

ENGINEERING INCENTIVES IN DISTRIBUTED SYSTEMS WITH HEALTHCARE
APPLICATIONS

A Dissertation

by

BRANDON REED POPE

Submitted to the Office of Graduate Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

August 2011

Major Subject: Industrial Engineering

Engineering Incentives in Distributed Systems with Healthcare Applications

Copyright 2011 Brandon Reed Pope

ENGINEERING INCENTIVES IN DISTRIBUTED SYSTEMS WITH HEALTHCARE
APPLICATIONS

A Dissertation

by

BRANDON REED POPE

Submitted to the Office of Graduate Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

Approved by:

Co-Chairs of Committee,	Andrew Johnson Abhijit Deshmukh
Committee Members,	Georgia-Ann Klutke J. James Rohack Rajiv Sarin
Head of Department,	Brett A. Peters

August 2011

Major Subject: Industrial Engineering

ABSTRACT

Engineering Incentives in Distributed Systems with Healthcare Applications.

(August 2011)

Brandon Reed Pope, B.S., Abilene Christian University; M.E., Texas A&M University

Co-Chairs of Advisory Committee: Dr. Andrew Johnson
Dr. Abhijit Deshmukh

U.S. healthcare costs have experienced unsustainable growth, with expenditures of \$2.5 trillion in 2009, and are rising at a rate faster than that of the U.S. economy. A major factor in the cost of the U.S. healthcare system is related to the strategic behavior of system participants based on their incentives. This dissertation addresses the challenge of designing incentives to solve problems in healthcare systems. Principal agent theory and Markov decision processes are the primary methods used to construct incentives.

The first problem considered is how to design contracts in order to align consumer and provider incentives with respect to preventive efforts. The model consists of an insurer contracting with two agents, a consumer and a provider, and focuses on the trade off between *ex ante* moral hazard and insurance. Two classes of efforts on behalf of the provider are studied: those which complement consumer efforts, and those which substitute with consumer efforts. The results show that the provider must be given incentives when the consumer is healthy to induce effort, and that inducing provider effort allows an insurer to save on incentives given to the consumer. The insurer can save on the cost of incentives by using a multilateral contract compared to the bilateral benchmark. These savings are illustrated by an example showing which model features affect the savings achieved.

The second problem addresses the decision to provide knowledge to consumers regarding the consequences of health behaviors. The model developed to address this second problem extends the literature on incentives in healthcare systems to consider dynamic environments and includes a behavioral model of healthcare consumers. By using a learning model of consumer behavior, a policy maker's knowledge provision problem is transformed into a Markov decision process. This framework is used to solve for optimal knowledge provision policies regarding behaviors affecting coronary health. Sensitivity analysis shows robust threshold features of optimal policies. The results show that knowledge about smoking should be provided at most health and behavior states. As the cost of providing knowledge increases or aptitude for behavioral change decreases, fewer states are in the optimal knowledge provision policy, with healthy consumers dropping out first. Knowledge about diet and physical activity is provided more selectively due to the uncertainty in the health benefits, and the time delay in accrued rewards.

DEDICATION

I dedicate this dissertation to my family and friends. Special recognition goes first to my wife and best friend, Kami, and second to my parents, Don and Beth, for their many years of love and support.

ACKNOWLEDGMENTS

I wish to thank my committee members for their significant contributions to my work and development. I feel the variety of expertise on my committee has aided me in strengthening many areas of my dissertation. Special thanks go to my advisors, Dr. Andrew Johnson and Dr. Abhijit Deshmukh, for their willingness to work with me to investigate a new frontier in engineering healthcare systems.

I would like to thank the U.S. Department of Education's Graduate Assistance in Areas of National Need Program (GAANN P200A060241), as well as the National Science Foundation's Civil, Mechanical and Manufacturing Innovation Division (CMMI-0944670) for their generous financial support in allowing me to pursue my personal research interests. I would also like to thank the Industrial and Systems Engineering Department at Texas A&M University for financial and technical support. I have had the pleasure of working with many excellent faculty and staff during my time in the department.

NOMENCLATURE

AAD	average absolute deviation
AHA	American Heart Association
AMI	acute myocardial infarction
ARIC	Atherosclerosis Risk In Communities
CHD	coronary heart disease
CIAD	conditional independence - additive dependence
CIMD	conditional independence - multiplicative dependence
CMH	Cochran-Mantel-Haenszel
EU	expected utility
EWA	experience-weighted attraction
FOC	first order condition
GDP	gross domestic product
HDL	high-density lipoprotein
IC	incentive compatibility
IND	independence
IR	individual rationality
FOC	first order condition
MAD	maximum absolute deviation
MH	Mantel-Haenszel
MDP	Markov decision process
NHLBI	National Heart, Lung, and Blood Institute
QALY	quality-adjusted life year
SSD	sum of squared deviations

TABLE OF CONTENTS

	Page
ABSTRACT	iii
DEDICATION	v
ACKNOWLEDGMENTS	vi
NOMENCLATURE	vii
TABLE OF CONTENTS	viii
LIST OF TABLES	x
LIST OF FIGURES	xii
1. INTRODUCTION	1
2. LITERATURE REVIEW	4
2.1 Distributed Systems	4
2.2 Healthcare Systems Engineering	6
2.2.1 Incentives in Healthcare Systems	7
3. MULTILATERAL CONTRACTING AND PREVENTION	10
3.1 Literature Review	10
3.2 Model	12
3.2.1 Basic Assumptions and Notations	13
3.3 Analysis	16
3.3.1 Private Insurer's Problem	17
3.3.2 Substitutive Efforts	28
3.3.3 Value of Multilateral Contracting	31
3.3.4 Sensitivity Analysis	33
3.4 Results	36
4. DYNAMIC KNOWLEDGE-BASED INCENTIVES IN HEALTHCARE	38

	Page
4.1 Literature Review	40
4.2 Model	42
4.3 MDP Control for a Policy Maker	45
5. KNOWLEDGE INCENTIVES FOR CORONARY HEART DISEASE	50
5.1 Literature Review	51
5.2 Model Specification	53
5.2.1 Consumer Health and Behavior	53
5.2.2 Consequences of CHD	56
5.2.3 Policy Maker Decisions and Rewards	56
5.3 Parameter Estimation	58
5.3.1 Health Transition Probabilities	58
5.3.2 Behavioral and Health Costs	59
5.3.3 Learning Rule Parameters	61
5.4 Model Validation	62
5.4.1 Ideal Validation: Design of Experiments	66
5.5 Results	68
5.5.1 Future Extensions	71
6. CONCLUSIONS	73
REFERENCES	75
APPENDIX A	83
APPENDIX B	99
VITA	122

LIST OF TABLES

TABLE	Page
3.1 Sensitivity Analysis of Multilateral Savings	34
5.1 Categorization of Health States	54
B.1 Summary of ARIC Behavioral Data	99
B.2 χ^2 Tests for Independence	100
B.3 CMH Tests for Independence	101
B.4 Comparing Conditional Independence Models	106
B.5 Model Fitness with Training and Testing Data	106
B.6 Comparing Attraction Regression Models	108
B.7 Comparing Attraction Regression Models: Positive Attractions	109
B.8 Regression with Sub-behavior Consistency	111
B.9 Attraction Changes Given Exercise	111
B.10 Regression with Sub-behavior Consistency: Positive Attractions	112
B.11 Pearson Correlations Between Simulated Annealing Probabilities	116
B.12 Maximum Absolute Deviations Between Simulated Annealing Probabilities	116
B.13 Average Absolute Deviations Between Simulated Annealing Probabilities .	117
B.14 Pearson Correlations - β_{EXR} Parameters	118
B.15 Pearson Correlations - β_{FAT} Parameters	118

TABLE	Page
B.16 Pearson Correlations - β_{SMK} Parameters	119
B.17 Pearson Correlations- μ_{BP} Parameters	119
B.18 Pearson Correlations - μ_{HDL} Parameters	120
B.19 Pearson Correlations - μ_{TC} Parameters	120
B.20 Pearson Correlations - All Parameters	121

LIST OF FIGURES

FIGURE	Page
3.1 Contract Timeline	15
3.2 Comparison of Consumer Incentives	26
3.3 Value of Multilateral Contracting	33
3.4 Sensitivity of Savings to Consumer's Linear Disutility Coefficient	35
4.1 Policy Maker's Markov Decision Process	47
5.1 Optimal Policy for Smoking Knowledge Provision	63
5.2 Optimal Smoking Knowledge Provision Policies by Cost of Knowledge	65
5.3 Optimal Diet Knowledge at $s = (4, 2, 10)$	69
5.4 Optimal Exercise Knowledge at $s = (3, 5, 10)$	70
A.1 Sensitivity to a with $b = 0.20$	88
A.2 Sensitivity to a with $b = 0.25$	89
A.3 Sensitivity to a with $b = 0.35$	90
A.4 Sensitivity to a with $b = 0.40$	91
A.5 Sensitivity to a with $b = 0.45$	92
A.6 Sensitivity to a with $b = 0.50$	93
A.7 Sensitivity to r with $q = 0.03$	94
A.8 Sensitivity to r with $q = 0.04$	95

FIGURE	Page
A.9 Sensitivity to r with $q = 0.05$	96
A.10 Sensitivity to r with $q = 0.06$	97
A.11 Sensitivity to r with $q = 0.07$	98
B.1 Dependence Pathways	102

1. INTRODUCTION

The U.S. healthcare system is a world leader in many dimensions, however cost efficiency is not one of them. U.S. healthcare costs have risen at a much higher rate than that of other developed nations, and in 2009 healthcare spending stood at 17.3% of the U.S. GDP (Truffer et al., 2010). By comparison, no other developed nation spends more than 11% of GDP on healthcare (Organisation for economic co-operation and development, 2009). These expenditures total \$2.5 trillion, and are rising at an unsustainable rate faster than that of the U.S. economy. Although part of the cost increase can be attributed to the development of new medications, availability of advanced diagnostic and surgical procedures, and an aging population base, these factors are present in other developed nations, and in fact the U.S. has one of the youngest median ages of the developed world (Rohack and Einboden, 2006). Estimates on the excess cost in the system consistently exceed \$500 billion (Institute of Medicine, 2010).

The rapid increase in costs has caused engineers to apply their tools and methods at the healthcare industry, focusing primarily on the delivery and operational aspects of the system. Such approaches typically neglect the inherently decentralized nature of the U.S. healthcare system. The strategic behavior of healthcare system participants is a major factor in the cost increases (Porter and Teisberg, 2006), and requires distributed solutions to control problems. Policy makers and payers have an interest in controlling the autonomous decisions of patients, providers, and other system participants, which can have enormous implications for healthcare expenditures. An estimated 80% of heart disease, stroke, and Type 2 diabetes could be prevented by controlling the risk factors of diet, physical inactivity, and tobacco use (World Health Organization, 2009).

This dissertation follows the style of *Management Science*.

The estimated cost of these first two diseases alone was \$503.2 billion in the U.S. for the year 2010 (Lloyd-Jones et al., 2010). Incentives are a potential solution garnering increasing interest as a means of distributed control for healthcare systems (Institute of Medicine, 2010; Valdez et al., 2010). Incentives can also perversely influence strategic behavior of participants in several ways. For example, insurance contracts can reduce incentives for preventive care by removing the burden of risk from individuals. Since modifying behavior through incentives is costly, the problem is to design incentives which balance the benefits from controlled behavior against the costs of constructing the incentives.

The overall theme for this dissertation is designing incentives in a distributed healthcare system. This research promises to be generalizable to other systems since as systems become larger and increase in complexity, centralized strategies become harder to implement and less feasible as means of system control. Other systems such as healthcare and education systems contain some level of inherent distribution in decision making autonomy and thus require distributed solutions to control problems. One means of distributed control is through the use of incentives to guide autonomous decisions. The task of engineering incentives has significant potential to improve system outcomes for distributed systems such as the healthcare system. Contracts, policies, and information are among the ways that incentives are conveyed to agents in distributed systems. This dissertation investigates the question of how these tools can be designed to provide incentives to autonomous agents in the healthcare system. The structure of this dissertation is as follows: Section 2 provides a general review of the literature on healthcare systems engineering research and the work on incentives and strategic behavior in healthcare. More specific literature related to each model is discussed within the respective sections. Section 3 studies the question of how multilateral contracts should be structured when healthcare consumers and providers can potentially take preventive efforts. Section 4 addresses some of the strong assumptions in the first model, and presents a model for designing incentives using information

for consumers who learn about healthcare behaviors in a dynamic setting. Section 5 takes the model introduced in Section 4, and applies the model to designing incentives to control costs from coronary heart disease. Section 6 provides conclusions.

2. LITERATURE REVIEW

2.1 Distributed Systems

Distributed systems pose challenges for engineers. Rather than optimizing or designing from a centralized perspective, in distributed systems the information, resources, and decision making authority are distributed throughout multiple agents. Shoham and Leyton-Brown (2009) provide an overview of research in distributed systems including optimization, equilibrium concepts, and welfare problems. Much of the focus on distributed systems has come from the computer science community, where distributed algorithms as well as parallel and distributed computation are studied (Bertsekas and Tsitsiklis, 1997). The classic example of a distributed system is an economy, in which agents (firms and consumers) make transactions according to their own preferences and resources. The foundational theory of Arrow and Debreu (1954) gives conditions under which an efficient outcome for the distributed system is realized. However, many features of real world distributed systems destroy such efficiencies. Among these are asymmetric information, hidden actions, missing markets, externalities, bounded rationality, and bounded computational abilities. In such environments, system engineers face a challenge to move from inefficient system outcomes towards efficient outcomes.

Sensor networks, supply chains, and healthcare systems are just a few other examples of distributed systems which have garnered the attention of engineers. Recently, education systems have become an interesting application of incentives. Fryer (2010) examines data from school systems where financial incentives were tested to enhance student performance. Education systems have interesting parallels to healthcare systems since at least two groups of decision makers (students and teachers or consumers and providers) need to have proper incentives for the system to achieve good outcomes.

An analysis of systems with multiple decision makers is often built upon game theoretic principles. Classic game theory was developed in the 1940's and 50's by Von Neumann and Morgenstern (1944) and Nash (1951). See Fudenberg and Tirole (1991) for a comprehensive treatment. The theory seeks to motivate and predict outcomes and prescribe decisions in situations of strategic interdependence. These situations arise when agents' optimal decisions depend on the decisions chosen by other agents. The concept of Nash equilibrium predicts outcomes in which agents best-respond to other agents' decisions and have mutually consistent beliefs about how other agents play. Distributed systems engineering makes use of game theoretic principles since designs must be evaluated in terms of the equilibria they sustain.

The theory of incentives in distributed systems has largely been developed within the economic community. The field of mechanism design (Mas-Colell et al., 1995, Ch. 23) studies the problem a designer faces in constructing the rules by which a distributed system functions. A mechanism designer faces an environment, $e \in E$, and an outcome space Y . Given the designer's objective correspondence $F : E \rightarrow Y$, and beliefs about equilibrium behavior, the goal is to design a mechanism $\Gamma = \langle M, h \rangle$ consisting of a message space M and an outcome function $h : M \rightarrow Y$ such that the mechanism implements the designer's objective correspondence. A primary application of mechanism design theory is auction theory (Krishna, 2009), where in some circumstances, the rules of the auction can be designed to meet the designer's goals of efficiency and/or truthful revelation of private information. The special bilateral case of mechanism design is called principal-agent theory (Laffont and Martimort, 2002). In principal-agent settings, such as the supervisor-employee, lawyer-client, and teacher-student relationships, the principal wants to contract with the agent regarding some task in order to ensure preferred outcomes are reached. In realistic cases when the agent has more information regarding the task at hand, or the principal cannot perfectly observe the agent's action, the full information efficient outcomes

are generally unattainable. These causes of inefficiency are called adverse selection and moral hazard, stemming from the insurance literature. Mechanism design and principal-agent theory have produced results and are increasingly applied in engineering and operations research Cachon and Zhang (2006); Fuloria and Zenios (2001); Gallien (2006); Su and Zenios (2006). Bolton and Dewatripont (2005) provide a comprehensive survey of the literature on contracting in principal-agent and more general multiagent circumstances. For a summary of the empirical evidence to support the theory of incentives, see Prendergast (1999). As a distributed system, healthcare systems are of particular interest because of the complexity, cost, and inefficiency.

2.2 Healthcare Systems Engineering

The recent explosion of healthcare costs in the U.S. has led to an equal explosion in healthcare research in systems engineering and related fields. A large portion of this growth has focused on applying existing methods to familiar engineering problems in healthcare settings such as scheduling (Patrick et al., 2008), logistics (Daskin and Dean, 2004), and supply chain management (Pierskalla, 2004). Another component of the growth can be largely classified as tackling problems in medical decision making, such as optimal timing of interventions (Denton et al., 2009; Shechter et al., 2008) and optimal procedure protocols (Lee and Zaider, 2008). Decision making at the policy level has also received some attention in the areas of substance control and treatment (Tragler et al., 2001), immunizations (Engineer et al., 2009), cost-effectiveness studies (Owens et al., 2004), and the diffusion and value of health information technology (Diana, 2009). Another major area of interest to policy makers (Antos et al., 2009), yet receiving little attention from the engineering community, is related to the decision making, as guided by incentives, of the distributed players in the healthcare system. Existing efforts include Fuloria and Zenios

(2001) who study incentives between a payer and a provider in a dynamic environment where treatment intensity is unobservable, and Yaesoubi and Roberts (2009), who study contracts to coordinate actions between a healthcare purchaser and a provider controlling the number of individuals receiving an intervention. The majority of the existing literature studying incentive issues in healthcare systems comes from the health economics community, where research has focused on incentives between consumers, providers, and insurers.

2.2.1 Incentives in Healthcare Systems

Consumers of healthcare introduce several inefficiencies into the healthcare system, stemming from both private information and non-contractible actions. When these are non-existent, the full information, or first best, insurance contract provides a more risk averse consumer with full insurance from a less risk averse insurer. However, when unobserved, or at least non-contractible actions (e.g. diet, exercise) which affect the consumer's probability of illness are present, insurance against healthcare costs leads to *ex ante* moral hazard. Uninsured consumers would take preventive actions when the expected benefits exceed the costs, but insured consumers will have reduced incentive to take such actions when they are costly. Ehrlich and Becker (1972) first considered this mode of moral hazard and gave conditions under which the second best contract leaves the consumer with some risk. These *ex ante* moral hazard circumstances arise frequently in health insurance since consumers have many actions which insurers cannot fully observe, yet affect their risk of healthcare expenditures. In order to leave the consumer with some risk, insurers write contracts with deductibles, co-insurance rates, and copayments. Additionally, if true health states of consumers are impossible or too costly to observe, health insurance contracts cannot be written on the state of the consumer's health, and therefore contracts are

frequently written on the basis of healthcare expenditures. *Ex post* moral hazard arises in this environment since consumers who face reduced marginal costs of treatment will demand treatment at excessive levels. The seminal work of Zeckhauser (1970) studies the trade off between this *ex post* moral hazard and efficiency of risk reduction, then designs optimal incentive contracts for efficient expenditures. Zweifel et al. (2009, ch.6) present models for both modes of moral hazard. Within *ex ante* models, they focus on binary preventive efforts and conditions for which full insurance or copayments are optimal. Within *ex post* models, they conclude that copayments should be used to control moral hazard, with higher copayments for more price elastic services. Other recent extensions to these models include Ellis and Manning (2007), who consider a consumer facing both *ex ante* and *ex post* moral hazard. They considers linear coverage for observable decisions of prevention and treatment and derive the optimal coinsurance rates for each type of good. Goldman and Philipson (2007) consider optimal insurance with a consumer under *ex post* moral hazard with multiple goods. Their conclusions suggest that ceterus paribus, insured goods which are substitutable will have lower copays, and those which are complements will have higher copays.

Consumers also have private information regarding their risk types. In these circumstances, adverse selection occurs since individuals can select their most favorable plan from a menu, with sicker consumers choosing more generous plans, and healthier consumers choosing more moderate plans. Adverse selection creates inefficiencies whether the equilibrium involves pooling or separating of types (Rothschild and Stiglitz, 1976). These inefficiencies can be controlled by designing an optimal menu of contracts, or by adjusting insurance parameters based on characteristics of the consumer. For a summary of the literature on adverse selection see Cutler and Zeckhauser (2000).

While the literature initially focused on demand-side incentives through the consumer's insurance contract, provider incentives, especially those specified by their remuneration or

payment contract, has been another focus of the literature on incentives in healthcare systems. When providers whose incentives are not perfectly aligned with consumers make quantity decisions (Ellis and McGuire, 1986), or have the ability to select patients (Ellis, 1998), the strategic behavior of providers driven by supply-side incentives in their remuneration contract can create inefficiencies. The consensus of the literature is that when reimbursement is contingent on expenditures, retrospective schemes (such as fee-for-service, the most common method) will lead to excessive services and up-coding (reclassifying patients into more lucrative diagnoses), while prospective schemes will lead to under provision of services and avoidance of high-severity patients, and that mixed schemes can balance these trade offs. Prospective schemes include capitation, under which a provider is paid a set amount per time period for each patient, regardless of services delivered. Models which consider both consumer and provider incentives are similar to the model presented in Section 3. This literature is reviewed and related to this dissertation's model in the following section.

3. MULTILATERAL CONTRACTING AND PREVENTION

This section analyzes the problem of designing incentives through contracts taking into consideration the multilateral interactions in healthcare systems. The primary focus here is on the trade off between *ex ante* moral hazard and insurance, considering both consumer and provider incentives to solve the problem of optimal contracting in the presence of hidden preventive efforts. Results show that if inducing the provider's effort is optimal, the provider must be given incentives when the consumer is healthy. That is, the provider must be better off when consumers are healthy rather than ill. Inducing the provider's effort allows an insurer to save on incentives given to consumers by distorting incentives from the bilateral benchmark. The interaction between consumer and provider efforts is modeled as both complementary and substitutive, showing the results to be robust. The optimal multilateral contract is compared with the optimal bilateral benchmark, and an example illustrates which model features and parameters affect the overall savings that the multilateral contract is able to achieve. Subsection 3.1 reviews literature related specifically to the model and approach taken in this section. Subsection 3.2 introduces the model, notations, and assumptions. Subsection 3.3 analyzes the insurer's optimal multilateral contract, studying the incentives given to each agent and the savings that the multilateral contract achieves relative to the bilateral contract. Finally, subsection 3.4 summarizes the results found.

3.1 Literature Review

The research considering the important problem of how incentives interact amongst agents in the healthcare system has focused primarily on the *ex post* dimension of moral hazard. Ellis and McGuire (1990) study incentives for a provider and consumer who

bargain over utilization and conclude that the optimal incentive system gives generous insurance coverage to consumers, but gives providers incentives to control costs. Ma and McGuire (1997) consider moral hazard when both a provider and insurer have non-contractible actions in the production of health, and the insurer sees only a report (possibly non-truthful) of treatment. Their results show that providers need cost-sharing incentives when these actions are substitutes, and cost-plus incentives when they are complements. Ma and Riordan (2002) study optimal contracts using both demand-side and supply-side incentives, and study the level of utilization incurred relative to the full information benchmark. This literature highlights the need to consider both supply-side and demand-side incentives in order to efficiently control agency in healthcare.

Prevention is an important topic to consider for reasons beyond moral hazard. Kenkel (2000) provides a general review of the literature on the economics of prevention, including moral hazard, externalities, lack of consumer information, and cost-effectiveness. The model introduced in this section distinguishes itself from these efforts by focusing on the *ex ante* dimension of moral hazard. One of the primary arguments diminishing the importance of *ex ante* moral hazard in health insurance has been that non-financial costs (e.g. pain, discomfort, suffering) associated with adverse health events are uninsurable, therefore even financially insured consumers will have incentives to exert preventive efforts. Even if this argument has been valid in the past, the trend in medical research and technology is progressing towards minimizing or completely eliminating these non-financial costs. For example, the Door-to-Balloon (D2B) Initiative of the American College of Cardiology encourages hospitals to strive towards reducing the time from when an acute myocardial infarction patient enters the door to the time the angioplasty balloon is in the chest of the patient to under 90 minutes. Removing the non-financial burden of disease will lead to increased *ex ante* moral hazard amongst insured consumers. An example of this unintended consequence could be argued from the case of coronary heart disease (CHD). From

1970 to 2006, hospital-case fatality rates for CHD in patients age 65 and over fell substantially from near 40% to just above 10%. This accomplishment of medical science led to a greatly diminished non-financial risk of CHD. Over the same time horizon, knowledge of and ability to prevent CHD improved greatly, which would have presumably contributed to lowering incidence rates. However, hospitalization rates for the same group of consumers remained roughly constant over the same time horizon (National Heart, Lung and Blood Institute, 2009). This type of argument echoes Kenkel (2000), who notes that as prevention and cure become more perfect substitutes, the *ex ante* moral hazard problem becomes larger.

Empirical evidence for *ex ante* moral hazard in healthcare has been mixed, but appears to be building. Courbage and de Coulon (2004), using U.K. data, find no evidence for *ex ante* moral hazard with respect to exercising, check-ups, and smoking. Stanciole (2008), using U.S. data, does find evidence of *ex ante* moral hazard in the choice of heavy smoking, lack of exercise, and obesity. In the work most closely related to this section, Dave and Kaestner (2009) find a *ex ante* moral hazard effect regarding physical activity and tobacco consumption, and present evidence that providers do in fact influence consumer decisions regarding preventive efforts. This influence of the provider highlights the need to consider multilateral incentives as modeled in the next subsection.

3.2 Model

In order to create a more efficient healthcare system, solving bilateral incentive problems is not enough. The interactions of incentives between consumers, providers, and insurers, and the resulting strategic behaviors must be studied. The main contribution of this section is the consideration of both the consumer's and provider's roles in prevention, and the implications for the optimal multilateral contract to control *ex ante* moral

hazard. While most preventive efforts are ultimately in the hands of consumers, many consumers look to providers for guidance and direction in prevention (Town et al., 2005), and thus the provider's incentives need to be considered. The first provider efforts considered are those which complement preventive efforts from the consumer. Examples of such provider efforts with regard to cardiovascular disease include: asking about tobacco use, advising every tobacco user to quit, encouraging 30 to 60 minutes of moderate intensity aerobic exercise most days of the week, assessing body mass index on each visit and consistently encouraging weight maintenance. More generally, these efforts can be thought of as counseling and promoting the consumer's preventive effort, explaining the benefits and consequences of prevention, and educating the consumer about how to best implement preventive efforts. Although providers are ethically motivated to keep consumers healthy, poorly designed incentives can put the ethical incentives and financial incentives in conflict. Ethical motivations instead should support financial incentives to keep patients healthy. Since the insurer acting as the principal, contracts with both the provider and the consumer, the solution to the optimal contracting problem should be viewed in the light of a principal-multiagent problem with heterogeneous agents, and optimal contracts must consider the interaction of incentives. The issue of adverse selection on the part of the provider is abstracted away from, taking the consumer to be representative of the population.


3.2.1 Basic Assumptions and Notations

The basic model considers a risk averse consumer facing an uncertain health state, which will be either healthy (h) or ill (i). That is, there are no varying degrees of illness, or at least there is a single clear cut treatment option which restores the consumer to full health, and does not vary with the level of illness. While this may seem like a strong

assumption, it is approximately valid for acute illness episodes, and this assumption allows focus on the *ex ante* dimension of moral hazard. The consumer may obtain insurance to alleviate risk stemming from uncertain health states and healthcare expenditures. In the case of illness, the consumer visits the provider for treatment which costs d to administer. Let the insurer (I) offer a contract to the provider (P) and the consumer (C). Both the consumer and the provider are modeled as having an effort action (e_C, e_P) , with $e_P, e_C \in [0, 1]$. These efforts are hidden from the insurer, and yet relevant for determining the level of prevention utilized. The provider's effort can be thought of as advocating or promoting the prevention to the consumer, and the consumer's action can be thought of as physically taking the preventive action. The provider's effort incurs a cost $c_P(e_P)$, which reflects the time and other resources required to exert the effort. The provider's cost is assumed to be increasing and convex, with $c_P(0) = 0$. This convexity can be explained by arguing that the provider can initially find the time to exert this effort without sacrificing much in the way of other activities. As the level of effort increases, increasingly attractive activities must be sacrificed which could have brought revenue or utility. The provider's utility is assumed to be separable in income and effort and initially assume the provider to be risk neutral. The provider's effort serves to lower the disutility experienced by the consumer when the prevention is taken. The consumer experiences disutility $\psi(e_C, e_P)$ from exerting effort. This disutility is increasing and convex in e_C ($\frac{\partial \psi}{\partial e_C} \geq 0, \frac{\partial^2 \psi}{\partial e_C^2} \geq 0$) and decreasing both absolutely and marginally in the provider's effort ($\frac{\partial \psi}{\partial e_P} \leq 0, \frac{\partial^2 \psi}{\partial e_C \partial e_P} \leq 0$). These decreases in consumer disutility can be thought of as the benefits of the provider's effort from the insurer's perspective, and making the assumption that marginal benefits of the provider's effort are decreasing ($\frac{\partial^2 \psi}{\partial e_P^2} \geq 0, \frac{\partial^3 \psi}{\partial e_C \partial e_P^2} \geq 0$). We normalize $\psi(0, e_P)$ to 0. The consumer's Bernoulli utility function over wealth is denoted by $u(\cdot)$, which is strictly increasing and concave, and consumer utility is taken as separable in income and effort.

Figure 3.1 shows the time line of the contracting problem. In the first stage, the insurer offers a contract which is accepted or rejected. In the second and third stages the provider and consumer respectively choose their efforts, with the consumer observing the effort level of the provider. In the fourth stage nature determines the consumer's health state, and finally in the fifth stage the contract is executed.

Figure 3.1. Contract Timeline



The consumer's effort impacts the probability distribution over health states. When the consumer exerts preventive effort e_C , the probability of being healthy is $\pi(e_C) \in [0, 1]$, where $\pi'(\cdot) > 0, \pi''(\cdot) < 0$. Notice that in this model, the provider's effort and the consumer's effort are not substitutes in the sense that no amount of effort from the provider can directly impact the probability over health states. This modeling assumption is geared towards capturing the preventive actions that providers have in influencing consumer choices such as those mentioned previously for cardiovascular disease. These actions are in contrast to preventive actions providers may take which directly substitute for the consumer's effort such as vaccinations. The effect of substitutive efforts is considered in subsection 3.4. After receiving provider effort e_P , exerting effort e_C , and facing incomes y_h, y_i in the case of health or illness, the consumer's expected utility will be

$$U(y_h, y_i, e_C, e_P) = \pi(e_C)u(y_h) + [1 - \pi(e_C)]u(y_i) - \psi(e_C, e_P). \quad (3.1)$$

The consumer is assumed to be educated regarding the health benefits of the preventive effort, and maximizes expected utility (EU) with respect to the prevention decision. Under these conditions, the EU maximizing consumer will exert effort until the point where marginal benefit of prevention is equal to the marginal cost,

$$\pi'(e_C)\Delta u = \frac{\partial \psi}{\partial e_C}(e_C, e_P). \quad (3.2)$$

Where $\Delta u = u(y_h) - u(y_i)$ is the risk the consumer faces, the marginal value of staying healthy. Denote this optimal level of consumer prevention as a function of the provider's effort and the risk the consumer faces by $e_C(e_P; \Delta u)$. Clearly only partially or uninsured consumers will exert positive effort. Then by assumptions on $\psi(\cdot)$ and $\pi(\cdot)$, when Δu is positive, $e_C(\cdot)$ is increasing in e_P and Δu (see details in Appendix). The insurer can increase consumer effort by two means; inducing more provider effort and thus lowering the consumer's disutility of effort, or increasing the risk the consumer faces. The two effects are found empirically by Dave and Kaestner (2009). Both these controls have costs. Exposing the consumer to greater risk will limit the transfers the insurer can extract from the consumer (i.e. how much the consumer is willing to pay for the insurance), and inducing the provider's effort requires more payments.

3.3 Analysis

The outcome of the system depends on the effort levels exerted by the agents. These autonomous decisions are products of the agents' incentives, which can be modified by a contract. How this contract should be optimally designed is the main focus of the analysis.

3.3.1 Private Insurer's Problem

The analysis is performed from a private insurer's perspective, with the insurer modeled as having contracting power. This is likely the most realistic assumption since in many healthcare markets the largest 2 insurers control the bulk of the market (American Medical Association, 2007). Using data from 44 states, 80% of states have the top two insurers serving greater than 60% of the market. The average share of the top two insurers across these 44 states is 70.23%. Under this assumption the insurer will offer a contract to maximize its own objective, taken to be profit (or positive margin in the case of a nonprofit insurer). The contract offered by the insurer will specify a set of transfers from the consumer and to the provider contingent upon whether the consumer is healthy or ill and the intended effort levels for the consumer and the provider ($\{t_h^C, t_i^C, t_h^P, t_i^P, e_C, e_P\}$), and guarantee the provider will treat the consumer in the case of illness. The strong assumption that transfers can be made contingent on health states is weakened by the minimal health state space ($\{h, i\}$). Under the mild conditions that healthy consumers do not seek treatment while ill consumers do seek treatment, observing expenditures is equivalent to observing the binary health state. This argument relates expenditures and health to a single condition and individual. This presents little conflict with the geographic variations literature, which casts doubt on the correlation between spending and illness based on data encompassing many individuals, diseases, and other complex factors. Notice that from the insurer's perspective, a contract establishes a random payment to be made to the provider and a random payment to be received from the consumer contingent on the health outcome. These random variables will be denoted by \tilde{t}^P and \tilde{t}^C . For comparison, the first case considered is when the insurer has complete information regarding the agents' efforts. Complete information could be obtained by a costly observation process. After establishing the complete

information results, the more interesting and realistic case of hidden information is considered.

Complete Information

In the case when efforts are observable, the individual rationality (IR) constraints of the agents are active in the insurer's problem. These constraints ensure that each agent receives in expectation at least a reservation level of utility. Each agent's reservation utility is the utility they could obtain without participating in the contract. Let $u_h = u(w - t_h^C)$ and $u_i = u(w - t_i^C)$ denote the utilities of the healthy or ill consumer, with initial wealth w . Also let $f(\cdot)$ denote the consumer's inverse utility function (that is $f(u(y)) = y$), which is guaranteed to exist by the assumptions on $u(\cdot)$. Assuming that uninsured consumers receive no effort from a provider, the consumer's reservation utility can be written as

$$U^0 = \pi(e_C(0; \Delta u^0))u(w) + (1 - \pi(e_C(0; \Delta u^0)))u(w - d) - \psi(e_C(0; \Delta u^0), 0),$$

where $\Delta u^0 = u(w) - u(w - d)$. The provider's reservation utility will be normalized to zero. Then the provider's and consumer's IR constraints are given as

$$\pi(e_C)t_h^P + (1 - \pi(e_C))(t_i^P - d) - c_P(e_P) \geq 0, \quad (3.3)$$

$$\pi(e_C)u_h + (1 - \pi(e_C))u_i - \psi(e_C, e_P) \geq U^0. \quad (3.4)$$

With complete information, the insurer can specify effort levels in the contract, observe the levels of effort exerted, and heavily penalize the agents if the contracted levels are not

followed. The insurer's objective is to maximize expected profits given by $\mathbb{E}[\tilde{t}^C - \tilde{t}^P]$.

Then the insurer's problem is given as

$$\max_{\{u_h, u_i, t_h^P, t_i^P, e_C, e_P\}} \pi(e_C)[w - f(u_h) - t_h^P] + (1 - \pi(e_C))[w - f(u_i) - t_i^P]$$

subject to (3.3) and (3.4). Letting λ and μ denote the Lagrange multipliers of (3.3) and (3.4), respectively, forming the Lagrangian, and taking derivatives w.r.t. u_h, u_i, t_h^P , and t_i^P yields the following conditions:

$$\lambda = 1 \Rightarrow \mathbb{E}[\tilde{t}^P] = c_P(e_P) + (1 - \pi(e_C))d, \quad (3.5)$$

$$\mu = f'(u_i) = f'(u_h) \Rightarrow u_i = u_h = U^0 + \Psi(e_C, e_P). \quad (3.6)$$

The provider's expected payments only cover the cost of effort plus the expected cost of treating the consumer. It is worth noting that if the provider were risk averse, an optimal contract would make him equally well off in each state of nature, similar to the consumer. Such a contract could only be accomplished by a zero cost sharing scheme, with $t_i^P - t_h^P = d$. The consumer obtains full insurance since marginal utility from income is assumed to be identical in all states of health. The insurer's problem then becomes

$$\max_{e_C, e_P} w - f(U^0 + \Psi(e_C, e_P)) - c_P(e_P) - (1 - \pi(e_C))d. \quad (3.7)$$

Inspecting this objective, when $\Psi(e_C, e_P)$ is convex (for which $\frac{\partial^2 \Psi}{\partial e_C^2} \frac{\partial^2 \Psi}{\partial e_P^2} \geq \left[\frac{\partial^2 \Psi}{\partial e_C \partial e_P} \right]^2$ is a sufficient condition), since $f(\cdot)$ is convex, then $f(U^0 + \Psi(\cdot))$ is as well. Then since

$-f(U^0 + \psi(\cdot))$, $-c_P(\cdot)$, and $\pi(\cdot)$ are all concave, the insurer's objective (3.7) is concave.

The first order conditions (FOC) then give that

$$\pi'(e_C)d = \frac{\partial \psi}{\partial e_C}(e_C, e_P) \cdot f'(U^0 + \psi(e_C, e_P)), \quad (3.8)$$

$$c'_P(e_P) = -\frac{\partial \psi}{\partial e_P}(e_C, e_P) \cdot f'(U^0 + \psi(e_C, e_P)). \quad (3.9)$$

Here (3.8) shows that the first-best efforts which solve the insurer's complete information problem equate the marginal savings in treatment payments to the provider by increasing consumer effort and the increased cost of ensuring the consumer's participation. Also, (3.9) shows that the increased cost of ensuring provider participation by increasing provider effort must be equated with the decreased cost of ensuring consumer participation. By offering a contract which satisfies (3.5),(3.6),(3.8), and (3.9), the insurer will maximize expected profit while ensuring participation by the consumer and the provider.

Incomplete Information

When efforts are unobservable, the insurer must write the contract to ensure that the agents exert the specified level of effort and accept the contract. In order for the contract to be followed by all parties, it must satisfy the incentive compatibility (IC) constraints in addition to the IR constraints. These IC constraints ensure that the agents' best actions are to exert the effort specified by the contract. The provider's IC constraint can be written as

$$e_P \in \arg \max_{\hat{e}_P} \pi(e_C(\hat{e}_P; \Delta u))t_h^P + [1 - \pi(e_C(\hat{e}_P; \Delta u))](t_i^P - d) - c_P(\hat{e}_P).$$

The consumer's IC constraint can be written as

$$e_C \in \arg \max_{\hat{e}_C} \pi(\hat{e}_C)u_h + [1 - \pi(\hat{e}_C)]u_i - \Psi(\hat{e}_C, e_P).$$

When these IC constraints are solutions to concave programs, they can be replaced by their FOC's. Earlier assumptions ensure that the consumer's IC constraint is concave, and a sufficient condition for concavity of the provider's IC constraint is that $e_C''(\cdot) \leq 0$. This intuitive condition is that the provider's effort has decreasing marginal ability to induce consumer effort. The concavity of both constraints is considered in more detail in the Appendix, where some mild technical conditions are provided to ensure the concavity of the IC constraints

$$\pi'(e_C(e_P; \Delta u))e_C'(e_P; \Delta u)[\Delta t^P + d] = c_P'(e_P), \quad (3.10)$$

$$\pi'(e_C)\Delta u = \frac{\partial \Psi}{\partial e_C}(e_C, e_P), \quad (3.11)$$

where $\Delta t^P + d = t_h^P - (t_i^P - d)$, is the provider's marginal value of keeping the consumer healthy. From (3.10) one can see that the provider must have $\Delta t^P + d > 0$ in order to exert any effort. Since $\pi'(\cdot), e_C'(\cdot), c_P'(\cdot)$ are all positive, if $\Delta t^P + d \leq 0$, then the solution to the provider's IC constraint will be to set $e_P = 0$. Therefore, in any situation where inducing provider effort is optimal, the provider's remuneration contract must make him better off when the consumer is healthy as compared to when the consumer is ill. This result is in contrast to the typical use of cost-plus and fee based schemes where providers are only reimbursed for procedures and services delivered to ill consumers. In circumstances when provider efforts have significant influence over a consumer's preventive behavior, inducing provider effort is likely desirable, and the insurer's optimal contract must create marginal value for the provider when the consumer is healthy.

The insurer's incomplete information problem then becomes

$$\max_{\{u_h, u_i, t_h^P, t_i^P, e_C, e_P\}} \pi(e_C)[w - f(u_h) - t_h^P] + (1 - \pi(e_C))[w - f(u_i) - t_i^P]$$

subject to (3.3),(3.4),(3.10),(3.11). Although consumer and provider efforts appear in the insurer's contracting problem, they are autonomously chosen, hidden and thus non-enforceable. However, the contract must be written according to the IC constraints of both agents. Therefore, the efforts could be dropped from the contract without any effect, the transfers determine the efforts chosen with the IC constraints are concave. Again let λ, μ denote the Lagrange multipliers of the IR constraints, and additionally let γ, δ denote the multipliers of (3.10), and (3.11) respectively. Forming the Lagrangian and taking derivatives with respect to u_h and u_i yields

$$\frac{1}{u'(w - t_h^C)} = \mu + \delta \frac{\pi'(\cdot)}{\pi(\cdot)} + \gamma \left(\frac{\Delta t^P + d}{\pi(\cdot)} \right) \left[\pi''(\cdot) e'_C(\cdot) \frac{\partial e_C(\cdot)}{\partial \Delta u} + \pi'(\cdot) \frac{\partial e'_C(\cdot)}{\partial \Delta u} \right], \quad (3.12)$$

$$\frac{1}{u'(w - t_i^C)} = \mu - \delta \frac{\pi'(\cdot)}{1 - \pi(\cdot)} - \gamma \left(\frac{\Delta t^P + d}{1 - \pi(\cdot)} \right) \left[\pi''(\cdot) e'_C(\cdot) \frac{\partial e_C(\cdot)}{\partial \Delta u} + \pi'(\cdot) \frac{\partial e'_C(\cdot)}{\partial \Delta u} \right]. \quad (3.13)$$

The first two terms on the right hand sides of (3.12) and (3.13) are the standard terms found in the second best bilateral contract between an insurer and consumer in the presence of *ex ante* moral hazard. The final term represents a distortion from the classic second best result due to the interaction between the consumer and the provider. The interesting factor in this term, $\frac{\partial e'_C(e_P; \Delta u)}{\partial \Delta u}$, can be interpreted as the change in the impact of provider effort due to a change in risk the consumer faces.

Proposition 3.3.1 *The consumer's effort decision is a submodular function of risk and the provider's effort, that is*

$$\frac{\partial e'_C(e_P; \Delta u)}{\partial \Delta u} \leq 0.$$

Proof in Appendix. It is tempting to think this term should be positive by following the logic that when the consumer is exposed to greater risk (large value of Δu), the provider's effort in promoting the prevention should have greater impact. Or similarly when the consumer faces very little risk (Δu is small), the provider's effort in recommending prevention should not be worth much. This thinking, however does not take into consideration the assumption of diminishing marginal returns to efforts. That is, the consumer already has strong incentives to exert preventive efforts when faced with great risk, and will prevent without any recommendation from the provider. Exerting more preventive effort makes less impact on the probability of illness, and will cost more. Similarly, when the consumer faces little risk, not much preventive effort will be exerted, in which case a recommendation from the provider makes more impact since the consumer's effort still has relatively significant impact on the chance of illness, and does not cause excessive disutility. Cast in terms of insurance completeness, this result can be interpreted as consumers with more incomplete (e.g. via higher cost sharing) insurance have strong incentives for prevention, and therefore will exert less incremental effort when encouraged by the provider.

Now going back to the Lagrangian and constructing the FOC's with respect to t_h^P and t_i^P yields

$$\lambda + \gamma \left[\frac{\pi'(e_C)}{\pi(e_C)} e'_C(e_P; \Delta u) \right] = 1, \quad (3.14)$$

$$\lambda - \gamma \left[\frac{\pi'(e_C)}{1 - \pi(e_C)} e'_C(e_P; \Delta u) \right] = 1. \quad (3.15)$$

If $\gamma \neq 0$, (3.14) and (3.15) cannot hold simultaneously, since both bracketed terms are positive. Thus, the optimal solution must have $\gamma = 0$ and $\lambda = 1$. Recalling that in the classical bilateral principal-agent theory, the agent's risk preferences cause the distortion

in the second-best bilateral contract, the effects of the provider's risk preferences on the optimal multilateral contract are now investigated.

Provider Risk Attitude

Assume now that the provider has risk averse preferences characterized by a strictly increasing and concave utility function $v(\cdot)$, the IR and IC constraints become

$$\begin{aligned}\pi(e_C)v(t_h^P) + (1 - \pi(e_C))v(t_i^P - d) - c_P(e_P) &\geq 0, \\ \pi'(e_C(e_P; \Delta u))e'_C(e_P; \Delta u)\Delta v &= c'_P(e_P),\end{aligned}$$

where $\Delta v = v(t_h^P) - v(t_i^P - d)$. The FOC's with respect to t_h^P and t_i^P of the Lagrangian from now become

$$\lambda + \gamma \left[\frac{\pi'(e_C)}{\pi(e_C)} e'_C(e_P; \Delta u) \right] = \frac{1}{v'(t_h^P)}, \quad (3.16)$$

$$\lambda - \gamma \left[\frac{\pi'(e_C)}{1 - \pi(e_C)} e'_C(e_P; \Delta u) \right] = \frac{1}{v'(t_i^P - d)}. \quad (3.17)$$

Since $v'(\cdot) > 0$, the provider's risk attitude will not change the result that $\Delta t^P + d > 0$ (equivalently $\Delta v > 0$) in order to induce effort. When the provider's effort is induced, by the concavity of $v(\cdot)$,

$$v'(t_h^P) < v'(t_i^P - d) \quad \Rightarrow \quad \frac{1}{v'(t_h^P)} > \frac{1}{v'(t_i^P - d)}.$$

In this case (3.16) and (3.17) must be solved with a positive value of γ . Similarly, if the provider is risk loving, then $\gamma < 0$. Then returning to the optimal multilateral contract with the consumer, the following result relates consumer incentives in the multilateral contract


to the first best solution of full insurance and the bilateral second best consumer incentives.

Corollary 3.3.1 *When provider effort is induced ($\Delta t^P + d > 0$),*

- *A risk neutral provider $\Rightarrow \gamma = 0 \Rightarrow$ there is no distortion from the second best bilateral incentives.*
- *A risk averse provider $\Rightarrow \gamma > 0 \Rightarrow$ there is a distortion from the second best back towards the first best.*
- *A risk loving provider $\Rightarrow \gamma < 0 \Rightarrow$ there is a further distortion away from the first best incentives.*

The corollary shows that there is no distortion in the consumer's incentives when the provider is risk neutral. In general, the second best multilateral contract will differ from the second best bilateral contract even in the case of provider risk neutrality. Since the provider's effort aides the insurer in inducing consumer effort, the optimal level of consumer effort chosen in the multilateral contract will be higher than that in the bilateral contract.

The result also gives that when the provider is risk averse, the insurer's optimal contract will shift the consumer's incentives back towards the first best bilateral contract of full insurance. Just as distorting incentives away from full insurance is costly to give to a risk averse consumer, distortions back towards full insurance provide a savings. This can be seen in Figure 3.2, in which u_h, u_i represent a full insurance contract, u_h^*, u_i^* show the classic bilateral distortion due to moral hazard, and u_h^{**}, u_i^{**} show the new distortion attainable by the multilateral contract with a risk averse provider. The figure shows that distortions given to a risk averse consumer are costly in expectation. Multiplying (3.16)

Figure 3.2. Comparison of Consumer Incentives

by $\pi(e_c)$, and adding (3.17) times $(1 - \pi(e_c))$ gives

$$\lambda = \frac{\pi(e_c)}{v'(t_h^P)} + \frac{1 - \pi(e_c)}{v'(t_i^P - d)} > 0.$$

Since $\lambda > 0$, the provider's participation constraint is assured to be binding. Similarly, multiplying (3.12) by $\pi(e_c)$, and adding (3.13) times $(1 - \pi(e_c))$ gives

$$\mu = \frac{\pi(e_c)}{u'(w - t_h^C)} + \frac{1 - \pi(e_c)}{u'(w - t_i^C)} > 0.$$

Similarly, $\mu > 0$ forces the consumer's participation constraint to be binding. Transfers for the consumer and provider can now be solved by the system of equations: (3.3),(3.4),(3.10),(3.11) all holding at equality. Solving leads to the following results:

$$t_h^C = w - f \left(U^0 + \Psi(e_C, e_P) + \frac{1 - \pi(e_C)}{\pi'(e_C)} \frac{\partial \Psi(e_C, e_P)}{\partial e_C} \right) \quad (3.18)$$

$$t_i^C = w - f \left(U^0 + \Psi(e_C, e_P) - \frac{\pi(e_C)}{\pi'(e_C)} \frac{\partial \Psi(e_C, e_P)}{\partial e_C} \right) \quad (3.19)$$

$$\Rightarrow \Delta u = \frac{1}{\pi'(e_C)} \frac{\partial \Psi(e_C, e_P)}{\partial e_C},$$

$$t_h^P = g \left(c_P(e_P) + \frac{c'_P(e_P)[1 - \pi(e_C)]}{\pi'(e_C)e'_C(e_P; \Delta u)} \right) \quad (3.20)$$

$$t_i^P = d + g \left(c_P(e_P) - \frac{c'_P(e_P)\pi(e_C)}{\pi'(e_C)e'_C(e_P; \Delta u)} \right), \quad (3.21)$$

where $g(\cdot)$ is the provider's inverse utility function (that is $g(v(y)) = y$). These transfers could be used by an insurer to create an insurance contract for consumers, and remuneration scheme for providers which gives optimal incentives for preventive efforts. Consumer transfers given in the second best multilateral contract can be reduced from the second best bilateral contract in two dimensions. Firstly, since the provider's effort directly makes the consumer better off, the insurer can extract more transfers in the contract. Secondly, since the provider's effort reduces the marginal disutility of effort for the consumer, the insurer can also reduce the consumer's incentives measured by distortions from full insurance, which are costly to give to a risk averse consumer. The provider's contract makes him better off when the consumer is healthy (when income is t_h^P) than when the consumer is ill (when income is $t_i^P - d$). This marginal utility required to induce the provider's effort will cost the insurer beyond the expected costs of treatment and effort, which would be

the costs associated with the second best bilateral contract. Since $t_h^P + d < t_i^P$, this contract with the provider imposes cost sharing by paying the provider less than d extra when the patient is ill.

3.3.2 Substitutive Efforts

Although the types of provider efforts discussed so far would seem to interact as complements with consumer efforts, other activities on the part of the provider would likely substitute with the consumer's effort. Examples of these efforts would include free samples of medicine and other elements of treatment not reported. These efforts are modeled as directly impacting the consumer's probability of illness ($\pi(e_C, e_P)$), rather than impacting the consumer's disutility from effort. The provider's effort contributes to preventing disease ($\pi'_{e_P}(\cdot) > 0$) but with decreasing effectiveness ($\pi''_{e_P}(\cdot) < 0$). The provider's and consumer's efforts are modeled as substitutes ($\pi''_{e_C e_P}(\cdot) < 0$). The following analysis investigates under what conditions do the primary results still hold: that the provider must be better off when the consumer is healthy to induce effort, and that inducing provider effort allows consumer incentives to be shifted back towards the first best of full insurance.

Facing risk Δu and effort e_P from the provider, the consumer chooses effort $e_C \in \arg \max_{\hat{e}_C} \pi(\hat{e}_C, e_P)\Delta u - \psi(\hat{e}_C)$. This objective is again concave, and the consumer's FOC imposes that

$$\pi'_{e_C}(e_C, e_P)\Delta u = \psi'(e_C). \quad (3.22)$$

Based on the assumptions above, it is easily shown that the consumer again increases effort in response to higher risk ($\frac{\partial e_C}{\partial \Delta u} \geq 0$), but now decreases effort in response to provider effort ($\frac{\partial e_C}{\partial e_P} \geq 0$). Heading straight to the case of incomplete information, the provider's IC is now

$$e_P \in \arg \max_{\hat{e}_P} v_i + \Delta v \cdot \pi(e_C(\hat{e}_P; \Delta u), \hat{e}_P) - c_P(\hat{e}_P).$$

Again focusing on the case when the provider's IC constraint is concave, the FOC gives that

$$\Delta v \left(\frac{\partial \pi(\cdot)}{\partial e_C} \frac{\partial e_C}{\partial e_P} + \frac{\partial \pi(\cdot)}{\partial e_P} \right) = c'_P(e_P). \quad (3.23)$$

Proposition 3.3.2 $\frac{\partial}{\partial e_C} \left(\frac{\pi'_{e_C}(\cdot)}{\pi'_{e_P}(\cdot)} \right) \leq 0 \Rightarrow$ *the provider must be better off when healthy than when ill to exert any effort.*

Proof in Appendix. This condition, that as the consumer's effort increases, the consumer's effort becomes less effective in prevention relative to the provider's effort, ensures that increasing provider efforts lead to a higher likelihood of health on behalf of the consumer.

The insurer's incomplete information contracting problem is

$$\max_{\{u_h, u_i, v_h, v_i, e_C, e_P\}} \pi(e_C, e_P)[w - f(u_h) - g(v_h)] + (1 - \pi(e_C, e_P))[w - f(u_i) - g(v_i)]$$

subject to (3.22),(3.23), and the individual rationality constraints

$$u_i + \pi(e_C, e_P)\Delta u - \psi(e_C) \geq 0 \quad (3.24)$$

$$v_i + \pi(e_C, e_P)\Delta v - c_P(e_P) \geq 0. \quad (3.25)$$

Using the same multipliers as before, differentiating the Lagrangian with respect to v_h, v_i gives that

$$\lambda + \frac{\gamma}{\pi(\cdot)} \left[\pi'_{e_C} \frac{\partial e_C}{\partial e_P} + \pi'_{e_P} \right] = \frac{1}{v'(t_h^P)}$$

$$\lambda + \frac{\gamma}{1 - \pi(\cdot)} \left[\pi'_{e_C} \frac{\partial e_C}{\partial e_P} + \pi'_{e_P} \right] = \frac{1}{v'(t_i^P - d)}.$$

The condition given in Proposition 3.3.2 ensures the bracketed term is positive, and therefore provider risk preferences ensure the same signs on γ as were found in the case of complementary efforts. Differentiating with respect to u_h, u_i gives that

$$\frac{1}{u'(w - t_h^C)} = \mu + \delta \frac{\pi'(\cdot)}{\pi(\cdot)} + \gamma \left(\frac{\Delta v}{\pi(\cdot)} \right) \left[\pi''_{e_C} \frac{\partial e_C}{\partial \Delta u} \frac{\partial e_C}{\partial e_P} + \pi'_{e_C} \frac{\partial^2 e_C}{\partial e_P \partial \Delta u} + \pi''_{e_C e_P} \frac{\partial e_C}{\partial \Delta u} \right]$$

$$\frac{1}{u'(w - t_i^C)} = \mu - \delta \frac{\pi'(\cdot)}{1 - \pi(\cdot)} - \gamma \left(\frac{\Delta v}{1 - \pi(\cdot)} \right) \left[\pi''_{e_C} \frac{\partial e_C}{\partial \Delta u} \frac{\partial e_C}{\partial e_P} + \pi'_{e_C} \frac{\partial^2 e_C}{\partial e_P \partial \Delta u} + \pi''_{e_C e_P} \frac{\partial e_C}{\partial \Delta u} \right].$$

The final term again represents the multilateral distortion due to the provider's influence. The direction of the distortion again depends upon the sign of the term $\frac{\partial^2 e_C}{\partial e_P \partial \Delta u}$.

Proposition 3.3.3 *When $\frac{\partial}{\partial e_C} \left(\frac{\pi'_{e_C}(\cdot)}{\pi'_{e_P}(\cdot)} \right) \leq 0$ and $\frac{\partial^2 e_C}{\partial e_P \partial \Delta u} \leq 0$, the conclusions of Corollary 3.3.1 hold for substitutive efforts.*

Proof in Appendix. This proposition shows that the previous findings are not unique to a single type of interaction between the patient and provider. Attention is now turned to the value of multilateral contracting. Since the primary conclusions hold for both types of provider efforts considered, only the first case of complementary efforts is presented.

3.3.3 Value of Multilateral Contracting

The optimal multilateral contract provides both a savings on consumer incentives and an additional cost on provider incentives when compared with the bilateral contract. Since the insurer can simply induce zero effort from the provider, the savings will always outweigh the costs in the optimal multilateral contract. Without knowing the consumer and provider utility functions, this is impossible to evaluate how much the insurer is able to save via multilateral contracting. Thus to investigate, the consumer and the provider are modeled by inverse utility functions $f(u) = u + ru^2$ and $g(v) = v + qv^2$. These functions are not meant to provide general solutions, but rather to illustrate the possible savings using plausible functions. These risk averse utility functions exhibit decreasing absolute risk aversion, and higher values of r and q are associated with higher levels of risk aversion. Then for given effort levels, the expected profits for the insurer under the multilateral contract given by (3.18)-(3.21) is

$$\begin{aligned}
& \pi(e_C)[t_h^C - t_h^P] + (1 - \pi(e_C))[t_i^C - t_i^P] = w - \mathbb{E}[\tilde{u} + r\tilde{u}^2] - \mathbb{E}[\tilde{v} + q\tilde{v}^2] - \bar{\pi}(e_C)d \\
& = w - \left[U^0 + \Psi(e_C, e_P) + r(U^0 + \Psi(e_C, e_P))^2 + r\pi(e_C)\bar{\pi}(e_C) \left(\frac{1}{\pi'(e_C)} \frac{\partial \Psi(e_C, e_P)}{\partial e_C} \right)^2 \right] \\
& \quad - \left[c_P(e_P) + q(c_P(e_P))^2 + q\pi(e_C)\bar{\pi}(e_C) \left(\frac{c'_P(e_P)}{\pi'(e_C)e'_C(e_P; \Delta u)} \right)^2 \right] - \bar{\pi}(e_C)d, \quad (3.26)
\end{aligned}$$

where \tilde{u}, \tilde{v} represent the random utilities facing the consumer and provider under the terms of the contract, and $\bar{\pi}(\cdot) = 1 - \pi(\cdot)$ for notational brevity. The insurer's optimal bilateral contract will induce effort from the consumer by setting

$$\begin{aligned} u_h^{BL} &= U^0 + \psi(e_C, 0) + \frac{1 - \pi(e_C)}{\pi'(e_C)} \frac{\partial \psi(e_C, 0)}{\partial e_C} \\ u_i^{BL} &= U^0 + \psi(e_C, 0) - \frac{\pi(e_C)}{\pi'(e_C)} \frac{\partial \psi(e_C, 0)}{\partial e_C} \end{aligned}$$

The insurer's expected profit from the bilateral contract will be


$$\left[U^0 + \psi(e_C, 0) + r(U^0 + \psi(e_C, 0))^2 + r\pi(e_C)\bar{\pi}(e_C) \left(\frac{1}{\pi'(e_C)} \frac{\partial \psi(e_C, 0)}{\partial e_C} \right)^2 \right] + w - \bar{\pi}(e_C)d. \quad (3.27)$$

Let e_C^* represent the optimal effort level induced in the optimal bilateral contract, and e_C^{**}, e_P^{**} the optimal efforts induced in the multilateral contract. Then the insurer's benefit from using a multilateral contract will be (3.26) evaluated at e_C^{**}, e_P^{**} minus (3.27) evaluated at $e_C^*, e_P = 0$. The analytical expression for this difference is too complicated to be of direct interest. To illustrate the savings, the remaining functions in the model are also parameterized. Let $\pi(e_C) = \sqrt{e_C}$, $c_P(e_P) = m \cdot e_P^2$, and $\psi(e_C, e_P) = e_C[a + b \cdot e_C - k\sqrt{e_P}]$. These forms are chosen for simplicity and to satisfy the earlier assumptions. The cost of information under the bilateral optimal contract is compared to the cost of information under the multilateral optimal contract.

Figure 3.3 shows the cost of information savings that are made possible by using the multilateral contract. The axes show k and m , which represent how much the provider influences consumer disutility, and how costly the provider's effort is to induce. The value represented on the vertical axis shows $\frac{\text{profit}_{ML}^{\text{incomplete}} - \text{profit}_{BL}^{\text{incomplete}}}{\text{profit}_{ML}^{\text{full}} - \text{profit}_{BL}^{\text{incomplete}}}$. When $m \approx 0$, the insurer

can save nearly all of the cost of information, since the provider's effort can be induced without cost. When $k \approx 0$, the insurer can save nearly nothing compared to the bilateral benchmark since the provider's effort makes no difference to the consumer.

Figure 3.3. Value of Multilateral Contracting $(a, b, d, q, r, w) = (0.3, 0.3, 5, 0.05, 0.10, 50)$



3.3.4 Sensitivity Analysis

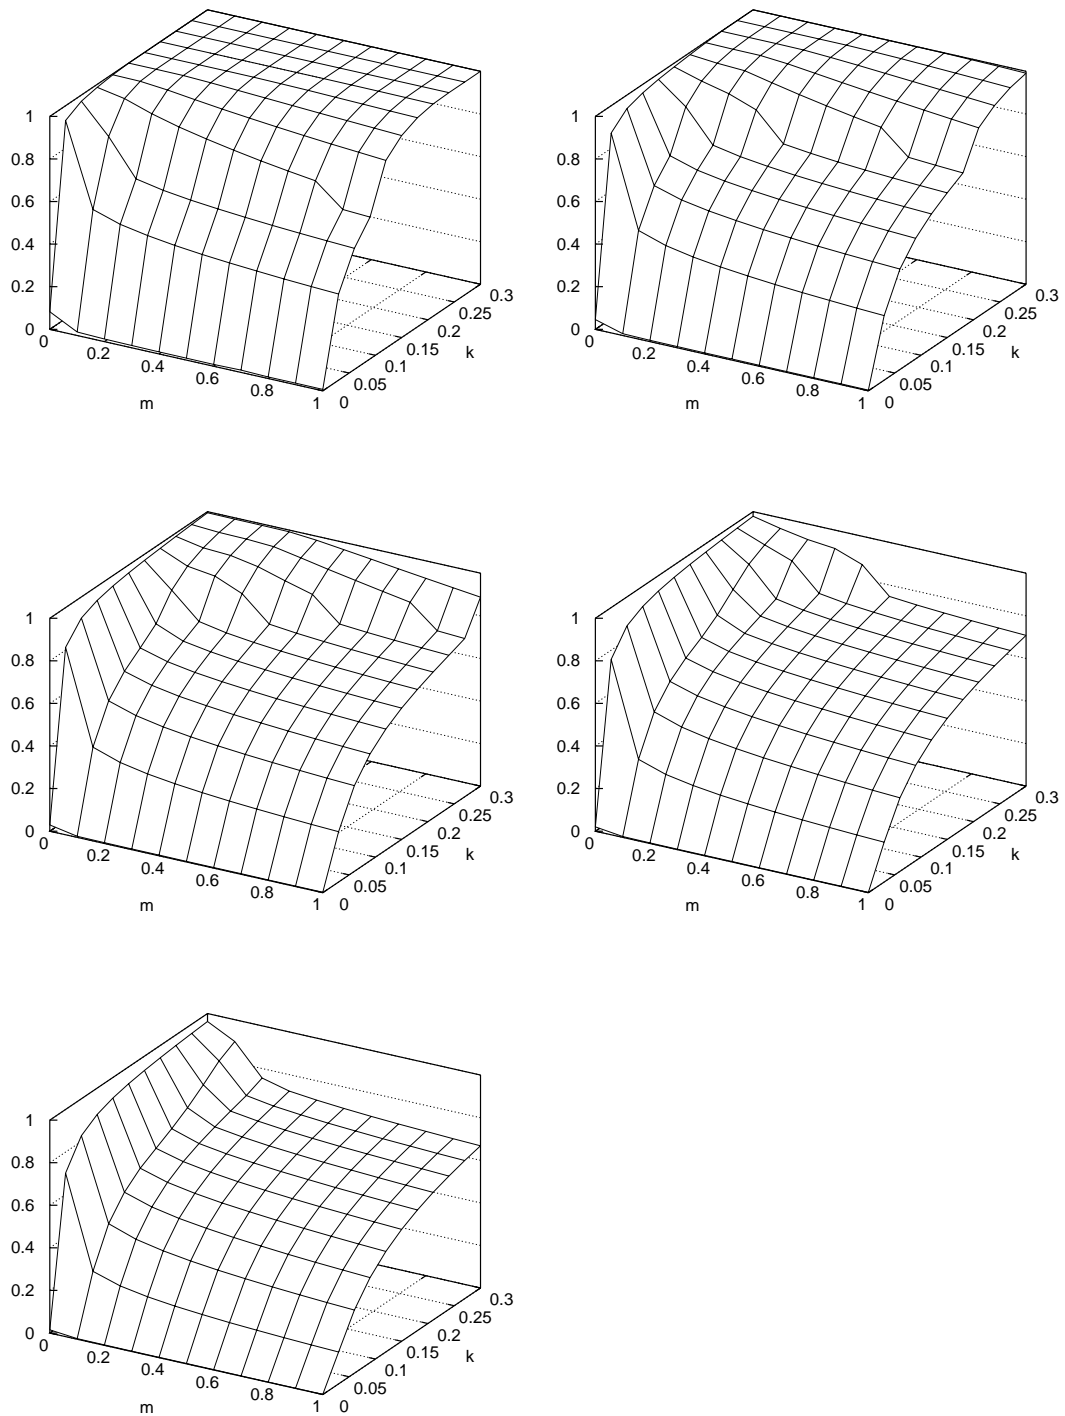
From Figure 3.3, there appear to be two distinct regions. One where the savings are relatively constant in m and rise in k^θ , with $\theta < 1$, and a second where approximately the entire cost of information can be saved by the multilateral contract. The first region would seem to be driven by the marginal disutility of the consumer's effort, $\frac{\partial \phi}{\partial e_C} = a + 2be_C - k\sqrt{e_P}$. This expression is constant in m , and as k increases, the profits attainable by the multilateral contract increase by the reduction in $\frac{\partial \phi}{\partial e_C}$ as seen in (3.26). The second region,

Table 3.1
Sensitivity Analysis of Multilateral Savings

Parameter (i)	Threshold Effect ($F_i(\cdot)$)
k	> 0
q	≈ 0
a, b, r	< 0

where the full cost of information can be recovered by the multilateral contract, appears to be driven by a threshold. Further analytical investigation of this region proves challenging, however by varying the parameters of the model, several insights are available. Figure 3.4 shows how this threshold changes as the parameter a , which captures the consumer's linear coefficient of disutility, varies. In the figure, a increases from the top left and moving to the right. As the parameter changes, the threshold for m clearly decreases. In addition to a , the parameters b, r , and q were varied to determine their effect on this threshold. From the base settings reported in Figure 3.3, a pair of two-way sensitivity analyses were performed with $(a, b) \in \{0.1, 0.2, 0.3, 0.4, 0.5\} \times \{0.20, 0.25, 0.30, 0.35, 0.40, 0.45, 0.50\}$ and $(q, r) \in \{0.05, 0.075, 0.10, 0.125, 0.15\} \times \{0.03, 0.04, 0.05, 0.06, 0.07\}$. The results show that the threshold below which m allows full savings of the cost of information, $m < F(a, b, q, r, k)$, is a decreasing function of the consumer's linear and quadratic coefficients of disutility from effort (a and b) and the consumer's risk aversion (r), an increasing function of the provider's influence on the consumer's disutility (k), and is nearly constant in the provider's risk aversion (q). Graphs showing the results of each trial of both two way sensitivity analyses are located in the Appendix. Table 3.1 summarizes the effects of each parameter on the threshold which defines the second region.

Figure 3.4. Sensitivity of Savings to Consumer's Linear Disutility Coefficient. $a = \{0.1, 0.2, 0.3, 0.4, 0.5\}$, $(b, d, q, r, w) = (0.30, 5, 0.05, 0.10, 50)$



3.4 Results

This model in this section considered a health insurer contracting with a healthcare consumer and provider whose efforts interact to stochastically produce the health outcome experienced by the consumer. Focusing on the trade off between *ex ante* moral hazard and insurance, the optimal insurance contract to induce preventive efforts was studied. Although this analysis was done under the assumption of a private, profit motivated insurer, a similar analysis could be undertaken to design optimal contracts with providers and consumers for a public insurer such as Medicare. The provider's effort interacted with the consumer's effort in both complementary and substitutive fashions to capture the various activities providers may perform. The multilateral model highlights the two options for controlling the preventive efforts of the consumer at the insurer's disposal: modifying the consumer's risk over health outcomes, and inducing the provider's effort. The results showed conditions under which any optimal contract where the provider exerts effort, both agents must be better off when the consumer is healthy. This finding supports recent interest in devising new payment systems for provider accountability (Antos et al., 2009) including mechanisms which focus on health outcomes rather than services. Comparing with the second best bilateral benchmark, a risk averse provider allowed the optimal multilateral consumer incentives to be shifted back towards the first-best contract of full insurance.

While guaranteed to dominate the bilateral benchmark, the optimal multilateral contract imposes both costs and savings when compared to the bilateral benchmark. By parameterizing functions from the general model, an example illustrated how the multilateral contract is able to save on information costs, and sensitivity analysis showed how the savings possible vary as parameters of the model changed.

This model makes several assumptions to maintain tractability. One particular assumption called into question by the structure of the provider's incentives is the absence of selection against unhealthy consumers. Risk-adjusting incentives to keep heterogeneous consumers healthy would be one extension to reduce this inefficiency. Contracts and the contingent payments they generate are just one method of designing incentives in health-care systems. Providing knowledge to encourage healthy behaviors would be another, and is addressed in the next section.

4. DYNAMIC KNOWLEDGE-BASED INCENTIVES IN HEALTHCARE

As discussed, incentives are a major area of interest to healthcare policy makers and administrators. The majority of the existing work studying healthcare systems incentives focuses on the interaction between consumers, providers, and insurers to control inefficiencies from information asymmetries and distributed autonomy. While the primary means of giving incentives studied in the literature has been contracts and the contingent payments they generate, incentives can be generally thought of as any mechanism that affects decision making. This broader view of incentives motivates interest in other means of modifying behavior besides payments. The incentives considered in this section are created by providing knowledge to healthcare consumers. Certainly a lengthy discussion could be presented about the relationship and distinctions between education, information, knowledge, wisdom, and behavior. These specific definitions and relationships are beside the point of this dissertation. The question addressed here is that since education/information/knowledge can modify behavior but is costly to provide, it is worth investigating under what conditions is such provision prudent. The term *knowledge* is preferred to the term *information* to avoid the significant and varied loaded meanings from various disciplines.

A general framework is presented for analyzing incentives in dynamic environments. The scenario considered is that of a principal giving incentives to an autonomous agent making repeated decisions which affect the agent's state. This framework is not limited to healthcare applications, the state of the agent could be quite general. The agent's state evolves stochastically according to the current state and decision, and the agent's payoffs are a random function of the state and decision made. Based on this general system, a principal with preferences over the agent's outcome would like to control the agent's decisions. However, the agent's autonomy prevents direct control and incentives must

be used. This general dynamic control problem can model a variety of principal-agent-decision-state scenarios for which incentives could be constructed. This section focuses on the case of a policy maker interested in controlling healthcare consumers' behavioral decisions affecting health states.

The model developed in this section makes three contributions to the study of incentives in healthcare systems. The first contribution is the study of incentives in a dynamic setting, whereas much of the existing literature on incentives in healthcare focuses on static environments. Dynamics settings are important in healthcare applications since health states are long lived and chronic diseases account for more than 75% of total healthcare expenditures (Centers For Disease Control and Prevention, 2009). The second contribution is the departure from the majority of the literature on incentives in healthcare which uses the expected utility (EU) framework by modeling consumer behavior via a learning rule. Departure from this classic paradigm is motivated by the demands of rationality, information, and intelligence imposed by the EU framework, which are especially strong in healthcare and other complex settings. Empirical studies have also cast doubt on the EU model as a description of healthcare consumer decision making (Hibbard et al., 1997), and have shown consumers to be better at learning valuable behaviors when payoffs or reinforcements are repeated as in the case of chronic diseases (Cutler and Lleras-Muney, 2010), implying the use of learning rules when facing such decisions. Finally, this section provides an analytical approach to designing knowledge-based incentives, rather than the financial incentives typically considered in the literature. This model provides a framework for describing how knowledge affects a consumer's decision making through a state of attractions to various behaviors. The consumer's attraction and health state evolution over time, and a policy maker's provision of knowledge combine to create a Markov decision process (MDP) from the policy maker's perspective.

The remainder of this section is organized as follows. Subsection 4.1 provides a review of the literature on modeling in medical decision making at the clinical and policy levels, consumer behavior and learning in healthcare and other settings, and the role of health education as an incentive. Subsection 4.2 presents a model of dynamic behavior, learning, and outcomes concerning a healthcare consumer, and subsection 4.3 discusses how a policy maker could control this process.

4.1 Literature Review

Modeling health states and diagnoses as stochastic processes is a well accepted practice in the medical decision making literature (Briggs et al., 2006). The Markovian assumption assumes that probabilistic transitions are based on only the current state information. The Markovian assumption in healthcare modeling has been popular since the seminal works of Beck and Pauker (1983) and Sonnenberg and Beck (1993). When there is a decision to be made at each epoch of the stochastic process, the decision maker's problem is known as a Markov decision process. MDP's (Puterman, 2005) are a well known and applied framework within operations research. This framework is used to model dynamic and stochastic decision problems where rewards are earned based on state-action pairs, and states evolve stochastically according to state-action pairs with the Markov property. Based on the Bellman Principle of Optimality, MDP's can be solved by well known algorithms such as value iteration and policy iteration. Applications of MDP's in medical decision making include Denton et al. (2009); Shechter et al. (2008). The assumption of Markovian transitions is at times a strong one. The Markovian assumption can always be satisfied by enlarging the state space of the model, but at a cost of increased computational burden. Strategies for dealing with large state spaces include decomposition (Hazen, 2011). Testing the Marko-

vian assumption can be accomplished by using data to compare the fit of models with varying dependence assumptions (Welton and Ades, 2005).

Another strong assumption that the majority of the literature makes is that healthcare consumers are EU maximizers. This assumption is often not a plausible description of reality when decision makers face complex decision problems in which probabilities and payoffs are unclear and complicated problems themselves (e.g. healthcare decisions). In response to this criticism, various non-EU models of decision making have arisen (Machina, 2004). Among these alternatives, learning rules are particularly intuitive and describe how choices evolve dynamically in individual decision frameworks and games. Flexible rules can serve as paradigms for various learning protocols such as choice reinforcement models (Roth and Erev, 1995) and belief based models (e.g. fictitious play (Brown, 1951)). Modeling consumer behavior through a learning rule provides a more reasonable description of healthcare consumer decision making, and also allows investigation of knowledge-based incentives.

Health education has been recognized as an important component of modifying consumer behavior. See Glanz et al. (2002) for a comprehensive theoretical treatment including theories of behavioral change for individuals and communities, and putting theory into practice. Of the models presented the theories of planned action and planned behavior (Montãno and Kasprzyk, 2002) are most similar to the models used in this section. Operationalization of these theories is based in attitude measurement, or attractions, through expectations concerning actions, very similar to learning rules. Maibach et al. (2002) discuss social marketing including mass media campaigns and its role as an incentive for producing behavioral change. The problem of a policy maker providing knowledge could be posed at several levels. At the individual level, this decision would be similar to the decision to provide counseling services. Community based educational programs and interventions have been considered a promising level of granularity for inducing behavioral

change (Merzel and D’Afflitti, 2003). At the state or national level, knowledge provision is roughly equivalent to mass media campaigns, and this is most similar to the level of knowledge provision addressed in this section. Even beyond choosing what level to place a knowledge provision policy, other distinctions are required including how to communicate the knowledge. Strategies include motivation, instruction, fear-based response, and others. While important, these issues are beyond the scope of this dissertation, and are discussed in more in the health education literature (Randolph and Viswanath, 2004; Rimer et al., 2001). The next subsection introduces the model used to design knowledge provision policies.

4.2 Model

The model presented here is of a healthcare consumer who chooses a behavior each period from a finite set, $a_t \in \mathcal{A}$. Consumer behavior in each period is stochastically dependent on a state of attractions to each behavior $s_t = (s_t^1, s_t^2, \dots, s_t^{|\mathcal{A}|}) \in \mathcal{S} = \mathbb{R}^{|\mathcal{A}|}$. Based on current attractions, the probability of choosing a given behavior is computed through a stochastic choice rule $f : \mathcal{S} \rightarrow \Delta(\mathcal{A})$. The probability of choosing action a is computed by the logit rule,

$$p(a|s_t) = \frac{e^{\lambda \cdot s_t^a}}{\sum_{k \in \mathcal{A}} e^{\lambda \cdot s_t^k}}.$$

In this rule λ represents the ability of the consumer to best respond to current attractions. The logit rule chooses more attractive behaviors with higher probability, but chooses less attractive behaviors with positive probability as well. If the attractions are interpreted as expected utilities, this decision making model relaxes the assumption from the EU framework that consumers always choose the optimal decision. In this case, as $\lambda \rightarrow \infty$, the logit rule approximates EU maximization, whereas $\lambda = 0$ implies uniform randomiza-

tion between behaviors. The logit rule has shown to compare favorably to other popular stochastic choice rules such as the power rule (Camerer and Ho, 1998). This stochastic choice model is not meant to be normative, instead the uncertain behavior is meant to capture heterogeneity in the population represented by a mean consumer, as well as random variations in behavior at the individual level due to factors such as random shocks to preferences or cognitive load.

The consumer's behavior impacts current payoffs, future attractions, and future health state. The consumer's payoffs in each period are comprised of two components, a health-based cost component, and a deterministic cost component,

$$\pi(a, \omega) = h(\omega) + c(a)$$

The health-based cost component, h , is a random variable reflecting costs from a realized illness state ω in a given period. These costs may include direct and indirect costs as well as costs from disutility. The illness state ω belongs to a finite set of possible states Ω . The probability distribution of these states in any period is governed by the consumer's current behavior-health state pair. For example, if $\Omega = \{\text{healthy, heart attack}\}$, a consumer who is smoking and in a poor health state would have a higher probability of experiencing a heart attack than a healthy non-smoker. The health-based component gives an unhealthy consumer a worse expected payoff than a consumer in healthier circumstances. The deterministic component simply reflects the direct cost of choosing the behavior (e.g. the cost of exercising). Based on the distribution of illness states conditional on the behavior-health state pair, the consumer's expected payoff of choosing behavior a in health state x is computed by

$$\Pi(a, x) = \mathbb{E}[\pi] = c(a) + \sum_{\omega} \rho(\omega|a, x)h(\omega).$$

The consumer's future attractions are determined by current attractions, behavior, and payoffs, and are computed through a learning rule, $L : \mathcal{S} \times \mathbb{R} \times \mathcal{A} \rightarrow \mathcal{S}$. The probabilities of future attractions given current attractions, health state, and behavior are given by

$$P(s_{t+1}|s_t, x_t, a_t) = \sum_{\omega(s_t, a_t, s_{t+1})} \rho(\omega|x_t, a_t),$$

where $\omega(s_t, a_t, s_{t+1}) = \{\omega : L(s_t, a_t, \pi(a_t, \omega)) = s_{t+1}\}$. The consumer's health state is assumed to evolve according to a stochastic process with Markovian transition probabilities, determined by health state and behavior in the previous stage,

$$P(x_{t+1}|x_t, a_t, x_{t-1}, a_{t-1}, \dots, x_1, a_1) = P(x_{t+1}|x_t, a_t)$$

The consumer's full state, consisting of attractions and health state, is then driven by two conditionally independent Markovian processes,

$$\begin{aligned} P(s_{t+1}, x_{t+1}|s_t, x_t) &= \sum_{a \in \mathcal{A}} p(a|s_t) \cdot P(s_{t+1}, x_{t+1}|s_t, x_t, a) = \\ &= \sum_{a \in \mathcal{A}} p(a|s_t) \cdot P(s_{t+1}|s_t, x_t, a) \cdot P(x_{t+1}|x_t, a). \end{aligned}$$

The assumption of conditional independence is a simplifying assumption, but is motivated by the sources of uncertainty driving the two components of the consumer's state. The uncertainty in attractions stems from the uncertainty over payoffs through the occurrence of random acute events. These acute events can be caused by stressful incidents and the short-term burden of behaviors. Whereas the uncertainty in health state stems from the uncertainty in how the consumer's health evolves in response to behaviors and other factors. This latter uncertainty could be diminished by obtaining better information (e.g. by genetic

testing) about the characteristics of the consumer. The following subsection discusses how this process can be controlled by a policy maker.

4.3 MDP Control for a Policy Maker

The previous subsection introduced a model in which a healthcare consumer's attractions and health state evolve according to stochastic processes governed by behavior and payoffs. A policy maker or payer with preferences over sample paths would be interested in controlling the consumer's process. This framework has been motivated by the experience of a single consumer, however it is unlikely that a policy maker could design incentives at the individual level for healthcare behaviors. Rather, the group of consumers the policy maker is concerned with will be treated as an aggregate consumer whose attractions and health evolve as described. Of course if a policy maker could reliably differentiate knowledge between different groups of consumers they could design different knowledge strategies for each group. Whether this objective is to maximize social welfare, or to minimize costs, the sample path taken greatly impacts the policy maker's objective value. Since the consumer's decisions are autonomous, a key complication in healthcare and other distributed systems, the policy maker cannot directly choose the consumer's behavior, mechanisms, regulation, or incentives must be designed and used to control consumer behavior.

One means of control would be through the consumer's payoffs. Modifying payoffs could be accomplished in a variety of ways: setting prices for treatment of acute events, setting insurance parameters such as coinsurance rates, or taxing and subsidizing particular behaviors. These financial incentives could all be designed to control the consumer's behavior with the policy maker's objective in mind. The new approach introduced here is to control behavior by adjusting parameters of the consumer's learning rule. If the

learning rule is a function of parameters under the policy maker's discretion, then setting these parameters can be another way of modifying consumer attractions and behavior. The particular incentive under investigation here is providing knowledge to consumers. Existing examples of providing knowledge to consumers are commercials on television and highway billboards encouraging consumers to exercise more, eat healthier, and quit smoking. These advertisements provide knowledge and educate consumers of the possible consequences of various behaviors.


In order to model how knowledge affects the consumer's attractions, the experience-weighted attraction (EWA) learning rule will be used (Camerer and Ho, 1998). EWA has been shown to be empirically flexible and economically valuable in a variety of settings, and generalizes other popular learning mechanisms. In this model, the consumer's attractions to each behavior, $\{s_t^j\}_{j \in \mathcal{A}}$, are updated according to the following rule,

$$s_t^j = \frac{\phi \cdot N_{t-1} \cdot s_{t-1}^j + [\delta^j + (1 - \delta^j) \mathbb{1}(a_{t-1} = j)] \cdot \pi(j, \omega_{t-1})}{N_t},$$

where N_t is a scalar representing the consumer's experience at time t , governed by

$$N_t = \rho \cdot N_{t-1} + 1, \quad t \geq 1.$$

There are three parameters in the EWA learning model: ϕ , ρ , and δ . The first parameter, ϕ conveys how experience translates into attractions, and the second, ρ characterizes how experience accrues over time. The remaining parameter, $\delta \in [0, 1]^{|\mathcal{A}|}$, characterizes how well the consumer gets payoff signals about behaviors not chosen, with possibly distinct values for every behavior. The attraction vector gets updated with the full impact of the behavior chosen, however for behaviors not chosen how strongly the attractions are updated depends on δ . These parameters are modeled as a decision variables of the policy

Figure 4.1. Policy Maker's Markov Decision Process

maker with the interpretation that the policy maker chooses how to provide knowledge to the consumer regarding various behaviors.

Using this framework for how knowledge can be used to control the consumer's stochastic process, a policy maker's incentive problem can be formulated as a MDP. Figure 4.1 shows graphically the process in which the consumer's health and attraction states in the rounded rectangle are stochastically modified and produce rewards based on the policy maker's decisions. The policy maker's reward each period could be total costs of health-care, consumer utility, or some other metric depending on the policy maker's objective. For generality write the reward as $r(s_t, x_t, \delta_t)$, making explicit the cost of the policy maker's

decision, providing knowledge δ_t . Taking an infinite horizon with discount factor γ , the policy maker's objective is

$$\max_{\delta_t(x_t, s_t)} \mathbb{E} \left[\sum_{t=1}^{\infty} \gamma^t r(s_t, x_t, \delta_t) \right].$$

If the state space of the Markov process, (S, X) , is observable and finite, solution procedures for the policy maker's MDP are well known. In general, the policy maker may not directly observe either of these variables. However the policy maker will likely have data, or at least a sample of data with which to estimate the both the attraction and health states. For example, if the policy maker observes the empirical distribution of behaviors chosen in each period, the empirical probabilities, $\{p(a = j)\}_{j \in \mathcal{A}}$, can be used to form an estimate of the attractions s_t . Using the logit choice rule the policy maker can solve a system of $|\mathcal{A}| - 1$ linearly independent equations of the form

$$\frac{1}{\lambda} \ln \left(\frac{p_j}{p_k} \right) = s_j - s_k$$

to estimate the attractions. Since there are too many degrees of freedom, the attraction corresponding to the behavior with the highest frequency could be set to 0. This normalization is reasonable since the behavior payoffs are costs (direct and health-related), and since the logit rule is indifferent to a constant shift in all of the attractions

$$\frac{e^{\lambda(s_i^a - q)}}{\sum_{k \in \mathcal{A}} e^{\lambda(s_i^k - q)}} = \frac{e^{\lambda s_i^a} e^{-\lambda q}}{\sum_{k \in \mathcal{A}} e^{\lambda s_i^k} e^{-\lambda q}} = \frac{e^{-\lambda q} e^{\lambda s_i^a}}{e^{-\lambda q} \sum_{k \in \mathcal{A}} e^{\lambda(s_i^k - q)}} = \frac{e^{\lambda s_i^a}}{\sum_{k \in \mathcal{A}} e^{\lambda s_i^k}}.$$

Data on health states and behaviors are available in existing data sets such as the National Health and Nutrition Examination Survey (NHANES). In order to meet the criteria of finiteness for standard iterative algorithms such as value iteration, the attraction and health

state spaces will require discretization. The following section specifies the model and estimates parameters in order to solve for knowledge provision policies to control costs from coronary heart disease.

As policy makers, healthcare payers, and engineers continue to search for cost saving solutions, incentives will become an increasingly important research domain. This section has presented a dynamic framework for designing incentives for non-EU maximizing consumers. This framework is not limited in its usefulness to healthcare incentives, but could be used in a variety of applications where indirect control of an agent's state is of interest. In particular, this framework can be used to design knowledge-based incentives, in addition to more standard monetary incentives through contracts, taxes, and subsidies.

5. KNOWLEDGE INCENTIVES FOR CORONARY HEART DISEASE

This section uses the Markov framework introduced in the previous section to design optimal knowledge incentives to control costs from coronary heart disease (CHD). Constructing incentives to control CHD is desirable because of its prevalence, cost, and preventability. After specifying the model and estimating parameters to control behaviors affecting coronary heart disease, iterative algorithms can be used to solve for the optimal knowledge incentive policies. Solving the MDP yields optimal strategies for a policy maker to give incentives by providing knowledge about consumer behaviors. Data from the Atherosclerosis Risk In Communities (ARIC) study conducted by the National Heart, Lung, and Blood Institute (NHLBI) is used to estimate parameters for the model.

This section is organized as follows. Subsection 5.1 reviews literature related to educational interventions for cardiovascular diseases. Subsection 5.2 develops the model specification of behaviors and health states pertinent for coronary heart disease. Subsection 5.3 presents the estimation of parameters, followed by a discussion of model validation in subsection 5.4. This subsection on model validation is comprised of three components. First, the full model is scaled down to knowledge provision for smoking behavior. This allows solution with many more parameter settings to check the results of the model. The validation focuses on the model results with respect to the experience of the consumer. The experience parameter modifying how willing consumers are to change behavior based on new payoff information, possibly including knowledge. After validating the results at the boundaries and matching results to the observed behavior found in the empirical literature, sensitivity analysis provides insights into how the results vary with changing cost parameters. The parameters varied are the direct costs of smoking, the consumer's cost of coronary heart disease, and the cost of providing knowledge. The interest in these parameters is further discussed later in the section. By validating the model's performance when

restricted to smoking behaviors, the full model can be partially validated without excessive computational burden. The results on optimal knowledge provision policies are presented in subsection 5.5, which also discusses possible extensions of the model.

5.1 Literature Review

Coronary heart disease is an important disease for policy makers to consider because of its prevalence, cost, and preventability through modifiable risk factors. CHD affected 17.6 million Americans in 2010, for a total direct and indirect cost of \$177.1 billion (Lloyd-Jones et al., 2010). In addition to financial costs, the costs of human life are severe as well, with 1 of every 6 deaths in the U.S. caused by CHD in 2006 (Lloyd-Jones et al., 2010). Despite the costs, consumers still engage in behaviors which are major risk factors for CHD. Prevalence of smoking (21%), inadequate physical activity (35%), and high saturated fat diets (46-67%, depending on sex and ethnic group), indicate the potential for alleviating some of the cost of CHD. Due to its importance, several studies have implemented knowledge-based interventions and considered best practices for designing the interventions, cost-effectiveness, and their effect on prevalence.

Smoking is the behavior which has received the most attention from mass-media, counter-advertising, and other knowledge-based interventions. Controlled experiments as well as implemented community and population level interventions have shown the ability of smoking cessation knowledge and promotion to change behavior. Results from various geographic regions including California (Farquhar et al., 1990), Texas (McAlister et al., 1992), Vermont and New Hampshire (Secker-Walker et al., 2000), Florida (Bauer et al., 2000), British Columbia (Gagne, 2007), and Finland (Puska et al., 1985) have all shown that providing knowledge about smoking can reduce prevalence or increase the cessation rate. The amount of behavior change and the subset of the population significantly af-

fects varies throughout these studies. These variations can be explained by differences in funding and penetration, secular trends and other concurrent media and programs, and the targeting and duration of the intervention (Davis et al., 2008). Excellent reviews on the effectiveness of mass-media and other knowledge-based incentives for smoking can be found in Chaloupka and Warner (2000), Davis et al. (2008), and Bala et al. (2008).

In addition to smoking, diet and exercise are two other behaviors which have been addressed by community or mass-media health education campaigns. Heimendinger et al. (2007); Toobert et al. (2005); Wendel-Vos et al. (2009) all report positive behavioral changes in dietary consumption and physical activity as a product of health education programs. Stern et al. (1976) also report significant reductions in saturated fat and cholesterol consumption as a result of a two-year bilingual mass-media campaign. Other studies report significant changes in cholesterol and blood pressure levels of consumers (Diehl, 1998; Farquhar et al., 1990; Puska et al., 1985) without explicitly tracking diet and exercise behavior. Lin et al. (2010) provide an extensive review of the literature and conclude that behavior counseling is associated to positive (but small) changes in diet and physical activity.

Recommendations for how to structure educational programs to reach their full potential include the need for programs to be comprehensive (Randolph and Viswanath, 2004), hard-hitting (World Health Organization, 2008), targeted to specific subgroups (Winkleby et al., 1994), and community/socially oriented (Schar and Gutierrez, 2001). General discussions on improving general, cardiovascular, and smoking-specific health education programs can be found in Rimer et al. (2001), Parker and Assaf (2005) and Davis et al. (2008)

The challenge in designing incentives to control CHD costs is that both the effectiveness of interventions and the potential savings from behavioral change vary within subgroups of any population. Ebrahim et al. (2011) review the literature on multiple risk factor interventions and conclude that the interventions are effective for high risk groups,

although not so for the general population. Research on the normative side of health education has focused on the cost-effectiveness of interventions. Cost effectiveness measures for anti-smoking campaigns can vary based on the parameters chosen. Tengs et al. (2001) use a dynamic model and vary cost, effectiveness, and recidivism to find that the cost per quality-adjusted life-year (QALY) of anti-tobacco education varies from \$4,900 to \$340,000. Ronckers et al. (2005) review and standardize many cost effectiveness studies and find that the cost per year of life saved ranges from \$1,000 to \$15,000 depending on the target and structure of the intervention. Tosteson et al. (1997) study educational approaches to reduce cholesterol levels and find that costs per QALY vary from zero (cost saving) to \$38,000 per QALY. These studies generally conclude that educational interventions should be implemented on the basis of their cost effectiveness (a typical cost threshold per QALY is \$50,000). The next subsection begins specifying the model for designing knowledge-based incentives.

5.2 Model Specification

5.2.1 Consumer Health and Behavior

The consumer's health state is modeled using sub-states of total cholesterol (TC), high-density lipoprotein (HDL) cholesterol, and systolic blood pressure (BP), three major risk factors for CHD. While still a significant abstraction from reality, this description of a consumer's health state extends the detail of many existing models in medical decision making which primarily model disease states Briggs et al. (2006). In order to retain a finite state space, each of these sub-health states are modeled as categorical variables. The categorizations used, seen in Table 5.1 are adapted from Wilson et al. (1998) and are similar to the definitions and guidelines from the Fifth Joint National Committee on

Hypertension and the National Cholesterol Education Program. This categorization results in $5 \times 4 \times 4 = 80$ possible health states for the consumer.

Table 5.1
Categorization of Health States

Category	BP (mm Hg)	HDL (mg/dL)	TC (mg/dL)
1	<130	≥ 60	<160
2	130-139	45-59	160-199
3	140-159	35-44	200-239
4	≥ 160	<35	240-279
5	-	-	≥ 280

Consumer behavior is also modeled by three sub-behaviors: choosing a low saturated fat diet, smoking, and getting a recommended level of exercise. Each sub-behavior is modeled as a binary decision, leading to a total of 8 possible combined behaviors by the consumer. The distinction of a low or high saturated fat diet is based on the percentage of total calories consumed coming from saturated fat. Using American Heart Association (AHA) guidelines, the recommended percentage is $\leq 10\%$. The AHA guideline for physical activity is getting at least 150 minutes of moderate physical activity each week. Each of the binary decisions is defined such that 0 represents the unhealthy choice, and 1 represents the healthy choice. Using the logit stochastic choice rule on each sub-behavior, and by normalizing the attraction $s^{a=0} = 0$ for each unhealthy sub-behavior, the consumer's attractions state can be reduced to three real numbers representing the attractions to $s^{EXR=1}$, $s^{FAT=1}$, $s^{SMK=1}$.

In order to create a finite state space, the range of each attraction is bounded to $[-1.5, 1.5]$ and approximated within this range by a 0.3-fine grid. As previously mentioned, the attractions are invariant to the addition of a constant, so the location of this range has no

effect other than the interpretation of costs previously discussed. The range of attractions may seem small, however even an observed probability ratio of 90 can be approximated. That is, suppose the most frequently observed behavior (say behavior i) has an empirical probability of 0.9 and the least observed behavior (j) has an empirical probability of 0.01. Using the logit rule with a λ value of 3, the difference in attractions would be $s_i - s_j = \ln(0.9/0.01)/3 \approx 1.5$. Therefore, neither the range nor the location of the discretization loses much generality. Increasing the fineness of the approximation grid would increase the ability to accurately model behavioral change on behalf of consumers, but at the cost of increased computational burden.

In order to model consumer behavior from attractions to each sub-behavior, one simplifying assumption would be that consumers decide whether or not to smoke, to exercise, and to consume a low fat diet independently. By comparison, this assumption seems at least as plausible as the assumption that consumers are making decisions about all three sub-behaviors simultaneously. Upon investigation, the assumption of behavioral independence failed to hold, however the behaviors of diet and smoking were found to be independent once conditioned on exercise behavior. An additive model of dependence and conditional independence was developed and found to fit the data much better than the assumption of independence. The details of this investigation and modeling are reported in the Appendix. The additive model implemented allows the probabilities of each sub-behavior to be computed using only knowledge of the consumer's attraction to each sub-behavior without assuming independence. This is desirable in order to save on the size of the state space for computation.

5.2.2 Consequences of CHD

Modeling acute events related to CHD is a challenging task. By AHA definitions, CHD consists of ICD-9 disease classifications 410.x-414.x, which is comprised of a number of conditions including the major categories of acute myocardial infarction (AMI), angina pectoris, and ischemic heart disease. Accurate modeling, prediction and cost estimation of these and other more detailed acute events is beyond the scope of this dissertation. Instead a single acute event is modeled which represents the incidence of acute CHD.

5.2.3 Policy Maker Decisions and Rewards

The general framework using the EWA learning rule with knowledge parameters characterizing policy decisions allows for distinct δ values for each behavior the consumer might take. The policy maker's decision of whether or not to provide knowledge about the expected payoffs from each sub-behavior, is simplified to $\delta = (\delta^{EXR}, \delta^{FAT}, \delta^{SMK}) \in \{0, 1\}^3$. The interpretation being that the policy maker will provide information about the benefits and costs of smoking, diet, and exercise separately. Since the consumer sub-behaviors are binary, choosing to provide knowledge conveys the expected payoffs of the sub-behavior decision. That is, setting $\delta^{SMK} = 1$ will update the consumer's attraction to smoke and to not smoke. This modeling feature could be generalized to allow for a policy maker to decide for example how much funding to provide for a mass media campaign. Relaxing this assumption would increase the already considerable computational burden of the model, as well as take careful modeling of the varying effect of funding on mass media campaigns, and is left for future research.

The policy maker's decision determines whether or not the consumer receives payoff knowledge about behaviors not chosen. Since healthcare consumers are likely unable

to internalize their payoffs (including health related payoffs) for actions not chosen, one alternative to the basic EWA setup is to set the payoff knowledge about actions not taken equal to the expected payoff for the state-behavior pair, $\Pi(a, x)$. This is problematic for policy makers in reality and in this CHD model since exercise and diet have no short-term impact on the risk of CHD. Therefore, the payoff knowledge the policy maker provides is set to reflect costs ten years in the future given today's health state and action chosen. This time frame is chosen since shorter horizons may not yield significant differences in health costs while longer horizons may lead to excessive discounting when interpreted by consumers. One option would be to send information about the total expected costs of the chosen behavior,

$$\Gamma(a, x) = \Pi(a, \mathbb{E}_{a,x}[x^{(10)}]),$$

where $\mathbb{E}_{a,x}[x^{(10)}]$ is the expected health state in 10 years given today's state and repeated behavior. However, this function provides knowledge about direct costs which are relatively well known to consumers, and also may not reflect favorably on healthy behaviors (such as the cost of exercising). Furthermore, this choice of Γ provides different knowledge to consumers based on their current behavior. Based on these concerns a more realistic modeling choice is to set

$$\Gamma(a, x) = [p(w = 1|a, \mathbb{E}_{a,x}x^{(10)}) - p(w = 1|a', \mathbb{E}_{a',x}x^{(10)})]h(w),$$

where a' represents the complement of a in whatever dimension knowledge is provided. For example, if $a = (a_{EXR}, a_{FAT}, a_{SMK}) = (1, 1, 0)$ and knowledge is provided regarding exercise, then $a' = (0, 1, 0)$. This Γ function provides knowledge about health-based consequences for the behavior chosen relative to the opposing behavior. This function also has the appealing feature that $\Gamma(a, x) = -\Gamma(a', x)$, so that all consumers hear the same mes-

sage. Consumers with healthy behavior will receive positive knowledge, while consumers with unhealthy behavior will receive negative knowledge. Providing this payoff information allows the policy maker to deter unhealthy behaviors related to diet and exercise despite the lack of short term consequences.

The policy maker's cost is modeled as proportional to the number of sub-behaviors for which knowledge is provided,

$$C(\delta) = k(\delta_{PM}^{EXR} + \delta_{PM}^{FAT} + \delta_{PM}^{SMK}).$$

The constant k , reflecting the cost per consumer of providing knowledge, will vary depending on the size and distribution of the population. Since national averages have been used as inputs for parameters, the \$130 million spent on the 2010 U.S. Census campaign for a cost of roughly \$.43 per individual is used as a baseline estimate of the cost of providing knowledge to an individual. The sensitivity of the results to this and other parameters is investigated later. The policy maker's objective is assumed to be minimizing the total expected costs of both CHD consequences and providing knowledge. Thus rewards in each period are given by

$$r(s_t, x_t, \delta_t) = \mathbb{E}_\omega[h(\omega) + C(\delta_t)] = C(\delta_t) + \sum_{a \in \mathcal{A}} \sum_{\omega \in \Omega} h(\omega) p(\omega|a, x_t) p(a|s_t).$$

5.3 Parameter Estimation

5.3.1 Health Transition Probabilities

The state transition probabilities are an important part of the model. By modeling the health state and behavior using three sub-components each, the accuracy and descriptiveness of the model has improved, but also have significantly increased the number of

transition probabilities to estimate. The currently specified model consists of 80 health states and 8 behaviors for a total of 51,200 possible state-state-action tuples. In order to estimate these transition probabilities, the ARIC data which contains cohort data on individuals with detailed information about behavior and health status over time is used. The data provides 20,604 state-state-action transitions with which to estimate the transition probabilities. A flexible parametric model is used with maximum likelihood estimation to compute the transition probabilities, without making unnecessarily strong assumptions such as probabilistic independence of sub-transitions. Details of the parametric model and estimation procedure can be found in the Appendix.

5.3.2 Behavioral and Health Costs

The cost of smoking is estimated by the consumption of the average smoker, a pack a day, for a cost of \$1,825 per year at \$5 per pack. The cost of getting the recommended level of exercise is not as clear, since most individuals get some exercise in their daily routine. The additional physical activity needed for an individual currently not meeting the AHA guidelines is estimated from the ARIC data and multiplied by the median wage rate to estimate the incremental cost of getting enough exercise per week. This cost comes to \$1,888 per year for the average under-exercising individual who needs to increase activity by 136.6 minutes per week at a cost of \$15.95 per hour, the median U.S. wage rate (Bureau of Labors Statistics, 2009). Accurately estimating the cost of eating a diet with less than 10% of calories coming from saturated fat is more difficult still. Diets vary greatly between regions of the country and socioeconomic groups. Food groups such as ground beef seem to indicate that a diet with less fat would be more expensive, since lean meat is more expensive than fatty meat. Other foods, milk for example suggest otherwise as skim milk is not more expensive than 2% or whole milk. Carlson et al. (2007) study whether or not it is

feasible to maintain a healthy and nutritious diet on a budget equal to that of the maximum food stamp allotment. Carlson and colleagues use prices paid by low-income people for food, consider the cost and nutrients of foods as-consumed (including preparation and cooking), and take an objective of minimizing deviations from current dietary habits to avoid bizarre recommendations. The conclusions of their study find that is feasible for each of 15 age-gender groups to achieve the 2005 Dietary Guidelines for Americans, including a diet with less than 10% of calories from saturated fat, on a minimal budget. Given that consumers can choose a nutritious low fat diet on a minimal budget, the direct cost of choosing a diet with less than 10% of calories from saturated fat is set to zero. The fact that many consumers prefer fat in their diet will be captured through the attractions.

The cost of this acute event is set to \$10,000, approximately the average cost of CHD for Americans in 2010, and a lower bound on the reported average costs of more specific diagnoses (\$14,009 for AMI, \$12,977 for coronary atherosclerosis, and \$10,630 for other ischemic heart disease Lloyd-Jones et al. (2010)). From the policy maker's perspective, this simplification makes little sacrifice, since the population of consumers will have a variety of acute CHD symptoms, and this average cost of treatment will represent the total consequences across consumers. Also, from the consumer's perspective, it seems plausible that any form of acute CHD would be enough to modify behavior, considering the relative cost of CHD to the behavioral costs and preferences. This interpretation can be seen in the model since whether the patient experiences a AMI, coronary atherosclerosis, or some other acute CHD diagnosis, the EWA model predicts the attraction to chosen behaviors becoming very small.

Each period represents one year. In order to calculate the 1-year probability of CHD from each health state-behavior pair, the prediction model of D'Agostino et al. (2000) is used. These authors incorporate the health factors considered in this model, and use data from the the Framingham Heart Study (Truett et al. (1967), www.framinghamheartstudy).

org) to fit a parametric model predicting the risk of CHD within a 1-4 year span based on systolic blood pressure, HDL cholesterol, and total cholesterol along with several other factors. Variables in the D'Agostino et al. model that are not captured in the state space or behavior in this model are entered at the population average level.

5.3.3 Learning Rule Parameters

Although it has been estimated for a variety of decision situations, EWA has not been estimated within a healthcare decision context. The parameter ϕ is set to 1.0, which is consistent with empirical estimates from other decision paradigms (Camerer and Ho, 1999). Setting $\phi < 1$ typically prevents attractions from growing without bound, but since the attractions in the model described are renormalized each period, this issue is not a concern. Since the consumer represents an aggregate consumer of the policy maker's population, with older individuals leaving the population and young individuals entering, The experience of the consumer is modeled as constant over time. Any level of experience can be considered while maintaining the assumption of constant population experience by setting $\rho = (N - 1)/N$. The value of experience, N determines the weighting of new information (payoffs and knowledge) relative to historical attractions in computing new attractions. New information is given a weight of $1/N$. The range of values of experience used in computations reflect the literature on how healthcare consumers make decisions, which reports that consumers 'give priority to their personal experience' (Moser et al., 2010), and that many consumers will not use information in decision making (Hibbard et al., 1997). Since the experience parameter controls the behavior change of consumers, it will be a focus of the sensitivity analysis to validate the model.


In order to choose an estimate for λ , evidence found in Camerer et al. (2002) is used. Camerer et al. estimate EWA parameters from a pooled data set of a variety of

games including patent race, continental divide, median action, p-beauty contest, traveler's dilemma, and pot games. Their pooled estimate of λ is 2.95. Recall that $\lambda \in [0, \infty)$ represents the ability or propensity of consumers to select more attractive actions. For example, if a consumer has an attraction to smoking of 0 ($s_{SMK=0} = 0$), and an attraction to not smoking of 1 ($s_{SMK=1} = 1$), then $\lambda = \{0, 1, 3.5, 10\} \Rightarrow p(SMK = 1) = \{0.5, 0.731, 0.950, 0.993, 0.999\}$. Two factors present in healthcare consumers that would tend to increase λ from 0 (uniform randomization) are the lack of tendencies to explore and experience in the decision environment. In light of these factors, Camerer et al.'s pooled estimate appears reasonable.

5.4 Model Validation

In order to validate the model for prescribing knowledge provision policies, the model is scaled down to consider only smoking behavior. This strategy is taken since the full model is computationally expensive to run and this smaller version allows a more complete validation. and gives some credibility to the larger model since the model is quite similar for each sub-behavior. Figure 5.1 shows the structure of the optimal policy as attraction and health states and experience vary. The black cells show the attraction and health state combinations at each experience level for which the policy maker's cost minimizing decision is to provide knowledge about the benefits of not smoking. The highlighted window shows the policy information conveyed at each attraction-experience combination. Recall that the health states are ordered from most to least healthy, so the threshold policy seen in the highlight pair, when $N = 50$, and $s_{SMK} = 2$, shows a policy in which it is optimal to provide knowledge to the least healthy consumers. At the boundaries, as experience goes to zero or becomes large, the model exhibits the expected behavior. That is, as experience decreases, consumer place increasing weight on new information through payoffs


Figure 5.1. Optimal Policy for Smoking Knowledge Provision. Highlighted at $s = 2$, Experience = 50



and knowledge. Thus as experience goes to zero, consumers easily change their behavior when presented with knowledge about the expected health benefits of not smoking. In this case, providing knowledge is optimal at all attraction and health states. Analogously, as experience grows, consumers place increasing weight on their historical attractions and less weight on new information. As experience becomes large enough, no behavior change arises from information provided, and costly knowledge provision is clearly sub-optimal. This prediction is realized seen in Figure 5.1 as experience reaches 170.

Sensitivity to the cost of smoking, the cost of providing knowledge, and the consumer's cost of an acute CHD event were also tested. Figure 5.2 shows how the states which are cost effective to provide knowledge in decreases as the cost of providing knowledge increases. The black region shows which states providing knowledge is the optimal decision if providing knowledge costs \$5.00 per individual per year. The dark gray region shows which additional states have cost saving knowledge provision if the cost of providing knowledge is \$2.00. The light gray region shows the additional states which make up Figure 5.1, where the cost of providing knowledge is \$0.43. Sensitivity was performed with respect to the consumer's cost of an acute CHD event to simulate consumer's insurance from financial risk. The policy maker's cost of an acute CHD event was left at the full cost of \$10,000, while the consumer's cost was set varied between \$10,000, \$5,000, and \$2,000. The optimal policies were found to be relatively constant in the consumer's cost of CHD on the range of experiences where providing knowledge is interesting. Finally, one result worth explaining is the feature that as experience rises, the first consumers to leave the optimal policy are those who are most likely to smoke. This result seems counter intuitive as those consumers would be the ones most in need of knowledge. This artifact of the optimal policy exists because of the fact that consumers who smoke experience the direct costs of smoking, and negatively impact their attraction to smoking. The consumer's cost of smoking was varied to investigate the possibility that consumer's are addicted or

Figure 5.2. Optimal Smoking Knowledge Provision Policies by Cost of Knowledge. $k = \{0.43 \text{ (Light gray)}, 2 \text{ (Dark gray)}, 5 \text{ (Black)}\}$



have internalized the cost of smoking. The results show that as the cost of smoking decreases, more knowledge should be provided to these most addicted smokers. This result agrees with the intuition that if smoking is very cheap, a policy maker cannot rely on direct costs alone to deter consumers from smoking. While the model exhibits the behavior which agrees with intuition and exhibits appropriate boundary behavior, in order to fully validate the model, experiments would need to be designed and performed for the particular population of consumers at hand. The design of such experiments is discussed in the following.

5.4.1 Ideal Validation: Design of Experiments

In order for a policy maker to have full confidence in the results of the model, experiments would need to be designed to estimate parameters more accurately. Parameters of particular interest are those of the learning rule (ϕ, λ, N) which model attractions and behavioral change in consumers, and the costs of each type of behavior as perceived by consumers. Prior to designing experiments, the model presented in this section could give a policy maker a sense of the value of information of a certain parameter. For example, consider the optimal policies from Figure 5.1. If a policy maker were fairly certain that the experience of the population under consideration had an experience between 60 and 100, the value of experimentation to determine the precise experience of the population would be of little value. This can be seen since the optimal policy governing what states should be provided smoking knowledge is quite constant over this range of experiences. On the other hand, a prior belief on the range of 100 to 150 could create a significant value of experimentation.

Estimation of learning rule parameters can be accomplished through maximum likelihood or minimum deviation procedures Cabrales and Garcia Fontes (2000); Camerer et al. (2002). A group of participants could be randomly assigned to both control and treatment groups where the treatment group receives knowledge about the expected consequences of health behaviors. This treatment could be specifically administered to treatment group participants through electronic and paper media. By taking repeated measurements of attractions to health behaviors and reported behaviors, the parameters could be estimated for the cohort of experiment participants. Attractions to behaviors should be measured in both absolute (e.g. through a willingness-to-pay assessment) and categorical scales (e.g. $\{0, 1, \dots, 10\}$). In order to accurately estimate these parameters, care must be taken to ensure that the treatment and control groups are representative of the population which is be-

ing considered for the intervention. This can be aided by stratifying the experiment groups for important possible confounding variables such as age, education, and income. Even if the groups are representative, unobserved heterogeneity can lead to bias in estimation of learning parameters Cabrales and Garcia Fontes (2000). In light of the heterogeneity of healthcare consumers, this may be a concern in estimating healthcare learning parameters. Cabrales and Garcia Fontes (2000) discuss estimating the distribution of parameters in order to deal with this bias.

In order to support the experimental evidence, quasi-experiments should also be designed. These experiments do not randomly assign participants to treatment groups, but rather the intervention could be administered to a treatment population and a control population. These populations would need to be measured for prevalence of the health behaviors identified both before and after the treatment period in order to observe the change in behavior due to the educational intervention. The treatment and control populations would likely be similar communities, closely matched on a number of factors which could be confounding factors to the treatment effect. These factors would include distribution of age, education, income, initial behavior prevalence, and geography. Geography is a challenge factor to control for because of the trade offs involved with proximity. Too great proximity would tend to produce significant climatological and cultural differences between the populations, while too small of a proximity would be difficult to ensure that the treatment is not having a peripheral effect on the control group. Even after controlling for these factors, other threats to the internal validity of such quasi-experiments include differences population history and predisposition and differences in population dynamics over the experimentation period.

A challenge in measuring the behavior prevalence is that self-reporting through surveys or interviews would be cheapest measurement tool. In the case of reporting the level of exercise or the percent of calories consumed from saturated fat, the experiment participants

may forget or simply not know what the true response should be. Additionally, all three healthy sub-behaviors may tend to be over-reported because of a social desirability bias. Smoking can be tested for through nicotine tests, however exercise and diet are more difficult behaviors to definitively assess. More precise information about model parameters could lead a policy maker or payer to better decisions about when to provide knowledge to consumers. Experiments to estimate parameters which control for potential biases and confounding factors such as those discussed here can provide decision makers with the information needed. The next subsection discusses the results from the full model for providing knowledge regarding exercise, diet, and smoking.

5.5 Results

The full model is solved by a parallelized value iteration algorithm. With such a large model, the results can be challenging to interpret. The results are presented with the parameter settings established earlier. The results for smoking knowledge provision are quite similar to the results presented earlier, with thresholds determining which consumers are sick enough to warrant knowledge about their smoking behavior. There is some interaction between the optimal policy for smoking and the attractions to other behaviors. As consumers become attracted to exercising and low-fat diets, some health states drop out of the optimal policy for smoking.

Very few health and attraction states yield a cost savings from providing knowledge about exercise or low-fat diet behavior. Knowledge about diet should be provided to non-smokers, who are attracted to high-fat diets. Within this set, knowledge is valuable to consumers with low blood pressure and total cholesterol, but high HDL cholesterol. See Figure 5.3, which shows a piece of the optimal policy at $(s_{EXR}, s_{FAT}, s_{SMK}) = (4, 2, 10)$, where each cell entry is read as a tuple $(\delta_{EXR}, \delta_{FAT}, \delta_{SMK})$, and in this table the cells in

which providing knowledge about diets is optimal are highlighted. Examining the transi-

Figure 5.3. Optimal Diet Knowledge at $s = (4, 2, 10)$

		TC				
		1	2	3	4	5
BP=1	HDL=1	0	0	1	1	1
	2	0	0	1	1	1
	3	11	1	1	1	1
	4	11	11	1	1	1

tion probabilities between health states per diet behavior reveals that a low saturated fat diet is most useful in lowering HDL cholesterol states (raising HDL cholesterol), and less so for other health sub-states. This result driven by the data makes sense because saturated fat is not the only major dietary control of blood pressure (sodium) and because a unsaturated fats can raise HDL levels. Therefore the results show that the health-attraction combinations where providing knowledge about a low saturated fat diet are few, and include non-smokers who could improve their diets, but whose health isn't already too poor to help.

The results for exercise knowledge provision are similar to those for diet. See Figure 5.4, which shows the optimal policy at $(s_{EXR}, s_{FAT}, s_{SMK}) = (3, 5, 10)$, and the cells where providing knowledge about exercise is cost saving are highlighted. Again knowledge goes to healthy non-smokers who have room to improve their health state. One curious result here is that exercise knowledge is not provided at health state $(BP, HDL, TC) = (1, 4, 1)$, while it is at both $(1, 3, 1)$ and $(1, 4, 2)$. The flexible approach to estimating transition probabilities makes the exact source of this gap hard to identify. The change in the policy from $(1, 4, 1)$ to $(1, 4, 2)$, where providing knowledge becomes optimal could be explained

Figure 5.4. Optimal Exercise Knowledge at $s = (3, 5, 10)$

		TC				
		1	2	3	4	5
BP=1	HDL=1	100	0	1	1	1
	2	100	0	1	1	1
	3	101	101	1	1	1
	4	1	101	101	101	1

by the increased opportunity to improve health. The change from $(1, 3, 1)$ to $(1, 4, 1)$ is harder to explain. One explanation that is partially supported by the data is if consumers in HDL state 4 are less likely to improve their HDL health. At first these seems unlikely since consumers in HDL state 4 have very low HDL counts, and therefore have significant room to improve. However, the fact that they are so low in HDL, makes it more likely that family history and genetics are contributing to the low HDL count, and that exercise is less likely to aid in this regard.

The results conclude that providing knowledge and diet behavior is rarely cost saving. Although this result would seem to contradict the increasing occurrence of these types of education and knowledge in media today, several factors may explain the difference. First, cost-effectiveness is likely a motivation for existing health education initiatives whereas this study has considered cost savings. A second and related factor is that this study models costs stemming from CHD only. Diet and exercise have less immediate effect on preventing acute CHD events than smoking behavior, thus the benefits of improved behavior are discounted. Poor diet and exercise behavior would likely lead to increased costs and decreased quality of life due to other diseases as well (e.g. diabetes). Finally, the model is somewhat limited by the health state categorizations and the input data. The less structured form of the exercise and diet policies tend to suggest the results may be more

tied to the state and behavior specifications and data. As can be seen from the correlations in the Appendix, the effects of healthy exercise and diet behavior on health transitions are the hardest to identify. Likely the granularity of the state categorization and noise in the data contribute to this complication. Whether the effect is truly noisy, or if the techniques used here can be significantly improved on, the effect is that a policy maker is less certain about the benefits of change consumer behavior in these dimensions. By modeling the effects of behavior on health in more sophisticated ways, and modeling total healthcare costs stemming from the health and attraction states, providing knowledge for these behaviors will become a better investment for a payer.

5.5.1 Future Extensions

Even just considering CHD, this model could be extended in many ways, to incorporate more of the vast information available on CHD. One example would be through differentiating knowledge between groups of the public. D'Agostino et al. (2000) provide gender specific CHD risk models. Presumably the optimal knowledge that should be provided for women would differ from that of men. If a policy maker could reliably choose knowledge paths to differentiate the knowledge presented to men and women, one-size fits all policies could be improved upon. Other extensions could include costs to consumers of acute CHD events above the direct financial costs due to pain and suffering. If this cost or disutility could be reliably estimated it could be incorporated in the total cost of an acute CHD event. This could be an important feature to model if consumers have very generous insurance eliminating the financial costs of acute events. Another interesting extension would be to consider community or herd influences on behavior in the population. This feature could be implemented by extending the learning rule and state space to consider

attractions to behaviors of the individual as well as the community. While interesting this extension would enlarge an already burdensome state space.

6. CONCLUSIONS

While operation practices will continue to be refined over and over, real efficiency gains in healthcare and other distributed systems must tackle the problem of controlling strategic behavior. This dissertation provided answers and analytical techniques to questions of how incentives should be designed in healthcare systems. Two types of incentives were considered: financial incentives constructed through insurance and remuneration contracts, and knowledge-based incentives for behavioral change.

The first problem considered the implications of interacting preventive efforts on behalf of consumers and providers. Optimal multilateral contracts show a distortion from the benchmark of the bilateral contract, and showed that consumers and providers must both be better off when a consumer is healthy as opposed to ill to exert preventive efforts. This feature of the optimal contracts has been suggested as a feature of new payment systems to improve system incentives (Antos et al., 2009). On the consumer side, this characteristic is common in policies with copays and coinsurance rates. The provider's incentives must be designed such that his payments ensure his reservation level of utility, but also make the desired level of effort his best choice. Achieving the balance of payments and incentives requires a mixed remuneration contract of prospective payments, retrospective payments, and bonuses for good health outcomes of consumers. This finding to control *ex ante* moral hazard strengthens the literature which recommends mixed incentives to solve various agency problems *ex post*, and extends the support for mixed provider incentives when such *ex post* concerns are negligible. To implement such a scheme, there would need to be some association between the consumer and a specific provider with responsibility for the particular condition under contract. Accountable care organizations and medical home arrangements are two examples where such conditions would exist.

The second problem investigated when should knowledge be provided to consumers to affect health behaviors, and subsequently health states and costs. The methodology used relaxed the strong assumption of expected utility maximization on behalf of healthcare consumers, and considers system dynamics, an improvement on the majority of incentives literature which uses static frameworks. The decision to provide knowledge or not was modeled as an MDP from a policy maker or payers and solved via a parallelized value iteration algorithm. The policies regularly showed threshold structures in which knowledge is only cost-saving when provided to sicker consumers. The model was validated to some extent through sensitivity analysis and boundary behavior, and a more thorough validation through design experiments was discussed.

Incentives can be constructed through a variety of tools and for a range of strategic behaviors to be controlled. Financial incentives for preventive efforts from consumers and providers and knowledge-based incentives for health behavior from consumers were considered here independently. Future work on incentives will likely include these mechanisms as well as other policy tools, and consider these participants in addition to hospitals and technology companies. Even more complex models could consider the interaction of financial, knowledge-based, and other incentives to jointly optimize incentive problems. By doing so, healthcare systems can hope to find significant cost savings through strategic behavior of system participants.

REFERENCES

- American Medical Association. 2007. Competition in health insurance: A comprehensive study of U.S. markets. Tech. rep. http://www.ama-assn.org/ama1/pub/upload/mm/368/compstudy_52006.pdf.
- Antos, J., J. Bertko, M. Chernew, D. Cutler, D. Goldman, M. McClellan, E. McGlynn, M. Pauly, L. Schaeffer, S. Shortell. 2009. Bending the curve: Effective steps to address long-term health care spending growth. Tech. rep. http://www.brookings.edu/~media/Files/rc/reports/2009/0826_btc/0826_btc_fullreport.pdf.
- Arrow, K.J., G. Debreu. 1954. Existence of an equilibrium For a competitive economy. *Econometrica* **22**(3) 265–290.
- Bala, M., L. Strzeszynski, K. Cahill. 2008. Mass media interventions for smoking cessation in adults. *Cochrane Database of Systematic Reviews* (1).
- Bauer, U.E., T.M. Johnson, R.S. Hopkins, R.G. Brooks. 2000. Changes in youth cigarette use and intentions following implementation of a tobacco control program - Findings from the Florida Youth Tobacco Survey, 1998-2000. *JAMA-Journal of the American Medical Association* **284**(6) 723–728.
- Beck, J.R., S.G. Pauker. 1983. The Markov process in medical prognosis. *Medical Decision Making* **3**(4) 419–458.
- Bertsekas, D.P., J.N. Tsitsiklis. 1997. *Parallel and Distributed Computation: Numerical Methods*. Athena Scientific, Nashua, NH.
- Bolton, P., M. Dewatripont. 2005. *Contract Theory*. The MIT Press, Cambridge, MA.
- Briggs, A., K. Claxton, M. Sculpher. 2006. *Decision Modeling for Health Economic Evaluation*. Oxford University Press, New York.
- Brown, G.W. 1951. Iterative Solutions of Games by Fictitious Play. Tjallingis Koopmans, ed., *Activity Analysis of Production and Allocation*. Wiley, New York.
- Bureau of Labors Statistics. 2009. Occupational employment statistics. http://www.bls.gov/oes/current/oes_nat.htm#00-0000, Accessed February 7, 2011.
- Cabrales, A., W. Garcia Fontes. 2000. Estimating learning models from experimental data. Economics Working Papers 501, Department of Economics and Business, Universitat Pompeu Fabra. <http://ideas.repec.org/p/upf/upfgen/501.html>.
- Cachon, G.P., F.Q. Zhang. 2006. Procuring fast delivery: Sole sourcing with information asymmetry. *Management Science* **52**(6) 881–896.

- Camerer, C., T.H. Ho. 1999. Experience-weighted attraction learning in normal form games. *Econometrica* **67**(4) 827–874.
- Camerer, C.F., T.H. Ho. 1998. Experience-weighted attraction learning in coordination games: Probability rules, heterogeneity, and time-variation. *Journal of Mathematical Psychology* **42**(2-3) 305–326.
- Camerer, C.F., T.H. Ho, J.K. Chong. 2002. Sophisticated experience-weighted attraction learning and strategic teaching in repeated games. *Journal of Economic Theory* **104**(1) 137–188.
- Carlson, A., M. Lino, W-Y. Juan, K. Hanson, P.P. Basiotis. 2007. Thrifty food plan, 2006 (CNPP-19). Tech. rep.
<http://www.cnpp.usda.gov/Publications/FoodPlans/MiscPubs/TFP2006Report.pdf>.
- Centers For Disease Control and Prevention. 2009. Chronic diseases - The power to prevent, the call to control: At a glance 2009.
<http://www.cdc.gov/chronicdisease/resources/publications/aag/chronic.htm>.
- Chaloupka, F.J., K.E. Warner. 2000. The economics of smoking. Anthony L. Culyer, Joseph P. Newhouse, eds., *Handbook of Health Economics*. North Holland, Amsterdam.
- Courbage, C., A. de Coulon. 2004. Prevention and private health insurance in the UK. *Geneva Papers on Risk and Insurance-Issues and Practice* **29**(4) 719–727.
- Cutler, D.M., A. Lleras-Muney. 2010. Understanding differences in health behaviors by education. *Journal of Health Economics* **29**(1) 1–28.
- Cutler, D.M., R.J. Zeckhauser. 2000. The anatomy of health insurance. Anthony L. Culyer, Joseph P. Newhouse, eds., *Handbook of Health Economics*. North Holland, Amsterdam.
- D'Agostino, R.B., M.W. Russell, D.M. Huse, R.C. Ellison, H. Silbershatz, P.W.F. Wilson, S.C. Hartz. 2000. Primary and subsequent coronary risk appraisal: New results from The Framingham Study. *American Heart Journal* **139**(2, Part 1) 272–281.
- Daskin, M., L. Dean. 2004. Location of health care facilities. Margaret L. Brandeau, Francois Sainfort, William P. Pierskalla, eds., *Operations Research and Health Care: A Handbook of Methods and Applications*. Kluwer Academic, Boston.
- Dave, D., R. Kaestner. 2009. Health insurance and ex ante moral hazard: Evidence from Medicare. *International Journal of Health Care Finance and Economics* **9**(4) 367–390.
- Davis, R.M., E.A. Gilpin, B. Loken, K. Viswanath, M.A. Wakefield. 2008. The role of the media in promoting and reducing tobacco use. Tech. rep.
http://www.cancercontrol.cancer.gov/tcrb/monographs/19/m19_complete.pdf.

- Denton, B.T., M. Kurt, N.D. Shah, S.C. Bryant, S.A. Smith. 2009. Optimizing the start time of statin therapy for patients with Diabetes. *Medical Decision Making* **29**(3) 351–367.
- Diana, M.L. 2009. Exploring information systems outsourcing in US hospital-based health care delivery systems. *Health Care Management Science* **12**(4) 434–450.
- Diehl, H.A. 1998. Coronary risk reduction through intensive community-based lifestyle intervention: The Coronary Health Improvement Project (CHIP) experience. *American Journal of Cardiology* **82**(10B, Sp. Iss. SI) 83T–87T.
- Ebrahim, S., F. Taylor, K. Ward, A. Beswick, M. Burke, G.D. Smith. 2011. Multiple risk factor interventions for primary prevention of coronary heart disease. *Cochrane Database of Systematic Reviews* (1).
- Ehrlich, I., G.S. Becker. 1972. Market insurance, self-insurance, and self-protection. *Journal of Political Economy* **80**(4) 623–648.
- Ellis, R.P. 1998. Creaming, skimping and dumping: Provider competition on the intensive and extensive margins. *Journal of Health Economics* **17**(5) 537–555.
- Ellis, R.P., W.G. Manning. 2007. Optimal health insurance for prevention and treatment. *Journal of Health Economics* **26**(6) 1128–1150.
- Ellis, R.P., T.G. McGuire. 1986. Provider behavior under prospective reimbursement - cost-sharing and supply. *Journal of Health Economics* **5**(2) 129–151.
- Ellis, R.P., T.G. McGuire. 1990. Optimal payment systems for health-services. *Journal of Health Economics* **9**(4) 375–396.
- Engineer, F.G., P. Keskinocak, L.K. Pickering. 2009. Catch-up scheduling for childhood vaccination. *Operations Research* **57**(6) 1307–1319.
- Farquhar, J.W., S.P. Fortmann, J.A. Flora, C.B. Taylor, W.L. Haskell, P.T. Williams, N. Maccoby, P.D. Wood. 1990. Effects of community-wide education on cardiovascular-disease risk-factors - The Stanford 5-city project. *JAMA - Journal of the American Medical Association* **264**(3) 359–365.
- Fryer, R.G. 2010. Financial incentives and student achievement: Evidence from randomized trials. *NBER Working Paper No. w15898*.
- Fudenberg, D., J. Tirole. 1991. *Game Theory*. MIT Press, Cambridge, MA.
- Fuloria, P.C., S.A. Zenios. 2001. Outcomes-adjusted reimbursement in a health-care delivery system. *Management Science* **47**(6) 735–751.
- Gagne, L. 2007. The 2005 British Columbia smoking cessation mass media campaign and short-term changes in smoking. *Journal of Public Health Management and Practice* **13**(3) 296–306.

- Gallien, J. 2006. Dynamic mechanism design for online commerce. *Operations Research* **54**(2) 291–310.
- Glanz, K., B.K. Rimer, F.M. Lewis. 2002. *Health Behavior and Health Education: Theory, Research, and Practice*. Jossey-Bass Publishing, San Francisco.
- Goldman, D., T.J. Philipson. 2007. Integrated insurance design in the presence of multiple medical technologies. *American Economic Review* **97**(2) 427–432.
- Hazen, G. 2011. Cohort decomposition for Markov cost-effectiveness models. *Medical Decision Making* **31**(1) 19–34.
- Heimendinger, J., T. Uyeki, A. Andhara, J.A. Marshall, S. Scarbro, E. Belansky, L. Crane. 2007. Coaching process outcomes of a family visit nutrition and physical activity intervention. *Health Education & Behavior* **34**(1) 71–89.
- Hibbard, J.H., P. Slovic, J.J. Jewett. 1997. Informing consumer decisions in health care: Implications from decision-making research. *Milbank Quarterly* **75**(3) 395–414.
- Institute of Medicine. 2010. *The Healthcare Imperative: Lowering Costs and Improving Outcomes: Workshop Series Summary*. The National Academies Press, Washington, DC.
- Kenkel, D.S. 2000. Prevention. Anthony L. Culyer, Joseph P. Newhouse, eds., *Handbook of Health Economics*. North Holland, Amsterdam.
- Krishna, V. 2009. *Auction Theory*. Academic Press, Burlington, MA.
- Laffont, J.J., D. Martimort. 2002. *The Theory of Incentives: The Principal-agent Model*. Princeton University Press, Princeton, NJ.
- Lee, E.K., M. Zaider. 2008. Operations research advances cancer therapeutics. *Interfaces* **38**(1) 5–25.
- Lin, J.S., E. O'Connor, E.P. Whitlock, T.L. Beil. 2010. Behavioral counseling to promote physical activity and a healthful diet to prevent cardiovascular disease in adults: A systematic review for the U.S. Preventive Services Task Force. *Annals of Internal Medicine* **153**(11) 736–W244.
- Lloyd-Jones, D., R.J. Adams, T.M. Brown, M. Carnethon, S. Dai, G. De Simone, T.B. Ferguson, E. Ford, K. Furie, C. Gillespie, A. Go, K. Greenlund, N. Haase, S. Hailpern, P.M. Ho, V. Howard, B. Kissela, S. Kittner, D. Lackland, L. Lisabeth, A. Marelli, M.M. McDermott, J. Meigs, D. Mozaffarian, M. Mussolino, G. Nichol, V.L. Roger, W. Rosamond, R. Sacco, P. Sorlie, R. Stafford, T. Thom, S. Wasserthiel-Smoller, N.D. Wong, J. Wylie-Rosett. 2010. Heart disease and stroke statistics-2010 update a report from the American Heart Association. *Circulation* **121**(7) E46–E215.

- Ma, C.T.A., T.G. McGuire. 1997. Optimal health insurance and provider payment. *American Economic Review* **87**(4) 685–704.
- Ma, C.T.A., M.H. Riordan. 2002. Health insurance, moral hazard, and managed care. *Journal of Economics & Management Strategy* **11**(1) 81–107.
- Machina, M.J. 2004. Nonexpected Utility Theory. JL Teugels, S Bjorn, eds., *Encyclopedia of Actuarial Science*. Wiley, New York.
- Maibach, E.W., M.L. Rothschild, W.D. Novelli. 2002. *Social Marketing*, Chap. XIX. Jossey-Bass Publishing, San Francisco.
- Mas-Colell, A., M. Whinston, J. Green. 1995. *Microeconomic Theory*. Oxford University Press, New York.
- McAlister, A.L., A.G. Ramirez, C. Amezcua, L.V. Pulley, M.P. Stern, S. Mercado. 1992. Smoking cessation in Texas-Mexico border communities: A quasi-experimental panel study. *American Journal of Health Promotion* **6**(4) 274–9.
- Merzel, C., J. D’Afflitti. 2003. Reconsidering community-based health promotion: Promise, performance, and potential. *American Journal of Public Health* **93**(4) 557–574.
- Montãno, D.E., D. Kasprzyk. 2002. *The Theory of Reasoned Action and the Theory of Planned Behavior*, Chap. IV. Jossey-Bass Publishing, San Francisco.
- Moser, A., I. Korstjens, T. van der Weijden, H. Tange. 2010. Themes affecting health-care consumers’ choice of a hospital for elective surgery when receiving web-based comparative consumer information. *Patient Education and Counseling* **78**(3, Sp. Iss. SI) 365–371.
- Nash, J. 1951. Non-cooperative games. *Annals of Mathematics* **54**(2) 286–295.
- National Heart, Lung and Blood Institute. 2009. Morbidity & mortality: 2009 chart book on cardiovascular, lung, and blood diseases. Tech. rep. http://www.nhlbi.nih.gov/resources/docs/2009_ChartBook.pdf.
- Organisation for economic co-operation and development. 2009. OECD Health Data 2009 <http://stats.oecd.org/Index.aspx?DataSetCode=HEALTH>.
- Owens, D.K., D.M. Edwards, J.F. Cavallaro, R.D. Shachter. 2004. The cost effectiveness of partially effective HIV vaccines. Margaret L. Brandeau, Francois Sainfort, William P. Pierskalla, eds., *Operations Research and Health Care: A Handbook of Methods and Applications*. Kluwer Academic, Boston.
- Parker, D.R., A.R. Assaf. 2005. Community interventions for cardiovascular disease. *Primary Care* **32**(4) 865–881.

- Patrick, J., M.L. Puterman, M. Queyranne. 2008. Dynamic multipriority patient scheduling for a diagnostic resource. *Operations Research* **56**(6) 1507–1525.
- Pierskalla, W.P. 2004. Supply chain management of blood banks. Margaret L. Brandeau, Francois Sainfort, William P. Pierskalla, eds., *Operations Research and Health Care: A Handbook of Methods and Applications*. Kluwer Academic, Boston.
- Porter, M., E. Teisberg. 2006. *Redefining Health Care: Creating Value-based Competition on Results*. Harvard Business School Press, Boston.
- Prendergast, C. 1999. The provision of incentives in firms. *Journal of Economic Literature* **37**(1) 7–63.
- Puska, P., A. Nissinen, J. Tuomilehto, J.T. Salonen, K. Koskela, A. McAlister, T.E. Kottke, N. Maccoby, J.W. Farquhar. 1985. The community-based strategy to prevent coronary heart-disease - conclusions from the 10 Years of the North Karelia Project. *Annual Review of Public Health* **6** 147–193.
- Puterman, M.L. 2005. *Markov Decision Processes*. Wiley-Interscience, New York.
- Randolph, W., K. Viswanath. 2004. Lessons learned from public health mass media campaigns: Marketing health in a crowded media world. *Annual Review of Public Health* **25** 419–437.
- Rimer, B.K., K. Glanz, G. Rasband. 2001. Searching for evidence about health education and health behavior interventions. *Health Education & Behavior* **28**(2) 231–248.
- Robert, C.P., G. Casella. 2010. *Introducing Monte Carlo Methods with R*. Springer, New York.
- Rohack, J.J., A.F. Einboden. 2006. The Ecology of Healthcare in Texas and the United States.
- Ronckers, E.T., W. Groot, A.J.H.A. Ament. 2005. Systematic review of economic evaluations of smoking cessation: Standardizing the cost-effectiveness. *Medical Decision Making* **25**(4) 437–448.
- Roth, A.E., I. Erev. 1995. Learning in extensive-form games - experimental data and simple dynamic models in the intermediate-term. *Games and Economic Behavior* **8**(1) 164–212.
- Rothschild, M., J. Stiglitz. 1976. Equilibrium in competitive insurance markets - essay on economics of imperfect information. *Quarterly Journal of Economics* **90**(4) 629–649.
- Schar, E., K. Gutierrez. 2001. Smoking cessation media campaigns from around the world: Recommendations from lessons learned. Tech. rep. http://www.ash.org.uk/files/documents/ASH_324.pdf.

- Secker-Walker, R.H., B.S. Flynn, L.J. Solomon, J.M. Skelly, A.L. Dorwaldt, T. Ashikaga. 2000. Helping women quit smoking: Results of a community intervention program. *American Journal of Public Health* **90**(6) 940–946.
- Shechter, S.M., M.D. Bailey, A.J. Schaefer, M.S. Roberts. 2008. The optimal time to initiate HIV therapy under ordered health states. *Operations Research* **56**(1) 20–33.
- Shoham, Y., K. Leyton-Brown. 2009. *Multiagent Systems: Algorithmic, Game-theoretic, and Logical Foundations*. Cambridge University Press, New York.
- Sonnenberg, F.A., J.R. Beck. 1993. Markov-models in medical decision-making - A practical guide. *Medical Decision Making* **13**(4) 322–338.
- Spall, J.C. 2003. *Introduction to Stochastic Search and Optimization*. Wiley, New York.
- Stanciole, A.E. 2008. Health insurance and lifestyle choices: Identifying ex ante moral hazard in the US market. *Geneva Papers on Risk and Insurance-Issues and Practice* **33**(4) 627–644.
- Stern, M.P., J.W. Farquhar, N. Maccoby, S.H. Russell. 1976. Results of a 2-year health-education campaign on dietary behavior - Stanford 3 Community Study. *Circulation* **54**(5) 826–833.
- Su, X., S.A. Zenios. 2006. Recipient choice can address the efficiency-equity trade-off in kidney transplantation: A mechanism design model. *Management Science* **52**(11) 1647–1660.
- Tengs, T.O., N.D. Osgood, L.L. Chen. 2001. The cost-effectiveness of intensive national school-based anti-tobacco education: Results from the Tobacco Policy Model. *Preventive Medicine* **33**(6) 558–570.
- Toobert, D.J., L.A. Strycker, R.E. Glasgow, M. Barrera, K. Angell. 2005. Effects of the Mediterranean Lifestyle Program on multiple risk behaviors and psychosocial outcomes among women at risk for heart disease. *Annals of Behavioral Medicine* **29**(2) 128–137.
- Tosteson, A.N.A., M.C. Weinstein, M.G.M. Hunink, M.A. Mittleman, L.W. Williams, P.A. Goldman, L. Goldman. 1997. Cost-effectiveness of populationwide educational approaches to reduce serum cholesterol levels. *Circulation* **95**(1) 24–30.
- Town, R., R. Kane, P. Johnson, M. Butler. 2005. Economic incentives and physicians' delivery of preventive care - A systematic review. *American Journal of Preventive Medicine* **28**(2) 234–240.
- Tragler, G., J.P. Caulkins, G. Feichtinger. 2001. Optimal dynamic allocation of treatment and enforcement in illicit drug control. *Operations Research* **49**(3) 352–362.
- Truett, J., J. Cornfield, W. Kannel. 1967. A multivariate analysis of risk of coronary heart disease in Framingham. *Journal of Chronic Diseases* **20**(7) 511–524.

- Truffer, C.J., S. Keehan, S. Smith, J. Cylus, A. Sisko, J.A. Poisal, J. Lizonitz, M.K. Clemens. 2010. Health spending projections through 2019: The recession's impact continues. *Health Affairs* **29**(3) 522–529.
- Valdez, R.S., E. Ramly, P.F. Brennan. 2010. Industrial and systems engineering and health care: Critical areas of research—final report. Tech. rep. http://healthit.ahrq.gov/portal/server.pt/gateway/PTARGS_0_3882_948220_0_0_18/IndustrialAndSystemsEngineeringAndHealthCare.
- Von Neumann, J., O. Morgenstern. 1944. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ.
- Welton, N.J., A.E. Ades. 2005. Estimation of Markov chain transition probabilities and rates from fully and partially observed data: Uncertainty propagation, evidence synthesis, and model calibration. *Medical Decision Making* **25**(6) 633–645.
- Wendel-Vos, G.C.W., A.E. Dutman, W.M.M. Verschuren, E.T. Ronckers, A. Ament, P. van Assema, J. van Ree, E.C. Ruland, A.J. Schuit. 2009. Lifestyle factors of a five-year community-intervention program The Hartslag Limburg Intervention. *American Journal of Preventive Medicine* **37**(1) 50–56.
- Wilson, P.W.F., R.B. D'Agostino, D. Levy, A.M. Belanger, H. Silbershatz, W.B. Kannel. 1998. Prediction of coronary heart disease using risk factor categories. *Circulation* **97**(18) 1837–1847.
- Winkleby, M.A., J.A. Flora, H.C. Kraemer. 1994. A community-based heart-disease intervention - predictors of change. *American Journal of Public Health* **84**(5) 767–772.
- World Health Organization. 2008. WHO report on the global tobacco epidemic, 2008: The MPOWER package. Tech. rep., Geneva. http://whqlibdoc.who.int/publications/2008/9789241596282_eng.pdf.
- World Health Organization. 2009. Ten facts about chronic disease http://www.who.int/features/factfiles/chp/10_en.html, Accessed November 10, 2009.
- Yaesoubi, R., S.D. Roberts. 2009. Coordinating contracts in a preventive health care system. <http://www4.ncsu.edu/~ryaesou/Documents/Papers/Contracts-MS2009.pdf>.
- Zeckhauser, R. 1970. Medical insurance - case study of tradeoff between risk spreading and appropriate incentives. *Journal of Economic Theory* **2**(1) 10–26.
- Zweifel, P., F. Breyer, M. Kifmann. 2009. *Health Economics*. Springer, New York.

APPENDIX A

Properties of $e_C(\cdot)$ with Complementary Efforts

Claim: $\frac{\partial e_C(e_P; \Delta u)}{\partial e_P} = e'_C(e_P; \Delta u) \geq 0$.

Proof: In (3.2), replace e_C by $e_C(e_P; \Delta u)$, then (3.2) becomes

$$\pi'(e_C(e_P; \Delta u))\Delta u = \frac{\partial \Psi}{\partial e_C}(e_C(e_P; \Delta u), e_P).$$

Taking derivatives with respect to e_P gives

$$\pi''(e_C(e_P; \Delta u))e'_C(e_P; \Delta u)\Delta u = \frac{\partial^2 \Psi}{\partial e_C^2}(e_C(e_P; \Delta u), e_P)e'_C(e_P; \Delta u) + \frac{\partial^2 \Psi}{\partial e_C \partial e_P}(e_C(e_P; \Delta u), e_P)$$

Rearranging gives

$$e'_C(e_P; \Delta u) = \frac{\frac{\partial^2 \Psi}{\partial e_C \partial e_P}(e_C(e_P; \Delta u), e_P)}{\left[\pi''(e_C(e_P; \Delta u))\Delta u - \frac{\partial^2 \Psi}{\partial e_C^2}(e_C(e_P; \Delta u), e_P) \right]} = \frac{-}{(-)(+) - (+)} \geq 0$$

Claim: $\frac{\partial e_C(e_P; \Delta u)}{\partial \Delta u} \geq 0$.

Proof: Again (3.2) gives that

$$\pi'(e_C(e_P; \Delta u))\Delta u = \frac{\partial \Psi}{\partial e_C}(e_C(e_P; \Delta u), e_P).$$

Taking derivatives with respect to Δu gives

$$\pi''(e_C(e_P; \Delta u))\frac{\partial e_C(e_P; \Delta u)}{\partial \Delta u}\Delta u + \pi'(e_C(e_P; \Delta u)) = \frac{\partial^2 \Psi}{\partial e_C^2}(e_C(e_P; \Delta u), e_P)\frac{\partial e_C(e_P; \Delta u)}{\partial \Delta u}$$

Solving for $\frac{\partial e_C(e_P; \Delta u)}{\partial \Delta u}$ yields

$$\frac{\partial e_C(e_P; \Delta u)}{\partial \Delta u} = \frac{\pi'(\cdot)}{\frac{\partial^2 \Psi(\cdot)}{\partial e_C^2} - \pi''(\cdot) \Delta u} = \frac{+}{(+)-(-)(+)} \geq 0.$$

Concavity of IC Constraints

The consumer's IC constraint (3.3.1) is clearly concave since

$$\frac{\partial^2}{\partial e_C^2} [\pi(\hat{e}_C)u_h + [1 - \pi(\hat{e}_C)]u_i - \Psi(\hat{e}_C, e_P)] = \pi''(\cdot) \Delta u - \frac{\partial^2 \Psi(\cdot)}{\partial e_C^2} = (-) - (+) \leq 0.$$

Considering the provider's IC constraint (3.3.1),

$$\begin{aligned} & \frac{\partial^2}{\partial e_P^2} [\pi(e_C(\hat{e}_P; \Delta u))t_h^P + [1 - \pi(e_C(\hat{e}_P; \Delta u))](t_i^P - d) - c_P(\hat{e}_P)] \\ &= [\pi''(\cdot)(e'_C(\cdot))^2 + \pi'(\cdot)e''_C(\cdot)] (\Delta t_P + d) - c''_P(\cdot) = [(-)(+) + (+)(?)](+) - (+), \end{aligned}$$

where the unknown sign comes from $e''_C(\cdot)$. If this sign is negative, the provider's IC constraint is concave. The reason that $e'_C(\cdot) \geq 0$ is that the provider's effort lowers consumer disutility, and thus aids in inducing consumer effort. Since the benefits from provider effort are assumed to be decreasing, it is expected that $e''_C(\cdot) \leq 0$. More rigorously, it was previously found that

$$e'_C(e_P; \Delta u) = \frac{\frac{\partial^2 \Psi}{\partial e_C \partial e_P}(e_C(e_P; \Delta u), e_P)}{\left[\pi''(e_C(e_P; \Delta u)) \Delta u - \frac{\partial^2 \Psi}{\partial e_C^2}(e_C(e_P; \Delta u), e_P) \right]}$$

Denote the numerator (denominator) of this expression by † (‡), which are both negative.

Then,

$$e_C''(\cdot) = \left(\frac{1}{\ddagger^2}\right) \left(\left[\frac{\partial^3 \Psi}{\partial e_C^2 \partial e_P} e_C'(\cdot) + \frac{\partial^3 \Psi}{\partial e_C \partial e_P^2} \right] \ddagger - \left[\pi'''(\cdot) e_C'(\cdot) \Delta u - \frac{\partial^3 \Psi}{\partial e_C^2 \partial e_P} - \frac{\partial^3 \Psi}{\partial e_C^3} e_C'(\cdot) \right] \dagger \right)$$

Then since $e_C'(\cdot) \geq 0$ and $\frac{\partial^3 \Psi(\cdot)}{\partial e_C \partial e_P^2} \geq 0$, the following technical conditions (A.1) are sufficient for the provider's IC constraint to be concave.

$$\begin{aligned} \pi'''(\cdot) &\leq 0 \\ \frac{\partial^3 \Psi(\cdot)}{\partial e_C^2 \partial e_P} &\geq 0 \\ \frac{\partial^3 \Psi(\cdot)}{\partial e_C^3} &\geq 0 \end{aligned} \tag{A.1}$$

Proof of Proposition 3.3.1

Previously it was found that

$$e_C'(e_P; \Delta u) = \frac{\frac{\partial^2 \Psi}{\partial e_C \partial e_P}(e_C(e_P; \Delta u), e_P)}{\left[\pi''(e_C(e_P; \Delta u)) \Delta u - \frac{\partial^2 \Psi}{\partial e_C^2}(e_C(e_P; \Delta u), e_P) \right]}.$$

Now taking the derivative with respect to Δu , (again using the same †, ‡ notation for the numerator and denominator, which are both negative)

$$\frac{\partial e_C'(\cdot)}{\partial \Delta u} = \left(\frac{1}{\ddagger^2}\right) \left[\left(\frac{\partial^3 \Psi(\cdot)}{\partial e_C^2 \partial e_P} \frac{\partial e_C(\cdot)}{\partial \Delta u} \right) \ddagger - \left(\pi''(\cdot) + \pi'''(\cdot) \frac{\partial e_C(\cdot)}{\partial \Delta u} \Delta u - \frac{\partial^3 \Psi(\cdot)}{\partial e_C^3} \frac{\partial e_C(\cdot)}{\partial \Delta u} \right) \dagger \right].$$

The term in the brackets needs to be examined. By the previous conditions (A.1), the signs of terms within the brackets are

$$(+)(+)(-) - [(-) + (-)(+) - (+)(+)](-) \leq 0.$$

Proof of Proposition 3.3.2

Differentiating and adding a positive term to the right hand side, gives that

$$\frac{\partial}{\partial e_C} \left(\frac{\pi'_{e_C}(\cdot)}{\pi'_{e_P}(\cdot)} \right) \leq 0 \Rightarrow \frac{\pi''_{e_C} \pi'_{e_P} - \pi''_{e_C e_P} \pi'_{e_C}}{(\pi'_{e_P})^2} \leq 0 \Rightarrow \frac{\pi''_{e_C} \pi'_{e_P} - \pi''_{e_C e_P} \pi'_{e_C}}{(\pi'_{e_P})^2} \leq \frac{\Psi''}{\Delta u \cdot \pi'_{e_P}}.$$

Rearranging terms gives that

$$\pi'_{e_P} \geq \pi'_{e_C} \frac{\Delta u \cdot \pi''_{e_C e_P}}{\Delta u \cdot \pi''_{e_C} - \Psi''}.$$

Finally substituting in $\frac{\partial e_C}{\partial e_P}$ gives

$$\pi'_{e_C} \frac{\partial e_C}{\partial e_P} + \pi'_{e_P} \geq 0.$$

Therefore, if $\Delta v \leq 0$, the provider maximizes utility by setting effort equal to 0.

Proof of Proposition 3.3.3

Substituting in $\frac{\partial e_C}{\partial e_P} = \frac{\Delta u \cdot \pi''_{e_C e_P}}{\Psi'' - \Delta u \cdot \pi''_{e_C}}$ transforms the bracketed term in equations (3.26) and (3.26) into

$$\left[\left(\frac{\pi''_{e_C e_P} \cdot \Psi''}{\Psi'' - \Delta u \cdot \pi''_{e_C}} \right) \frac{\partial e_C}{\partial \Delta u} + \pi'_{e_C} \frac{\partial^2 e_C}{\partial e_P \partial \Delta u} \right]$$

Under the assumptions on $\pi(e_C, e_P)$, and the second stipulation of the proposition, this term is negative. The first stipulation of the proposition ensure the sign of γ is consistent with Corollary 3.3.1, and the result follows from the form of equations (3.26) and (3.26).

Sensitivity Analysis of Multilateral Savings

Figure A.1. Sensitivity to a with $b = 0.20$. $a = \{0.1, 0.2, 0.3, 0.4, 0.5\}$,
 $(b, d, q, r, w) = (0.20, 5, 0.05, 0.10, 50)$

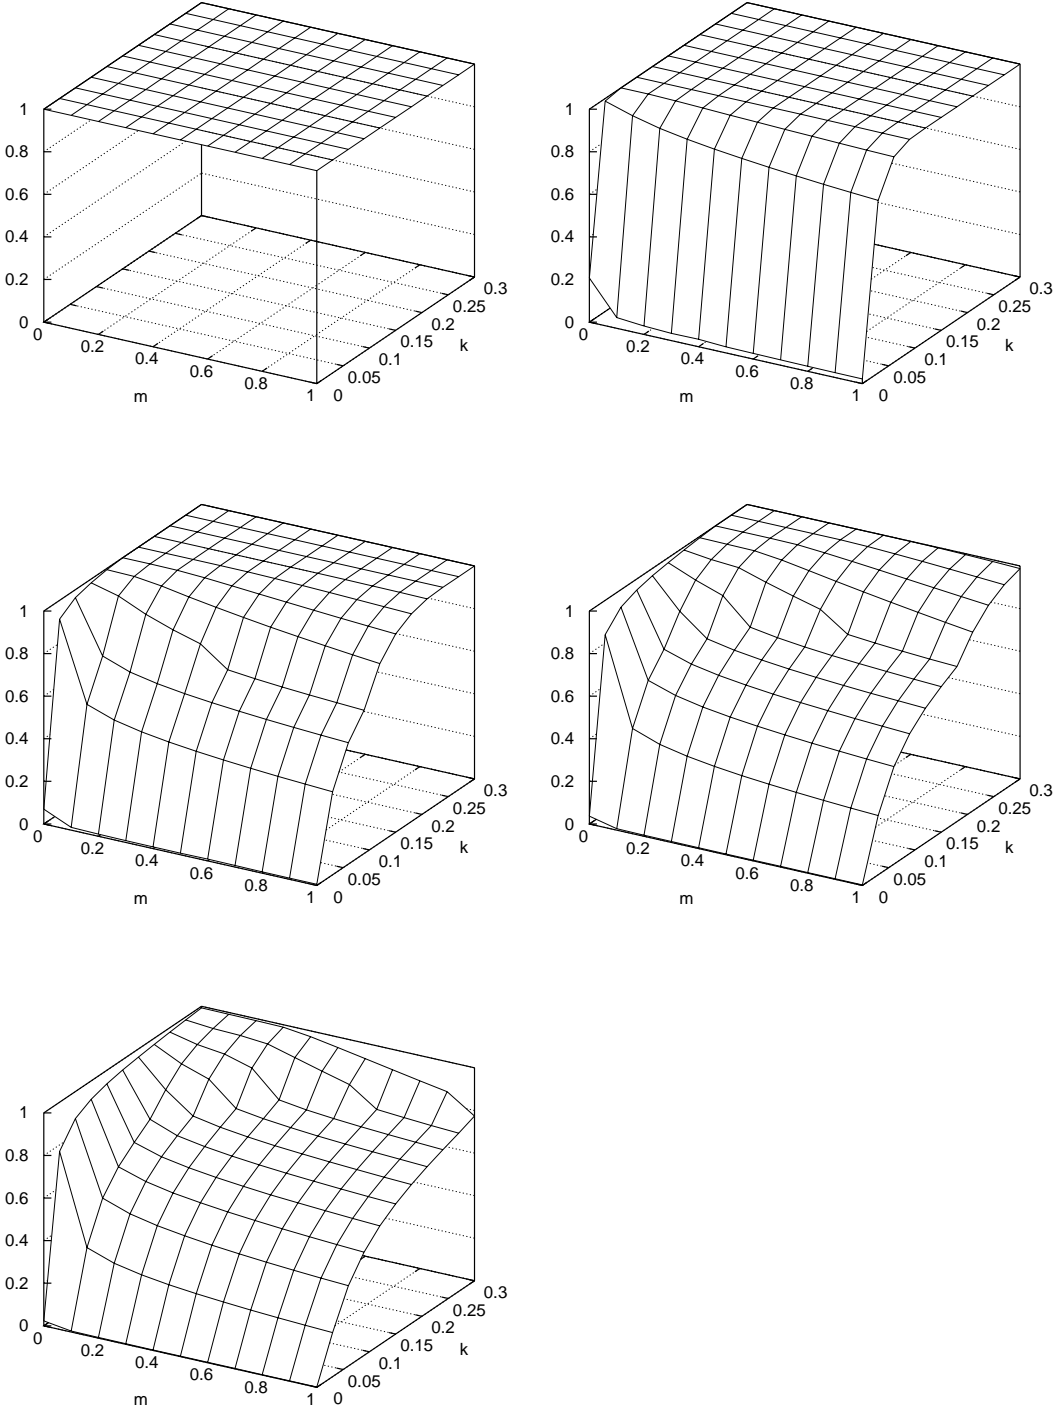


Figure A.2. Sensitivity to a with $b = 0.25$. $a = \{0.1, 0.2, 0.3, 0.4, 0.5\}$,
 $(b, d, q, r, w) = (0.25, 5, 0.05, 0.10, 50)$

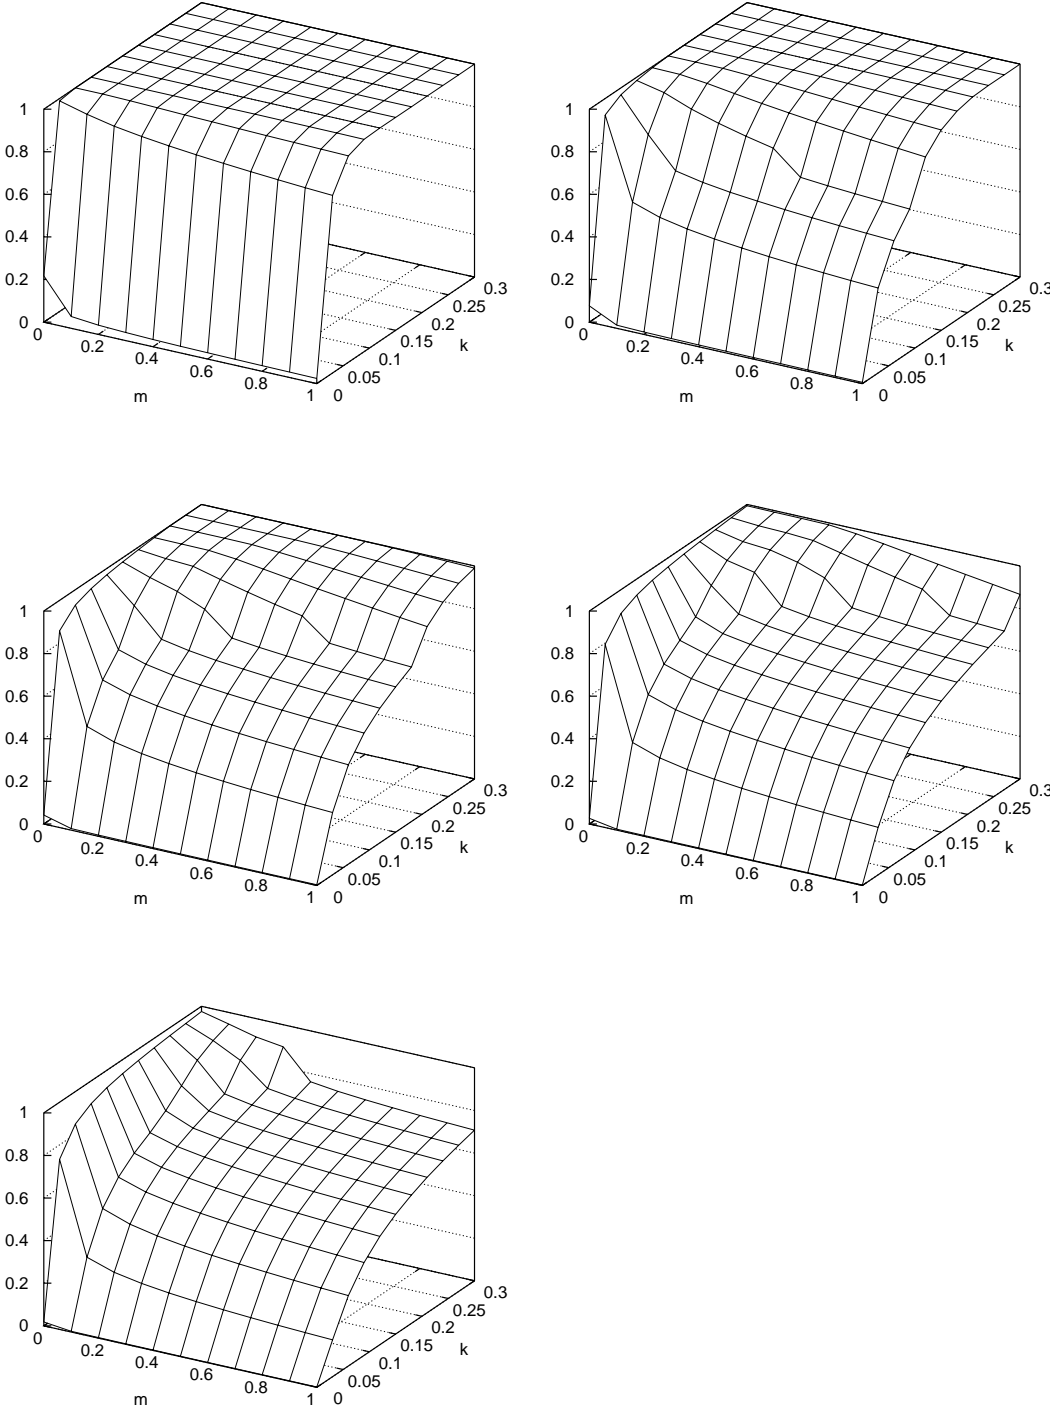


Figure A.3. Sensitivity to a with $b = 0.35$. $a = \{0.1, 0.2, 0.3, 0.4, 0.5\}$,
 $(b, d, q, r, w) = (0.35, 5, 0.05, 0.10, 50)$

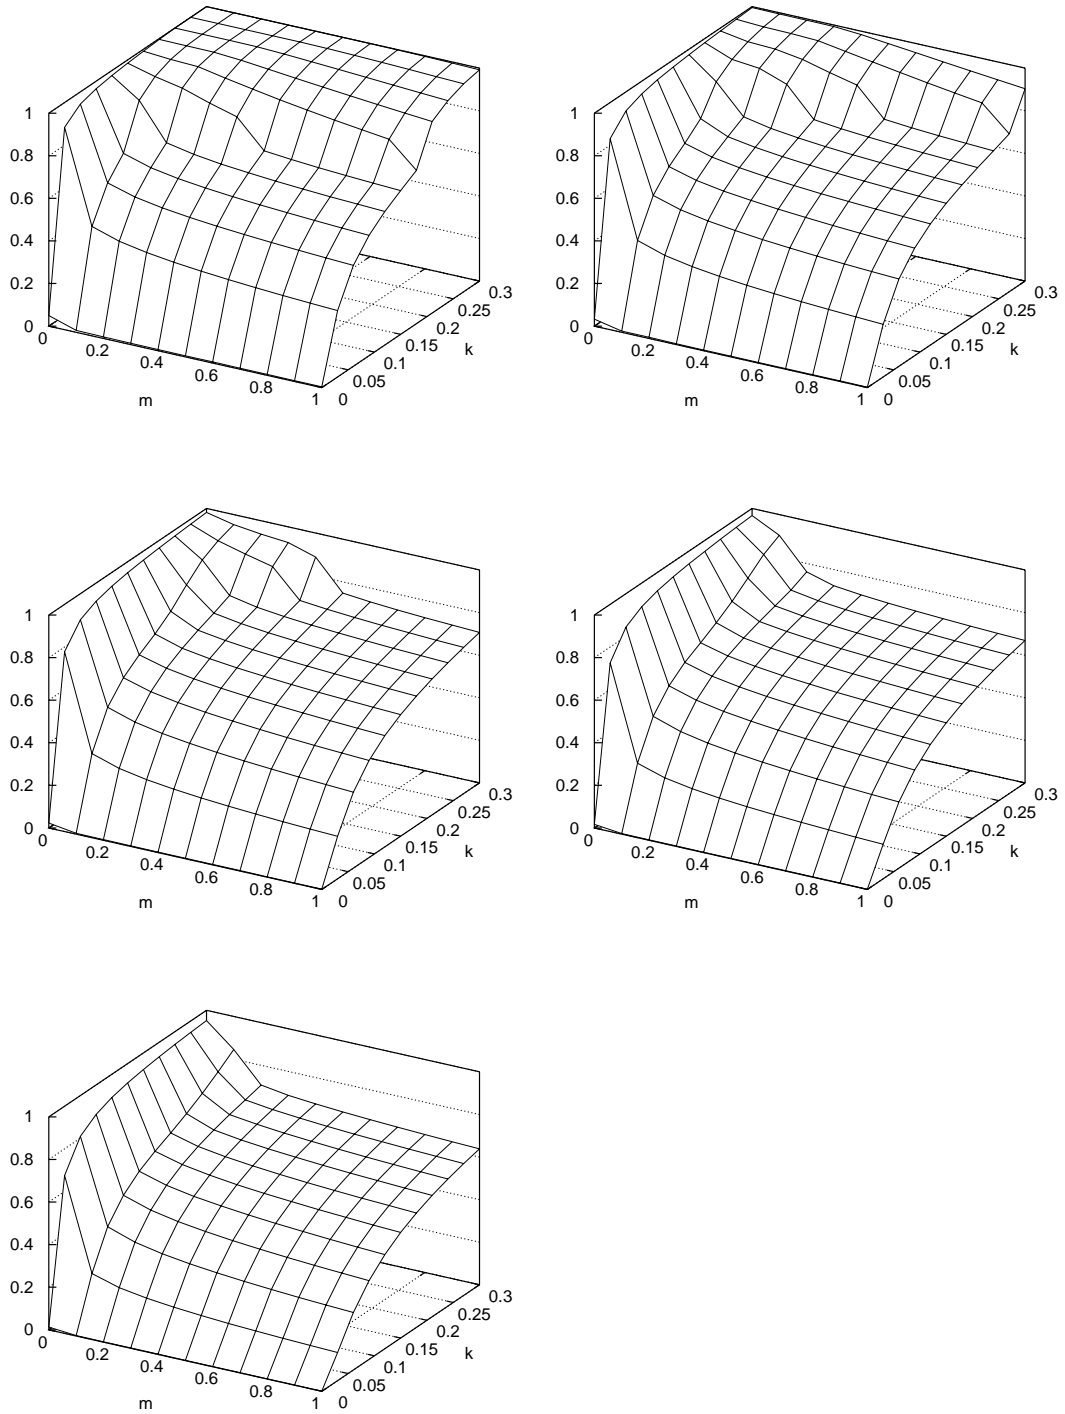


Figure A.4. Sensitivity to a with $b = 0.40$. $a = \{0.1, 0.2, 0.3, 0.4, 0.5\}$,
 $(b, d, q, r, w) = (0.40, 5, 0.05, 0.10, 50)$

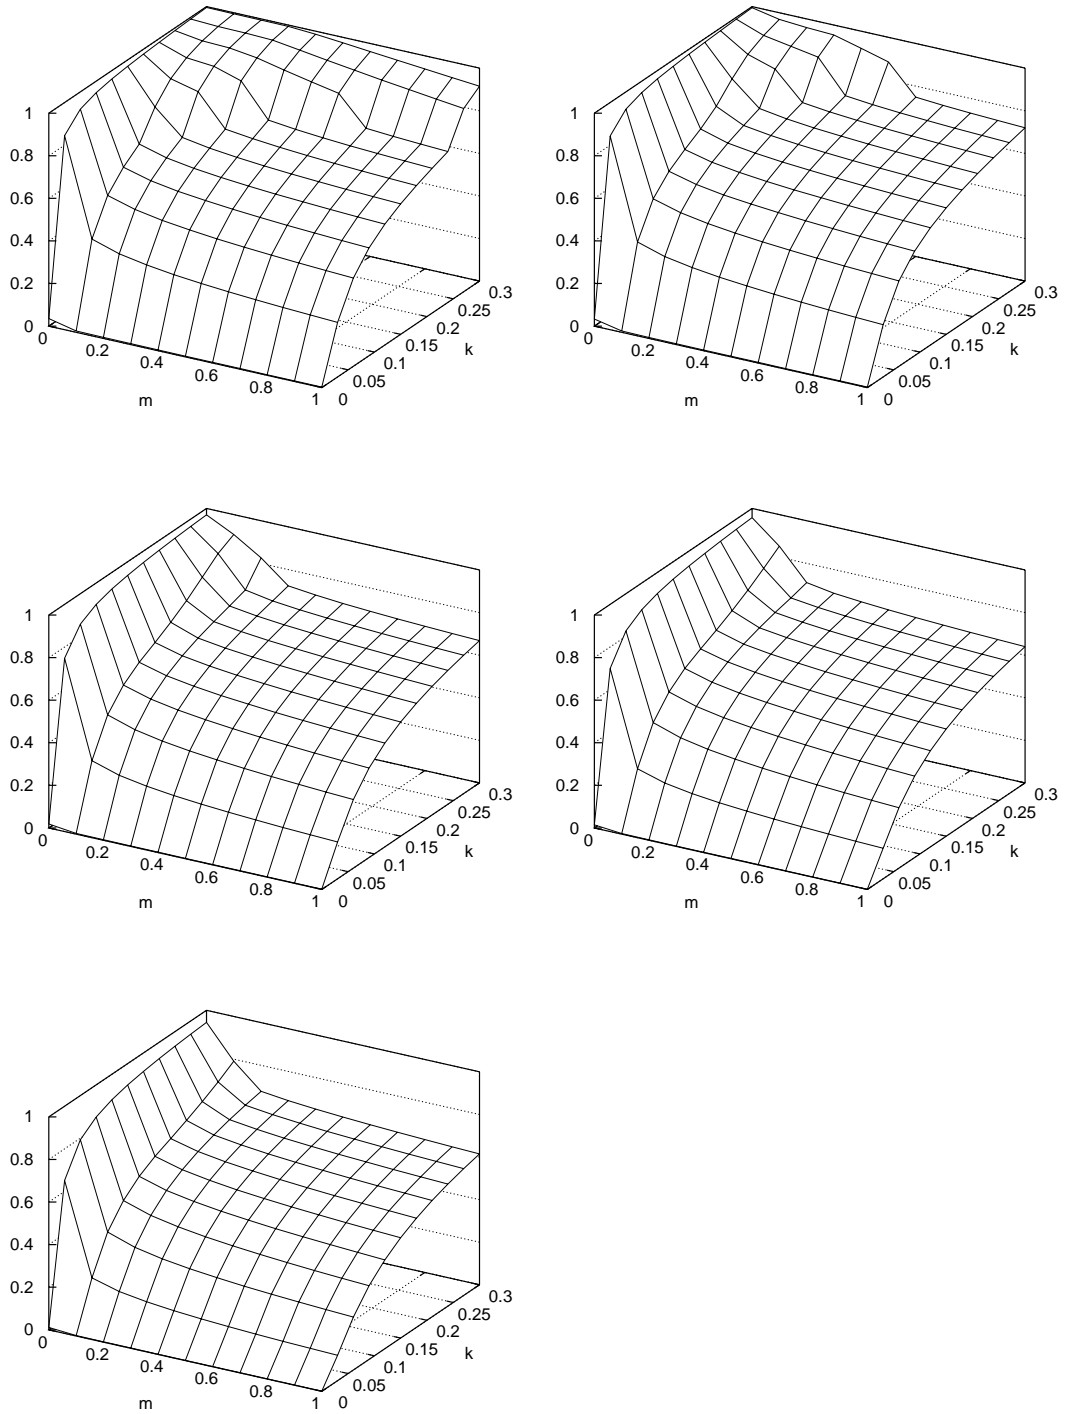


Figure A.5. Sensitivity to a with $b = 0.45$. $a = \{0.1, 0.2, 0.3, 0.4, 0.5\}$,
 $(b, d, q, r, w) = (0.45, 5, 0.05, 0.10, 50)$

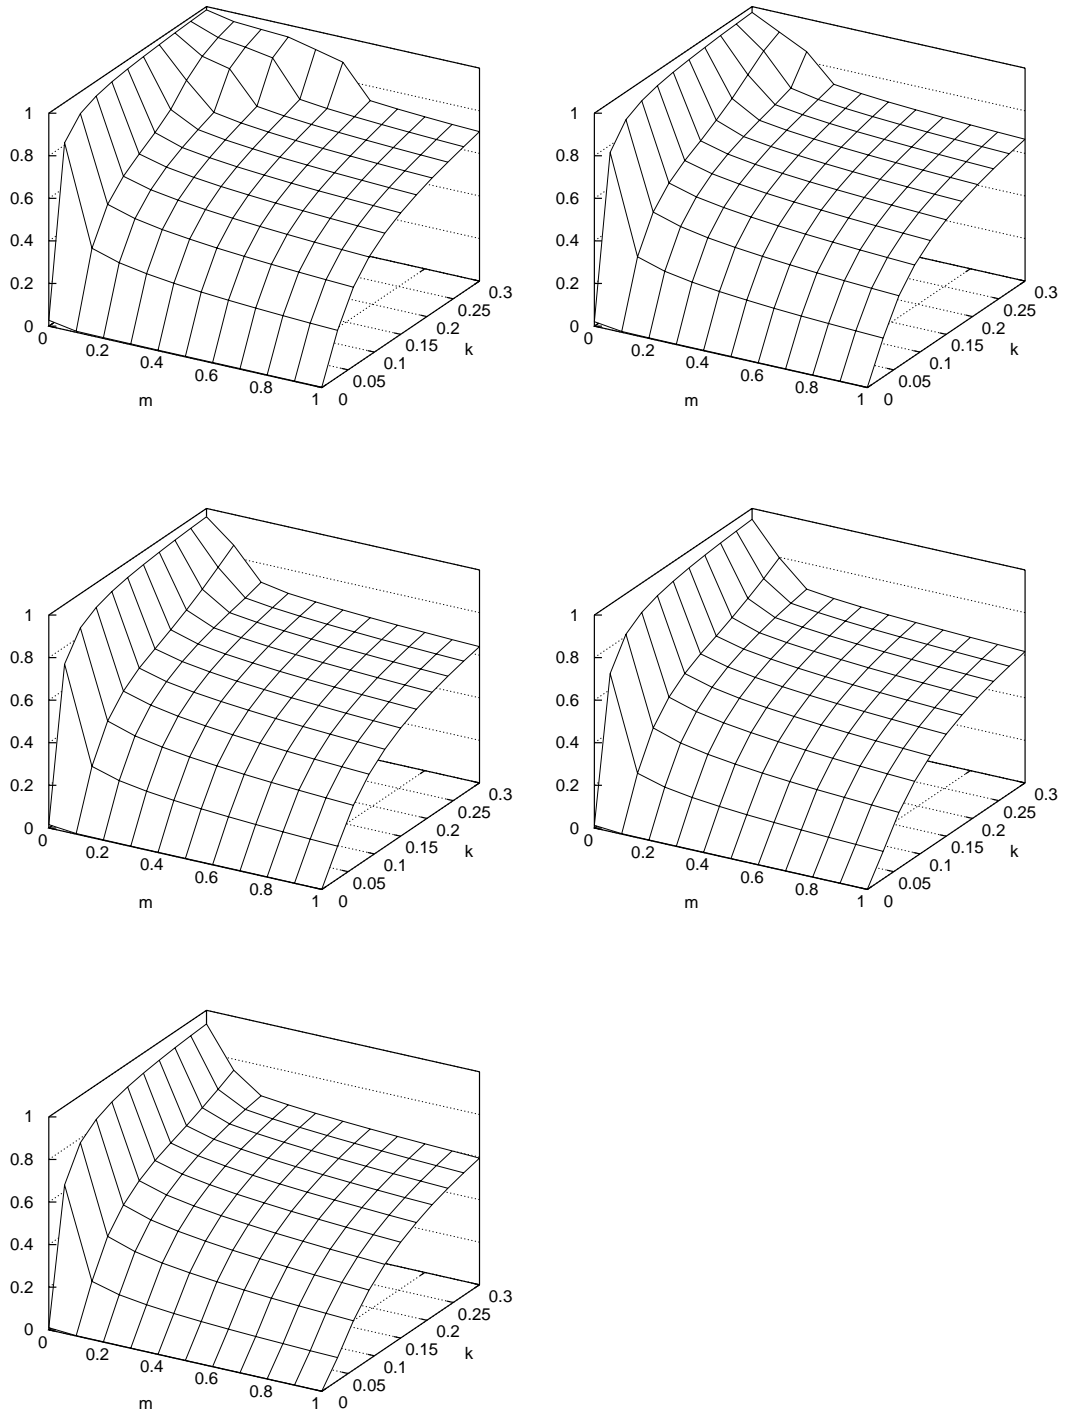


Figure A.6. Sensitivity to a with $b = 0.50$. $a = \{0.1, 0.2, 0.3, 0.4, 0.5\}$,
 $(b, d, q, r, w) = (0.50, 5, 0.05, 0.10, 50)$

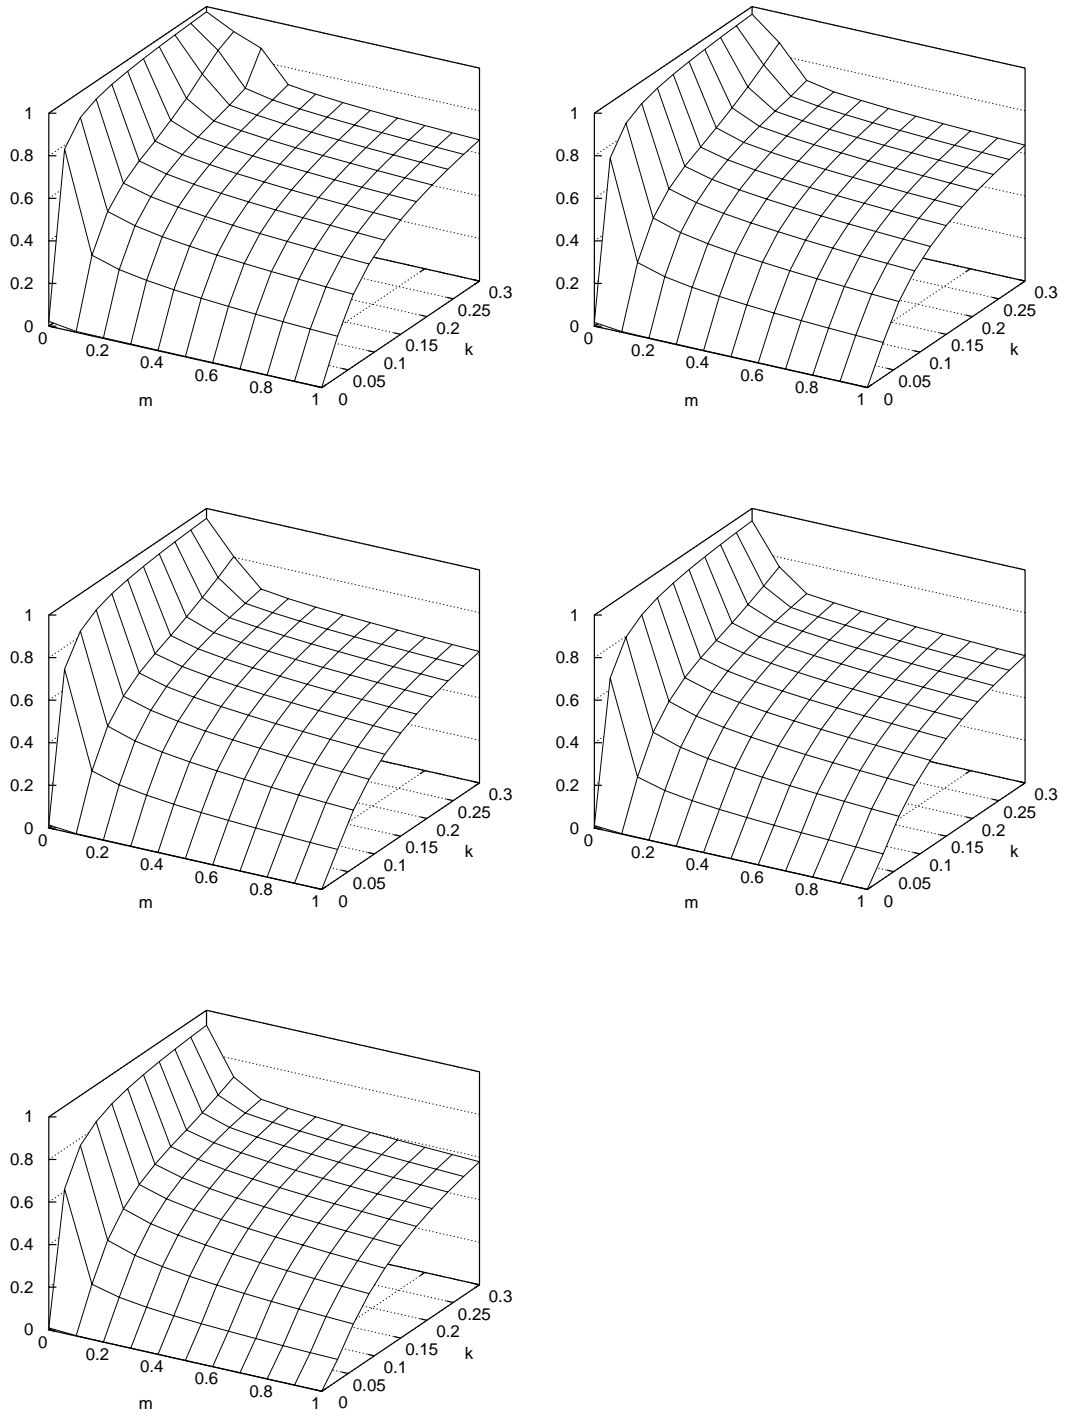


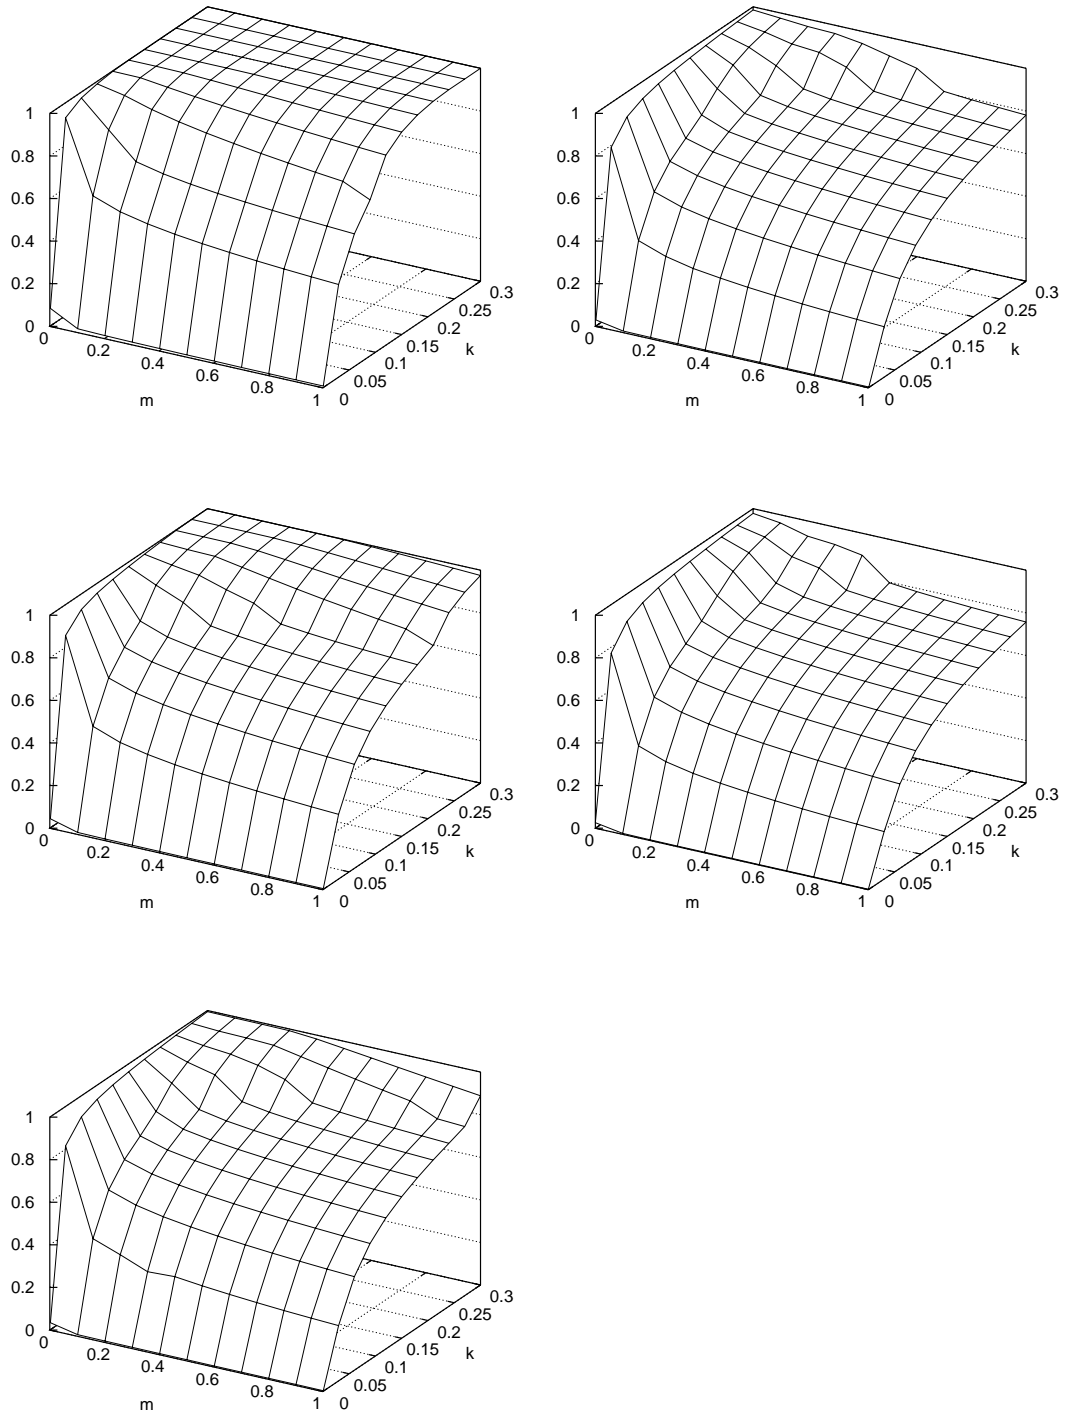
Figure A.7. Sensitivity to r with $q = 0.03$. $r = \{0.05, 0.075, 0.10, 0.125, 0.15\}$, $(a, b, d, q, w) = (0.30, 0.30, 5, 0.03, 50)$ 

Figure A.8. Sensitivity to r with $q = 0.04$.
 $r = \{0.05, 0.075, 0.10, 0.125, 0.15\}$, $(a, b, d, q, w) = (0.30, 0.30, 5, 0.04, 50)$

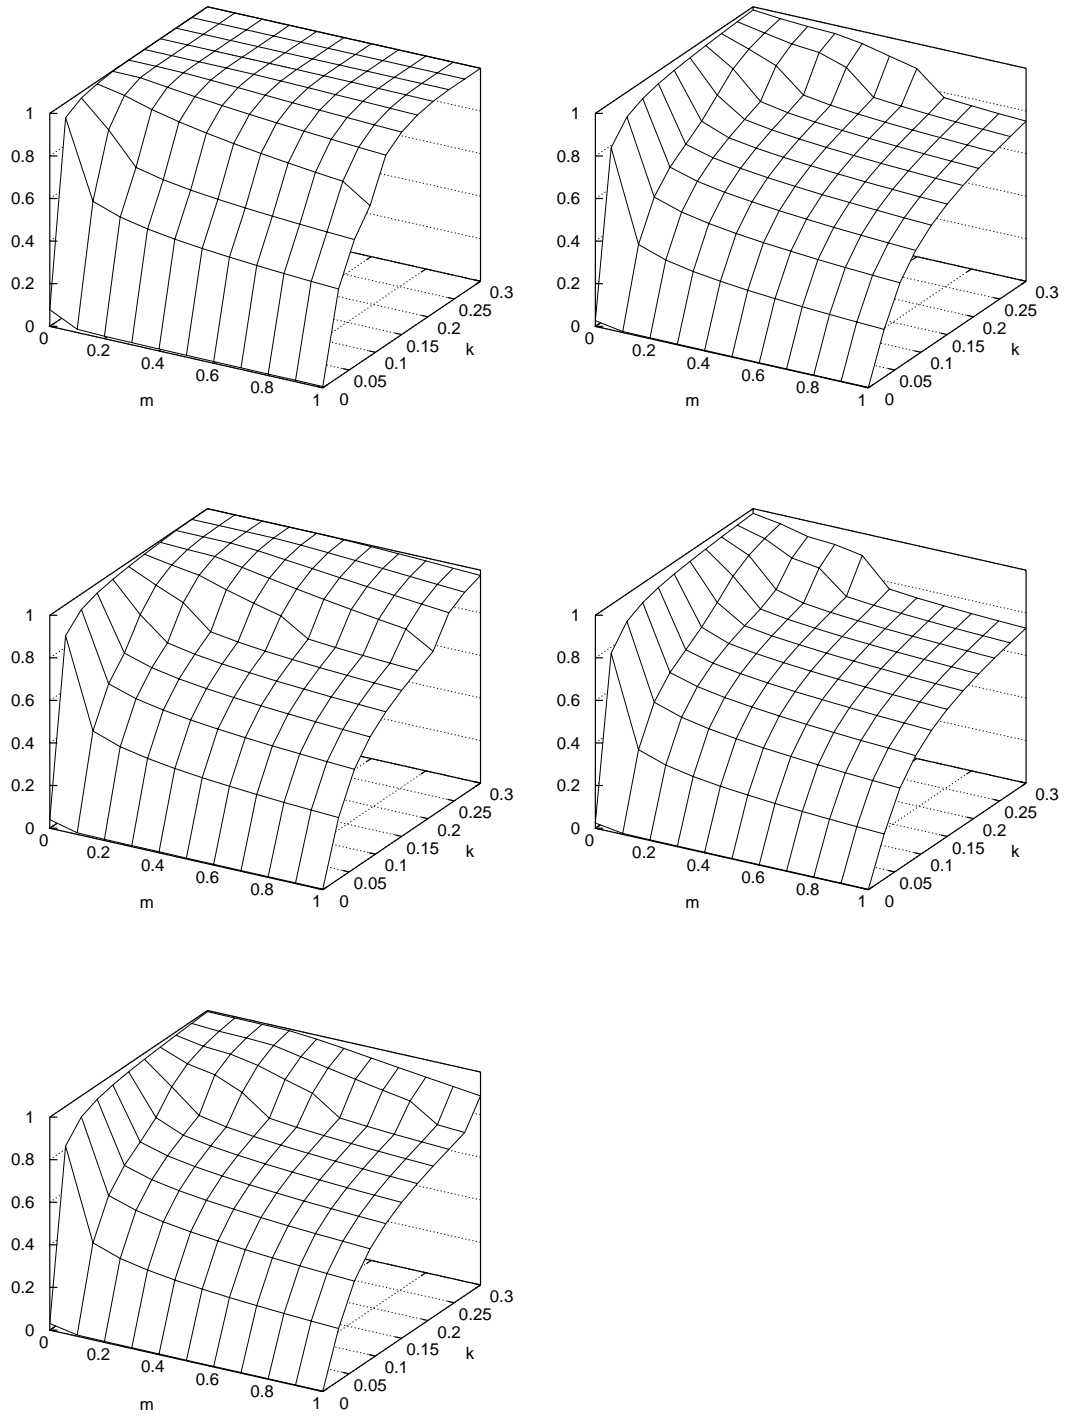


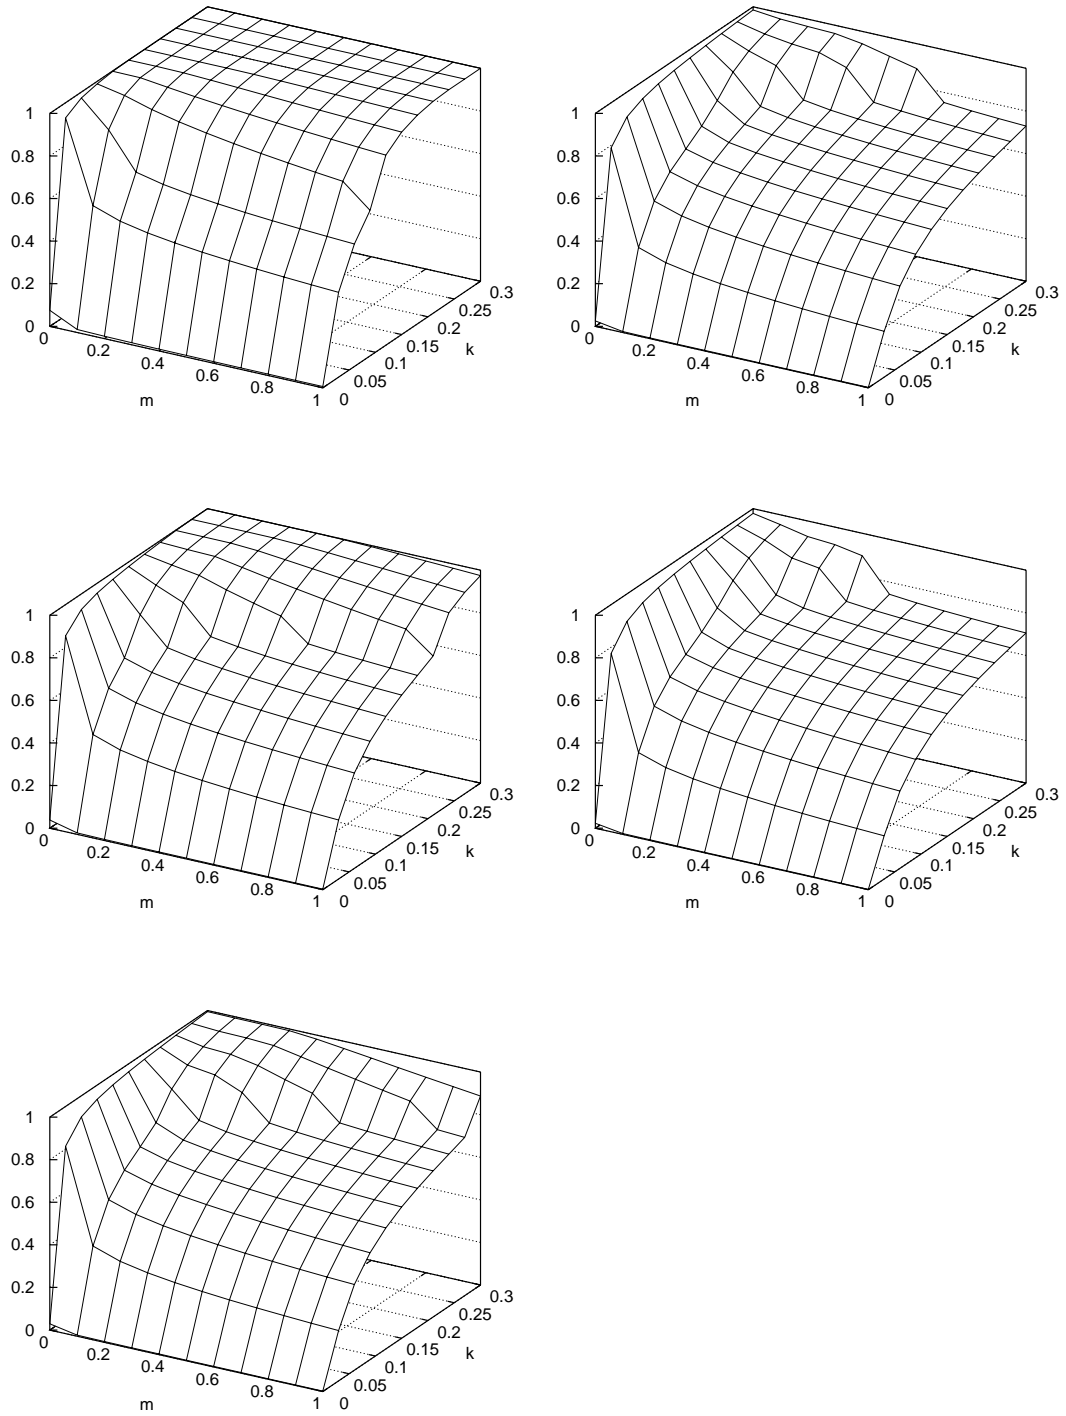
Figure A.9. Sensitivity to r with $q = 0.05$. $r = \{0.05, 0.075, 0.10, 0.125, 0.15\}$, $(a, b, d, q, w) = (0.30, 0.30, 5, 0.05, 50)$ 

Figure A.10. Sensitivity to r with $q = 0.06$.
 $r = \{0.05, 0.075, 0.10, 0.125, 0.15\}$, $(a, b, d, q, w) = (0.30, 0.30, 5, 0.06, 50)$

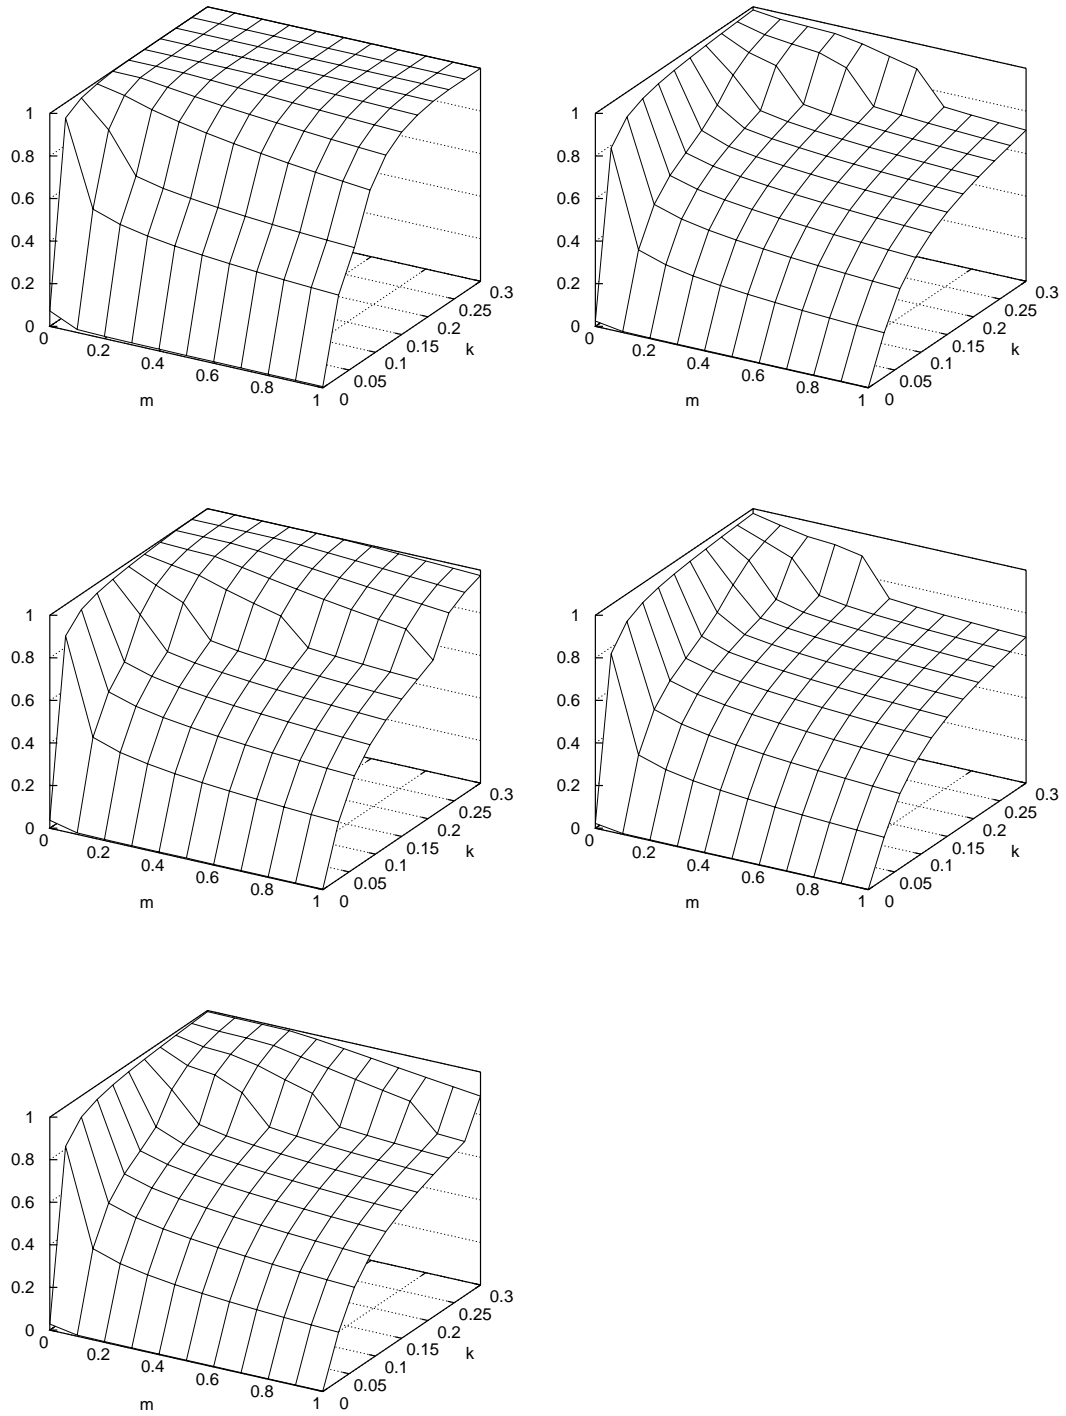
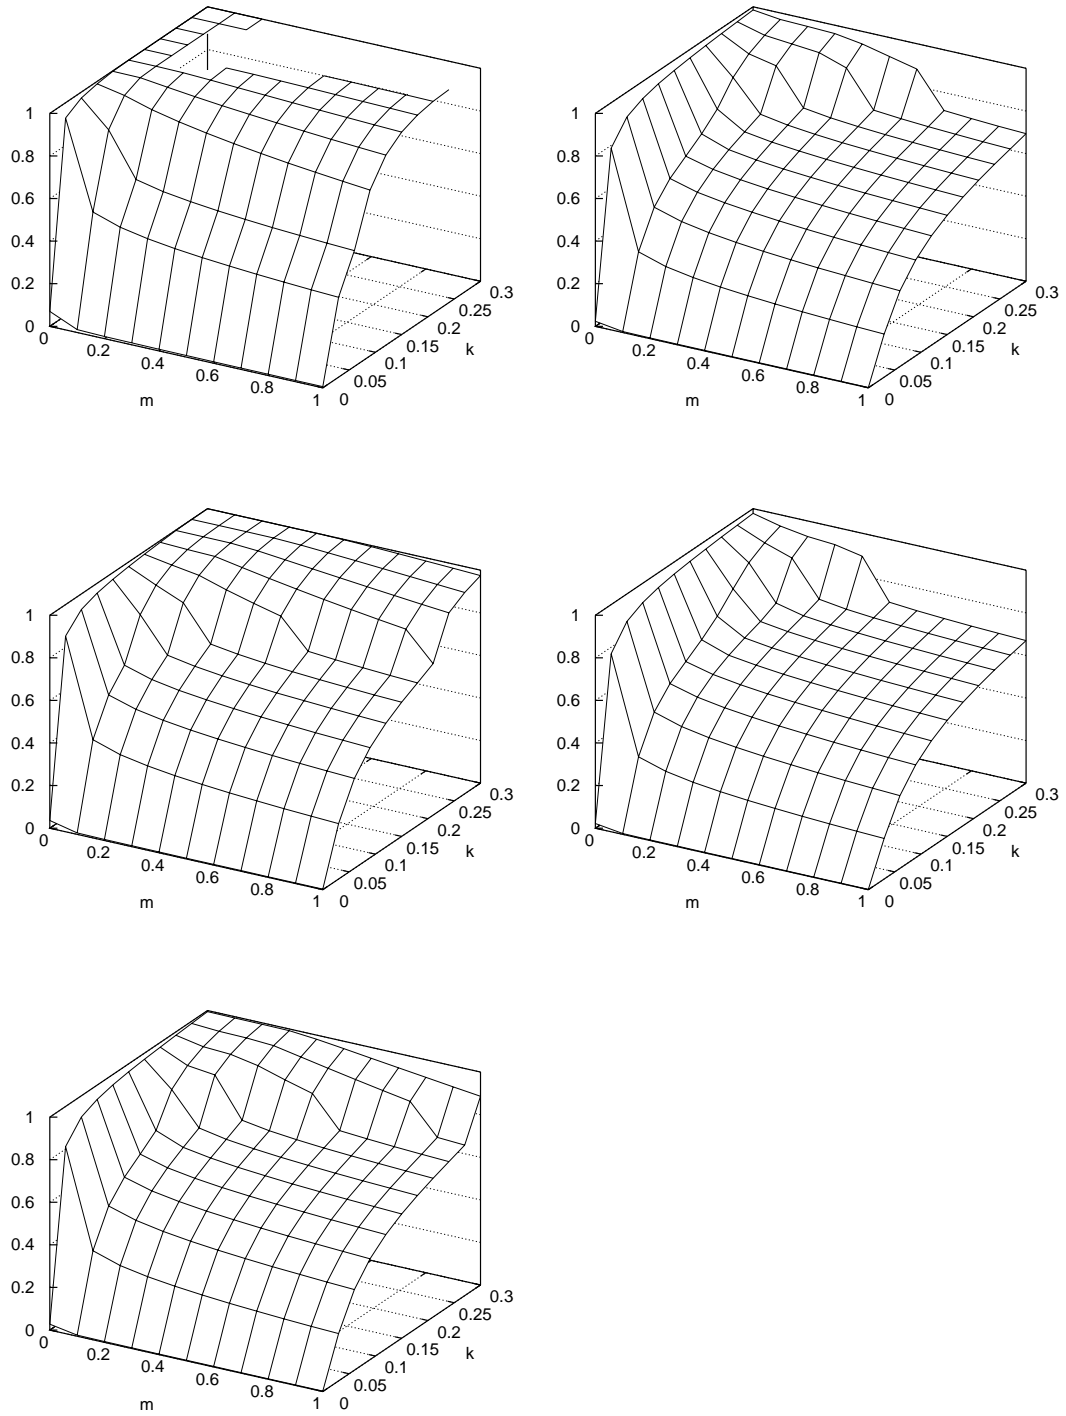


Figure A.11. Sensitivity to r with $q = 0.07$.
 $r = \{0.05, 0.075, 0.10, 0.125, 0.15\}$, $(a, b, d, q, w) = (0.30, 0.30, 5, 0.07, 50)$



APPENDIX B

Behavioral Dependence

The data used in this appendix is a subsample of the Atherosclerosis Risk In Communities (ARIC) data, cleaned for missing data, from a single time period, so that independence between the observations is reasonably assumed. The data consists of 10,309 individuals who choose to exercise (E_1) or not (E_0), consume a low (F_1) or high (F_0) fat diet, and smoke (S_0) or not (S_1). Table B.1 provides a summary table of the data.

Table B.1

Summary of ARIC Behavioral Data. $pr(E_1) = .46$, $pr(F_1) = .24$, $pr(S_1) = .79$

	E_0		E_1	
	S_0	S_1	S_0	S_1
F_0	1054	3323	641	2817
F_1	264	913	207	1090

The first tests performed were χ^2 tests for independence. For $n \times n$ contingency tables, the test statistic

$$\chi^2 = \sum_{i,j} (O_{ij} - \hat{E}_{ij})^2 / \hat{E}_{ij},$$

where O_{ij} is the observed frequency of cell ij , and \hat{E}_{ij} is the expected frequency of cell ij based on the marginal probabilities of i and j , has the χ^2 distribution with 1 degree of freedom. Low p-values would lead to reject the null hypothesis of independence. Independence was tested for between all three pairs of behaviors at three levels each (3rd variable = 0, 3rd variable = 1, combined across 3rd variable). These tests are performed

in R (where the Yates' continuity correction is automatically performed) and verified via spreadsheet. Table B.2 details the results.

Table B.2
 χ^2 Tests for Independence

Test Pair	3rd Level	χ^2	p-value
FS	0	1.3064	0.2530
FS	1	4.1001	0.04288
FS	C	7.4779	0.006246
ES	0	34.6925	3.861e-09
ES	1	16.3409	5.291e-05
ES	C	53.3189	2.836e-13
EF	0	5.5629	0.01834
EF	1	43.7788	3.677e-11
EF	C	51.6666	6.577e-13

The relationship between fat and smoking is the only pair for which the null hypothesis of independence is not easily rejected. Two more tests based on the odds-ratio, the Cochran-Mantel-Haenszel (CMH) test, and the Mantel-Haenszel (MH) test, were performed to check for conditional independence. In data where there is more than two variables, these tests check for conditional independence (CMH) and the strength of association (MH). The procedure `cmh.test()` in R computes the CMH test and MH stat, summarized in Table B.3.

Here again, there is some evidence of conditional independence between diet and smoking, as the null hypothesis fails to be rejected at the 0.02 significance level.

Table B.3
CMH Tests for Independence

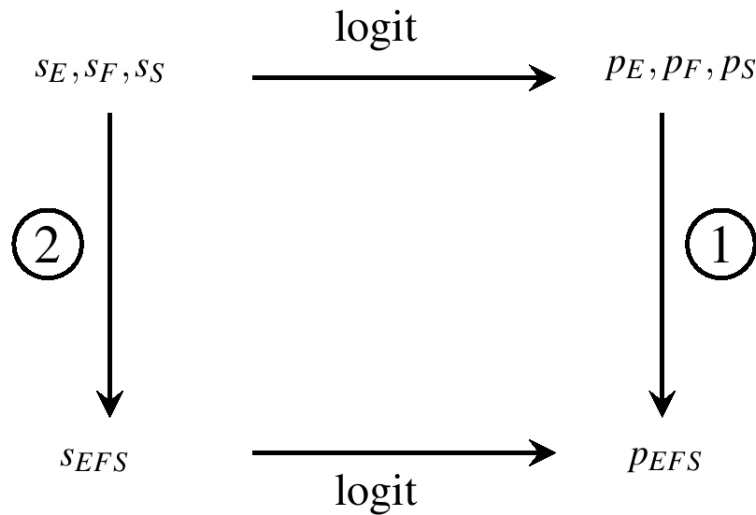
Test Pair	CMH stat	p-value	MH OR	pooled OR	3rd=0 OR	3rd=1 OR
FS	5.126	0.024	1.141	1.174	1.097	1.198
ES	51.133	0.000	1.422	1.434	1.394	1.523
EF	49.417	0.000	1.385	1.395	1.289	1.408

Models of Dependence for Prediction

Each sub-behavior is modeled as a 0-1 variable with 1 representing the healthy choice. The purpose of modeling dependence is to predict probabilities of each type of combined behavior (e.g. exercising, high-fat diet, no smoking) from just three attractions, one for each sub-behavior. Using only three inputs keeps the state space manageable for solving the MDP iteratively. The relationship between attractions and probabilities is computed through the logit rule, where given a set of actions, \mathcal{A} , and attractions s_a , $a \in \mathcal{A}$, the probability of choosing an action is given by

$$pr(a) = \frac{e^{\lambda \cdot s_a}}{\sum_{a' \in \mathcal{A}} e^{\lambda \cdot s_{a'}}}$$

Using the logit rule to translates attractions into probabilities, there are two ways which dependence could be incorporated: through the probabilities directly, or through with the attractions. This concept can be seen in Figure B.1.

Figure B.1. Dependence Pathways

Before modeling dependence, notice that assuming independence between the sub-behaviors implies that $E_i F_j = E_i + F_j$, that is, the interaction ② is additive. Since this function makes the following chain of equalities true.

$$pr(E_i F_j) = \frac{e^{\lambda \cdot s_{E_i F_j}}}{e^{\lambda \cdot s_{E_0 F_0}} + e^{\lambda \cdot s_{E_0 F_1}} + e^{\lambda \cdot s_{E_1 F_0}} + e^{\lambda \cdot s_{E_1 F_1}}} \quad (\text{B.1})$$

$$= \frac{e^{\lambda \cdot s_{E_i}} e^{\lambda \cdot s_{F_j}}}{(e^{\lambda \cdot s_{E_0}} + e^{\lambda \cdot s_{E_1}})(e^{\lambda \cdot s_{F_0}} + e^{\lambda \cdot s_{F_1}})} \quad (\text{B.2})$$

$$= pr(E_i) \cdot pr(F_j) \quad (\text{B.3})$$

Conditional Independence - Additive Dependence (CIAD)

The first pathway to model dependence is pathway ①. Since the tests have shown that independence does not hold, the three marginal probabilities cannot simply be multiplied

together to get a joint probability. However, the tests have shown evidence that conditional on exercise, fat and smoking are independent.

$$pr(E_i F_j S_k) = pr(E_i) \cdot pr(F_j S_k | E_i) = pr(E_i) \cdot pr(F_j | E_i) \cdot pr(S_k | E_i) \quad (\text{B.4})$$

The question then becomes how to model $pr(F|E) \neq pr(F)$. Note also that care must be taken to ensure the probabilities are always non-negative, and sum to 1 when appropriate. The first approach introduced is an additive effects model.

$$pr(F_1 | E_1) = pr(F_1) + \alpha_{F_1 | E_1}$$

$$pr(F_0 | E_1) = pr(F_0) - \alpha_{F_1 | E_1}$$

$$pr(F_1 | E_0) = pr(F_1) + \alpha_{F_1 | E_0}$$

$$pr(F_0 | E_0) = pr(F_0) - \alpha_{F_1 | E_0}$$

$$pr(S_1 | E_1) = pr(S_1) + \alpha_{S_1 | E_1}$$

$$pr(S_0 | E_1) = pr(S_0) - \alpha_{S_1 | E_1}$$

$$pr(S_1 | E_0) = pr(S_1) + \alpha_{S_1 | E_0}$$

$$pr(S_0 | E_0) = pr(S_0) - \alpha_{S_1 | E_0}$$

All the α values here are easily estimated from the data. For example, $\hat{\alpha}_{F_1 | E_1} = \hat{pr}(F_1 | E_1) - \hat{pr}(F_1)$. Performing these computations leads to $\hat{\alpha}_{F_1 | E_1} = 0.0328$, $\hat{\alpha}_{S_1 | E_1} = 0.0318$, $\hat{\alpha}_{F_1 | E_0} = -0.0281$, $\hat{\alpha}_{S_1 | E_0} = -0.0272$. It is interesting that $\hat{\alpha}_{F_1 | E_1} \approx \hat{\alpha}_{S_1 | E_1} \approx -\hat{\alpha}_{F_1 | E_0} \approx -\hat{\alpha}_{S_1 | E_0} \approx .03$. The only concern about the meaningfulness of the probabilities would be if $pr(E_i), pr(S_j) \notin$

[.03, .97], which would lead to negative and higher than 1 probabilities. Based on these α values, the conditional independence - additive dependence model can be used to predict occurrences of each type of behavior. Evidence for the goodness-of-fit of this model is presented after introducing a conditional independence - multiplicative dependence model.

Conditional Independence - Multiplicative Dependence (CIMD)

Another approach for modeling $pr(F|E) \neq pr(F)$ would be a multiplicative approach, $pr(F|E) = \beta \cdot pr(F)$. In order to do this, extra care must be taken to ensure the meaningfulness of the probabilities. To tackle this problem, it can be observed from the data and from the additive approach that when $i \neq j$, $pr(F_i|E_j) < pr(F_i)$ and similarly for smoking. This would result in a value of $\beta < 1$. This insight proves useful since a probability multiplied by a scalar less than 1 is still in $[0, 1]$. Then for the other behavior, $pr(F_j|E_j) = 1 - pr(F_i|E_j)$. Using this model, there are 4 β values to estimate

$$pr(F_1|E_1) = 1 - pr(F_0|E_1)$$

$$pr(F_0|E_1) = \beta_{F_0|E_1} \cdot pr(F_0)$$

$$pr(F_1|E_0) = \beta_{F_1|E_0} \cdot pr(F_1)$$

$$pr(F_0|E_0) = 1 - pr(F_1|E_0)$$

$$pr(S_1|E_1) = 1 - pr(S_0|E_1)$$

$$pr(S_0|E_1) = \beta_{S_0|E_1} \cdot pr(S_0)$$

$$pr(S_1|E_0) = \beta_{S_1|E_0} \cdot pr(S_1)$$

$$pr(S_0|E_0) = 1 - pr(S_1|E_0)$$

Again, all the β values are estimable from the data, resulting in $\hat{\beta}_{F_0|E_1} = .96$, $\hat{\beta}_{F_1|E_0} = .88$, $\hat{\beta}_{S_0|E_1} = .85$, and $\hat{\beta}_{S_1|E_0} = .97$. This multiplicative model is guaranteed to produce meaningful probabilities. Now both the CIAD and CIMD models are used to predict outcomes, judging both versus the real data, with independence (IND) as a benchmark, the results are reported in Table B.4, where the sum of squared deviations (SSD) measures differences from observed counts.

It appears each dependence model predicts the data much better than pure independence. The additive approach fairs slightly better, and is appealing in that all four α values are the same, but has a small concern of inconsistent probabilities. The 10,309 observations composing the full data set are split to form a training data set ($\approx 80\%$) and a testing set ($\approx 20\%$). The training set was used to estimate the α 's and β 's, and the testing set

Table B.4
Comparing Conditional Independence Models

Behavior	Data	IND	CIAD	CIMD
000	1054	888	1055	1026
001	3323	3342	3342	3365
010	264	281	281	275
011	913	1055	888	901
100	641	757	623	618
101	2817	2847	2839	2842
110	207	239	230	229
111	1090	899	1050	1053
SSD	-	100143	4199	5798

to compute SSD, with independence (IND) as a bench mark. The results are reported in Table B.5.

Table B.5
Model Fitness with Training and Testing Data

Behavior	Data	IND	CIAD	CIMD
000	217	176	203	203
001	701	702	704	703
010	58	53	53	53
011	166	212	183	184
100	113	148	122	123
101	585	591	588	586
110	32	45	43	44
111	233	179	209	209
SSD	-	8242	1317	1365

Having presented two models of dependence through pathway ①, attention is now turned to modeling dependence through pathway ②.

Attraction Interactions

In order to work in attraction space, a λ value of 2.95 is used. Attractions $\widehat{s}_{E_i F_j S_k}$ can then be calculated by reversing the logit rule with the attraction $\widehat{s}_{E_0 F_0 S_0}$ normalized to zero. Since independence \Rightarrow pure additive relationship from $s_{E_i}, s_{F_j}, s_{S_k}$ to $s_{E_i F_j S_k}$, then relaxing this assumption by adding coefficients and cross terms to the pure additive form will relax the assumption of independence. The regression equation

$$\widehat{s}_{E_i F_j S_k} = \theta_E \widehat{s}_{E_i} + \theta_F \widehat{s}_{F_j} + \theta_S \widehat{s}_{S_k} + \theta_{EF} \widehat{s}_{E_i} \widehat{s}_{F_j} + \theta_{ES} \widehat{s}_{E_i} \widehat{s}_{S_k} + \theta_{FS} \widehat{s}_{F_j} \widehat{s}_{S_k} + \theta_{EFS} \widehat{s}_{E_i} \widehat{s}_{F_j} \widehat{s}_{S_k}$$

was used to estimate the parameters θ by minimizing the sum of squared errors $\epsilon_{EFS} = \widehat{s}_{EFS} - s_{EFS}$, while constraining some of the θ 's to zero. There are eight equations of this form, and since under the logit rule, attractions are non-unique up to a scalar, normalizing leaves seven equations to estimate θ 's. The results are displayed in Table B.6, where $\theta = (\theta_E, \theta_F, \theta_S, \theta_{EF}, \theta_{ES}, \theta_{FS}, \theta_{EFS})'$. In this table and following tables, where relevant, significance has been denoted by *** for the .01 level, ** for the .05 level, and * for the .10 level. From the table it can be seen that the pure additive model, where $\theta = (1, 1, 1, 0, 0, 0, 0)'$ gives the same results as the independence assumption from Table B.4. When the model is given full flexibility (seven θ values) the model can perfectly match the observed attractions, as expected. The final three columns of Table B.6 show three special cases of the regression: when only coefficients on the linear terms are allowed to be non-zero, when only the three-way interaction coefficient is restricted to zero, and when both the three-way interaction and two-way FS coefficients are restricted to zero.

The results show that giving flexibility on the linear terms alone actually decreases the ability of the model to predict the observed behavior. This result seems counterintuitive since the regression could have selected $\theta = (1, 1, 1, 0, 0, 0, 0)'$ and done better, however

Table B.6
Comparing Attraction Regression Models

Behavior	Data	θ_E	1	3.10	1.73	3.18**	3.38**
		θ_F	1	1.20	1.11***	1.21**	1.17***
		θ_S	1	0.87	0.97***	0.86**	0.89***
		$\theta_{EF} = 0$	0	4.05	0	4.66*	5.01*
		θ_{ES}	0	-4.62	0	-5.14*	-5.44*
		θ_{FS}	0	-.18	0	-.25	0
		θ_{EFS}	0	3.14	0	0	0
000	1054	888		1054	994	1056	1039
001	3323	3342		3323	3601	3287	3379
010	264	281		264	275	261	268
011	913	1055		913	996	926	873
100	641	757		641	753	634	604
101	2817	2847		2817	2727	2857	2906
110	207	239		207	208	210	214
111	1090	899		1090	754	1078	1027
SSD	-	100143		0	221029	3334	18165

the regression seeks to minimize squared errors from the attractions, not the probabilities themselves. Forcing $\theta_{EFS} = 0$ does not hurt the model much, increasing the sum of squared differences by only 3% of the error of the independence model. Forcing additional coefficients to zero hurts the fit of the behavior prediction considerably. The only reasonable fit occurs when forcing θ_{EFS} and θ_{FS} to zero, the final column in Table B.6. While this finding concurs with the analysis showing that diet and smoking are independent given exercise, the dependence model does not predict the data as well as either of the previous models.

This regression approach to predict aggregate attractions given sub-behavior attractions suffers from several flaws. The first is related to the normalization required by the logit rule to compute attractions from probabilities. The attractions are currently normal-

Table B.7
Comparing Attraction Regression Models: Positive Attractions

Behavior	Data	θ_E	1	1.52	1.34*	1.44***	1.66**
		θ_F	1	.98	1.06***	.97***	.93***
		θ_S	1	1.25	1.07***	1.24***	1.21***
		$\theta_{EF} = 0$	0	4.06	0	4.66**	4.28*
		θ_{ES}	0	-5.84	0	-5.31***	-5.64**
		θ_{FS}	0	-.35	0	-.28**	0
		θ_{EFS}	0	3.14	0	0	0
000	1054	888	1054	861	1068	1019	
001	3323	3342	3323	3545	3283	3364	
010	264	281	264	254	261	267	
011	913	1055	913	1045	925	882	
100	641	757	641	695	633	597	
101	2817	2847	2817	2861	2854	2960	
110	207	239	207	205	207	205	
111	1090	899	1090	843	1077	1015	
SSD	-	100143	0	169677	3562	31757	

ized so that attractions to the unhealthy behaviors are 0, and the non-zero attraction to healthy sub-behaviors is computed from the data. This results in $\hat{s}_{E_0} = \hat{s}_{F_0} = \hat{s}_{S_0}$, forcing much of the ‘data’ in the regression equations to zero. This combined with the fact that the attractions to healthy sub-behaviors can be negative or positive, and are multiplied by each other, ruins the interpretation of the coefficients, and the generality of the regression results. This point is to be highlighted further later.

The regressions were repeated using a different normalization scheme. The normalization used was to set attractions to all the less frequently chosen actions to zero. This normalization has the affect of causing all the attractions to be positive. The same set of regressions was performed, and found the results were similar and are reported in Table B.7. Normalizing such that all the attractions are positive seems to improve the signifi-

cance of the regression output. This evidence suggests that the fit of these models is not too sensitive to the normalization chosen.

Consistency Coefficients

Building on the previous results, and seeking to improve on the clarity of the results from the previous regression, regressions were also run of the form

$$\widehat{s}_{E_i F_j S_k} = \theta_E \widehat{s}_{E_i} + \theta_F \widehat{s}_{F_j} + \theta_S \widehat{s}_{S_k} + \theta_{E=F} \mathbf{1}_{i=j} + \theta_{E \neq F} \mathbf{1}_{i \neq j} + \theta_{E=S} \mathbf{1}_{i=k} + \theta_{E \neq S} \mathbf{1}_{i \neq k}$$

to capture the increased likelihood of (un)healthy smoking and diet behaviors when exercise behavior is (un)healthy. The results of this regression are reported in Table B.8, where now $\theta = (\theta_E, \theta_F, \theta_S, \theta_{E=F}, \theta_{E \neq F}, \theta_{E=S}, \theta_{E \neq S})'$.

Since $\theta_{E=F}, \theta_{E \neq F}, \theta_{E=S}, \theta_{E \neq S}$ are a linearly dependent set, one of the coefficients must be zero. The coefficients make sense given the observation about how consistent sub-behaviors are more likely while inconsistent sub-behaviors are less likely. Given that an individual exercises ($i = 1$), the combined diet and smoking behaviors receive the following attraction changes:

While it may seem strange that all of these changes are negative, since the logit rule normalizes a monotonic function of attractions to find probabilities, the relative changes show the impact of exercising on diet and smoking behavior. Again the attractions are re-normalized to make all attractions positive, and the regression is run again.

Interestingly the predicted observations are exactly the same despite the new normalization, again showing the results have little sensitivity to the normalization chosen. The signs and size of the coefficients are different, although the total relative effect is very similar, and the significance of the regression has decreased.

Table B.8
Regression with Sub-behavior Consistency

Behavior	Data	θ_E	1.13*
		θ_F	1.03***
		θ_S	1.04***
		θ_{EF}	= -0.07*
		θ_{ES}	-0.12**
		θ_{FS}	0.06*
		θ_{EFS}	0
000	1054		1020
001	3323		3367
010	264		268
011	913		883
100	641		607
101	2817		2918
110	207		214
111	1090		1032
SSD	-		18938

Table B.9
Attraction Changes Given Exercise

FS Behavior	Δ_s
F_1S_1	-0.01
F_1S_0	-0.06
F_0S_1	-0.07
F_0S_0	-0.12

Discussion

In deciding whether to model dependence through pathway ① or ②, there should be some consideration of the ‘physical’ process being modeled. That is, consumers are

Table B.10
Regression with Sub-behavior Consistency: Positive Attractions

Behavior	Data	θ_E	θ_F	θ_S	θ_{EF}	θ_{ES}	θ_{FS}	θ_{EFS}
		1.14*	1.03***	1.04***	= 0.02	-0.04	0.06*	0
000	1054				1020			
001	3323				3367			
010	264				268			
011	913				883			
100	641				607			
101	2817				2918			
110	207				214			
111	1090				1032			
SSD	-				18938			

attracted to varying degrees to a finite set of sub-behaviors. These behaviors interact in a dependent fashion. Does it seem more plausible that attractions to joint-behaviors are adjusted according to the consistency of their composing sub-behaviors, or that consumers form probabilities of each sub-behavior, and then adjust these probabilities based on the outcome of a ‘first-move’ sub-behavior. For me, both alternatives are plausible. The first possibility, ②, is a sort of premeditated dependence, whereby the consumer thinks ahead and says “Well, if I exercise, there is no point in smoking.” or “If I’m going to eat right, I might as well exercise too.” The second type of dependence ①, is based on the reality that each sub-behavior is not chosen simultaneously, and that based on the sequential nature of the decisions, there is some preference for consistency with regards to the healthiness of behaviors. That is, the consumer thinks “Now that I’ve exercised, I might as well not

waste my efforts by eating unhealthily”. After considering the plausibility of each model of dependence, and the ability to characterize the data, pathway ① and the CIAD model is implemented.

Estimating Transition Probabilities

The large number of transition probabilities to be estimated requires use of a parametric model to increase the power of the data. For this purpose, a multinomial logit model is used to estimate the transition probabilities for consumer health-health-action tuples. Letting T denote the set of possible transitions, the model calculates the probability of a given transition j by

$$p(j|x, a; \mu, \beta) = \frac{e^{\mu_{x,j}^{BP} + \mu_{x,j}^{HDL} + \mu_{x,j}^{TC} + \beta_j^{EXR} a^{EXR} + \beta_j^{FAT} a^{FAT} + \beta_j^{SMK} a^{SMK}}}{\sum_{i \in T} e^{\mu_{x,i}^{BP} + \mu_{x,i}^{HDL} + \mu_{x,i}^{TC} + \beta_i^{EXR} a^{EXR} + \beta_i^{FAT} a^{FAT} + \beta_i^{SMK} a^{SMK}}}. \quad (\text{B.5})$$

To reduce the number of unknown parameters to be estimated, the set of possible health state transitions modeled is limited to those where no sub-state increases or decreases by more than one category. The notation $\{\uparrow, \downarrow, \rightarrow\}$ is used to designate the increase, decrease, or constant transition of a health sub-state category. Considering the size of the categories and the length of one time period, this assumption limiting the observable transitions is violated in only the most extreme cases. The β parameters capture the influence of the consumer’s behavior on state transitions including the interaction of sub-state transitions. The μ parameters capture the trend of each sub-health dimension to transition up (becoming less healthy), down (becoming more healthy), or staying the same, given the level of the sub-health state. The model is prevented from predicting transitions outside the state space (e.g. a HDL state decrease from current state HDL-1) by setting the $\mu_{0,\downarrow}^i$ parameters for $i \in \{BP, HDL, TC\}$ equal to $-M$, a large negative number. The same is done for

$\mu_{5,\uparrow}^{TC}$ and $\mu_{4,\uparrow}^j$ for $j \in \{BP, HDL\}$. This model results in a total of 114 parameters¹ to be estimated. The parameters are estimated by maximizing the log-likelihood function of the observed transitions Z .

$$\hat{\mu}, \hat{\beta} = \max_{\mu, \beta} LL(\mu, \beta),$$

where $LL(\mu, \beta) = \log \prod_{z \in Z} p(z; \mu, \beta)$ uses (B.5). In order to optimize in such a large dimensional space, simulated annealing, a heuristic statistical optimization approach is used.

Simulated annealing (see Spall (2003)) searches increasingly narrow neighborhoods of a current solution, moving to new solutions probabilistically based on the relative fitness of the new solution. In what follows, let $\theta = (\mu, \beta)$. The updating step is performed by the Metropolis-Hastings algorithm with updating step $\zeta \sim U[-l^{-4}, l^{-4}]$, the local candidate deviation for each dimension of θ during iteration l .

$$\theta_{l+1} = \begin{cases} \theta_l + \zeta & \text{with probability } \rho = e^{\Delta LL/T_l} \wedge 1, \\ \theta_l & \text{otherwise,} \end{cases} \quad (\text{B.6})$$

where $\Delta LL = LL(\theta_l + \zeta) - LL(\theta_l)$. This wide and slowly narrowing local neighborhood is used to avoid getting stuck in local optima near the starting solution. The schedule $T_l = [\log(1 + l)]^{-1}$ is used following the recommendation of Robert and Casella (2010). Ten different starting solutions are used to provide a robustness check for the solution to the maximum likelihood problem.

In order to obtain each starting solution, 100,000 simple Monte Carlo draws are taken from the parameter space, and the best of each set becomes a starting solution. Rather than allow each of the parameters to range across the entire real line, the parametric model and

¹The β parameters total $\#\{\uparrow, \downarrow, \rightarrow\}^{\#\{BP, HDL, TC\}} \cdot \#\{EXR, FAT, SMK\} = 3^3 \cdot 3 = 81$, and the μ parameters total $\#\{\uparrow, \downarrow, \rightarrow\} \cdot [\#\{BP \text{ states}\} + \#\{HDL \text{ states}\} + \#\{TC \text{ states}\}] - \text{fixed } \mu\text{'s} = 3 \cdot (5 + 4 + 4) - 6 = 33$.

knowledge of the effect of the consumer's behaviors can provide intuition for coming up with initial guesses for the parameters.

$$p(\downarrow\downarrow\downarrow | a^{EXR} = 1, \cdot) = p(\downarrow\downarrow\downarrow | a^{EXR} = 0, \cdot) \cdot e^{\beta_{\downarrow\downarrow\downarrow}^{EXR}},$$

therefore, β_j^{EXR} should be positive for healthy transitions that exercising would encourage and negative for unhealthy transitions. The same logic would hold for smoking and eating a low fat diet. Further, $\beta_j^{EXR} = 2$ would mean that ceteris paribus, exercising would create a 7.4-fold increase in the probability of transition j . The following guesses of β are used to guide the Monte Carlo draws

$$\tilde{\beta}^{EXR} = \tilde{\beta}^{FAT} = \tilde{\beta}^{SMK} = \begin{array}{c|ccc|ccc|ccc} & & \rightarrow & & & \uparrow & & & \downarrow & & \\ & \rightarrow & \uparrow & \downarrow & \rightarrow & \uparrow & \downarrow & \rightarrow & \uparrow & \downarrow & \\ \rightarrow & 0 & -.33 & .33 & -.33 & -.66 & 0 & .33 & 0 & .66 & , \\ \uparrow & -.33 & -.66 & 0 & -.66 & -1 & -.33 & 0 & -.33 & .33 & \\ \downarrow & .33 & 0 & .66 & 0 & -.33 & .33 & .66 & .33 & 1 & \end{array}$$

where the labeling of the dimensions as $\Delta BP, \Delta HDL, \Delta TC$ is indifferent. Similarly, initial guesses of $\tilde{\mu}_{x,\rightarrow} = 0, \tilde{\mu}_{x,\uparrow} = .1, \tilde{\mu}_{x,\downarrow} = -.1$ for all sub-health states are based on the nature of health to deteriorate, rather than improve over time. Based on these guesses Monte Carlo draws for starting solutions to the simulated annealing are taken by sampling uniformly from the range centered at the guess, and spanning 4 multiples of the guess. For example, $\beta_{\downarrow\rightarrow\downarrow}^{EXR} \sim U[-.66, 1.98]$. Concerning parameters for which the guess is zero, draws are taken from $U[-1, 1]$. Simulated annealing (B.6) was performed for 100,000 replications on each of the 10 starting solutions to arrive at 10 estimates of the maximum likelihood parameters.

Table B.13
Average Absolute Deviations Between Simulated Annealing Probabilities

	2	3	4	5	6	7	8	9	10
1	0.00	0.00	0.00	0.00	0.00	0.01	0.00	0.00	0.00
2		0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
3			0.00	0.00	0.00	0.01	0.00	0.00	0.00
4				0.00	0.00	0.00	0.00	0.00	0.00
5					0.00	0.01	0.00	0.00	0.00
6						0.00	0.00	0.00	0.00
7							0.00	0.00	0.01
8								0.00	0.00
9									0.00

The probabilities produced from the simulated annealing solutions appear to agree with one another to a high degree as evident from the high correlations and low average absolute deviations. Comparing the parameters from the simulated annealing solutions, more variation between solutions can be found. Part of this variation can be attributed to the fact that the parameters of the multinomial logit model are non-unique in their prediction of probabilities. This characteristic means that correlation may be a more meaningful measure of closeness than absolute deviations. The following tables report the correlations, and maximum and absolute deviations for the sets of parameters from the multinomial logit model.

VITA

Name: Brandon Reed Pope

Address: Industrial and Systems Engineering
Emerging Technology Building
Rm. 4062
3131 TAMU
College Station, TX 77843-3131

Email Address: brandon_pope@tamu.edu

Education: B.S., Mathematics, Abilene Christian University, 2006
M.E., Industrial Engineering, Texas A&M University, 2008
Ph.D., Industrial Engineering, Texas A&M University, 2011