

Using Prediction to Facilitate Patient Flow in a Health Care Delivery Chain

by

Jordan Shefer Peck

S.M., Technology and Policy, Massachusetts Institute of Technology, 2008

B.A., Physics, Binghamton University, 2006

B.S, Mechanical Engineering, Binghamton University, 2006

Submitted to the Engineering Systems Division
in Partial Fulfillment of the Requirements for the degree of

Doctor of Philosophy
at the
Massachusetts Institute of Technology

February 2013

© 2013 Massachusetts Institute of Technology. All rights reserved.

The author hereby grants to MIT permission to reproduce and to distribute publicly paper and electronic copies of this thesis document in whole or in part in any medium now known or hereafter created.

Signature of Author
Engineering Systems Division
December 31, 2012

Certified by.....
Deborah J. Nightingale
Professor of Practice, Aeronautics and Astronautics and Engineering Systems
Thesis Supervisor

Certified by.....
Stephen C. Graves
Abraham J. Siegel Professor of Management Science Professor of
Engineering Systems and Mechanical Engineering

Certified by.....
James C. Benneyan
Professor of Mechanical, Industrial & Manufacturing Engineering
Northeastern University

Accepted by.....
Oliver L. DeWeck
Associate Professor of Engineering Systems and Aeronautics and Astronautics
Chair, Engineering Systems Division Education Committee

This dissertation is dedicated to George Ginsberg and Abraham Peck, patrons of the Veterans Health Administration and beloved role models for hard work and perseverance in the face of adversity.

Using Prediction to Facilitate Patient Flow in a Health Care Delivery Chain

by
Jordan Shefer Peck

Submitted to the Engineering Systems Division on
December 31, 2012 in Partial Fulfillment of the Requirements for the Degree of
Doctor of Philosophy in Engineering Systems: Technology Management and Policy

Abstract

A health care delivery chain is a series of treatment steps through which patients flow. The Emergency Department (ED)/Inpatient Unit (IU) chain is an example chain, common to many hospitals. Recent literature has suggested that predictions of IU admission, when patients enter the ED, could be used to initiate IU bed preparations before the patient has completed emergency treatment and improve flow through the chain. This dissertation explores the merit and implications of this suggestion.

Using retrospective data collected at the ED of the Veterans Health Administration Boston Health Care System (VHA BHS), three methods are selected for making admission predictions: expert opinion, naïve Bayes conditional probability and linear regression with a logit link function (logit-linear regression). The logit-linear regression is found to perform best.

Databases of historic data are collected from four hospitals including VHA BHS. Logit-linear regression prediction models generated for each individual hospital perform well based on multiple measures. The prediction model generated for the VHA BHS hospital continues to perform well when predictive data are collected and coded prospectively by nurses.

For two weeks, predictions are made on each patient that enters the VHA BHS ED. This data is then summarized and displayed on the VHA BHS internet homepage. No change was observed in key ED flow measures; however, interviews with hospital staff exposed ways in which the prediction information was valuable: planning individual patient admissions, personal scheduling, resource scheduling, resource alignment, and hospital network coordination.

A discrete event simulation of the system shows that if IU staff emphasizes discharge before noon, flow measures improve as compared to a baseline scenario where discharge priority begins at 1pm. Sharing ED crowding or prediction information leads to best patient flow performance when using specific schedules dictating IU response to the information.

This dissertation targets the practical and theoretical implications of using prediction to improve flow through the ED/IU health care delivery chain. It is suggested that the results will have impact on many other levels of health care delivery that share the delivery chain structure.

Thesis Supervisor: Deborah Nightingale

Title: Professor of Practice, Aeronautics and Astronautics and Engineering Systems

Acknowledgements

This dissertation is the culmination of over 23 years of continuous education. During that time I have received the assistance and advice of countless individuals. I cannot possibly acknowledge all of the people that have contributed to this final achievement; nonetheless, I have some in particular that I would like to identify here.

My time at MIT has been enriched by great professors, my committee in particular. Professor Deborah Nightingale has served as my advisor and committee chair. Professor Nightingale's unwavering support ensured that I had the knowledge and resources I needed to pursue my research without hesitation or financial delay. Professor Nightingale's positive attitude and inquisitive spirit is contagious and I am sincerely grateful for her advice on life and academics. I would also like to thank Professor Stephen Graves who had the uncanny ability to change my entire perspective through a single conversation; it was a great privilege to learn from him. Finally, I would like to thank Professor James Benneyan. Professor Benneyan's excellence as a health systems researcher and educator continues to be recognized across the industry and I am truly grateful for his mentorship. I would also like to recognize Professor Sang-Gook Kim, who was not a member of my committee but provided valuable advice and support throughout my time at MIT.

In general I would like to thank the teachers I have had from grade school until now, who did not give up on me, despite my "poor spelling." Their patience with my learning disabilities allowed me to overcome these personal difficulties and this achievement is a testament to the heroics they perform every day as educators.

While not always directly contributing to my research, I would like to acknowledge the staff and students of the MIT Lean Advancement Initiative for their administrative, moral, and academic support: Sarah Benson, Paula Buick, Dr. John Carroll, Clinton Campbell, Dr. Erkan Ceyhan, Adam Glick, Nicolene Hengen, Dick Lewis, Prof. Earll Murman, Juliet Perdichizzi, Mark Prendergast, Dr. Donna Rhodes, Dr. Jayakanth Srinivasan, Prof. Joseph Sussman, Stephanie Toews-Moeling, and many others.

For continued moral and intellectual support directly related to my research, as well as to my life, I would like to recognize the members of my MIT health care and living community: Shahed Al-Haque, Julie Andren, Donna Denoncourt, Ruth Fishbein, Dr. Jorge Fradinho, Dr. Wiljeana Glover, Dr. Howard Heller, Annette Jacobs, Ozge Karanfil, Dr. William Kettyle, Maryanne Kirkbride, Dr. Judy Maro, Dr. Robert Rudin, Ashley Wessendorf and the staff of MIT's Baker House. They made an otherwise isolating experience into one filled with fascinating conversations and joyful memories.

This research would not have been possible without the contributions of many individuals at the Veterans Health Administration. I would like to thank all of the staff and engineers of the Veterans Engineering Resource Center, in particular Lynne Cannavo, Joan Clifford (and the patient flow committee), Janis Hersh, Dr. Anand Kartha, Dr. James Schlosser, Chiragi Shah, Dr. David Thornton (and his medical residents), and Stephanie Triplett.

The research presented in this dissertation relied heavily on the curiosity and devotion to improvement that characterized the staff members of VHA Boston Health Care System's Emergency Services. If all hospitals had such amazing and committed staff, the world would be a much better place. In particular I would like to thank Nurse Manager Jack Marinello for his leadership and flexibility. A special thanks to the Director of Emergency Services, Dr. Stephan Gaehde. Dr. Gaehde was an integral part of this research; his contributions spanned all conceivable aspects of the project. Dr. Gaehde also served as an eternally patient advisor and mentor for my healthcare education; I am truly enriched by working with him. Thanks also to my primary contacts at the other emergency departments included in this study: Dr. Eric Dickson, Dr. David Gelman, Dr. David Huckins, Dr. Mark Lemons and Virginia Mangolds.

In closing I acknowledge the members of my family. My older brother, Ron Peck, and his wife, Corinne, have eased the difficulties of graduate school through their generosity and uplifting personalities. It has been a true blessing to have them nearby. My younger brother, Ethan Peck, is working towards his own PhD and has been a continual source of inspiration. I am unbelievably proud of both of my siblings; together we define the concept of brotherhood in all of its meanings.

I cannot even begin to thank my parents. A common hypothetical question is "what would you attempt if you knew you would not fail," the unwavering support provided by my parents has allowed me to answer that question. They believed in me when I did not, they fought for me when others would not, they listened to me when the world seemed deaf, they taught me when I did not understand, they led by example when words were insufficient and they consistently put their children before themselves. My father's unending pursuit of knowledge, fierce dedication, and stubborn passion inspired me, from an early age, to follow my dreams and seek to fulfill my academic and life potential. My mother's grace under pressure, humble wisdom, and sense of humor taught me how to get through the most difficult of times and emerge, smiling. All children should be so lucky.

Finally I would like to recognize my wife, Emily Peck, who has been with me since I began this academic journey. When I lost the strength to continue my work, she restored it, when I wavered from my commitment, she was my anchor, and when I lost track of my goals, she gave me purpose. I am filled with excitement for the future that we will build together.

Table of Contents

Abstract.....	3
Acknowledgements.....	4
List of Figures.....	8
List of Tables.....	10
List of Acronyms.....	11
Chapter 1: An Introduction to Health Care Delivery Chains.....	13
1.1 The health care “problem” and selected top down solutions.....	13
1.2 Selected methods for hospital performance improvement.....	16
1.2.1 Electronic Health Records/Health Information Technology.....	17
1.2.2 Decision Support Tools.....	19
1.2.3 Human Factors.....	20
1.2.4 Health Care Supply Chains/Operations Management.....	21
1.2.5 Patient Flow.....	23
1.2.6 Organizational Change and Systems Re-design.....	24
1.3 Pulling methods together through frameworks.....	24
1.3.1 Continuous improvement frameworks.....	24
1.3.2 Looking at health care delivery systems through the lenses of Enterprise Architecture.....	25
1.4 Defining health care delivery chains.....	31
1.5 Dissertation goals and outline.....	35
Chapter 2: The Emergency Department and Inpatient Unit Delivery Chain.....	38
2.1 Background and Motivation for Emergency Department Improvement.....	38
2.1.1 The role of the Emergency Department.....	38
2.1.2 The Emergency Department System.....	40
2.1.3 Measuring the quality of care in the emergency department.....	43
2.2 Past Solutions.....	48
2.2.1 Input.....	48
2.2.2 Throughput.....	50
2.2.3 Output.....	54
2.3 Prediction.....	58
2.3.1 Prediction in Health Care.....	59
2.3.2 Sample prediction based studies in health care.....	62
2.3.3 Prediction in the Emergency Department.....	64
Chapter 3: Predicting Emergency Department Admissions.....	68
3.1 Methods.....	69
3.1.1 Study Design.....	69
3.1.2 Study Setting.....	69
3.1.3 Study Protocol.....	71
3.1.4 Data Analysis/Measures.....	76
3.2 Results.....	78
3.3 Discussion.....	84
3.4 Limitations.....	85
3.5 Conclusions.....	86
Chapter 4: Generalizability of the Emergency Department Prediction Model.....	88

4.1	Methods.....	89
4.1.1	Study Design.....	89
4.1.2	Study Setting.....	89
4.1.3	Study Protocol.....	90
4.1.4	Data Analysis/Measures.....	94
4.2	Results.....	94
4.3	Discussion.....	98
4.4	Limitations.....	101
4.5	Conclusions.....	102
Chapter 5: Implementation of the Emergency Department Prediction Model.....		104
5.1	Methods.....	106
5.1.1	Study Design.....	106
5.1.2	Study Protocol.....	106
5.2	Results.....	111
5.2.1	Admission Planning.....	115
5.2.2	Resource Scheduling.....	116
5.2.3	Personal Scheduling.....	118
5.2.4	Resource Alignment.....	119
5.2.5	Hospital Network Management.....	120
5.3	Discussion.....	121
5.4	Limitations.....	125
5.5	Conclusions.....	125
Chapter 6: Simulation of the Emergency Department Prediction Model.....		127
6.1	A conceptual model of the ED/IU “Pull” System.....	127
6.2	Methods.....	131
6.2.1	Simulation Scenarios.....	134
6.2.2	Calibration and Validation.....	136
6.3	Results.....	141
6.4	Discussion.....	149
6.5	Limitations/Conclusions.....	151
Chapter 7: Conclusions, Contributions, and Future Work.....		153
7.1	Conclusions and contributions of studies.....	153
7.1.1	Question 1.....	154
7.1.2	Question 2.....	155
7.1.3	Question 3.....	156
7.2	Future Work on the ED/IU Health Care Delivery Chain.....	158
7.2.1	Question 1.....	158
7.2.2	Question 2.....	158
7.2.3	Question 3.....	159
7.3	Health Care Delivery Chains.....	160
7.3.1	Contribution to/Invention of Health Care Delivery Chain Management.....	160
7.3.2	Next Steps.....	162
Chapter 8: References.....		163
Appendix A: Probability values for admission prediction models – VHA West Roxbury.....		179
Appendix B: Probability values for admission prediction models - Four sample hospitals.....		183

List of Figures

Figure 1 National health expenditures per capita, 1960-2010	13
Figure 2 Percentage of people postponing or foregoing care due to cost.....	14
Figure 3 Personal health care expenditures per capita by service type.....	16
Figure 4 Percentage of hospitals and office based physicians that have adopted even basic EHR systems.....	18
Figure 5 Health care development and supply chain	21
Figure 6 Map of EA framework views	29
Figure 7 Health care supply chain with multiple treatment steps.....	32
Figure 8 Health care delivery life cycle chain	33
Figure 9 ED visits vs. Operating EDs.....	39
Figure 10 Simple health care delivery chain.....	40
Figure 11 ED/IU health care delivery chain	40
Figure 12 General patient flow model through ED/IU system.....	41
Figure 13 Emergency department in-department level delivery chain.....	41
Figure 14 ED/IU delivery chain with time element.....	43
Figure 15 Growing waiting time to see an ED physician	47
Figure 16 Timeline conception of using prediction to reduce boarding time.....	57
Figure 17 Patient visits to VHA West Roxbury ED by year	70
Figure 18 Monthly patient arrival pattern.....	71
Figure 19 Monthly average waiting time pattern.....	71
Figure 20 Expert opinion triage questionnaire.....	72
Figure 21 Conceptual illustration of real time bed demand forecast (Running expected number of admissions and standard deviation)	78
Figure 22 Categorized predictions of patient admissions versus percent of patients admitted from each category	80
Figure 23 Real time, statistically predicted-expected and actual number of cumulative admissions	81
Figure 24 Correlation between predicted admissions and actual admissions based on expert opinion (top), naïve Bayes (middle), and logit-linear regression (bottom) approaches	82

Figure 25 Prediction residuals (predicted minus actual) based on expert opinion (top), naïve Bayes (middle), and logit-linear regression (bottom) approaches	83
Figure 26 Comparison of predicted vs. actual admissions by probability decile	96
Figure 27 VHA 1 and VHA 2: daily actual vs. predicted admissions and residuals	97
Figure 28 Small and Large Public hospital: quarter daily actual vs. predicted admissions and residuals	98
Figure 29 Data entry form at patient arrival	107
Figure 30 Public display of prediction variables	109
Figure 31 Normalized predicted and current bed orders compared with normalized ED admits already on floors during trial period	112
Figure 32 Average (top) and Standard deviation of (bottom) ED waiting times by week	113
Figure 33 Average (top) and Standard deviation of (bottom) ED boarding times by week	114
Figure 34 a) Traditional pull system, b) current ED/IU push system, c) potential ED/IU information based pull system	130
Figure 35 Discrete event simulation model logic	131
Figure 36 Doctor Decision Cycle	134
Figure 37 Simulated and real VHA Boston daily patient arrival rates by hour	137
Figure 38 Simulated and real VHA Boston daily admission request rates by hour	138
Figure 39 Simulated and real VHA Boston daily, normalized, IU discharge rates by hour	139
Figure 40 Simulated and real VHA Boston patient IU LOS	139
Figure 41 Average IU Boarding time (and 95% confidence intervals) with shifting 3 hour discharge priority start times without an Case 1 (top), Case 2 (middle), Case 3 (bottom)	142
Figure 42 Optimized sensitivity schedule using Crowding index for Case 1 (top), Case 2 (middle), and Case 3 (bottom)	144
Figure 43 Optimized sensitivity schedule using Imperfect index for Case 1 (top), Case 2 (middle), and Case 3 (bottom)	145
Figure 44 Optimized sensitivity schedule using Perfect Index for Case 1 (top), Case 2 (middle), and Case 3 (bottom)	146
Figure 45 IU wait time and ED wait time for optimized scenarios and time based scenarios ..	147
Figure 46 ED/IU health care delivery chain	160
Figure 47 Generic health care delivery chain	160

List of Tables

Table 1 Basic patient characteristics between development, validation and test datasets.....	73
Table 2 Factors tested for admission prediction ability and the empirical probabilities of occurrence.....	74
Table 3 Model parameters for best fitting logit-linear regression	80
Table 4 Characteristics of participating hospitals based on development datasets	90
Table 5 List of factors collected at triage.....	91
Table 6 Dataset attributes.....	91
Table 7 Example admission probabilities given selected factors	93
Table 8 Chosen models and coefficients for each hospital.....	95
Table 9 Quality of predictive models when applied to test sets	95
Table 10 Historical average and standard deviation values of prediction variables provided as reference for hospital staff during implementation period	109
Table 11 Pay off matrix of prioritizing discharge versus treatment	124
Table 12 Simulation output data vs. VHA Boston monthly average data for validation	140
Table 13 Difference in average IU boarding times between scenarios: Case 1.....	148
Table 14 Difference in average IU boarding times between scenarios: Case 2.....	148
Table 15 Difference in average IU boarding times between scenarios: Case 3.....	148

List of Acronyms

ACO: Accountable Care Organization
ADI: Advanced demand information
ARIMA: Autoregressive moving average time series analysis
AUC: Area under the Receiver Operating Characteristic Curve
CDSS: Clinical Decision Support Systems
CDC: Center for Disease Control and Prevention
CMS: Centers for Medicare and Medicaid Services
CQI: Continuous Quality Improvement
DES: Discrete Event Simulations
EA: Enterprise Architecture
ED: Emergency Department
EHR: Electronic Health Record
EMS: Environmental Management Services
EMTALA: Emergency Medical Treatment and Active Labor Act
ESI: Emergency Severity Index
GAO: Government Accountability Office
GOF: Hosmer-Lemeshow goodness of fit
HCD: Health Care Delivery System
HIT: Health Information Technology
HMSS: Health Care Management Support Systems
IOM: Institute of Medicine
ITO: Input/Throughput/Output
IU: Inpatient Unit
KMS: Knowledge Management Systems
LOS: Length of Stay

LWBS: leave/left without being seen

NVA: Non-Value Added

PPACA: Patient Protection and Affordable Care Act

TQM: Total Quality Management

USA: United States of America

VHA: Veterans Health Administration

VHA 1: VHA West Roxbury

VHA 2: VHA Medical Center 2

VHA BHS: VHA Boston Health Care System

VHA West Roxbury: VHA hospital in West Roxbury, Massachusetts, part of the VHA Boston Health care System

Chapter 1: An Introduction to Health Care Delivery Chains

In this chapter, the concept of health care delivery chains will be introduced. The chapter will begin by discussing the impetus for improving how care is delivered in the US health care system. There will be a discussion of the ratio between cost and quality of health care and how others have attempted to improve this ratio.

The generalized discussion of the health care system will lead into the specific approach taken in this dissertation, namely, a health care delivery chain framework to improve the provision of care. Finally, there will be a generalized description of the primary research questions addressed by this dissertation and an outline will be provided that explains how the chapters of the dissertation seek to address these research questions.

1.1 The health care “problem” and selected top down solutions

Providing health care to growing populations has become one of the world’s most prominent issues. One major issue associated with health care is its rising cost. Data has shown that health expenditures per capita continue to rise (Figure 1). While national costs continue to rise there is little evidence that increasing costs correlate to increased quality [Feldstein 2003, Fisher et al. 2009].

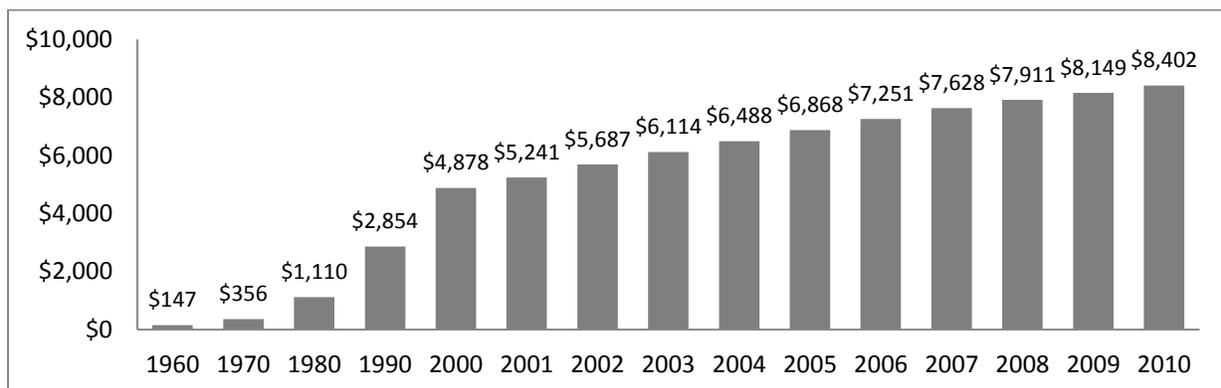
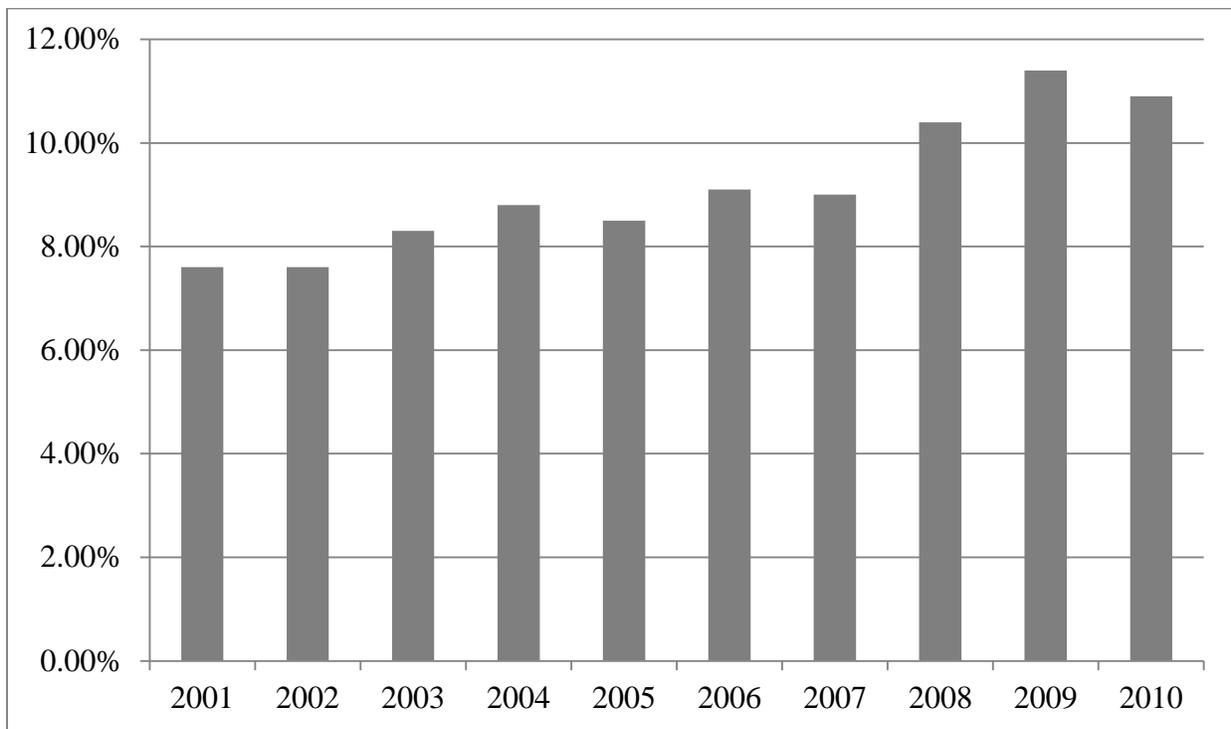


Figure 1 National health expenditures per capita, 1960-2010 [KFF 2012]

There have been many top down attempts to improve the cost/quality ratio of the health care system. One example has been making government level policy changes to affect how insurance,

providers, and customers interact and are permitted to work. This is most recently exemplified by the Patient Protection and Affordable Care Act (PPACA) [US Senate 2010]. This act included many provisions to increase the number of US residents with insurance coverage, in order to reduce the number of patients who postpone or forego care due to cost. These patients often seek care only in emergencies, leading to higher cost treatments for issues that may have been addressed for less money earlier on. Figure 2 shows how the number of patients postponing care has changed over time based on a survey by the US Department of Health & Human Services. This insurance provision is simply one part of the act which includes many other provisions seeking to improve care by exercising more control over aspects of the health care market.



**Figure 2 Percentage of people postponing or foregoing care due to cost
[Health system Measurement Project 2012b]**

Another attempt to improve the cost/quality ratio is to change how providers are reimbursed. It is possible for public and private insurance agencies to use incentives to affect provider decisions. For example, in 2007 the Centers for Medicare and Medicaid Services (CMS) began an initiative to reduce the occurrence of “never events.” These events were a list of injuries or infections that a patient would receive while being cared for at a hospital. It was found that reimbursing for the treatment of never events was costing CMS millions of dollars per year. In order to reduce this

number, a policy was created that treatment for such events would not be reimbursed and therefore the treatment would be at cost for the hospital [Zhan et al. 2006, Pear 2007].

Another way to reduce costs and improve quality is through organizational change and data sharing. An example of this was the creation of Accountable Care Organizations (ACOs) by CMS. These organizations are specifically targeted at improving the quality of care by enhancing the connectedness of health care providers through increased communication, data sharing, and joint decision making. [National Public Radio 2011, CMS 2012a]

The above are just a few top down options that have been explored for improving the cost/quality ratio of health care. There are in fact many examples, which have been segmented into three categories: provider change, payment change, and market change [Lee and Mongan 2009]. However when looking at the initiatives described above, it is clear that the focus has been placed on providers. Whether it is controlling provider behavior, enabling the provider to make better decisions, or influencing how providers communicate with one another. The reason for this is clear when looking at Figure 3 which shows personal health care expenditures per capita by service type [Health System Measurement Project 2012c]. As the figure shows a large portion of spending is on hospital care and other areas of direct patient provider interaction. The implication of this, is that the areas where health care delivery takes place are reasonable targets for high impact attempts to improve the cost/quality ratio.

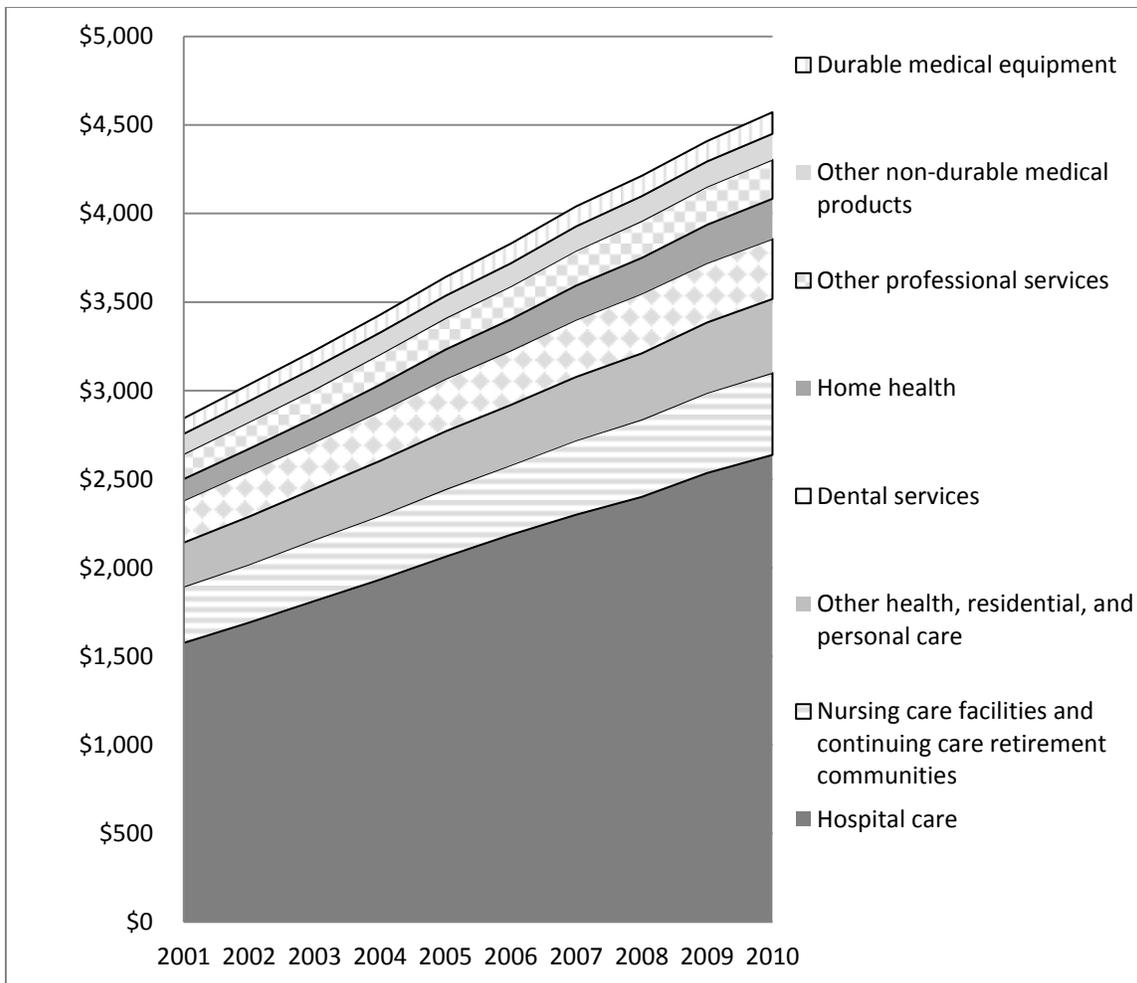


Figure 3 Personal health care expenditures per capita by service type
 [Health System Measurement Project 2012c]

1.2 Selected methods for hospital performance improvement

The above discussion has established that areas of health care delivery, particularly hospitals, are high impact targets for improving the cost/quality ratio of health care systems. This message was very strongly enforced by the Institute of Medicine (IOM) report “Crossing the Quality Chasm” [IOM 2001]. In the IOM report general principles were proposed as guidelines for creating a better system, however it did not actually introduce many tools or specific solutions for achieving improvement. With that said, there are many methods that have been employed. Discussing all of these methods is beyond the scope of this dissertation; however this section introduces some of the more common/popular tools. The tools and methods employed in future

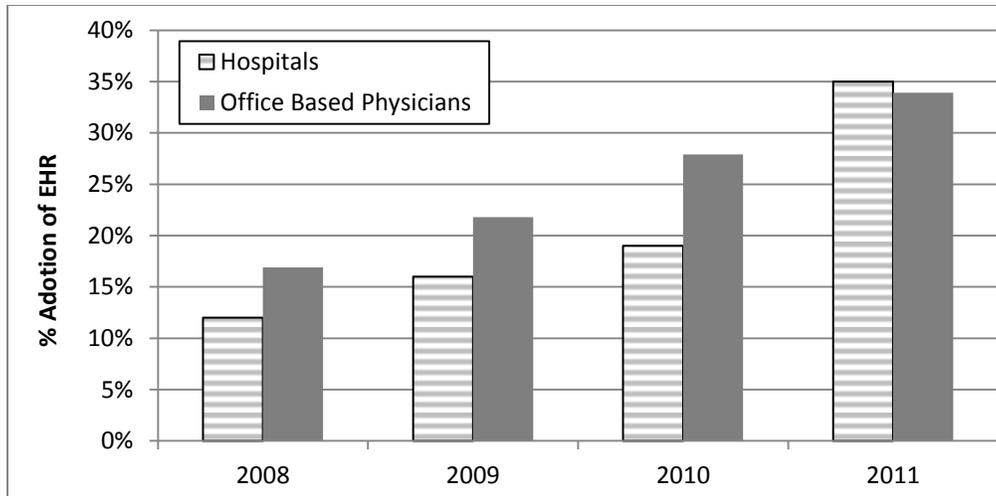
chapters build upon elements of many of these solutions. In order to enable the discussion the solutions are categorized:

- Electronic Health Records and Health Information Technology,
- Decision Support Tools,
- Human Factors,
- Health Care Supply Chains/Operations Management,
- Patient Flow,
- Organizational Change and Systems Re-design,
- Improvement Frameworks.

The next subsections will discuss each category, however it is worth noting that there is some overlap between them and therefore the categorical lines may not be as clear-cut as this treatment suggests.

1.2.1 Electronic Health Records/Health Information Technology

Electronic health records (EHR) have become one of largest areas for development, study, and industry growth in the health system. Despite the amazing growth in the EHR market, Figure 4 shows that only a minority of office based physicians and hospitals have adopted even basic Electronic Health Record systems [Jha et al. 2009, Health System Measurement Project 2012a]. Nevertheless, providers that have adopted EHR systems have claimed to receive many benefits including enhanced: access to patient records, patient communication, pharmaceutical prescription and medicine reconciliation, streamlined claims processing, reduced operational time and fiscal costs, and ability to spend more time at the patients bedside [Gans et al. 2005, Poissant et al. 2005].



**Figure 4 Percentage of hospitals and office based physicians that have adopted even basic EHR systems
[Health system Measurement project 2012a]**

Section 1.1 discussed CMS’s negative incentive towards improving care by not paying for “never events.” CMS has also used positive incentives. An example of this is the EHR Incentive Programs.

“The Medicare and Medicaid EHR Incentive Programs provide incentive payments to eligible professionals, eligible hospitals and critical access hospitals as they adopt, implement, upgrade or demonstrate meaningful use of certified EHR technology. Eligible professionals can receive up to \$44,000 through the Medicare EHR Incentive Program and up to \$63,750 through the Medicaid EHR Incentive Program.” [CMS 2012b]

The goal of this program is to encourage the spread of EHRs around the country and capture the quality improvements and financial savings that can come from the implementation of an EHR system.

While EHR is often implemented by itself, EHR is merely a way of storing patient data on a computer rather than in paper files. The use of information technology in health care is more extensive than that. These other uses, and EHR, fall under the more general term of Health Information Technology (HIT). The adoption of HIT has also been encouraged by government initiatives, such as the American Recovery and Reinvestment Act of 2009 [Blumenthal 2009]. HIT has the ability to go beyond recording EHRs, it can dynamically monitor patients, prescribe

medications, detect potential medical errors, act as a platform for decision support tools, enable tele-health, and has many other uses [Chaudhry 2006].

1.2.2 Decision Support Tools

The treatment of patients is often seen as a combination of science and art. The health care system depends upon practitioners who make diagnoses often using uncertain data and relying on the knowledge gained by experience and education [Gawande 2002]. As medical research and knowledge continue to grow and medical professions continue to specialize it is becoming more difficult for practitioners to keep track of all there is to know about medicine. At the same time, this information is being shared online with patients on a growing basis.

In order to enable medical professionals to access state of the art data and best treatment practices for a disease, decision support tools are becoming increasingly popular. The types of decision support tools that have been discussed in health care are clinical decision support systems (CDSS) and knowledge management systems (KMS). Often these systems are defined using electronic terms, however in a more general sense they do not necessarily have to be built into an electronic system (thus making the decision support tools distinct from HIT).

CDSS can be defined as any system “designed to aid directly in clinical decision-making, in which characteristics of individual patients are used to generate patient-specific assessments or recommendations that are then presented to clinicians for consideration”[AHRQ 2012]. KMS can be defined as “a tool that selectively provides information relevant to the characteristics or circumstances of a clinical situation but which requires human interpretation for direct application to a specific patient” [AHRQ 2012]. CDSSs and KMSs have been implemented during health care delivery using many different outlets and for many different purposes. Some purposes include preventative care (such as suggesting immunizations if a patient is travelling), diagnosis (given certain symptoms outputting a probable cause), treatment planning (suggesting drug dosages and schedules, alerts for potential drug interactions, suggesting dates of treatment steps), and cost reduction (suggesting alternative treatments that can be effective but less expensive) [Berber 2009]. Some outlets for decision support tools include being built into an HIT system, paper based tools, smart phone applications, web based applications and more [Pearson 1994, Hunt 1998, IMPROVE 2011, AHRQ 2012].

Decision support systems are used in many industries for various purposes. Recent studies on the use of decision support systems in health care have been primarily focused on those described above. Despite this, there have been some efforts towards developing tools that fit in a third kind of decision support system for health care, management support systems. There is no universal name for these systems, but here they will be called Health Care Management Support Systems (HMSS). HMSSs are defined in this paper as tools for facilitating operational decisions and aligning resources with the goal of achieving increased health care delivery quality in a timely and cost effective manner [Forgionne 1996]. Examples of HMSS are common though they may not be consistently categorized as HMSS. Some of these examples are resource scheduling tools, bed boards, and hospital bed assignment tools [Clerkin et al. 1995, Bose 2003, van Merode et al. 2004, van Essen et al. 2012].

1.2.3 Human Factors

Human factors, the understanding of interactions among humans and other elements of a system, are fundamental in health care. In health care delivery, humans are not only the employees; the health outcome of a human is also the product and that same human is the customer. It is therefore not surprising that human factors engineering is often seen as the first tool applied to improving hospital efficiency and quality. This application of human factors engineering is often attributed to the early 1900s work of Frank and Lillian Gilbreth who performed motion study techniques to reduce motion and improve efficiency of health care workers in the hospital setting [Smalley 1982].

While it is important to understand the limitations of the processes and tools employed in health care delivery, it is also necessary to understand how these factors interact with the people who use them and the people upon whom they are used [Vicente 2005]. Besides simply looking at how people act as agents of a process, each person in a health system derives a different type of value from the system. Understanding these values is the key to avoiding conflicting goals and outcomes based on perverse incentives [Kolker and Story 2012].

An example of the importance of human factors can be seen in the implementation of decision support tools. When implemented, these tools have shown effectiveness when carefully worked into the normal process flow of the hospital. This integration makes them more easily

remembered and accessed by the humans that would use them [AHRQ 2012]. A similar understanding of human factors must also be included in HMSS in order to ensure that the expected efficiency gains are not compromised by designs that make it impractical for the human users to access or act upon the recommendations of an HMSS.

Human factors have also been utilized when looking at variation in health care delivery quality. For instance, quality can be affected by clinician fatigue. One example is a case where blood sugar medication was provided to a patient who needed an anti-coagulant. The vials for both drugs were similar and it was late at night. As a result, the tired nurse accidentally confused the two medications and the patient died. Human factors engineering can suggest ways to improve the system by making the bottles different shapes, not storing them next to each other, having bright labels, etc. [Spear 2009]

1.2.4 Health Care Supply Chains/Operations Management

Adopting a previously developed graphic one can look at the development/supply chain process as shown in Figure 5. Adding the final step of “Treat Patient” makes Figure 5 into a depiction of a health care supply chain [Simchi-Levi et al. 2003].

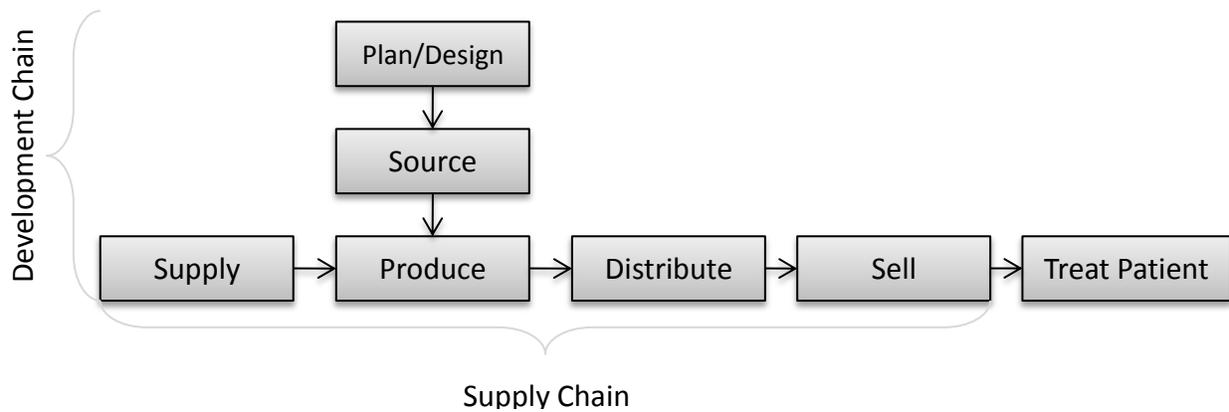


Figure 5 Health care development and supply chain

If a person works for a company that plans or designs, sources or produces products that end up being used in the treatment of a patient they may be considered a part of the health care product development chain. Likewise, one can be part of a company that makes supplies for producers of medical products, actually produces medical products, or distributes or sells medical products,

and people in these companies can be considered a part of the health care product supply chain. The need to manage the process depicted above to decrease health care costs has been noted [McKone-Sweet 2005]. Also related to this chain is the cost/benefit trade off of seeking new products and diagnostic machines to maximize quality, but at greater financial expense. As new products enter the market, hospitals compete to have the most up to date treatment methods. The cost of new products increases as old products are discarded in the name of improving quality [Feldstein 2003, Fuchs and Emanuel 2005]. It is possible to take an even wider view of the health care supply chain and consider the insurance providers who must be willing to pay for certain supplies or treatments for a patient. The complexity further increases when considering the large number of products that are used in any one patient engagement. This means that there are many development chain/supply chain combinations leading into a single patient treatment step.

Figure 5 is a simplified representation of a development/supply chain. In reality these chains are often networks of chains and whole fields of study are devoted to the management of these complex networks. The field of Operations Management focuses on how to best design these chains, where to place inventory and to invest resources in order to make these chains optimally efficient [Cachon and Terwiesch 2009]. In the health care context, such optimization becomes more complicated as hospitals often need to hold inventory for critical products.

Another way to take a wider perspective on health care supply chains is to go beyond the supply of products and instead consider the supply and scheduling of larger resources such as hospital beds, medical staff, and expensive diagnostic machines [Uzsoy 2005]. These resources are more equivalent to managing the production equipment and staff in a factory as opposed to the products that flow through them. This is still an important part of operations management, and is similarly important in health care/hospital operations management. Many tools, including HMSS, have been developed to improve hospital efficiency (the cost/quality ratio) by optimizing the scheduling of more expensive resources. Using the principles that have been employed in factory optimization and management has led to improvements in the operations of hospitals and other health care delivery systems. There have been applications of the theory of constraints towards improving general flow in the hospital, optimization techniques for scheduling staff and other resources, forecasting of demand to improve scheduling, the application of statistical

process control, the development of “takt” times for steps within health care treatment processes, as well as other methods and tools [Carey 2003, McLaughlin 2008, Kolker and Story 2012].

1.2.5 Patient Flow

The study of patient flow is focused on effecting how patients move through the system as if they are a product. When studying a health care delivery system it is possible to assume that the arrival rate and service rate of patients cannot be changed. With those assumptions in place a hospital manager can employ hospital supply chain and operations management techniques in order to designate resources to meet the need of those patients in the least expensive way possible. In contrast, the study of patient flow assumes that arrival and service rates of patients can be controlled to a certain degree. Patient flow is focused less on cost and more on the reduction of unnecessary waiting, movement, and processing of patients. The metrics of patient flow often have direct correlation to quality (as will be discussed in the context of the Emergency Department later). The connection between quality and flow arises when patients do not receive timely access to care, often leading to an exacerbation of their symptoms or illness. Alternatively a relatively healthy patient who is waiting for prolonged periods in the hospital has an increasing risk of acquiring a new condition from within the hospital. Not only do these cases have bad implications for quality, they also have bad financial implications. Hospital revenue is generally based on the number of patients treated. If patient flow is poor, then that number of treated patients (and associated revenue) decreases, but the overhead costs of the hospital do not. If poor patient flow leads to a patient acquiring a new illness, that patient must be treated. Often these hospital acquired issues are considered “never events” by CMS and the hospital may not be reimbursed for the extended treatment of that patient. The value of the study of patient flow comes in the understanding that patients moving through the system are the primary driver of the system. History has shown that improving how patients flow will improve cost, quality, and patient satisfaction [Hall 2006, McLaughlin 2008].

Common methods that have been applied specifically to improving patient flow include: queuing theory, process mapping, point-to-point diagrams, computer simulation, staffing and scheduling tools, and forecasting [Smalley 1982, Hall 2006, Graban 2008, McLaughlin 2008].

1.2.6 Organizational Change and Systems Re-design

How a system is organized and designed is a key to its performance. It is possible for two health care organizations with the same resource levels, staff of equal competence, and the same input demands to perform differently. This can be explained by considering organizational and design issues. Like a machine or a computer program, an organization will perform the way it is designed and built to perform, whether that was the intention or not. The ways that information is shared in an organization, how performance in that organization is measured and rewarded, the unique values of each of the organizations stakeholders, and other factors can enable or limit the final performance of the organization, regardless of how much money is spent on resources [Nightingale and Srinivasan 2011]. How different organizational and communication structures affect performance has been studied for many years, including in hospitals [Longest 1974, Shortell et al. 1976]. Despite this there is still a great deal more to be learned about the comparative effectiveness hospital organization structures and how to translate organizational successes from one hospital to another [Fradinho 2011].

1.3 Pulling methods together through frameworks

1.3.1 Continuous improvement frameworks

The above discussion focused on the range of tools that are currently employed to improve the cost/quality ratio of providing care in the hospital environment. It is also common to integrate and sustain the use of these methods, and others, through continuous improvement frameworks. One long standing example is Continuous Quality Improvement (CQI) or Total Quality Management (TQM). One study described the use of CQI/TQM in hospitals through the application of five principles:

“(1) a focus on underlying organizational processes and systems as causes of failure rather than blaming individuals; (2) the use of structured problem-solving approaches based on statistical analysis; (3) the use of cross-functional employee teams; (4) employee empowerment to identify problems and opportunities for improved care and to take the necessary action; and (5) an explicit focus on both internal and external customers.” [Shortell et al. 1995]

Looking at these principles it is possible to see a relationship to the categories of improvement methods that have been suggested in the prior sections; “organizational processes and systems” as in section 1.2.6, “statistical analysis” which is used in operations management and “employee empowerment” and focus on “customers” which relates strongly to human factors. In this way CQI/TQM enables the collection of these tools. A limitation of CQI/TQM is that it does not necessarily encourage large scale, long time frame projects that may have the largest impacts. Instead this framework relies on encouraging and sustaining smaller projects performed by general staff members, which may limit the complexity of the tools that are applied to a specific problem.

Lean is another improvement framework that has been applied in the hospital setting, and high profile successes such as Virginia Mason Hospital have increased interest in this approach [Kenney 2010]. Lean was created as a derivation of the Toyota Production System, based on the efforts of the MIT International Motor Vehicle Program, with the goal of reducing waste and focusing on delivering value [Womack et al. 1991, Womack and Jones 2003]. Since then it has taken on many forms which vary with hospitals that apply it. In general, lean tools fit into the categories that were described above; value stream/process mapping, standard work (like that suggested by the Gilbreths), inventory management, statistical analysis, patient flow analysis tools, and employee engagement through organizational re-design [Graban 2008]. One description of lean in health care health care identifies seven critical flows: Flow of patients, flow of clinicians, flow of medication, flow of supplies, flow of information, flow of equipment, flow of process engineering. The tools and methods for dealing with these flows fit into the categories that have already been discussed [Black and Miller 2008].

1.3.2 Looking at health care delivery systems through the lenses of Enterprise Architecture

The National Academy of engineering and Institute of Medicine joint publication, Building a Better Delivery System, introduced the concept of a health care delivery system (HCD) in a paper entitled “Changing Health Care Delivery Enterprises.” As part of this introduction the full complexity of HCDs is described.

“HCD enterprises are very large complex operational systems comprised of numerous people and machine elements. Tens of thousands of people are involved as providers patients support staff and managers organized into specialties, departments, laboratories, and other organizations... Perhaps most important, these processes involve large numbers of interactions within units, among units, and across processes... We need better ways of analyzing systems of this magnitude.” [Bonder 2005]

Analyzing complex systems is a difficult task. An HCD is a complex adaptive system, comprised of many different interacting parts including a significant human element. The action in one part of the system may have unforeseen consequences later in the system. “Health Care Systems are adaptive because unlike mechanical systems, they are composed of individuals... their actions are not always predictable, and... can be seen as contributing to huge variation in the delivery of health care” [IOM 2001]. It is for this reason that while all of the methods for hospital performance improvement that have been described above are valuable, they become most useful when employed in combination with one another as is done in the continuous improvement frameworks.

Continuous improvement on its own does not address some of the more complex issues that require a full systems view. The tools of Systems Architecture allow systems engineers to better understand, predict or even control the systems, to some degree, despite their complexity. “System architecture is an abstract description of the entities of a system and the relationships between those entities” [ESD 2004]. In other words, system architecture is a representation of a complex system, and by creating this abstraction one can begin to understand the many different connections in the system and begin to predict and control emergent behavior.

Many of the human systems that exist today, particularly in service sectors such as health care, can also be described as an Enterprise. Enterprises have been defined as “complex, highly integrated systems comprised of processes, organizations, information, and supporting technologies, with multifaceted interdependencies and interrelationships across their boundaries” [Nightingale and Rhodes 2004]. The approach to abstracting, understanding, designing and controlling enterprises is called Enterprise Architecture (EA). Given the definition of the complexity of HCDs provided earlier and this definition of an enterprise there can be no question that a hospital can be considered a complex enterprise system and the tools of EA are applicable.

One EA framework uses a holistic approach to representing a complex system and is comprised of 8 views:

1. Strategy: “Business model, business strategies, internal/external strategic drivers, enterprise metrics, and objectives” [Nightingale and Rhodes 2012]. Health care delivery enterprises may be for profit, not-for-profit, or government hospitals. Many academic hospitals choose a specialization strategy and become well known for a specific area of health care delivery, such as cardiac care or cancer care. In contrast many community hospitals have the strategy of providing the best broad spectrum of care so that they better serve the entire community. These are just a few examples of how hospitals can differ in strategy, which can then influence many other elements of the enterprise [Lee and Mongan 2009].

2. Policy: “Policies that impact the enterprise as well as policies internal to the enterprise that affect performance” [Nightingale and Rhodes 2012]. Earlier in this chapter policies for improving the health care system were mentioned. These policies have direct impacts on health care delivery enterprises [Feldstein 2003]. One policy mentioned earlier encourages the purchasing of HIT systems, but also regulates how the HIT systems must be used. As will be discussed in the next chapter, there is a policy that mandates every emergency department to treat any patient that presents. Such policies force the hospitals to create strategies, processes, etc. in order to respond to the external demands. Hospitals may also have internal policies; these policies may focus other enterprise views towards the achievement of one specific task. A famous example was the Pittsburgh Regional Health Care Initiative where member hospitals adopted the policy of eliminating central line infections, focused all views on achieving this goal, and had great success [Spear 2009].

3. Process: “Key business processes, and activities that capture, manipulate, and manage the business information to support business operations” [Nightingale and Rhodes 2012]. Hospitals have many different processes; there are treatment processes, administrative processes, data sharing processes, etc. How these are managed and how well the processes integrate with the other enterprise views can strongly influence the performance of the hospital. In the next chapter there will be some discussion of how processes are changed in the emergency department to improve flow [Graban 2008].

4. Organizational: “The organizational structure of the enterprise, major operations performed by organizations, types of workers, work location, and distribution of organizations to locations” [Nightingale and Rhodes 2012]. Organizational redesign was already discussed as a method for improving how a hospital performs. Indeed how a hospital organizes itself physically, through the hierarchy of staff, or the assignment of duties can have major implications on performance [Spear 2008, Fradinho 2011]. It is partly for this reason that there is an emergence of mid-level providers with more authority in hospitals [Brown et al. 2012]. Similarly reorganizations in how care is provided have led to the rise of team based care in hospitals [IOM 2001, Carter et al. 2009].

5. Knowledge: “All information and knowledge needed to perform the enterprise business operations and relationships among that information” [Nightingale and Rhodes 2012]. Knowledge is of the utmost importance in health care. Providers on all levels go through extensive training build a knowledgebase of symptoms, diseases and treatments. As mentioned earlier, there are attempts at better standardizing and documenting knowledge through decisions support systems yet, for the moment, knowledge continues to remain primarily with practitioners, this sometimes hinders reforms in other enterprise views.

6. Information Technology: “Key IT infrastructure (both hardware and software) that supports the enterprise” [Nightingale and Rhodes 2012]. HIT has already been described as an emerging tool for improving the quality and organization of health care delivery. In hospitals, HIT systems are being implemented to manage/organize significant portions of the care process. These tools have become primary methods for communication, in some cases eliminating the need for two staff members to communicate directly. EHRs have been used in order to improve communication, continuity, and quality in health care delivery organizations and the US government is continuing to invest in HIT with the belief that more gains are yet to come [Chaudhry 2006, CMS 2012b]

7. Product: “Products are developed by the enterprise; key platforms; modular vs. integral architectures, etc.” [Nightingale and Rhodes 2012]. While a hospital’s ultimate output is focused on the service of health care delivery there are many products that are produced in a health care delivery system. For example, the hospital’s internal pharmacy mixes and distributes medications or a podiatrist uses machinery for creating custom orthotics. A health care delivery

enterprise produces many such products which contribute to a patient’s future health. Health care delivery organizations also consume a great deal of products. The quality and availability of these products can have significant impacts on other factors of a health care delivery enterprise.

8. Services: “Services delivered and or supplied by the enterprise” [Nightingale and Rhodes 2012]. One of the primary goals (if not the only primary goal) of a health care delivery enterprise is to provide the service of improved health for a patient. It is this goal that drives the other views of the enterprise. If a view is not properly aligned to improve the service of health care delivery this is often a prime target for system redesign.

The enterprise views allow a user of the framework a certain amount of structure in order to collect the necessary information and understand the key elements of their enterprise. However the high level nature of the views allows a user to bring in tools from other fields and frameworks and apply them in the context of the larger framework. Finally in order to help the user of the framework draw conclusions about how all of the views interact, EA integrates the views into a high level map. This map shows a flow of influence between the different views and allows the user to derive a certain amount of emergent behavior, similar in effect to System Dynamics, though with less quantitative overhead and consequently less specificity [Sterman 2000]. The map is shown in Figure 6 below.

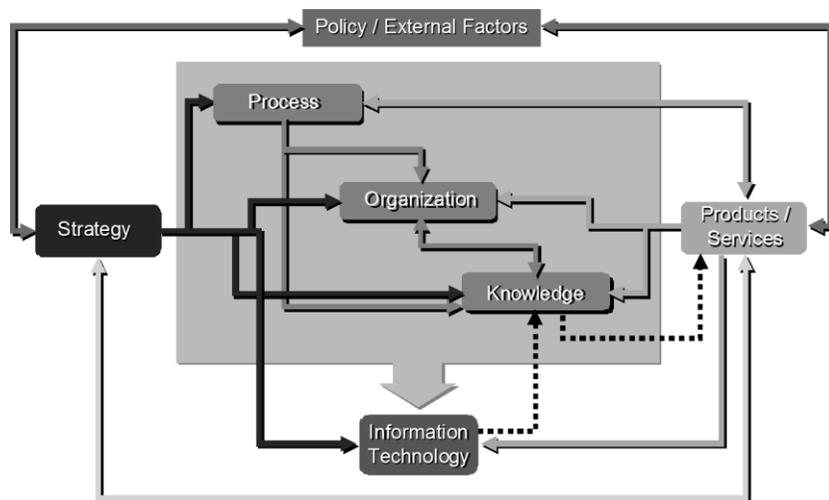


Figure 6 Map of EA framework views [Nightingale and Rhodes 2012].

Looking at Figure 6 it can be understood that any improvement that is to be made to a health care delivery enterprise must take all of the views into account. Policies that are made must be aware

of which services/products they are trying to influence and how the policy will impact the strategies of delivery enterprise. Those who set the strategies for a health care delivery enterprise must be aware of how process, organization, and the exchange/location of knowledge interact in order to enable the delivery of a quality product or service. Finally, information technology can have an overarching influence on all of these views.

Supply chain management is an example of how all of the lenses can be brought together in order to improve a system, however it can also fail when not all views are considered. By looking at the flow of products, supply chain management improves processes such as manufacturing and distribution. Areas of supply chain management also study the organizations that perform those processes, the knowledge that was held and shared between parts of the chain and what information should be collected throughout the chain. Each of these areas of supply chain management have the potential to improve total outcomes, however if the entirety of the enterprise views are not considered in conjunction, it is possible to only achieve pockets of improvement and local optimizations. Successful management of an entire supply chain must also consider enterprise wide strategies that take all of these lenses into account in order to optimize the performance of the supply chain as a whole.

It is common in health care delivery improvement to see each treatment step as a single occurrence that can be optimized. However just as in a supply chain, optimization of performance metrics for a single step may not be best for the entire enterprise. Hospital improvement should also be approached on an enterprise level. Using the EA descriptions above to consider a hospital it is reasonable to start with the service view as the basis for organizing the other views. The service of health care delivery is not provided in a single moment but occurs over time in a series of patient encounters. Understanding the health care delivery organization as a connected enterprise, makes it clear that studying patient flow goes beyond modeling “unit processes” as is usually done, but must include all of the lenses of EA in order to quantify and control flow rather than react to it [Uzsoy 2005]. In other words, improvements to patient flow are an improvement to the service of the hospital, however to do this properly it is important to understand the entire patient flow process on the enterprise level, rather than in a single patient interaction. This approach includes understanding how the organization of the hospital is equipped to react to enterprise patient flow and how HIT can be used to generate knowledge

about flow and facilitate the processes to respond to flow needs. Similarly strategies can be developed that take into account the higher level policy issues (like those discussed in section 1.1) and drive flow improvement on an enterprise level. Finally, in an ideal enterprise, the experiences that come from the improvements should feedback to future policy decisions in order to enable more improvement.

With all of the above in mind, it is necessary to find a solid starting point for an EA approach to improving health care delivery. It was suggested that the service of health care delivery and its relationship to patient flow are the core of a health care delivery enterprise; therefore the next section introduces health care delivery chains as a method for mapping and understanding the flow of patients through the services of the health care delivery enterprise. Then, the chapters to follow will describe a study of an example chain where all enterprise views are considered in developing a method for improving patient flow.

1.4 Defining health care delivery chains

An article on enterprise resource planning for hospitals discusses the need to control a hospital and plan resources by not only understanding long term trends, but also by looking at patients in the system at a moment in time [van Merode et al. 2004]. Although some work continues on the use of enterprise resource planning in hospitals, the common assumption in operations and supply chain management as well as patient flow studies in hospitals is that patients are not controllable, they arrive when they arrive and leave when their treatment complete. The key to using an enterprise view of hospital improvement is to understand that in many cases the flow of the patient is indeed knowable and can therefore be anticipated and controlled.

In Figure 5 above, a key simplification is made in the characterization of the health care delivery portion of the chain as one block, “treat patient.” In fact the treatment of a patient is a long process that involves many treatment steps, each with supply and development chains attached. Figure 7 shows this more complex image of the health care supply chain with multiple treatment steps.

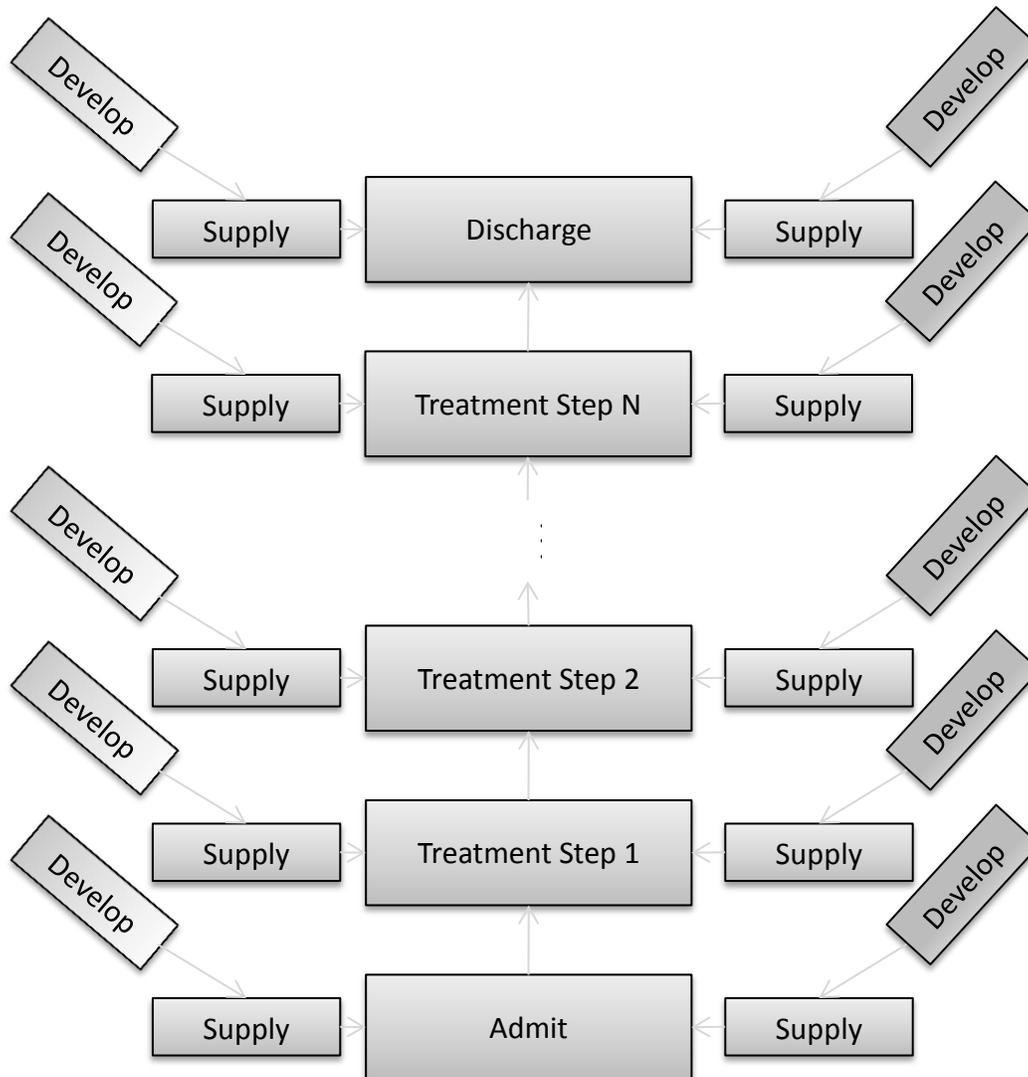


Figure 7 Health care supply chain with multiple treatment steps

The chain of health care treatment steps that runs through the middle of Figure 7 can be called the “health care delivery chain” and when it is combined with the supply and development chains one can see a larger conception of the flows in a health care delivery enterprise. This network of flows is similar in concept to the seven flows in health care discussed earlier; however it takes them to another level by showing the interaction of these flows rather than dealing with each separately [Black 2008]. A concept similar to this was called for in a recent article on supply chain management in health services [Vries and Huijsman 2011]. When taking an enterprise architecture view of the system it is clear that understanding the interaction between the flows is a key aspect of improving the system as a whole.

The health care delivery chain can be summarized in many cases as: Admit, Treatment Step 1, Treatment Step 2, through treatment step N and Discharge. It is the connection of steps/processes through which a patient flows in order to receive their treatment, and offers a new dimension of analysis for optimizing health care delivery. Like the supply chain depiction in Figure 5, simplifying this into one chain is not quite correct. A health care delivery enterprise actually contains many networks of health care delivery chains, supply chains, and product development chains. The time it takes for a product or patient to flow through each of these many chains can vary from seconds to years. In many cases it is also possible to connect the end of the chain “discharge” to a block called “home care” which then feeds back to admission, creating what can be called the health care delivery lifecycle chain (Figure 8). This may be similar to product supply chains that use recycled parts which feed back into the supply chain continuously. However in this case, rather than recycling parts, it is the humans that move through the chain continuously throughout their, literal, life cycle. Truly optimizing health care delivery means optimizing the flow of the patient across this entire lifecycle chain.

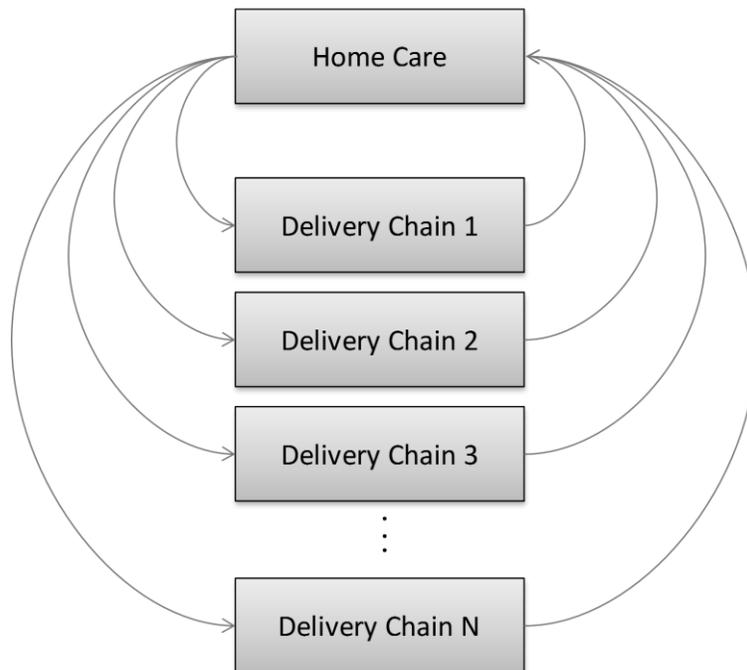


Figure 8 Health care delivery life cycle chain

Often in the organization of a patient’s care, the immediate unit processes or steps in the delivery chain are the key considerations of medical personnel. However it is becoming increasingly

common to find care organizations focusing on how the patient cares for themselves while at home. This new emphasis is often referred to as “the medical home” and is an acknowledgement that the care at home is a significant part of the health care delivery lifecycle chain, and may have large quality and cost impacts [Starfield and Shi 2004, Rosenthal 2008].

Examples of health care delivery chains can be seen on many abstracted levels of the organization:

In department level: Where a patient moves between treatment steps in one department. Such as an exam which leads to ordering a test (such as an x-ray, blood test or CT scan), preparation time/patient waiting time, and then the test. A trip to the hospital may begin with a visit to an emergency department. This trip has many treatment steps that follow each other sequentially from the time the patient’s information is taken to the time that a doctor makes the decision to have the patient admitted to the hospital. The time scale for this level is on the order of minutes to hours.

Cross department level: Where a patient moves between two departments of one organization. Such as when a patient is discharged from an emergency department and is waiting for the inpatient unit bed preparation and admission. The patient then receives treatment in the inpatient unit leading to a discharge out of the hospital. The time scale for this level is on the order of days.

Cross organizational level: Where a patient has multiple points of contact with the health system through administratively unrelated (or loosely related) providers, such as primary care followed by a specialist. Similarly a cross organization chain may be the patient flowing through the hospital and being discharged to a nursing home. The time scale for this level may be on the order of months or years.

On all levels, the path has a similar series of steps, all that changes is the level of abstraction. Similarly, methods developed for improving one level may be applicable to improving another level. The key aspect of working with any of these health care delivery chains is the understanding that state information in the earlier steps of the chain can be used to plan and control steps later in the chain. However doing so properly requires the interaction of many lenses of the health care delivery enterprise that contains the chain of interest.

1.5 Dissertation goals and outline

This dissertation seeks to further develop the concept of health care delivery chains by approaching a well-known issue in health care delivery and patient flow from the health care delivery chain perspective. The chain that will be studied is that which flows through a hospital and connects the emergency department and inpatient unit. This is an example of a cross department chain. It is hoped that the methods that are described will be applicable to chains of other abstraction levels and provide insight into how other chains can be identified and improved.

Chapter 2 will set the stage for the rest of the dissertation. In this chapter, the details of the Emergency Department/Inpatient Unit chain will be discussed. This discussion will include a characterization of the quality metrics that are studied in the emergency department. It will be seen that these quality metrics are heavily influenced by the rate of patient flow in the emergency department and that this patient flow is tightly linked to the performance of the inpatient unit. Past studies that focused on improving emergency department patient flow will be reviewed and placed in the context of the improvement methods described in section 1.2. These past studies will conclude with a recent suggestion: If it were possible to predict that a patient would require admission early in their treatment, then this prediction could be passed to the inpatient unit, improving the timeliness of response when the patient is ready for admission. This in turn would improve flow and quality metrics in the emergency department. This is followed by a discussion of prediction in the health care delivery setting, with some examples of what methods are used to make predictions and how those predictions have been applied. Prediction studies specific to the emergency department will then be discussed.

Studying prediction in the emergency department leads to three primary research questions that will be answered by the research presented in the following chapters:

1. What predictive methods work best to predict downstream demand in the context of a single Emergency Department/Inpatient Unit health care delivery chain?
2. How portable or robust are these prediction methods to multiple hospital contexts?
3. Given advance demand predictions, what possible adaptive actions can a hospital system take to improve flow given (a) perfect and (b) imperfect downstream demand prediction?

Chapter 3 will be focused on answering Question 1. In this chapter, the reader will be introduced to the Veterans Health Administration hospital in West Roxbury, Massachusetts, (VHA West Roxbury) where the primary research for this dissertation has taken place. After this introduction, there will be a discussion of the methods chosen to study prediction at VHA West Roxbury: expert opinion, naïve Bayesian conditional probability, and logit-linear regression. While introducing methods for making predictions, this chapter will also introduce concepts for using these predictions in a practical setting. These practical applications will drive the evaluation of the predictive measures. The results of this chapter will show that the logit-linear regression was the best performing prediction method and worth attempting in other settings.

Chapter 4 will build upon the conclusions of chapter three and describe a study aimed at answering Question 2. In this chapter, three other hospitals will be introduced. These hospitals will have varying sizes and economic strategies. For each hospital the process for creating a logit-linear regression, prediction model in VHA West Roxbury will be repeated in order to judge the portability of this methodology to the new contexts. There will also be a dataset collected for VHA West Roxbury where live nurses performed prospective data coding, this will be used to test how robust the early prediction model is when applied to a live implementation. This chapter will conclude that prediction is indeed possible in multiple contexts.

Chapter 5 will describe a two week live implementation of the prediction system at VHA West Roxbury. For this time period, admission predictions were made in real time, based on codes input by triage nurses, and shared with hospital staff. Studying the results of this implementation is a step towards answering the third research question. When the study was designed it was not known exactly how sharing predictions would influence hospital performance, therefore quantitative and qualitative data were collected. The results of these two methods of data collection will be presented followed by a discussion of their interpretation and value.

Chapter 6 will further seek to understand the potential quantitative results of using prediction in the Emergency Department/Inpatient Unit chain. While Chapter 5 will show some interesting qualitative outcomes of using prediction, no quantitative improvements were found during the short implementation time period. In order to further answer Question 3 and explore the value of perfect and imperfect prediction over a long time period, a discrete event simulation was developed and tested. The simulation was based on VHA West Roxbury. This chapter will begin

with a description of the use of computer simulation in the health care delivery. This will be followed by a description of the simulation that was developed for this study and a validation and verification of the simulation using data from VHA West Roxbury. Finally multiple scenarios will be tested for relative hospital performance when driving hospital behavior using emergency department crowding information, imperfect prediction information, perfect prediction information, and time based strategies. The outcomes of these scenarios will also be tested for sensitivity to changes in the hospital process design and resource levels.

Chapter 7 will close the dissertation with a discussion of the results from the previous chapter and their contributions to the study of emergency department/inpatient unit patient flow improvement, as well as health care delivery chain management. This chapter will include a discussion of the limitations of the study designs as well as future work that could stem from this research.

Chapter 2: The Emergency Department and Inpatient Unit Delivery Chain

Having established the key role played by hospitals in accounting for cost in the health care system this chapter adopts the health care delivery chain framework to discuss one of the core processes of hospital care delivery: the flow of patients from the emergency department (ED) to the inpatient unit (IU). The chapter will begin by discussing the importance of EDs to the US health care system and how ED crowding is affecting the system. Past solutions to ED crowding will be discussed in an input/throughput/output (ITO) paradigm that has been introduced in the ED literature. These solutions will also fit into the solution categories introduced in Chapter 1. The result of this literature review is the recognized need to focus more on the output side of the ED. In other words, solutions that deal with the total ED/IU delivery chain are expected to have higher impact on ED patient flow. This will set the stage for a discussion of the ED/IU system as a health care delivery chain. Together, these discussions will enable an introduction to the studies described in the chapters to follow, focused on improving the management of this chain using prediction. The chapter will end with a discussion of prediction methods in health care and the ED in particular setting the stage for Chapter 3.

2.1 Background and Motivation for Emergency Department Improvement

2.1.1 The role of the Emergency Department

The Emergency Department (ED) is one of the most commonly studied parts of a hospital. In many hospitals, the majority of patients that are admitted, enter through the ED. Despite the large amount of patients that enter the hospital from the ED, typically only small percentages of ED patients are admitted to the hospital (20-35%). This means that the ED deals with a significant amount of patients that are never even seen by the rest of the hospital.

The US health care system relies heavily on EDs. Often a patient is concerned about their health and feels that the issue is too urgent to wait for an appointment with a primary care physician. This is not a rare occurrence given that a primary care an appointment can take up to 44 days for a new patient in Massachusetts [MMS 2012]. In this case the patient goes to an ED. Similarly patients will go to the ED when faced with true emergencies such as traumatic injuries, cardio-

vascular events, acute mental illness, and other issues that cannot wait for any length of time, let alone days. EDs are open 24/7, they are conspicuous in their communities, and it is possible to call transport (ambulances) if needed. Furthermore patients know that if they arrive at the ED they cannot be turned away as “the Emergency Medical Treatment and Active Labor Act (EMTALA) mandates that any individual who presents to a hospital ED must receive a medical screening examination and, if an emergency medical condition is identified, be offered treatment to stabilize that condition or offered safe transfer to an appropriate facility” [Asplin 2006]. Given the convenience of the ED and the confidence that a patient has of being seen, it is little wonder that the volume of ED visits continues to rise, as seen in Figure 9. Despite the increase in ED visits over recent years, the number of operating EDs is actually dropping. This in conjunction with hospital budget cuts means that operating EDs are required to treat more patients with unchanging or reduced resources [US GAO 2003, CDC 2011].

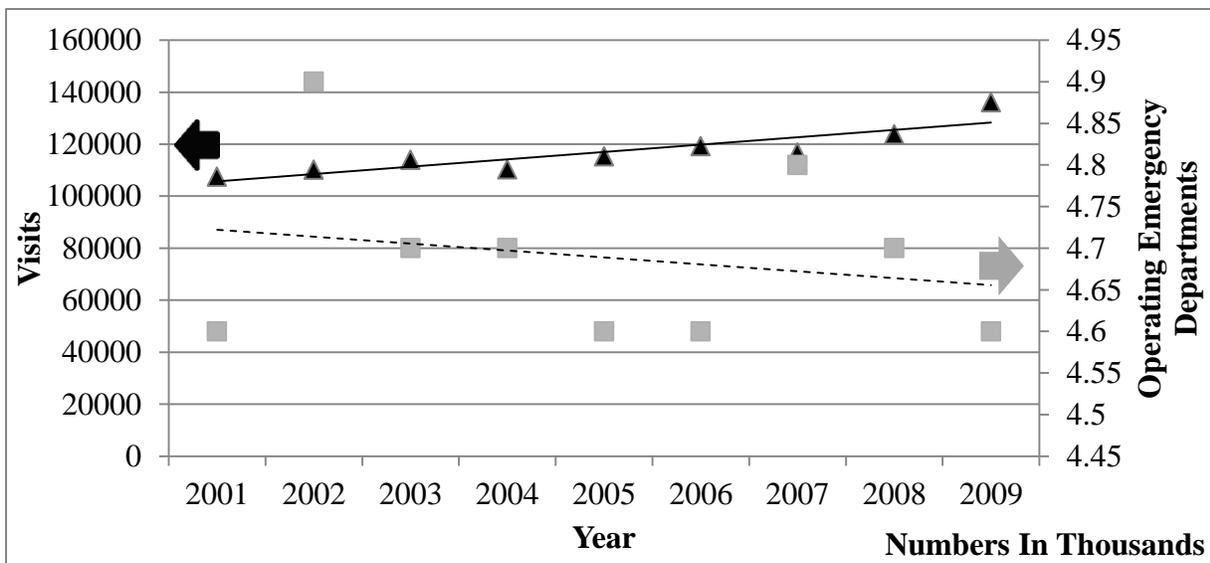


Figure 9 ED visits vs. Operating EDs [CDC 2011]

It is also interesting to note, that the increase in ED usage can be attributed to patients who do have private insurance. This usage is a testament to the perceived need for emergency/immediate care and the convenience of the ED [Weber et al. 2008, Cunningham 2011]. Nonetheless, it is with uninsured and vulnerable populations in mind that many have referred to the ED as the “safety net” of the health care system [Fields 1999, Asplin et al. 2003]. Although not the original intent, this title can also be based on the fact that it is the last chance for the health care system to compensate for a lack of accessibility, even for patients with resources who have otherwise

navigated the system correctly. Despite the importance of the ED to the health care system, the ratio of supply and demand continues to be mismatched.

2.1.2 The Emergency Department System

Until this point, the ED/IU system has been referred to in general terms, before continuing, it is worthwhile defining the system more clearly. An understanding of the system will facilitate the definition of terminology and make improvement targets and quality measures more tangible in later discussions. Recalling the health care delivery chain generalized picture shown in Figure 7, of Chapter 1, the supply chains can be removed leaving the simple health care delivery chain shown in Figure 10.

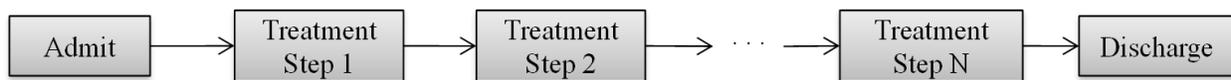


Figure 10 Simple health care delivery chain

The ED/IU system depicted in this way can be seen in Figure 11. The figure is a very simplified representation, but emphasizes that this is indeed the ED/IU health care delivery chain with two primary steps.



Figure 11 ED/IU health care delivery chain

Figure 12 depicts this chain with increased complexity by adding more detail to the flow of patients through the system. As can be seen in the figure, there are two ways that patients can arrive at the ED, by walking in or by an ambulance. While patients wait to get a bed in the ED they may leave without being seen (LWBS). Also if the ED declares, to local authorities, that it is crowded, in some states, ambulances will be diverted to other EDs (this will be discussed more in the next section). When the patient completes their ED treatment they may get discharged to their home, they may be transferred to another hospital, they may leave against medical advice, they may pass away, or they may be admitted to the IU.

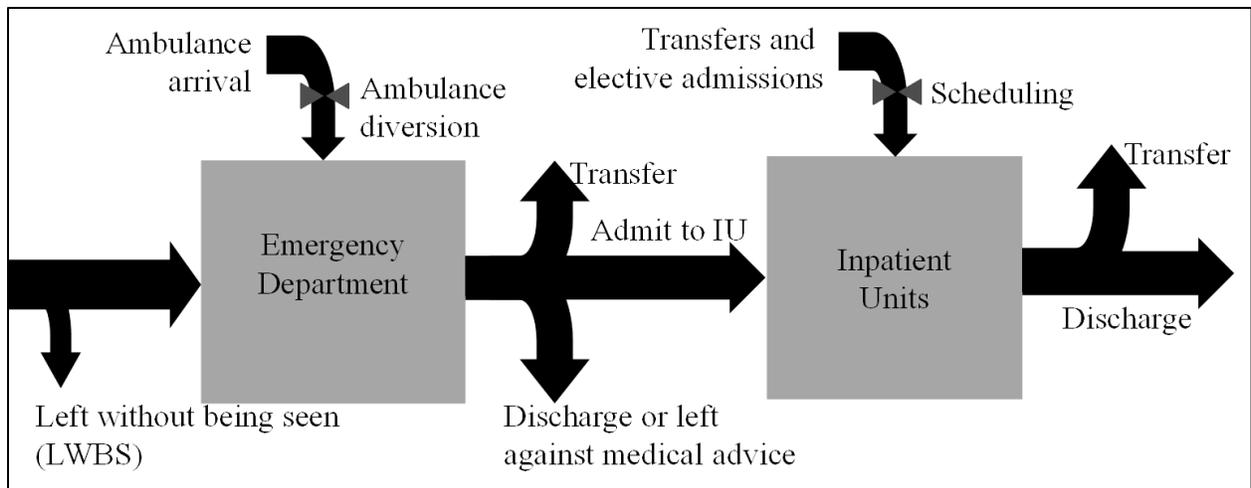


Figure 12 General patient flow model through ED/IU system

In most EDs, approximately one third of the patients get admitted to the IU. Although this becomes the majority of patients who are admitted to the IU, patients can also enter the IU through transfers from other hospitals or elective admissions, for procedures that were scheduled ahead of time. Finally patients leave the IU through discharge, transfer, or when they pass away.

Looking into the ED process in more detail it is possible to see the in-department level delivery chain. This is shown in Figure 13. Improving this delivery chain is the focus of many studies, some of which will be discussed in Section 2.2. However, looking at this chain is also useful in order to better understand quality metrics that are used in the ED. These metrics will be discussed in more detail in the next section.

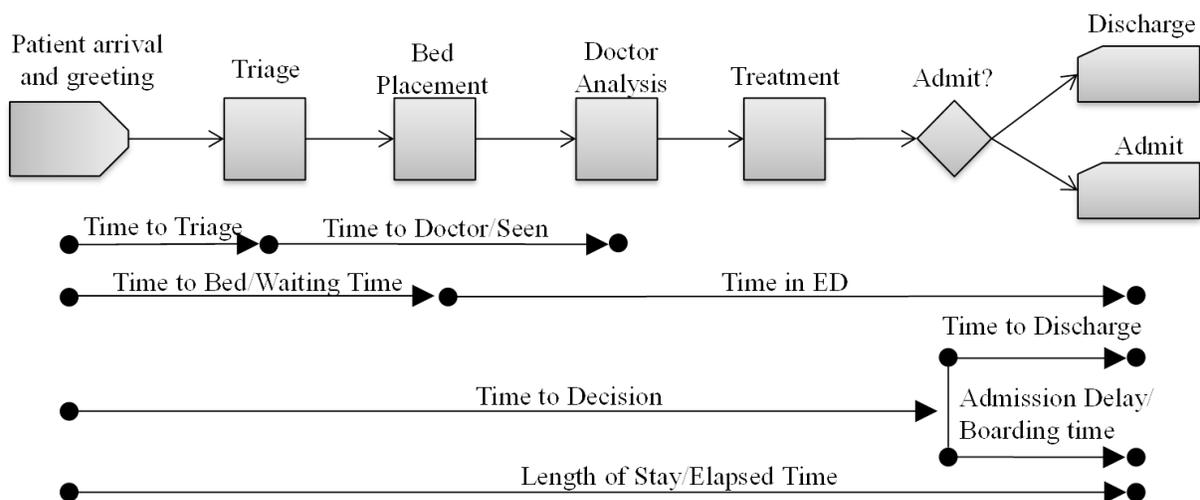


Figure 13 Emergency department in-department level delivery chain

The micro processes in each ED and IU can differ, however in order to better understand the process, the following is a description of patient flow through the ED/IU system for the VHA West Roxbury system which Chapter 3 will introduce in more detail. As described earlier, a patient may enter the ED by walking in or through ambulance delivery. These patients are entered into the computer system by a greeter, which begins the recording of time for that patient. If it is clear that the patient cannot wait, the greeter can get ED staff to expedite the patient into a bed and triage is done at the bed side. Otherwise the patient waits in a waiting area until they are brought to a triage room by a medical practitioner.

In triage, the practitioner will do a basic exam and potentially order some preliminary tests. The practitioner in triage also will assign a triage level to the patient. The initial purpose of triage levels was to prioritize patients based on urgency or acuity. In practice triage is also used for assigning patients to other units such as an urgent care or fast track [Gilboy et al., 2005, Hauswald 2005, Peck and Kim 2010]. The most common system for assigning triage levels is known as the Emergency Severity Index (ESI). ESI Level is assigned by the triage staff based upon the medical urgency and expected resource usage of the patient [Gilboy et al. 2005]. Patients are then assigned beds in an order based on the judgment of a charge nurse who is taking into account: ESI level, order of arrival, and distribution of nurse work load.

Once the patient is in their bed they will receive a nurse and doctor analysis, the “time to doctor” is a commonly measured treatment milestone as will be described in the next subsection. The patient then goes through a series of exams, tests, and treatments, until finally a doctor decides whether the patient requires admission to the hospital or not. If so, the ED doctor will enter an order for an IU bed into the computer system; this order appears on the bed management system that is being monitored by hospital bed managers.

Upon receiving a bed order, the bed managers find an appropriate staffed bed for the patient or wait for an acceptable opening. Nurses are assigned to a specific unit within the hospital and treat the patients in beds on that unit. Doctors can be assigned a patient on many different wards, these assignments are based on a system in the hospital that accounts for the doctor’s current patient panel, loads the doctor has had in recent days, and (to some extent) location of the patient. If the patient is not chosen for admission to the hospital, arrangements are made for a transfer, patient pickup, or patient walk out.

The simple ED/IU chain depiction in Figure 11 can also be made to contain more information by adding a time dimension. This is shown in Figure 14. This abstraction of the system exposes the level of overlap between the ED step and IU step caused by the need to coordinate beds in the IU to accommodate a patient who is being admitted from the ED as described above. As can be seen, while the coordination is taking place the patient holds a bed in the ED, even if they are no longer receiving emergency treatment.

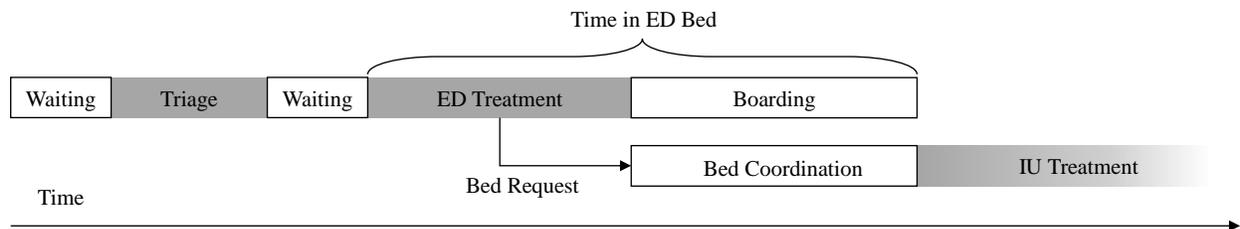


Figure 14 ED/IU delivery chain with time element

2.1.3 Measuring the quality of care in the emergency department

Having provided an understanding of the ED/IU system, it is worth looking at the relevant metrics of quality for this delivery chain. There are metrics for quality related to all diseases that will pass through a hospital, however when considering a chain and the flow of patients that go through that chain, quality is related to timeliness of patient flow. In the emergency environment, disease specific and operational metrics of quality are primarily defined by how quickly a patient gets to and through required treatment [Graff et al. 2002, Bernstein et al. 2009, Horwitz et al. 2010].

Deficiencies in patient flow through the ED/IU chain are often summarized as “Emergency Department Crowding.” While crowding implies an issue of too many patients per unit area, the term has gone well beyond this in the literature. Crowding is a flow issue and to define the problem as crowding is misleading. When a sink is filling with water we do not claim that the problem is a filled sink; we claim the problem is a clogged drain. This concept has not been completely lost in the ED literature [Asplin et al. 2003]. However, despite a growing understanding that ED crowding is really an issue of ED flow the literature continues to struggle to define crowding. The definition has settled upon a mix of its consequences and causes [US GAO 2009]. Whether looking at EDs as being crowded or as having poor flow, there are many quality measures that have been studied in the ED that are influenced by this issue.

An example of a specific disease related metric is “time to antibiotics” for patients with pneumonia. The amount of time it takes for a patient to receive an antibiotic for their pneumonia is directly related to the quality of their treatment outcome. This direct link has caused studies to focus on how the flow of patients through the ED can impact this “time to antibiotic” and consequently the quality of care for the patient. While this measure is disease specific, the goals of studies using this metric were to study the link between flow/crowding and quality of care as a whole, the authors simply choose quality of pneumonia care as a proxy for total care [Fee et al. 2007, Pines et al 2007].

Another common example of a quality measure that gets its roots in flow, but focuses on a specific disease, is time to echocardiogram and balloon inflation. In this case, patients who may be suffering from a myocardial infarction are the target patients and there is an established quality benefit from ensuring that they are diagnosed and treated quickly [Braunwald et al. 2002, Antman et al. 2004, Diercks et al. 2007].

On a similar note, a measure that uses a specific group to generally judge ED quality and flow is the time to diagnosis and treatment of critically ill patients. As was mentioned earlier, there are patients who absolutely, medically cannot wait to enter the ED (often designated ESI 1 or 2). Due to physical limitations of the ED space, finding an area to place one of these patients and getting a practitioner to find a safe moment to stop their current action and move to the new urgent patient can take varying amounts of time. How quickly an ED responds to these urgent patients is an important measure of quality as it touches on the true purpose of the ED, to treat those in an emergency situation [Cowan and Trzeciak 2005, Clark and Normile 2007].

While patients who are in a severe amount of pain may not be in mortal danger, they are also an important group to diagnose quickly and provide initial treatment. Time to pain assessment and analgesic has become a commonly studied metric of ED flow and quality. What makes this metric distinct from others, is that these patients may not necessarily need a bed immediately, just treatment. So, while some of the other measures target how quickly a patient gets to a bed, this metric simply targets how quickly the patient is seen [Hwang et al. 2008, Pines and Hollander 2008]. This is a justification for having practitioners who can prescribe and dispense medication available at triage.

Not all quality/flow metrics are disease specific. Some use more systematic measures that affect all patients. One such quality measure is the number of patients who leave without being seen (LWBS). LWBS was mentioned in the flow model of the ED in Figure 12. In essence this is when a potential patient balks from the ED bed queue. LWBS may not seem to be tied to how quickly a patient gets to and through treatment; however it is in fact closely linked to the average waiting time. Patients will LWBS for many reasons such as; frustration with the wait, reconsidering the need for treatment, concern for the cost of treatment, actual waiting room crowding, etc. The major concern about LWBS is that it does not follow the intuition that only low acuity patients (patients who are not urgent) LWBS. In fact sometimes patients with high acuities will leave, causing them to return to an ED later with their emergency exacerbated by the delay [Baker et al. 1991, Bindman et al. 1991, Weiss et al. 2005, Asaro et al. 2007].

When flow through the ED is delayed, patients will remain in treatment for a longer period of time. Thus another systematic quality issue related to the flow of patients through the ED is exposure to safety risks. As described earlier when flow is poor, the ED becomes crowded and this increases the opportunity for a patient to catch a disease from another patient. Physically crowded conditions can also increase chances of accidental physical injuries. Also when patients are not flowing quickly through the system, and the ED is getting crowded, practitioners must continually change which patient they are attending, increasing the chance for medical errors [Trzeciak and Rivers 2003, Hollander and Pines 2007].

As mentioned earlier, in many local hospital systems, it is possible for an ED to declare ambulance diversion status. This means that they will not accept ambulances unless that ambulance is carrying someone who cannot survive being diverted to a farther hospital. The idea behind ambulance diversion is to turn off the faucet when the sink gets full. However it was found that this didn't really help the system (since many hospitals in one area would go on diversion at the same time) and many states have outlawed this practice. For those areas that still practice diversion it is often used as a proxy measure for being crowded, which in turn is a proxy measure for bad flow. Since diversion can be harmful to patients, the amount of time that a hospital spends in diversion status can be seen as another quality measure that has its roots in flow issues [Kelen et al. 2001, Asplin 2003, Schull et al. 2003, McConnell et al. 2005, Patel et al. 2006].

As can be seen, many of the flow/crowding measures, are rooted in symptoms as opposed to direct measures. This is because most of the studies have been done by practitioners. Practitioners are more interested in proving how flow affects quality, rather than studying flow itself. Despite the fact that flow was not the target of these studies, thanks to the research that has been done by these practitioners, it has been proven that flow has a direct implication on quality of care. It is therefore possible to study the health care delivery chain by purely using flow measures, knowing that improvements in flow will indeed impact care.

One example of a more direct measure of flow is general time to provider. This measure can have two different meanings, some may consider the time to provider to be the time to triage where the patient will often see a nurse, physician assistant and sometimes a physician. Alternatively time to doctor is also a commonly used measure that is more concerned with when the patient is seen by a physician only. Measures like this cause some hospitals to invest in placing a physician in triage in order to improve performance. In truth, this is merely playing with the metric and does not necessarily have impact on the flow issues that may have been delaying the patient from getting to a bed and the traditional physician exam [CDC 2011, Hing and Bhuiya 2012]. While the measure can possibly be gamed by a hospital, with a clear definition it could be more valuable as it closely relates to the disease specific measures discussed above. There was a recent study that collected data from the Center for Disease Control's National Hospital Ambulatory Medical Care Survey (leading up to 2009) which shows that average waiting times to see a provider has changed over time. The results are redisplayed in Figure 15. As can be seen, time to provider seems to be rising, which may correlate to the shrinking resource/patient visit ration data provided in Figure 9.

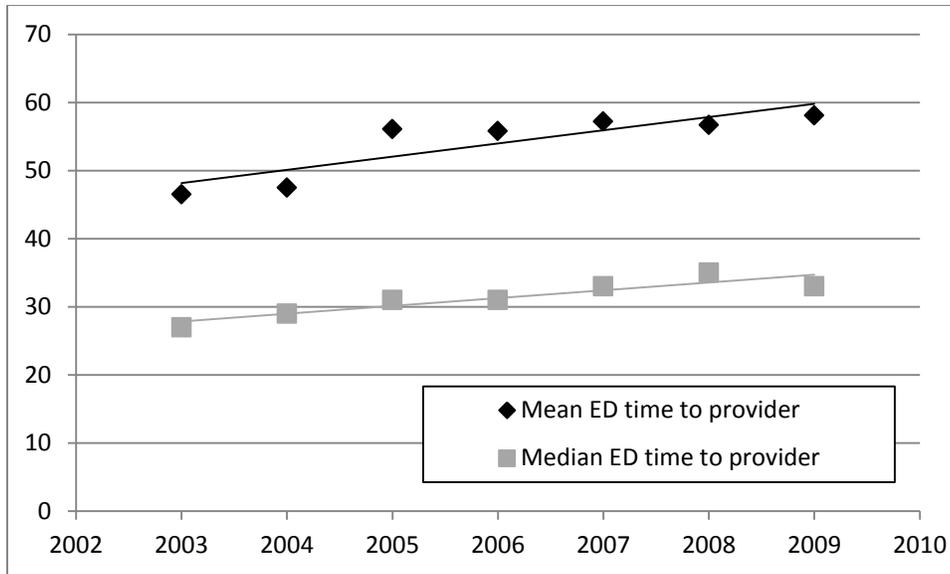


Figure 15 Growing waiting time to see an ED physician [Hing and Bhuiya 2012]

All of the measures discussed above are important depending on what perspective one has when looking at ED flow and crowding. However if one considers the higher level view where the ED/IU delivery chain is considered as a whole, connected system, it is important to look at the metrics that directly tie the links in the chain. The first such metric that will be used throughout this dissertation is “waiting time.” For this study, waiting time is defined as the time period beginning when a patient enters the chain (in the admit phase of the generic chain or arrival in the ED specific chain) and when they enter the first treatment step (for this study, the ED bed). The second treatment step of the ED/IU chain is entering the IU, therefore the next measure that is important to consider in this chain is “boarding time,” or the waiting time between when an ED practitioner decides that the patient will need admission, to the time that the patient actually leaves the ED and is placed in an IU bed. Indeed these measures have been discussed as important in this system [Solberg et al. 2003, Wilson and Nguyen 2004, Olshaker and Rathlev 2006, Falvo 2007, US GAO 2009, Wong et al. 2010]. The next section will discuss solutions to ED crowding that have been studied leading to an understanding that those which focus on reducing boarding time have some of the greatest impact on waiting time as well.

2.2 Past Solutions

Studies in the ED have not only focused on measuring the causes and effects of poor patient flow. There have also been many studies focused on solutions to the issue. The solutions have been categorized as “increased resources, demand management, and operations research” [Hoot and Aronsky 2008]. To some degree these solutions and the above categories reflect the ITO model of ED patient flow [Asplin et al. 2003]:

- Input (demand management),
- Throughput (increased ED resources and operations research),
- Output (increased IU resources and operations research)

The ITO model was mentioned earlier when discussing the flow of patients through the ED. This model is credited with being one of the first clinically originated models to frame ED flow as a complex problem that could be improved through better understanding of the larger system in which the ED operates. This is in contrast to believing that solutions can be generated within the ED alone. Many solutions that fall into the ITO categories can be found in summary articles that have been published [US GAO 2003, Olshaker and Rathlev 2006, Bernstein and Asplin 2006, ACEP 2008, Moskop et al. 2009, US GAO 2009]. For the sake of understanding, some of these solutions as categorized by the ITO framework are presented here. Many of the presented solutions also fall into the categories that were presented in Chapter 1: Electronic health records and health information technology, Decision support tools, Human factors, Health care supply chains/Operations management, Patient flow, Organizational change and systems redesign, Improvement frameworks.

2.2.1 Input

The input side of ED improvement solutions is focused on better controlling what patients arrive at the ED, how those patients are organized once they arrive at the ED, and better mitigating the negative effects of ambulance diversion. One example of an input solution that was suggested is adding more registration clerks to the front/greeting desk. These clerks would help answer patient questions, better prepare patients for their ED treatment and improve overall flow [McGuire 1994]. This is also an example of a system redesign.

Another solution used decision support tools to suggest exams that a nurse could order for a patient when they are first seen in triage. The system would facilitate having the exams completed or ready when the patient enters their bed [Kirtland et al. 1995]. This is actually delivery chain solution as it takes the ED in-department chain into account by using in input treatment step to improve a later treatment step.

Earlier the concept of having ED physicians perform triage was mentioned as a way of reducing the time to doctor, however another benefit was found to this in terms of better controlling what patients actually enter the ED. Having a doctor in triage satisfies the EMTALA legal requirements that every patient who enters the ED be seen but also allows the doctor to send a patient to a less resource limited treatment area before they add to the ED load, such as an acute care unit or other non-emergency walk in option [Kelen 2001, ACEP 2008]. This may be considered a patient flow based system redesign.

Another common suggestion for controlling the input of patients is to increase the availability of primary care through increased insurance [Richardson 2002]. As mentioned earlier it is unclear whether this will in fact reduce the number of patients using the ED, since the majority of patients who enter the ED are insured [Weber et al. 2008]. In the end the value from increasing capability to pay for primary care is limited until the capacity of primary care is increased; the combination of interventions may increase access [Parchman and Culler 1999, Paradise et al. 2011]. This solution was discussed in Chapter 1 as a large policy level issue.

Earlier, ambulance diversion was mentioned as a measure of quality affected by patient flow. To some degree diversion was also initiated as a solution to patient flow issues. Diversion tries to cut off a portion of the ED input in order to give the ED a chance to reduce its crowding. Some studies accept diversion as a useful tool but understand that it has some quality implications. These studies focused on how to allow diversion but reduce its impact by improving geographic coordination or being more specific about where to send patients based on their needs [Wilson and Nguyen 2004, Patel et al. 2006, Shah et al. 2006, US GAO 2009]. These improvements use a mix of system redesigns and decision support tools.

While it is possible to make more primary care available to patients through increased insurance and capacity, this does not necessarily mean that the patients will still choose primary care over

ED care. Often patients go to the ED simply because they don't know that other options exist. For this reason some studies have looked at how community education can reduce the number of patients that go to the ED. These education interventions result in some patients choosing another treatment option during a particular event, but also make use of primary care options that reduce the need for the ED [Michelen et al. 2006, Gawande 2011]. This is another larger policy based solution.

2.2.2 Throughput

While many input interventions may help reduce who is coming into the ED which may reduce the crowding symptom of poor patient flow, there would still be the issues that caused flow problems in the first place. Turning off ones sink does not mean the drain is no longer clogged. Therefore while input solutions may help in the short run, as the population grows and ED capacity does not, EDs would be faced with the crowding issue again. With this in mind there are many studies that focus on how to improve throughput in the ED. For the sake of this discussion, throughput can be defined as the amount of time it takes between when the patient has been triaged and the ED completes the tests and procedures involved with the patient's treatment, culminating in a discharge/admit decision. Many of the studies about ED crowding are initiated by ED personnel. Therefore they are limited to making changes within the ED. For this reason throughput based solutions are more common than input or output based solutions. The following are a sampling of such solutions.

From a cost perspective, it is ideal to be able to leave staffing levels alone and change processes in order to improve flow, however, ED flow may be limited by the availability of staffed beds. This means that a viable solution to ED flow issues may be redistribution of staff or changing staffing levels entirely. Queuing theory is a mathematical method that uses the average rate and variability of patient arrivals and service in order to calculate a staffing level necessary in order to achieve a target average waiting time. For organizations that have the resources and flexibility to respond to these staffing recommendations, this may be the most straightforward method for improving flow. Similarly it may be possible to use queuing theory by holding staffing constant and seeing what average service speed would be needed to achieve a waiting time goal. This service speed can then be set as a target for other improvements. This application of queuing theory may be necessary when the limitation on staffed beds is not the staff, but the number of

beds or rooms physically available or when finances are limited [Vassilacopoulos 1985, Green et al. 2006].

In order to make the problem mathematically tractable, the application of queuing theory often requires some simplifying assumptions. In many cases these assumptions are reasonable; however in cases where more detail about the department needs inclusion, computer simulation can also be used for the same purpose [Rossetti et al. 1999, Samaha et al. 2003]. More discussion on the use of simulation in health care and the ED/IU system in particular will take place in Chapter 6. Queuing and simulation fall into the operations management and patient flow categories of solutions.

Studies of patients who were waiting for long periods of time in the ED, and causing crowding in the literal sense, found that low acuity patients waited the longest. This is unsurprising, given the nature of the triage process. While letting patients of lower acuity wait may be medically and morally justified, as mentioned earlier, there are health hazards associated with patients waiting too long without being seen and having physically crowded spaces. Hospitals also desire a reputation for a positive customer experience. One popular solution to alleviate the buildup of low acuity patients is Fast Track. Fast Track is a set of resources designated for treating low acuity patients that are not expected to need extended treatment. Fast Track allows an ED to quickly process the low acuity patients and thereby relieve the potential complications of having patients wait for too long. While Fast Track has been found to reduce average waiting times, it has been noted that the resources set aside for low acuity patients must be minimal (such as a nurse and bench), if Fast Track resources are capable of treating higher acuity patients (bed, nurse and doctor) then the ED may have a moral obligation to accept patients of higher acuity and let low acuity patients wait as before [Meislin 1988, Rubino 2007, ACEP 2008, Peck and Kim 2010]. This solution can be seen as organizational change and systems redesign.

A great deal of the work is done in the ED, independent of the hospital to which it is attached. It is cost effective to keep the ED attached to the hospital as the connection enables access to a wider variety of testing and diagnostic tools. This connection can sometimes mean that hospital support services, used by the ED, may not make the ED a first priority and may delay the treatment of ED patients [Peck et al. 2010]. For this reason flow can sometimes be improved by designating some support services that are strictly used by the ED. One example of designating

support services to the ED, in order to improve flow, is building a mini-laboratory within the ED. This mini-laboratory could be used to handle more common tests ordered by the ED (enough to keep the mini-laboratory utilization high) and reduce reliance on the hospital laboratory which may have conflicting priorities [Lee-Lewandrowski et al. 2003]. This is another example of a systems redesign.

While it is often the case that the delay for an ED bed is the bottleneck in the system, there are other times where the ED may have been quiet for some time and then receives a sudden increase in arrivals. In this case, the resources set aside for triage and registration of a patient may act as the bottleneck for giving a patient a bed, rather than the limitation of available ED beds. For this reason, some systems redesign solutions focus on eliminating this bottleneck. One example of this is moving patient registration to the bedside. An ED is legally bound to see a patient irrespective of their ability to pay; consequently, there is no requirement to do a full registration of the patient before they are in a bed. By performing registration at the bed side, it can be done while a patient is waiting for treatment reasons and therefore does not disrupt flow [Gorelick 2005, ACEP 2008]. Taking this to a greater level, if hospital staff are willing to enable an organizational redesign that allows for variable actions, all front end operations (triage and registration) can be bypassed when there are opened beds. This is called a “direct to bed” methodology, however the staff would have to be constantly aware of the ED bed situation and be willing to move between offering front end triage or not based on ED’s state. This lack of standard flow may be frustrating for staff but is working well in the VHA West Roxbury ED and at other hospitals [Bertoty et al. 2007].

One more common operations management technique that has been applied to improving ED flow is assigning standard times to specific processes in the ED. While much of what is done in the ED is variable (based on practitioner and patient) there are some tasks such as lab tests that have tighter bounds on how long they should take. This makes it possible to look for best performance levels for a task and assign a maximum allowable time that this task should take. This approach may be used to reduce flow times or variability in flow times [McGuire 1994].

It has been shown in manufacturing that having electronic displays can be used to improve communication and illuminate areas that need attention. This has similarly been found to be the case in the ED. Many EDs have adopted HIT systems that include an electronic tracking board.

These tracking boards enable all nursing staff to gain a quick understanding of how the ED is performing. Often these boards include indicators of how long a patient has been waiting and will chime, flash, and/or change the color of a patient's name when the patient has been waiting for longer than a predetermined time. The boards often color code patients based on their triage urgency level which can be updated as the patient is treated. The boards also include information about the patient that make it easier for nursing managers to quickly make decisions about moving a patient or asking patients to share a room [Boger 2003]. It is possible to take the boards even further by combining the information on the boards with other information in the hospital and in patient records to create patient flow based decision support tools for hospital and ED management [Gordon and Asplin 2004]. Another example of a quick way to get the state of the ED using a decision support system is an ED crowding index. Multiple types of crowding indexes or measures have been explored that are calculated using multiple different methodologies and definitions of crowding. Some EDs use these crowding indexes in order to display whether the ED is crowded or not throughout the hospital. This ED state can be used to enable decision making by hospital management with a systematic view, or merely serve as a warning to potential future ED visitors [Bernstein et al. 2003, Weiss et al. 2005, Jones et al. 2006, Epstein and Tian 2006, McCarthy et al. 2008].

One of the most popular tools of the improvement frameworks described in chapter one is the use of process mapping. Often when a problem is identified, the first task taken by a team is to map the process around it and look for areas of waste. Excess processing, excess communication, excess movement, and other areas of waste can be identified when looking at a process map and then a new process can be created which does not include wasteful steps. Often guidelines of types of waste are written in order to help facilitate the process mapping activity. Process mapping and waste identification tools have been implemented in EDs as well as other health care systems [King et al. 2006, Graban 2008, Black and Miller 2008, Dickson et al. 2009]. In many cases the identified waste is that of excess movement. ED staff may walk back and forth across the ED many times a day between supply rooms and patient rooms. A way to fix this that has shown great improvements in productivity is changing the physical layout of the ED or by creating mobile supply carts [Miro et al. 2003].

2.2.3 Output

While there has been a great deal of effort put into the input and throughput aspects of the ED, often the newest solutions focus on the output side of the ED. Indeed it is now understood that the interface between the ED and where patients flow out of the ED, the IU, is “the single most important factor” [Olshaker and Rathlev 2006] attributed to flow problems experienced by the ED [US GAO 2003, 2009].

As discussed earlier, when a patient is approaching the end of their emergency treatment, an ED doctor may recommend the patient for admission to the IU, and bed coordination begins. The time it takes to assign a bed to a patient can be long. While a patient waits for admission they hold a bed in the ED, also known as boarding. Patient boarding delays the ED from taking a new patient even if the treatment on the current patient is complete [US GAO 2003]. It is for this reason that the wait time between the ED and IU is directly tied to waiting time for other patients to enter the ED [Falvo et al. 2007]. With the impact of ED boarding in mind, many techniques have been studied to reduce this output flow delay.

The concept of using queuing models to make staffing and resource decisions has already been introduced in the context of improving ED throughput. However when considering the ED/IU system, ED boarding time can also be re-defined as the waiting time to get into the IU system. Thus, just as queuing theory was applied to the ED to reduce ED waiting time it can equally be applied to the IU in order to reduce IU waiting time/ED boarding time [Green and Nguyen 2001, de Bruin et al. 2007]. Once again, similar studies can be performed using simulations if mathematical models are undesirable [Levin et al. 2008]. Just as before, the use of these tools is an acceptable approach when the hospital has the financial resources to act upon the recommendations and shift or add staffing/beds. To that end, the simple addition of resources is also a solution even when not analyzed using models [McConnell et al. 2005]. As described earlier, without financial resources and staff flexibility, the models can only act as a goal setting tool.

When a system has resources to expend on the issue of output there are other operations management based solutions that can be considered. While increasing IU resources may actually improve output flow, sometimes funding and administration may be disconnected and the ED is on its own to reduce the symptoms of poor output flow. In this case one solution that has been

studied is adding a buffer of beds between the ED and IU. Buffers reduce the impact of boarding patients, but do not eliminate the quality issues of having extra patients in the ED. Buffers are also subject to overflowing and result in the same boarding issues if output flow is badly blocked [McGuire 1994, Kolb et al. 2008, ACEP 2008]. In order to reduce some of the quality issues of having extra patients in the ED, it is also possible to invest in transition teams of midlevel providers who watch boarding patients and unburden ED physicians [Ganapathy and Zwemer 2003].

As has already been suggested, in the case of beds and hospital staff, the absence of funds to increase staffed bed capacity can be compensated with some flexibility to match scheduled capacity with expected demand at different times of day. Increased IU capacity can be reached by moving staffing schedules around or it may be accomplished through better coordinating the flows of patients into the IU, that compete with the flow coming from the ED. Many hospitals accept patients for elective and surgical admissions that are not emergencies; these patients use IU beds just like patients coming from the ED but are scheduled for arrival. These schedules can be better controlled to match expected bed demand from the ED or can be cancelled when unexpected demand from the ED occurs [ACEP 2008]. Better planning of elective admissions increases effective capacity for patients being admitted from the ED, however many of these elective admissions can become emergencies if delayed for too long, they are also significant revenue generators. With this in mind, hospital administration may have a difficult decision between ED quality measures and other important hospital performance measures when controlling elective admission schedules.

Just as in the ED, process mapping, improvement frameworks, and organizational redesign can be equally useful when improving flow between the ED and the IU. By looking at the process of ordering an IU bed for an ED patient and how that bed eventually gets assigned, it is possible to find many areas of waste that will reduce the efficiency of this flow. One noted example of waste is the policy to have ED physicians request an IU consult to visit the ED to decide whether a patient can be admitted. Waste can be reduced by having the ED physician make the decision. While the original policy may increase the likelihood that a patient only gets admitted if they really need it, or that the patient is admitted to the best ward for their illness, it also significantly reduces the flow rate out of the ED. Instead, it is possible to use a set of hospitalists whose job is

to facilitate admissions, this reduces variability, increases the speed of the process and it also overcomes incentives that consultants may have to keep their ward less crowded [Howell et al. 2004]. Another approach to improving the flow of patients between the ED and IU is to create teams that facilitate the patient's admission [Moss et al. 2002]. An administrative way to reduce admission delays is to institute hospital bed managers. These managers can be used in many ways. Bed managers may pressure inpatient staff to fill empty beds, search the hospital for beds that could be emptied soon, or enforce an ED physician's decision to admit a patient. The use of such bed controllers continues to grow and show reductions in ED boarding time [Moskop et al. 2009].

Often flow from the ED into the IU is delayed due to a mismatch in the time that ED based demand for IU beds grows and the time when the supply of beds is replenished through discharges [Williams 2006]. One solution that has been suggested to this problem is to increase awareness of crowding in the ED by displaying crowding measures (discussed earlier) throughout the hospital. These methods are often accompanied by a system of colors to express the urgency of ED crowding. Another method is to simply increase communication between the ED and IU through regular updates [Howell et al. 2008]. While these methods may work in some hospital cultures, increasing communication can be time consuming for inpatient staff that have other work to do. IU staff also may not necessarily begin to work more urgently knowing that the ED is busy if they do not internalize the connection between ED business and future IU business. For this reason, other solutions have been created that more directly specify IU staff actions in order to improve ED flow.

One popular heuristic solution to this problem is called the discharge by noon system. In this system, doctors are encouraged or incentivized to discharge any patient (that can be safely discharged) before noon. The early discharge policy enables beds to be ready for the characteristic surge of ED admissions just after noon. This system has proved successful in many hospitals however it is often displeasing to doctors who may feel that forcing early IU discharge means that they cannot spend their mornings treating patients, this may be counter to their incentives and to treatment quality. IU staff may also feel that their efforts would be wasted if the number of opened beds is greater than future need [Rubino et al. 2007, ACEP 2008].

Another method to encourage IU staff to work faster and alleviate the pressure of ED boarding is the controversial hallway admissions strategy. In this case, rather than have patients board in the ED they are placed in beds in the IU hallways. This strategy puts pressure on IU staff to get these patients into rooms, has been shown to successfully reduce ED boarding and waiting times, however it only shifts the burden of waiting patients and does not alleviate it from the whole system [ACEP 2008, Viccellio et al. 2009].

In response to the limitations of current solutions to improve flow between the ED and IU, recent literature has suggested that if IU admission could be predicted and communicated to the rest of the hospital, when a patient enters the ED, then the IU could begin preparations before the patient has completed emergency treatment. Overlapping emergency treatment and bed coordination should reduce boarding time and consequently ED waiting time [Yen and Gorelick 2007]. The concept of using prediction in this way is shown in principle in Figure 16. By making a prediction of likely admission at an early step such as triage, the bed prediction could serve as a signal of demand to the IU. This would cause the bed coordination process to begin while the patient is still undergoing their emergency treatment. While this would reduce the bottleneck that exists in the current system, it would add uncertainty to the system. The IU would have to decide how often to respond to predictions that could be incorrect, as the ED provider would still make the final decision whether to really admit the patient or not.

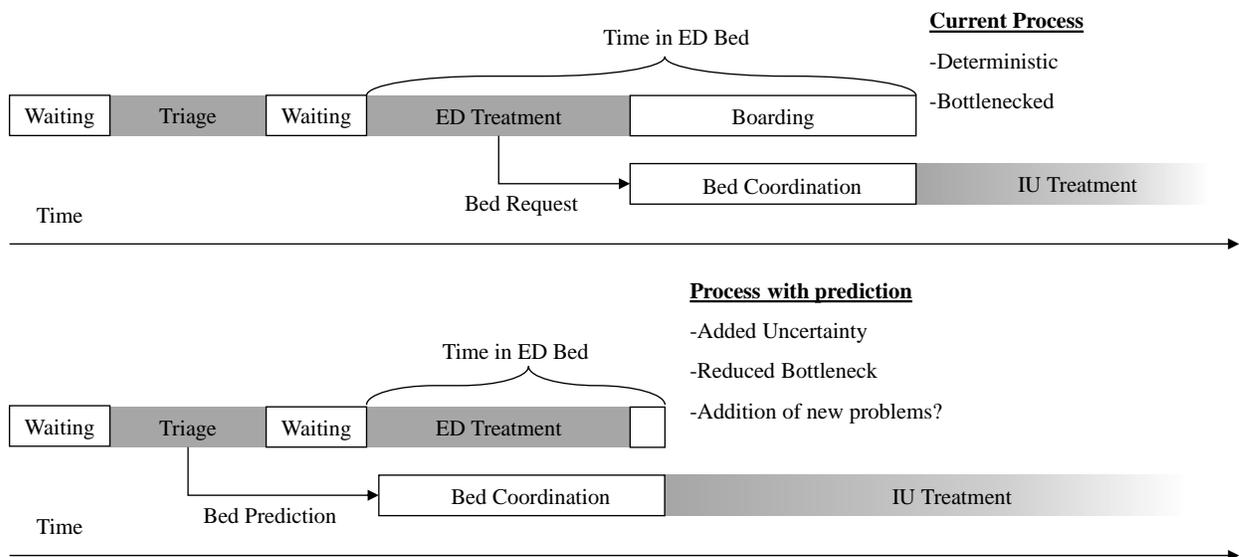


Figure 16 Timeline conception of using prediction to reduce boarding time

The suggestion to use prediction was made in other studies in an off handed way and it is natural for a medical practitioner to feel a little discomfort with the idea of using prediction. There is a great deal of uncertainty in health treatment. When a medical practitioner makes a diagnosis, 100% confidence is rarely, if ever, achieved. In an environment of such uncertainty it may be desirable to have moments of absolutism. However waiting for the concrete decision (such as the decision to admit a patient), though comfortable, is not necessarily what is best for flow. This delay for certainty allows the bottleneck of assigning patients between the ED and IU to become significant. On the other hand by introducing some uncertainty, it is possible to improve flow and productivity. Looking at a simple inventory model, it is possible to wait until customers order a product before beginning production; however that means long wait times. To solve this, prediction can be used to guess product demand, start production early and supply the predicted amount of customers in a timely manner.

The logic behind using prediction is a byproduct of the fact that the ED and IU are indeed parts of a health care delivery chain. It is therefore hoped that the study of using prediction to improve flow in the ED/IU chain can lead to similar work in other chains which are similarly connected. While the study will focus specifically on the ED/IU chain the goal is to be an example of how this approach may be used in other chains. To begin the exploration of prediction in the ED/IU chain, there will now be a discussion of prediction in health care as a whole, followed by a discussion of prediction in the ED in particular. These discussions will set the stage for the rest of the dissertation, focused on studying the application of prediction to an ED.

2.3 Prediction

Although formalized prediction is not particularly common in the daily activities of a health care delivery organization, ED admission is not the first health care context in which prediction was a recommended solution. The use of prediction to allow proactive care has been listed as one of the “new rules” for redesigning and improving care in the IOM report, “Crossing the Quality Chasm” [IOM 2001]. Specifically, the report calls for the anticipation of needs:

“Under the current approach, health care resources are marshaled when they are needed. The system works largely in a reactive mode, awaiting complications and underinvesting in prevention. The new system would not wait for trouble. It would use patient registries to track patients and draw them into care. It would use predictive models to anticipate

demand and allocate its resources according to those predictions, thereby smoothing flow.” In short this means that we should “organize health care to predict and anticipate needs based on knowledge of patients, local conditions, and a thorough knowledge of the natural history of illness” [IOM 2001].

The “anticipation of needs” as described by the IOM has two components. The first is the actual prediction. The prediction can be based on a known, deterministic, series of actions (for example, if a patient is diagnosed with an issue then there is a series of steps that all patients must follow), or the prediction may be based on a forecast of need (for example a situation where not all patients go through the same exact steps but a prediction can be made for each individual patient, such as the ED). The second piece of “anticipation of needs” is the actual action that follows the anticipation. Whether the future steps are deterministic or forecasted, the system must be set up in order to respond. The IOM specifically suggests developing responses based on the allocation of resources (such as medical supplies and staff), but other actions may include changes in decision making, improved scheduling, increased coordination between providers, etc. While it is not always done, responding to predetermined needs is relatively straight forward, however responding to predicted needs is not. This section will discuss more what is meant by predicting needs by mentioning methods for making these predictions in the context of their applications.

2.3.1 Prediction in Health Care

Prediction in health care can be broken down into multiple dimensions. The first dimension is how the prediction is to be used. Just as decision support systems were broken down into those that are used for clinical purposes (CDSS) and those that were used for operational purposes (HMSS), so too can predictions, which are often the heart of the decision support systems.

A clinical prediction is also often called a prognostic. Given a set of symptoms that a patient has, the predictor can provide an assessment of likely conditions that may be causing the symptoms. A CDSS may take this likelihood assessment and assist a medical professional in choosing tests and other diagnosis tools in order to narrow down the patient’s true condition.

An operational prediction can focus on a patient but can also focus on a health care delivery system as a whole. In the case of a patient, the operational prediction can assess the likely future state of a patient, such as needing a transfer, being admitted to a long term care facility, needing

re-admission in the future, missing a scheduled appointment, etc. These predictions are used (often as part of an HMSS) to help managers choose actions that facilitate or prevent future steps the patient may make. In the case of a health care delivery system, data from historical patient visits, historical system performance, or aggregation of current system states could be used in order to predict potential long term or near term states of the system, which can then be applied towards facilitating or avoiding these states.

The second dimension of predictions is based on the type of data used in order to make the prediction. The data can be exogenous, a type of data that is unrelated to the system, such as time of day, day of year, temperature outside etc. Alternatively data can be endogenous, a type of data that is specific to the system, such as number of patients currently in the system, current values of key metrics, number of staff who are on duty, current number of available beds, patient attributes, etc. Two hospitals that are geographically adjacent to one another would always have the same exogenous variables but not necessarily the same endogenous variables. When one considers predictions (or forecasting, which is often used interchangeably) one or both of the types of variables can be applied to make a prediction.

Often, supply chain management is concerned with matching supply with demand by optimally coordinating all parts of the supply chain. To better plan the supply chain ahead of the arrival of actual demand, prediction is a common tool. If one considers the interactions of two parts of the chain (just as the ED and IU are being considered here) there are studies of how predictions can be shared in order to improve coordination and performance [Simchi-Levi et al. 2003, Kurtulus et al. 2011]. Often these predictions rely simply on historical demand over time which may be considered an exogenous variable. However other supply chains that have returning customers, contract agreements, or subscription members are able to use the endogenous variable, number of customers tied to the supply chain by a standing agreement, in order to plan future production. This is known as advanced demand information (ADI) because the system knows in advance some level of demand that it will have, however it must contend with the fact that some orders may be cancelled and some new orders may be generated. Thus, although the data are based on information that is known, the future outcomes of this advanced demand is not always known perfectly, thus ADI is often studied in terms of its usefulness when information is perfect vs. imperfect [Chen 2001, Gayon et al. 2009].

In the medical field a clinical prediction is usually based on endogenous variables. These are the symptoms that the patient has when they present to the practitioner. However it is also possible to mix exogenous variables into the diagnosis. For example, a doctor is much more likely to predict that a patient with a fever, stuffy nose and weakness has the flu during flu season versus during the summer. When considering the operations based predictions in health care, exogenous variables are far more commonly used in order to predict demand and control resource scheduling, examples will be provided in the next sub-section.

The use of endogenous system variables to predict and control future states is done in manufacturing systems which have longer lead times, as in the ADI studies, similar work can be done for health care delivery chains that are on the cross organizational level. However when considering the daily or hourly time frame, ADI and short term forecasting are less popular in supply chain management as these systems are more deterministic for those times frames. In a hospital there is a significant amount of variability, decisions made and actions taken in short time frames. These short term activities have impact on whole system performance, thus there is room for more study into how forecasting can improve performance in these short time frame systems, which may not have direct analogies in supply chain literature. The study presented in the following chapters is one such example. The study does have similarities to the long lead time studies performed with ADI in supply chains in the general approach taken, but also differs somewhat in results due to the volatility and time frame of demand.

The third dimension in prediction, and final dimension described here, is how the prediction is made. The fields of statistics, data mining, machine learning, and artificial intelligence all contain tools for taking data and making conclusions based on the data. While each of these fields have some different tools and methods associated with them, the exact distinctions between the fields are not completely clear [Witten and Frank 2005, Shmueli et al. 2007, Montgomery et al. 2011]. When discussing health care prediction, the most common method is not a part of any of the above fields, expert opinion. Every day medical personnel are making a form of prediction. Medical professionals use historical observations of patients, lists of symptoms learned in text books and gut feelings in order to make predictions about what condition a patient has. Over the course of a patients treatment these predictions are verified or discounted based on the efficacy of treatments or through diagnostic tests/exams. This personal

ability to predict a patient's condition is what contributes to the belief that medicine is a combination of art and science [Gawande 2002].

In training, medical students will be faced with constant differential diagnosis sample cases, these cases teach the students likely outcomes based on sets of symptoms and test results. Over time, the students learn the different conditions that are related to these results, begin to form impressions of the likelihood of each condition, and recognize the further tests needed to distinguish between two conditions with similar symptoms. These intrinsic understandings of likelihood or probability are in some way based on the sheer number of times that an event occurs, however they can also be based on less rational beliefs formed by recent events. The fields of data analysis, mentioned earlier, seek to mimic and improve upon the decision making process that occurs within a practitioners mind. In creating a tool, there is no accidental emphasis placed on recent outcomes or specific symptoms, any method for analyzing the data must be programmed in, and emphasis on specific events only occurs when it is shown to improve the accuracy of future predictions. Just like training a medical student, mathematical methods use historical data in training sets in order to draw probabilistic connections between different symptoms and specific conditions. Similar to clinically based predictions, management based expert opinions are also derived from experiences and are subject to bias. Again mathematical management predictions will be programmed to be bias or unbiased based on performance. There are many different mathematical tools for making predictions; therefore rather than explain each tool on its own, the next sub-section describes different studies that used prediction and the tools that they employed.

2.3.2 Sample prediction based studies in health care

As mentioned earlier, prediction is regularly used in health care when a doctor makes an initial diagnosis of a patient. However as early as the 1950s, studies began to emerge that compared the ability of a statistical predictor to suggest the correct final diagnosis versus an expert [Meehl 1954]. The study of computer/mathematics based diagnosis continues to evolve today and the studies that have been published use a wide variety of tools, such as Neural Networks, Multiple Linear Regression, Bayesian Networks, Fuzzy Logic, etc. [Gustafson et al. 1971, Szolovits et al. 1988, Long 1989, Szolovits 1995, Grove et al. 2000, Patel et al. 2009]. Despite the fact that diagnosis tools are being developed, they are largely used as decision support tools rather than

replacing a clinician and it is unclear how much these tools are used even for support. For example one study implemented a simple probability based prediction rule to identify patients at risk of acute myocardial infarction amongst patients with chest pain, however, it was found that without being coerced, doctors rarely consulted the tool [Pearson et al. 1994]. Another study suggested the use of a Bayesian belief network assist providers in triage to assess the urgency of a patient with greater consistency, however, it was not implemented [Sadeghi et al. 2006, Paul and Sambhoos 2010].

While the use of prediction for clinical applications is interesting, it is less pertinent to health care delivery chain management, except perhaps for using predictions of diagnoses upstream to better staff downstream delivery steps. However, there is a long history in health care of using predictions for operational purposes as well. In 1966, one study proposed using the expert opinions of doctors and nurses about the length of stay of patients in order encourage staff to focus on patients who are expected to be discharged and thereby reduce artificial variability in patient length of stay [Robinson et al. 1966]. In 1968 one study attempted to predict the length of stay of patients in a hospital to assist future planning of elective admissions as well as resource scheduling. This study explored five different methodologies for prediction: expert opinion, multiple linear regression, Bayesian conditional probability, historical means, and direct posterior odds estimation [Gustafson 1968]. Research on predicting discharges and length of stay for the purpose of improving resource allocation and budgeting continues to this day. The methods applied vary from well-established to relatively new, such as: generalized stochastic models [Trivedi 1980], hazard models [Liu et al. 1991], neural networks [Walczak et al. 1998, Walczak et al. 2003, Adams and Wert et al. 2005], Bayesian belief networks [Marshall et al. 2001, Michalowski et al. 2006], multiple linear regression [Omachonu et al. 2004], discrete Markov process [Perez et al. 2006], autoregressive moving average time series analysis (ARIMA) [Rathlev et al. 2007], binary logistic regression [Park et al. 2009], and quantile regression [Ding et al. 2009].

Another significant area of research that uses prediction in hospitals is the prediction of readmissions. This research area seeks to identify patients who are likely to return to the hospital after being discharged recently. Although predicting readmissions can allow a system to better organize its resources and prepare for the admission operationally, the literature tends to make

these predictions for quality purposes. In other words if a readmission can be predicted, steps can be taken just after the patient is discharged to maintain their health and avoid the readmission. Many of the studies are specific to certain disease categories, rather than the entire population. These studies employ a variety of prediction tools such as: logistic regression [Anderson and Steinberg 1985, Boulton et al. 1993, Lyons et al. 1997, Friedmann et al. 1997], proportional hazards regression [Luchansky et al. 2000], and expert opinion [Allaudeen et al. 2011].

There are many other areas where prediction may be used in hospitals and in health systems in general. The examples shown here are specific to hospitals but in any area where there are trends in demand these trends can be used to predict future demand. This can be done in primary care, in flu shot demand, demand for small clinic services, and almost any other health service one can imagine [Smalley 1982].

2.3.3 Prediction in the Emergency Department

There are few consistently studied applications of prediction in health care for operational purposes. While length of stay and readmissions are important topics, they are just the beginning of the many health care delivery chains that exist in a hospital system, or in the entire health system. One area that does attract attention is the ED/IU health care delivery chain. This is likely due to the prominence of this system as described earlier in this chapter.

Clinically speaking, diagnosis tools can be particularly useful in the ED, considering the amount of time it takes to reach a diagnosis is directly linked to quality. Going beyond that, the ESI triage system described earlier takes expert opinion of a patient's urgency and likely resource usage into account [Gilboy et al. 2005]. This in itself is a mix of a clinical and operational prediction. The clinical benefits of ESI triage are achieved by knowing which patients need to be treated first. The operational benefits come from how the ESI levels are used to assign staff to a patient. Although it is not a formalized process, nursing managers will often use the ESI levels to manage the work load of their nurses by giving a mix of acute and less acute patients. It was mentioned earlier that studies have sought to improve triage accuracy through statistical decision support using Bayesian networks. These studies were framed in the clinical context however they also provide value by creating consistent triage performance, which can be used to formalize processes for optimally assigning patients to medical staff in the ED [Sadeghi et al. 2006, Paul and Sambhoos 2010].

Other operational predictions have been suggested in the ED based on the desire to prepare for crowding. Crowding indexes discussed earlier are not just based on the amount of patients that are in the ED but they are also based on the acuity and resource demands of the patients in the ED. This means that the indexes can be used or modified to create work indexes. An ED with a high work index level is likely to become delayed and therefore this index can predict future crowding which can then be used by management to take preventative actions [Epstein and Tian 2006]. While crowding measures use endogenous variables to predict crowding it is also possible to use a mix of endogenous variables and exogenous variables. For example one study made an ARIMA model to directly predict long term crowding trends [Schweigler et al 2009]. Another study used a discrete event simulation that takes a current ED census and service rates into account but then uses exogenously based, expected arrival rates to predict conditions into the future [Hoot et al. 2008], similar work has been done using a Markov chain [Au et al. 2009], another study used a simulation method called Petri-nets to study the flows in the ED which lead to a crowded condition [Chockalingham et al 2010].

An exogenous variable built into many of the tools that predict crowding is the forecast of patient arrival. The crowding prediction tools use the forecast to influence short term resource decisions, but the forecast itself is often also used for long term resource and staff scheduling. This has been done using many different models, with varied success, over different time periods. Some examples of methods used to make long term forecasts of ED demand are: ARIMA [Jones et al. 2002], Poisson regression [McCarthy et al. 2008], multivariate time series [Jones et al. 2009], review of many methods [Wargon et al. 2009], and a general linear model [Wargon et al. 2010]. In one case, binary logistic regression was applied and was combined with a practical application for avoiding surge based overcrowding by comparing the forecast to staff scheduling [Chase et al. 2012]. Another study made the translation between the ED and the IU and made an ARIMA model that predicts long term IU demand from the ED based on future ED demand [Abraham et al. 2009].

The above studies forecasting ED arrivals over time are useful for long term planning of generalized capacity and scheduling, however they do not allow for planning reactions to daily surges. To this end it is worth returning to the concept that initiated the prediction discussion: predicting the likelihood of individual patients being admitted to the IU from the ED. This

involves making predictions of short term demand based on the endogenous variables associated with each individual patient, and can be seen as a type of ADI in the ED/IU system. Some attempts at individual prediction have had clinical objectives. These studies focused on predicting whether specific categories of patients will be admitted: Neural network for children presenting with bronchiolitis [Walsh et al. 2004], Expert opinion on patients with acute coronary symptoms [Arslanian-Engoren 2004], Expert opinion on patients arriving by ambulance [Levine et al. 2006, Clesham et al. 2008].

Some recent studies have seen the operational benefits to making predictions and focused on the entire ED population. These studies used the following methods: Bayesian network [Leegon et al. 2005], support vector machines, naïve Bayes [Li et al. 2009], and logistic regression [Sun et al. 2011]. The studies are valuable for developing prediction models; however they fall short from describing how the models could be used in a practical manner to improve flow. Only one of the studies looks at multiple methods for making the prediction. While the study shows that one method performs a little bit better than others, by not addressing the practical value of the predictions, it is unclear whether it is worth investing in more complicated methods to achieve more accurate predictions [Li et al. 2009]. Another of the studies uses a relatively simple method but uses patient variables that may not commonly be available in hospitals when they are not part of a nationalized health care system [Sun et al. 2011]. In each of the studies, there is a reliance on the historic conclusion that predictions should be useful but they do not explore how the models can be applied.

While demonstrating the ability to predict IU admission from the ED is a contribution towards studying the ED/IU delivery chain, it falls short of showing how prediction can be used to manage the chain. Similarly, these studies do not fulfill both parts of anticipation of need as suggested by the IOM: making the prediction and acting on the prediction. The rest of this dissertation will focus on applying simple methods of predicting admission to the IU from the ED with a focus on how to apply these predictions to meaningfully influence decisions in the IU and improve key metrics in the ED/IU delivery chain. This will begin with the exploration of three prediction methods in Chapter 3 leading to the selection of linear regression as a high performing prediction method in this context. This will be followed, in Chapter 4, by a study to expand the practicality of the findings of Chapter 3 by showing that the regression method and

variables it used are generalizable to other hospitals. Having shown the potential for the regression to predict admission in many hospitals Chapter 5 will describe an implementation study that seeks to understand the more practical aspects of using prediction in the ED/IU delivery chain. Finally Chapter 6 will apply discrete event simulation towards showing the potential for prediction to improve flow in a controlled hospital environment.

Chapter 3: Predicting Emergency Department Admissions¹

As discussed in Chapter 2, Emergency Department (ED) crowding is a major problem nationally and occurs when there is a mismatch between the demand and supply of the resources needed to evaluate, treat and discharge patients from the ED. Resource constraints may be related to resources controlled within the ED such as nurse and provider staffing, or from resource constraints external to the ED such as the availability of support services capacity or the availability of open inpatient beds. In Chapter 2 it was described that availability of inpatient beds to receive ED patients is arguably “the single most important factor” related to ED flow problems [Asplin et al. 2003, US GAO 2003, Olshaker and Rathlev 2006, Williams 2006, Falvo et al. 2007, Hoot and Aronsky 2008, US GAO 2009].

Organizational solutions to address this problem can be categorized as static ones such as “discharge by noon” procedures and dynamic ones that are activated based on specific situations within the ED. Examples of dynamic solutions include placing boarding patients in inpatient unit (IU) hallways, encouraging IU staff to schedule discharges to match historical patterns of expected admissions, and activation of inpatient resources based on the level of ED crowding [Rubino et al. 2007, ACEP 2008, Viccellio et al. 2009].

As was suggested in Section 2.2.3 flow in the ED may be improved by estimating the likely number of patients who will be admitted at a point in the near future and sharing this information with IU staff who may then mobilize resources before crowding becomes an issue. Most studies that predict admission have focused on either the entire ED population [Li et al. 2009, Sun et al. 2011] or specific categories of patients [Arslanian-Engoren 2004, Walsh et al. 2004, Levine et al. 2006, Clesham et al. 2008] and treat admissions as binary in the sense of estimating “yes” or “no” at the patient level. This approach may be less useful when the goal is to predict aggregate demand. Pooling of patient admission probabilities across all ED patients should theoretically provide more precise predictions of near future aggregate demand for inpatient beds [Hopp and Spearman 2001]. The primary objective of this chapter is to describe and evaluate three simple

¹ The majority of the material in this chapter has been previously published as [Peck et al. 2012]

methods for generating admission predictions based on patient characteristics, available at the time of patient triage. The secondary objective of this chapter is to introduce a new method for using predictive information by aggregating the individual patient predictions into a summative measure of near future IU bed demand, rather than sharing single patient predictions.

3.1 Methods

3.1.1 Study Design

Three methods to predict IU admission at the time of ED triage were developed and tested: expert opinion, naïve Bayes conditional probability and generalized linear regression with a logit link function (logit-linear regression). Retrospective patient visit data was collected to form two datasets. Statistical models were created using a development dataset. To avoid overestimation of model performance due to over fitting, the models were assessed using a separate validation dataset and final logit-linear regression and naïve Bayes models were identified. A third test dataset was developed during a study of triage nurse expert opinion predictions. The performance of the final statistical models was then assessed on this test dataset, which allowed for direct comparison of expert opinion and the two statistical models.

All portions of this study were approved by the Institutional Review Board of VHA BHS. All analysis was performed using The MATLAB (R2011b-7.13.0.564), MathWorks, Inc., and Microsoft Office Excel 2007.

3.1.2 Study Setting

This study took place at VHA West Roxbury. VHA West Roxbury is a federal tertiary care, referral hospital devoted to the care of the USA veteran population. It serves both the local community and acts as a referral hospital for the six other VA Medical Centers in the New England region. The hospital is affiliated with two medical schools and has house staff from affiliated programs. VHA West Roxbury has a 13 bed ED with an annual volume in 2010 of 12,672 visits; there are six inpatient wards and four specialty care units comprising approximately 170-180 staffed beds. The hospital receives a capitation based budget but can also receive compensation from private insurers. The ED receives local ambulances carrying patients from the surrounding communities and will accept all patients whether they are veterans or not.

This study was performed in partnership with the staff of the VHA West Roxbury ED, including the direct involvement of the ED director Dr. Stephan Gaehde and the Nurse Manager John Marinello. The project was also performed with the consistent feedback and support of the entire ED staff. The enthusiasm for improvement that is demonstrated by the participation of the staff is also demonstrated by the consistently high performance of the ED.

This performance can be seen through key metrics in the ED. Figure 17 shows the how the volume of ED visits has grown from 2006 to 2011, yet at the same time the waiting time to be seen has decreased.

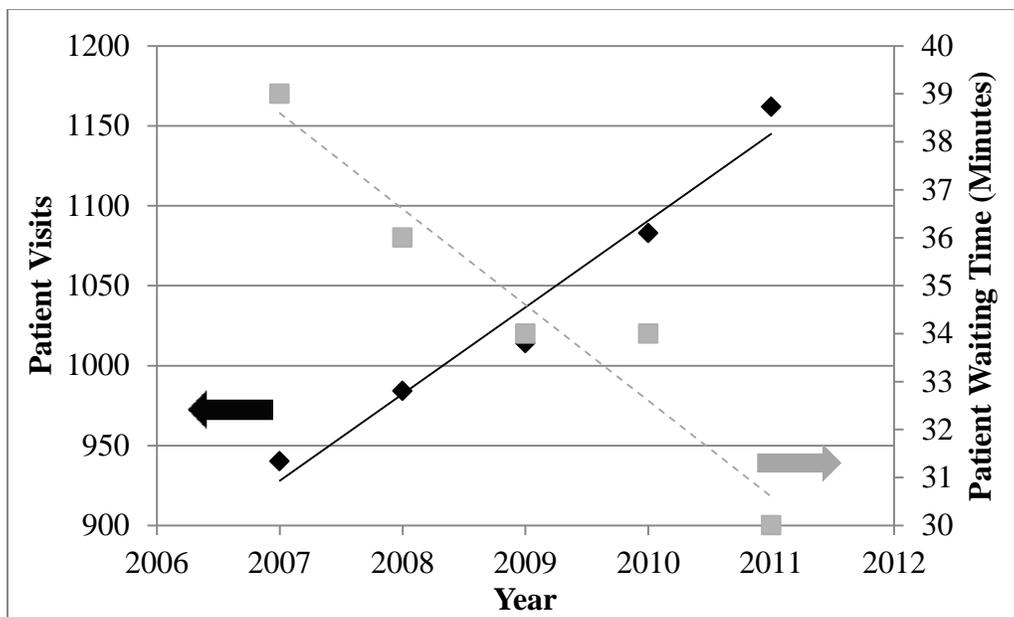


Figure 17 Patient visits to VHA West Roxbury ED by year

Figure 18 shows ED visits broken down by month, exposing a seasonal pattern in arrivals, however Figure 19 shows that there was no seasonal trend seen in waiting times. This suggests that ED resources may not be working at a utilization level that would lead to visible sensitivity to fluctuations in arrivals. The explanation for continued improvements in performance may originate from within the ED itself. It may also come from flow improvements from the inpatient unit, or (most likely) some combination of the two. Nevertheless, the ED staff continues to work towards improved performance using proven techniques such as those discussed in Chapter 2, as well as new possibilities such as prediction.

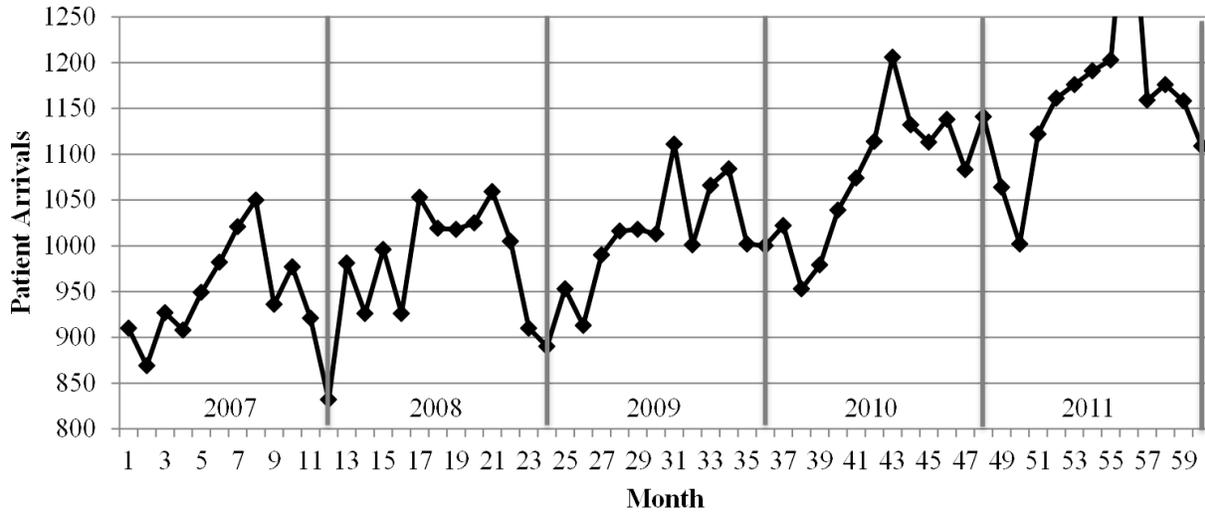


Figure 18 Monthly patient arrival pattern

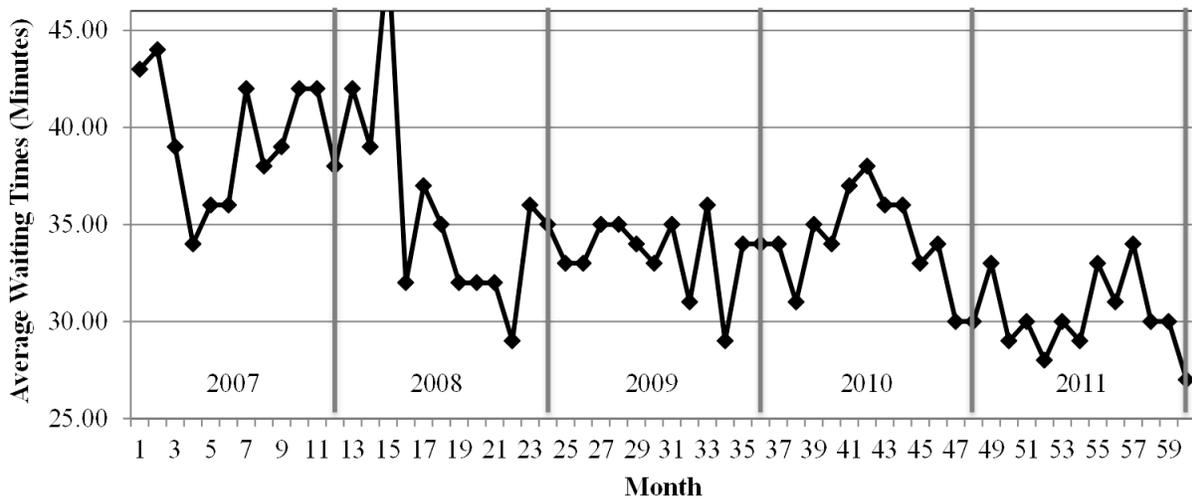


Figure 19 Monthly average waiting time pattern

3.1.3 Study Protocol

The expert opinion portion of the study was conducted from September 22, 2010 to November 26, 2010 between the hours of 7am and 5pm. During the study hours, separate triage rooms were operational. Due to lower patient volume, between 5pm and 7am, patients are sent directly to ED beds and are triaged at the bedside, consequently, bypassing the expert opinion process which took place at the triage stations. During the study period 1160 patients entered the ED, 767 of

the 1160 patients were triaged in the triage stations. Triage nurses classified each patient's likelihood of admission, using their expert judgment, into one of 6 categories (Figure 20).

How likely is it that the patient will need admission to the hospital?
<input type="checkbox"/> Definitely Yes (95-100%)
<input type="checkbox"/> Highly Likely (75-94%)
<input type="checkbox"/> Likely (50-74%)
<input type="checkbox"/> Unlikely (25-49%)
<input type="checkbox"/> Highly Unlikely (5-24%)
<input type="checkbox"/> Definitely No (0-4%)

Figure 20 Expert opinion triage questionnaire

Nursing staff were treated as an IRB defined vulnerable population and their predictions were not shared with any other ED staff or supervisors. Triage nurses were blinded to the specific purpose of the study but were aware that it was being conducted to improve ED patient flow.

Structured interviews conducted with each of the triage nurses identified 6 possible patient characteristics available at the time of triage for possible inclusion in the predictive model. These were:

- Patient Age: Continuous range of values
- Primary Complaint: Free text entered by triage nurse
- ED Provider: Provider assigned to the patient
- Designation: Fast track or standard ED bed
- Arrival Mode: Stretcher, Wheelchair, or Ambulatory
- Urgency (ESI) Level: 1, 2, 3, 4, 5

For the development of the statistical models, retrospective triage data on all patients who entered the ED was collected from January 1, 2010 to May 6, 2010 totaling 4187 patient visits. Using this development dataset, ANOVA analysis was performed for each of the selected factors identified from the expert opinion study and they were found to be significantly associated with hospital admission. A validation dataset for model assessment and selection of the optimal mode was composed of ED visits between May 7, 2010 to May 31, 2010 and September 1, 2010 to

September 21, 2010, totaling 1614 patient visits. Table 1 summarizes basic patient characteristics for the patients included in the development dataset, validation dataset and test dataset.

Table 1 Basic patient characteristics between development, validation and test datasets

Patient Counts:	Development Dataset (24 Hours)	Validation Dataset (24 Hours)	Test Dataset (7am-5pm)
Urgency			
1	7	1	1
2	56	26	7
3	2441	892	585
4	1347	388	318
5	336	302	249
Arrival Mode			
Ambulatory	2844	1139	860
Stretcher	895	308	156
Wheelchair	448	167	144
Age			
10-19	1	0	0
20-29	248	103	37
30-39	196	70	55
40-49	311	123	66
50-59	697	309	215
60-69	1052	403	315
70-79	770	293	209
80-89	779	278	228
>90	133	35	35
Sex			
Female	200	78	53
Male	3987	1536	1107

The statistical approaches make use of event probabilities and conditional probabilities which require categorical data. Age was categorized into decades. Primary complaint was coded using a previously established system slightly modified to remove the free text options, resulting in 62 complaint categories [Aronsky et al. 2001]. All other factors were already categorical. Table 2 lists examples of categories from each factor and their corresponding empirical probabilities estimated from the data, where $P(X)$ means the unconditional probability of event X (used as the independent variable values in the logit-linear regression models) and $P(X|Y)$ is the conditional

probability of event X given that event Y has occurred (used in the naïve Bayes models). For example, reading from the fifth row, historically 10.69% (or 0.1069) of all patients (admitted to an IU or not) arrive by wheelchair, whereas 15.49% (0.1549) of all admitted patients arrive by wheelchair and 49.11% (0.4911) of those patients arriving by wheelchair were admitted. The complete table can be found in Appendix A.

Table 2 Factors tested for admission prediction ability and the empirical probabilities of occurrence

Factor/code	Probability of code	Probability of code given admit	Probability of admit given code
Designation	P(Designation)	P(Designation Admit)	P(Admit Designation)
ER	0.6237	0.9888	0.5383
Fast Track	0.3763	0.0112	0.0101
Arrival Mode	P(Mode)	P(Mode Admit)	P(Admit Mode)
Ambulatory	0.6793	0.4035	0.2014
Stretcher	0.2138	0.4415	0.7006
Wheelchair	0.1069	0.1549	0.4911
Urgency Level	P(Urgency)	P(Urgency Admit)	P(Admit Urgency)
1	0.0017	0.0042	0.8571
2	0.0134	0.0211	0.5357
3	0.5830	0.9415	0.5477
4	0.3217	0.0267	0.0282
5	0.0802	0.0063	0.0268
Patient Age	P(Age)	P(Age Admit)	P(Admit Age)
10-19	0.0002	0.0000	0.0000
20-29	0.0592	0.0070	0.0403
30-39	0.0468	0.0134	0.0969
40-49	0.0743	0.0479	0.2186
50-59	0.1665	0.1606	0.3271
60-69	0.2513	0.2690	0.3631
70-79	0.1839	0.2085	0.3844
80-89	0.1861	0.2458	0.4480
90-99	0.0318	0.0479	0.5113
Provider	P(Provider)	P(Provider Admit)	P(Admit Provider)
1	0.0262	0.0126	0.1636
2	0.0160	0.0134	0.2836
3	0.0105	0.0112	0.3636
4	0.0086	0.0112	0.4444
5	0.1150	0.1019	0.3008
...
Primary Complaint	P(Complaint)	P(Complaint Admit)	P(Admit Complaint)
Abdominal pain	0.0480	0.0685	0.4850
Abdominal	0.0504	0.0749	0.5048

problems			
Abnormal Labs	0.0134	0.0275	0.6964
...
Cardiac arrest	0.0065	0.0148	0.7778
Cardio-vascular complaint	0.0310	0.0516	0.5659
Chest pain	0.0480	0.0862	0.6100
Cold/Flu	0.0595	0.0106	0.0605
...
Fainting/syncope	0.0074	0.0155	0.7097
Fall	0.0250	0.0311	0.4231
Fever	0.0158	0.0297	0.6364
...
Joint Problems	0.0353	0.0056	0.0544
Kidney and Liver Failure	0.0151	0.0219	0.4921
Laceration	0.0096	0.0035	0.1250
Medication refill	0.0247	0.0000	0.0000
...
Psychiatric/social problems	0.0429	0.0523	0.4134
Respiratory problems	0.0909	0.1801	0.6728
Skin complaint/trauma	0.0420	0.0162	0.1314
Total Probability of Admit	0.3395		

A naïve Bayesian model and a logit-linear regression model were then created for each of the 63 possible combinations of the 6 identified factors. For instance, a naïve Bayes Model and a logit-linear regression model were created for the case where just patient age is used as a predictive factor, then patient age and primary complaint are used, then patient age, primary complaint and mode of arrival are used, etc. These models are then applied to the validation dataset in order to evaluate their performance. A final model that has a balanced performance in each measure was then selected and applied to the test dataset enabling the comparison of predictions for each of the three prediction methods.

3.1.4 Data Analysis/Measures

Each of the logit-linear regression and naïve Bayes models were constructed using the development dataset of 4187 historical patient points and evaluated for predictive ability using the 1614 patient points that were included in the validation dataset.

To illustrate how the naïve Bayesian method works [Witten and Frank 2005, Shmueli et al. 2007] given three hypothetical factors “F1”, “F2”, and “F3”, the admission probability for any particular patient is estimated as

$$P[\text{Admit} | F1, F2, F3] = \frac{P[F1 | \text{Admit}] * P[F2 | \text{Admit}] * P[F3 | \text{Admit}] * P[\text{Admit}]}{P[F1] * P[F2] * P[F3]}$$

If the model is a combination of patient age and complaint the equation would only use those two factors, if the model is the combination of all six of the identified triage factors the equation would use all six factors. The data for each factor is calculated using the development dataset and a sample is displayed in Table 2. The naïve Bayes models were calculated using Microsoft Excel.

The logit-linear regression method that was employed uses the conventional log-odds link function and is calculated as

$$\text{Log}(P[\text{Admit}]/1-P[\text{Admit}]) = \beta_0 + \beta_1 * P[\text{Admit}|F1] + \beta_2 * P[\text{Admit}|F2] + \beta_3 * P[\text{Admit}|F3]$$

The admission probability then is estimated via the inverse logit as

$$P[\text{Admit} | F1, F2, F3] = \frac{e^{\beta_0 + \beta_1 * P[\text{Admit}|F1] + \beta_2 * P[\text{Admit}|F2] + \beta_3 * P[\text{Admit}|F3]}}{1 + e^{\beta_0 + \beta_1 * P[\text{Admit}|F1] + \beta_2 * P[\text{Admit}|F2] + \beta_3 * P[\text{Admit}|F3]}}$$

The size of the β coefficients represent the amount of influence of each factor has on admission probability. Both methods can be calculated in standard spreadsheet or statistical software. The logit-linear regression models were calculated using the statistical package built into Matlab.

As described in the introduction, other published models for making admission predictions in the ED seek to assign a yes/no value to the patient. This use of predictions can indeed facilitate early admission coordination, by simply placing the admission order sooner. One common method to evaluate a prediction model that has the goal of suggesting a yes or no prediction is the receiver

operating characteristic (ROC) plot's area under the curve (AUC). This value was calculated for each model and allows the user to calibrate the model to reduce false orders.

A qualitative method was also applied for evaluating model accuracy by categorizing patients into probability groups and judging whether the model accurately categorizes patients. For instance if 20-30% of patients assigned an admission probability in the 20-30% range are actually admitted, then the model is seen as accurate in that range.

Binary prediction of admission increases estimation error by forcing the computed probability of admission from a fractional value to 1 or 0. While this may be useful strategy for early communication of likely IU admission for an individual patient, it increases estimation error when the predictions are summed across a group of patients to provide an estimate of aggregate near-future IU bed demand. Instead an ED can maintain an aggregate measure of future bed demand based on the summation of raw probabilities.

This 'running bed demand' can be calculated using any method that generates an admission probability, such as those applied in this study. The resultant probabilities are totaled across all patients currently in the ED as shown in Figure 21 to produce a total momentary predicted bed demand. For example, given n ED patients each with IU admission probabilities of $p_1, p_2,$ and so on, the estimated total number of admissions, $E(T)$, to expect is $E(T) = p_1 + p_2 + \dots + p_n$. Since the actual number may be higher or lower, the standard deviation of total admissions, $\sigma(T)$ can be estimated as $\sigma(T) = \sqrt{(p_1(1-p_1) + p_2(1-p_2) + \dots + p_n(1-p_n))}$ and, for more advanced applications, this can be used to generate confidence bounds on the number of predictions. Using these calculations, at any moment of a day, bed demand information can be compared with hospital wide availability and appropriate actions taken.

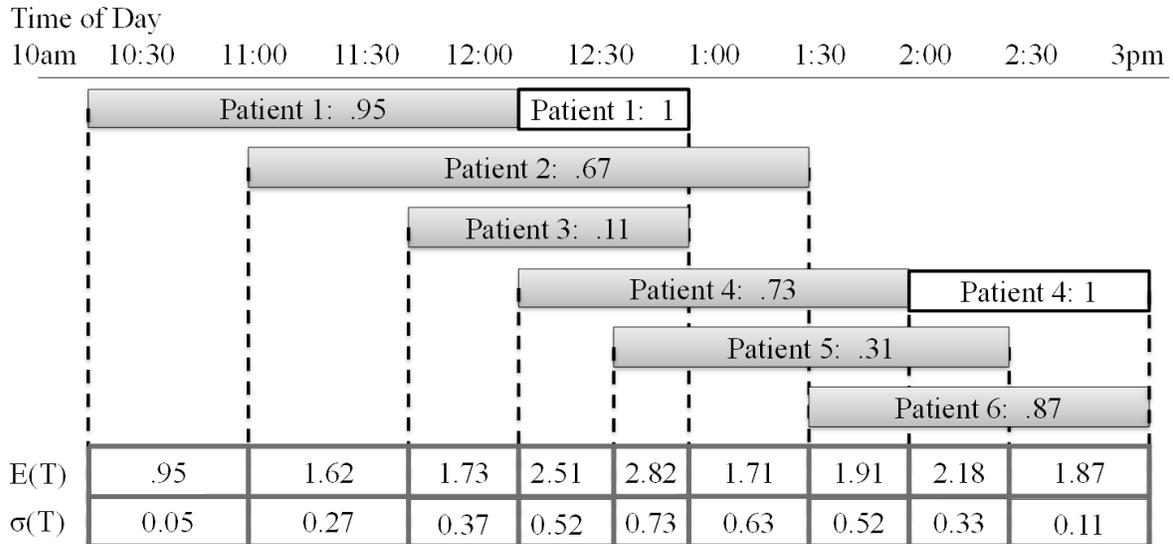


Figure 21 Conceptual illustration of real time bed demand forecast (Running expected number of admissions and standard deviation)

Two methods are employed in order to evaluate how accurately each model generates the running bed demand. The first method is to use visual inspection. Over the course of the day, the running bed demand and the cumulative admissions are plotted side by side and it can be seen whether the two are well correlated. The peak value of the bed demand for each day can then be compared to the peak number of admissions to see how well informed the IU staff were when they received this value.

Another mathematical way to assess model performance at generating the running bed demand is to simply add all of the predictions for each day and to compare these predictions to the actual number of patients that were admitted each day by generating an R^2 correlation value. Noting that an R^2 correlation does not reflect errors in magnitude, this can be combined with a study of model residuals to achieve a better understanding of how well the model aggregates predictions. None of the methods described above are perfect evaluators on their own; however in combination they provide a good sense of how well the model performs.

3.2 Results

The 63 naïve Bayes and 63 logit-linear regression models created with the development dataset were applied to the validation dataset. Using, ROC AUC, R^2 , residual analysis, and goodness of

fit into prediction categories, final models were selected for application to the test dataset. Although multiple models performed well in some evaluative measures, a few performed consistently well in all. Consequently the final models chosen are not the only options but provide a basis for comparing methodologies and a sense of model potential. It is likely that the unique traits of a hospital exploring this methodology will influence the weighting of factors that emerge as better predictors and the chosen model for that specific hospital.

When applied to the validation dataset, the logit-linear regression model that performed consistently high in all analyses (and highest in some) comprised of patient age, primary complaint, designation, and mode of arrival. In contrast the naïve Bayes methodology incurred many tradeoffs and the final model was chosen for consistent high performance in all categories, though it was the best in none. When applied to the test dataset, these final models had AUCs of 0.841 and 0.887 for the naïve Bayes model and the logit-linear regression model respectively. In contrast, the worst performing models were those that just used the ED provider as the predictive factor with an AUC of 0.5 for both the naïve Bayes and logit-linear regression versions of the model.

Figure 22 compares how well the triage nurse predictions, the final naïve Bayes model, and the final logit-linear regression model assign patients into admission probability categories, using the test dataset. For the latter two cases' probabilities, which were continuously assigned by the models, were grouped into the same ranges used in the expert opinion for comparison; e.g. all patients assigned a probability between 0-4% were put in the "definitely no" category. As shown, the logit-linear regression results best fit the mid-point of each category in all but the "definitely no" tail (where naïve Bayes appears better), whereas expert opinion significantly under-estimates admission in all categories.

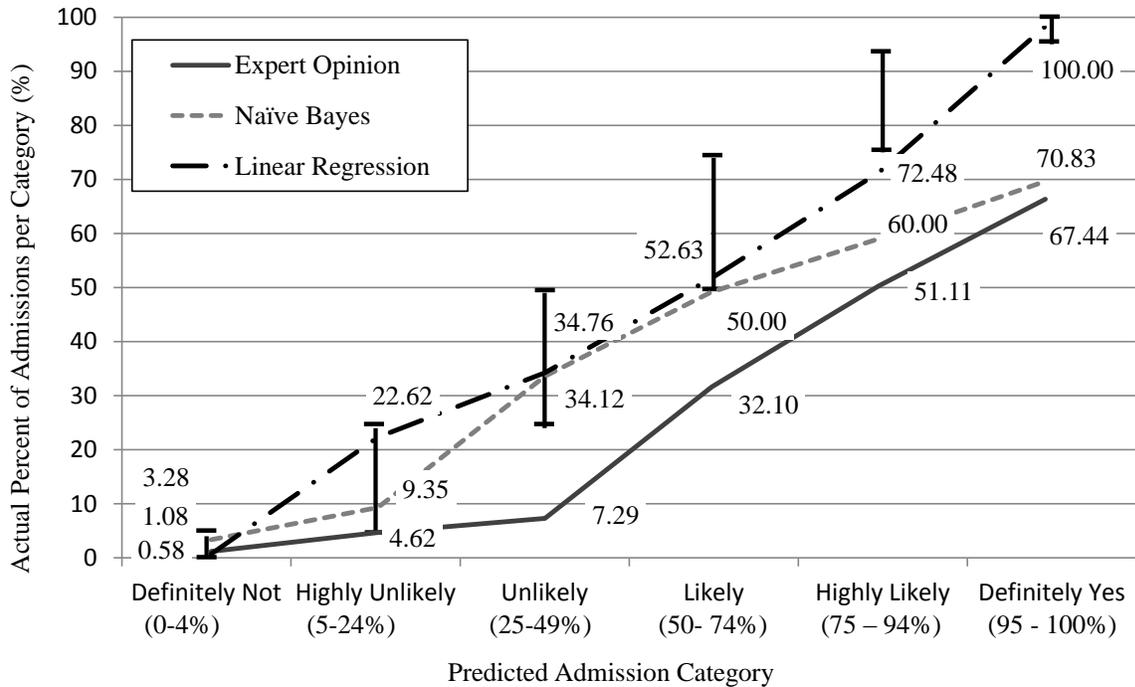


Figure 22 Categorized predictions of patient admissions versus percent of patients admitted from each category

Table 3 summarizes the significant factors for the best performing logit-linear regression model. Note that all listed primary factors (patient age, primary complaint, designation, arrival mode) are highly significant statistically ($p \ll .0001$ in all cases).

Table 3 Model parameters for best fitting logit-linear regression

Factor or Interaction	Coefficient (β)	Significance (p -value)
Constant	-7.02	1.4e-56
Designation (fast track or not)	5.48	1.0e-21
Primary Complaint	2.89	5.3e-24
Patient Age	3.39	1.9e-05
Mode of Arrival	2.69	2.8e-21

Figure 23 compares continuous actual versus predicted bed census (as described by Figure 21) for 15 days using expert predictions (top), logit-linear regression, and naïve Bayes (bottom). This data was generated by breaking up the test dataset into hourly ED census. For each hour, the model predictions of admission probability for each patient were added together, with the probability of boarding patients taken to be a 1. For expert opinion, admission probability was

taken as the midpoint for each category (ie. 84.5% for patients in the highly likely category) and based only on patients who physically went through triage as opposed to all patients in the ED leading to reduced numbers in the chart. As shown, the logit-linear regression method appears to match actual admit volumes most accurately, with all three methods providing several hours advance notice. Over all 2 months of data used in the validation dataset, the difference between predicted peak bed demand and actual demand for the expert opinion, naïve Bayes, and logit-linear regression methods on average were 0.82, 0.69, and -0.26, respectively (with standard deviations of .93, 1.81, and 1.59). These predicted peaks occurred on average 3.0, 3.7, and 3.52 hours before the actual peaks, respectively (with standard deviations 1.96, 2.20, and 1.96).

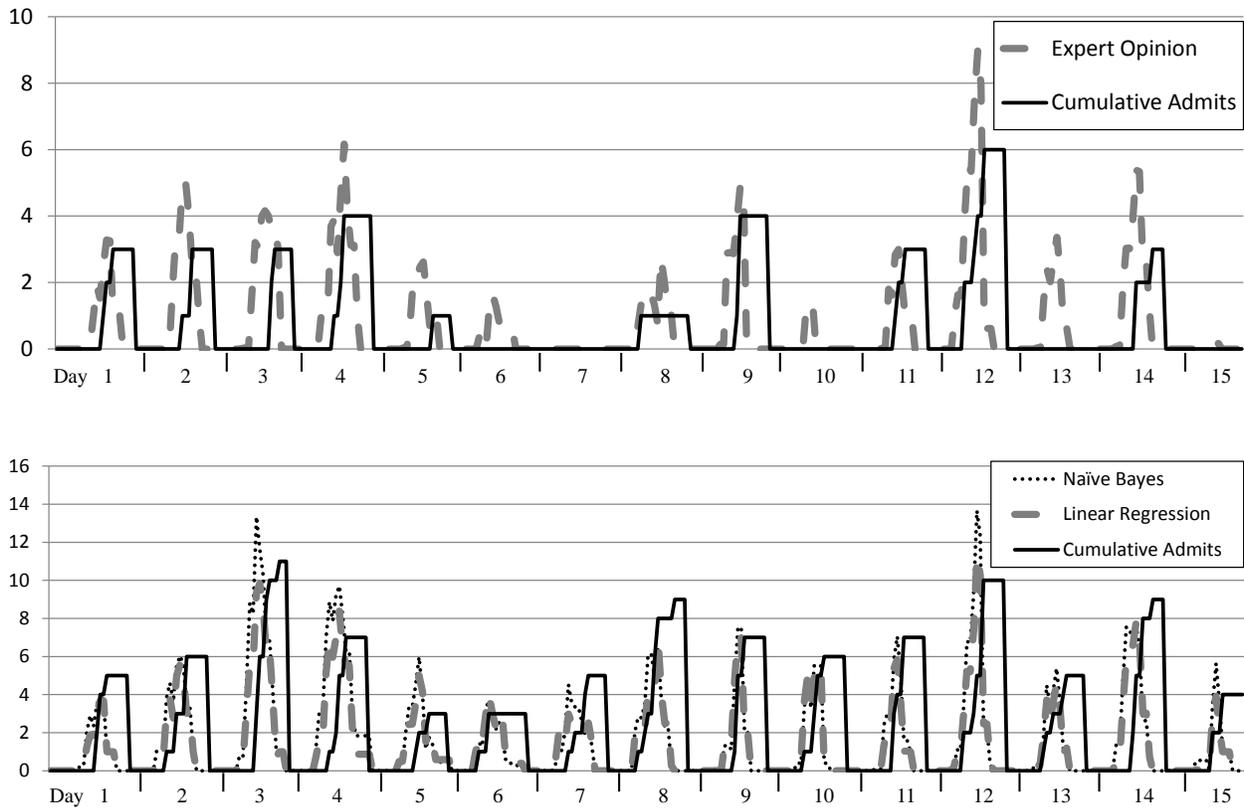


Figure 23 Real time, statistically predicted-expected and actual number of cumulative admissions

Figure 24 compares actual and predicted total daily admissions, for the test dataset, using expert opinion, naïve Bayes, and logit-linear regression respectively. The R^2 value for the logit-linear regression is the greatest at 0.5826 followed by 0.5775 for the naïve Bayes model, and 0.5243 for expert opinion. None of the methods perform well at predicting small admission volumes (since ideal fits would pass close to the origin as demonstrated by the horizontal line in each figure).

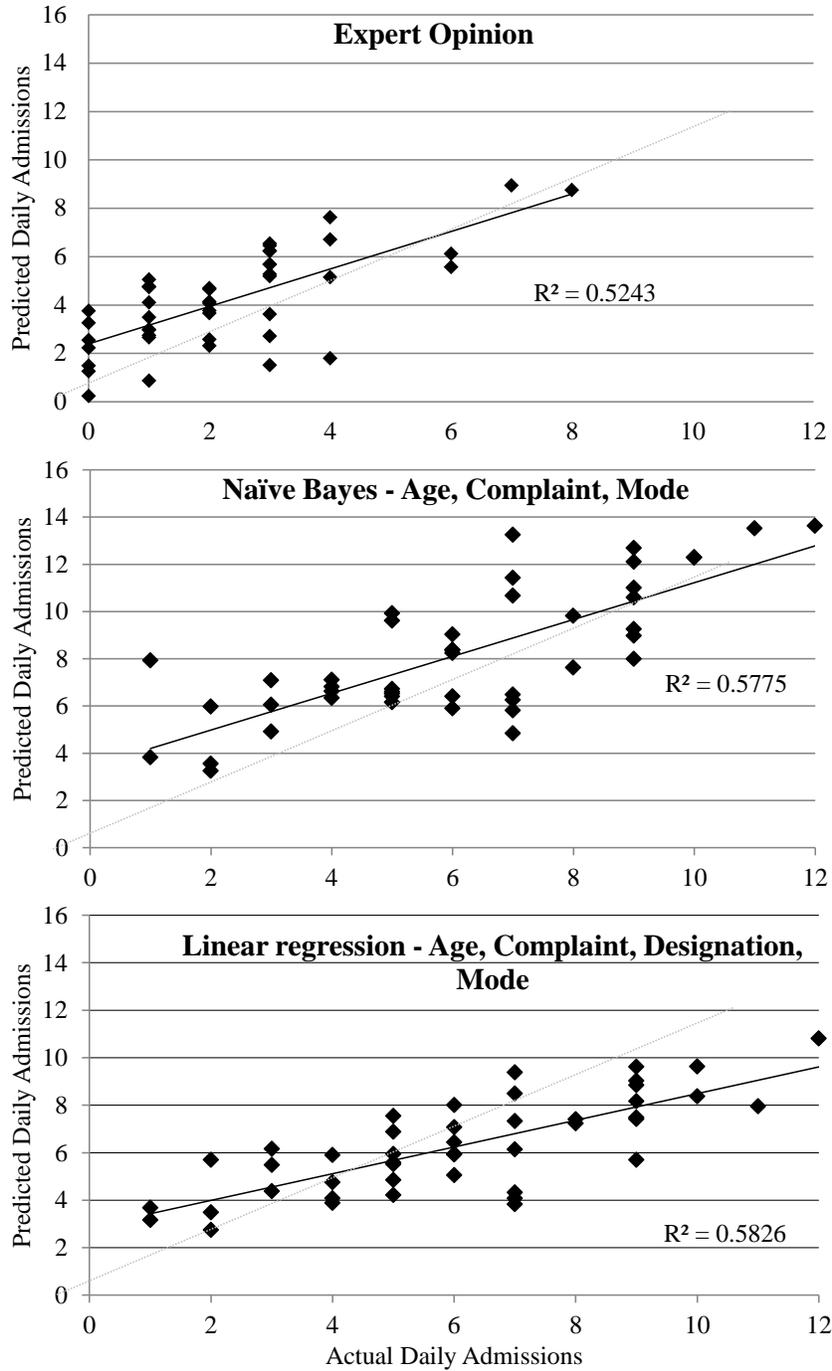


Figure 24 Correlation between predicted admissions and actual admissions based on expert opinion (top), naïve Bayes (middle), and logit-linear regression (bottom) approaches

As mentioned earlier, R^2 is a measure of how well the prediction trend follows the actual trend. On its own R^2 does not prove the accuracy of a model, it is therefore valuable to analyze the residuals of the models. Figure 25 illustrates the residuals (predicted minus actual) for each

model. The residuals expose the tendency of each method to over-predict to some degree. Logit-linear regression appears to perform the best having consistent performance while the other models seem to increase in error as predicted values increase.

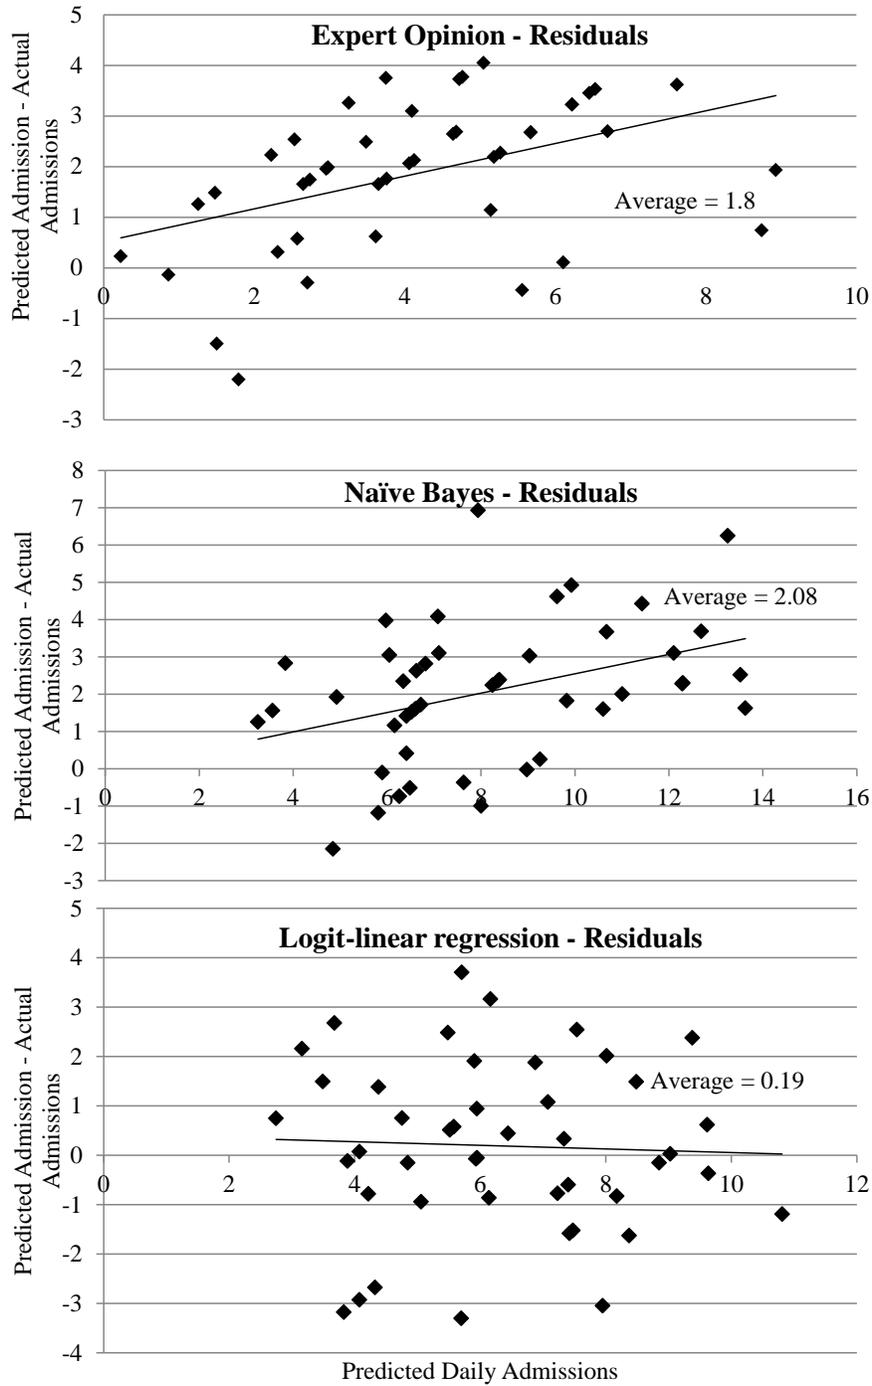


Figure 25 Prediction residuals (predicted minus actual) based on expert opinion (top), naïve Bayes (middle), and logit-linear regression (bottom) approaches

3.3 Discussion

As discussed in Chapter 2, timely patient ED discharge to IU remains a major contributor to ED crowding. Most prior studies have focused on predicting individual patient admission or have focused on methods to predict longer term admission trends. Common approaches, for example, focus on resource planning and staffing for future days [Tandberg and Qualls 1994, Jones et al. 2002, Jones et al. 2008, Abraham et al. 2009], predict short term ED visit surges [Hoot et al. 2009, Schweigler et al. 2009] or use ED crowding indexes to predict ED congestion in the near future [Bernstein et al. 2003, Weiss et al. 2004, Epstein and Tian 2006]. While forecasting can help set a baseline staff level, these forecasts are not based on same-day demand and therefore do not sufficiently inform real-time bed management and encourage behavior based on immediate, direct incentives. Alternatively, while predictors of short term ED demand surges or measures of crowding may inform hospital staff and increase the pace of work and sense of urgency, these measures do not necessarily translate to high IU demand/admissions, and therefore can mislead IU staff who choose their actions based on these measures.

In response to this, one suggestion to improve ED-to-IU flow is to predict admission demand when patients arrive to the ED [Yen and Gorelick 2007]. In contrast, to ED crowding measures, admission predictions are a more direct measure of incoming IU demand and can be used to more accurately inform the actions of IU staff. This chapter described a method to aggregate individual patient admissions predictions into a summative measure of near future IU bed demand which may be useful for informing hospital wide decisions on a daily, real-time basis.

The three prediction methods discussed in this chapter are fairly easy to implement, with logit-linear regression being the most accurate in the test setting, followed by the naïve Bayes approach. ROC curve results suggest that these models could be used as part of the current work flow, where orders for admission are made for specific patients using predictions rather than waiting for the final provider's order. However, all three methods that have been explored in this chapter also enable risk pooling of individual admission probabilities and thus may be more accurate at the aggregate level than methods that dichotomously classify each patient as “admit” or “not admit”. (More complex approaches - e.g. Bayesian belief networks, neural networks, others – also tend to fall in this latter category.) For example, three ED patients each with a 45% IU admission probability might each be classified as “no admit” by such a method, spurring no

action, whereas expected admissions under the proposed approach is $0.45 + 0.45 + 0.45 = 1.35$ with a standard deviation of 0.86, suggesting the IU probably should open at least one bed and perhaps as many as three (using the mean plus 2 standard deviations).

The results in Figure 23 also suggests that predicted admission information can allow bed managers to start planning for peak demand significantly earlier than what currently occurs. Other benefits may result from sharing these data hospital-wide, such as allowing medical staff to better prioritize clinical activities, discharge ready patients in timelier manners, or manage bed preparations and room assignments for specific kinds of patients.

A practical question concerns how many beds to prepare relative to the expected demand, standard deviation, and likely range. That is, if a demand for 8.7 beds is predicted with a 95% interval ranging from 5.8 to 11.6 beds, it is not clear if a bed manager should plan for 6, 9, 12, or some other number of beds. This decision might be based on the relative costs of being over versus under prepared. This decision also may evolve as a day progresses and knowledge is gained as to which early ED patients in fact were admitted. Admission likelihood estimates also could be updated during a patient's ED visit, such as based on test results, doctor evaluations, and changes in physiologic status. Additionally, it could be useful to predict each ED patient's length of stay in order to better estimate IU bed demand timing (e.g. estimated ED-arrival-to-ED-discharge time) over the course of each day.

The final models that were chosen in this chapter may give rise to questions of face validity, given that many would consider patient urgency/triage level as the likely candidate for best predictor. Although models that used this factor did perform well, they may not be the best performers because age and primary complaint and mode of arrival (which were in both of the final models) strongly influence ESI level and therefore it could be acting as a surrogate variable that may then include other less predictive parts such as predicted resource usage. Another possible explanation is that the study site may not assign ESI levels the same way as other hospitals where ESI data would be more predictive.

3.4 Limitations

While the methods described above are simple and effective, a few limitations exist. The simplicity of the models allows for a reduction in the data requirements necessary to achieve

useful results. This makes the models and methods easily implemented by hospital staff with limited knowledge and software. However this simplicity may also lead to reduced performance compared to more complicated models, such as those used in some of the other studies discussed in Chapter 2 [Witten and Frank 2005].

Predictive models only remain accurate if the underlying behavior of the system being modeled remains stationary. Therefore models may need to be recalibrated when there are substantial shifts in admission patterns. For instance such change may occur due to introduction of more effective treatment methods, treatments that shift care from the inpatient to the outpatient setting, changes in insurance practices, or payment structures. Similarly the methodology for applying ESI in another site may lead to it becoming a more (or less) predictive factor as described in the discussion. This will be explored further in Chapter 4.

From an implementation perspective, both probability methods require initial effort to develop a coded dataset, including coding primary complaints, and to calculate the probabilities and coefficients used in the logit-linear regression and naïve Bayes methods. Additionally, while this study adapted a previously published coding scheme for convenience, it is unclear whether this scheme is best for prediction purposes. Any coding method also may suffer from inter-coder reliability; the coding in this study was all performed by the same investigator. When implementing the proposed methodology in an ED setting, multiple people would be entering codes which may reduce or improve the functionality of the chosen models. How this implementation effects model performance, and whether implementing predictions does indeed improve patient flow are other important directions for future research.

3.5 Conclusions

This chapter described and evaluated models for using data available at the time of triage to predict ED-to-IU admissions using expert opinion and two simple statistical models. This chapter also introduced a method for combining these predictions into a summative measure of near term ED demand for IU beds. The logit-linear regression model performed the best, with an AUC of .887 and an R^2 of 0.58 and a daily average estimation error for the summative model of 0.19 beds. This method was based on four readily available inputs (patient age, primary complaint, designation, and arrival mode). Recent studies have suggested that ED flow can be

improved by anticipating IU bed demand. The proposed summative measure provides a reliable estimate of near-future IU bed demand that replaces traditional ED crowding measures for influencing IU staff behavior and decisions. This is in contrast to a yes/no predictor that seeks to preempt provider bed orders in current work flow paradigms.

The prediction models in this chapter were developed from data at one site and the above results have not been demonstrated to generalize to other EDs. Furthermore the ED where the model was developed receives a low patient volume and resides in a small tertiary care VHA hospital, providing care to a specialized population. The set of factors that leads to an admission at a small hospital should be similar in larger hospitals, but that cannot be known without replicating the study in a larger site. Chapter 4 will describe a study where the logit-linear regression method is applied to three more hospitals. Additionally, a new test dataset is generated for VHA West Roxbury's logit-linear regression model where nurses perform the necessary coding. This will begin to establish the potential generalizability of the results found in here in Chapter 3.

Chapter 4: Generalizability of the Emergency Department Prediction Model²

As discussed in the previous chapters, prediction can be used to improve organizational factors, without increasing resources, by offering information that helps hospital staff prioritize their work. While a long term predictive method may be accurate over time, on any one day the natural variability of the hospital system may cause a spike in demand that is not accounted for in long term predictions. These spikes can be mitigated through the use of real time methodologies such as predicting admissions as patients enter the ED. Chapter 2 introduced some studies that used data collected early in a patient’s treatment (such as in triage or even in an ambulance) to predict whether that patient would eventually require admission to the hospital IU. This prediction allows hospital staff to reduce ED boarding times by preemptively mobilizing inpatient admission resources while the patient is still receiving their emergency treatment. Consequently, when the patient finally is ready for admission, downstream resources have already been aligned and the patient’s boarding time would be reduced [Levine et al. 2006, Clesham et al. 2008, Li et al. 2009, Sun et al. 2011, Peck et al. 2012].

Chapter 2 described the development and study of a linear regression with a logit transform (logit-linear regression) model for assigning a probability of hospital admissions at the time of ED triage. That study focused on predicting patient admission at VHA West Roxbury and focused on answering Question 1 of this dissertation “what predictive methods work best to predict downstream demand in the context of a single Emergency Department/Inpatient Unit health care delivery chain?” Given that all hospitals have organizational differences and two hospitals with identical resource levels and community demographic may still perform differently, it is important to explore the generalizability of prediction models to other settings [Hoot et al. 2009b]. This is the basis of Question 2 of the dissertation, “How portable or robust are these prediction methods to multiple hospital contexts?”

² *The majority of the material in this chapter is being prepared for publication as [White Paper et al. 2012].

The objective of this chapter is to study the generalizability of the prediction model/methodology introduced in Chapter 2. A second objective of this chapter is to explore how the model performs when data coding is performed by nurses in real time rather than by a single investigator retrospectively.

4.1 Methods

4.1.1 Study Design

In order to study the generalizability of the logit-linear regression approach discussed in Chapter 2, retrospective patient visit data was collected from four hospitals: VHA West Roxbury (VHA 1), VHA Medical Center 2 (VHA 2), a Small Private hospital, and a Large Public hospital. Separate development, validation and test datasets were collected for each hospital. In each case, the development set was used to create logit-linear regression models using a variety of combinations of factors. The models were then applied to the validation set, used to select a best performing factor combination. Finally the selected model was applied to the test set which was used to generate final performance results for comparison. Each dataset was coded by the author of this dissertation. In order to study the effect of generalizability due to live implementation of the model, an additional dataset was collected and coded by triage nurses prospectively at VHA 1.

All portions of the study were approved or granted exemption by the respective institutional review boards of each hospital. Approval for complete study design and implementation was provided by the institutional review board of VHA 1.

4.1.2 Study Setting

Four hospitals were chosen for this study, all located in the northeastern United States. This study sample included two small public hospitals, VHA 1 and VHA 2, one small private community hospital, and one large public teaching hospital. Based on the development datasets collected for this study, Table 4 summarizes the characteristics of the participating hospitals and corresponding EDs.

Table 4 Characteristics of participating hospitals based on development datasets

Characteristics	VHA 1	VHA 2	Small Private	Large Public
Medical Center Properties				
Bed Count	170	181	313	386
Trauma Level	Level 3	Level 3	Level 3	Level 1
Population	Adult Veterans	Adult Veterans	Adults and Children	Adults and Children
Funding	Public	Public	Private	Public
Community	Urban	Urban	Suburban	Urban
Emergency Department				
Bed Count	13	9	36	53
Triage System	5 level ESI	5 level ESI	5 level ESI	5 level ESI
~ Monthly Volume	1200	1200	4700	5200
Admission Percentage	32%	28%	26%	28%

4.1.3 Study Protocol

In Chapter 2, six factors collected at the VHA West Roxbury/VHA 1 triage were found to have value towards predicting whether a patient will eventually require admission to the hospital. These factors were: patient age, primary complaint, ED provider, designation, arrival mode, and Urgency (ESI).

In order to apply the logit-linear regression and choose a final model, three retrospective datasets were collected for each hospital. A development set was used for creating a logit-linear regression model using each combination of factors (listed in Table 5) [Witten and Frank 2005, Shmueli et al. 2007]. Not all of the hospitals in the sample for this study collect the same data at ED triage. Table 5 summarizes which factors were collected by each hospital and the data options for each factor. Even amongst the two VHA hospitals different data are being collected.

Table 5 List of factors collected at triage

Factors	VHA 1	VHA 2	Small Private	Large Public
Patient Age	Continuous	Continuous	Continuous	Continuous
Primary Complaint	Free Text	Free Text	Free Text	Free Text
ED Provider	Provider Set	Provider Set	Provider Set	Provider Set
Designation	Fast Track ED	Not Tracked	Not Tracked	North Ward South Ward Urgent Care
Mode of Arrival	Ambulatory Stretcher Wheelchair	Ambulance/Police Ambulatory Clinic Nursing Home Police Transfer Other	Ambulatory Stretcher Wheelchair Other	Not Tracked
Urgency	ESI Level 1-5	ESI Level 1-5	ESI Level 1-5	ESI Level 1-5

Table 6 Dataset attributes

Dataset Attributes	Dates	Collection Hours	Number of Patients	Coder
VHA 1				
Development Set	1/1/2010 - 5/6/2010	24hr	4187	Investigator
Validation Set	5/7/2010 - 5/31/2010 and 9/1/2010 - 9/21/2010	24hr	1614	Investigator
Test Set - Retrospective	9/22/2010 - 11/26/2010	7am - 5pm	1160	Investigator
Test Set - Prospective	6/13/2012 – 7/13/2012	7am - 5pm	910	ED Nurses
VHA 2				
Development Set	5/9/2011 - 8/31/2011	24hr	4077	Investigator
Validation Set	10/13/2011 - 11/30/2011	24hr	1648	Investigator
Test Set	1/1/2012 - 2/6/2012	24hr	1270	Investigator
Small Private				
Development Set	1/17/2007 - 2/28/2007	24hr	4910	Investigator
Validation Set	3/1/2007 - 3/15/2007	24hr	1712	Investigator
Test Set	3/20/2007 - 3/31/2007	24hr	1394	Investigator
Large Public				
Development Set	3/1/2011 - 3/24/2011	24hr	4020	Investigator
Validation Set	3/25/2011 - 3/31/2011	24hr	1150	Investigator
Test Set	6/1/2011 - 6/10/2011	24hr	1723	Investigator

It was unclear which combination of factors would generate the best performing model for a specific hospital. Therefore a model was created for every possible combination of the 6 or less factors collected at triage in each hospital. In other words, for each hospital, a linear model was created for each factor on its own, then every combination of two factors, three factors, etc. A validation set was used for each site to study the performance of each of the generated models. The results of this performance evaluation lead to the selection of a final model. The final model was then applied to a test set which is used for the reported model performance. Table 6 shows the attributes of each of the datasets for the four hospitals as well as a second test dataset that was prospectively generated for VHA 1, to test the effects of live implementation on model accuracy.

In order to create these models, some data needed to be categorized. The coding for the retrospective datasets was performed by the author and the prospective dataset was coded by triage nurses in real time. Age was categorized into decades and primary complaint was categorized using a modified version of a previously published ED complaint coding system [Aronsky et al. 2001]. All other factors were already in a categorized format.

The probability of admission was calculated for each value of each categorized factor. This probability of admission given a factor can be represented as $P[\text{Admit}|\text{Factor}]$. These probability values are used as the independent variable values in the logit-linear model. When generating the model the historical dependent variable has the values 0 (no admit) and 1 (admit). The model was fit using a logit-link function to ensure that predictions remain between 0 and 1. Table 7 shows examples of these probability values for factors that are tracked in each hospital. The complete table can be found in Appendix B.

Table 7 Example admission probabilities given selected factors

Probabilities of admission given factor	VHA 1	VHA 2	Small Private	Large Public
Urgency Level: P(Admit Urgency)				
ESI 1	0.86	0.25	0.82	0.96
ESI 2	0.54	0.55	0.46	0.57
ESI 3	0.55	0.34	0.26	0.29
ESI 4	0.03	0.13	0.03	0.02
ESI 5	0.03	0.11	0.00	0.00
Patient Age: P(Admit Age)				
<20	0.00	0.00	0.12	0.06
20-29	0.04	0.09	0.09	0.08
30-39	0.10	0.09	0.13	0.14
40-49	0.22	0.10	0.17	0.22
50-59	0.33	0.26	0.22	0.32
60-69	0.36	0.27	0.30	0.46
70-79	0.38	0.32	0.44	0.58
80-89	0.45	0.36	0.52	0.71
>90	0.51	0.39	0.62	0.74
Primary Complaint (sample): P(Admit Complaint)				
Abdominal pain	0.49	0.28	0.29	0.36
Abdominal problems	0.50	0.39	0.26	0.35
Abnormal Labs	0.70	0.66	0.54	0.56
...				
Cardio-vascular complaint	0.57	0.40	0.38	0.49
Chest pain	0.61	0.46	0.41	0.38
Cold/Flu	0.06	0.16	0.14	0.08
...				
Fainting/syncope	0.70	0.48	0.48	0.40
Fall	0.42	0.24	0.24	0.34
Fever	0.64	0.61	0.28	0.22
...				
Joint Problems	0.05	0.07	0.04	0.07
Kidney and Liver Failure	0.79	1.00	1.00	0.88
Laceration	0.10	0.00	0.04	0.01
...				
Psychiatric/social problems	0.41	0.40	0.63	0.36
Respiratory problems	0.67	0.40	0.47	0.64
Skin complaint/trauma	0.13	0.07	0.04	0.10
Total Probability of Admit	0.34	0.28	0.24	0.28

4.1.4 Data Analysis/Measures

Chapters 2 and 3 introduced two methods for utilizing predictions in order to improve patient flow. The first method was to use predictions within the current work flow of an ED and preempt a doctor's order for a specific patient. For this method, a prediction model would assign a 'yes admit' or 'no admit' to a patient which would be used as a coercive prediction to order an inpatient bed. This is then confirmed or overturned by the doctor at the end of the patient's ED treatment. The models that were generated for each hospital in this study were evaluated for how well they can perform within this work flow using area under the Receiver Operating Characteristic Curve (AUC) and the Hosmer-Lemeshow goodness of fit test (GOF) [Hosmer and Lemeshow 2000].

As described in Chapter 3, although it is important how well a model assigns the admission probability of a single patient, this method does not take complete advantage of risk pooling. In order to enable risk pooling, it was proposed that adding the probabilities of all patients in the ED can create a non-coercive running expected bed demand that can inform the behavior and decisions of IU staff. The summation of probabilities also allows for practical nuances involved with variance and confidence as discussed earlier.

Just as in Chapter 3, the models generated in this study were evaluated for how well they would create a running expected bed demand by calculating the total predicted bed orders for each day of test data and comparing that demand to the actual number of admissions for each day. These values could be compared using an R^2 correlation along with the residuals created by the difference between the predictions and the actual admission values. For the purpose of comparison, the larger volume hospitals are evaluated for R^2 by quarter days. As in Chapter 3, neither R^2 nor residuals are an accurate measure of prediction accuracy on their own, however, the combination provides useful insight.

4.2 Results

For each hospital, no single combination of factors created a model that performed best in all measurement categories: AUC, GOF, R^2 and daily residuals. Therefore, a model was chosen for each hospital based on high performance in each category. The chosen models are as follows:

- VHA 1: Patient age, primary complaint, designation, and mode of arrival
- VHA 2: Patient age, primary complaint, and mode of arrival
- Small Private: Patient age, primary complaint, mode of arrival, and urgency
- Large Public: Primary complaint, location, urgency

The coefficients for each of these models are shown in Table 8.

Table 8 Chosen models and coefficients for each hospital

Model Coefficients	VHA 1	VHA 2	Small Private	Large Public
Constant	-7.02	-5.06	-5.5	-6.15
Patient Age	3.39	4.99	4.52	-
Primary Complaint	2.89	5.76	3.98	4.1
ED Provider	-	-	-	-
Designation	5.48	NA	NA	8.53
Mode of Arrival	2.69	2.92	2.64	NA
Urgency	-	-	4.77	3.68
Significance $P \ll 0.01$ for all values				

Table 9 shows the quality of the prediction models when applied to the test sets, using the measurements that were described earlier. As can be seen, the models perform well using any measure. An exception is the GOF for VHA 2, which had a $P > 0.03$ when applied to the validation set and $P = 0.002$ when applied to the test set. Figure 26 shows a plot of how well each model categorizes patient admission likelihoods into probability deciles.

Table 9 Quality of predictive models when applied to test sets

	VHA 1 Retrospective	VHA 1 Prospective	VHA 2	Small Private	Large Public
AUC	0.89	0.86	0.80	0.86	0.82
R²	0.58	0.82	0.90	0.68	0.84
Average daily residual	0.50	-0.33	-0.96	-0.37	-0.31
GOF	$P = .02$	$P = 0.04$	$P = 0.002^*$	$P = 0.09$	$P = 0.07$
* $P > 0.03$ in validation set, explanation provided in discussion					

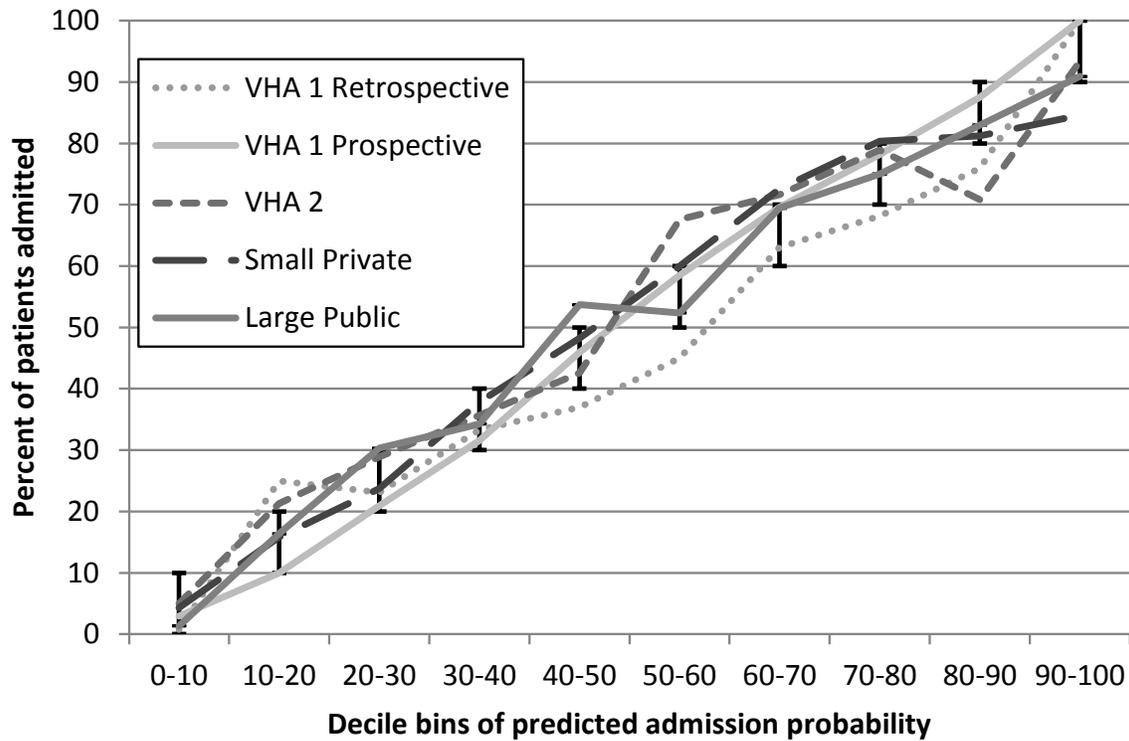


Figure 26 Comparison of predicted vs. actual admissions by probability decile

Although the total values for the R^2 and the average daily residuals are useful, it is also important to analyze the distribution of points that lead to these values. Figure 27 shows the plot of actual vs. predicted daily admissions as well as the resultant residuals for VHA 1 (retrospective and prospective) and VHA 2. Figure 28 shows actual vs. predicted quarter daily admissions as well as the resultant residuals for the Small Private and Large Public Hospitals.

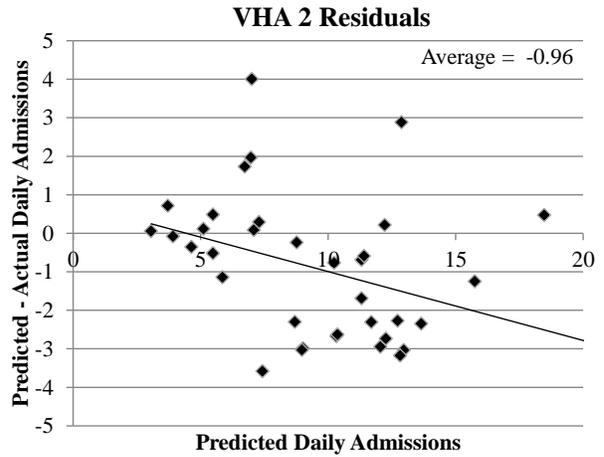
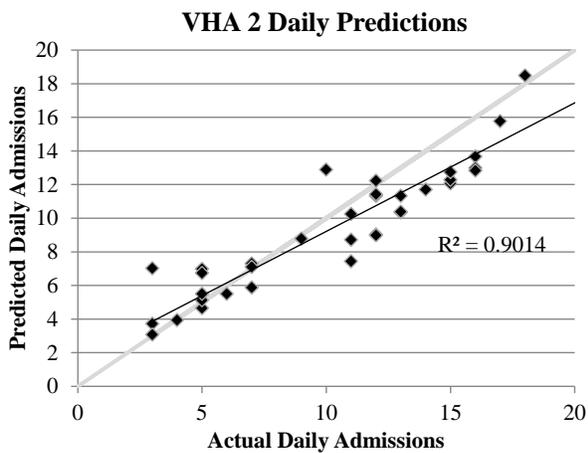
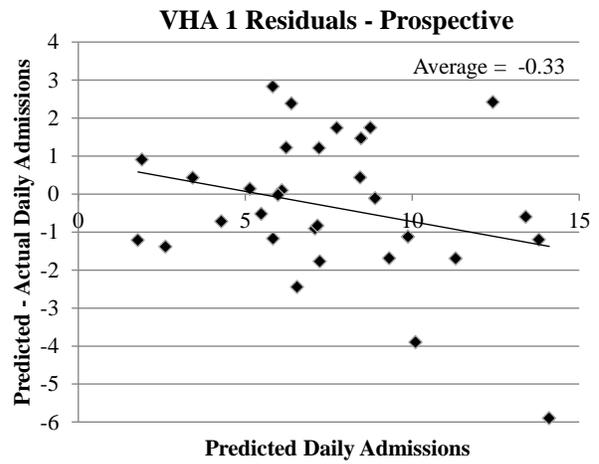
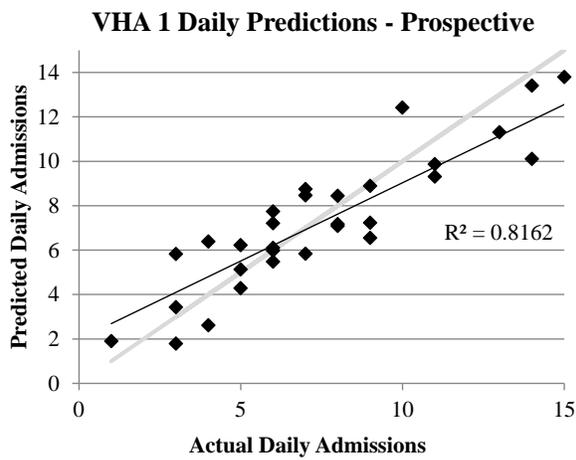
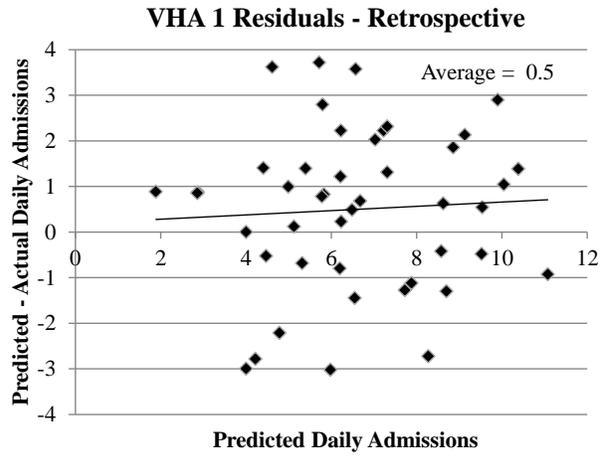
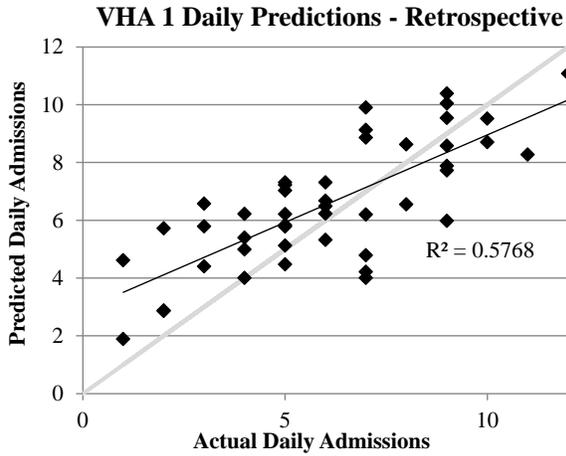


Figure 27 VHA 1 and VHA 2: daily actual vs. predicted admissions and residuals

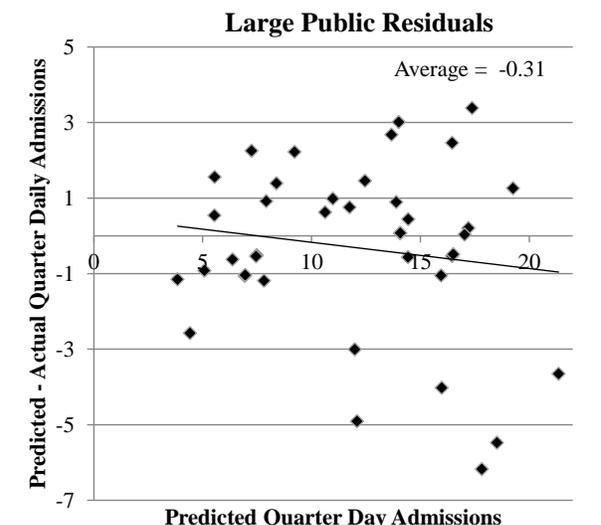
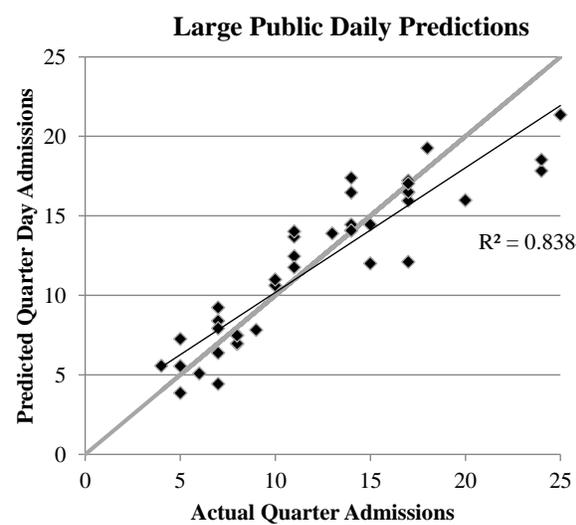
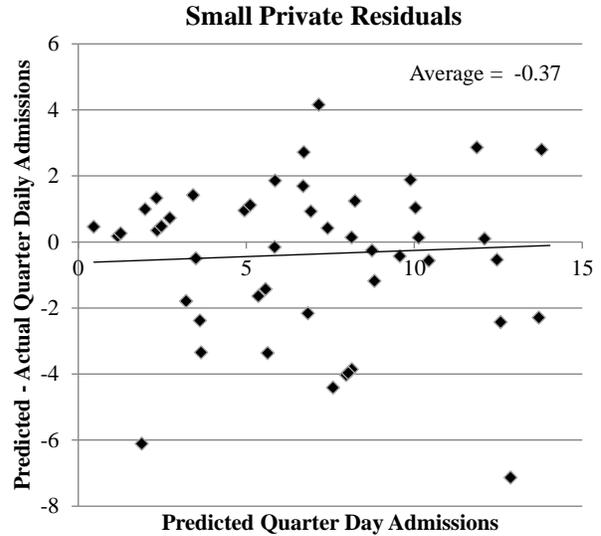
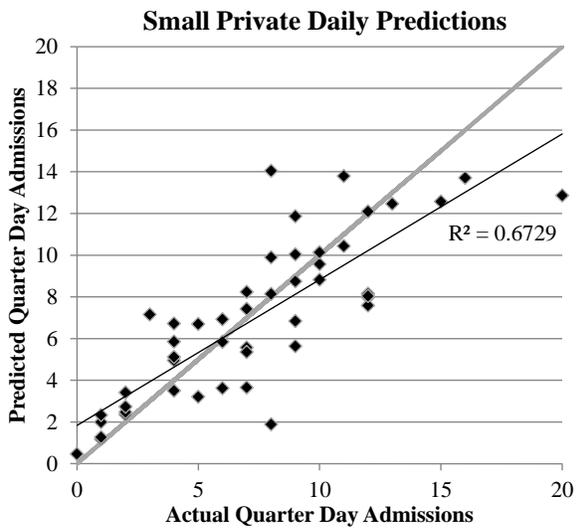


Figure 28 Small and Large Public hospital: quarter daily actual vs. predicted admissions and residuals

4.3 Discussion

In the past, models to predict whether a patient will require admission to the hospital from the ED have been developed with the goal of being coercive [Li et al. 2009, Sun et al. 2011]. These models assign a yes/no value of admission to a patient and can be used to place preemptive bed orders to the hospital. The models studied in this chapter show potential for being used in a coercive manner at the respective hospitals. When applied to the test datasets, all of the models had areas under the ROC curve of greater than 80%; this means that the models can identify patients who will need admission, reasonably well. The results of the GOF analysis showed that all of the models except the one applied to VHA 2 had a high GOF. Visual inspection of Figure

26, shows that the model created for VHA 2 can categorize patients into probability decile bins relatively well. In fact, the GOF for VHA 2 when applied to the validation set was in an acceptable range. It is possible that the GOF for VHA 2 shifted between sets because the value is sensitive to how many measurements are taken in each decile and in the case of VHA 2's test set, the model had some decile bins with a low number of patients. While this understanding may be a reason for moving forward with the model developed for VHA 2, it could also bring the GOF results for the other models into question. Therefore, while GOF is a useful measure, it must be used alongside the other measures presented in the data analysis section and the visual analysis of Figure 26. With that in mind, the combination of measures does show potential for all models to generate useful coercive predictions.

Implementing the models for coercive use would require further study of the ROC curves, in order to choose a prediction probability cutoff point for ordering a bed that best balances the positive effects of early action against the negative effects of early incorrect actions. The use of a coercive prediction measure also introduces questions of health decision quality where predictions become self-fulfilling. A doctor may believe that a bed is already open for a patient, or rely too much on the prediction model and therefore simply admit that patient, even if they would not have done this without the prediction. This is considered automation bias, which is a common issue with decision support systems [Skitka et al. 1999, Cummings 2004]. The details of this situation are a question of provider psychology and medical decision making which is beyond the scope of this study.

In order to avoid the issues associated with a coercive prediction model, Chapter 3 introduced the creation of a running bed demand index that would encourage hospital staff to make treatment and discharge decisions that are mindful of the current needs of the ED [Peck et al. 2012]. This index would not be used to order beds for specific patients but would replace ED crowding scales as a method of informing inpatient staff of ED demand [Bernstien et al. 2003, Jones et al. 2006]. This aggregate measure would be a more direct connection between ED demand and IU demand. Aggregating predictions is also a means of risk pooling which should improve overall accuracy.

There is no single statistic to measure accuracy of the models when demand is aggregated. The measures used to judge individual patient prediction accuracy also are a basis for evaluating the

use of the aggregate measure (if it is accurate for one patient, it will be accurate for many). However aggregation also enables the use of the R^2 statistic which measures whether an increase in actual daily bed demand is met with a proportional increase in predicted bed demand. This is a good measure of whether the model is dynamically working correctly but does not measure the raw prediction accuracy. For this reason the R^2 analysis is accompanied by a residuals analysis. This analysis shows that, as one would expect, errors get larger as predictions get larger. The analysis also shows that the models have ranges for which they work best but generally perform well on the average. Although model inaccuracy is a concern, it actually would be mitigated in practice since the measure would be updated in real time. In other words, as the day goes by a patient with an incorrect prediction would leave the ED and the model would be corrected, while in the R^2 analysis errors add up for the entire day.

While all of the hospitals studied in this paper have a similar system where the ED feeds into the rest of the hospital, not all hospitals follow the same processes for facilitating this flow. Variations occur on all levels of the organization from the actual medical decision making to the logistical decision making. In this way, despite the similarity in the macro system, significant performance variation can be expected between hospitals. The results of this are evident in the creation of the prediction models. It was seen that the predictive value of each factor varied by hospital. This can be the result of how practitioners use each factor. For example, the ESI system may be implemented differently at each hospital, which may explain why the ESI 1 admission probability was strangely low for VHA 2. It can also be a difference in how the factors are collected, for example the use of ED versus fast track in VHA 1 and the use of North Ward, South Ward and Urgent care at the Large Public hospital. Whatever the cause, absolute generalizability of predictive factors was not found across the hospitals. Instead generalizability has been shown here to be based on the process of creating a prediction model (using multiple factors collected at triage to create a logit-linear regression), rather than the specific resultant models themselves.

In this study all retrospective data was coded by a single investigator, while in practice the person doing the coding would likely be a medical practitioner. It was just discussed that the type of data and method of collection for each factor could impact that factor's predictive value. This is also true for complaint coding and, therefore, the person performing the coding can have an

impact on the success of a model that uses primary complaint. It is for this reason that the study included a test dataset which was prospectively coded by nurses in real time. The results of this piece of the study showed that the VHA 1 model continued to perform well when applied to data coded by nurses in real time which suggests that the predictive models generated by the process described in this paper are likely generalizable to live implementations, however this must be done in other sites to be proven.

Given basic triage level data, a prediction model was developed for each hospital that performed fairly similarly. Generalizability of the factors that went into the model could not be proven since the actual information collected by each hospital was different. For this reason it is impossible to prove (or disprove) that, if all hospitals collected the same data, in the same format, then they would all use the same model. However this seems unlikely, since Table 7 shows that even factors that are the same, such as ESI level and age, have different admission probabilities at each hospital.

The above are descriptive conclusions where the prediction models characterize the behavior of the respective hospitals by explaining which variables are predictive. However, creating prediction models for multiple hospitals also can lead to prescriptive conclusions. By looking at the data it can be seen that both non-VHA hospitals had urgency in their selected prediction models. This may be coincidence, it may be a reflection of the patient population, or it may also suggest that the implementation of the ESI scale is different at these hospitals and is more effective for predictive purposes. It can also be seen that both the Large Public Hospital and VHA 1 tracked patient location, and both had this information in their predictive models. Therefore one may conclude that it would be valuable for the other hospitals to track this information (if separate designations exist in those hospitals). Similarly all hospitals but the Large Public Hospital has arrival mode in their chosen models, which suggests that it may be useful for the Large Public Hospital to begin tracking this information.

4.4 Limitations

The prediction methodology that was adopted here is a relatively simple method which was able to provide results with a reasonable amount of data. This is particularly important when an investigator is coding each primary complaint. It also means the study may be repeatable by a

practitioner with limited statistical experience. Nonetheless, there are many other methods that exist and can be implemented by those with more advanced experience; these may require a significantly larger database [Shamuli et al. 2007, Whitten and Frank 2005].

Any conclusions that can be drawn by comparing the four hospitals in this study are limited because four sites do not comprise a complete sample. By demonstrating that accurate predictive models can be generated for all four hospitals, it may be reasonable to assume that similarly accurate models can be generated for any hospital. The limited sample is enough to prove a negative: not all hospitals with similar characteristics will necessarily perform the same, given that the two VHA hospitals of similar size had different predictive models. Finally, although the model was applied to a smaller private hospital and a larger public hospital, the conclusions made are specific to those institutions and should not necessarily be applied to all large and small private/public hospitals.

A previously existing primary complaint coding methodology was adopted. How the datasets are coded can strongly affect how models using those codes perform. The system used, adopts a mix of terms that encompass symptoms (such as chest pain) and other terms that are diagnoses (such as cardiac arrest). Another system may require coding cardiac arrest as chest pain, which would significantly reduce the estimated probability of admission for a patient who enters the ED while suffering a cardiac arrest, reducing the quality of the predictor. Just as it is clear that predictions can be made worse, based on a different coding system, it is also possible that predictions can be improved by using a different coding system.

4.5 Conclusions

The study in this chapter showed that logit-linear prediction models can be developed for multiple different hospitals of varying size and administrative structures. Generalizability is shown for the methodology rather than for specific models that were derived by the methodology. In one hospital it was shown that the prediction model continues to perform well even when coding is performed by triage nurses prospectively. These prediction models can be used in a coercive system for driving specific behaviors or a non-coercive system for sharing information and encouraging resource allocation decisions that are based on larger system knowledge. The next two chapters will discuss studies focused on characterizing the potential

value of sharing the ED admission information in a non-coercive way. Chapter 5 will introduce a live implementation study where the running bed demand measure, introduced in Chapter 3, is actually calculated and shared in real time at VHA West Roxbury/VHA 1 for two weeks. Understanding that short term, live implementations have limitations; Chapter 6 will describe a computer based simulation study that seeks to further understand the potential benefits of using prediction by studying a controlled environment.

Chapter 5: Implementation of the Emergency Department Prediction Model

In Chapters 3 and 4 a methodology was introduced for predicting and sharing the likelihood that a patient would be admitted from an ED to an IU based on information gathered in ED triage. It was shown that a logit-linear regression model can be used to make relatively accurate predictions compared to the naïve Bayes and expert opinion methods. It was also shown that the logit-linear regression method could be consistently applied in multiple hospitals, as well as in a live implementation, where nurses perform the coding. The studies described in the previous chapters were focused on answering the first two questions proposed in Section 1.5. The next two chapters will be focused on answering the final question: “Given advance demand predictions, what possible adaptive actions can the hospital system take to improve flow given (a) perfect and (b) imperfect downstream demand prediction?”

This question resonates with the second of two components of the IOM “anticipation of need” described in Section 2.3 [IOM 2001]:

1. The Prediction – making forecasts about the future needs, progress, processes or steps of the patient.
2. The Response – taking definite actions in response to the prediction.

As discussed in Chapter 2, in supply chain management the use of prediction in order to improve coordination between multiple components is well established [Simchi-Levi et al. 2003]. The previous chapters have suggested that the coordination of care between a hospital’s ED/IU health care delivery chain is comparable to the coordination between a two parts supply chain. It was suggested that better coordination, through prediction, can reduce the amount of time between an ED provider’s admission decision and the patient’s assignment of an IU bed, also known as boarding time. This boarding time has been identified as a major barrier to improving ED crowding and quality [Asplin et al. 2003, US GAO 2003, Falvo et al. 2007, Hoot and Aronsky 2008, US GAO 2009, Viccellio et al. 2009].

It has been discussed that some ED forecasting models have been developed with coercive actions in mind, whether it is long term forecasts to define resource allocation decisions [Jones et al. 2002, Jones et al. 2008, Abraham et al. 2009, Wargon et al. 2010] or predictions to expedite admission procedures for a specific patient [Arslanian-Engoren 2004, Walsh et al. 2004, Sun et al. 2011, Li et al. 2009].

While predictions may seem inherently useful, the question of how to best use them remains unanswered. Strong leadership and systematic studies can be used to develop an optimal coercive response to individual forecasts. However, such high handed actions may not be well accepted in hospitals, deterring adoption. For this reason it is desirable to find a prediction system that still relies on the ability of IU staff to weigh the costs and benefits of acting upon this prediction, as they understand them. Other studies used ED crowding scales and predictors in order to influence inpatient staff decisions in a non-coercive way [Bernstein et al. 2003, Jones et al. 2006]. These studies suggested that IU staff will respond appropriately when they see that the ED is crowded. However there is not always a direct connection between ED crowding and IU demand. This disconnect may make IU staff less likely to respond to a crowding index. Chapter 3 described use of admission predictions in the ED in order to create a running bed demand index. This bed demand index is comprised of demand directly related to the IU and would more directly inform inpatient staff of incoming demand, in order to influence priorities and decisions [Peck et al. 2012].

When using the bed demand index in a non-coercive way, the decision to prioritize patients who need discharge or those that need treatment is based on the individual perspective of the decision maker. While a hospital manager may value a reduction in waiting times and lengths of stay, a practitioner may have other values, both systematic and selfish. How the practitioner weighs these values when provided with information of a distinct format influences the effectiveness and value of sharing that information. This chapter seeks to better understand the individual and systematic effects of a live implementation of the running bed demand index prediction system. The goal of this implementation is to understand the results of sharing the prediction, both in terms of quantitative measures that are tracked by the hospital as well as qualitative measures captured by surveys and interviews. The results will inform possible improvements that can be made to the prediction model for future long term implementations.

5.1 Methods

5.1.1 Study Design

The study took place at VHA West Roxbury which is described in Section 3.1.2. A trial implementation was run for two business weeks, 10 Days, July 2nd 2012 through July 13th 2012. During this trial the predicted bed demand was calculated every half hour from 7am through 7pm. This demand was then shared on the hospital internal homepage, which is displayed each time a staff member opens the default web browser, Microsoft Internet Explorer 8. The quantitative measurements taken during this trial included waiting time for patients entering the ED and the boarding time between admit decision and inpatient bed assignment. This data was analyzed through a time series analysis comparing performance during the experiment, data from the year leading up to the experiment, data for two months after the trial, and data from 2011 (to rule out seasonal effects and long term trends).

Qualitative data was focused on understanding the staff actions and thoughts during the trial period. The qualitative data included specific ways that staff used the predictions towards their own priorities as well as the thoughts of staff members on how to make a prediction model more useful. This data was collected through a mix of interviews and surveys. Surveys were sent out to all hospital employees in order to capture any unexpected effects of the predictor, interview requests were for a sample of hospital management, nurse managers, in-patient ward residents, hospitalists, environmental management services (including housekeeping) and hospital bed managers.

The protocol for the study including interview questions, survey questions, data tracking mechanisms and data sharing mechanisms were approved by the VHA BHS institutional review board. All numerical analysis is performed in Microsoft Excel.

5.1.2 Study Protocol

This study implemented the logit-linear regression prediction model and the summative bed demand measure discussed in Chapter 3. It will be recalled that the prediction method assigns a probability of admission to each patient that enters the ED. This probability is based on four factors: patient age, primary complaint, designation (emergency room or fast track), and mode of

arrival (stretcher, wheel chair, or ambulatory) [Peck et al. 2012]. For each patient that entered the VHA West Roxbury ED during the trial period (7/2/2012 – 7/13/2012), a nurse would input the patient’s data into a computerized data entry form shown in Figure 29. This form had dropdown menus where nurses could select the appropriate categorized data for each patient. The full list of these categories is in Appendix A. The last four digits of the patient’s social security number were combined with the patient’s date of entry, to create a unique patient ID for data reconciliation purposes. All identifiable data remained on a closed secure network.

The image shows a screenshot of a web-based data entry form. At the top, there is a header with the text "VA Boston Healthcare System" and "Emergency Department Admission Information for Prediction". Below the header, there are five input fields, each with a label and a dropdown menu. The labels are "Last Four:", "Age (Decade):", "Complaint:", "Location:", and "Mode of Arrival:". Each field has a small asterisk to its right. Below the input fields is a "Submit" button.

Figure 29 Data entry form at patient arrival

The data entry form was used to populate a web-based Microsoft SharePoint database that would use the logit-linear regression formula introduced in Section 3.1.4, with the coefficients in Table 3, to assign a probability of admission to each patient based on the data entered by the nurse. On trial days, every half hour from 7am to 7pm an investigator would access a Microsoft Excel spreadsheet that automatically downloaded the information from the SharePoint web-database. The investigator would then pull an updated ED record from the VHA health information system and enter this data into the Excel spreadsheet which would cross reference the two data sources to update the current status of each patient to one of four options:

1. ED Patient,
2. Admit – ED Boarding,

3. Admitted,
4. Discharged

A patient's status is routinely updated in the VHA health information system as nurses enter data for patients in the connected ED information system. This process ensured that no extra data entry was required of nurses after the initial patient entry.

The spreadsheet then calculated three variables to be shared with hospital staff. These were the "Predicted Bed Orders," "Current Beds Ordered," and "ED Admits Already on Floors." Noting that the number of patients entered into the computer system can exceed the 13 ED beds, due to the use of chairs and hallway beds, the variables were calculated as follows:

1. Predicted Bed Orders – The sum of the calculated admission probability for all patients who have the status 'ED Patient.'
2. Current Beds Ordered – The count of patients who are in the ED and have the status 'Admit – ED Boarding.'
3. ED Admits Already on Floors – All patients who entered the ED during that day and have the status 'Admitted.'

Note that patients with status 'Admit – ED Boarding,' do not contribute to the Predicted Bed Orders, despite currently residing in the ED. Also note that ED Admits Already on Floors is a cumulative measure that resets to 0 each day of the trial at 12am while the other measures are variable based on current ED status. Finally note that patients with the status of discharged had no effect on any measure.

The variables described were shared with all staff in the hospital through a prominent display on the VHA West Roxbury intranet homepage. This was the website that appears every time a staff member opens Internet Explorer (the default web-browser for all computers on the VHA network). The homepage featured the following description:

“The VHA West Roxbury Emergency Department is working on a crystal ball:

Wouldn't it be great if we could predict what is coming in the future, even just a few hours? The West Roxbury campus Emergency Department (ED) is working on just that. Over the past year ED staff members, along with the New England Veterans Engineering

Resource Center (NE VERC), have gained national attention for working on a system that will predict, hours in advance, how many patients will need admission to the hospital. However, it remains unclear how this information can be used to improve quality of care for our Veterans. So help us out. Take a look at the predicted numbers shared here:”

This was followed by the real time updating display of the variables described above, shown in Figure 30.

From West Rox ED as of 7/5/2012 5:29 PM		
Predicted Bed Orders	Current Beds Ordered	ED Admits Already on Floors
3.2	5	3

Figure 30 Public display of prediction variables

Table 10 Historical average and standard deviation values of prediction variables provided as reference for hospital staff during implementation period

Hour	Avg. Beds Predicted	+ 1 Std. Dev.	Avg. Total Admits	+ 1 Std. Dev.
6am	0.17	0.55	2.46	5.16
7	0.18	0.55	2.83	5.83
8	0.21	0.62	3.15	6.41
9	0.21	0.62	3.15	6.41
10	0.33	0.85	3.44	6.97
11am	0.54	1.26	3.80	7.58
12pm	1.16	2.52	4.36	8.58
1	1.32	2.82	4.65	9.08
2	1.41	3.06	4.89	9.52
3	1.41	3.06	4.89	9.52
4	1.49	3.19	5.11	9.89
5	1.41	2.99	5.34	10.37
6	1.51	3.13	5.89	11.39
7	0.86	1.90	6.10	11.78
8pm	0.69	1.60	6.31	12.19

*Based on data from: 1/1/2010 - 5/31/2010

The image was followed by the contact information of the investigators as well as a link to a supplemental webpage with more information about the prediction model and its development. Also included on the supplemental page was a chart of historical data for the variables in order to give staff some information for comparison, shown in Table 10.

Before the beginning of the experimental period the author of this dissertation presented at the nurse manager meeting (attended by the head nurse of each ward in the hospital), resident morning rounds (an educational session for the medical residents that was held each morning), and the patient flow committee meeting (a group of hospital staff tasked with studying and managing patient flow including the hospital bed managers).

The homepage content was posted for 10 business days, spanning from 7/2/2012 – 7/13/2012. During this time, the variables were updated every 30mins. Patient flow data was routinely collected by hospital software and this served as the data collection mechanism for the experiment. This data was analyzed for any quantitative effects of the trial. At the end of the trial a link to a web-based survey was sent to all staff members who attended one of the pre-experiment presentations. In addition to the survey, individual interviews were performed with a sample of hospital staff in order to generate information about the value of the predictions. Interview sampling was a mix of purposive and snowball sampling. Specific members of the hospital staff were chosen for interviews. These staff members then suggested others who had shown interest in the prediction system or who were in positions that may derive value from the prediction system.

The sampling technique led to 10 semi-structured interviews. These interviews consisted of three sets of questions. The first set was to establish the expertise of the interviewee, the second set of questions were to establish whether the interviewee was aware of the prediction trial and if they used the information that had been shared. The final section of the interview was to better understand the potential value of predictions from the ED and what changes could be made to the tool, in order to achieve this potential value. Some of the questions also enabled an open ended dialog. The survey had the same format as the interviews, with essay boxes for the open ended questions.

5.2 Results

The results collected from the intervention were both quantitative and qualitative. The purpose of capturing both types of data was to get a more complete understanding of how the prediction information could be used. In particular, efforts were made to capture unexpected benefits or issues that arise from sharing the prediction information, besides facilitating flow between the ED and IU.

Before analyzing the results of sharing the prediction information, it was worth evaluating the accuracy of the information shared. This is the dataset that was used as the prospective test dataset for VHA 1 in Chapter 4. It will be recalled that predictions were made prospectively using nurses as coders during an extended time period including the trial, 6/13/2012 – 7/13/2012. The area under the receiver operating characteristic curve for this time period was 0.89, the R^2 correlation between daily total predicted beds needed and actual beds needed was 0.82 with an average daily residual of -0.33 beds. Finally the individual patient predictions had a Hosmer-Lemeshow goodness of fit of $p = 0.04$ [Hosmer and Lemeshow 2000]. All of these measures show that the model was working accurately during the trial period.

As a visual indicator of the performance of the predictor during the trial period, Figure 31 shows the normalized sum of the predicted and current bed orders throughout the trial and the normalized amount of current bed orders already on the floors throughout the trial. As can be seen in the figure, there is a close relationship between high peaks in predicted beds and high peaks in cumulative admissions for a particular day. It is also possible to see that the peak in predictions occurs, on average, 3 hours prior to the peak in actual admissions.

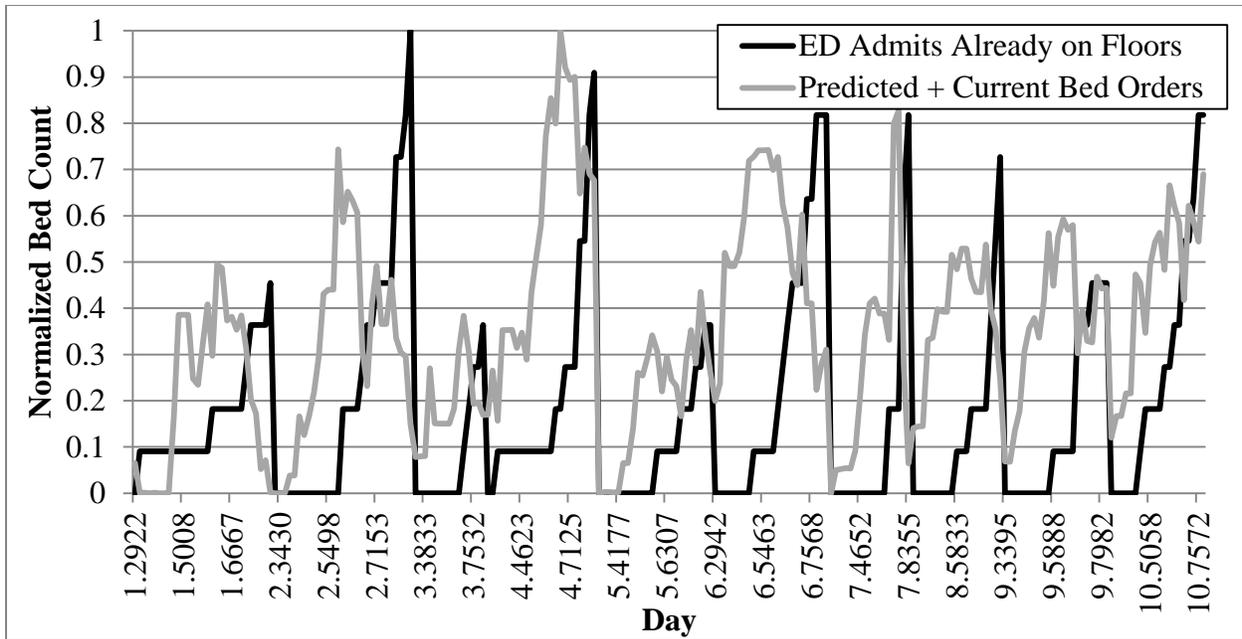


Figure 31 Normalized predicted and current bed orders compared with normalized ED admits already on floors during trial period

Having established that the prediction data had some accuracy, it is now worth exploring the results of sharing this data. Two primary measurements were collected for the trial time period. First was the waiting time between when a patient enters the ED and when they are assigned an ED bed. Figure 32 shows the weekly average and standard deviation of waiting time for the time period of 6/1/2011 – 8/31/2012, from 7am to 7pm. The weeks of the trial are highlighted.

Due to a complication with data entry, true boarding time is not tracked perfectly, thus boarding time was estimated as the difference between the average length of stay of admitted and non-admitted patients. Figure 33 shows the average and standard deviation of boarding times, for the same time period as Figure 32.

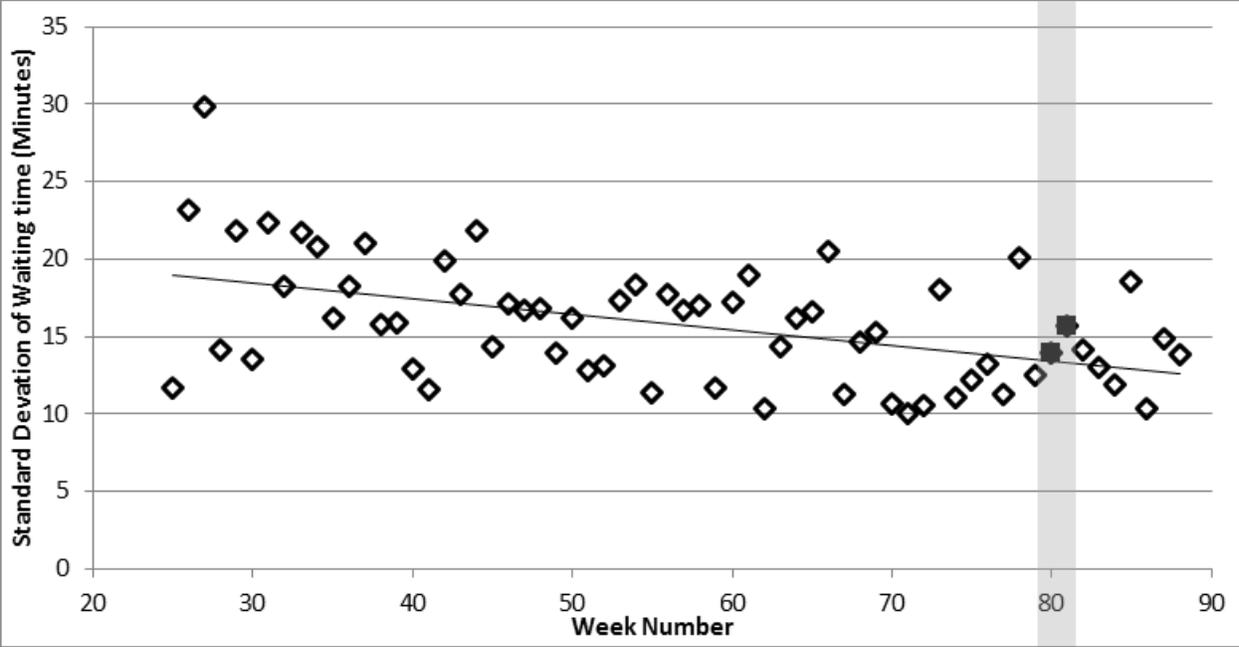
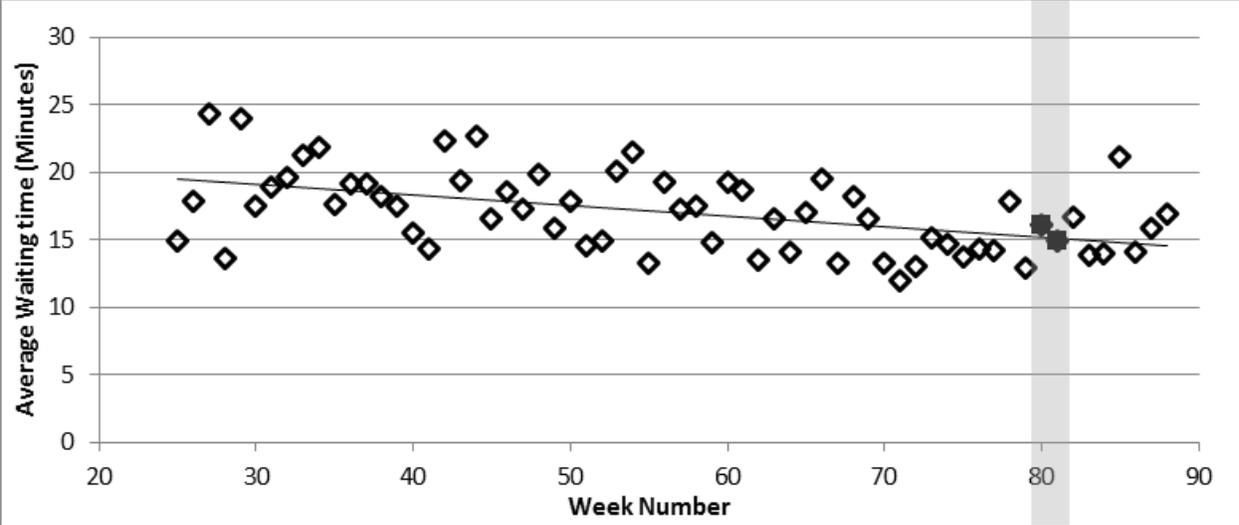


Figure 32 Average (top) and Standard deviation of (bottom) ED waiting times by week

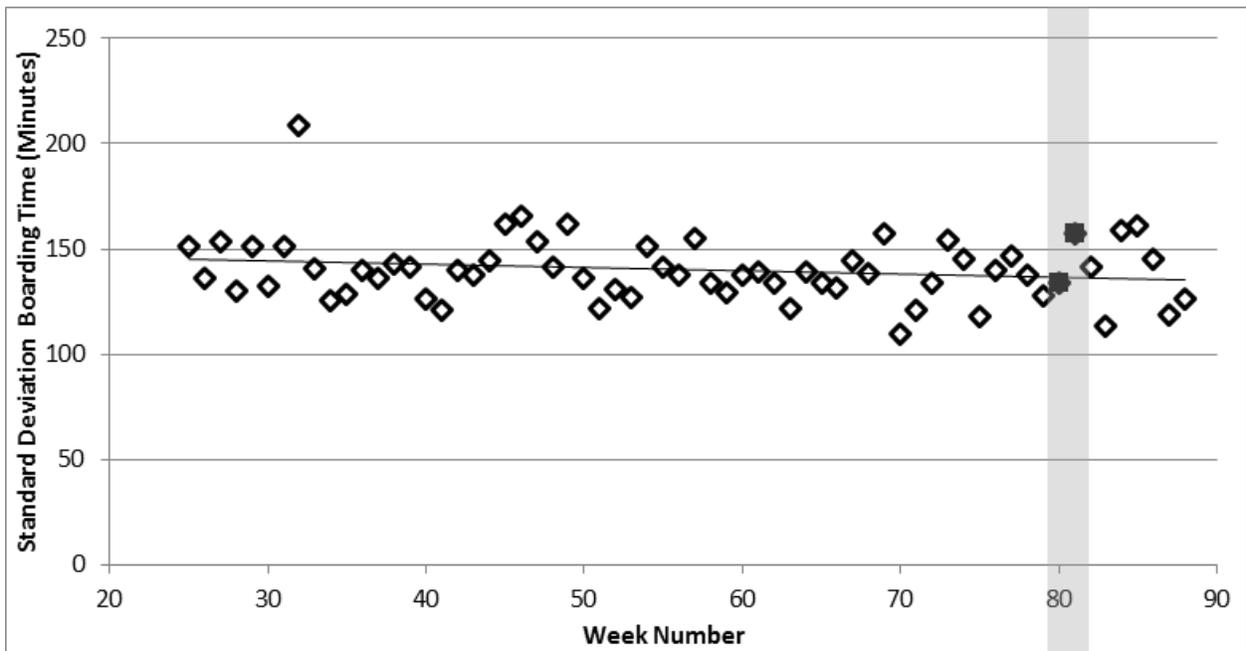
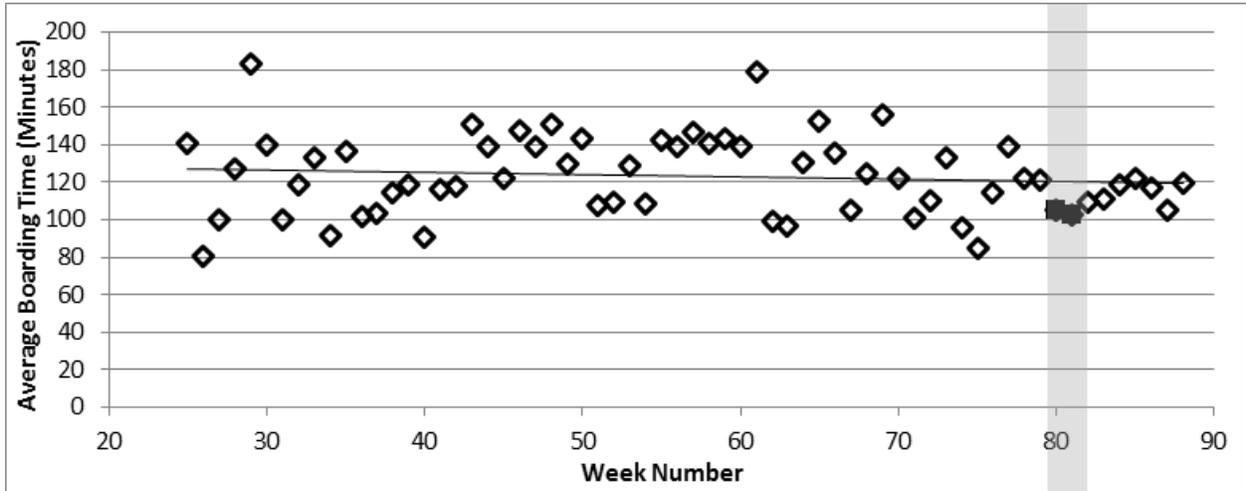


Figure 33 Average (top) and Standard deviation of (bottom) ED boarding times by week

The 10 interviews and survey data were summarized into comments related to how the prediction information was, or could be, used. Five categories for uses of the prediction information emerged from an analysis of the comments:

1. Admission Planning
2. Resource Scheduling
3. Personal Scheduling
4. Resource Alignment
5. Hospital Network Management

The interview notes were then coded into these five categories. All comments used for the following data are kept anonymous.

5.2.1 Admission Planning

The primary intended goal of using a prediction system, as described by previous chapters, was to enable admission planning. During the trial period there were some cases of this occurring. For example, when the “predicted bed orders” number got high (in the opinion of a senior resident at that moment) senior residents would sometimes walk to the wards, let the residents know about the prediction value, suggest that the residents start focusing on discharges and admissions. Some of the residents looked at the prediction number themselves. This turned out to be a new incarnation of behavior that sometimes takes place anyway. Bed controllers and residents both suggested that, in the past, they would look at the ED bed board and try to guess which patients would be admitted and this would inform their decisions. The prediction number in some ways provided an easier method for getting the information that would otherwise require accessing the ED system.

Despite the fact that those involved in admitting a patient sought out the prediction information, there were some aspects of the information, and sharing method, that undermined the final value of the predictions. While the aggregation of admission prediction information creates a more accurate measure of incoming bed demand, it also creates a level of abstraction. Doctors must see this single number and then understand how this translates into crowding or into future work load. Without being exposed to the model for a longer period of time, the practitioners did not have a good sense of how a specific ED prediction state, as displayed by Figure 30, would

translate into future work load. While historic data for the prediction variables was provided (Table 10), the practitioners did not intrinsically trust it. The practitioners needed to gain their own sense of the current ED prediction state relative to other ED prediction states that they had experienced firsthand.

The lack of ability to directly translate a prediction number into a full understanding of future business became an issue when staff members were considering taking preemptive admission or discharge actions. These staff members were concerned with the opportunities for wasted work caused by overreacting to a prediction state. Comments on this issue were always clarified with the addendum that if the staff member had acted on the prediction and it was indeed correct, then time would be saved and the actions would be worthwhile, but they were not comfortable with the risk.

Finally, acting upon a prediction for the purpose of admission planning was further undermined by the lack of specific predictions tied to individual patients. While the hospital staff did find it useful to see aggregate numbers when they wanted a quick understanding of the system, they also desired specific predictions for each patient, when they had time to look at the data in more detail. These specific predictions become important when one considers patients with special needs such as medical vs. surgical, telemetry, negative pressure, contagion precautions etc. Such needs limit which rooms a patient could be assigned too and make it necessary to know if the prediction number included patients that would need special conditions and the specific prediction number for those patients. Putting work into expediting an admission could lead to negative consequences and wasted effort if the patient that gets admitted needs specialty care and the newly opened bed is not appropriate.

5.2.2 Resource Scheduling

As discussed in Chapter 2, scheduling has been recognized as a key use for long term prediction models. However, that is in terms of baseline schedules. The interview and survey results showed a possibility for using real time predictions for short term resource scheduling. The residents' weekend schedule is often difficult to manage and requires some guess work for the senior residents on Friday at the end of the day. During the trial period, senior residents used the prediction model on Friday afternoon to get a better sense of what the hospital would look like

over the weekend. In the past they only knew who was in the IU at the moment they were doing the scheduling, during the trial, at the very same time of day, they had a more accurate sense of what the hospital would look like, because they also knew what demand was incoming.

Resident teams on the wards also used the prediction for short term staffing issues. The main teams go off duty around 7pm, a small group of night doctors and some residents manage the later shift. At 6pm, the teams were able to look at incoming demand from the ED and better prepare for their hand off. Alternatively, leadership could use the predicted demand to ask teams to stay longer in order to smooth the transition. The demand on night staff was also felt to be directly correlated to demand on morning staff. Consequently, it was suggested that knowing incoming demand at the end of the day could be used to better plan for the morning. Before this trial, demand was less clear when the administrative staff left for the day; during the trial, administrative staff felt more capable of making decisions for the next day, before leaving.

Another use for predictions in resource scheduling was suggested for Environmental Management Services (EMS). These are the staff members who clean rooms after patients have been discharged and prepare rooms before patients are admitted. These staff members can only clean a room when the order has been placed in the computer system. While they are scheduled to be evenly distributed around the hospital, it is possible to move staff around the hospital to accommodate surges of patients into certain wards. However, they can only plan for this if they are given an idea of where the surge will be going. In this way, the current system relies on others to use the predictions to make preemptive orders to which EMS would have the flexibility to respond.

The uses above have some limitations that were mentioned during the interviews. The first is in terms of timing. While the prediction model does predict well about 3-3.5 hours ahead of time, when a staff member looks at the displayed prediction, they don't have an exact sense of what will be happening in the very near future. The number does not suggest how many of the ED patients only recently entered the ED, how many have been receiving treatment for some time, and how many will likely request admission very soon. Even if time was known, the flexibility of staffing becomes an issue. While EMS felt that it is possible to move staff around, this is not as easily done on the medical wards. One interviewee expressed frustration at seeing the predictions and not having the power to call in more staff or move staff around.

5.2.3 Personal Scheduling

Personal scheduling relates to how predictions cause staff members to make decisions related to their own time, rather than adjust the times of others. At VHA West Roxbury, practitioners in the IU were often shielded from misalignments between bed availability and demand. This is because they only heard about a patient when bed control assigned that patient to an IU bed. Often, IU practitioners were unaware of the patient that was waiting for bed control to be given a free bed to work with. The prediction state, as shown in Figure 30, showed ward staff how many patients were waiting for beds, it also showed how many already had been given beds earlier in the day. This provided them with a sense of the limitations on bed supply and whether it would meet incoming demand.

Knowledge of incoming demand was important to the medical residents. These staff members felt that it was often forgotten that physical beds were not the only potential bottleneck in the system, but that they could only perform a certain amount of tasks in a given time frame. Each admission requires a process, which takes time; therefore, even when beds are available, staff can only admit a limited number of patients in a certain time frame. The residents felt that by knowing the prediction state of the ED and how tight resources may be, they could better choose how to prioritize their activities. This was important in the morning when discharged patients seemed low priority compared to patients that have not been seen that day. The prediction knowledge may cause residents to prioritize morning discharge and prepare an admission in order to avoid a bottleneck later on. Similarly throughout the day, there were educational sessions for the residents, these sessions could be missed occasionally and the residents may use the ED prediction state as a basis for choosing whether to attend or not. Decisions of this nature, again, lead to the issue mentioned earlier in terms of potential to do unnecessary work. Another issue that arose when considering personal scheduling is that a staff member may merely hope that the incoming demand wouldn't be assigned to them and choose to proceed with their normal routine. This is reflective of the human element in this system. There are some practitioners who like the potential for control over the system, others are comfortable dealing with issues as they come and would avoid looking for information on the future and trying to act upon it.

Nursing staff similarly found the predictions to be useful for personal scheduling. They also used it for prioritization. Nursing managers felt that the prediction state of the ED at the end of the day

may cause them to stay later in order to smooth the transition into the night shift. This was a work around for the fact that they didn't have the option to schedule other people to stay later or come in, but they were able to make the personal decision to stay.

A final, personal scheduling based, value to sharing the prediction model was the mental preparedness that occurs when the future was better known. All interviewees felt that simply by knowing that the future would be busy they were better prepared to deal with it. This may not directly translate into any ED performance improvements but may have effects on job satisfaction and other quality measures.

5.2.4 Resource Alignment

A common point that was made by all interviewees was the issue of coordinating with other staff. There were times when the prediction state of the ED encouraged some staff members to prioritize a patient for discharge but they were awaiting the actions of another person or office. A doctor could not discharge a patient without knowing that there was a place for the patient to go; this required the efforts of the patient care coordinators, social workers and discharge planners. Often these other key personnel were not given advanced notice that a patient would be ready to leave and they would first start to work with families or other facilities to arrange for the patients discharge, when they received word that the patient was ready to leave. While doctors are waiting for some parts of the hospital to coordinate the patients discharge, nurses are not able to perform the discharge procedures until the doctor has written a discharge note. Similarly, getting outgoing prescription orders filled by the hospital pharmacy, or having out bound tests performed, required coordination with further support services. After these process delays, that some interviewees suggested could take days, the empty bed would finally be entered into the computer, enabling a bed cleaner prepare the bed for the next admission.

Each interviewee identified at least one other service within the hospital as needing to use the prediction in order to better drive behavior. This information should not be seen as placing the blame, but rather, as recognizing true process steps that prevent a staff member from completing a discharge or performing an admission. A bed controller noted that the prediction system was useful for overcoming this systematic grid lock, by providing more data to bring to morning discharge meetings, that would enable communication and future planning across departments.

5.2.5 Hospital Network Management

While the initial primary goal of the prediction model was to improve flow between the ED and IU, actually sharing predictions with the hospital uncovered the importance of understanding flow across a hospital network. VHA West Roxbury was part of three hospital networks. The first, more direct network is the VHA BHS, this system was comprised of four hospital campuses; VHA West Roxbury has the ED, the other three campuses had urgent care units that also sent a significant amount of patients for admission to the VHA West Roxbury IU. Similarly VHA West Roxbury was a part of the greater New England Veterans Integrated Service Network, which was comprised of 10 hospital campuses (including the 4 that comprise the VHA BHS). Each of these hospitals may seek to transfer a patient for admission to VHA West Roxbury, either from their ED/Urgent Care departments or from their own IUs. Finally VHA West Roxbury was part of the greater network of hospitals in the local region. Private hospitals in the region would regularly transfer Veterans that have been brought to their EDs, or have been admitted, to VHA West Roxbury once the patient is stabilized. Each of these sources of admission did not individually outweigh the load sent from the VHA West Roxbury ED, but together comprised a significant number of admissions that do not necessarily follow the same trend as the VHA West Roxbury ED. It was therefore suggested that prediction numbers be generated for as many of these other sources as possible (most easily the three urgent care units that are part of VHA BHS). These additional predictions would avoid occasions when the prediction coming from the VHA West Roxbury ED gave the impression that the day would be lighter and, to the contrary, the IU ended up getting a heavy load from the VHA BHS urgent care units.

Another suggested network management use for the prediction was as a determinant for choosing to deny potential transfers from one of the many networked hospitals. Transfer requests would sometimes arrive before the IU was busy, however, after the transfer was accepted the ED bed requests would begin to rise and the bed manager would regret accepting the transfer. The prediction state enabled the bed managers to postpone or reject transfers based on knowledge of incoming demand. Finally it was suggested that the prediction states could encourage IU staff to connect with hospital outlets such as assisted living homes, rehabilitation centers, and other long

term care facilities, to improve flow out of the hospital, by engaging the outward end of the network.

5.3 Discussion

The earlier chapters of this dissertation have suggested that patient flow in the hospital setting could be improved by predicting a patient's likelihood of IU admission when they enter the ED sharing this prediction information with the hospital. Many studies have developed such prediction models; few have taken the next step, by discussing how the predictions should actually be used. Rather these studies default to the use of predictions as being coercive and replacing or preempting the decisions of ED physicians. In contrast this dissertation was inspired by the health care delivery chain conception of the ED/IU system to suggest a method for sharing predictions by aggregating patient admission probabilities into a total expected bed demand [Peck et al. 2012]. Based on the literature review performed for this study, no cases were found where a real time prediction system had actually been tested. One case did perform an interview based study, with a hypothetical prediction system, seeking an understanding of what would be needed to make it useful [Jessup et al. 2010]. However, based on preliminary interviews performed by the author of this dissertation, it was clear that without actually being faced with a prediction, hospital staff were unsure about what they would do with it. This was a complication, as an ideal prediction model would be tailored to the ways that it would be used, leading to a "chicken and egg" scenario when considering how to implement a prediction model.

For this reason, the prediction generation and sharing method described in this chapter was developed in order to provide a simple prediction system that favored less data rather than more. In this way practitioners would not be overwhelmed with data but would still have the opportunity to judge what value the predictions provided. Despite the simplicity of the model used, there was some complication in choosing the data categories that were displayed in Figure 30, as many people use different terminology. For example, while some would say a patient is boarding, others would call it a bed order. Similarly while some people refer to the inpatient units, others refer to the floors, or the inpatient wards. The final terminology used in the display was chosen by consensus of the investigation team and hospital management. Despite the complications, the trial was permitted to proceed with the belief that the data collected could act as a guide for future refinements of the prediction model and data sharing methodology.

Looking at the results of the trial shown, in Figures 32 and 33, it can be seen that there is a trend of improvement over the entire data collection period (6/1/2011 – 8/31/2012). This is consistent with the trends described in Chapter 3. Also, as in Chapter 3 there did not seem to be any notable seasonal effects on waiting or boarding time that would need to be accounted for in an analysis. With this in mind, looking at the trial period itself, there is no statistically significant difference between the two weeks of trial period, the weeks before the trial, and the weeks after the trial. In fact, the ED performance during the trial period seem to almost perfectly match that which would be predicted based on a historic trend, when no trial was attempted. This is not surprising given the short duration of the trial. As the interviews later confirm, staff members did start to find ways to use the predictions, however, the staff members did not have enough time to fully systematize these methods. Also those that were using the predictions did not have enough time to encourage colleagues to use them. However, the interview and survey data provide reason to believe that future work with a prediction method may yield the desired improvements.

Looking in more detail at the interview and survey results, there are some notable concepts that arise, and future work to be performed. The initial purpose of using a prediction model in the hospital system is to better enable admission planning. Based on the interviews it is clear that hospital staff did indeed consider using the prediction information for this purpose. However, as was discussed, many of the hospital staff members who attempted to pre-emptively work on the admission and discharge processes, found it difficult due to missing information, and were hesitant about taking actions that may be wasted effort. When faced with the decision to prioritize admission and discharge processes, versus treating patients, the health practitioner is in a situation similar to the newsvendor model that is commonly studied in operations management. The newsvendor model receives its name from an example case: in the morning a newsvendor must choose how many newspapers to stock for the day. The optimal solution to this issue comes by comparing the relative cost of overstocking and the opportunity cost of under-stocking to the probability that a certain number of newspapers will be purchased by the end of the day [Cachon and Terwiesch 2009]. Similarly a health practitioner must consider the probability that a patient will be admitted and decide whether that admission probability warrants risking the cost of over reacting (opening a bed when no bed ends up being needed) versus the opportunity cost of under-reacting (not opening a bed that does end up being needed) and causing a bottleneck until they

eventually get to it. Finding this optimal balance of over and under reacting is a clear direction for future work.

The value of prediction may increase when one considers how to better enable the potential benefits of resource planning, that were identified by the hospital staff. As was mentioned in the results of the interviews, many managers would have liked to move their staff, but did not have the power to do so. The goal of a real time prediction model is to enable real time reactions to the data that it generates. Reactions require resources. While hiring more staff to be on standby may be possible, it may also be possible to cross train staff, creating flexibility, such that a nurse or doctor can move between units as needed. In the case of physical resources, it may be worthwhile investing in more flexible treatment areas that can be made to fit a larger diversity of patient needs and reduce some of the potential costs of opening an inappropriate bed.

Personal planning was a common way of using the predictions expressed in the interviews. All interviewees occasionally looked at the prediction information, wanted to use it to make decisions about their own scheduling, and to prioritize their own duties. However, there was a definite variety in risk aversion to over reacting and under reacting as described above. The perceived value of making changes to one's schedule is unique to that person. Therefore, without a method of suggesting specific reactions to specific ED prediction states, the personal planning responses to predictions will remain inconsistent, reducing effectiveness.

Issues with resource alignment when sharing prediction information are tied to the variety in personal planning responses. This is because many groups of staff members must rely on each other in order to complete the discharge and admission processes. This reliance on others leads to all members of the hospital staff feeling that "you are telling the wrong person" when they are being pushed to facilitate a process. This situation can be summarized by considering a utility matrix as shown in Table 11 [Fudenberg and Tirole 1991, Pindyck and Rubinfeld 2009]. In this table we can simplify the hospital into two groups, the medical teams in the IU (comprised of residents, nurses, and interns) and the hospital support services (such as radiology, social workers, case managers, and the pharmacy). Each group can choose to prioritize discharge or treatment. If we assume that treatment is always prioritized in the current state, then we can say that the payoff of prioritizing treatment is 0, the base line. Then we can consider the case (bottom left of Table 11) where the medical teams react to predictions and prioritize discharge, but the

support services don't. Despite the efforts of the medical teams, the patient remains in the IU occupying a bed. In this case the support services still accomplished treatment work, thus they get the baseline pay off, but the medical teams wasted time by discharging when they could have been treating, thus they get a payoff of -1. This works in the opposite direction as well. When medical teams prioritize treatment (0) and support services prioritize discharge (-1). However, following the initial hypothesis of the value of prediction, if both groups prioritize discharge based on predictions and open a bed for predicted patients, then there may be a benefit above the baseline state for all parties, with a payoff of 1.

Table 11 Pay off matrix of prioritizing discharge versus treatment

		Medical Teams	
		Prioritize Discharge	Prioritize Treatment
<i>Hospital Support Services</i>	<i>Prioritize Discharge</i>	<i>1,1</i>	<i>-1,0</i>
	<i>Prioritize Treatment</i>	<i>0,-1</i>	<i>0,0</i>

As logical actors, it is likely that each of the groups will want to avoid the -1 payoff; therefore, they both stick to treatment, leaving the hospital at the baseline. However, it may be possible to force collusion, by having hospital management dictate specific instances when both groups should prioritize discharge. For example, patients who will soon be ready for discharge can be given red flags in the computer system. If a certain prediction state is reached in the ED, hospital management can force all groups to prioritize patients with red flags. This makes it clear how to expedite the discharge patients for all parties and may create the optimal situation. This is similar to expediting concepts used in manufacturing and the theory of constraints [Goldratt and Cox 2004]. Actually finding the correct ED prediction state that warrants encouraging all parties to expedite discharges would require more research and is the focus of Chapter 6.

Finally, the potential for hospital network management is an unexpected outcome of the prediction system implementation. Currently, network effects may be limited for many hospitals. However, the size of hospital networks continue to grow, in the US, as hospitals seek to gain economies of scale [Cuellar and Gertler 2003]. Therefore, the use of real time predictive models

can enable a whole new level of bed control on the network level. This could be used to inform when and to where patient transfers should occur and improve the ability for related hospitals to level their loads and share resources in an optimal way. The study of optimal hospital network management is another significant area for future research.

5.4 Limitations

There are some notable limitations with the study described here. The first comes from the limited length of the prediction trial. One of the consistent difficulties with research in a hospital environment is the potential to hinder the treatment of patients. The ability to try a “what if” scenario without damaging the system is the reason simulation methodologies are popular [Banks et al. 2010]. While simulations are helpful, a live implementation often provides insights that a simulation cannot. However, the implementation experiment was limited due to the need for nurses to enter extra data into the prediction system. Ideally, all of the data would be captured by the current health care information system used at the hospital and this would automatically feed into a prediction algorithm, thus removing the need for extra data entry and enabling a longer trial. Efforts are being made at VHA West Roxbury to achieve this.

Another limitation of this study is the fact that sampling was used for choosing interviewees. Ideally all hospital staff members could be interviewed and this would provide a more complete understanding of how the prediction information was or could be used. However, such a study may also place an undue burden on the hospital. Nevertheless, the current study protocol was able to capture a wide range of opinions and derive value from the data set.

5.5 Conclusions

Chapters 2 through 4 were based on the premise set forth by past studies, suggesting that systems to predict, when a patient enters the ED, whether that patient will require admission to the hospital could be used to encourage IU staff to facilitate that future admission and improve flow/quality. Some studies were found that made these predictions and Chapters 2 and 3 described how such predictions could be made and shared. However, no studies were found that focused on establishing whether such a system would indeed improve flow, nor understand the realities of developing such a system.

The study described in this chapter was an initial attempt towards implementing and studying an admission prediction system. For two weeks predictions were made in the ED and shared with the IU, during this time no change was found in key patient flow metrics. This was likely due to the short time period of the trial.

In order to capture qualitative effects of sharing predictions a sample of hospital staff in multiple different positions around the hospital was taken and interviewed. The results of these interviews showed potential uses for prediction in terms of admission planning, resource scheduling, personal scheduling, resource alignment, and hospital network management.

The results of this study suggest that the initial belief that prediction could be used to improve flow in the ED/IU chain may be true. While the results do not actually quantify any flow benefit from prediction, the interviews suggested some human and system factors that prevented the achievement of any potential benefits. These factors may be overcome through future work. Chapter 6 will seek to gain a fuller understanding of the use of prediction in this system through the use of computer simulation. Chapter 6 will describe the use of simulation in health care as well as the development of a model for capturing the behavior of the ED/IU system, with some simplifying assumptions. This model will show the maximum potential value for prediction in the ED/IU delivery chain and offer insights into how to achieve these results in reality.

Chapter 6: Simulation of the Emergency Department Prediction Model

In Chapter 2 and throughout this dissertation it has been established that metrics of quality in an ED are defined by how quickly a patient gets to and through their emergency treatment [Graff et al. 2002, Bernstein et al. 2009, Horwitz et al. 2010]. This characterization of quality in the ED has led to this dissertation's adoption of ED waiting time as a primary measure of ED flow improvement. It was also discussed that the output process of admitting a patient to an IU has significant impact on ED waiting time [US GAO 2003, Olshaker and Rathlev 2006, Falvo et al. 2007, US GAO 2009]. This led to the adoption of ED/IU boarding time as the second primary measure used in this dissertation. The two measures combined, provide a sense of how patients flow through the entire ED/IU health care delivery chain.

The previous chapters adopted the suggestion that prediction of IU admission early in a patient's ED treatment can be used to improve flow in the ED/IU chain. Chapters 3 and 4 discussed how to make these predictions and the potential to operationalize these predictions, by generating a running bed demand value that combines the predictions assigned to all patients in the ED. Chapter 5 was focused on how this running bed demand may be interpreted and used in a live implementation. The results of the implementation suggested some interesting applications for the prediction system, however, there was no improvement found in the key flow metrics. It was suggested that the lack of notable flow improvements could be due to the limited implementation length. It was also suggested that lack of notable improvements could be due to a lack of alignment between the many treatment and support services teams throughout the hospital. The goal of this chapter is to gain a fuller understanding of the potential for prediction in the ED/IU chain by using computer-based, discrete event simulation (DES) to create a model of the system that can be studied without time restraints and where resources can be better controlled.

6.1 A conceptual model of the ED/IU “Pull” System

Creating a DES of a system is a process. This process begins with the problem entity, or system to be modeled, which is the ED/IU chain. The next step in the process is the conceptual model. The conceptual model “is the mathematical/logical/verbal representation of the problem entity

developed for a particular study” [Sargent 2000]. In Chapter 2, Figures 12 through 14 were depictions of the ED/IU system with varying details. Each of these figures can be considered a different conceptual model of the ED/IU system and each would result in a different computerized model, the creation of which is the last step in the modeling process. With this in mind, it is necessary to discuss the conceptual model of the ED/IU system that enables the use of predictions; this lays the conceptual structure for the computer model that was generated and will be discussed.

Those who have studied the ED/IU health care delivery chain and who have some understanding of process improvement methodologies often suggest that the ED/IU system be improved by becoming a pull system. A traditional pull system (Figure 34a) is based on sharing information to closely match upstream production with downstream demand [Hopp and Spearman 2001, Simchi-Levi et al. 2003]. In typical supply systems, such as grocery stores, an order originates downstream (at the store) and then the downstream step pulls a new item from the resource buffer at the upstream step (manufacturing or storage facility), triggering the production of a new product to replace the one that was taken. If this method were to be directly adopted in the ED/IU system, it would be expected that the IU would not wait for a bed to be ordered by the ED. Instead, the IU would preemptively seek to open beds with the plan of pulling a patient from the ED. Although this may initially sound reasonable, there is a practical limitation. The ED does not have a controlled stock of patients and without knowing that patients are waiting in the ED, the IU staff does not have the incentive to open a bed and pull a patient. Therefore, flow in the ED/IU health care delivery chain lacks the information and incentive structure to enable a traditional pull system.

By graphically representing the flow of patients through the ED/IU system (Figure 34b) it can be seen that demand originates upstream rather than downstream. Therefore the downstream step, IU, would not know to pull a patient; IU staff does not know the patient exists. Instead the IU staff will focus primarily on the real demands of patients already in the wards, who need treatment. In response to this issue, the health information technology systems within a hospital can be leveraged in order to create an information based pull system. This type of pull system is really an information-based push system, where the upstream step sends a signal to the downstream step that it will be producing a product/patient and that the downstream step should

be prepared to pull the patient when ready. This system differs from a normal push system in that it gives the downstream step some time to prepare and control the flow to a certain degree. This is the conceptual basis for using prediction in the ED/IU chain and is depicted in Figure 34c. The running IU bed demand value, proposed in the earlier chapters, serves as the signal that suggests that the IU pull patients from the ED. It has been proposed in this dissertation and in previous studies that sharing this prediction information to create an information-based pull system could improve the bed coordination process and improve flow [Yen and Gorelick 2007, Sun et al. 2011, Peck et al. 2012]. Using the conceptual model in Figure 34c the rest of this chapter focuses on characterizing the patient flow benefits of sharing predictive information between an ED and IU. This characterization begins by discussing the development and structure of the DES model. Then there will be a discussion of the logic and results of scenarios for exploring how a hospital's health information system can be leveraged in order to share ED state information and how sharing this information with the hospital may affect ED patient flow and, consequently, treatment quality.

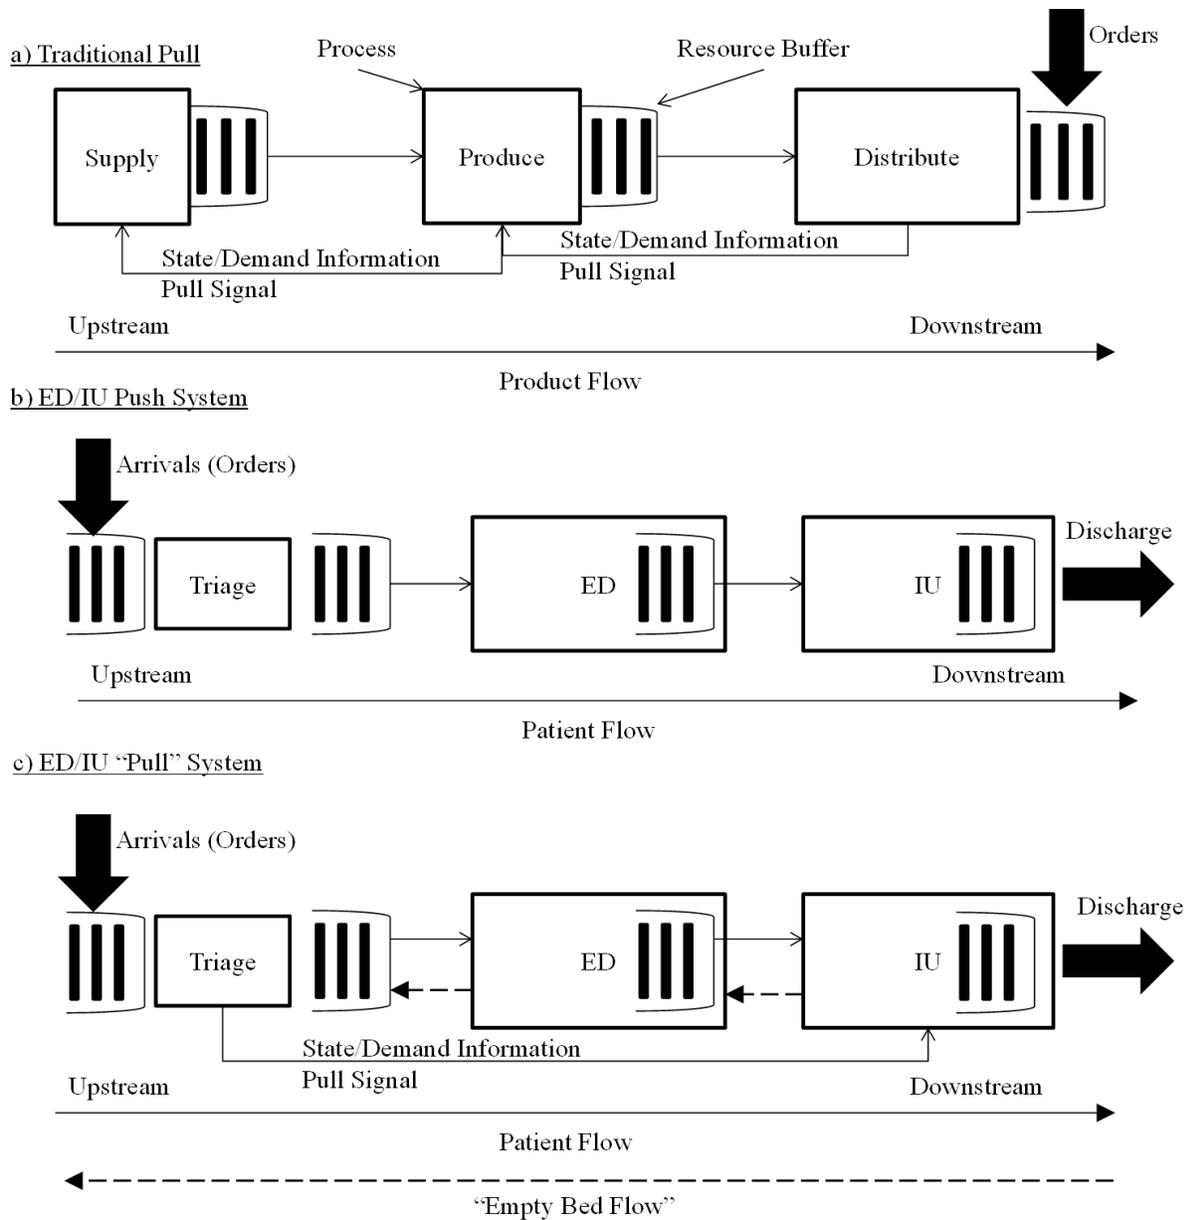


Figure 34 a) Traditional pull system, b) current ED/IU push system, c) potential ED/IU information based pull system

6.2 Methods

To some, creating a prediction based pull system in the ED may have intuitive benefits. However since human lives are at risk when changes are made in health systems, simulation is a popular tool for exploring “what if” scenarios [Jacobson et al. 2006]. One type of simulation that is commonly used for studying the ED/IU system is DES [Baesler et al. 2003, Connelly and Bair 2004, Kolb et al. 2008, Li and Howard 2010, Paul et al. 2010, Peck and Kim, 2010]. The DES model in this study was built in Rockwell Automation, Inc.’s ARENA DES software version 13.5. The model is based on the ED/IU chain at VHA West Roxbury. VHA West Roxbury has a 13 bed ED which received approximately 1200 patients per month in 2010. The hospital IU has a varying resource level of approximately 170-180 staffed beds. The IU beds are shared between elective admissions from local VHA clinics, transfers from other VHA hospitals, veteran transfers from non-VHA hospital facilities, and local veteran and non-veteran emergency patients.

The simulation model assumes 100 beds are reserved for ED patients. The logic of the model is shown in Figure 35 and is comprised of four primary sub-models: Arrival, the Emergency Department, the Inpatient Unit and Bed Management.

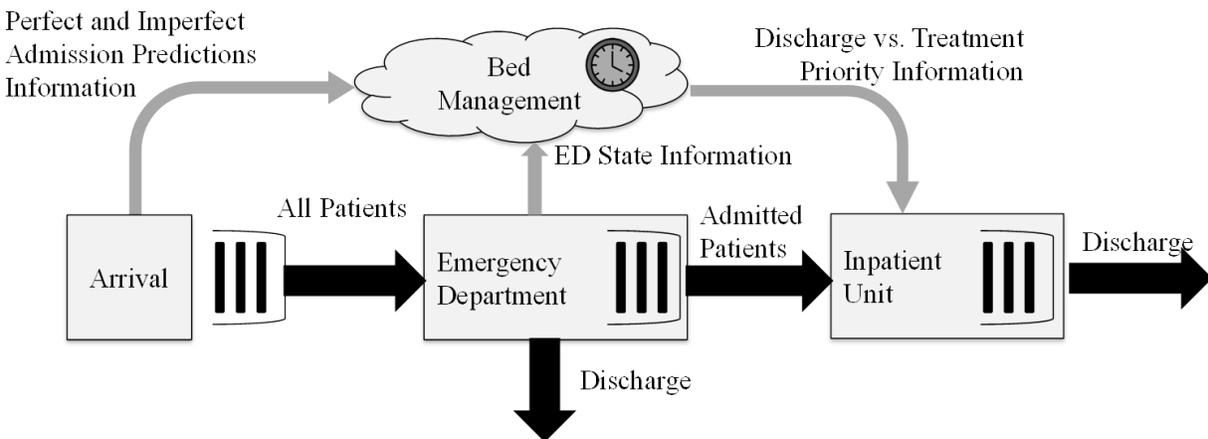


Figure 35 Discrete event simulation model logic

The arrival sub-model consists of a creation module that generates patients based on the actual patient arrival pattern, derived from the VHA West Roxbury data. After a patient is created, they are assigned a probability of being admitted. This was done by creating a probability distribution

of probabilities of being admitted, based on the logit-linear regression predictions on the test dataset from Chapter 3. A simulated patient is assigned a probability of admission from this distribution which was fit, using Arena's input analyzer, to be a beta distribution with the following equation:

$$P(\text{admission}) = .94 * \text{BETA}(0.345, 0.878)$$

Rather than use predictions as a yes/no diagnostic tool, the model uses the suggestion of the earlier chapters and adds patient probabilities together, generating a running expected bed demand. The running bed demand could be shared with the IU, enabling the accuracy benefits of risk pooling. The model calls the running bed demand, based on the logit-linear regression predictions, the imperfect predicted index of demand for IU beds from the ED. Patients then move through a decision module which assigns whether the patient will indeed require admission using that patient's assigned probability value. This predetermined decision is used to generate a perfect predicted bed demand index by counting all patients who were chosen for admission. Both of the bed indexes are shared with the bed management sub model described later.

In summary, when the arrival process is complete, patients have two attributes, and the system has two running variables.

Attributes:

1. Imperfect predicted Bed Need, a continuous probability from 0-1
2. Perfect predicted Bed Need, 1 or 0

Running Variables, which are the sums of the respective prediction attributes over all patients in the ED:

3. Total imperfect bed demand
4. Total perfect bed demand

Upon receiving their admission predictions, patients enter the ED which is comprised of 13 beds, just like the VHA West Roxbury ED. Patients then seize a bed for their treatment duration, again the length of this treatment is based on the distribution of treatment times observed at the actual

hospital, using the test dataset from Chapter 3. Arena's input analyzer fit an Erlang distribution to this data as follows:

$$\text{ED treatment time} = -0.001 + \text{ERLA}(78.2, 2)$$

While the simulation is running, the number of ED beds that are full is being tracked and shared with the bed management module. After completing their ED treatment, those patients who were predestined for admission enter a queue to seize an IU bed while continuing to hold an ED bed. At this point the patient's imperfect admission prediction is updated to a 1 while discharged patient admission predictions are reduced to 0 and the admission indexes are updated accordingly.

The IU sub model contains 100 beds based on the assumption that a significant number of VHA West Roxbury's 170-180 beds are reserved for elective admissions. To capture how information can affect decisions, and consequently flow, the model also assumes that the doctors are the decision maker and limited resource in the IU and that all other support services have unlimited capacity. This eliminates the issue of coordinating priority between the medical teams and support services. Figure 36 is a representation of the model logic. As can be seen in the figure, a patient first seizes a doctor for treatment. The patient can only be treated or discharged by this unique doctor from that point on. The patient then releases the doctor and goes through some randomized amount of value added treatment. At the end of the cycle, the amount of time the patient spent is deducted from the patient's total value added IU length of stay (LOS). The value added length of stay was calculated by analyzing 32,156 patient visits to the VHA West Roxbury IU, spanning all visits from 10/02/2008 to 6/30/2011. The Arena input analyzer fit the length of stay data to a log-normal distribution as follows:

$$\text{Patient IU LOS} = -0.001 + \text{LOGN}(8.89, 17.1)$$

The LOS dataset was based on the true LOS of patients and included non-value added (NVA) waiting times. Since the model needed to capture changes to this NVA time, the LOS assigned to a patient was divided by a LOS reduction variable. The value of this variable was chosen to be 2.5 based on calibration efforts in the validation stage, using a baseline scenario (Scenario 1) that will be described later.

The simulated patient continues to go through this cycle, of doctor treatment and value added treatment, until they have depleted their assigned LOS. At this point, rather than re-enter the treatment queue, the patient enters the queue to seize a doctor for discharge orders. In this way those patients waiting for discharge are in direct competition for doctors with patients who are still receiving treatment. To manage this competition, the bed management module has the ability to shift priority between the two processes, when certain conditions are met; these conditions are based on the scenarios described below. Doctors only accept patients from 7am to 8pm, at 8pm the doctors will finish processing patients that are in the queues but all others are held back until the next day.

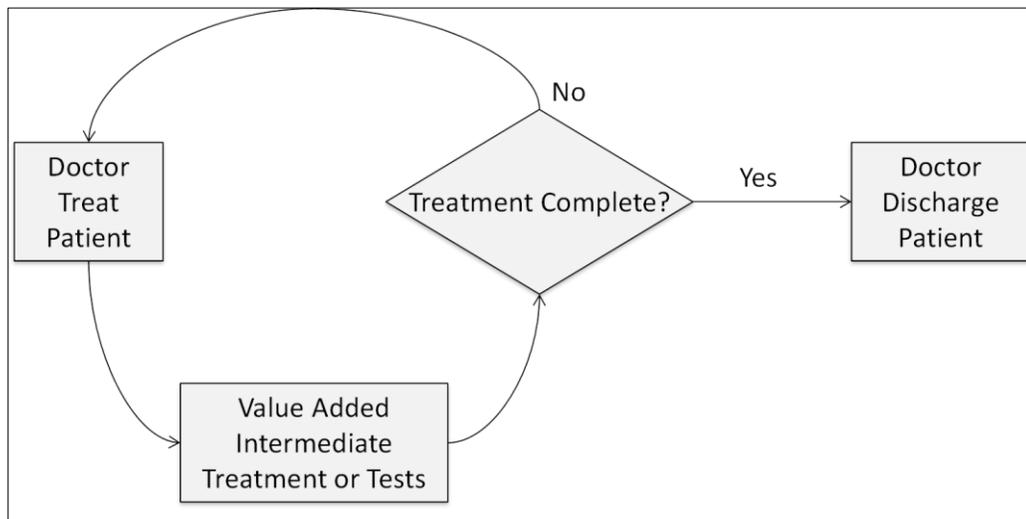


Figure 36 Doctor Decision Cycle

6.2.1 Simulation Scenarios

Six primary scenarios were studied using this simulation model, below these are described in words and in equation form where:

$T = \text{time}$

$ED_n = \text{designates ED bed } n \text{ where } n = \{1,2,3,4,5,6,7,8,9,10,11,12,13\}$

$I = \text{number of patients in the ED or waiting room}$

$PT_i = \text{designates patient } i \text{ in the ED or ED waiting room } i = \{1,2,3 \dots I\}$

$IUBeds = \text{available IU beds}$

$S(t)$ = designated time varying sensitivity level

$$Prio = \begin{cases} 0 & \text{IU doctors' priority is set to treatment} \\ 1 & \text{IU doctors' Priority is set to discharge} \end{cases}$$

Scenario 1: This is the baseline scenario where priority is set to discharge beginning at 1pm and ending at midnight.

$$\text{if } T \geq 13 \text{ then } Prio = 1 \text{ else } Prio = 0$$

Scenario 2: At a predetermined time of day priority is set to discharge for three hours.

for $t = 0$ through 23

$$\text{if } T \geq t \text{ or } T < t + 3 \text{ then } Prio = 1 \text{ else } Prio = 0$$

Scenario 3: Priority is set to discharge while a time varying designated difference between the number of occupied ED beds and available IU beds, or Crowding Index, has been reached or exceeded.

$$F(ED_n) = \begin{cases} 0 & \text{if ED bed } n \text{ is occupied} \\ 1 & \text{if ED bed } n \text{ is empty} \end{cases}$$

$$\text{Crowding Index} = \sum_1^{13} F(ED_n) - IUBeds$$

$$\text{while } \text{Crowding Index} \geq S(t), Prio = 1 \text{ else } Prio = 0$$

Scenario 4: Priority is set to discharge while a time varying designated difference between the imperfectly predicted IU bed demand and IU bed availability, or Imperfect Index, is reached or exceeded.

$P(PT_i)$ = Imperfectly predicted probability that patient 'i' will be admitted, 1 if patient has completed ED treatment and is awaiting admission.

$$\text{Imperfect Index} = \sum_1^I P(PT_i) - IUBeds$$

$$\text{while } \text{Imperfect Index} \geq S(t), Prio = 1 \text{ else } Prio = 0$$

Scenario 5: Priority is set to discharge while a time varying designated difference between the perfectly predicted IU bed demand and IU bed availability, or Perfect Index, is reached or exceeded.

$$Admit(PT_i) = \text{Perfect prediction that patient 'i' will } \begin{cases} 0 & \text{not be admitted} \\ 1 & \text{be admitted} \end{cases}$$

$$Perfect\ Index = \sum_1^I Admit(PT_i) - IUBeds$$

while $Perfect\ Index \geq S(t)$, $Prio = 1$ else $Prio = 0$

Scenario 6: The current best practice of discharging by noon where discharge is prioritized for any time before noon.

for $t = 0$ through 23

if $T < 12$ then $Prio = 1$ else $Prio = 0$

Each of these scenarios can be tested for sensitivity, using factors that have been built into the model. Sensitivity was studied through three cases. Case 1 had no non-value added (NVA) admission delay and 25 IU doctors, making a four to one patient to provider ratio. The second case had the same patient to provider ratio but had a variable NVA delay, between the ED and IU, which is normally distributed with a mean of 30 minutes and a standard deviation of 15. This delay occurs after an IU bed is assigned, but before the ED bed is released. This delay can be interpreted as delay of ED staff in receiving the assignment, delay of hospital bed managers from communicating the assignment, extra cleaning requirements, room set up delay, transportation delay, or many other possible sources of delay. The third case, has no NVA delay but changes the patient to provider ratio to five to one by reducing IU doctor capacity to 20.

6.2.2 Calibration and Validation

There are multiple frameworks for validating a simulation model [Balci 1995, Sargent 2004, Banks et al. 2010]. To validate the model in this study, the authors relied primarily on face validity and historical data validation. Face validity was established by presenting the model to medical experts to get their opinion on the logic. Historical Validity is established by looking at the outputs of the model and comparing them to the true VHA Boston data.

As the simulated system processes input data, it is likely that the outputs will begin to diverge from the true system data. While it is unlikely to achieve a perfect match between the real and simulated systems for all outputs it is important to judge whether the two remain close. Figure 37 shows the pattern of patient arrivals to the ED, these are, and should be, almost exactly the same as the input data, since no processing has been done to the patients at that point.

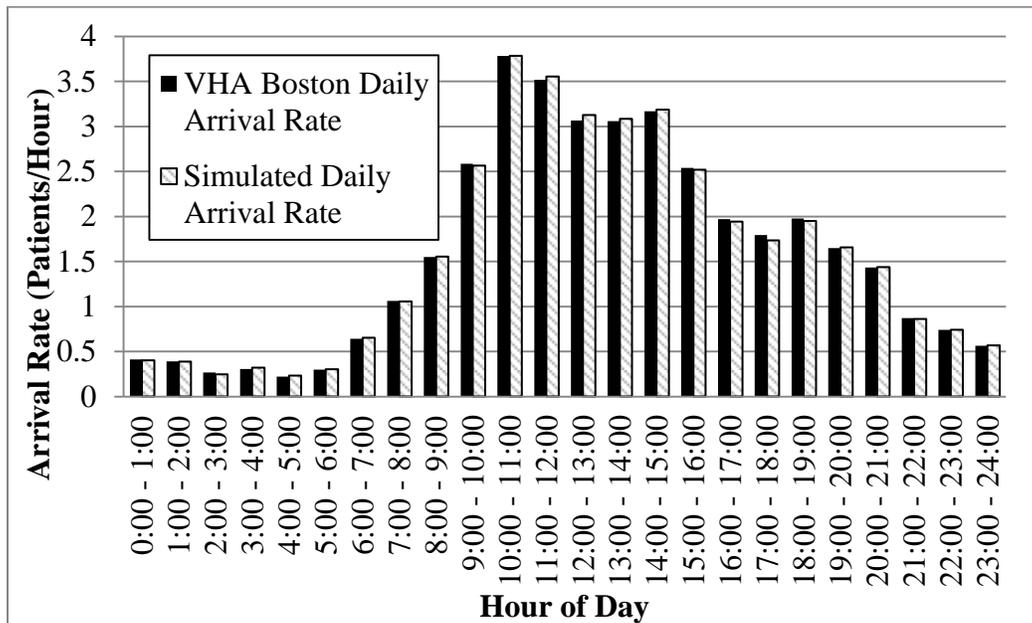


Figure 37 Simulated and real VHA Boston daily patient arrival rates by hour

After some processing in the ED, admission requests from the ED to IU are generated. Figure 38 shows the pattern of requests for the simulated and real system. As can be seen there is a difference between the simulated and real admission request rates. This difference is likely due to the fact that the model sought simplicity in the ED module and did not differentiate ED LOS for admitted patients and non-admitted patients when creating the LOS distribution.

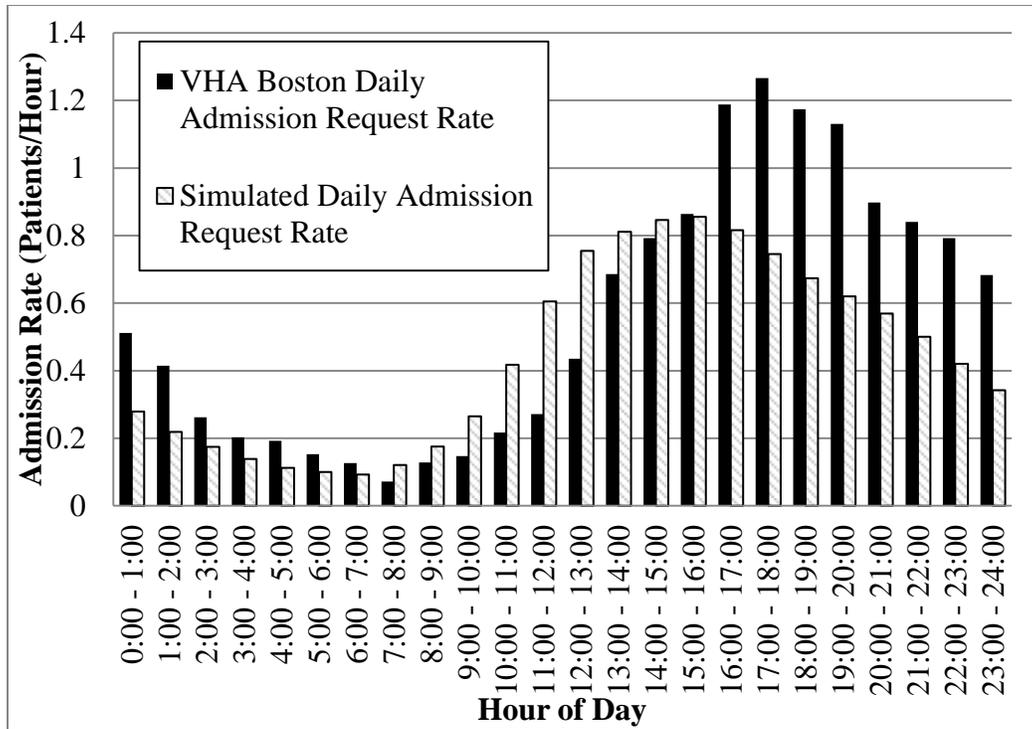


Figure 38 Simulated and real VHA Boston daily admission request rates by hour

Those patients who are admitted to the IU proceed through the process described in Figure 36. The simulation is a simplified model of the true IU system and therefore is likely to distort resulting simulated outputs from the true outputs. Figure 39 shows the hourly IU discharge rates for the real and simulated systems. Despite the simplified model of the IU, these patterns seem close in shape.

Figure 40 shows the IU LOS pattern for the real and simulated systems. Note that Figures 39 and 40 are normalized. The VHA West Roxbury discharged patient data collected did not include patient origins, thus it was not known how many came from the ED versus other sources, making a direct comparison of rates impossible.

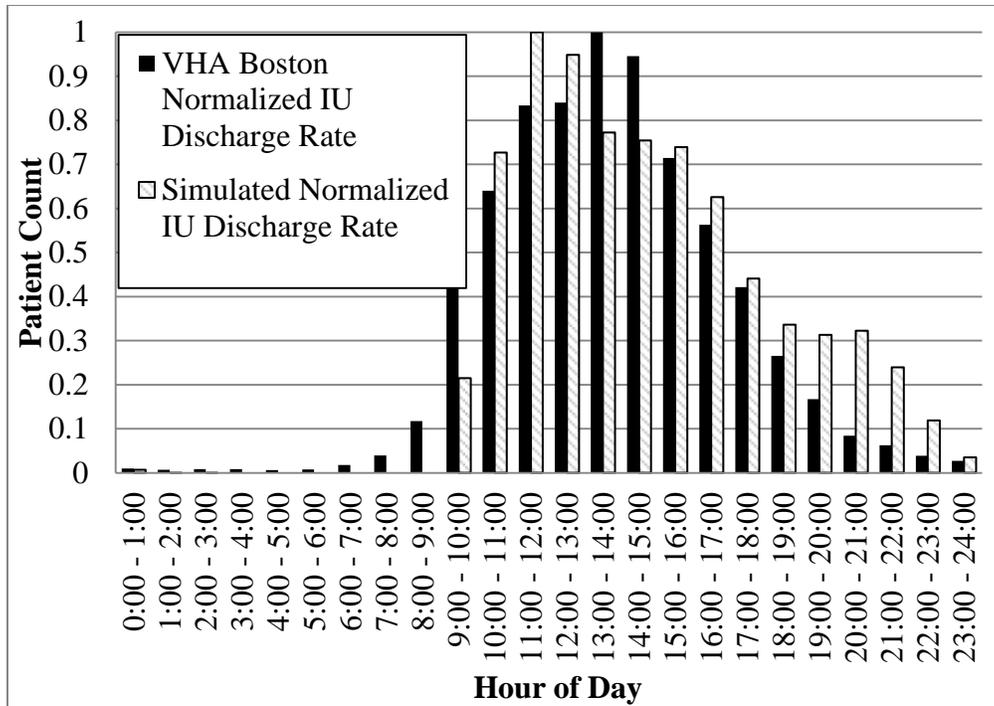


Figure 39 Simulated and real VHA Boston daily, normalized, IU discharge rates by hour

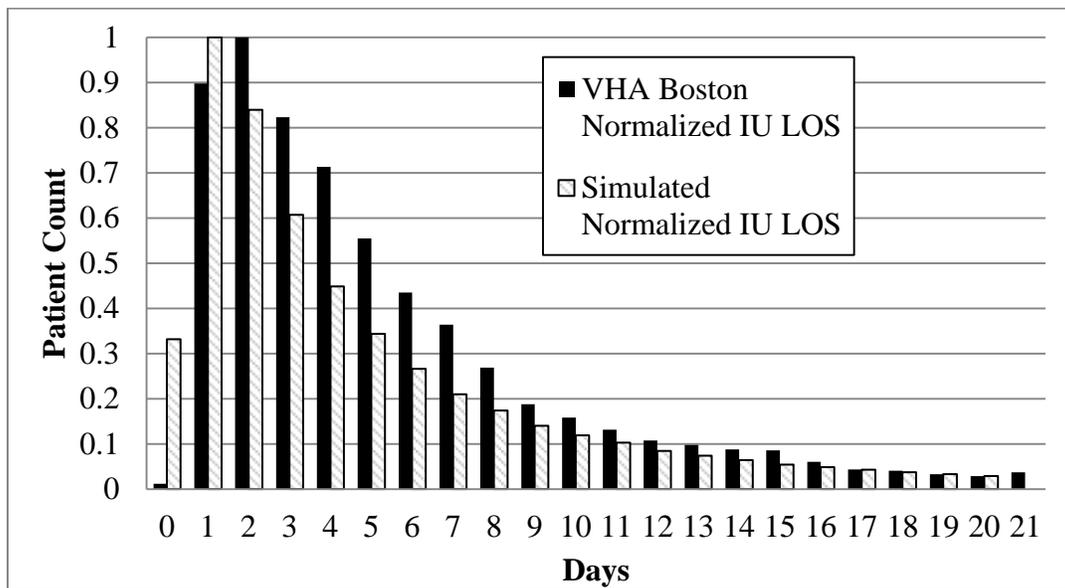


Figure 40 Simulated and real VHA Boston patient IU LOS

Finally, Table 12 compares key performance and demographic values for the real VHA West Roxbury system and the simulation, for the case when discharge priority begins at 1pm and ends at midnight (Scenario 1).

Table 12 Simulation output data vs. VHA Boston monthly average data for validation

	West Roxbury	Model	Units
Total Patients	1240.5	1137.7	Patients
Percent Admitted	34	28	%
ED Wait Time	0:17	0:11	Hours
ED LOS Admitted	4:01	3:05	Hours
ED LOS No Admit	2:19	2:36	Hours
Boarding time	0:28	0:28	Hours
IU LOS	10.2	7.53	Days

The simulated results closely match or are similar to the actual values. This is an important confirmation that the simulation is using the correct internal data and assumptions to generate its final results. The Simulation results are dependent on the emergent behavior of the cyclical IU model and other internal processes. Therefore, although the model is a simplification of the true system, the data being generated by the model reflects that of the true ED/IU system and serves as evidence that the model is a reasonable representation of the system. The validated baseline suggests that the results of other scenarios would also have some resemblance to the results that would be generated by running these scenarios in the real ED/IU system. Naturally, with all models, there is no guarantee of this and one must make the tradeoff between value of the model and cost of further development and refinement [Sargent 2004].

The validation figures and table show that the pattern for the real hospital and for the simulated hospital are not exactly the same, however, the simulated pattern is not unreasonable for a realistic fictional hospital. Although this means that the model is not a perfect fit for the VHA West Roxbury ED/IU system, the results are close enough to suggest that the real and simulated systems have similar dynamics, therefore simulation scenario results may be applicable to the real system, with adjustment.

6.3 Results

As described earlier, all results are shown for three cases:

Case 1: the default case, where there was full IU doctor capacity of 25 and no NVA delay.

Case 2: where there was a 30min NVA delay after a patient was assigned an IU bed but before they release their ED bed.

Case 3: where IU doctor capacity was reduced to 20 but there was no NVA delay.

All scenarios were run for each case. The error bars in each figure represent the 95% confidence interval for the data point based on 5000 replications of the simulation.

For each of the three cases, Figure 41 shows how IU boarding time changes with the time of day that discharges are emphasized. Each discharge emphasis shift lasts for 3 hours, this is Scenario 2 described earlier. The figure suggests that emphasizing discharge is more valuable early in the day and detrimental later in the day, just as has been asserted by the popular discharge by noon heuristic. Thus, when considering sensitivity levels to the crowding and predicted indexes (as described in section 6.2.1), it would be logical to want a greater sensitivity earlier in the day and a lesser sensitivity later, however this may have exceptions based on the dynamics of the simulated hospital.

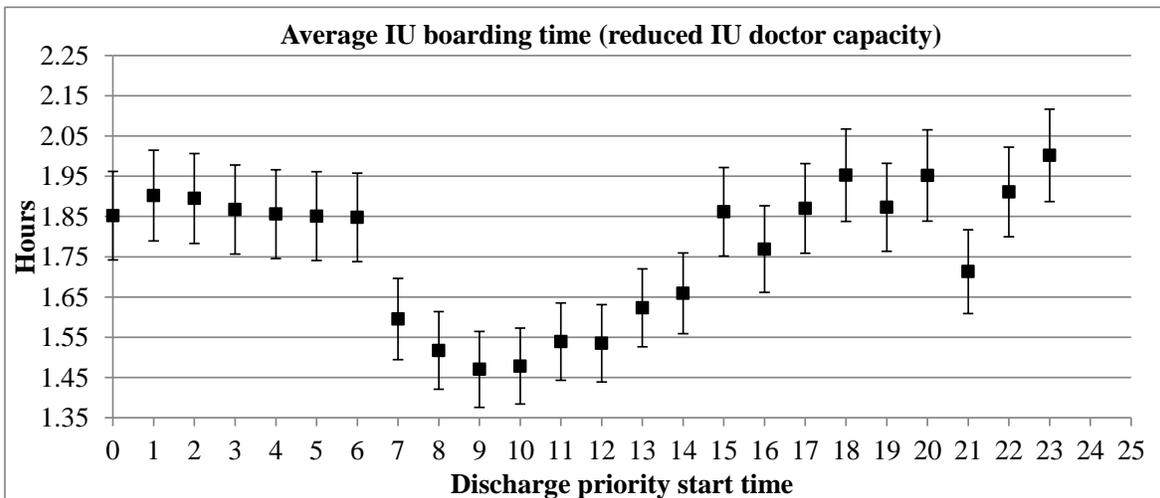
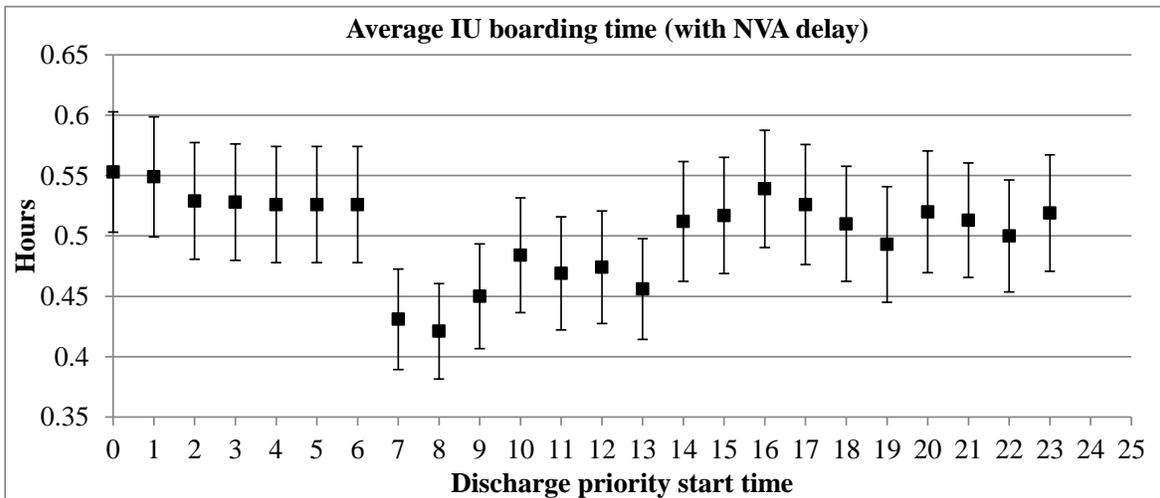
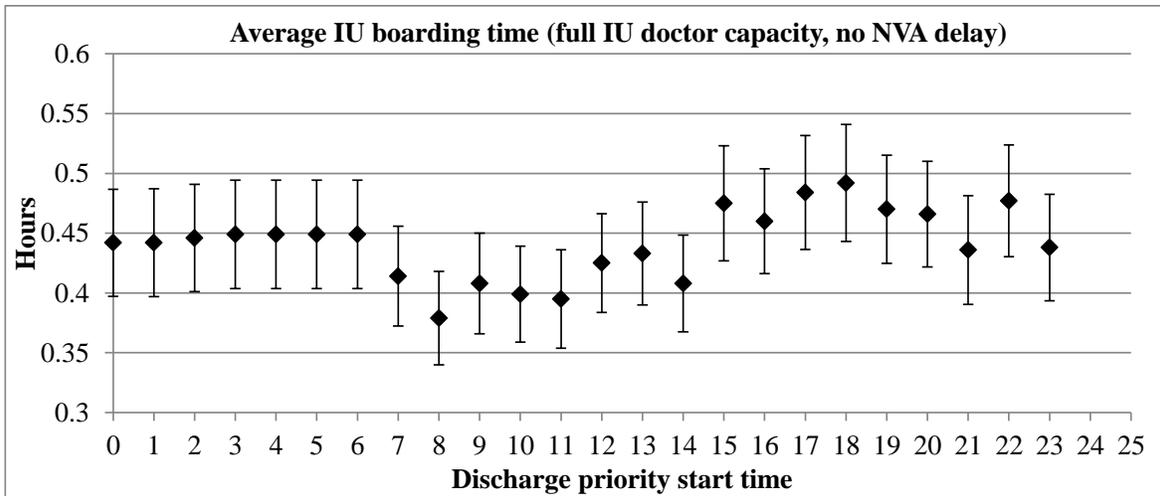


Figure 41 Average IU Boarding time (and 95% confidence intervals) with shifting 3 hour discharge priority start times without an Case 1 (top), Case 2 (middle), Case 3 (bottom).

To explore the potential for varying sensitivity to crowding and predicted admissions throughout the day, the simulation was built to enable variation in index sensitivity, $S(t)$, where a different S can be chosen for each hour. In order to approach the optimal daily schedule of sensitivities, the sensitivity level for each hour was entered as a separate variable into the simulation. These variables were then entered into the optimization software built into Arena, OptTek Systems Inc.'s OptQuest for Arena. The optimization objective was to minimize boarding time (the time that an ED patient waits in the queue to receive an IU bed). Figures 42, 43 and 44 show the optimized sensitivity schedules when using the ED Crowding Index, Imperfect Index and Perfect Index (Scenarios 3-5 above) in each of the three cases.

For comparison of quality outcomes, Figure 45 shows the IU boarding and ED wait times in each of the three cases for the baseline scenario (Scenario 1), the optimized index scenarios (Scenarios 3-5), and the discharge by noon scenario (Scenario 6). Tables 13, 14, and 15 show the results of evaluating the significance in the difference between the average IU boarding times for the optimized, time based and baseline scenarios.

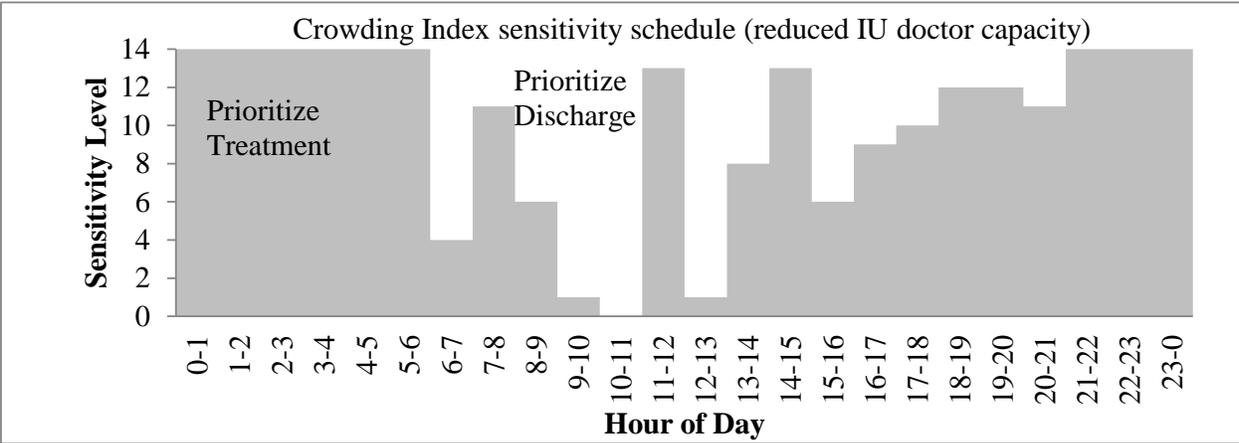
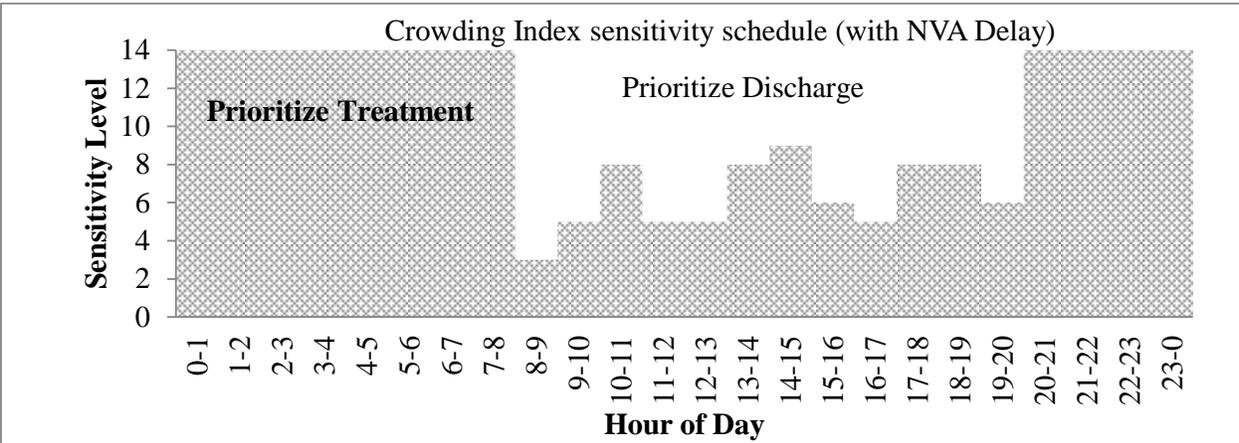
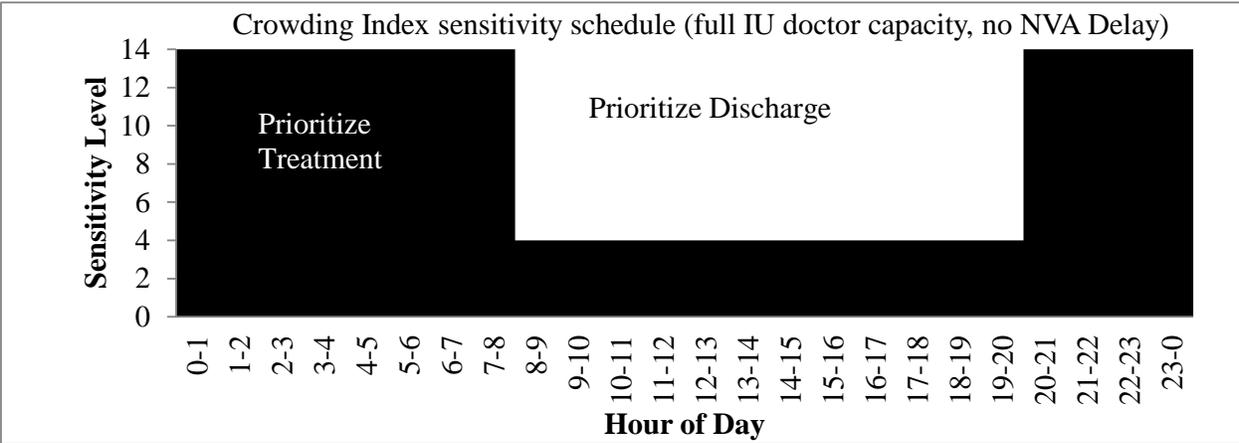


Figure 42 Optimized sensitivity schedule using Crowding index for Case 1 (top), Case 2 (middle), and Case 3 (bottom)

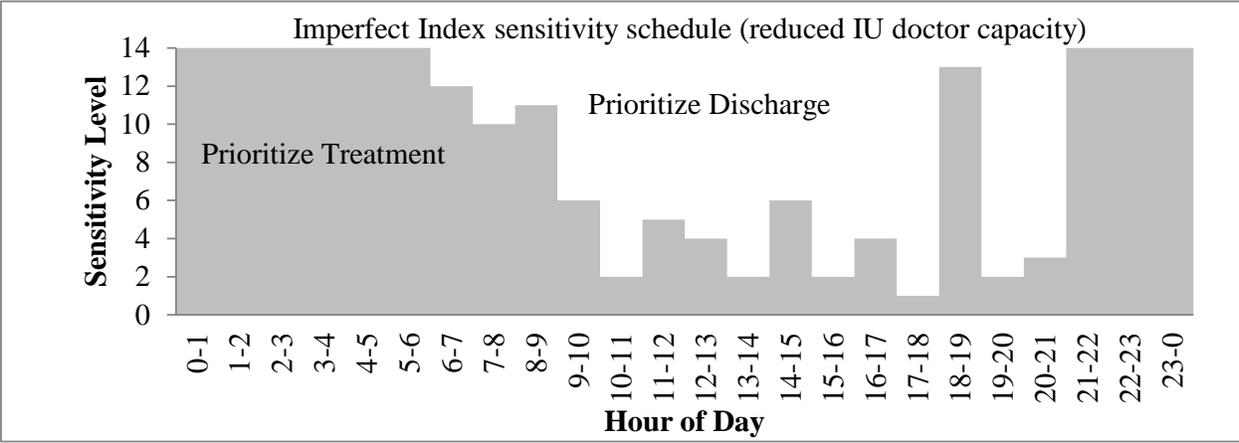
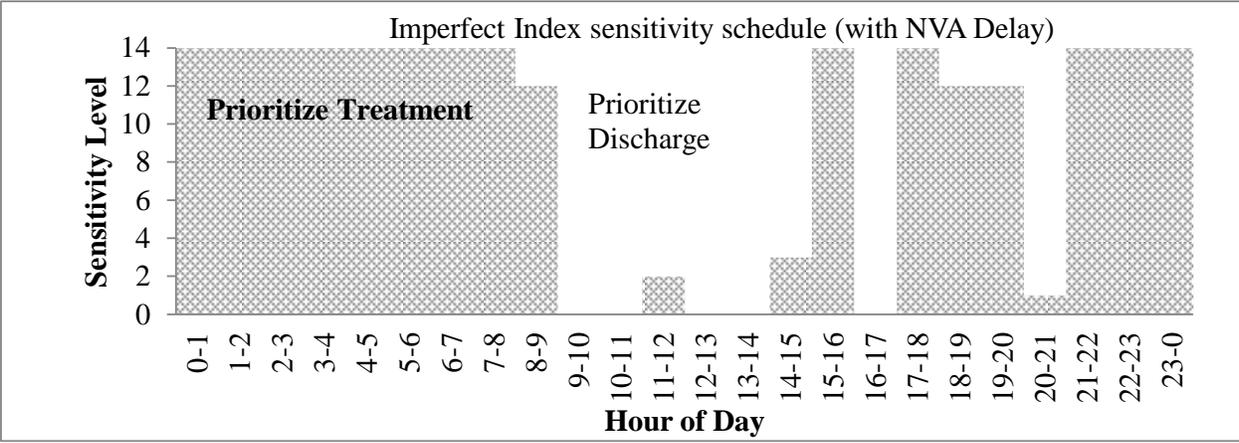
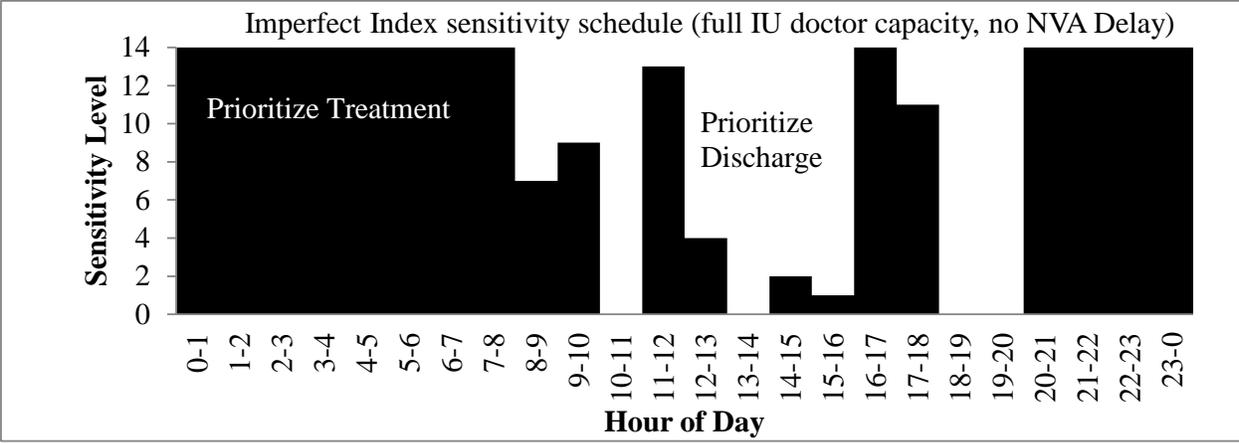


Figure 43 Optimized sensitivity schedule using Imperfect index for Case 1 (top), Case 2 (middle), and Case 3 (bottom)

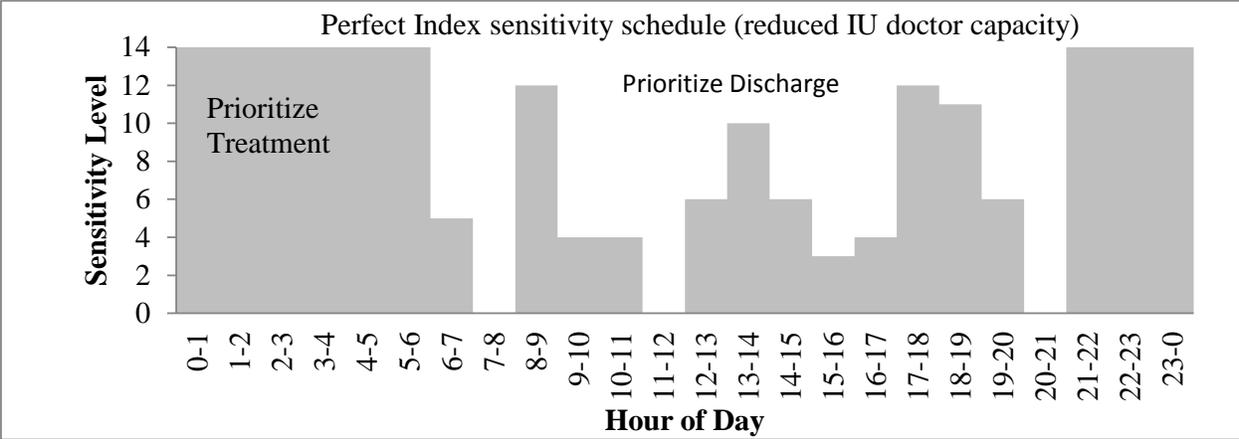
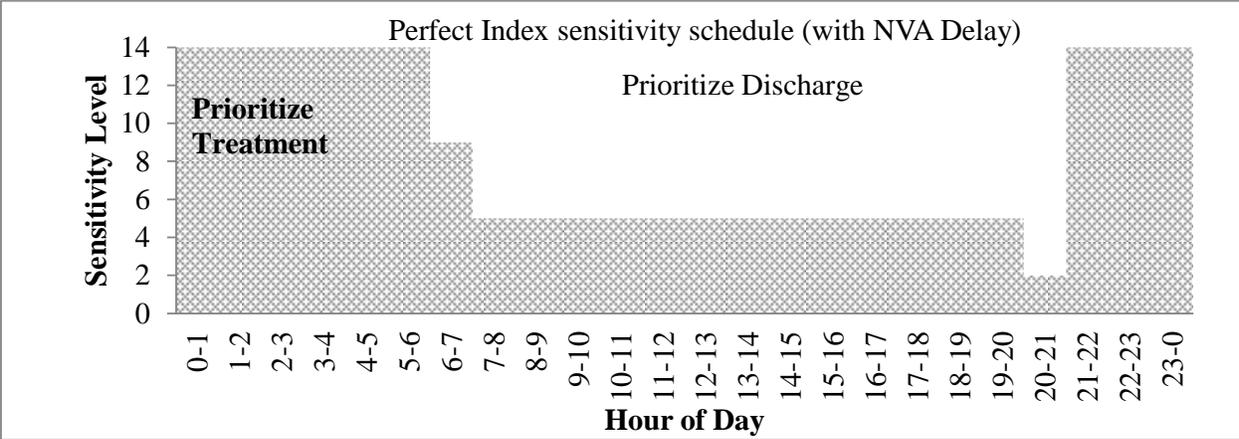
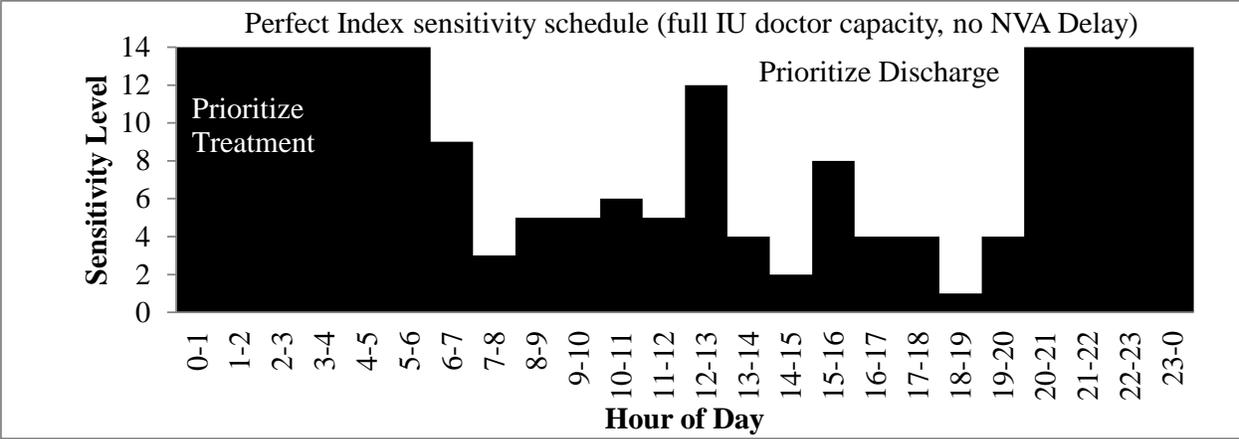


Figure 44 Optimized sensitivity schedule using Perfect Index for Case 1 (top), Case 2 (middle), and Case 3 (bottom)

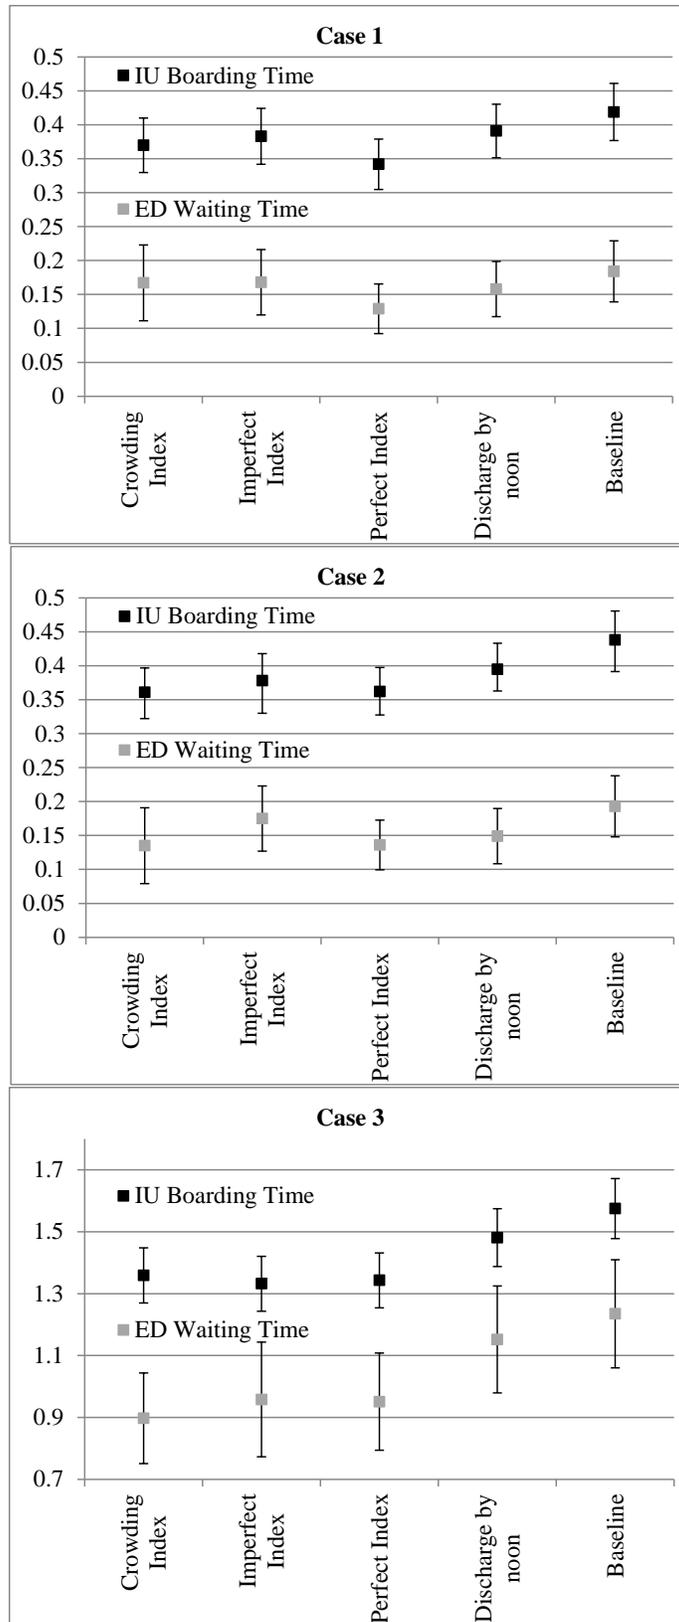


Figure 45 IU wait time and ED wait time for optimized scenarios and time based scenarios

Table 13 Difference in average IU boarding times between scenarios: Case 1

Case 1: Average time difference (minutes)*, p-value**					
* $\mu_x - \mu_y$, x = column, y = row, μ = average wait time					
**Hypothesis: $\mu_x = \mu_y$ for IU boarding times with no NVA delay, rejecting $p < 0.05$					
	Crowding Index	Imperfect Index	Perfect Index	Time Based	Baseline
Crowding Index	=				
Imperfect Index	NA, 0.32	=			
Perfect Index	1.68, 0.02	2.46, << 0.05	=		
Discharge by Noon	NA, 0.10	NA, 0.54	-2.94, << 0.05	=	
Baseline	-2.94, << 0.05	-2.16, << 0.05	-4.62, << 0.05	-1.68, 0.03	=
Improvement over Baseline	11.69%	8.59%	18.38%	6.68%	=

Table 14 Difference in average IU boarding times between scenarios: Case 2

Case 2: Average time difference (minutes)*, p-value**					
* $\mu_x - \mu_y$, x = column, y = row, μ = average wait time					
**Hypothesis: $\mu_x = \mu_y$ for IU boarding times with no NVA delay, rejecting $p < 0.05$					
	Crowding Index	Imperfect Index	Perfect Index	Time Based	Baseline
Crowding Index	=				
Imperfect Index	NA, 0.16	=			
Perfect Index	NA, 0.93	NA, 0.19	=		
Discharge by Noon	-2.04, << 0.05	NA, 0.18	-1.98, << 0.05	=	
Baseline	-4.62, << 0.05	-3.60, << 0.05	-4.56, << 0.05	-2.58, << 0.05	=
Improvement over Baseline	17.58%	13.70%	17.35%	9.82%	=

Table 15 Difference in average IU boarding times between scenarios: Case 3

Case 3: Average time difference (minutes)*, p-value**					
* $\mu_x - \mu_y$, x = column, y = row, μ = average wait time					
**Hypothesis: $\mu_x = \mu_y$ for IU boarding times with no NVA delay, rejecting $p < 0.05$					
	Crowding Index	Imperfect Index	Perfect Index	Time Based	Baseline
Crowding Index	=				
Imperfect Index	NA, 0.26	=			
Perfect Index	NA, 0.27	NA, 0.90	=		
Discharge by Noon	-7.32, << 0.05	-8.94, << 0.05	-8.28, << 0.05	=	
Baseline	-12.96, << 0.05	-14.58, << 0.05	-13.92, << 0.05	NA, 0.14	=
Improvement over Baseline	13.71%	15.43%	14.73%	NA	=

6.4 Discussion

From the results above, it is clear that the simulation has a great deal of variability. Although this makes analysis of data more difficult, this was planned to make the system more realistic. Beginning with the results in Figure 41 it can be seen that waiting time does decrease as the emphasis of discharge is moved earlier, up to the point where doctors are no longer on duty. This is in line with the discharge by noon heuristic that has become a popular method for reducing waiting times [ACEP 2008].

As would be expected from the cyclical nature of the simulation, there is some resonance in the system that will either dampen or exacerbate waiting time depending on when discharge is prioritized. This can be seen in Figure 41 where the benefits of early discharge are not quite achieved when beginning discharge priority at 7am as opposed to 8am. This may be due to the fact that early in the morning, doctors begin with discharges, causing patients to wait for treatment until later in the day. This delays the patient's entry into the treatment cycle and may be reducing how many times a patient can be seen each day, thus hurting performance. This is directly analogous to the real hospital system where a doctor will likely want to see their treatment patients first, so that they can schedule a series of tests throughout the day. If they miss this meeting, the patient may not be able to get all of the necessary tests that day and their length of stay will be extended. It is therefore worth noting that it may be unwise to be overly strict in enforcing the discharge by noon heuristic and incur this negative situation.

There are also boarding time peaks when discharge priority begins in the afternoon; these may be the result of doctors not seeing treatment patients in the afternoon which causes them not to be seen until the next day, thus hurting performance. These peaks may purely be a result of how the model was designed, however an argument can be made that a real hospital may also have schedules in place that cause resonance. Staffing schedules, lunch hours, clinic hours, patient arrival patterns, educational sessions, and other recurring factors may mean that efforts at improving flow by encouraging discharges at a specific time may negatively affect flow by interacting with the hospitals emergent schedule.

The resonant peaks may also be the cause of the spikes in the sensitivity schedules shown in Figures 42 through 44. The optimization algorithm uses an advanced searching mechanism. The

searches likely identified times where sensitivities were harmful and created schedules that avoided those times. These schedules changed when the NVA delay was added or when resources were reduced because these systematic changes affected the resonance times for the entire system. This flexibility in scheduling, when using predictive and crowding indexes allowed the system to compensate for the NVA delay and reduced IU doctor capacity. This compensation lead to consistently larger improvements in waiting/boarding times than the discharge by noon scenario, when compared to the baseline. Also, while discharge by noon does seem to have some added benefit in the case of the NVA delay, it was not effective at managing the system when resources were reduced, leaving it statistically similar to the baseline case.

It is worth noting that the schedules presented may not be the true optimal schedules, and a schedule may exist that is more logical for each index. However the schedules were found to be high performing by the software and were verified as locally optimal based on manual adjustments. The manual adjustments however only changed one hour at a time. Likely this missed the complex interactions that would occur by changing two or more hours simultaneously as was done by the software and would be required if further optimization was desired.

The optimized schedules perform very well using the IU boarding time metrics. Average values and variability for the index scenarios are significantly reduced compared to the baseline and time based scenarios. This means that, when combined, an index and a carefully chosen sensitivity schedule have the potential for greater performance than the discharge by noon heuristic that is currently the industry standard. When looking at ED waiting time, the improvements are less clear in Case 1, however they become more pronounced in Cases 2 and 3. This may mean that in Case 1 ED performance is less impacted by the inpatient unit. However when adding the delays and resource reductions the IU begins to have more impact on the ED and managing the connection between the two units becomes more important.

Finally it is worth noting that, using the optimized schedules, all three index types (crowding, perfect, and imperfect) were capable of generating superior performance. It is unclear that one index was significantly better than another. This means that using the imperfect prediction method described in Chapters 3 and 4 may be good enough in the true hospital and investment in a more perfect method could be a waste of resources. Similarly it means that a crowding metric

could be used, however if the system is not guided by a sensitivity schedule and associated definitive actions, it is unlikely to have the same intuitive pull incentive as the prediction values.

6.5 Limitations/Conclusions

The study presented in this chapter has some inherent limitations in that it is simulation based. Despite calibrations made to the system, the validation procedure shows that the simulation does not directly match the true hospital system. To that end, the exact dynamics of the simulation are different than the true hospital and therefore the schedules that were created to optimize flow, based on these dynamics, are likely not directly transferable to the true hospital. The primary take away from the results is that schedules can be made to optimize flow, using prediction and crowding indexes. A limitation also exists in terms of variability built into the system and its effects on the results. Even using 5000 repetitions of the simulation, there is still a great deal of uncertainty about the true average waiting times and boarding times in each scenario. Limitations in computing power meant that the optimization tool had to use only 1000 repetitions to create the schedules and only the final result was tested at 5000 repetitions. Had 5000 or more repetitions been used by the optimization software, the schedules would likely be even closer to the true optimal; this would have taken an unacceptable amount of computing time.

While there is a clear benefit to finding the optimized schedules and showing that such a schedule may exist in the real hospital, it is unlikely that an hourly schedule can ever be truly found in VHA West Roxbury. Instead, more practical value will be derived from applying this simulation to finding semi-optimal simplified solutions where sensitivity is held at a specific level for 2, 3, or even 4 hours, rather than varying on an hourly basis. While this is interesting, the goal of this chapter was to find the maximum potential of using prediction, so the search for practical schedules is saved for future work.

The suggestion explored in the earlier chapters of this dissertation (prediction can be used to improve flow in the ED/IU health care delivery chain) naturally leads to the development of an information based pull system in the chain. This chapter showed that, in a simulated hospital, the creation of such a system does indeed have the ability to improve flow and reduce the effects of non-value added delays and resource limitations. The prediction based scenarios showed a consistent improvement of 8-18% in ED/IU boarding time, compared to baseline scenarios. The

impacts of this boarding time improvement on ED waiting time do not emerge until Cases 2 and 3 where the IU must be having a greater impact on the ED. Achieving this improvement required schedules that dictate hourly sensitivity to ED crowding and admission prediction indexes, that do not negatively interact with the emergent schedule of the hospital.

Chapter 7: Conclusions, Contributions, and Future Work

7.1 Conclusions and contributions of studies

Recent studies recommended exploring methods for formally translating methods from supply chain management into improving systems of connected health treatment steps [Vries and Huijsman 2011]. The primary goal of this dissertation was to introduce and study health care delivery chains. In order to do this, the Emergency Department (ED)/Inpatient Unit (IU) health care delivery chain was selected as a sample. This sample was selected because it is well known and commonly studied. The volume of studies focused on this chain made it possible to draw a more clear distinction between a health care delivery chain approach and other approaches that have been taken in the past.

In Chapter 2, background was provided about the ED/IU chain. This chapter described how flow of patients is directly tied to quality of care in the ED/IU chain. Chapter 2 also identified the primary measures to be used in this dissertation: waiting time (the time between when a patient registers at the ED welcome desk and when they are placed in an ED bed) and boarding time (the time between when an ED physician decides to have a patient admitted to the IU and they are placed in an IU bed). Chapter 2 also described historic approaches to improving flow in the ED/IU chain by targeting the input, output, and throughput elements of the chain. This discussion led to a common theme in recent literature, that the reduction of boarding time should be a primary goal of ED flow research. To this end, some studies suggested that if a patient's likelihood of admission could be predicted when they entered the ED then this information could be shared with the IU in order to allow the hospital to prepare for the admission. This preparation would reduce the administrative and process delays that impact boarding time.

The suggestion that prediction can be used to improve flow shows that the literature approached the concept of treating the ED/IU system as a chain without necessarily defining it as such. The suggestion has a direct correlation to the use of predictive information to improve product flow in a supply chain. Specifically, the idea that predictions can be made on patients who have already arrived in the ED, has direct comparison to the use of advanced demand information

(ADI) to manage supply chains. This is in contrast to the use of long term forecasting predictions that are also used in both supply and health systems.

Chapter 1 introduced the primary research questions of this dissertation, targeting the concept of using prediction to improve flow in the ED/IU health care delivery chain. Through the chapters that followed, studies were described that sought to answer these questions. Below is a summary of the results and contributions of each study, structured by how these results apply to one of the three questions.

7.1.1 Question 1

What predictive methods work best to predict downstream demand in the context of a single Emergency Department/Inpatient Unit health care delivery chain?

Before studying exactly how prediction could be used to manage the ED/IU health care delivery chain, it was first necessary to understand how admission predictions can be made on patients. Question 1 is based upon this necessity.

Chapter 3 was targeted specifically at answering Question 1. In Chapter 3, a study at VHA West Roxbury was described that used three simple methods for making admission predictions on patients, using data that can be collected at ED triage (the first time a medical professional sees that patient and collects their data). The three methods were: expert opinion, naïve Bayes conditional probability and a linear regression with a logit link function (logit-linear regression). It was found that the logit-linear regression performed best at making admission predictions.

The answer to Question 1, provided in Chapter 3, represents a contribution to the field of health systems research and to quantitative data analysis fields such as artificial intelligence, machine learning, data mining, and statistics. On the most basic level, the conclusions of Chapter 3 contribute to the study of the ED/IU health care delivery chain by providing an understanding of the applicability of simple, replicable methods for making predictions in the ED. The study also contributes on a broader level. Despite the early successes in applying prediction methods to health care systems, as described in Chapter 2, there remain many new areas where predictive methods can be applied. The study described in Chapter 3 deepens the knowledgebase of health care systems that can use predictions. It also serves as an example of how such a system is

structured, so future health care researchers may identify similar systems with prediction potential.

In terms of the quantitative science of prediction, Chapter 3 contributes to the number of successful applications of prediction methods to new areas; this means that future researchers can safely consider using these methods in similar situations. Furthermore, the results of Chapter 3 deepen the knowledge of how three methods for prediction perform compared to one another, which can inform future method selection. Finally, those who wish to develop and study more complex methods, have a new domain for application and a baseline model performance for comparison.

7.1.2 Question 2

How portable or robust are these prediction methods to multiple hospital contexts?

Showing that logit-linear regression was successful in one hospital was an important first step. However, it was worthwhile exploring whether the methodology to develop the model and its accuracy was truly valid, rather than a random coincidence due to the unique nature of the original hospital. This exploration was the goal of Chapter 4, which varied the context of the prediction in two dimensions: (1) the hospital to which the model was applied and (2) whether the model is applied to retrospective, investigator coded data or prospective, nurse coded data.

The results of Chapter 4 showed that the specific combination of factors in the final model of Chapter 3 were not generalizable; however the process of using basic triage level data to create a linear regression model to predict admissions was generalizable to different hospital settings. It was also found that the prediction model developed in Chapter 3 continued to perform well when applied to prospectively collected nurse coded data. Despite not having the conditions for a perfect experiment, the results of the study in Chapter 4 did suggest some prescriptive conclusions, about data that would be useful for a hospital to collect, and some descriptive conclusions on the potential effects that different internal processes may have on the accuracy and development of that hospital's model.

While answering Question 2, the results of Chapter 4 make a significant contribution to the study of health care systems and to the study of quantitative prediction methods. There have been

many published studies that show single successes using quantitative methods in health care systems. This leads some people to believe that many of these successes are isolated incidents and cannot be generalized to other contexts. The study in Chapter 4 contributes to the study of health systems by better describing exactly what is meant by “generalization” and by showing the process of searching for generalizability in this system. The results are a contribution to the quantitative study of prediction models by exploring the robustness of the logit-linear regression methodology that was selected.

7.1.3 Question 3

Given advance demand predictions, what possible adaptive actions can a hospital system take to improve flow given (a) perfect and (b) imperfect downstream demand prediction?

The answer to Question 3 has evolved throughout the dissertation, with the final conclusions in Chapters 5 and 6. Chapter 2 used figures to explain the general idea of how prediction could be used in the ED/IU chain in order to improve flow; however the chapter did not discuss the practical implications of implementing a prediction system. Chapter 3 began the discussion of using predictions on an individual basis or by aggregating predictions to create a running bed demand value that IU staff could monitor. Chapter 4 further discussed this differentiation and introduced measures that could be used to explore the quality of an individual based prediction, which can be used coercively to order beds, or an aggregate prediction, that could be used to inform inpatient staff decisions based on awareness of the entire system. While Chapter 3 and 4 suggested that the aggregate prediction tends to have more accuracy, it was unclear how it could actually be used in a live implementation. Gaining insight into this was the goal of Chapter 5.

Chapter 5 described a live implementation of the prediction system where aggregate ED to IU bed demand was shared in real time with the staff of VHA West Roxbury. While there was no significant improvement in the key flow measures, the qualitative results showed a variety of people who found the prediction system useful for five different application categories. However, it was also found that the variety of potential applications required a variety in data details, some applications required patient specific predictions, rather than the aggregate prediction that was shared. It was also found that the complex interactions between the multiple

stakeholders in the patient care process could result in reduced effectiveness in responding to predictions when there is a lack of alignment.

In order to gain a better quantitative understanding of how prediction could be used in the ED/IU chain, Chapter 6 described the development and study of a discrete event simulation (DES) model of the system. The value of the DES model is that it was able to simplify the system by removing the complex interactions between multiple stakeholders and some resource restrictions. This makes it possible to get a “perfect world” understanding of how much improvement could be expected using prediction, simply given the overarching dynamics of the hospital and patient arrival patterns. The simulation results were promising, showing an 8-18% reduction in ED/IU boarding time when prediction measures are shared.

While answering Question 3, the results of Chapters 3-6 provide contributions to many areas of research. As with the results from Questions 1 and 2, the results of answering Question 3 make a significant contribution to the area of health systems research. In particular the results show new methods for managing the ED/IU health care delivery chain, in a way that may be implemented. These methods include the concept of aggregating predictions to make decisions in a health care system, which was not found in any other study and has the potential to change how health systems researchers choose to use predictive methods in health care.

Another contribution is made towards the field of implementation science. The methods for studying the implementation of a research based system, in order to achieve effective and reliable results, are continuing to develop. The study in Chapter 5 serves as a case study of an implementation and for learning from interviews. The case study also exposes the complications that arise from the short time period and necessary limitations in experimental design. These results can inform future implementation studies, whether they are focused on a health care system or not. Chapter 5 also contributed to the understanding of the organizational behavior in the hospital, by exposing the potential for misalignment between stakeholder decisions, that could prevent the optimization of patient flow. This is a key concept that must be understood clearly before any truly effective solutions can be developed.

Chapter 6 makes significant contributions towards the field of discrete event simulation (DES). The first contribution is the application of the DES method to a complex system. While DES has

been applied to the ED in other studies, by once again applying the tool to this system and by including a discussion of the validation of the model, Chapter 6 further solidifies the value of DES for studying this and similar systems. The model described in Chapter 6 also contributes to the field of DES by describing a method for modeling the complex flows and provider decisions in a hospital as influenced by a centralized control system. The cyclical model logic described in Figure 36 is a method for representing a system where one resource serves two queues and where non-value added time is affected by how the resource is managed. This model logic may be applicable to many other models of health care and non-health care systems. Finally, the use of optimization packages in DES models continues to grow. By applying the optimization software for finding optimized crowding/prediction sensitivity schedules within the simulation, the study showed a new application of optimization.

7.2 Future Work on the ED/IU Health Care Delivery Chain

7.2.1 Question 1

While the results of the study in Chapter 2 approach a solution to Question 1, there are some limitations that arise from the study design. It was found that the logit-linear regression approach performed the best of the three approaches that were chosen. This does not mean that logit-linear regression is the absolute “best” method for making predictions in the ED. There are many other methods that exist or are being developed in the fields of artificial intelligence, machine learning, data mining, and statistics. The methods that were chosen in the study were selected due to their ease of use, which increases the likelihood that they can be replicated in practice. This was listed as one of the main contributions of the study. Nevertheless, future research should be devoted towards finding the absolute best method for making admission predictions. These other methods would have to be accompanied by tools that enable hospitals to use the models without requiring an advanced knowledge of quantitative methods, without this, the models may be ignored in practice.

7.2.2 Question 2

The value of the logit-linear regression model was further exposed by exploring its generalizability while answering Question 2. Fully understanding the generalizability of the model would require a greater sample of hospitals, including multiple hospitals that have exactly

the same characteristics. It would also be useful to create a more controlled experiment where data are collected at each hospital in the exact same format.

Despite the short comings in perfect experimental design, that are expected when working with large complex human systems, the investigators were comfortable enough with the results of the study in Chapter 4 to explore the possibilities of actually implementing the system. Although the results were enough to satisfy the investigators, there remains future work that could focus on creating a more complete experiment of generalizability. Such an experiment should continue to focus on the multiple levels of generalizability and use a more complete sample of hospitals, which would result in more conclusive results. The goals of such research would be to suggest the best data and data format to collect at triage in order to make predictions of admission. The study would also seek to identify properties that may make two hospitals have similar predictive models and other hospitals have different models.

7.2.3 Question 3

While the results of Chapter 5 and 6 begin to answer Question 3, there is still a great amount of work to be done refining these approaches. Future work can be performed to refine the live implementation in Chapter 5. It is worthwhile finding a method to collect and share the predictions in a more passive way. This would reduce the amount of extra effort that must be done during an implementation and enable longer term experimentation. With more automation it may be possible to run experiments across locations where the type of shared information is varied. This experimental design would seek to find the most effective data sharing format and system for reducing the key flow metrics. Also, while implementing over a long period of time, it is possible to identify, in real time, the organizational issues that reduce the usefulness of the predictions and address them while the implementation is still running. This may lead to a more complete understanding of these issues, rather than studying them after the fact and trying again. This approach may lead to dynamic solutions that may be applicable elsewhere. If improvements are performed one at a time, it may be possible to watch data trends to identify which ones had the largest impact.

The DES model introduced in Chapter 6 can also continue to be refined to show more realistic methods for controlling the hospital using prediction data, which will be more easily translated

into a live hospital. For example it may be possible to create simpler crowding/prediction sensitivity schedules made with multi-hour time blocks rather than allowing for an hourly change in sensitivity level. It may also be worth further developing the model to better reflect the limitations that arise from the complexity of the real system. A more complicated model would allow for the exploration of solutions to these limitations as they are discovered during the long term, live implementation.

7.3 Health Care Delivery Chains

7.3.1 Contribution to/Invention of Health Care Delivery Chain Management

It is possible to structure the series of studies presented in this dissertation as approaches specifically to improving Emergency Department patient flow, and ignore the conceptual development of Health Care Delivery Chains. Indeed by answering the questions, as described above, this dissertation has made a valuable impact on this issue. However, Chapter 1 established that the overarching goal of this collection of studies was to explore the concept of Health Care Delivery Chain Management.

The purpose of focusing on a specific example of a health care delivery chain is to better define the concept. Seeing “a thing” helps one know “a thing” and makes it possible to recognize another in the future. It is the hope of the author that, by describing the example ED/IU health care delivery chain (Figure 46), it will be possible for other researchers or practitioners to identify when an issue in a health system originates from the interaction between a series of health care delivery steps, linked together by a flow of patients. This may be represented by the generic health care delivery chain (Figure 47).

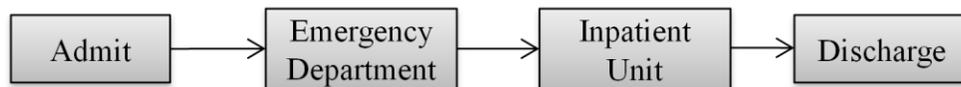


Figure 46 ED/IU health care delivery chain

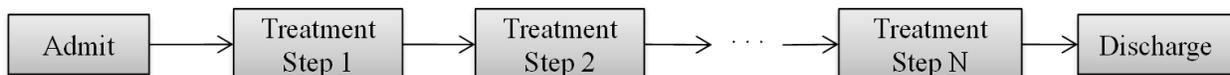


Figure 47 Generic health care delivery chain

Once this recognition is made it is then possible to consider the types of improvement methods that have been applied to other examples of these chains, rather than just tools that have been applied to isolated delivery units. While the literature on methods specific to health care delivery chains remains small, or not recognized as applying to the domain, it is the hope of the author that the studies described here provide a repeatable example of a method that can be applied to other chains. It is also hoped that the studies here provide an example of seeking comparable supply chain methodologies and converting them to health care delivery chains (note converting as opposed to simply applying).

By exploring the ED/IU delivery chain, this dissertation introduced one method for improving flow in all health care delivery chains, using predictions on incoming patients to better prepare downstream resources. A similar situation can be imagined for chains of the three different abstraction levels presented in Chapter 1 Section 4:

In Department: Predictions can be made on patients entering a specific unit that will define resource demands within that unit later. One example of this was described in Chapter 2 where the emergency severity level (ESI) of a patient is a prediction based on acuity and likely resource requirement. To assign this level, triage nurses actually predict the specific resources that the patient will need. Therefore triage is a step in a chain upstream from ED treatment and predictions associated with ESI level can be used to better manage testing and treatment resources in the ED Treatment step. In fact, current practice often allows a triage nurse to preemptively order tests so that all set up is complete, or results are available when the patient arrives in their bed. This practice has impact on total patient flow.

Cross department: Just as the flow of patients from the ED to the IU is important, so too is the flow from the IU to a long term stay facility. Such a facility may benefit from receiving predicted information from the hospital in order to better prepare for future patient arrivals or enable the scheduling of more efficient patient pick-ups.

Cross Organizational: In an integrated health system, specialty care physicians may be able to expect incoming future demand, based on the properties of patients being seen by the primary care physicians in that system, months ahead of time. This interaction re-emphasizes the concept

explained in Figure 8 where the knowledge of patients currently in one chain can feed back into planning for other chains over an entire lifecycle.

7.3.2 Next Steps

It is worth noting that methods currently used in Supply Chain Management cannot easily be translated to health care delivery chain management. In fact it may not always be easy to translate a tool from one health care delivery chain to another. This is due to the complex interactions between the different views of a health care delivery enterprise as described in Chapter 1, Section 3. A health care delivery chain is more than just the flow of patients; this flow is directly linked to the organization, processes, knowledge exchange etc. within the enterprise. While in a more controllable supply chain where the product is not human and the resource/capacity limitations are machine based, some assumptions can be made, resulting in models that remain true when implemented in reality. This is not the case in a health care delivery chain. As was seen in this dissertation, the results in Chapter 6 were not realized during the implementation in Chapter 5. It was for this reason that the studies presented here and any study into health care delivery chains must include some dedication to understanding the true system and the complex interactions between enterprise views, or else the results will have less meaning and applicability. The methods for understanding the human piece of a system are central to health care delivery chain management, while they may be considered a peripheral piece of supply chain management.

With the above considerations in mind, next steps in exploring the field of Health Care Delivery Chain Management should focus on developing many example cases of health care delivery chains. These examples should include the enterprise context of the chain and how improvement methods applied to the chain took this context into account. Through studying these many cases it may be possible to create a health care delivery chain tool kit containing clear methods that can be applied in specific contexts.

Chapter 8: References

- Abraham, G., Byrnes, G.B. & Bain, C.A., 2009. Short-Term Forecasting of Emergency Inpatient Flow. *IEEE Transactions on Information Technology in Biomedicine*, 13(3), pp.380–388.
- ACEP et al., 2008. Emergency Department Crowding: High Impact Solutions. no. American College of Emergency Physicians.
- Adams, J.B. & Wert, Y., 2005. Logistic and neural network models for predicting a hospital admission. *Journal of Applied Statistics*, 32(8), pp.861–869.
- AHRQ: Agency for Health care Research and Quality, 2012. Enabling Health Care Decisionmaking Through Clinical Decision Support and Knowledge Management - NCBI Bookshelf. Available at: <http://www.ncbi.nlm.nih.gov/books/NBK97318/> [Accessed September 12, 2012].
- Allaudeen, N. et al., 2011. Inability of Providers to Predict Unplanned Readmissions. *Journal of General Internal Medicine*, 26(7), pp.771–776.
- Anderson, G.F. & Steinberg, E.P., 1985. Predicting hospital readmissions in the Medicare population. *Inquiry: a journal of medical care organization, provision and financing*, 22(3), pp.251–258.
- Antman, E.M. et al., 2004. ACC/AHA guidelines for the management of patients with ST-elevation myocardial infarction: a report of the American College of Cardiology/American Heart Association Task Force on Practice Guidelines (Committee to Revise the 1999 Guidelines for the Management of Patients with Acute Myocardial Infarction). *Circulation*, 110(9), pp.e82–292.
- Aronsky, Dominik et al., 2001. A Comprehensive Set of Coded Chief Complaints for the Emergency Department. *Academic Emergency Medicine*, 8(10), pp.980–989.
- Arslanian-Engoren, C., 2004. Do emergency nurses' triage decisions predict differences in admission or discharge diagnoses for acute coronary syndromes? *The Journal of Cardiovascular Nursing*, 19(4), pp.280–286.
- Asaro, P.V., Lewis, L.M. & Boxerman, S.B., 2007. The Impact of Input and Output Factors on Emergency Department Throughput. *Academic Emergency Medicine*, 14(3), pp.235–242.
- Asplin, B.R. et al., 2003. A conceptual model of emergency department crowding. *Annals of Emergency Medicine*, 42(2), pp.173–180.
- Asplin, B.R., 2003. Does ambulance diversion matter? *Annals of Emergency Medicine*, 41(4), pp.477–480.
- Asplin, B.R., 2006. Hospital-Based Emergency Care: A Future Without Boarding? *Annals of Emergency Medicine*, 48(2), pp.121–125.

- Au, L. et al., 2009. Predicting overflow in an emergency department. *IMA Journal of Management Mathematics*, 20(1), pp.39–49.
- Baesler, F.F., Jahnsen, H.E. & DaCosta, M., 2003. Emergency departments I: the use of simulation and design of experiments for estimating maximum capacity in an emergency room. In *Proceedings of the 35th conference on Winter simulation: driving innovation. WSC '03. Winter Simulation Conference*, pp. 1903–1906. Available at: <http://dl.acm.org/citation.cfm?id=1030818.1031081> [Accessed January 11, 2012].
- Baker, D.W., Stevens, C.D. & Brook, R.H., 1991. Patients Who Leave a Public Hospital Emergency Department Without Being Seen by a Physician: Causes and Consequences. *JAMA*, 266(8), pp.1085–1090.
- Balci, O., 1995. Principles and techniques of simulation validation, verification, and testing. In *Simulation Conference Proceedings, 1995. Winter. Simulation Conference Proceedings, 1995. Winter*. pp. 147 –154.
- Banks, J. et al., 2010. *Discrete-event system simulation*, Upper Saddle River, N.J.; Singapore: Prentice Hall.
- Berber, E., 2009. *Clinical Decision Support Systems: State of the Art*.
- Bernstein, S.L. et al., 2003. Development and validation of a new index to measure emergency department crowding. *Academic Emergency Medicine: Official Journal of the Society for Academic Emergency Medicine*, 10(9), pp.938–942.
- Bernstein, S.L. et al., 2009. The Effect of Emergency Department Crowding on Clinically Oriented Outcomes. *Academic Emergency Medicine*, 16(1), pp.1–10.
- Bernstein, S.L. & Asplin, B.R., 2006. Emergency Department Crowding: Old Problem, New Solutions. *Emergency Medicine Clinics of North America*, 24(4), pp.821–837.
- Bertoty, D., Kuszajewski, M. & Marsh, E., 2007. Direct-to-room: one department's approach to improving ED throughput. *Journal of Emergency Nursing*, 33(1), pp.26–30.
- Bindman, A.B. et al., 1991. Consequences of Queuing for Care at a Public Hospital Emergency Department. *JAMA*, 266(8), pp.1091–1096.
- Black, J.R. & Miller, D., 2008. *The Toyota way to health care excellence increase efficiency and improve quality with Lean*, Chicago: Health Administration Press. Available at: <http://www.books24x7.com/marc.asp?bookid=34369> [Accessed September 13, 2012].
- Blumenthal, D., 2009. Stimulating the Adoption of Health Information Technology. *N Engl J Med*, 360(15), pp.1477–1479.
- Boger, E., 2003. Electronic tracking board reduces ED patient length of stay at Indiana Hospital. *Journal of Emergency Nursing*, 29(1), pp.39–43.
- Bonder, S., 2005. Changing Health Care Delivery Enterprises. In *Building a Better Delivery System*. National Academies Press, pp. 149–152.

- Bose, K.S. & Sarma, R.H., 1975. Delineation of the intimate details of the backbone conformation of pyridine nucleotide coenzymes in aqueous solution. *Biochemical and biophysical research communications*, 66(4), pp.1173–1179.
- Bose, R., 2003. Knowledge management-enabled health care management systems: capabilities, infrastructure, and decision-support. *Expert Systems with Applications*, 24(1), pp.59–71.
- Boult, C. et al., 1993. Screening elders for risk of hospital admission. *Journal of the American Geriatrics Society*, 41(8), pp.811–817.
- Braunwald, E. et al., 2002. ACC/AHA Guideline Update for the Management of Patients With Unstable Angina and Non-ST-Segment Elevation Myocardial Infarction--2002: Summary Article: A Report of the American College of Cardiology/American Heart Association Task Force on Practice Guidelines (Committee on the Management of Patients With Unstable Angina). *Circulation*, 106(14), pp.1893–1900.
- Brindley, P.J. et al., 1991. Characterization of a programmed alteration in an 18S ribosomal gene that accompanies the experimental induction of drug resistance in *Schistosoma mansoni*. *Proceedings of the National Academy of Sciences of the United States of America*, 88(17), pp.7754–7758.
- Brown, D. et al., 2012. Continued rise in the use of mid-level providers in US emergency departments, 1993–2009. *International Journal of Emergency Medicine*, 5(1), pp.1–5.
- de Bruin, A. et al., 2007. Modeling the emergency cardiac in-patient flow: an application of queuing theory. *Health Care Management Science*, 10(2), pp.125–137.
- Cachon, G. & Terwiesch, C., 2009. *Matching supply with demand: an introduction to operations management*, McGraw-Hill.
- Carey, R.G., 2003. *Improving health care with control charts : basic and advanced SPC methods and case studies*, Milwaukee, WI: ASQ Quality Press.
- Carter, B.L. et al., 2009. The Potency of Team-based Care Interventions for Hypertension. *Archives of internal medicine*, 169(19), pp.1748–1755.
- CDC: Centers for disease control and prevention, 2011. NAMCS/NHAMCS - Questionnaires, Datasets, and Related Documentation. Available at: http://www.cdc.gov/nchs/ahcd/ahcd_questionnaires.htm [Accessed September 30, 2012].
- Chase, V.J. et al., 2012. Predicting Emergency Department Volume Using Forecasting Methods to Create a “Surge Response” for Noncrisis Events. *Academic Emergency Medicine*, 19(5), pp.569–576.
- Chaudhry, B. et al., 2006. Systematic Review: Impact of Health Information Technology on Quality, Efficiency, and Costs of Medical Care. *Annals of Internal Medicine*, 144(10), pp.742–752.
- Chen, F., 2001. Market Segmentation, Advanced Demand Information, and Supply Chain Performance. *Manufacturing & Service Operations Management*, 3(1), pp.53–67.

- Chockalingam, A., Jayakumar, K. & Lawley, M.A., 2010. A stochastic control approach to avoiding Emergency Department overcrowding. In Simulation Conference (WSC), Proceedings of the 2010 Winter. Simulation Conference (WSC), Proceedings of the 2010 Winter. pp. 2399–2411.
- Clark, K. & Normile, L.B., 2007. Patient Flow in the Emergency Department: Is Timeliness to Events Related to Length of Hospital Stay? *Journal of Nursing Care Quality*, 22(1). Available at: http://journals.lww.com/jncqjournal/Fulltext/2007/01000/Patient_Flow_in_the_Emergency_Department__Is.16.aspx.
- Clerkin, D., Fos, P.J. & Petry, F.E., 1995. A decision support system for hospital bed assignment. *Hospital & health services administration*, 40(3), pp.386–400.
- Clesham, K. et al., 2008. Can emergency medical service staff predict the disposition of patients they are transporting? *Emergency Medicine Journal: EMJ*, 25(10), pp.691–694.
- CMS: Centers for Medicare and Medicaid Services, 2012. Accountable Care Organizations. Available at: <http://www.cms.gov/Medicare/Medicare-Fee-for-Service-Payment/ACO/index.html?redirect=/ACO/> [Accessed September 4, 2012].
- CMSb Centers for Medicare and Medicaid Services, 2012. EHR Incentive Programs. Available at: <http://www.cms.gov/Regulations-and-Guidance/Legislation/EHRIncentivePrograms/index.html?redirect=/ehrincentiveprograms/> [Accessed September 4, 2012].
- Connelly, L.G. & Bair, A.E., 2004. Discrete Event Simulation of Emergency Department Activity: A Platform for System-level Operations Research. *Academic Emergency Medicine*, 11(11), pp.1177–1185.
- Cowan, R.M. & Trzeciak, Stephen, 2005. Clinical review: Emergency department overcrowding and the potential impact on the critically ill. *Critical Care*, 9(3), pp.291–295.
- Cuellar, A.E. & Gertler, P.J., 2003. Trends In Hospital Consolidation: The Formation Of Local Systems. *Health Affairs*, 22(6), pp.77–87.
- Cummings, M.L., 2004. Automation Bias in Intelligent Time Critical Decision Support Systems. In *AIAA 3rd Intelligent Systems Conference*. AIAA, pp. 2004–6313.
- Cunningham, Peter, 2011. Nonurgent Use of Hospital Emergency Departments.
- DesRoches, C.M. et al., 2008. Electronic Health Records in Ambulatory Care — A National Survey of Physicians. *N Engl J Med*, 359(1), pp.50–60.
- Dickson, E. et al., 2009. Use of Lean in the Emergency Department: A Case Series of 4 Hospitals. *Annals of emergency medicine*, 54(4), pp.504–510.
- Diercks, D.B. et al., 2007. Prolonged Emergency Department Stays of Non-ST-Segment-Elevation Myocardial Infarction Patients Are Associated With Worse Adherence to the

- American College of Cardiology/American Heart Association Guidelines for Management and Increased Adverse Events. *Annals of Emergency Medicine*, 50(5), pp.489–496.
- Ding, R. et al., 2009. Predicting Emergency Department Length of Stay Using Quantile Regression. In *International Conference on Management and Service Science, 2009. MASS '09. International Conference on Management and Service Science, 2009. MASS '09*. pp. 1–4.
- Epstein, S.K. & Tian, L., 2006. Development of an emergency department work score to predict ambulance diversion. *Academic Emergency Medicine: Official Journal of the Society for Academic Emergency Medicine*, 13(4), pp.421–426.
- ESD: The Engineering Systems Division Architecture Committee, 2004. *The Influence of Architecture in Engineering Systems*.
- van Essen, J. et al., 2012. Decision support system for the operating room rescheduling problem. *Health Care Management Science*, pp.1–18.
- Falvo, T. et al., 2007. The Opportunity Loss of Boarding Admitted Patients in the Emergency Department. *Academic Emergency Medicine*, 14(4), pp.332–337.
- Fee, C. et al., 2007. Effect of Emergency Department Crowding on Time to Antibiotics in Patients Admitted With Community-Acquired Pneumonia. *Annals of Emergency Medicine*, 50(5), pp.501–509.e1.
- Feldstein, P.J., 2003. *Health policy issues : an economic perspective*, Chicago; Washinbgton, DC: Health Administration Press ; AUPHA Press.
- Fields, W., 1999. *Defining America's Safety Net*. Report of the American College of Emergency Physicians.
- Fisher, E. et al., 2009. *Health Care Spending, Quality, and Outcomes: More Isn't Always Better*. Available at: http://www.dartmouthatlas.org/downloads/reports/Spending_Brief_022709.pdf.
- Forgionne, G.A. & Kohli, R., 1996. HMSS: a management support system for concurrent hospital decision making. *Decision Support Systems*, 16(3), pp.209–229.
- Fradinho, J., 2011. *Towards high performing hospital enterprise architectures : elevating hospitals to lean enterprise thinking*. Cambridge, MA: Massachusetts Institute of Technology. Engineering Systems Division. Available at: <http://hdl.handle.net/1721.1/67757>.
- Friedmann, J.M. et al., 1997. Predicting early nonelective hospital readmission in nutritionally compromised older adults. *The American Journal of Clinical Nutrition*, 65(6), pp.1714–1720.
- Fuchs, V.R. & Emanuel, E.J., 2005. Health Care Reform: Why? What? When? *Health Affairs*, 24(6), pp.1399–1414.
- Fudenberg, D. & Tirole, J., 1991. *Game theory*, Cambridge, Mass.: MIT Press.
- Ganapathy, S. & Zwemer Jr, F., 2003. Coping with a crowded ED: an expanded unique role for midlevel providers. *The American journal of emergency medicine*, 21(2), pp.125–128.

- Gans, D. et al., 2005. Medical Groups' Adoption Of Electronic Health Records And Information Systems. *Health Affairs*, 24(5), pp.1323–1333.
- Gawande, A., 2002. *Complications: a surgeon's notes on an imperfect science*, New York: Metropolitan Books.
- Gawande, A., 2011. *The Hot Spotters*. *The New Yorker*. Available at: http://www.newyorker.com/reporting/2011/01/24/110124fa_fact_gawande [Accessed October 3, 2012].
- Gayon, J.-P., Benjaafar, S. & de Véricourt, F., 2009. Using Imperfect Advance Demand Information in Production-Inventory Systems with Multiple Customer Classes. *Manufacturing & Service Operations Management*, 11(1), pp.128–143.
- Gilboy, N. et al., 2005. *Emergency severity index, version 4: implementation handbook*. Rockville (MD): AHRQ Publication.
- Goldratt, E. & Cox, J., 1989. *The goal*, Gower.
- Gordon, B. & Asplin, BR, 2004. Using online analytical processing to manage emergency department operations. *Academic Emergency Medicine*, 11(11), pp.1206–1212.
- Gorelick, M.H., Yen, K. & Yun, H.J., 2005. The effect of in-room registration on emergency department length of stay. *Annals of Emergency Medicine*, 45(2), pp.128–133.
- Graban, M., 2008. *Lean hospitals: Improving quality, patient safety, and employee satisfaction*, Productivity Press.
- Graff, L. et al., 2002. Measuring and Improving Quality in Emergency Medicine. *Academic Emergency Medicine*, 9(11), pp.1091–1107.
- Green, LV et al., 2006. Using queueing theory to increase the effectiveness of emergency department provider staffing. *Academic Emergency Medicine*, 13(1), pp.61–68.
- Green, LV & Nguyen, V., 2001. Strategies for cutting hospital beds: the impact on patient service. *Health Services Research*, 36(2), p.421.
- Grove, W. et al., 2000. Clinical Versus Mechanical Prediction: A Meta-Analysis. *Psychological Assessment*, 12(1), pp.19–30.
- Gustafson, D.H. et al., 1971. Initial Evaluation of a Subjective Bayesian Diagnostic System. *Health Services Research*, 6(3), pp.204–213.
- Gustafson, D.H., 1968. Length of Stay. *Health Services Research*, 3(1), pp.12–34.
- Hall, R.W., 2006. *Patient flow: reducing delay in health care delivery*, Springer.
- Hauswald, M., 2005. Triage: better tools but the wrong problem. *Academic emergency medicine: official journal of the Society for Academic Emergency Medicine*, 12(6), p.533.

- Health System Measurement Project, Health Information Technology. Socrata. Available at: <https://healthmeasures.aspe.hhs.gov/topic/D> [Accessed September 5, 2012a].
- Health System Measurement Project, Percentage of People Who Did Not Receive or Delayed Needed Care Due to Cost in the Past 12 Months. Socrata. Available at: <https://healthmeasures.aspe.hhs.gov/measure/7> [Accessed September 4, 2012b].
- Health System Measurement Project, Personal Health Care and Federal Program Expenditures. Socrata. Available at: <https://healthmeasures.aspe.hhs.gov/measure/9> [Accessed September 4, 2012c].
- Hing, E. & Bhuiya, F., 2012. Wait Time for Treatment in Hospital Emergency Departments: 2009.
- Hollander, J.E. & Pines, J.M., 2007. The Emergency Department Crowding Paradox: The Longer You Stay, the Less Care You Get. *Annals of emergency medicine*, 50(5), pp.497–499.
- Hoot, N.R. et al., 2008. Forecasting emergency department crowding: a discrete event simulation. *Annals of Emergency Medicine*, 52(2), pp.116–125.
- Hoot, N.R., Leblanc, L.J., et al., 2009. Forecasting emergency department crowding: a prospective, real-time evaluation. *Journal of the American Medical Informatics Association: JAMIA*, 16(3), pp.338–345.
- Hoot, N.R., Epstein, S.K., et al., 2009. Forecasting emergency department crowding: an external, multicenter evaluation. *Annals of Emergency Medicine*, 54(4), pp.514–522.e19.
- Hoot, N.R. & Aronsky, Dominik, 2008. Systematic Review of Emergency Department Crowding: Causes, Effects, and Solutions. *Annals of Emergency Medicine*, 52(2), pp.126–136.e1.
- Hopp, W. & Spearman, M., 2001. *Factory physics: foundations of manufacturing management*, McGraw-Hill/Irwin.
- Horwitz, L.I., Green, J. & Bradley, E.H., 2010. US Emergency Department Performance on Wait Time and Length of Visit. *Annals of Emergency Medicine*, 55(2), pp.133–141.
- Hosmer, D.W. & Lemeshow, S., 2000. *Applied Logistic Regression*, John Wiley & Sons.
- Howell, E et al., 2008. Active bed management by hospitalists and emergency department throughput. *Annals of internal medicine*, 149(11), p.804.
- Howell, EE, Bessman, ES & Rubin, H., 2004. Hospitalists and an innovative emergency department admission process. *Journal of general internal medicine*, 19(3), pp.266–268.
- Hunt DL, H.R., 1998. Effects of computer-based clinical decision support systems on physician performance and patient outcomes: A systematic review. *JAMA: The Journal of the American Medical Association*, 280(15), pp.1339–1346.

- Hwang, U. et al., 2008. Emergency Department Crowding and Decreased Quality of Pain Care. *Academic emergency medicine: official journal of the Society for Academic Emergency Medicine*, 15(12), pp.1248–1255.
- IMPROVE Investigators et al., 2011. Predictive and Associative Models to Identify Hospitalized Medical Patients at Risk for Venous Thromboembolism. IMPROVE Publications. Available at: <http://www.outcomes-umassmed.org/IMPROVE/>.
- IOM - Institute of Medicine, Q., 2001. *Crossing the Quality Chasm*.
- Jacobson, S.H., Hall, S.N. & Swisher, J.R., 2006. Discrete-Event Simulation of Health Care Systems. In R. W. Hall & F. S. Hillier, eds. *Patient Flow: Reducing Delay in Health care Delivery*. International Series in Operations Research & Management Science. Springer US, pp. 211–252. Available at: <http://www.springerlink.com/content/q55m67703wk0qu3m/abstract/> [Accessed October 24, 2012].
- Jessup, M. et al., 2010. Implementing an emergency department patient admission predictive tool: Insights from practice. *Journal of Health Organization and Management*, 24(3), pp.306–318.
- Jha, A.K. et al., 2009. Use of Electronic Health Records in U.S. Hospitals. *N Engl J Med*, 360(16), pp.1628–1638.
- Jones, S.A., Joy, M.P. & Pearson, J., 2002. Forecasting Demand of Emergency Care. *Health Care Management Science*, 5(4), pp.297–305–305.
- Jones, S.S. et al., 2009. A multivariate time series approach to modeling and forecasting demand in the emergency department. *Journal of Biomedical Informatics*, 42(1), pp.123–139.
- Jones, S.S. et al., 2006. An Independent Evaluation of Four Quantitative Emergency Department Crowding Scales. *Academic Emergency Medicine*, 13, pp.1204–1211.
- Jones, S.S. et al., 2008. Forecasting Daily Patient Volumes in the Emergency Department. *Academic Emergency Medicine*, 15(2), pp.159–170.
- Kelen, G., Scheulen, J. & Hill, P., 2001. Effect of an emergency department (ED) managed acute care unit on ED overcrowding and emergency medical services diversion. *Academic Emergency Medicine*, 8(11), pp.1095–1100.
- Kenney, C., 2010. *Transforming health care: Virginia Mason Medical Center's pursuit of the perfect patient experience*, Portland, Or.; London: Productivity; Taylor & Francis [distributor].
- KFF: Kaiser Family Foundation, 2012. *Health Care Costs: A Primer* - Kaiser Family Foundation. Available at: <http://www.kff.org/insurance/7670.cfm> [Accessed September 3, 2012].
- King, D., Ben-Tovim, D. & Bassham, J., 2006. Redesigning emergency department patient flows: Application of Lean Thinking to health care. *Emergency Medicine Australasia*, 18(4), p.391.

- Kirtland, A. et al., 1995. Simulating an emergency department “is as much fun as...” In Simulation Conference Proceedings, 1995. Winter. pp. 1039–1042.
- Kolb, E.M.W. et al., 2008. Reducing emergency department overcrowding: five patient buffer concepts in comparison. In Proceedings of the 40th Conference on Winter Simulation. Miami, Florida: Winter Simulation Conference, pp. 1516–1525. Available at: <http://portal.acm.org/citation.cfm?id=1516744.1517010> [Accessed April 6, 2010].
- Kolker, A. & Story, P., 2012. Management engineering for effective health care delivery principles and applications, Hershey, PA: Medical Information Science Reference. Available at: <http://search.ebscohost.com/login.aspx?direct=true&scope=site&db=nlebk&db=nlabk&AN=389718> [Accessed September 12, 2012].
- Kurtulus, M., Ulku, S. & Toktay, B., 2011. The Value of Collaborative Forecasting in Supply Chains. *Manufacturing & Service Operations Management*, 14(1), pp.82–98.
- Lee, TH & Mongan, J., 2009. *Chaos and organization in health care*, Mit Pr.
- Leegon, J. et al., 2005. Predicting hospital admission for Emergency Department patients using a Bayesian network. *AMIA ... Annual Symposium Proceedings / AMIA Symposium. AMIA Symposium*, p.1022.
- Lee-Lewandrowski, E. et al., 2003. Implementation of a point-of-care satellite laboratory in the emergency department of an academic medical center. *Journal Information*, 127(4).
- Levin, S. et al., 2008. Optimizing cardiology capacity to reduce emergency department boarding: A systems engineering approach. *American Heart Journal*, 156(6), pp.1202–1209.
- Levine, S.D. et al., 2006. How well do paramedics predict admission to the hospital? A prospective study. *The Journal of Emergency Medicine*, 31(1), pp.1–5.
- Li, Jiexun, Guo, L. & Handly, N., 2009. Hospital Admission Prediction Using Pre-hospital Variables. In *IEEE International Conference on Bioinformatics and Biomedicine, 2009. BIBM '09. IEEE International Conference on Bioinformatics and Biomedicine, 2009. BIBM '09*. pp. 283 –286.
- Li, Jingshan & Howard, P.K., 2010. Modeling and analysis of hospital emergency department: An analytical framework and problem formulation. In *2010 IEEE Conference on Automation Science and Engineering (CASE). 2010 IEEE Conference on Automation Science and Engineering (CASE). IEEE*, pp. 897–902.
- Liu, K., Coughlin, T. & McBride, T., 1991. Predicting Nursing-Home Admission and Length of Stay: A Duration Analysis. *Medical Care*, 29(2), pp.125–141.
- Long, W., 1989. Medical Diagnosis Usina a Probabilistic Causal Network. *Applied Artificial Intelligence: An International Journal*, 3(2), p.367.
- Longest, B., 1974. Relationships Between Coordination, Efficiency, and Quality of Care in General Hospitals. *Hospital Administration*, 19(1), pp.65–86.

- Luchansky, B. et al., 2000. Predicting readmission to substance abuse treatment using state information systems: the impact of client and treatment characteristics. *Journal of Substance Abuse*, 12(3), pp.255–270.
- Lyons, J. et al., 1997. Predicting readmission to the psychiatric hospital in a managed care environment: implications for quality indicators. *The American journal of psychiatry*, 154(3), pp.337–340.
- Marniemi, J. & Parkki, M.G., 1975. Radiochemical assay of glutathione S-epoxide transferase and its enhancement by phenobarbital in rat liver in vivo. *Biochemical pharmacology*, 24(17), pp.1569–1572.
- Marshall, A.H. et al., 2001. Developing a Bayesian belief network for the management of geriatric hospital care. *Health Care Management Science*, 4(1), pp.25–30.
- McCarthy, Melissa L., Zeger, Scott L., et al., 2008. The Challenge of Predicting Demand for Emergency Department Services. *Academic Emergency Medicine*, 15(4), pp.337–346.
- McCarthy, Melissa L., Aronsky, Dominik, et al., 2008. The Emergency Department Occupancy Rate: A Simple Measure of Emergency Department Crowding. *Annals of Emergency Medicine*, 51(1), pp.15–24.e2.
- McConnell, K. et al., 2005. Effect of increased ICU capacity on emergency department length of stay and ambulance diversion. *Annals of emergency medicine*, 45(5), pp.471–478.
- McGuire, F., 1994. Using simulation to reduce length of stay in emergency departments. *Journal of the Society for Health Systems*, 5(3), pp.81–90.
- McKone-Sweet, K.E., Hamilton, P. & Willis, S.B., 2005. The Ailing Health care Supply Chain: A Prescription for Change. *Journal of Supply Chain Management*, 41(1), pp.4–17.
- Mclaughlin, D.B. & Hays, J.M., 2008. *Health care Operations Management*, Health Administration Press.
- Meehl, P.E., 1954. *Clinical vs. statistical prediction: a theoretical analysis and a review of the evidence*, Minneapolis, MN: University of Minnesota Press.
- Meislin, H.W. et al., 1988. Fast track: Urgent care within a teaching hospital emergency department: Can it work? *Annals of Emergency Medicine*, 17(5), pp.453–456.
- van Merode, G.G., Groothuis, S. & Hasman, A., 2004. Enterprise resource planning for hospitals. *International Journal of Medical Informatics*, 73(6), pp.493–501.
- Michalowski, W. et al., 2006. Using a Bayesian belief network model to categorize length of stay for radical prostatectomy patients. *Health Care Management Science*, 9(4), pp.341–348.
- Michelen, W. et al., 2006. Reducing frequent flyer emergency department visits. *J Health Care Poor Underserved*, 17(1 Suppl), pp.59–69.
- Miro, O. et al., 2003. Analysis of patient flow in the emergency department and the effect of an extensive reorganisation. *Emergency Medicine Journal*, 20(2), p.143.

- MMS: Massachusetts Medical Society, 2012 MMS Patient Access to Care Studies, Available at: <http://www.massmed.org/AM/Template.cfm?Section=Home6&CONTENTID=75485&TEMPLATE=/CM/ContentDisplay.cfm>.
- Montgomery, D.C., Runger, G.C. & Hubele, N.F., 2011. *Engineering statistics*, Hoboken, NJ: John Wiley.
- Moskop, J.C. et al., 2009a. Emergency Department Crowding, Part 1--Concept, Causes, and Moral Consequences. *Annals of Emergency Medicine*, 53(5), pp.605–611.
- Moskop, J.C. et al., 2009b. Emergency Department Crowding, Part 2--Barriers to Reform and Strategies to Overcome Them. *Annals of Emergency Medicine*, 53(5), pp.612–617.
- Moss, J. et al., 2002. A multidisciplinary Care Coordination Team improves emergency department discharge planning practice. *The Medical Journal of Australia*, 177(8).
- National Public Radio, *Accountable Care Organizations, Explained*: NPR. NPR.org. Available at: <http://www.npr.org/2011/04/01/132937232/accountable-care-organizations-explained> [Accessed September 4, 2012].
- Nightingale, D.J. & Srinivasan, J., 2011. *Beyond the Lean Revolution: Achieving Successful and Sustainable Enterprise Transformation*, AMACOM.
- Nightingale, Deborah & Rhodes, D., 2004. *Enterprise Systems Architecture: Emerging Art and Science within Engineering Systems*.
- Nightingale, Deborah & Rhodes, D., 2012. *Lectures: Enterprise Architecting*.
- Olshaker, JS & Rathlev, NK, 2006. Emergency department overcrowding and ambulance diversion: the impact and potential solutions of extended boarding of admitted patients in the emergency department. *Journal of Emergency Medicine*, 30(3), pp.351–356.
- Omachonu, V.K. et al., 2004. Predicting length of stay for Medicare patients at a teaching hospital. *Health Serv Manage Res*, 17(1), pp.1–12.
- Paradise, J., Dark, C. & Bitler, N., 2011. *Improving Access to Adult Primary Care in Medicaid: Exploring the Potential Role of Nurse Practitioners and Physician Assistants*. Available at: <http://www.kff.org/medicaid/upload/8167.pdf>.
- Parchman, M. & Culler, S., 1999. Preventable hospitalizations in primary care shortage areas. An analysis of vulnerable Medicare beneficiaries. *Archives of family medicine*, 8(6), pp.487–491.
- Park, J. et al., 2009. Factors Associated with Extended Length of Stay for Patients Presenting to an Urban Psychiatric Emergency Service: A Case-Control Study. *The Journal of Behavioral Health Services and Research*, 36(3), pp.300–308.
- Patel, P. et al., 2006. Ambulance diversion reduction: the Sacramento solution. *The American journal of emergency medicine*, 24(2), pp.206–213.

- Patel, V. et al., 2008. The coming of age of artificial intelligence in medicine. *Artificial Intelligence in Medicine*.
- Paul, J. & Sambhoos, K., 2010. Models for solving emergency room crisis. In SHS/ASQ Annual Conference. SHS ASQ Conference and Expo 2010. Atlanta, GA: SHS/ASQ.
- Paul, S.A., Reddy, M.C. & DeFlicht, C.J., 2010. A Systematic Review of Simulation Studies Investigating Emergency Department Overcrowding. *SIMULATION*, 86(8-9), pp.559–571.
- Pear, R., 2007. Medicare Says It Won't Cover Hospital Errors. *The New York Times*. Available at: <http://www.nytimes.com/2007/08/19/washington/19hospital.html> [Accessed September 3, 2012].
- Pearson, S. et al., 1994. Physician response to a prediction rule for the triage of emergency department patients with chest pain. *Journal of General Internal Medicine*, 9(5), pp.241–247.
- Peck, J., Nightingale, D. & Kim, S.-G., 2010. Axiomatic approach for efficient health care system design and optimization. *CIRP Annals - Manufacturing Technology*, 59(1), pp.469–472.
- Peck, J.S. et al., 2012. Predicting Emergency Department Inpatient Admissions to Improve Same-day Patient Flow. *Academic Emergency Medicine*, 19(9), pp.E1045–E1054.
- Peck, J.S. & Kim, Sang-Gook, 2010. Improving patient flow through axiomatic design of hospital emergency departments. *CIRP Journal of Manufacturing Science and Technology*, 2(4), pp.255–260.
- Pérez, A., Chan, W. & Dennis, R.J., 2006. Predicting the length of stay of patients admitted for intensive care using a first step analysis. *Health Services and Outcomes Research Methodology*, 6(3-4), pp.127–138.
- Pindyck, R.S. & Rubinfeld, D.L., 2009. *Microeconomics*, Upper Saddle River, N.J.: Pearson/Prentice Hall.
- Pines, J.M. et al., 2007. The Impact of Emergency Department Crowding Measures on Time to Antibiotics for Patients With Community-Acquired Pneumonia. *Annals of Emergency Medicine*, 50(5), pp.510–516.
- Pines, J.M. & Hollander, J.E., 2008. Emergency Department Crowding Is Associated With Poor Care for Patients With Severe Pain. *Annals of Emergency Medicine*, 51(1), pp.1–5.
- Poissant, L. et al., 2005. The Impact of Electronic Health Records on Time Efficiency of Physicians and Nurses: A Systematic Review. *Journal of the American Medical Informatics Association*, 12(5), pp.505–516.
- Rathlev, N.K. et al., 2007. Time Series Analysis of Variables Associated With Daily Mean Emergency Department Length of Stay. *Annals of Emergency Medicine*, 49(3), pp.265–271.

- Richardson, L.D., Asplin, B.R. & Lowe, R.A., 2002. Emergency department crowding as a health policy issue: Past development, future directions. *Annals of Emergency Medicine*, 40(4), pp.388–393.
- Robinson, G.H., Louis E. Davis & Leifer, R.P., 1966. Prediction of Hospital Length of Stay. *Health Services Research*, 1(3). Available at: <http://www.pubmedcentral.gov/articlerender.fcgi?artid=1067345> [Accessed April 11, 2010].
- Rosenthal, T.C., 2008. The Medical Home: Growing Evidence to Support a New Approach to Primary Care. *The Journal of the American Board of Family Medicine*, 21(5), pp.427–440.
- Rossetti, M.D., Trzcinski, G.F. & Syverud, S.A., 1999. Emergency department simulation and determination of optimal attending physician staffing schedules. In *Simulation Conference Proceedings, 1999 Winter*. Simulation Conference Proceedings, 1999 Winter. IEEE, pp. 1532–1540 vol.2.
- Rubino, L., Stahl, L. & Chan, M., 2007. Innovative Approach to the Aims for Improvement: Emergency Department Patient Throughput in an Impacted Urban Setting. *The Journal of Ambulatory Care Management*, 30(4), p.327.
- Sadeghi, S. et al., 2006. A Bayesian model for triage decision support. *International Journal of Medical Informatics*, 75(5), pp.403–411.
- Samaha, S., Armel, W. & Starks, D., 2003. Emergency departments I: the use of simulation to reduce the length of stay in an emergency department. In *Proceedings of the 35th conference on Winter simulation: driving innovation*. Winter Simulation Conference, pp. 1907–1911.
- Sargent, R.G., 2004. Validation and verification of simulation models. In *Simulation Conference, 2004. Proceedings of the 2004 Winter*. Simulation Conference, 2004. Proceedings of the 2004 Winter. p. 2 vol. (xliv+2164).
- Sargent, Robert G., 2000. Verification, validation, and accreditation: verification, validation, and accreditation of simulation models. In *Proceedings of the 32nd conference on Winter simulation. WSC '00*. San Diego, CA, USA: Society for Computer Simulation International, pp. 50–59. Available at: <http://dl.acm.org/citation.cfm?id=510378.510390> [Accessed October 30, 2012].
- Schmoltdt, A., Benthe, H.F. & Haberland, G., 1975. Digitoxin metabolism by rat liver microsomes. *Biochemical pharmacology*, 24(17), pp.1639–1641.
- Schull, M et al., 2003. Emergency department contributors to ambulance diversion: A quantitative analysis. *Annals of Emergency Medicine*, 41(4), pp.467–476.
- Schweigler, L.M. et al., 2009. Forecasting Models of Emergency Department Crowding. *Academic Emergency Medicine*, 16(4), pp.301–308.
- Shah, M. et al., 2006. Description and evaluation of a pilot physician-directed emergency medical services diversion control program. *Academic Emergency Medicine*, 13(1), pp.54–60.

- Shmueli, G., Patel, N.R. & Bruce, P.C., 2007. Data mining for business intelligence: concepts, techniques, and applications in Microsoft Office Excel with XLMiner, Wiley-Interscience.
- Shortell, S., Becker, S. & Neuhauser, D., 1976. The Effects of Management Practices on Hospital Efficiency and Quality of Care. *Organizational Research in Hospitals*, pp.90–107.
- Shortell, S.M. et al., 1995. Assessing the impact of continuous quality improvement/total quality management: concept versus implementation. *Health Services Research*, 30(2), pp.377–401.
- Simchi-Levi, D., Kaminsky, P. & Simchi-Levi, E., 2003. Designing and managing the supply chain: concepts, strategies, and case studies, McGraw Hill Professional.
- SKITKA, L.J., MOSIER, K.L. & BURDICK, M., 1999. Does automation bias decision-making? *International Journal of Human-Computer Studies*, 51(5), pp.991–1006.
- Smalley, H.E., 1982. Hospital management engineering : a guide to the improvement of hospital management systems, Englewood Cliffs, N.J.: Prentice-Hall.
- Solberg, L.I. et al., 2003. Emergency department crowding: Consensus development of potential measures. *Annals of Emergency Medicine*, 42(6), pp.824–834.
- Spear, S.J., 2009. Chasing the rabbit how market leaders outdistance the competition and how great companies can catch up and win, New York: McGraw-Hill. Available at: <http://search.ebscohost.com/login.aspx?direct=true&scope=site&db=nlebk&db=nlabk&AN=246693> [Accessed September 12, 2012].
- Starfield, B. & Shi, L., 2004. The Medical Home, Access to Care, and Insurance: A Review of Evidence. *Pediatrics*, 113(Supplement 4), pp.1493–1498.
- Sterman, J., 2000. Business dynamics : systems thinking and modeling for a complex world, Boston: Irwin/McGraw-Hill.
- Sun, Y. et al., 2011. Predicting Hospital Admissions at Emergency Department Triage Using Routine Administrative Data. *Academic Emergency Medicine*, 18(8), pp.844–850.
- Szolovits, P, 1995. Uncertainty and decisions in medical informatics. *Methods of Information in Medicine-Methodik der Information in der Medizin*, 34(1), pp.111–121.
- Szolovits, Peter, Patil, R. & Schwartz, W., 1988. Artificial Intelligence in Medical Diagnosis. *Annals of Internal Medicine*, 108(1), pp.80–87.
- Tandberg, D. & Qualls, C., 1994. Time Series Forecasts of Emergency Department Patient Volume, Length of Stay, and Acuity. *Annals of Emergency Medicine*, 23(2), pp.299–306.
- Trivedi, V.M., 1980. A stochastic model for predicting discharges: applications for achieving occupancy goals in hospitals. *Socio-Economic Planning Sciences*, 14(5), pp.209–215.
- Trzeciak, S & Rivers, E.P., 2003. Emergency department overcrowding in the United States: an emerging threat to patient safety and public health. *Emergency Medicine Journal*, 20(5), pp.402–405.

- US GAO, 2003. Hospital Emergency Departments: Crowded Conditions Vary among Hospitals and Communities. GAO-03-460.
- US GAO, 2009. Hospital Emergency Departments: Crowding Continues to Occur, and Some Patients Wait Longer than Recommended Time Frames. GAO-09-347.
- US Senate, 2010. The Patient Protection and Affordable Care Act, Available at: <http://democrats.senate.gov/pdfs/reform/patient-protection-affordable-care-act-as-passed.pdf>.
- Uzsoy, R., 2005. Supply-Chain Management and Health Care Delivery: Pursuing a System-Level Understanding. In *Building a Better Delivery System*. National Academies Press, pp. 143–145.
- Vassilacopoulos, G., 1985. Allocating Doctors to Shifts in an Accident and Emergency Department. *The Journal of the Operational Research Society*, 36(6), pp.517–523.
- Viccellio, A. et al., 2009. The Association Between Transfer of Emergency Department Boarders to Inpatient Hallways and Mortality: A 4-Year Experience. *Annals of Emergency Medicine*, 54(4), pp.487–491.
- Vicente, K., 2005. The Human Factor in Health Care Systems Design. In *Building a Better Delivery System*. National Academies Press, pp. 147–148.
- Vries, J. de & Huijsman, R., 2011. Supply chain management in health services: an overview. *Supply Chain Management: An International Journal*, 16(3), pp.159–165.
- Walczak, S., Pofahl, W.E. & Scorpio, R.J., 2003. A decision support tool for allocating hospital bed resources and determining required acuity of care. *Decision Support Systems*, 34(4), pp.445–456.
- Walczak, S., Scorpio, R. & Pofahl, W., 1998. Predicting Hospital Length of Stay with Neural Networks. In *Eleventh International FLAIRS Conference*. FLAIRS. merican Association for Artificial Intelligence, pp. 333–337.
- Walsh, P., Rothenberg, S.J., et al., 2004. A validated clinical model to predict the need for admission and length of stay in children with acute bronchiolitis. *European Journal of Emergency Medicine: Official Journal of the European Society for Emergency Medicine*, 11(5), pp.265–272.
- Walsh, P., Cunningham, Pdraig, et al., 2004. An artificial neural network ensemble to predict disposition and length of stay in children presenting with bronchiolitis. *European Journal of Emergency Medicine: Official Journal of the European Society for Emergency Medicine*, 11(5), pp.259–264.
- Wargon, M et al., 2009. A systematic review of models for forecasting the number of emergency department visits. *Emergency Medicine Journal*, 26(6), pp.395–399.
- Wargon, Mathias, Casalino, E. & Guidet, Bertrand, 2010. From Model to Forecasting: A Multicenter Study in Emergency Departments. *Academic Emergency Medicine*, 17(9), pp.970–978.

- Weber, E.J. et al., 2008. Are the Uninsured Responsible for the Increase in Emergency Department Visits in the United States? *Annals of Emergency Medicine*, 52(2), pp.108–115.e1.
- Weiss, S.J. et al., 2004. Estimating the Degree of Emergency Department Overcrowding in Academic Medical Centers: Results of the National ED Overcrowding Study (NEDOCS). *Academic Emergency Medicine*, 11(1), pp.38–50.
- Weiss, S.J. et al., 2005. Relationship between the National ED Overcrowding Scale and the number of patients who leave without being seen in an academic ED. *The American Journal of Emergency Medicine*, 23(3), pp.288–294.
- White Paper et al., 2012. White Paper: Exploring the generalizability of a model for predicting hospital admission from an emergency department.
- Williams, M., 2006. Hospitals and Clinical Facilities, Processes and Design for Patient Flow. *Patient Flow: Reducing Delay in Health care Delivery*, pp.45–77.
- Wilper, A.P. et al., 2008. Waits To See An Emergency Department Physician: U.S. Trends And Predictors, 1997–2004. *Health Affairs*, 27(2), pp.w84–w95.
- Wilson, M. & Nguyen, K., 2004. Bursting at the Seams: Improving Patient Flow to Help America's Emergency Departments. *Urgent Matters*. Available at: http://www.rwjf.org/files/research/Learning_Network5.pdf.
- Witten, I.H. & Frank, E., 2005. *Data mining: practical machine learning tools and techniques*, Morgan Kaufmann.
- Womack, J.P. & Jones, D.T., 2003. *Lean thinking: banish waste and create wealth in your corporation*, New York: Free Press.
- Womack, J.P., Jones, D.T. & Roos, D., 1991. *The machine that changed the world : how Japan's secret weapon in the global auto wars will revolutionize western industry*, New York, NY: HarperPerennial.
- Wong, H.J. et al., 2010. Understanding hospital and emergency department congestion: an examination of inpatient admission trends and bed resources. *CJEM: Canadian Journal of Emergency Medical Care = JCMU: Journal Canadien De Soins Médicaux D'urgence*, 12(1), pp.18–26.
- Yen, K. & Gorelick, M.H., 2007. Strategies to Improve Flow in the Pediatric Emergency Department. *Pediatric Emergency Care*, 23(10), pp.745–749.
- Zhan, C. et al., 2006. Medicare Payment For Selected Adverse Events: Building The Business Case For Investing In Patient Safety. *Health Affairs*, 25(5), pp.1386–1393.

Appendix A: Probability values for admission prediction models – VHA West Roxbury

Factor/code	Probability of code	Probability of code given admit	Probability of admit given code
Designation	P(Designation)	P(Designation Admit)	P(Admit Designation)
ER	0.624	0.989	0.538
Fast Track	0.376	0.011	0.010
Mode of Arrival	P(Mode)	P(Mode Admit)	P(Admit Mode)
Ambulatory	0.679	0.404	0.202
Stretcher	0.214	0.441	0.701
Wheelchair	0.107	0.155	0.491
Urgency Level	P(Urgency)	P(Urgency Admit)	P(Admit Urgency)
1	0.003	0.008	0.857
2	0.018	0.029	0.538
3	0.587	0.918	0.520
4	0.240	0.033	0.045
5	0.152	0.012	0.027
Patient Age	P(Age)	P(Age Admit)	P(Admit Age)
10	0.000	0.000	0.000
20	0.059	0.007	0.040
30	0.047	0.013	0.097
40	0.074	0.048	0.219
50	0.167	0.161	0.328
60	0.251	0.269	0.364
70	0.184	0.208	0.384
80	0.186	0.246	0.448
90	0.032	0.048	0.511
Physician	P(Physician)	P(Physician Admit)	P(Admit Physician)
1	0.026	0.013	0.164
2	0.016	0.013	0.284
3	0.010	0.011	0.364
4	0.009	0.011	0.444
5	0.115	0.102	0.301
6	0.000	0.000	0.000
7	0.000	0.001	1.000
8	0.005	0.002	0.130
9	0.017	0.020	0.408
10	0.060	0.063	0.356

11	0.095	0.129	0.460
12	0.037	0.038	0.344
13	0.033	0.025	0.255
14	0.000	0.001	1.000
15	0.026	0.019	0.248
16	0.004	0.004	0.333
17	0.000	0.000	0.000
18	0.016	0.011	0.224
19	0.003	0.001	0.077
20	0.004	0.004	0.294
21	0.000	0.001	1.000
22	0.004	0.001	0.125
23	0.015	0.011	0.254
24	0.024	0.032	0.460
25	0.048	0.069	0.488
26	0.007	0.016	0.742
27	0.010	0.005	0.171
28	0.019	0.010	0.175
29	0.000	0.000	0.000
30	0.002	0.000	0.000
31	0.000	0.000	0.000
32	0.157	0.139	0.300
33	0.029	0.025	0.292
34	0.006	0.006	0.308
35	0.009	0.010	0.359
36	0.002	0.002	0.300
37	0.032	0.039	0.412
38	0.020	0.031	0.518
39	0.091	0.074	0.274
40	0.027	0.038	0.478
41	0.017	0.025	0.500
Primary Complaint	P(Complaint)	P(Complaint Admit)	P(Admit Complaint)
Abdominal pain	0.048	0.069	0.485
Abdominal problems	0.050	0.075	0.505
Abnormal Labs	0.013	0.028	0.696
Admission	0.001	0.003	1.000
Allergies/hives/med reaction/sting	0.004	0.000	0.000
Assault, rape	0.001	0.000	0.000
Back pain	0.039	0.013	0.111
Bites	0.004	0.001	0.056
Body aches	0.026	0.008	0.111
Burns	0.002	0.000	0.000
Cardiac arrest	0.006	0.015	0.778
Cardio-vascular complaint	0.031	0.052	0.566

Chest pain	0.048	0.086	0.610
Cold/Flu	0.060	0.011	0.060
Convulsions, seizures	0.006	0.013	0.792
Dental, toothache	0.005	0.000	0.000
Diabetic problems	0.008	0.013	0.514
Dizzy	0.015	0.013	0.290
Ear/nose/throat problems	0.013	0.007	0.185
EDEMA/Swelling	0.024	0.022	0.307
Eye problem	0.009	0.001	0.027
Fainting/syncope	0.007	0.016	0.710
Fall	0.025	0.031	0.423
Fever	0.016	0.030	0.636
Flank pain	0.008	0.006	0.235
Fluid/nutrition alteration	0.003	0.008	0.917
Follow-up/Health Maintenance	0.049	0.005	0.034
Foreign body	0.002	0.000	0.000
Genito-urinary problem	0.026	0.016	0.211
Gun-shot wound	-	-	-
Gynecological problem	-	-	-
Headache	0.010	0.006	0.220
Hemorrhage	0.003	0.001	0.091
Industrial/machinery accidents	0.000	0.000	0.000
Infection	0.027	0.032	0.402
Ingestion (accidental)	-	-	-
Joint Problems	0.035	0.006	0.054
Kidney and Liver Failure	0.015	0.022	0.492
Laceration	0.010	0.004	0.125
Medication refill	0.025	0.000	0.000
Neck pain	0.006	0.001	0.077
Needle Stick/Exposure	0.004	0.000	0.000
Neurological complaint	0.007	0.008	0.400
Obstetrical problem	0.000	0.000	0.000
Orthopedic injury	0.012	0.013	0.353
Other (FT)	0.007	0.004	0.200
Overdose (intentional)	0.000	0.001	1.000
Peripheral vascular/leg pain	0.045	0.013	0.102
Procedure	0.011	0.004	0.111

Psychiatric/social problems	0.043	0.052	0.413
Respiratory problems	0.091	0.180	0.673
Skin complaint/trauma	0.042	0.016	0.131
Stabbing	-	-	-
Stroke/CVA	0.005	0.010	0.700
Substance abuse	0.021	0.033	0.528
Temperature related	-	-	-
Traffic injury	0.004	0.001	0.067
Traumatic injuries	0.000	0.000	0.000
Unconsciousness /unresponsive	0.000	0.001	1.000
Unknown Problem/Lethargy	0.002	0.006	0.800
Vision problems	0.004	0.007	0.667
Weakness	0.024	0.045	0.646
Total Probability of Admit	0.339		

Appendix B: Probability values for admission prediction models - Four sample hospitals

Probabilities of admission given factor	VHA 1	VHA 2	Small Private	Large Private
Total Probability of Admit	0.34	0.28	0.24	0.28
Urgency Level: P(Admit Urgency)				
ESI 1	0.86	0.25	0.82	0.96
ESI 2	0.54	0.55	0.46	0.57
ESI 3	0.55	0.34	0.26	0.29
ESI 4	0.03	0.13	0.03	0.02
ESI 5	0.03	0.11	0.00	0.00
Patient Age: P(Admit Age)				
>20	0.00	0.00	0.12	0.06
20-29	0.04	0.09	0.09	0.08
30-39	0.10	0.09	0.13	0.14
40-49	0.22	0.10	0.17	0.22
50-59	0.33	0.26	0.22	0.32
60-69	0.36	0.27	0.30	0.46
70-79	0.38	0.32	0.44	0.58
80-89	0.45	0.36	0.52	0.71
>90	0.51	0.39	0.62	0.74
Primary Complaint: P(Admit Complaint)				
Abdominal pain	0.48	0.28	0.29	0.36
Abdominal problems	0.50	0.39	0.26	0.35
Abnormal Labs	0.70	0.66	0.54	0.56
Admission	1.00	0.95	0.75	1.00
Allergies/hives/med reaction/sting	0.00	0.00	0.07	0.10
Assault, rape	0.00	0.19	0.00	0.10
Back pain	0.11	0.13	0.15	0.04
Bites	0.06	0.05	0.00	0.07
Body aches	0.11	0.16	0.16	0.17
Burns	0.00	0.00	0.00	0.00
Cardiac arrest	0.78	-	1.00	1.00
Cardio-vascular complaint	0.57	0.40	0.38	0.49
Chest pain	0.61	0.46	0.41	0.38

Cold/Flu	0.06	0.16	0.14	0.08
Convulsions, seizures	0.79	0.35	0.40	0.42
Dental, toothache	0.00	0.00	0.03	0.02
Diabetic problems	0.51	0.27	0.52	0.43
Dizzy	0.29	0.22	0.23	0.13
Ear/nose/throat problems	0.19	0.06	0.03	0.05
EDEMA/Swelling	0.31	0.24	0.17	0.25
Eye problem	0.03	0.00	0.00	0.03
Fainting/syncope	0.68	0.48	0.48	0.40
Fall	0.42	0.24	0.24	0.34
Fever	0.62	0.61	0.28	0.22
Flank pain	0.24	0.11	0.18	0.17
Fluid/nutrition alteration	0.92	0.53	0.50	0.50
Follow-up/Health Maintenance	0.03	0.01	0.15	0.05
Foreign body	0.00	0.06	0.09	0.10
Genito-urinary problem	0.22	0.13	0.13	0.17
Gun-shot wound	-	-	-	-
Gynecological problem	-	0.00	0.17	0.00
Headache	0.22	0.09	0.09	0.10
Hemorrhage	0.09	0.50	0.29	0.57
Industrial/machinery accidents	0.00	-	-	-
Infection	0.40	0.39	0.13	0.29
Ingestion (accidental)	-	-	0.00	0.00
Joint Problems	0.05	0.07	0.04	0.07
Kidney and Liver Failure	0.79	1.00	1.00	0.88
Laceration	0.13	0.00	0.04	0.01
Medication refill	0.00	0.02	0.00	0.06
Neck pain	0.08	0.12	0.12	0.03
Needle Stick/Exposure	0.00	0.00	0.00	0.00
Neurological complaint	0.40	0.32	0.24	0.56
Obstetrical problem	0.00	0.00	0.13	0.20
Orthopedic injury	0.35	0.03	0.08	0.16
Other (FT)	0.20	0.25	0.43	0.55
Overdose (intentional)	1.00	0.50	0.63	0.38
Peripheral vascular/leg pain	0.10	0.22	0.16	0.08
Procedure	0.11	0.15	0.04	0.27
Psychiatric/social problems	0.41	0.40	0.63	0.36
Respiratory problems	0.67	0.40	0.47	0.64
Skin complaint/trauma	0.13	0.07	0.04	0.10
Stabbing	-	-	-	0.45

Stroke/CVA	0.70	0.40	0.74	0.77
Substance abuse	0.53	0.56	0.26	0.20
Temperature related	-	0.00	0.33	1.00
Traffic injury	0.07	0.00	0.03	0.09
Traumatic injuries	0.00	0.00	0.00	0.46
Unconsciousness /unresponsive	1.00	0.00	0.43	0.67
Unknown Problem/Lethargy	0.80	0.73	0.60	0.66
Vision problems	0.67	0.29	0.38	0.06
Weakness	0.65	0.53	0.67	0.58
Designation: P(Admit Designation)				
ER	0.54	-	-	-
Fast Track	0.01	-	-	-
North	-	-	-	0.33
South	-	-	-	0.33
Urgent	-	-	-	0.01
Designation: P(Admit Mode of Arrival)				
Ambulatory	0.20	0.24	0.15	-
Stretcher	0.70	-	0.43	-
Wheelchair	0.49	-	0.36	-
Ambulance/Police	-	0.52	-	-
Clinic	-	0.38	-	-
Nursing Home	-	0.64	-	-
Transfer	-	0.60	-	-
Other	-	0.29	0.25	-
Designation: P(Admit Provider)				
1	0.17	0.17	0.23	0.00
2	0.28	0.27	0.22	0.43
3	0.36	0.27	0.29	0.22
4	0.44	0.28	0.20	0.00
5	0.30	0.44	0.08	0.19
6	0.00	0.30	0.23	0.29
7	1.00	0.33	0.28	0.00
8	0.13	0.26	0.14	0.16
9	0.41	0.21	0.21	0.40
10	0.36	1.00	0.31	0.12
11	0.46	0.31	0.08	0.00
12	0.34	0.28	0.09	0.32
13	0.26	0.26	0.13	0.38
14	1.00	0.13	0.19	0.00
15	0.25	0.33	0.28	0.39

16	0.33	0.19	0.30	0.38
17	0.00	1.00	0.20	0.31
18	0.22	0.29	0.05	0.27
19	0.08	0.00	0.28	0.00
20	0.29	0.30	0.27	0.31
21	1.00	0.24	0.26	0.00
22	0.13	0.29	0.24	0.27
23	0.25	0.46	0.46	0.00
24	0.46	0.14	0.26	0.31
25	0.49	0.50	0.01	0.49
26	0.74	0.44	-	0.02
27	0.17	0.35	-	0.33
28	0.18	0.00	-	0.00
29	0.00	0.15	-	0.30
30	0.00	0.28	-	0.00
31	0.00	0.33	-	0.44
32	0.30	0.25	-	0.16
33	0.29	0.17	-	0.00
34	0.31	0.35	-	0.00
35	0.36	0.34	-	0.38
36	0.30	0.23	-	0.37
37	0.41	0.23	-	0.13
38	0.52	0.20	-	0.24
39	0.27	0.20	-	1.00
40	0.48	0.15	-	0.02
41	0.50	0.36	-	0.00
42	-	0.31	-	0.34
43	-	0.32	-	0.17
44	-	0.24	-	0.25
45	-	0.36	-	0.00
46	-	0.28	-	0.26
47	-	0.45	-	0.42
48	-	0.19	-	0.07
49	-	0.16	-	-
50	-	0.33	-	-
51	-	0.15	-	-
52	-	0.23	-	-
53	-	0.12	-	-
54	-	0.23	-	-
55	-	0.67	-	-
56	-	0.00	-	-

