model is further weakened in light of the post-September 11 issues proposed by the problem. Among the four issues—fewer flights, heightened security, passengers' fear, and losses—this model can account only for increased passenger fear (indicated by a change in probability that a passenger shows up). Clearly this Single-Flight Model is not a proper solution to the airline-overbooking problem.

Derivation of the Two-Flight Model

The Two-Flight Model begins with updating both the probability and revenue functions. Unlike the Single-Flight Model, the new functions reflect the possibility that passengers bumped from one flight fill seats on the next. By this assumption, the probability function for the second flight, $P_2(k)$, changes, because k may now also be expressed as a combination of people ticketed for the second flight and bumped passengers from the first flight. Since the revenue function depends on the probability function, it too must change.

$$\begin{aligned} P_2(k) &= Pr(k \text{ people show up for flight 2}) \\ &= Pr(k \text{ regular passengers arrive}) Pr(\text{no one bumped from flight 1}) \\ &\quad + Pr(k-1 \text{ passengers arrive}) Pr(1 \text{ passenger bumped}) + \cdots \\ &\quad + Pr(k-j \text{ arrive}) Pr(j \text{ passengers bumped}) + \cdots \\ &\quad + Pr(k-(b-c) \text{ arrive}) Pr(b-c \text{ passengers bumped}) \\ &= P_1(k)\left[1 - \sum_{i=c+1}^{b} P_1(i)\right] + \sum_{j=1}^{b-c} P_1(k-j) P_1(c+j). \end{aligned}$$

A maximum of $b - c$ people can be bumped from flight 1, since at most b people show up for it and we assume that no passengers are carried over from any previous flight. The probability that 1 passenger is bumped from flight 1 is exactly the probability that $c + 1$ people are present for it. Thus we have $Pr(j \text{ passengers bumped}) = P_1(c + j)$. As long as b, p, and c remain the same, the probability that new (prebooked, non-bumped) passengers arrive never changes; it is independent of the number of bumped passengers from a previous flight. (We assume that there is no way for a passenger to know how many people have been bumped onto his or her flight from a previous one.) Thus, $Pr(k - j$ regular passengers arrive) will always be computed by $P_1(k - j)$, our original probability function for the Single-Flight Model.

In the second summation of $P_2(k)$, the term $k - j$ could be negative for small k. If so, we define the probability of a negative number of people showing up from a previous flight to be 0 (empty seats on a flight cannot be filled by passengers from later flights!).

We now express the second revenue function in terms of the second probability function:

$$\text{Revenue}_2(b, c, r, x, p) = \sum_{k=0}^{c+2(b-c)} P(k)\big[r \min(k, c - x \max(k - c, 0)\big].$$

A passenger bumped from one flight is automatically booked on the next flight and seated before regular passengers, so as to have almost no chance of being bumped again. For the second flight, we assume that the number of people who show up is affected only by that flight and the previous flight, and that there is enough demand to fill the predetermined number of bookings.

The summation now has $c+2(b-c)$ as it maximum value. The second flight must not only account for b passengers but must also account for the number of people possibly bumped from the first flight.

Application of the Two-Flight Model

By introducing a second flight, we more accurately model the situation. The optimal overbooking strategy and maximum revenue should either remain the same or slightly decrease.

Using the Revenue function for the Two-Flight Model, we now calculate maximum revenue and the associated overbooking strategy for the same plane capacities as for the Single-Flight Model. Again, the values of the ticket price r, the voucher x, and the arrival probability p are held constant at 300, 300, and .9. The results are in **Table 2**.

Table 2.

Results for the Two-Flight Model.

Capacity	Optimal overbooking	Revenue	Bump probability
10	11	\$2,745	34%
30	33	\$8,551	42%
100	111	\$29,107	57%

In each case, the optimal booking strategy is the same as the Single-Flight Model, but the maximum revenues is lower, and bump probability is higher. Since both flights are overbooked, the probability that someone is bumped should only increase.

The n-Flight Model

We generalize to n flights. We allow each flight to be influenced by the $(n-1)$ flights before it. We still assume that a passenger bumped from one flight is given preferential seating on the next. However, giving seats to bumped passengers who are already at the airport decreases the number of seats for pre-booked passengers. The n-flight model allows this domino effect of bumping to ripple through $n-1$ successive flights. As n gets large, our model becomes a better and better approximation of the real case, in which every flight is affected by many previous flights. Our probability function becomes recursive:

$$P_n(k) = P_1(k)\left[1 - \sum_{i=c+1}^{c+(n-1)(b-c)} P_{n-1}(i)\right] + \sum_{j=1}^{(n-1)(b-c)} P_1(k-j)P_{n-1}(c+j),$$

$$\text{Revenue}_n(b,c,r,x,p) = \sum_{k=0}^{c+n(b-c)} P_n(k)\big[r\min(k,c)\big) - x\max(k-c,0)\big].$$

For the first summation, zero people show up from the previous flight, meaning that there are enough seats for everyone on that flight and anyone bumped from a previous flight. If the total possible number of people who can show up to the current flight is $b+(n-1)(b-c)$ (as is explained in a moment), then the total number of people who can show up for the previous flight must be $b + (n-2)(b-c)$, which we use as the upper limit of the summation.

For the second summation, we use $(n-1)(b-c)$ instead of $(b-c)$, since now there can be at most $(n-1)(b-c)$ passengers bumped from flight $n-1$. This upper bound for bumped passengers can be proved by mathematical induction. [EDITOR'S NOTE: We omit the authors' proof.]

The revenue function for the 2-flight model can also be extended to n flights in a straightforward way. Note that at most $n(b-c)$ people can be bumped from the nth flight. We have:

$$\text{Revenue}_n(b) = \sum_{k=0}^{c+n(b-c)} P_n(k)\big[r\min(k,c) - c\max(k-c,0)\big].$$

We now consider booking strategies to optimize revenue.

Computation of the n-Flight Model

The Recursive Method

We can create documents in Maple to compute the probability and revenue functions. To compute $\text{Revenue}_n(b)$, we must evaluate $P_n(k)$ a total of $b + (n-1)(b-c)$ times. In turn, $P_n(k)$ must evaluate $P_{n-1}(k)$ at total of $(2n-1)(b-c)$ times, $P_{n-1}(k)$ must evaluate $P_{n-2}(k)$ a total of $[2(n-1)-1](b-c) = (2n-3)(b-c)$ times, and so on. Thus, without even accounting for all the evaluations of $P_1(k)$ in each iteration, we make

$$[b+(n-1)(b-c)](2n-1)(b-c)(2n-3)(b-c)\cdots[2n-(2k+1)](b-c)\cdots(1)(b-c)$$
$$= [b+(n-1)(b-c)\,\frac{(2n-1)!}{2(n-1)!}\,(b-c)^{n-1}$$

function calls. With $b = 105$ and $c = 100$, $\text{Revenue}_2(k)$ requires 1650 function calls, $\text{Revenue}_3(k)$ requires 86,250 calls, and $\text{Revenue}_4(k)$ requires more than 6.3 million function calls. The computation time is proportional to the number of function calls: $\text{Revenue}_2(105)$ takes less than 1 s, $\text{Revenue}_3(105)$ takes 13 s, and $\text{Revenue}_4(105)$ takes 483 s.

Of course, this is a very inefficient method. A more efficient method would be to store all probability values in an array, beginning with the values for $P_1(k)$ and working upwards to $P_n(k)$. However, Maple makes array manipulation difficult. Instead, we turn to another method.

[EDITOR'S NOTE: Mathematica (and perhaps Maple too) provides an easy-to-use capability for computation of such probabilities via dynamic programming. For an example of its use, see "Farmer Klaus and the Mouse," by Paul J. Campbell, *The UMAP Journal* 23 (2) (2002) 121–134.

Monte Carlo Simulation

We develop a Monte Carlo computer simulation coded in Pascal that runs the n-flight model numerous times and determines the average revenue for a large number of trials. Instead of obtaining precise probabilities using the functions developed above, we flip a (electronic) weighted coin to determine whether each individual passenger shows up for the flight. We tell the program how many trials to run, give it values for n, p, c, r, and x and tell it the largest value of b to check. The program begins with $b = c$. It flips numerous weighted coins to determine how many passengers show up for the first flight. It bumps any excess passengers to the second flight and flips coins again to see how many prebooked passengers arrive. The excess is bumped to the third flight and the process continues until the nth flight. Revenue is evaluated by adding an amount equal to the ticket price for each passenger who flies and deducting a penalty for each passenger who is bumped. The program iterates for successive values of b until it reaches the preassigned upper bound.

The output includes, for each b value, the mean revenue over all trials and the corresponding percentage standard error. Percentage standard error was usually less than 2% and often less than 1%.

Optimization Strategies for the n-Flight Model

We will never earn more than the ticket price (r) times the number of seats (c), so the gain from overbooking is limited—but the possible costs are not. At some point, the costs of overbooking outweigh the benefits; there should be a unique maximum for revenue.

To find the maximum revenue, we evaluate the revenue function at different booking values, beginning with $b = c$, until we find a b with Revenue$(b-1) <$ Revenue(b) and Revenue$(b+1) <$ Revenue(b). This will be our b_{opt}.

The obvious booking strategy is to book every flight with b_{opt} passengers. While this method maximizes flight revenue, it yields a high percentage of flights with bumped passengers. For a plane with 100 seats, the maximum revenue occurs at $b = 108$, with 34% of flights bumping passengers. Because our model does not account for changes in demand due to the airlines' behavior, this might not be the truly optimal value of b in the long run. Bumping large numbers of passengers will drive customers away; decreased demand will depress the price that we can charge and we reduce revenue in the long term. Similarly, an especially low percentage of bumped flights may increase demand, allow us to raise prices, and increase revenue. Thus, our model accounts only for short-term effects, not long-term ones.

Moving away from maximum revenue lowers expected revenue by a small amount but decreases the bump probability by a large amount. For convenience, we set both the price and penalty to \$1, to avoid large numbers. While \$1 is unrealistic, the value does not change the optimal booking strategy from the case where both price an penalty are both \$300, because it is the ratio of price to penalty—and not their actual values—that changes the optimal booking. Our example considers a 50-flight sequence of planes with capacity 100 each; if everyone showed up and there was no overbooking, the revenue would be \$5,000. At the optimal $b = 108$ for $p = .9$, the expected revenue is \$4,806 with bump probability of 33%. If we move down just 1 to $b = 107$, the revenue is \$4,791 and the bump probability drops to 21%.

Table 3.

Results of simulation: for each number for bookings, 100 trials with 50 flights per trial.

Bookings	Revenue	% Bump	Delta(%Bump)	Delta(%Rev)	D(Bmp/D(Rev)
100	\$4,501	0.00%	0.00%	0.02%	0.00
101	\$4,539	0.02%	0.02%	0.84%	0.02
102	\$4,589	0.04%	0.02%	1.09%	0.02
103	\$4,631	0.04%	0.58%	0.92%	0.63
104	\$4,677	0.62%	2.24%	0.97%	2.30
105	\$4,722	2.86%	5.66%	0.97%	5.82
106	\$4,758	8.52%	12.50%	0.76%	16.45
107	\$4,791	21.02%	12.64%	0.68%	18.64
108	\$4,807	33.66%	22.10%	0.34%	65.67
109	\$4,800	55.76%	17.04%	−0.15%	−115.28

We adopt as a criterion to compare two values of b the ratio of the relative change in the bump probability and the relative change in revenue:

$$\frac{\dfrac{\Delta P_{\text{bump}}}{P_{\text{bump}}}}{\dfrac{\Delta\text{Revenue}}{\text{Revenue}}} = \frac{\Delta(\%\text{Bump})}{\Delta(\%\text{Revenue})}.$$

The process goes: A maximum revenue is found, along with its high bump probability. The optimizer now considers a lower value of b and looks at the ratio of the change in the bump probability to the change in revenue. If this ratio is above a certain value k, the optimizer accepts the lower b. The optimizer continues to do this until the ratio is no longer greater than the constant. In **Table 3**, with $k = 20$, the new optimum b would be 107, because the ratio 18.64 is not greater than $k = 20$.

Table 4 shows three different optimization values for different plane capacities.

Application of the n-Flight Model

The problem specifically mentions four issues to be addressed by our model: fewer flights, heightened security, passengers' fear, and revenue losses.

Why are airlines offering fewer flights? If the airlines had kept offering the same number of flights, the question of an optimal overbooking strategy would be moot, because the planes would not fill. The huge drop in demand

Table 4.

Optimal bookings using different criteria ($p = .9$, $r = 1$, $x = 1$). For each number for bookings, 100 trials with 50 flights per trial.

c	Maximum revenue			$k = 20$			$k = 1$		
	b	Rev	$P_{\text{bump}}(\%)$	b	Rev	$P_{\text{bump}}(\%)$	b	Rev	$P_{\text{bump}}(\%)$
10	10	450	0	10	450	0	10	450	0
30	32	1390	45	31	1385	12	30	1350	0
50	54	2360	50	53	2360	25	51	2290	0.8
100	109	4810	50	107	4790	19	104	4670	0.7
150	163	7270	37	161	7220	14	157	7070	0.3
200	219	9740	47	215	9660	9	211	9490	0.4
280	307	13690	46	303	13610	13	299	13450	1.3

has reduced supply but could also result in slashed prices. Since the value of the compensation involuntarily bumped ticket-holders is tied to the ticket price (though with a ceiling), changes in ticket prices should affect the optimum booking level little if at all.

However, the fewer flights, the longer people who are denied boarding must wait for the next flight; being denied boarding is less convenient. Since compensation is usually offered in a kind of auction to induce voluntary relinquishing of seats, the airline will have to offer more. Therefore, longer delays between flights will increase the ratio of compensation amount to ticket price, tending to decreasing the optimal booking level.

How do heightened security measures affect our model? They mean more security checks, longer lines, longer waits, and an increased chance of missing a flight, particularly a connecting flight. Unfortunately, people who miss their connecting flight and thus are guaranteed a spot on the next flight are not included in our model explicitly; but they do have an implicit effect. If more people miss connecting flights, they put additional stress on the system: They increase the chance that the next and subsequent flights will have too many people. Therefore, in our booking strategy, we want a low bump probability. To attain it, we should decrease the ratio k, which decreases optimal booking level b.

Passenger fear leads not only to decreased demand (which we have already considered above) but also to a decreased probability p of a passenger showing up, which in turn increases the optimal booking level b.

However, the hardest to deal with is the huge revenue loss. Less profitable airlines may fold; but presumably if there is excess demand, other airlines will either add flights or raise the price. Hence, though the huge financial loss may change the industry as a whole, it doesn't affect the optimal booking strategy. It merely leads to fewer flights (already addressed) and may change prices (which we argued would have no effect).

We summarize these effects in **Table 5**.

Table 5.
Effects of post-September 11 factors.

Factor	Direct effect	Effect on optimal booking level
Fewer flights	$g \uparrow$	$b_{\text{opt}} \downarrow$
Heightened security measures	$\Delta P_{\text{bump}}/\Delta\text{Revenue} \downarrow$	$b_{\text{opt}} \downarrow$
Passenger fear	$p \downarrow$	$b_{\text{opt}} \uparrow$
Financial losses	—	—

Verification and Sensitivity of the Model

Since at least 100 trials were used per calculation, the Central Limit Theorem assures us that the distribution of the sample mean approximates well a normal curve and we can be 95% confident that the true value we are approximating is within two standard errors of the sample mean. Often this means we cannot be completely sure of the optimal b, because the maximum revenue is within two standard errors of the revenues of the values for b immediately above and below.

Convincing for us is that for small n and small c, the simulation provides values very close to those from the exact solutions processed in Maple. Because of the agreement, we are confident that our simulation is coded correctly and that the simulations are accurate, even for higher n and c.

That the simulation may be off by 1 for the optimal value of b is not much of a problem. For large c, though Maple may be too slow to calculate over a large range of values for b, Maple can be used to spot-check the value of b from the simulation, along with the ones immediately above and below.

In fact, we need not be much concerned about $n \geq 5$. A bumped person affects a second flight and may also affect a third and possibly a fourth flight. But the effect diminishes, so while the effect on flights close by cannot be discounted, ignoring her effect on a tenth flight does no great damage.

One might expect that changes in both booking level b and capacity c would significantly change the behavior of the model. But around b_{opt}, the revenue curve is fairly flat. For example, for $n = 50$, $c = 100$, and $r = x$, using $b_{\text{opt}} + 1$ instead of b_{opt} decreases revenue by only 0.12%, whereas adopting $b_{\text{opt}} - 1$ instead decreases revenue by only 0.21%. This insensitivity is important because one of our more limiting assumptions is constant p. Since slightly changing b only slightly changes revenue, the effect of varying p should not be too detrimental.

What is sensitive to changes in b is the bump probability. Using the same example as before, moving to $b_{\text{opt}} + 1$ increases the bump probability by 15 percentage points, while moving to $b_{\text{opt}} - 1$ decreases it 11 percentage points. While the smallest percentage changes in revenue are grouped around b_{opt}, the largest percentage changes in bump probability are grouped there.

Strengths, Weaknesses, and Extensions

Strengths

- The strong correspondence between the Maple calculations and the data from the simulation is quite heartening.

- Around b_{opt}, the revenue is insensitive compared to the bump probability. Variations on the n-flight mode provide a small range of near-optimal bs with similar results for revenue and a fairly wide range bump probability. The range allows an airline a choice.

Weaknesses

- The most obvious defect of our model is that many overbooking strategies are in use—and none of the them is ours! Our model is very restrictive because it assumes a constant booking strategy, as well as constant levels of p and c. In reality, most airlines use a dynamic system in which the overbooking level is not constant but instead is varied based on conditions that change from day to day and flight to flight.

- We replace the nation's vastly complicated network of intermeshing flights with a single flight path.

- We simplify the oligopoly of airlines to a single airline.

- We do not account for no-shows such as missed connections that are the fault of the airline or due to circumstances beyond its control (e.g., weather). In such circumstances, a flight's chance of being full is influenced by previous flights even if there is no overbooking.

- In assuming a binomial distribution, we assume people do not travel in groups, and thus their showing up are independent events.

Potential Extensions

- The bump probability could affect revenue in a way that we have not allowed for, namely, in terms of price. An airline that consistently offers better service should be able to charge a higher price. A way to incorporate this effect is to make price a function of bump probability, perhaps inversely proportional to it.

- It might be desirable to make the compensation x a function of the percentage of people that must be excluded from the plane. If 50% of the ticket-holders had to be excluded, then the incentives would have to be greater than if only 5% had to be excluded. At some point the airline would stop raising the incentive and resort to involuntary denied boarding, but these would also have costs resulting from customer satisfaction. One could experiment with setting x equal to some constant times the ratio of those to be bumped, $m-c$, to the total number of people m.

- The probability function could easily be generalized to variable p; in that case, $P(k)$ would become $P(k, p_n)$. The equation could be generalized to the planes having different values of c and b by changing the upper limit of the summations from $(n-1)(b-c)$ to $\sum_{i=1}^{n-1}(b_i - c_i)$.

References

About.com. 2000. Voluntary and Involuntary ("Bumped") Denied Boardings (Jan-Sep 2000).http://airtravel.about.com/library/stats/blrpt8.htm.

Airlines Overbooking Project. n.d. http://math.uc.edu/~brycw/classes/361/overbook.htm .

Airlines Transport Association. 2001. ATA Annual Report 2001.http://www.airlines.org/public/industry/bin/2001FactFig.pdf .

Conway, Joyce, and Sanderson Smith. 1997. Overbooking airline flights. http://www.keypress.com/tswt_projects/projects97/conwaysmith/airline_overbooking.html .

U.S. Department of Transportation, Office of Aviation Enforcement and Proceedings. n.d. Passengers denied boarding by U.S. airlines.http://www.dot.gov/airconsumer/984yovs.htm .

Memo

To: CEO, TopFlight Airways
From: Models R Us
Re: Optimal Overbooking Strategy

Dear Sir/Madam:

We have heard of your company's financial hardships in the wake of September 11. We offer you our assistance. We are a team of students who have dedicated four intense days to understand the problem of airline overbooking. While many have been working on this problem for years, we feel our approach will give your company the extra edge you are seeking.

Because only 90% of passengers arrive for their scheduled flights, an overbooking strategy is necessary to maximize revenue. However, there is a penalty for overbooking. As you know, airlines offer vouchers and other incentives to passengers to entice them to give up their seats. The airline is also responsible for finding bumped passengers a later flight.

Our model incorporates these features. We consider the effect on a given flight of any number of preceding flights. If too many passengers arrive from a previous flight, they can set off a domino effect; when these passengers are rescheduled on a later flight, they increase the chance that this flight, too, will be overbooked.

Our model allows you to combat this effect by finding the optimal booking level for a plane of a given capacity. We did computations for the model in two different ways: once using the mathematical software package Maple and again using a Monte Carlo simulation developed in Pascal. We found the values for these two computational approaches to be in very close agreement.

We also allow for the fact that maximizing revenue is not enough. If you maximize your revenue now but bump too many passengers, you could find demand for your services decreasing. You could be forced to charge a lower price, and your revenue might decrease in the long run. We offer you a way to establish a trade-off between revenue and percentage of flights with bumped passengers. You tell us how important it is to you to have few bumped flights, and we can tell you how many passengers to book.

Even using three different optimization strategies to account for the effects of fluctuating demand, we find that optimal values fall in a very narrow range. For a 100-seat plane, this range is 104 to 108.

We also evaluated the effect of the September 11th crisis on the airline industry. Our model predicts that, with a decreased number of flights, you should decrease the level of overbooking. If security delays many passengers from reaching their flights on time, you should also decrease the number of bookings. Increased passenger fear will decrease the probability that passengers show up for their flights, so in this case you should increase your booking number.

We have given you only a taste of what our model can do. We hope you will agree that contracting for our services will be of the highest benefit to your esteemed company.

Sincerely,

Models R Us

Author-Judge's Commentary: The Outstanding Airline Overbooking Papers

William P. Fox
Dept. of Mathematics
Francis Marion University
Florence, SC 29501
bfox@fmarion.edu

Introduction

Once again, Problem B proved to be a bigger challenge than originally considered, both for the students and the judges.

The students had a wealth of information for the basic model from the Web and from other resources. Students could consider and refine the basic information to fit the proposed post-9-11 scenario.

The judges had to read and evaluate many diverse (yet sometimes similar) approaches in order to find the "best" papers. Judges found mistakes—errors in modeling, assumptions, mathematics, and/or analysis—even in these "best" papers; so it is important to note that "best" does not mean perfect. The judges must read and apply their own subjective analysis to evaluate critically both the technical and expository solutions presented by the teams.

No paper analyzed every element nor applied critical validation and sensitivity analysis to all aspects of their model. Judges found many papers *with the exact same model (down to the exact same letters used for the variables) and none of these clearly cited the universal source anywhere in the submission*. The failure to properly credit the original source critically hurt these papers; it was obvious their basic model was not theirs but came from a published source.

Advice

At the conclusion of the judging, the judges offered the following comments:

The UMAP Journal 23 (3) (2002) 365–370.

- Follow the instructions
  - Clearly answer all parts.
  - List all assumptions that affect the model and justify your use of those assumptions.
  - Make sure that your conclusions and results are clearly stated.
  - In the summary, put the "bottom line and managerial recommendation" results—not a chronological description of what you did.
  - Restate the problem in your words.
- A CEO memorandum
  - Be succinct.
  - Include "bottom line and managerial results" answers.
  - Do not include methods used or equations.
- Clarity and Style
  - Use a clear style and do not ramble.
  - A table of contents is very helpful to the judges.
  - Pictures, tables, and graphs are helpful; but you must explain them clearly.
  - Do not include a picture, table, or graph that is extraneous to your model or analysis.
  - Do not be verbose, since judges have only limited time to read and evaluate your paper.
- The Model
  - Develop your model—do not just provide a laundry list of possible models.
  - Start with a simple model and then refine it.
- Computer Programs
  - If a program is included, clearly define all parameters.
  - Always include an algorithm in the body of the paper for any code used.
  - If running a Monte Carlo simulation, be sure to run it enough times to have a statistically significant output.
- Validation
  - Check your model against some known baseline.
  - Check sensitivity of parameters to your results.
  - Check to see if your recommendation/conclusions make common sense.

  - Use real data.
  - The model should represent human behavior and be plausible.

- Resources
  - All work needs to be original or referenced; a reference list at the end is not sufficient!
  - Teams can only use inanimate resources—no real people or people consulted over the Internet.
  - Surf the Web but document sites where obtained information is used.
  - This problem lent itself to a literature search, but few teams did one.

- Summary
  - This is the first piece of information read by a judge. It should be well written and contain the bottom-line answer or result.
  - This summary should motivate the judge to read your paper to see how you obtained your results.

Judging

The judging is accomplished in two phases. Phase I, at a different site, is "triage judging." These are generally only 10-minute reads with a subjective scoring from 1 (worst) to 7 (best). Approximately the top 50% of papers are sent on the final judging.

Phase II is done with different judges and consists of a calibration round and another subjection round based on the 1–7 scoring system. Then the judges collaborate to develop a 100-point scale to enable them to "bubble up" the better papers. Four or more longer rounds are accomplished using this scale, followed by a lengthy discussion of the last final group of papers.

Reflections of Triage

- Lots of good papers made it to the final judging.
- The initial summary made a significant difference in the papers (results versus an explanation).
- Report to the CEO also made a significant difference in papers.

Triage and Final Judges' Pet Peeves

- Tables with columns headed with Greek letters or acronyms that are not immediately understood.
- Definition and variable lists that are imbedded in a paragraph.
- Equations used without explaining terms and what the equation accomplished.
- Copying derivations from other sources; cite the reference and briefly explain is a better approach.

Approaches by the Outstanding Papers

Six papers were selected as Outstanding submissions because they:

- developed a workable, realistic model from their assumptions that could have been used to answer all elements;
- made clear recommendations;
- wrote a clear and understandable paper describing the problem, their model, and results; and
- handled all the elements.

The required elements, as viewed by the judges, were to

- develop a basic overbooking model that enabled one to find optimal values,
- consider alternative strategies for handling overbooked passengers,
- reflect on post-9-11 issues, and
- contain the CEO report of finding and analysis.

Most of the better papers did an extensive literature and Web search concerning overbooking by airlines and used this information in their model building.

> *The poorest section in all papers, including many of the Outstanding papers, was the section on assumptions with rational justification.*

Many papers just skipped this section and went directly from the problem to model-building!

Most papers used a stochastic approach for their model. With interarrival times assumed to be exponential, a Poisson process was often used to model passengers. Teams moved quickly from the Poisson to a binomial distribution with p and $1 - p$ representing "shows" and "no-shows" for ticket-holders.

Many teams started directly with the binomial distribution without loss of continuity. Some teams went on to use the normal approximation to the binomial. Revenues were generally calculated using some sort of "expected value" equation. Some teams built nonlinear optimization models, which was a nice and different approach.

Teams usually started with a simple example: a single plane with a fixed cost and capacity, one ticket price, and a reasonable value for no-shows based on historical data. This then became a model from which teams could build refinements (not only to their parameters) but also to include the changes based on post-9-11.

Teams often simulated these results using the computer and then made sense of the simulation by summarizing the results.

Wake Forest had two Outstanding papers. Team 69, with their paper entitled "ACE is High," was the INFORMS winner because of its superior analysis. Both papers began using a binomial approach as their base model. Team 273 developed a single-plane model, a two-plane model, and generalized to an n-plane model. Team 69 did a superb job in maximizing revenue after examining alternatives and varying their parameters.

The Harvey Mudd team, the MAA winner, had—by far—the best literature search. They used it to discuss existing models to determine if any could be used for post-9-11. Their research examined many of the current overbooking models that could be adapted to the situation.

The University of Colorado team used Frontier Airlines as their airlines. They began with the binomial random variable approach, with revenues being expected values. They modeled both linear and nonlinear compensation plans for bumped passengers. They developed an auction-style model using Chebyshev's weighting distribution. They also consider time-dependency in their model.

The Duke University team, the SIAM winner, had an excellent mix of literature search material and development of their own models. They too began with a basic binomial model. They considered multiple fares and related each post-9-11 issue to parameters in their model. They varied their parameters and provided many key insights to the overbooking problem. This paper was the first paper in many years to receive an Outstanding rating from each judge who read the paper.

The Bethel College team built a risk assessment model. They used a normal distribution as their probability distribution and then put together an expected value model for revenue. Their analysis of Vanguard Airlines with a plane capacity of 130 passengers was done well.

Most papers found an "optimal" overbooking strategy to be to overbook between 9% and 15%, and they used these numbers to find "optimal" revenues for the airlines. Many teams tried alternative strategies for compensation, and some even considered the different classes of seats on an airplane.

All teams and their advisors are commended for the efforts on the Airline Overbooking Problem.

About the Author

Dr. William P. Fox is Professor and the Chair of the Department of Mathematics at Francis Marion University. He received his M.S. in operations research from the Naval Postgraduate School and his Ph.D. in operations research and industrial engineering from Clemson University. Prior to coming to Francis Marion, he was a professor in the Department of Mathematical Sciences at the United States Military Academy. He has co-authored several mathematical modeling textbooks and makes numerous conference presentations on mathematical modeling. He is a SIAM lecturer. He is currently the director of the High School Mathematical Contest in Modeling (HiMCM). He was a co-author of this year's airline overbooking problem.

Statement of Ownership, Management, and Circulation

1. Publication Title	2. Publication Number	3. Filing Date
The UMAP Journal	0 1 9 7 _ 3 6 2 2	9/30/2002
4. Issue Frequency	**5. Number of Issues Published Annually**	**6. Annual Subscription Price**
quarterly	4 + annual collection	$75.00

7. Complete Mailing Address of Known Office of Publication *(Not printer) (Street, city, county, state, and ZIP+4)*

57 Bedford Street, Suite 210, Lexington, MA 02420

Contact Person: Kevin Darcy

Telephone: 781-862-7878 x.31

8. Complete Mailing Address of Headquarters or General Business Office of Publisher *(Not printer)*

same

9. Full Names and Complete Mailing Addresses of Publisher, Editor, and Managing Editor *(Do not leave blank)*

Publisher *(Name and complete mailing address)*

Solomon Garfunkel, 57 Bedford Street, Suite 210, Lexington, MA 02420

Editor *(Name and complete mailing address)*

Paul J Campbell, Beloit College, 700 College Street, Beloit, WI 53511

Managing Editor *(Name and complete mailing address)*

Pauline Wright, 57 Bedford Street, Suite 210, Lexington, MA 02420

10. Owner *(Do not leave blank. If the publication is owned by a corporation, give the name and address of the corporation immediately followed by the names and addresses of all stockholders owning or holding 1 percent or more of the total amount of stock. If not owned by a corporation, give the names and addresses of the individual owners. If owned by a partnership or other unincorporated firm, give its name and address as well as those of each individual owner. If the publication is published by a nonprofit organization, give its name and address.)*

Full Name	Complete Mailing Address
The Consortium for Mathematics and its Applications	57 Bedford Street
COMAP, Inc.	Suite 210
	Lexington, MA 02420

11. Known Bondholders, Mortgagees, and Other Security Holders Owning or Holding 1 Percent or More of Total Amount of Bonds, Mortgages, or Other Securities. If none, check box ☐ None

Full Name	Complete Mailing Address

12. Tax Status *(For completion by nonprofit organizations authorized to mail at nonprofit rates) (Check one)*

The purpose, function, and nonprofit status of this organization and the exempt status for federal income tax purposes:

☒ Has Not Changed During Preceding 12 Months

☐ Has Changed During Preceding 12 Months *(Publisher must submit explanation of change with this statement)*

13. Publication Title UMAP Journal			14. Issue Date for Circulation Data Below 11/15/2002	
15.		Extent and Nature of Circulation	Average No. Copies Each Issue During Preceding 12 Months	No. Copies of Single Issue Published Nearest to Filing Date
a. Total Number of Copies *(Net press run)*			1340	1300
b. Paid and/or Requested Circulation	(1)	Paid/Requested Outside-County Mail Subscriptions Stated on Form 3541. *(Include advertiser's proof and exchange copies)*	1040	1080
	(2)	Paid In-County Subscriptions *(Include advertiser's proof and exchange copies)*	0	0
	(3)	Sales Through Dealers and Carriers, Street Vendors, Counter Sales, and Other Non-USPS Paid Distribution	75	70
	(4)	Other Classes Mailed Through the USPS	0	0
c. Total Paid and/or Requested Circulation *[Sum of 15b. (1), (2),(3),and (4)]*		▶	1115	1150
d. Free Distribution by Mail *(Samples, compliment ary, and other free)*	(1)	Outside-County as Stated on Form 3541		
	(2)	In-County as Stated on Form 3541		
	(3)	Other Classes Mailed Through the USPS	80	60
e. Free Distribution Outside the Mail *(Carriers or other means)*			0	0
f. Total Free Distribution *(Sum of 15d. and 15e.)*		▶	80	60
g. Total Distribution *(Sum of 15c. and 15f)*		▶	1195	1210
h. Copies not Distributed			145	90
i. Total *(Sum of 15g. and h.)*		▶	1340	1300
j. Percent Paid and/or Requested Circulation *(15c. divided by 15g. times 100)*			93	95

16. Publication of Statement of Ownership

☒ Publication required. Will be printed in the third issue of this publication. ☐ Publication not required.

17. Signature and Title of Editor, Publisher, Business Manager, or Owner [signature] Date 9/30/2002

I certify that all information furnished on this form is true and complete. I understand that anyone who furnishes false or misleading information on this form or who omits material or information requested on the form may be subject to criminal sanctions (including fines and imprisonment) and/or civil sanctions (including civil penalties).

Instructions to Publishers

1. Complete and file one copy of this form with your postmaster annually on or before October 1. Keep a copy of the completed form for your records.
2. In cases where the stockholder or security holder is a trustee, include in items 10 and 11 the name of the person or corporation for whom the trustee is acting. Also include the names and addresses of individuals who are stockholders who own or hold 1 percent or more of the total amount of bonds, mortgages, or other securities of the publishing corporation. In item 11, if none, check the box. Use blank sheets if more space is required.
3. Be sure to furnish all circulation information called for in item 15. Free circulation must be shown in items 15d, e, and f.
4. Item 15h., Copies not Distributed, must include (1) newsstand copies originally stated on Form 3541, and returned to the publisher, (2) estimated returns from news agents, and (3), copies for office use, leftovers, spoiled, and all other copies not distributed.
5. If the publication had Periodicals authorization as a general or requester publication, this Statement of Ownership, Management, and Circulation must be published; it must be printed in any issue in October or, if the publication is not published during October, the first issue printed after October.
6. In item 16, indicate the date of the issue in which this Statement of Ownership will be published.
7. Item 17 must be signed.

Failure to file or publish a statement of ownership may lead to suspension of Periodocals authorization.