Inference and Decision Models for Regulatory and Business Challenges in Low-Income Countries

by

Michael Francis Beeler

HBSc., University of Toronto (2010)
MASc. University of Toronto (2012)

Submitted to the Sloan School of Management
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Operations Research

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2019

Author .....
Sloan School of Management
August 15, 2019

Certified by .....
Cynthia Barnhart
Chancellor and Ford Professor of Civil and Environmental Engineering
Thesis Supervisor

Certified by .....
David Simchi-Levi
Professor of Civil and Environmental Engineering
Thesis Supervisor

Accepted by .....
Georgia Perakis
William F. Pounds Professor of Management
Co-Director, Operations Research Center

Inference and Decision Models for Regulatory and Business Challenges in Low-Income Countries

by

Michael Francis Beeler

Submitted to the Sloan School of Management
on August 15, 2019, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Operations Research

Abstract

This thesis develops inference and decision models to address challenges of particular relevance in low-income countries (LICs). The areas studied include intelligent tutoring systems (ITS), network infrastructure pricing, and anti-counterfeiting.

The ITS chapter identifies previously unknown and serious limitations to Bayesian Knowledge Tracing and Deep Knowledge Tracing, which are two highly-cited methods designed to aid adaptive educational software. The work on Deep Knowledge Tracing led to new data augmentation methods for training recurrent neural networks to be robust in the face of unseen input sequences. We propose a statistically consistent, efficient, and unbiased alternative inference method for questions engaging one skill at a time.

The network infrastructure pricing chapters examine how to allocate the cost of a future infrastructure network whose structure depends on the price-taking decisions of potential users. In a multi-period setting, strategic joining delay by users typically leads to lower utility. We develop a cost-allocation rule that uses rebates to prevent strategic delay. In the single-period setting, we derive closed-form solutions to the expected value of offering to build a simple 1D network and use the 1D solution to establish a lower-bound estimate for more complex 2D networks.

The anti-counterfeiting chapter investigates the strategic procurement of counterfeits by retailers and the effects of shared retailer reputation on equilibrium procurement decisions using models that are more flexible and tractable than those previously appearing in the literature.

Thesis Supervisor: Cynthia Barnhart
Title: Chancellor and Ford Professor of Civil and Environmental Engineering

Thesis Supervisor: David Simchi-Levi
Title: Professor of Civil and Environmental Engineering

Acknowledgments

I am grateful to the many people who have supported my work and development as a scholar during my many years of graduate school. First and foremost, I wish to thank my thesis advisors, Chancellor Cynthia Barnhart and Professor David Simchi-Levi, who gave me the latitude to pursue such a wide-range of topics, overseas fieldwork, and participate in the MIT social entrepreneurship ecosystem, even when it meant a less efficient path to graduation. The learning and growth I have gained from this diverse and unusual set of experiences has been invaluable. Your flexibility, patience, and wise advice have made all the difference. In addition to research advising, I have especially appreciated the collegiality David deliberately fostered amongst his many students through delightfully catered group seminars and social events over the years. Likewise, Cindy's appointment as Chancellor allowed me to witness and learn from her incredible example as a leader; she has juggled more than I could have imagined, yet always made time for me and brought a positive, encouraging energy to every meeting, no matter the stresses and demands from her own work. My advisors are supported by wonderful administrators, Janet Kerrigan and Shannon McCord, who exemplify the incredible staff at MIT who make the Institute run so well.

My thesis committee members, Dr. Robert Stoner and Professor Tauhid Zaman, also played a defining role in my time at MIT. I am grateful to Dr. Stoner for connecting me to many researchers at the MIT Energy Initiative, and also for his leadership of the MIT Tata Center for Technology and Design, which supported my fieldwork in India, contributing greatly to 2 and 5, as well as to ongoing work on education software and language arts assessment technology that I look forward to sharing in the future. I had the pleasure of being Professor Zaman's teaching assistant in my first summer at MIT, and his brilliant use of humor, creativity and charisma during lectures set a fantastic example for me to follow in my own teaching. I would also like to thank my General Exam committee chair, Professor Asu Ozdaglar, whose course on game theory and feedback during my general exam improved my work in Chapter 3.

The instructors I worked with as a teaching assistant during my PhD were wonderful. I am thankful that they entrusted me with the development and delivery of new teaching materials, which I found very stimulating and valuable to my professional growth. Thank you professors Roy Welsch, Rob Freund, and Arm Farahat.

My work on the education chapter was motivated and improved by collaboration and advice from a wide range of individuals and organizations: the IIT Bombay (IIT-B) Center for Indian Language Technology (Professor Malhar Kulkarni Professor Pushpak Bhattacharyya, Hanumant Redkar, Sayali Khare, Dr. Nilesh Joshi), the IIT-B Tata Center for Technology and Design (Gayathri Thakoor), IIT-B Professor Preethi Jyothi, Arijit Mukherjee, Khushboo Thaker of the University of Pittsburg, LearnLab Summer School organizers at Carnegie Mellon University (Professor Kenneth Koedinger), the Ek Step Foundation (Shankar Maruwada, Deepika Mogilishetty, Jai Bannur, Raghu Tenkayala, and Prasad Sandbhor), Pratham (Anamara Baig and Renu Seth), the Sir Dorabji Tata Trust, and Doosra Dashak Pisangan (Neelu and Salim), Technirmiti (Shrikrishna Parab) and Dr. Priyanka Dhuliya.

The general industry insight shared by staff at Sproxil in Cambridge, MA (Alden Zecha, Danielle Goldschneider and Jennifer Campos), and by PharmaSecure in New Delhi (Nakul Pasricha) was very helpful in my approach taken in the chapter on counterfeit medicine.

I had the pleasure of working with several graduate students and researchers at the MIT Energy Initiative and Engineering Systems Division (now IDSS) who taught me a great deal about energy systems and policy, and/or shared useful data and code templates relating to energy systems research: Professor Ignacio Perez-Arriaga, Professor Mort Webster, Dr. Reja Amatya, Yael Borofsky, Douglas Ellman, Zachary Accuardi, and Patricia Levi.

I am grateful to have received fellowship funding from the United States Department of State through the International Fulbright Science and Technology Award Program, from the MIT Legatum Mastercard Foundation Fellowship, as from the MIT Tata Center for Technology and Design, as well as support and professional development from the staff at each program.

My time at MIT has been all the more enjoyable thanks to the many wonderful classmates that I worked with on course projects, qualifying exam preparation, or who offered feedback on my work: Alex Weinstein, Chiwei Yan, Chong Yang Goh, Christina Epstein, Clark Pixton, Colin Pawloski, Kris Ferreira, Iain Dunning, Jerry Kung, Louis Chen, Miles Lubin, Nataly Youssef, Shujing Wang, Swati Gupta, and Zachary Owen. I also greatly enjoyed serving as a teaching assistant alongside my advisor David Simchi-Levi, as well as Professors Roy Welsch, Rob Freund, and Amr Farahat.

My journey to MIT would not have been possible without the exceptional mentorship and encouragement I received during my Masters program at the University of Toronto under Professor Michael Carter and Professor Dionne Aleman. Mike and Dionne took a gamble by taking me on as their research student in Industrial Engineering: I was Peace and Conflict major with no coding experience and fairly modest mathematical and statistical training. I will be forever grateful for the opportunity they provided me to transition into the world of quantitative problem-solving.

I am especially grateful to the many volunteers and staff who deserve most of the credit for keeping my related non-profits running smoothly (SID, PfID and 1Room) during the busiest periods of my PhD. I hope that our work has put in place the foundation to eventually deploy in Kenya many of the educational technologies and algorithms I have studied during my PhD. A special thanks to Vivian Hui, Kate McCullough, Alexander Girard, Nathan Hernandez, Steven Lehman, Asem Azar, Boaz Ongote, Sharon Owala, Akoko Wilson, Wellington Odhiambo, and Pamela Gero, as well as all of my SID-PfID summer interns, of which there are too many to mention.

Finally, I am thankful to friends and family who supported me during the bumpier periods of graduate school, including but not limited to: Sadia Rafiquddin and Renjie Butalid, Abigail Geer, Dr. Selma Duhovic, Dr. Douglas Staple, Stephanie Currie, Alicia Landry, Jahvari Junior, Carl and Jean Beeler, and my extended family in Nova Scotia.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

This thesis examines three areas of broad public interest, with particular relevance to firms and regulators in low-income countries (LICs): educational software, infrastructure cost allocation, and counterfeiting. The problem areas and methodologies deployed in each chapter are admittedly distinct. This thesis does not put forth general guidelines or methodology for undertaking "operations research" or "management science" problems in LICs. It does, however, weave a common contextual thread through each chapter, which is the need to adapt models and solution techniques to the different behavioral, regulatory, and/or technological constraints in LICs.

Chapter 2 deals with inference challenges in the rapidly growing educational software sector. With falling hardware costs and growing web access, intelligent tutoring systems (ITS) have tremendous potential to help developing countries cope with acute teacher shortages and wide-ranging student learning levels. A key challenge for ITS is to estimate students' latent skill levels quickly, preferably on low-cost client devices, such as tablets and smart phones, since the large majority of households and schools cannot afford personal computers and electricity may be limited or unreliable.

The literature review, Section 2.2, provides a fairly comprehensive overview of the three largest, and rather distinct, methodological traditions for representing and estimating student knowledge, with a brief discussion of their limitations. Section 2.3 and 2.4 provide an in-depth examination of serious and previously unknown limitation of two important methods, Bayesian Knowledge Tracing (BKT) and Deep Knowledge

Tracing (DKT). More specifically, we show that Bayesian Knowledge Tracing (BKT) should not be used to estimate partial skill mastery and can produce worst-case predictive performance asymptotically, as well as an unacceptably high false-positive rate even when deviations from its model assumptions appear very minor. We then show that Deep Knowledge Tracing (DKT), a much-cited BKT alternative free from this limitation, can lose much of its predictive power when there are small departures from the question sequence on which it was trained.

In Section 2.5, we develop a training method that preserves almost all of DKT's performance in the face of unseen question sequences. The training method may be helpful in fitting other types of recurrent neural networks modeling data sequences generated from a small set of latent variables. It involves partitioning the ordered training set sequence into contiguous blocks, and, for each training batch, randomly shuffling the order of questions within each block. We also examine a method that uses DKT to cluster questions by common skill dependence and show previously published results to be non-reproducible and implausible. We find that these problems reveal unfixable errors in the DKT question clustering method.

Although psychometric methods based on item response theory avoid the problems we identified in BKT and DKT, their use of an arbitrary interval scale limits interoperability and interpretability. In Sections 2.6 to 2.8, we develop a novel latent skill model using a [0,1] continuous skill scale, with a view towards interpretability and facilitating information exchange between ITS platforms. We outline a procedure to estimate latent skills from questions of varying difficulty and derive a closed-form solution for a statistically efficient, consistent, and unbiased estimator in the case of questions that engage one skill at a time. Computationally, the method is suitable for use on low-cost client devices.

Chapters 3 and 4 deal with pricing and cost allocation problems in infrastructure networks whose shape and size depends on the joining decisions of price-sensitive potential users. The problem is motivated by micro-grid and electric grid expansion in developing countries. In Chapter 3, we show that the worst-case performance of average cost pricing, marginal cost pricing, and Shapley value pricing is total utility

loss. In a multi-period setting, strategic joining delay by users typically leads to lower utility. We develop a cost-allocation rule that uses rebates to prevent strategic delay.

In Chapter 4 we analyze the setting where regulations require the network builder to extend the network to anyone who accepts its fee proposal. We first suppose users are arranged on a 1D lattice and then characterize the builder's participation decision and pricing decision as a function of the reference price distribution. We show expected profit can be positive even when the expected marginal profit per potential user is negative. We show substantial profit improvements occur, but with lower total utility, if the builder may serve only the optimal subset of price-takers. Finally, we derive a lower bound for expected profit in the 2D setting using the 1D results.

In Chapter 5, we examine public interest problem of a very different flavor: the deliberate procurement of counterfeit goods by retailers, and especially of medicines in developing countries. Our interviews with executives in the pharmaceutical security industry suggest widespread deliberate procurement and sale of counterfeit or grey-market products by retailers, as well as a surprising reluctance by some manufacturers to adopt low-cost product authentication technologies. We develop a game theoretic model of the strategic decisions of manufacturers, retailers, and consumers to gain insight into the type of products, market structures, and consumer population traits that are likely to influence equilibrium counterfeit levels. The brief conclusion in Chapter 6 discusses the most promising directions for future work.

Inquiries relating to this thesis can be sent to beeler [at] alum [dot] mit [dot] edu. User-friendly copies of code used in this thesis will be posted to GitHub by October 30, 2019: https://github.com/MichaelBeeler?tab=projects.

Chapter 2

Knowledge Modeling and Inference

2.1 Introduction

The ongoing growth of digital learning platforms will produce more and more educational data over wider domains and over a wider range of learners. One implication of this general trend is that psychometricians and educational data-mining (EDM) researchers will be able to use datasets that will give them more statistical power to detect and model complex nuances in human learning. A second implication is that the social and economic benefits of improving the science of assessment will grow significantly with the rising use of digital learning tools.

The impact of digital learning tools may be especially pronounced in resource-constrained communities and developing countries, particularly those facing a shortage of qualified educators and overcrowded classrooms. In 2016, the UNESCO Institute of Statistics reported that 70% of countries in Sub-Saharan Africa have an acute shortage of primary teachers, and 90% have an acute shortage of secondary teachers [UNESCO Institute of Statistics, 2016]. With falling hardware costs, widespread phone usage, and growing web access, phone-based or tablet-based intelligent tutoring systems (ITS) have great potential to help families and schools in low-income countries. An important consideration for ITS in these settings is quick computation, possibly on low-cost client devices without reliable internet access.

Effective assessment can greatly improve teaching and learning, whether it be

through intelligent tutoring systems (ITS), or by helping teachers adapt to individual student needs. As early as the 1980s, meta-analyses on the impact of individualized tutoring by experts vs. group-based instruction have found large effects favoring expert tutoring and individualization [Bloom, 1984, Cohen et al., 1982]. The question of how to best conduct assessment is central to the development of quality ITS. The question is also complex because it necessarily leads to other open-ended questions: how to model knowledge, which diagnostic questions to ask, and how to mathematically relate the data to the knowledge model.

For each of these problems a wide range of approaches have been proposed, each with its own trade-offs with respect to granularity, computational tractability, data requirements, ease of interpretation, and suitability for different types of knowledge domains. Section 2.2 provides an overview of these methods and their main limitations. Section 2.3 shows that a popular ITS method, Bayesian Knowledge Tracing (BKT), should not be used to estimate partial skill mastery and can produce worst-case predictive performance asymptotically. In Section 2.4, it is then shown that Deep Knowledge Tracing (DKT), a much-cited BKT alternative free from this limitation, can lose much of its predictive power under small departures from the question sequence on which it was trained. A training method is developed that preserves most of DKT's performance in the face of unseen question sequences. The section also examines a method that uses DKT to cluster questions by common skill dependence and shows previously published results to be not only non-reproducible but also mathematically impossible. These problems are shown to reflect inherent problems with the DKT question clustering method.

Finally, Section 2.6.1 presents a novel latent skill model using a [0,1] continuous skill scale, with a view towards interpretability and facilitating information exchange between ITS platforms. We outline a procedure to estimate latent skills from questions of varying difficulty and derive a closed-form solution for a statistically efficient, consistent, and unbiased estimator in the case of questions that engage one skill at a time. Computationally, the method is suitable for use on low-cost client devices.

2.2 Literature Review

There appear to be three separate disciplinary clusters that deal with quantitative methods for skill assessment. The oldest is the psychometric tradition, whose main methods are item response theory (IRT) and diagnostic classification models (DCM). These methods have been developed primarily by psychologists and typically use logistic functions to relate the probability of a particular response to one or more latent skills and/or question features.

The second cluster, Knowledge Tracing, seems to originate more from computer science and human-computer interaction departments, with a focus on live use in ITS with learners whose skill mastery can change in real time. The third cluster, Knowledge Space Theory (KST), was developed by UC Irvine mathematicians in the 1990s and has been deployed in commercial software by McGraw-Hill. These traditions have developed somewhat independently and with less cross-pollination than might be expected, given the common problems they aim to solve.

This section's goal is not to comprehensively summarize the models of each tradition, but to highlight their salient theoretical differences and practical trade-offs. In particular, the methods differ along the following dimensions: (1) whether knowledge/skill components are treated as binary or continuous; (2) a focus on representing broad latent traits vs. specific identifiable skills, rules, declarative knowledge or procedures; (3) in their estimation methods and heuristics; (4) in representing knowledge states as static or dynamic.

2.2.1 Psychometric Traditions

Item response theory (IRT) is the dominant paradigm in psychometrics for measuring latent cognitive variables, such as skills, attitudes, or personalty traits. Dichotomous IRT represents the probability of a correct response to an "item" (e.g., a test question) with a logistic function that depends on (i) the true value of the respondent's latent variable(s), (ii) the difficulty of the item, and (iii) other item-specific parameters that can control the center, scale, and asymptotes of the logistic function. Variants of

IRT that model the probability of more than two discrete responses are known as polytomous models and they typically use multinomial logistic functions.

Notation in the IRT literature tends to be consistent. By convention, θ_ℓ represents the learner's latent ability, a_i represents an optional horizontal scaling factor (a.k.a. discrimination parameter), b_i represents item difficulty, and g_i is an optional pseudo-guessing probability, which can be fixed *a priori*. In dichotomous IRT, the probability $p_{i\ell}$ of a correct response by learner ℓ to item i often takes the form of a three-parameter logistic (3PL) model [Birnbaum, 1968]:

$$p_{i\ell} = g_i + \frac{1 - g_i}{e^{-a_i(\theta_\ell - b_i)}}. \tag{2.1}$$

An important special case of 3PL is obtained when $g_i = 0$, producing a more mathematically convenient two-parameter logistic model (2PL) with a single logistic function term [Birnbaum, 1957]:

$$ln(\frac{p_{i\ell}}{1 - p_{i\ell}}) = a_i(\theta_\ell - b_i). \tag{2.2}$$

Another important special case is the one-parameter logistic model (1PL), also known as the Rasch model [Rasch, 1960]. In 1PL, $g_i = 0$ and $a_i = 1$ or $a_i \approx 1.7$; when $a_i \approx 1.7$, the logistic function approximates the shape of cumulative density function of a normal distribution with variance 1 and mean b_i, which can be helpful if one wishes to ultimately interpret the units of θ and b as z-scores . To do so, one can fix the average θ across a population of learners, or the average b across a set of items, to be 0 during the model fitting process since $p_{i\ell}$ is invariant to horizontal translations on θ and b's shared dimension. In other words, only the difference between θ_ℓ and b_i matters, not their individual values.

Although IRT models involve a logistic function, they differ substantially from logistic regression because the predictor variables are unobserved. The ultimate purpose of the IRT models, and of most knowledge-modeling approaches in general, is not to correctly predict test performance, but to sensibly represent and correctly estimate the latent variables. Predictive performance should be seen as a means of

evaluating the quality of latent variable models and estimates. When either b or θ are "known", traditional maximum likelihood estimation generates a statistically consistent estimator for the unknown counterpart [Birnbaum, 1968].

However, when both b and θ are unknown, IRT models are typically fit using conditional maximum likelihood estimation (CMLE) [Anderson, 1972] or marginal maximum likelihood estimation (MMLE) [Bock and Aitkin, 1981] as a sub-routine within an expectation-maximization (EM) algorithm. These IRT fitting processes are computationally intensive and are typically implemented using specialized packages [Baker and Kim, 2004]. For example, using CMLE, which is simpler than MMLE, a set of values for one set of parameters is fixed so that a maximum likelihood estimate, conditional on the fixed parameters, can be obtained for the free-ranging counterpart. Under EM, the CMLE procedure is then performed repeatedly while alternating the type of parameters that are fixed, and the process stops after reaching a maximum number of iterations or a gradient-based stopping criterion. For logistic functions with latent variables, EM offers no guarantees of convergence or global optimality. As a result of the complexity of the fitting process, these models are not well-suited for fast computation on tablets or phones, nor for applications requiring near-instant results.

Another challenge for IRT is that the latent item and ability parameters live on an interval scale, not a ratio scale, and can therefore be shifted by an arbitrarily large constant. One option to improve the interpretability of the scale is to center either the mean question difficulty at zero, or the mean skill level at zero. Even when this happens, however, the fitted parameters for new learners or new questions must be interpreted with respect to the parameter estimates obtained from a particular reference population that was given a particular set of questions.

Latent Class Models

Success on some items may depend on multiple independent skills. Such situations motivated the development of logistic models with multiple skills and/or components of item difficulty.

Diagnostic Classification Models (DCM), for example, assume that the learner belongs to one of a discrete number of classes. Each class is defined by a certain combination of binary latent variables, which represent whether a particular skill is possessed or not. For the benefit of any reader wishing to engage the DCM literature, it is important to mention that the family of models described here as DCM also goes by many other names. These names include: skills assessment models, cognitive diagnosis models, cognitive psychometric models, latent response models, restricted latent class models, multiple classification models, structured located latent class models, and structured item response theory.

A very general DCM that effectively encompasses simpler DCM models is the Log-linear cognitive diagnosis model (LCDM) [Henson et al., 2009]. In LDCM, $\mathbf{q}_i$ is a vector of indicators specifying which skills the item requires, $\boldsymbol{\alpha}_\ell$ is a vector of indicators for whether the learner possesses the skill, $\mathbf{w}_i$ is a vector of weights controlling the importance of each skill for success on the item, and $\mathbf{h}(\alpha_\ell, \mathbf{q_i})$ is a vector to vector function that produces main effects or interaction terms, of any degree, for the skills that are engaged by the item. As before, b_i is item difficulty, though π_i is frequently used in the DCM literature. To simplify notation, we relate these items to $p_{i\ell}$ through z in the logistic function $p_{i,\ell} = \frac{1}{1+e^{-z}}$:

$$z = \mathbf{w}_i^T \mathbf{h}(\boldsymbol{\alpha}_\ell, \mathbf{q}_i) - b_i. \tag{2.3}$$

Certain functional forms of LCDM, in particular the Deterministic Input Noisy AND gate (DINA) model [de la Torre, 2011], have deep mathematical connections with knowledge space theory (KST) [Heller et al., 2015], which is discussed in Section 2.2.3. The main limitations of DCM are that they do not model partial skill mastery, that their complexity increases both the risk of over-fitting and the size of datasets needed to fit properly, and that they can be even more computationally difficult to estimate than continuous-parameter IRT models.

Multidimensional Item Response Theory (MIRT)

MIRT models use multiple continuous latent trait variables instead of categorical latent variables, and are usually more parsimonious and easier to estimate than DCM [Embretson, 2013].

A basic example of a model in the MIRT family is the multi-dimensional logistic model [Reckase and McKinley, 1991]. In this model there are M trait (skill) dimensions, a_{im} is the discrimination parameter for item i on dimension m, $\theta_{\ell m}$ is learner's latent ability trait on dimension m, and π_i is unidimensional item difficulty. Defining z as the argument in the logistic function, the simplest MIRT model takes the following form:

$$z = \sum_m a_{im}\theta_{\ell m} - b_i. \tag{2.4}$$

This model is compensatory rather than conjunctive; it assumes that not all skills are required, and that a stronger value in one skill can compensate for lower values in other skills.

The Linear Logistic Test Model (LLTM), by contrast, typically uses one dimension for skill and treats item difficulty as a linear function of M item attributes [Fischer, 1972]. Let q_{im} be an indicator variable (typically given *a priori*) for whether item i has difficulty attribute m, and let η_m be the weight controlling the attribute's contribution towards component difficulty. Let η_0 represent baseline difficulty. The basic LLTM model is the same as the 1PL or 2PL Rasch models except for the following substitution:

$$b_i = \sum_m^M \eta_m q_{im} + \eta_0. \tag{2.5}$$

Multicomponent Latent Trait Model (MLTM)

The mulicomponent latent trait model (MLTM) was developed by Susan Embretson [Whitely, 1980, Embretson, 2013, Embretson and Yang, 2015] to reflect the fact that complex tasks may involve sequentially-dependent problem-solving stages. These stages could be modeled as components of the overall task, with the probability of task completion being the product of the probability of success on each component. The

component probabilities are estimated by a latent trait model. For MLTM notation, let K_i be the set of indices for components of item i, $\theta_{k\ell}$ be the learner's ability in component k, and b_{ik} be the item's difficulty in component k, then the basic MLTM model takes the following form:

$$p_{i\ell} = \prod_{k \in K_i} \frac{1}{1 + exp(-(\theta_{k\ell} - b_{ik})}. \tag{2.6}$$

This is an example of a conjunctive model. Conjunctive models such as MLTM are at a greater risk of being non-estimable because they require certain types of variations in item design and respondent skill in the datasets they fit. Whereas a correct response suggests that all components for an item are "known", an incorrect response suggests that one or more components is deficient. The only way to pinpoint a deficient component is if that component is also present in other test items, either by itself or with a sufficiently wide variety of other components.

General Component Latent Trait Model (GLTM)

The General Component Latent Trait Model [Embretson, 1984] is a generalization of LLTM in which there are M_k additive/compensatory attributes within each conjunctive component that affect the component's difficulty. Let q_{ikm} be an indicator variable (typically specified *a priori*) for whether component k of item i has difficulty attribute m, and let η_{km} be the weight controling the attribute's contribution towards component difficulty. Let η_0 represent baseline difficulty. Then GLTM is the same as MLTM except for the following substitution

$$b_{ik} = \sum_{m}^{M_k} \eta_m q_{ikm} + \eta_0 \quad \forall k \in K_i. \tag{2.7}$$

It is important that the data used to fit MLTM and GLTM contain various combinations of skill strength. Although two skills may be conceptually independent in the sense that there is no pre-requisite relationship between them, if they happen to be taught very close in time in a particular curriculum, and thus highly correlated,

multicollinearity can occur. Furthermore, latent variables that are correlated because of local curriculum idiosyncrasies could give the impression that the skills or components can be merged without losing predictive power. However, such a skill merger would prove counterproductive if the model is used for learners whose curriculum decouples the two skills that are otherwise unrelated (i.e., by teaching the skills far apart in time).

Guessing and Slippage in IRT

Pseudo-guessing was modeled in the 3PL IRT model with the parameter g_i, which is the limiting probability of success as the latent trait value approaches negative infinity. Likewise, one may wish to have a limiting probability of success that is slightly less than 1 to account for careless errors that may persist even among experts. Let s represent the probability of slippage (a.k.a. a careless error) by a learner who has mastered the question. Then the IRT models considered so far could be modified as follows:

$$p_{i\ell} = (s_i - g_i)P(\text{success given by IRT model}) + g_i. \tag{2.8}$$

This is known as the 4PL IRT model, first proposed by Barton and Lord [Barton and Lord, 1981]. The 4PL IRT model has been shown to result in more robust and statistically efficient estimation using real ITS datasets [Liao et al., 2012, Yen et al., 2012].

The use of guessing and slippage parameters is essential to Knowledge Space Theory and Bayesian Knowledge Tracing, and is also used in an original model presented in Section 2.6.1.

2.2.2 Knowledge Tracing

Knowledge tracing (KT) deals with the task of modeling a person's knowledge or skill level over a short time interval during which improvement may occur as a result of their interaction with an ITS. Most KT models have been developed by computer scientists with an interest in psychology and human-computer interactions, rather

than by psychologists. KT models estimate the probability of item success using the learner's performance on past items.

KT models are developed for live use in ITS, which typically engage a relatively limited and related set of skills per session rather than providing a broad skill diagnosis across a large curriculum. This section of the literature review introduces three types of knowledge tracing: (1) dynamic IRT-based models; (2) Bayesian Knowledge Tracing; and (3) a recurrent neural network model called Deep Knowledge Tracing (DKT). In most empirical KT studies, with the exception of DKT, the educational software targets skills that are intended to be near to the student's knowledge frontier. As a result, the formulas used for several KT methods implicitly assume that a learner could not possibly remain stuck and will eventually master the target material.

Dynamic IRT-based models

The first family of KT models considered all use a logistic function with latent variables, much like classic IRT, but include additional variables that depend on the learner's response history up to that point in time. They use a matrix Q of indicator variables to represent which test items possess certain well-defined features that are thought to affect difficulty. For example, a feature could be a discrete skill requirement, such as how to calculate the area of a circle using its radius. In the paragraphs that follow, the Q matrix is treated as known, though identifying Q through domain expertise or data is a challenge in itself. All of these models decompose item difficulty parameter b_i into a sum of skill-difficulty parameters β_k for skills/features relevant to the question, as indicated by Q.

The most basic dynamic IRT-based model for KT is the Additive Factors Model (AFM) [Cen et al., 2006], proposed in 2006. AFM presents a latent variable logistic regression model in which the learner still has a single latent variable for their overall skill. Just like LLTM, which was proposed in 1973, AFM decomposes the item difficulty parameter into a weighted sum of parameters for the item's attributes (skills tested). Unlike LLTM, the learner's probability of success on a particular item is automatically incremented as the learner practices identical or similar items that involve

skills relevant to the item. The probability that learner ℓ answers item i correctly is modeled as

$$ln(\frac{p_{i\ell}}{1-p_{i\ell}}) = \theta_\ell + \sum_k \beta_k Q_{ik} + \sum_k Q_{ik}\gamma_k N_{k\ell}, \tag{2.9}$$

where Q_{ik} is an indicator variable that indicates whether item i engages skill k, $N_{k\ell}$ is the number of practice opportunities completed by learner ℓ on skill k, and γ_k is the effect of additional practice opportunities on the odds ratio of success for questions that involve skill k. Generally, β_k are negative; a positive value would mean that attribute k makes the item easier.

Observe that AFM approaches 1 as practice opportunities increase, regardless of performance, so this model will be statistically inconsistent and upwardly biased if the student is not learning, or learning at a slower rate than the general population.

A minor modification to AFM, called the Performance Factors Model (PFM) [Cen et al., 2008], was developed in 2008 to address this limitation. Instead of automatically increasing the odds of success with practice opportunities, PFM links the odds of success on an item to the number of past successes and failures on questions that involve one or more skills tested by the item. Defining $S_{k\ell}$ as the number of successes by learner ℓ on items that require skill k, and $F_{k\ell}$ as the number of failures, the PFM model is:

$$ln(\frac{p_{i\ell}}{1-p_{i\ell}}) = \theta_\ell + \sum_k \beta_k Q_{ik} + \sum_k Q_{ik}(\mu_{k\ell} S_{k\ell} + \rho_k F_{k\ell}). \tag{2.10}$$

Evidence presented in the original PFM paper suggests that PFM performs better than AFM. One refinement to PFM, known as the Instructional Factors Model (IFM) [Chi et al., 2011], adds an extra factor $T_{k\ell}$ that counts the number of instructional steps that have occurred relating to skill k and for which no user response is graded (e.g., hints, instruction, etc), since these ungraded steps may increase the chance of success. The IFM model is:

$$ln(\frac{p_{i\ell}}{1-p_{i\ell}}) = \theta_\ell + \sum_k \beta_k Q_{ik} + \sum_k Q_{ik}(\mu_{\ell k} S_{k\ell} + \rho_k F_{k\ell} + v_k T_{k\ell}). \tag{2.11}$$

The KT models presented so far assume that additional item difficulty features or skill requirements have an additive effect on the log odds ratio of success. However, for some items, especially ones involving distinct components / modules, a conjunctive model may be more appropriate. A conjunctive version of AFM is also proposed in the PFM paper [Cen et al., 2008]; it is basically the same as MLTM, except that it includes the $N_{k\ell}$ factor from the AFM model in each component's logistic function:

$$p_{i\ell} = \prod_k \left[\frac{1}{1+exp(-(\theta_\ell + \beta_k + \gamma_k N_{k\ell}))}\right]^{Q_{ik}}. \tag{2.12}$$

The models' authors did not find a substantial difference in performance between PFM and the conjunctive AFM [Cen et al., 2008] on the datasets tested, though that does not preclude one model having advantages over the other on different tests or populations.

All of these KT models suffer from several common limitations. First, they use a single skill variable, which makes them unsuitable for diagnosing errors on multi-skill questions and limits their usefulness on tests comprised of a wide-range of questions or complex questions. Second, the time-dependent factors $N_{k\ell}$ and $T_{k\ell}$ are counts that strictly increase per observation and drive the estimate for $p_{i\ell}$ to 1, even if no learning is occurring. Third, by using the count of correct and incorrect responses, as opposed to their ratio or a count of recent observations with some time decay applied, the PFM and IFM models' estimates for $p_{i\ell}$ will tend towards 0 and 1 as the number of observations grows even if the true $p_{i\ell}$ is somewhere in between. This can be seen by observing the limiting behavior as the number of observations, denoted by t, increases: $\lim_{t\to\infty} |\mathbb{E}_t[S_{k\ell} - F_{k\ell}]| = \infty$ when $p_{i\ell} \neq 0.5$ and that even when $p_{i\ell} = 0.5$, we still have $\lim_{t\to\infty} Var_t(S_{k\ell} - F_{k\ell}) = \infty$. The third problem is shared by Bayesian Knowledge Tracing.

Figure 2-1: Hidden Markov Model for BKT

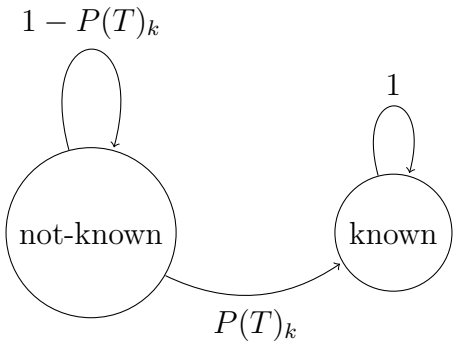

Bayesian Knowledge Tracing (BKT)

BKT was first proposed by Corbett and Anderson in a much-cited 1994 paper [Corbett and Anderson, 1994], and has been developed further in many subsequent publications, e.g., [Pardos and Heffernan, 2010, Yudelson et al., 2013, Doroudi and Brunskill, 2016]. BKT represents knowledge using discrete skills that are either known or not. A skill may be tested by multiple questions, but a question is only supposed to be associated with a single skill. The mastery of each skill is estimated independently using only the questions that target the skill.

The standard BKT model has four parameters for each skill k: $P(T)_k$ is the probability of a student's knowledge of skill k transitioning from the *not-known* to *known* state after an opportunity to apply the skill; s_k is the probability of an error (slippage) when applying known skill; g_k is the probability of a correct answer (lucky guess) when applying not-known skill k; $p(L_0)_k$ is the *a priori* probability of the student knowing skill k before data pertaining to that skill are generated by the intelligent tutoring system.

BKT uses these four parameters per-skill to generate a two-state hidden Markov Model (HMM) shown in Figure 2-1. Let $(L_t)_k$ be the event that the learner knows skill k when generating the t^{th} response to questions pertaining to skill k. Let $R_t = 1$ denote the event that the t^{th} response is correct, and $R_t = 0$ denote an incorrect response. BKT models $P(L_t)_k$ as follows (subscript k is removed for compactness):

$$P(L_1) = P(L_0) \tag{2.13}$$

$$P(L_t|R_t = 1) = \frac{P(L_{t-1})(1-s)}{P(L_{t-1})(1-s)) + (1-P(L_{t-1}))g} \tag{2.14}$$

$$P(L_t|R_t = 0) = \frac{P(L_{t-1})s}{P(L_{t-1})s + (1-P(L_{t-1})(1-g)} \tag{2.15}$$

$$\forall t \geq 2,\ P(L_t) = P(L_{t-1}|R_{t-1}) + P(T)(1-P(L_{t-1}|R_{t-1})). \tag{2.16}$$

The equations for $P(L_t|R_t = 1)$ and $P(L_t|R_t = 0)$ are Bayesian updates of the prior $P(L_{t-1})$ probability that the learner is in a *known* state given the most recent response, but without considering the possibility of a state transition. The equation for $P(L_t)$ is the probability of previously being in the absorbant *known* state, plus the probability of being in the *not known* state transitioning. After observation $t-1$, the predicted probability $P(C_t)$ that question t will be answered correctly is given by:

$$P(C_t) = g(1-P(L_t)) + P(L_t)(1-s). \tag{2.17}$$

By convention, BKT implementations deem a student to be in the *known* state when $P(L_t) \geq 0.95$, or some similar threshold, at which point the ITS will switch to teaching a different skill.

Some variants of BKT propose making one or more of the four parameters not only skill-specific, but also student-specific [Pardos and Heffernan, 2010, Yudelson et al., 2013]. Such models have produced small improvements in predictive performance for student-specific $P(T), P(L_0)$. BKT parameters are typically fit using expectation maximization (EM), or gradient-based methods over a differentiable loss function $L(\mathbf{P}(\mathbf{C}), \mathbf{R})$ on all predictions and responses across all questions and learners. In both cases, the BKT parameters are typically initialized at realistic values and constrained such that $s < 0.25,\ g < 0.25$. The BKT parameter-fitting process can terminate at local minima and is generally too slow for onine implementation; however, once BKT

parameters are fit for each question (if not student-specific), $P(L_t)$ can be computed very quickly.

BKT was developed for intelligent tutoring systems (ITS) rather than for assessment. Consequently, most ITS data-sets and empirical BKT research involve data that are collected in contexts that have teaching and learning as the objective, rather than assessment. In ITS, the learner is offered repeated practice, sometimes with remediation, hints, or explanations in between attempts, until they master the targeted skill. As a result, the learners ITS sessions' typically involve a relatively narrow domain of skills that are chosen by the human instructor to roughly match the learning level of the students. BKT was not designed to efficiently diagnose a highly varying student population's level of mastery across a wide range of skills. Section 2.3 examines BKT's limitations as method for skill assessment in greater detail.

Deep Knowledge Tracing (DKT)

Deep knowledge tracing (DKT), first introduced in 2015 [Piech et al., 2015], uses a recurrent neural network (RNN) to generate a vector $\mathbf{y}_t$ of success probability estimates,

$$\{\hat{P}(L_t)_1, ..., \hat{P}(L_t)_k, ..., \hat{P}(L_t)_K\},$$

for K questions (or homogeneous question types) at each time step t. RNNs are considered a state-of-the-art method for time series tasks with large amounts of training data, such as speech to text [Graves et al., 2013] and translation [Mikolov et al., 2010]. DKT's authors report the method achieving a 25% improvement in area under the curve (AUC) vs. the best previously published predictive performances on datasets from Khan Academy Math and Assistments math software [Piech et al., 2015]. DKT's strong predictive performance has since been replicated on other datasets [Wilson et al., 2016, Khajah et al., 2016], and methodological refinements [Yeung and Yeung, 2018] and applications to adaptive ITS have been proposed [Minn et al., 2018]. However, it is too early to conclude that DKT will consistently dominate other known methods in terms of predictive performance over a wide range of educational datasets.

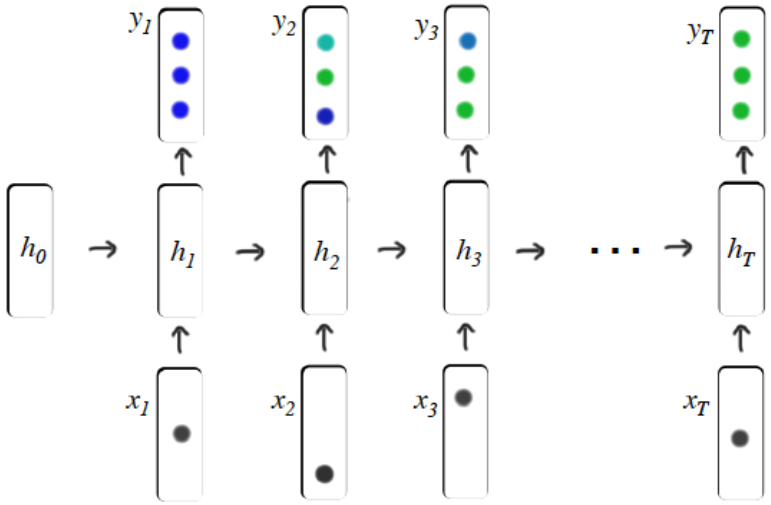

Figure 2-2: DKT RNN architecture
DKT Recurrent Neural Network Structure (reproduced from [Piech et al., 2015])

A recent replication study found that time-dependent IRT-based methods are at least as good as DKT [Wilson et al., 2016], and another, also using the Assistments dataset, showed that an augmented variant of BKT was able to achieve similar predictive performance [Khajah et al., 2016]. Despite these mixed results, DKT appears to be an intriguing new method whose potential warrants further investigation.

The RNN relates a sequence of input vectors $\mathbf{x}_1, ..., \mathbf{x}_T$ to a sequence of prediction vectors $\mathbf{y}_1, ..., \mathbf{y}_T$ through a layer of hidden states, denoted by vectors $\mathbf{h}_0, ..., \mathbf{h}_T$, according to the following two equations in the simple RNN case:

$$\mathbf{h}_t = \vec{\tanh}(\mathbf{W}_{hx}\mathbf{x}_t + \mathbf{W}_{hh}\mathbf{h}_{t-1} + \mathbf{b}_\mathbf{h}), \tag{2.18}$$

$$\mathbf{y}_t = \vec{\sigma}(\mathbf{W}_{yh}\mathbf{h}_t + \mathbf{b}_y), \text{ where } \sigma() \text{ is the logistic function.} \tag{2.19}$$

The parameters learned during the training process are the matrices $\mathbf{W}_{hh}$, $\mathbf{W}_{hx}$, $\mathbf{W}_{yh}$ and vectors $\mathbf{h}_0$, $\mathbf{b}_h$, $\mathbf{b}_y$. The input vectors can either be a one-hot encoding ($\mathbf{x}_t \in \{0,1\}^{2Q}$), where Q is the number of unique questions, or a compressed encoding ($\mathbf{x}_t \in \mathbb{R}^N, N < Q$). The compressed encoding is necessary when the number of unique

questions is large, and is implemented by assigning a randomly generated vector from $\mathcal{N}(0, I)$ for each (question, score) tuple.

The hidden states are estimated using the previous time period's estimate and the new data; however, the weighting of those two information sources does not account for the changing level of certainty surrounding the hidden states, which may result in over-weighting the importance of the previous hidden state early on, and under-weighting as time goes on. A popular but much more complex RNN variant called the Long Short-Term Memory (LSMT) model may overcome this shortcoming by providing greater ability to capture a broader range of potential influences from past observations [Piech et al., 2015] (e.g., from the quantity of observations, or particular permutations of observations), though LSTM does not explicitly model uncertainty either.

The poor interpretability of the hidden states is both a scientific and practical limitation. Performance prediction is often not the main purpose of modeling. Rather, it is a tool for validating a model whose primary purpose is to intelligibly represent the state of a person's mastery of identifiable skills, and/or to identify the components and magnitude of the difficulty of a task. An important aspect of statistical modeling is to improve human understanding of the knowledge domain, of the learner, and of the test item. RNNs, while good at prediction, are generally not well-suited for these explanatory ends.

DKT is also not designed to predict performance on new questions that were not part of the training data, or to help generate new questions adaptively. A low-dimensional representation of the latent skills and approximate difficulty of a question may be ascertained by an expert or known in advance by software if the question is generated from code. In these cases, it would be possible to roughly predict a student's performance on the new question using a cognitive diagnostic model that shares the low-dimensional representation of the latent skill space. By contrast, the ambiguous meaning of the hidden states makes it unrealistic for an expert to identify the correct new column in $\mathbf{W}_{hx}$ associated with a new question. It also makes it difficult for software engineers to meaningfully use hidden RNN node values, or other

RNN model attributes, as inputs into question generation software.

Another challenge is that DKT will capture non-cognitive, non-conceptual relationships between questions that may be present idiosyncratically in the data, such as topic orderings in a particular jurisdiction that do not generalize to other states or countries. This problem could be overcome by labeling students' jurisdictions and using training data from a range of states and countries. Finally, the arguments within tanh() and $\sigma()$ are matrix products, reflecting the assumption of a linear compensatory relationship between hidden states rather than a conjunctive one. This assumption may be reasonable for the simple math questions on which DKT was tested, but perhaps less so for long multi-step problems. Section 2.4 provides an in-depth replication analysis of DKT and identifies a number of more serious problems with the method.

2.2.3 Knowledge Space Theory (KST)

KST is a well-established mathematical framework for modeling and estimating student knowledge using set theory. It was developed primarily by Falmagne and Doignon in the mid-1980s [Doignon and Falmagne, 1985, Falmagne and Doignon, 1988a,b, Falmagne, 1989, Falmagne et al., 1990b,a, Doignon and Falmagne, 1999] and has led to a couple hundred papers by many other scholars who have expanded on the original work [Hockemeyer, 2018]. KST has commercial application in ALEKS (Assessment and Learning in Knowledge Spaces), a math education platform developed at UC Irvine in the mid-1990s and later purchased by McGraw-Hill.

Given a set of question types Q, a student's knowledge state is defined as the subset of the Q question types that the student can answer with high probability of success. A knowledge structure $\mathcal{K}$ is defined as the set of all feasible knowledge states. If $\mathcal{K}$ is closed under unions, it is referred to as a knowledge space.

To identify a student's knowledge state, the knowledge space should be much smaller than the power set; otherwise, the number of possible knowledge states will grow exponentially in $|Q|$. Domains with a strong prerequisite structure, such as mathematics and physics, have helpful natural constraints on the knowledge space

and are therefore more disposed to have manageable knowledge space sizes.

It was originally proposed that knowledge spaces be constructed by querying experts in a systematic manner to reveal prerequisite relationships [Koppen and Doignon, 1990, Koppen, 1993]. This is a reasonable strategy when $|Q|$ is sufficiently small, but does not scale well with large $|Q|$. Since querying experts is expensive, this method is also impractical in applications that generate new questions regularly or on-demand. The techniques will also not work well in complex domains that have fuzzy prerequisite relationships, such as biology, language arts, social science, because (i) experts may not know the ground truth about the feasible sets, (ii) even if experts could know, the weaker restrictions on the combinations of questions that could be mastered may cause the knowledge space to still grow exponentially as Q increases.

Another method for constructing a knowledge space is to propose a number S of discrete skills that comprehensively characterize the question domain ($S < Q$), and then identify which questions require which skills. This procedure produces a knowledge space containing at most 2^S possible knowledge states. So long as every combination of skills produces a unique set of answerable questions, this knowledge space's states have a 1-1 relationship with the classes of a Diagnostic Classification Model with a single full interaction term for all of the categorical variables for the skills that the question requires [Schrepp, 2004, Heller et al., 2015].

Data-driven methods for identifying the set of all possible knowledge states, without expert querying or skill-modeling, have also been proposed [Schrepp, 1999]. A challenge for these methods is that real-world data from a particular student population may not contain a sufficiently large sample of students from each knowledge state to confirm the existence of states that may occur with higher frequency in other populations. Data-driven methods may also rule-out conceptually possible states that are not occurring in the data as a result of widespread but arbitrary patterns of curriculum ordering across large school systems.

Unlike IRT and LCDM, KST does not explicitly model or estimate latent cognitive variables; however, human intuition about latent cognitive variables or question attributes may be implicitly codified in the expert-querying process. It would not be

unreasonable to think that when the "true" cognitive model is mathematically complex an expert might have more modeling success with the ground-up framework of KST, where rules and relationships are specified through many small cases, than with an effort to arrive at an explicit mathematical expression that attempts to represent the "true" cognitive model. This may be a practical benefit of using KST, especially when (i) there is less data than desired for model selection and over-fitting prevention, and (ii) when there is a large space of candidate models because of the complexity of question-skill dependencies in the domain. However, unless there are codified links between knowledge states and latent variables, KST becomes a self-contained framework with limited portability and interpretability.

Another limitation of standard KST arises when performance on a question is most reasonably scored along a continuum or with ratings. Although KST treats question mastery as binary, data from ALEKS show that more advanced students have lower rates of error on easier questions [Beeler, 2018]. While KST can be generalized to allow for multiple response qualities along a linear scale [Schrepp, 1997], it appears that this has not been done in practice in software or in published empirical work. Doing so would risk making the size of the knowledge space, and complexity of specifying it, intractably large.

Knowledge State Inference

The central function of KST is not to characterize the set of possible knowledge states, but to infer the particular knowledge state of a learner at a given point in time. Such inference is generally done under probabilistic assumptions: (i) that student who does not know how to answer question i may still succeed by guessing with probability g_i, and (i) that students who know how to answer question i may make "careless errors" with probability s_i. "Careless errors" are simply KST's terminology for "slippage". [1]

Let $X_{i\ell}$ be a random variable for learner ℓ's scored response on question i. Let $K_{i\ell}$ be an indicator variable for whether question i is in learner ℓ's knowledge state,

[1] KST papers often use other letters by convention to denote guessing and slippage, such as η and β. This section uses g and s for consistency with non-KST models since the parameters have identical definitions and mathematical roles.

K_ℓ. Then

$$\mathbb{P}(X_{i\ell} = 1) = (1 - s_i)^{K_{i\ell}} \cdot g_i^{(1-K_{i\ell})} \tag{2.20}$$

$$\mathbb{P}(X_{i\ell} = 0) = s_i^{K_{i\ell}} \cdot (1 - g_i)^{(1-K_{i\ell})}. \tag{2.21}$$

The basic linear independence model (BLIM) holds that, after conditioning on the knowledge state, a learner's responses to various questions are independent. Let R_ℓ be a 0-1 vector containing the scores of learner ℓ's responses, and suppose that vectors $\mathbf{s}$ and $\mathbf{g}$ have been estimated separately using data from many learners. Denote the prior probability of a student being in knowledge state K as π_K. Then the probability of observing R_ℓ, given $K, \mathbf{s}, \mathbf{c}$ is given by

$$P(R_\ell|K, g, s) = \prod_i \left[R_{i\ell}\mathbb{P}(X_{i\ell} = 1) + (1 - R_{i\ell})\mathbb{P}(X_{i\ell} = 0)\right], \tag{2.22}$$

and the probability of observing R_ℓ is given by

$$P(R_\ell) = \sum_K P(R_\ell|K, g, s) \cdot \pi_K. \tag{2.23}$$

The notation above uses vectors and indicator variables instead of set notation to make the equations more readily implementable for readers from a statistical or optimization background. The convention in most KST literature is to use set theoretic notation; for clarity to readers with a KST background, alternate versions of equations 2.22 and 2.23 are presented below:

If we let r_q be the a learner's response to question q, and K be the set of questions

the learner can answer, then

$$P(r_q|K) = \begin{cases} \beta_q & if\ r_q = 0, q \in K \\ 1 - \eta_q & if\ r_q = 0, q \notin K \\ 1 - \beta_q & if\ r_q = 1, q \in K \\ \eta_q & if\ r_q = 1, q \notin K. \end{cases}$$

If we define R and W as the set of correctly and incorrectly answered questions by a learner, respectively, then Equations 2.22 and 2.24 are equivalent.

$$P(R, W|K) = \left[\prod_{q \in K \cap W} \beta_q\right] \left[\prod_{q \in K \cap R} (1 - \beta_q)\right] \left[\prod_{q \in R \setminus K} (\eta_q)\right] \left[\prod_{q \in W \setminus K} \beta_q\right]. \tag{2.24}$$

Under BLIM, the most likely knowledge state(s) is $\text{argmax}_K P(R, W|K)\pi_K$, which can be found in $O(|Q||\mathcal{K}|)$ time. Generalizations of BLIM to handle missing data have been proposed [Anselmi et al., 2016].

A related group of KST inference methods described as "the continuous procedure" [Falmagne and Doignon, 1988a] can be used to choose which question to pose next and when inferential certainty is sufficient to terminate the assessment process. The continuous procedure starts with a prior likelihood for each knowledge state, $\mathcal{L}_0(K)$. The prior could simply be uninformative (discrete uniform), or it could reflect empirical distributions from relevant student populations [Anselmi et al., 2016]. There are many possible question-selection steps, which typically take the form of a heuristic to minimize the number of questions needed to reach a stopping criterion. One such heuristic is the half-split rule, which chooses a question q for which the probability of a correct answer is as close as possible to 0.5; mathematically, this is achieved by minimizing $|0.5 - \sum_{K \in \mathcal{K}_q} \mathcal{L}_n(K)|$, where $\mathcal{K}_q$ is the set of knowledge states that contain q. The procedure typically stops when $\exists\ K^*\ s.t.\ \mathcal{L}(K^*) \geq \tau$ for some $\tau > 0.5$. The likelihood function can be updated by Bayes Rule [Anselmi et al., 2016] as follows:

$$\mathcal{L}_{n+1}(K) = \frac{P(r_q|K)\mathcal{L}_n(K)}{\sum_{K' \in \mathcal{K}} P(r_q|K')\mathcal{L}_n(K')}. \tag{2.25}$$

When all questions have been administered the equality $\mathcal{L}_n(K) = P(K|R, W)$ holds, where by Bayes Rule

$$P(K|R, W) = \frac{P(R, W|K)\mathcal{L}_0(K)}{\sum_{K' \in \mathcal{K}} P(R, W|K')\mathcal{L}_0(K')}. \tag{2.26}$$

The big-O computation time to find K^* is $O(|\mathcal{K}||Q|^2)$. Much faster computation may happen in practice if the sequence of questions chosen through sequential bifurcation of the probability mass identifies K^* within $O(log(|Q|))$ steps; however, such behavior will depend on the structure of $\mathcal{K}$.

2.3 Statistical Properties of BKT

Although some knowledge modeling paradigms, such as KST and BKT, represent knowledge as a binary state, it is widely acknowledged, even by the earliest proponents of KST and BKT, that skill mastery often expresses itself along a continuum [Corbett and Anderson, 1994]. Amongst those familiar with BKT, there is a widespread belief that the quantity $P(L_t)$ represents not only the probability that the learner is in the learned state, but that it also serves as an estimator for the learner's true level of partial mastery. Contrary to this belief, $P(L_t)$ is not an estimator for partial skill mastery; it is a quantity that reflects which of two possible states are most favored by the evidence, and can in fact be quite misleading when the true generative model is intermediary between the two states. This section will formally prove that $P(L_t)$ should not be used as an estimator for partial mastery, and then illustrate its behavior using simulated data.

Suppose that the learner's true mastery of a skill is parameterized by $\lambda \in (0, 1)$ such that $P(C_t) = g(1 - \lambda) + \lambda(1 - s)$. Observe that λ takes the same place as $P(L_t)$ in Equation 2.17; one could be forgiven for thinking that $P(L_t)$ is therefore an

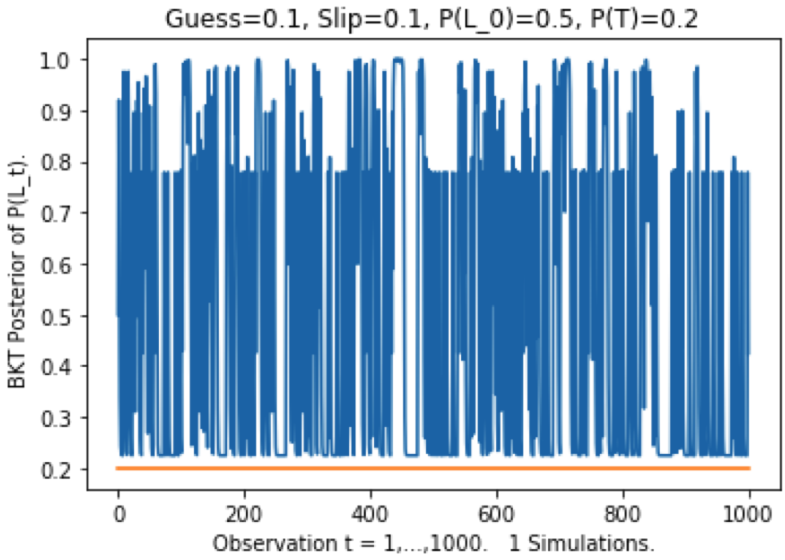

Figure 2-3: $P(L_t)$ does not converge. P(T) creates a lower bound for $P(L_t)$ and biases it upwards. This is one simulation with 1000 observations and $\lambda = 0.2$, indicated by the horizontal line.

estimator for the true λ. However, it will be shown by both simulation, and later by formal proof, that $P(L_t)$ does not converge to to λ even when the model is given the true values of g, s and $P(T)$.

The non-convergence, even after 1000 observations, is shown in Figure 2-3. The BKT equations cause $P(T)$ to act as a lower bound for $P(L_t)$; since $\lambda = 0.2$ in this example, all of the $P(T)$ estimates are above λ. To show that this behavior is not the result of cherry-picked simulation, Figure 2-4 presents 200 equally noisy and non-convergent simulations with the same settings as in Figure 2-3.

Since the true generative model is stationary, we could set $P(T) = 0$ to eliminate this particular source of bias. Figure 2-5 shows the paths of $P(L_t)$ for 200 simulations when $P(T)$ is set to 0 in the equation for $P(L_t)$. The paths, each with their own color in the figure, are highly overlapping and so they are not all individually visible. Nonetheless, it is clear that the paths are converging towards $P(L_t) = 0$ when $\lambda = 0.2$. Despite this simulated learner's true mastery state being very close to non-mastery, 26% of the simulated time series had $P(L_t) \geq 0.9$, a common mastery threshold, before the tenth observation. When $\lambda = 0.5$, 74.2% of 1000 simulations surpassed had $P(L_t) \geq 0.9$ before the tenth observation.

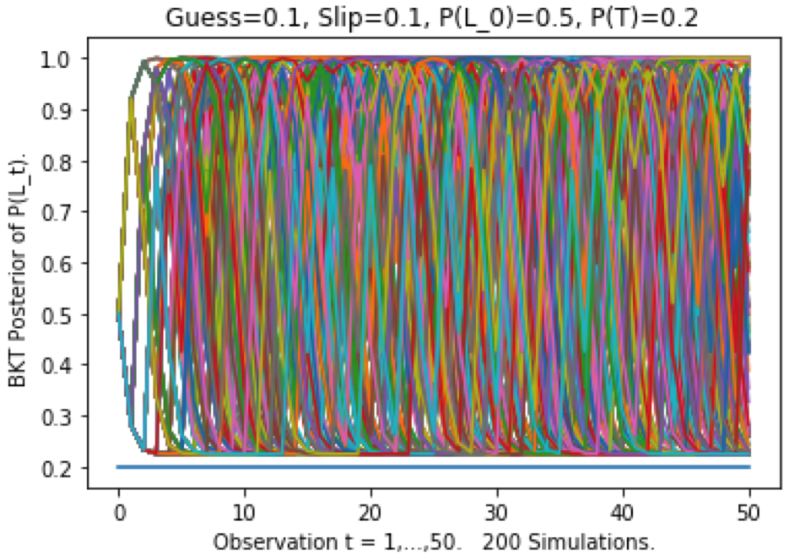

Figure 2-4: 200 simulations with the same settings as in Figure 2-3. Average value at t=10 was $P(L_t) = 0.48$. 47.5% of simulations had $P(L_t) \geq 0.9$ before the tenth observation, even though $\lambda = 0.2$.

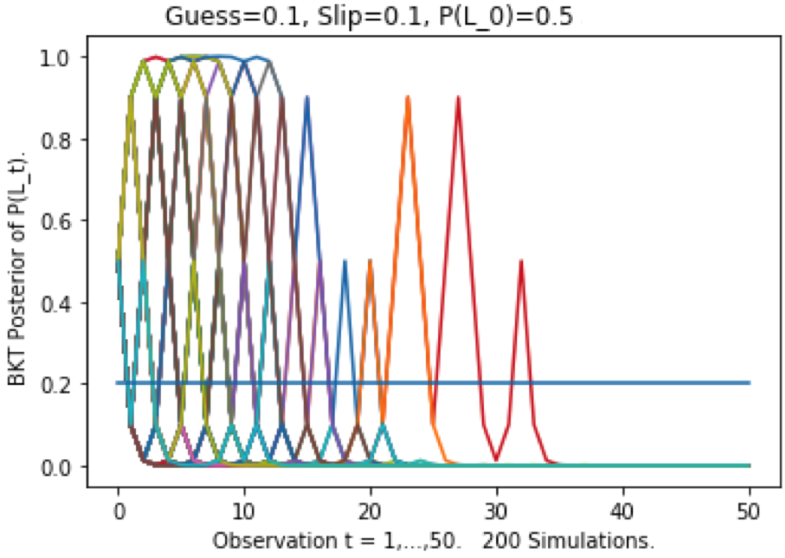

Figure 2-5: $P(L_t)$ from t=1,...,50 with P(T) = 0, g=s=0.1, and $P(L_0) = 0.5$, simulated 200 times with $\lambda = 0.2$, indicated by the horizontal line. $P(L_t)$ crossed 0.9 within $t \leq 10$ in 26% of simulations.

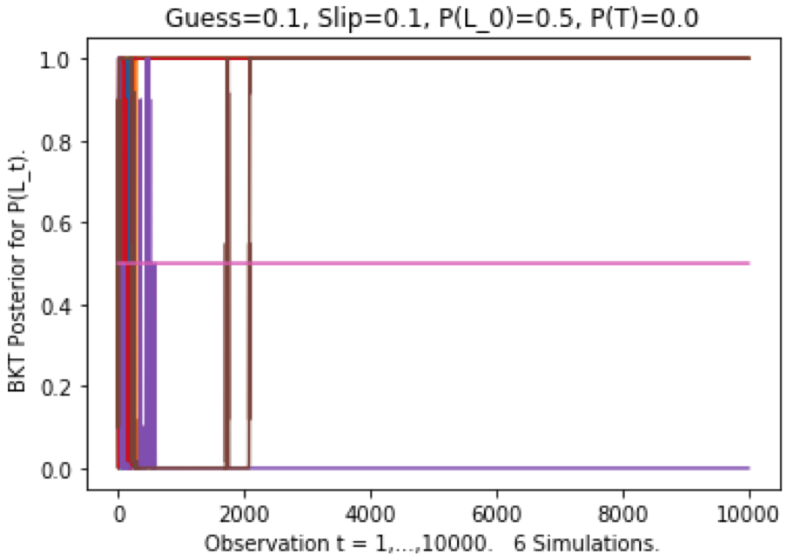

Figure 2-6: The horizontal line at 0.5 is to contrast the true value $\lambda = 0.5$ with the simulated paths for $P(L_t)$, which converge to 0 or 1 in this case.

When $\lambda = 0.5, P(T) = 0$, $g = s = .1$, then $P(L_t)$ converges to either 0 or 1 (Figure 2-6). If we allow $P(T) > 0$, the upward bias causes $P(L_t)$ to converge to 1 (Figure 2-7) when $s = g$. If $s < g$ while $\lambda = 0.5, P(T) = 0$, then the sequence $\{P(L_t)\}$ has downward drift, and $s > g$ produces upward drift. The drift from $s \neq g$ may be counteracted, or reinforced, by $P(T) > 0$ or $\lambda \neq 0.5$ in ways that are difficult to characterize analytically. For example, in the realistic case of low slippage s=0.02 and multiple choice questions with four options, g=0.25, the limiting behavior of $P(L_t)$ is sensitive to small changes in λ. For example, when $\lambda = 0.65$ the average value of $P(L_{1000})$ was 0.50 (20 simulations in Figure 2-8), though most samples were either very close to 1 or very close to 0. Increasing λ slightly to 0.7, the average value of $P(L_{1000})$ becomes 0.99 (20 simulations in Figure 2-9).

2.3.1 Formal Proofs of BKT Properties

Property 1. *When response data are generated by a stationary $\lambda \in (0, 1)$, for the posterior estimate $P(L_t)$ in BKT $lim_{t\to\infty} P(L_t) \neq \lambda$.*

Proof. We can use a proof by contradiction. First, observe that the updating rule for $P(L_t)$ has a Markov property; the exact history of observations does not matter.

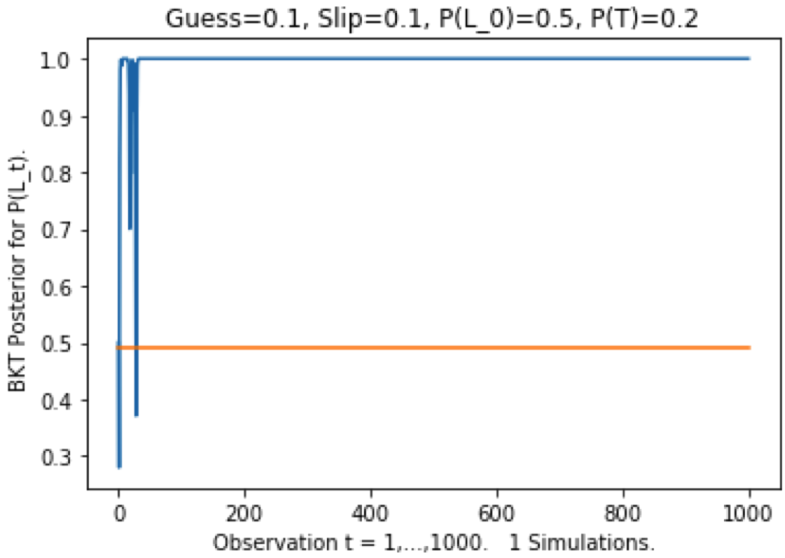

Figure 2-7: P(T)=0.2 produces enough upward bias for $P(L_t)$ to converge to 1 when $\lambda = 0.5$.

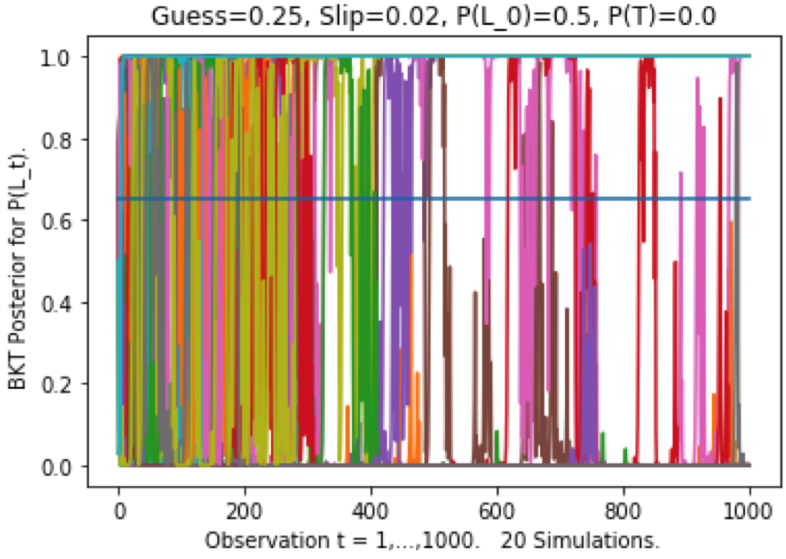

Figure 2-8: When $\lambda = 0.65$, g = 0.25 (e.g., multiple choice with four options) and s = 0.02 (question is easy if relevant concepts are known), $P(L_t) > 0.9$ within $t < 10$ in 19 out of 20 simulations. At $t = 1000$ the average value of $P(L_t)$ was 0.50.

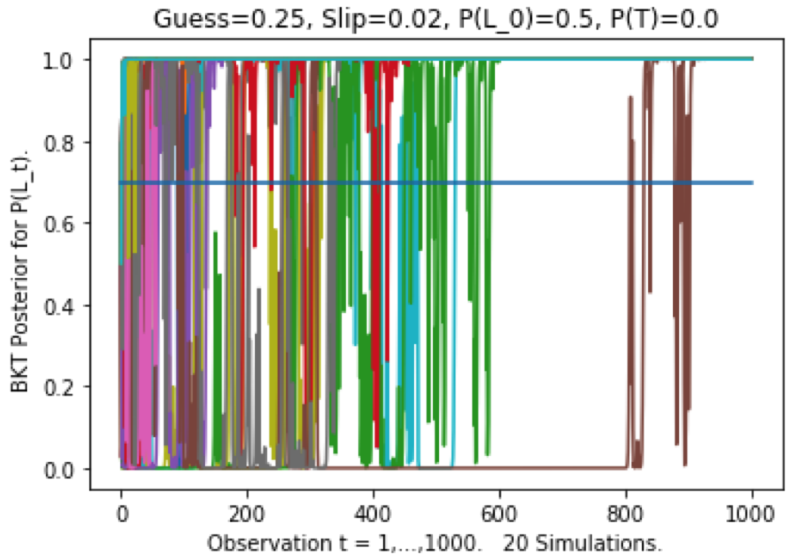

Figure 2-9: When $\lambda = 0.70$, g = 0.25 (e.g., multiple choice with four options) and s = 0.02 (question is easy if relevant concepts are known), $P(L_t) > 0.9$ within $t < 10$ in 20 out of 20 simulations. At $t = 1000$ the average value of $P(L_t)$ was 0.99.

The transition to $P(L_t)$ depends only on $P(L_{t-1})$ and the most recent observation. Next, suppose $\exists s, g \in (0, 0.5), \lambda \in (0, 1)$ such that $lim_{t\to\infty} P(L_t) = \lambda$. For this to be true, $lim_{t\to\infty} |P(L_{t+1}) - P(L_t)| = 0$. Since the Markov property applies, the difference $|P(L_{t+1}) - P(L_t)|$ does not directly shrink as a function of t itself, so it would have to do so as $P(L_t)$ gets closer to λ over time. Observe that $\forall t$ if $P(L_t) = \lambda$ then $P(L_{t+1}|R_{t+1} = 1) = \frac{\lambda(1-s)}{\lambda(1-s)+(1-\lambda)g} \neq \lambda$ and $P(L_{t+1}|R_{t+1} = 0) = \frac{\lambda s}{\lambda s+(1-\lambda)(1-g)} \neq \lambda$. These differences will not diminish over time, and so the supposition is incorrect and convergence to $\lambda \in (0, 1)$ cannot occur. □

Not only does $P(L_t)$ fail to converge to stationary $\lambda \in (0, 1)$, it is easy to show for many values of s, g, λ that $P(L_t)$ will converge to 0 almost surely ($g > s, p \leq 0.5$), to 1 almost surely ($g < s, P(C_t) \geq 0.5$), or to either 0 or 1 almost surely ($s = g, P(C_t) = 0.5$).

Property 2. *When response data are generated by a stationary $\lambda \in (0, 1)$ and $s = g \in (0, 0.5)$ the posterior estimate $P(L_t)$ in BKT is a Markov Chain with an isomorphic relationship to an unbounded random walk.*

Proof. Let X_t be 1 if the t^{th} response is correct and -1 otherwise. Let $S_\tau = \sum_{t=1}^{\tau} X_t$.

S_τ increases with probability $p = \lambda g + (1-\lambda)(1-s)$ and decreases with probability $(1-p)$ at every step. S_τ therefore represents the state of an unbounded random walk after τ observations. Given well-known properties of random walks, for $p > 0.5$, $lim_{\tau\to\infty} S_\tau = \infty$ almost surely (a.s.), and for $p < 0.5$, $lim_{\tau\to\infty} S_\tau = -\infty$ a.s. For $p = 0.5$, S_τ has no limit, but $lim_{\tau\to\infty} Var(S_\tau) = \infty$, which means there is no stationary distribution for the series $\{S_\tau\}$; hence, for any finite lower bound L and upper bound U, the fraction of time $S_\tau \in [L, U]$ approaches 0 as $\tau \to \infty$. The possible values of S_τ have an isomorphic ordered relationship with the possible values of $P(L_\tau)$. This relationship is summarized by the following equivalencies:

$$
\begin{aligned}
S_0 = 0 &\iff P(L_0) = \rho \text{ (an assumed prior)} \\
\forall\tau, S_{\tau+1} = S_\tau + 1 &\iff P(L_{\tau+1}) = \frac{P(L_\tau)(1-s)}{P(L_\tau)(1-s) + (1-P(L_\tau))s} \\
\forall\tau, S_{\tau+1} = S_\tau - 1 &\iff P(L_{\tau+1}) = \frac{P(L_\tau)s}{P(L_\tau)s + (1-P(L_\tau))(1-s)}.
\end{aligned}
$$

From the previous three equivalencies, we can obtain the following Lemma:

Lemma 1.

$$\forall\tau, k, S_{\tau+k} = S_\tau \iff P(L_{\tau+k}) = P(L_\tau) \tag{2.27}$$

To prove Lemma 1 we show that $P(L_{t+2}|X_{t+1} + X_{t+2} = 0) = P(L_t), \forall t$. In other words, that the change in $P(L_t)$ from a correct answer is perfectly reversed by a

subsequent wrong answer, and vice-versa (since $s = g$ creates symmetry).

$$\text{Given } P(L_{t+1}|X_{t+1} = 1) = \frac{P(L_t)(1-s)}{P(L_t)(1-s) + (1-P(L_t))s} \text{and}$$
$$P(L_{t+2}|X_{t+2} = -1) = \frac{P(L_{t+1})s}{P(L_{t+1})s + (1-P(L_{t+1}))(1-s)},$$

substitute $P(L_{t+1}|X_{t+1} = 1)$ for $P(L_{t+1})$, and for notation compactness let $y = P(L_t)(1-s) + (1-P(L_t))s$ to obtain $P(L_{t+2}|X_{t+1} + X_{t+2} = 0) =$

$$\begin{aligned}
&\frac{P(L_t)(1-s)s}{y\left[\frac{P(L_t)(1-s)s}{y} + \left(1 - \frac{P(L_t)(1-s)}{y}\right)(1-s)\right]} \\
&= \frac{P(L_t)(1-s)s}{P(L_t)(1-s)s + \left[y - \left[\frac{y(P(L_t)(1-s)}{y}\right]\right](1-s)} \\
&= \frac{P(L_t)s}{P(L_t)s + \left[y - P(L_t)(1-s)\right]} \\
&= \frac{P(L_t)s}{P(L_t)s + \left[P(L_t)(1-s) + (1-P(L_t))s - P(L_t)(1-s)\right]} \\
&= \frac{P(L_t)s}{P(L_t)s + (1-P(L_t))s} \\
&= \frac{P(L_t)s}{s} \\
&= P(L_t).
\end{aligned}$$

Lemma 1 produces Property 2 by induction. □

Property 2 can be seen in Figure 2-10 via the limited and recurring set of points (states of the Markov Chain) through which the 1000 simulated $\{P(L_t)\}$ sequences pass. When $\lambda = 0.5$ and $s = g \in (0, 0.5)$, it follows that $p = 0.5$; in this case, the sequence will diverge from the prior ρ towards either 0 or 1 with equal probability, which means that although $E\left[lim_{t\to\infty} P(L_t)\right] = 0.5$, the asymptotic error of the BKT estimator is also 0.5. This is the maximum possible error for $\lambda = 0.5$.

When $s \neq g$, Lemma 1 no longer holds, and for each possible value of S_τ, the corresponding value of $P(L_\tau)$ drifts upwards ($s > g$) or downwards ($s < g$), as shown in (Figure 2-12). For some values of s and g, this drift can counteract the drift induced

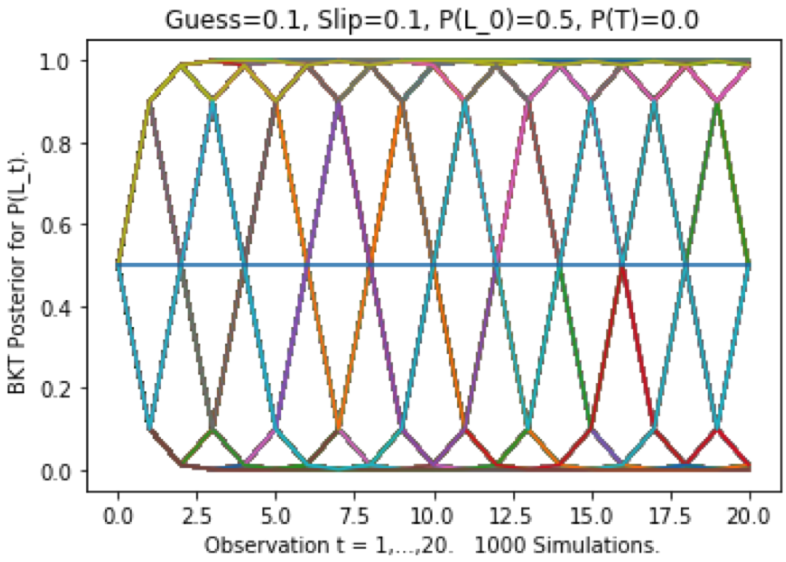

Figure 2-10: Property 2 reflected in the limited set of $P(L_t)$ values $\forall t$ when $s = g$.

on S_τ from $p \neq 0.5$, and can cause the asymptotic error to exceed 0.5. For example, in all 100 simulations when $g = 0.8, s = 0.01$ and $\lambda = 0.53$ the value of $P(L_t)$ converged to zero even though $p = 0.5623 > 0.5$, resulting in an asymptotic error of 0.53 (see Figure 2-11).

BKT's bias is even more pronounced when the method is applied to conjunctive (i.e., multi-step) tasks, as suggested in the original BKT paper [Corbett and Anderson, 1994]. Applying BKT to a task with m steps could produce $lim_{t\to\infty} \prod_m P(L_{t,m})$ $= 1$ with stationary data despite the true probability of success being $(0.5 + \epsilon)^m$. Clearly, BKT is not a suitable modeling framework for domains in which partial mastery occurs, and the BKT posterior should not be used to estimate partial mastery.

2.4 Deep Knowledge Tracing - Question Cluster Detection

This section examines the claim that DKT is useful for detecting hidden relationships between questions and clusters of questions connected by a common hidden skill. This claim plays an important role in advancing the idea that the RNN finds a

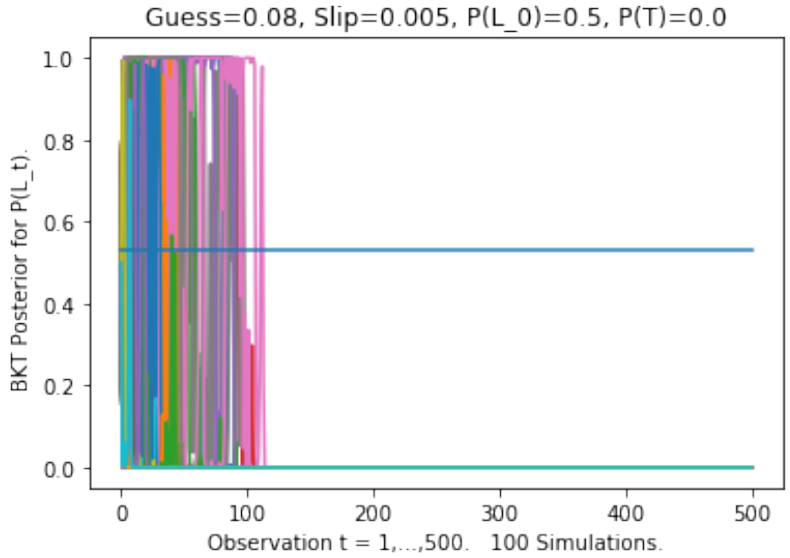

Figure 2-11: When $s \neq g$ the error between λ and $P(L_\infty)$ can exceed 0.5 and the direction of the error can be the opposite of the drift of the random walk representing the score.

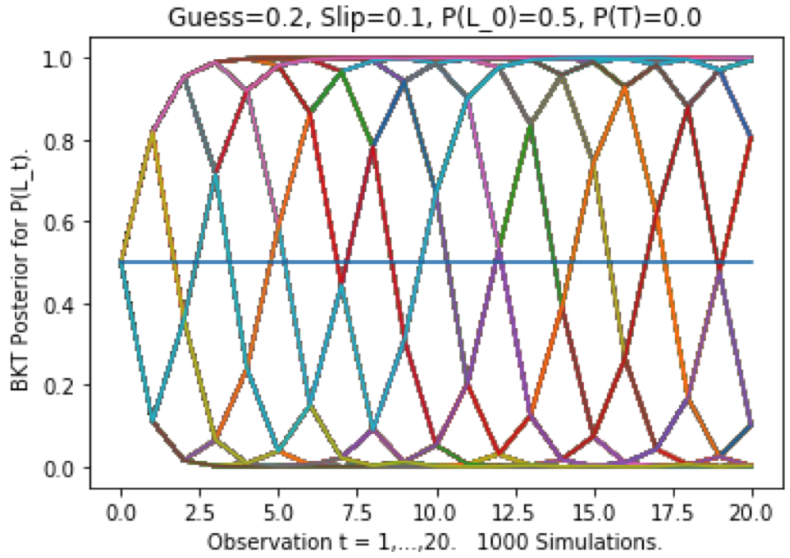

Figure 2-12: When $s < g$, the possible values $P(L_t)$ corresponding to a particular score S_t drift downwards with t.

representation of the true underlying statistical structure of the data; in other words, it is a black box that can be trusted. The claim, if true, could make DKT a valuable tool for automatically identifying skill dependencies in a large body of questions, e.g., from national high-school exam systems in China and India. Such labeling would make it easier to integrate past questions into the skill ontology (if any) of a digital learning platform.

According to Piech et al.,

> The DKT model can further be applied to the task of discovering latent structure or concepts in the data, a task that is typically performed by human experts. We approached this problem by assigning an influence J_{ij} to every directed pair of exercises i and j,
>
> $$J_{ij} = \frac{y(j|i)}{\sum_k y(j|k)}, \tag{2.28}$$
>
> where $y(j|i)$ is the correctness probability assigned by the RNN to exercise j on the second timestep, given that a student answered exercise i correctly on the first. We show that this characterization of the dependencies captured by the RNN recovers the pre-requisites associated with exercises (p. 5).

This methodology is reported to have generated "a perfect clustering of latent concepts in the synthetic data" (p. 8) when influences above a threshold of 0.1, in either direction, was treated as positive identification of a connection between questions. Intuitive, useful clustering was also presented for data from Khan Academy, though that dataset is not public.

Unfortunately, this clustering result using this influence metric could not be replicated, and it appears this failure may be due to problems with both the metric and the RNN fitting process rather than a lack of signal.

2.4.1 Model Replication

To attempt replication, the DKT model was constructed and fit in Pytorch using the same parameters and methods described in the original paper: 200 hidden nodes,

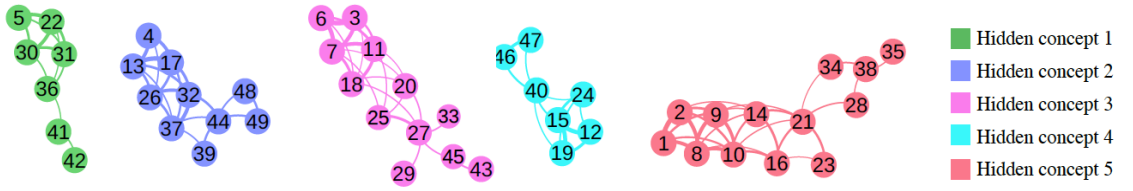

Figure 2-13: Copy of original image in [Piech et al., 2015] showing perfect clustering of questions by hidden concept using influence metric with threshold 0.1 to form connections.

batches of 100, binary cross-entropy (BCE) as the loss function, dropout (50%) in the training sets, and gradient clipping (though mostly not needed). The learning rate and number of epochs were not stated in the original paper. We used a learning rate of 0.0033 and 25 epochs, which was sufficient to achieve apparent statistical convergence in the loss function within two minutes on a laptop. No GPU processing was required. We used the Adam algorithm [Kingma and Ba, 2015] instead of standard stochastic gradient descent (SDG).

The authors' synthetic datasets were made public, comprising of 20 datasets with 5 hidden skills and 20 with 2 hidden skills, and with each dataset comprising a training set and test set of 2000 learners and 50 questions apiece. As in the original paper, the models for the synthetic datasets were fit on "the same sequence of 50 questions." The paper reported an out-of-sample AUC of 0.75 for the datasets with five hidden concepts. The AUC for two hidden concepts was not stated but accuracy appears to be 81% for a particular dataset presented in a figure. In our replication, the AUC was higher than in the original paper. The difference between the test set and training set AUC was small enough to suggest little overfitting, as summarized in Table 2.4.1. However, it is unclear why we obtained a better AUC. In any case, the high out-of-sample AUC gives confidence that the RNN was fit at least as well as in the original paper.

Hidden Skills	Train AUC	Test AUC	Original Result
5	0.832-0.842	0.821 - 0.833	AUC 0.75
2	0.866-0.883	0.865-0.882	Accuracy 81%

Table 2.1: Predictive performance of model on the original paper's 40 synthetic datasets was consistently better than the original results.

2.4.2 Using Influence to Find Clusters - Main Problems

Having established that the fitted RNN produces a high AUC, it may seem puzzling that the influence metric did not produce the expected clusters. A closer inspection of the mathematical properties of the influence metric suggests that the metric itself is not well constructed as a signal for influence, and that attaining an influence level of 0.1 may be implausible on a set of 50 questions.

Observe that if all 50 questions influence each other equally (i.e., have the same $y(\cdot|\cdot)$), then $J_{ij} = \frac{1}{|Q|} = 0.02 \; \forall i, j$, whether there be a perfect correlation between the questions, or no relationship at all. The metric should distinguish between these two cases. It is unclear how the authors chose 0.1 as a threshold. A statistical procedure for choosing a threshold was not proposed. The threshold's dependence on $|Q|$ should be formalized or eliminated.

In the case of a cluster comprising of α % of the questions, let $y(i|j) = \beta$ when i and j belong to different clusters, and let $y(i|j) > \beta$ when i and j belong to the same cluster. For equal-sized clusters, $(1 - \alpha)\%$ of the influence values J_{ij} will be distributed with a central tendency between $\frac{1}{|Q|}$ and $\frac{\beta}{(1-\alpha)\beta|Q|+\alpha|Q|}$, and $\alpha\%$ of the values will have influences distributed with a somewhat higher mean. However, it is straightforward to show that same-cluster influences cannot be expected to reach the threshold of 0.1 when $|Q| = 50$. When all same-cluster influences are as high as possible, $y(i|j) = 1$, then $J_{ij} = \frac{1}{\alpha\beta|Q|+(1-\alpha)|Q|}$. In the case of $\alpha = 0.2$ and $\beta = 0.5$, $J_{ij} = 0.0\overline{33}$, which is well below the threshold of 0.1. Now consider $\alpha = 0.2$, $\beta = 0.2$, which could represent a five-skill scenario in which the model is trained on students who tend not to know the material initially (i.e., random guessing on multiple choice with five answers). Then the maximum for J_{ij} becomes $\frac{1}{0.36|Q|} = 0.0\overline{55}$, which is still well below the threshold of 0.1. The way to maximize J_{ij}, given α, β, and subject

to the requirement that a connected cluster be formed, is to suppose that within each hidden skill cluster each question has $y = 1$ on only one other question, $y = \beta$ otherwise, and they are linked in a chain. In this setting, the theoretical maximum J_{ij} is $\frac{1}{0.2(|Q|-1)+1} = \frac{1}{9.8+1} = 0.09\overline{259} < 0.1$. Therefore, it appears that the clustering results displayed in Figure 2-13 are not mathematically plausible.

The original authors were contacted by email and provided a copy of the Lua code they used to generate influence metrics, and a confirmation of the hyperparameters used to train the RNN. As of the time of writing, we were not able to identify an error in our Pytorch implementation of the RNN or in our calculation of the influence metric, though we remain in touch with the authors.

The histogram of influences J_{ij} generated in our replication study, shown in Figure 2-14, confirm that the central tendency is around 0.02. No values attained 0.1. In fact, no values exceeded 0.026. In order for perfect clustering to occur, the influence metric would need to have a bimodal distribution such that above some threshold τ there is not a single false positive edge (otherwise two unrelated clusters would merge), and a sufficient number of true positives surviving the threshold to connect the same-skill questions to each other.

We developed alternative influence metrics to see whether the fitted RNN might still be useful for question cluster identification. The first change in the metric was to include signals from both correct and incorrect answers: failure on a basic pre-requisite skill could strongly imply failure on a harder skill, whereas success on the pre-requisite skill may only mildly raise the predicted chance of success on the harder skill. Formally, we substituted $y(i|j)$ for $z(i|j) =: \max\big(y(i|j \text{ correct}), |y(i|j \text{ incorrect}) - 1|\big)$. The second change was to standardize the metric by using $\text{median}_k z(i|k)$ instead of $\sum_k z(i|k)$ in the denominator, and to use $z(i|k) - \text{median}_k z(i|k)$ in the numerator. A more rigorous method, though not implemented, would have been to use the $\frac{(1-\alpha)}{2}$-percentile instead of the median as the denominator, and to subtract it in the numerator, as it would be an unbiased estimator of $z(i|k)$ when i and k are in different clusters. The third change was to increase the treatment dosage of influence from j on i in the RNN by inputting 2 or 3 correct (or incorrect) responses on

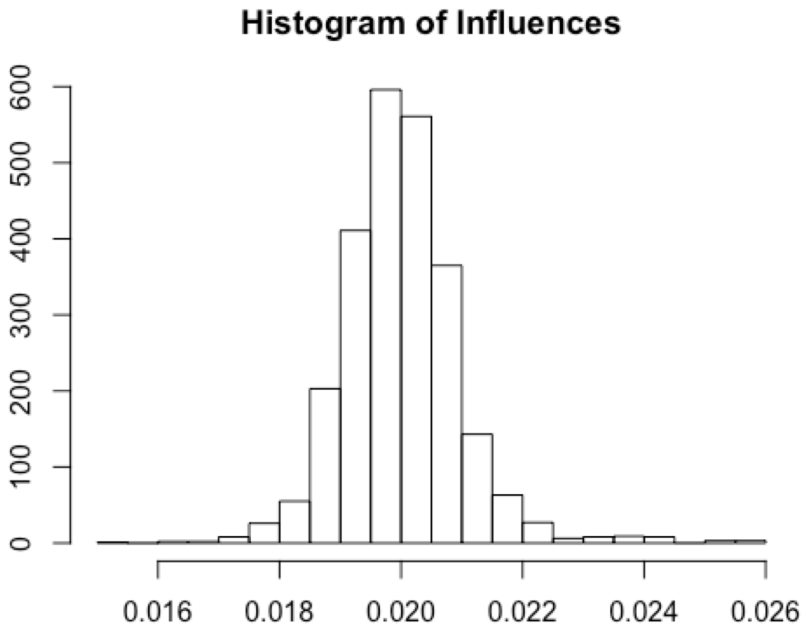

Figure 2-14: Histogram of influence values. Dataset #20, 5 hidden skills.

question j in a row, rather than just one. Finally, the RNN was fit with and without L1 regularization on the model parameters. The L1 penalty parameter was chosen such that the L1 term would be only 5%-15% the size of the BCE loss. As shown in Figure 2-15, the re-defined influence metrics succeeded in producing a heavier right tail, and regularization brought about more of a bimodal distribution, albeit not with the perfect separation one would want to confidently avoid false-positive connections between questions.

Surprisingly, though for reasons that will be made clear in Section 2.5, for each of ten thresholds tried between 0.15 and 0.85, the questions consistently formed a single giant component instead of five distinct clusters. To see if cluster-like communities existed within each of these giant components, a community detection algorithm was implemented. The number of communities identified was almost always less than 5, and the modularity of the community assignments was under 0.25 (very low). This result suggests that the largest influence measures extracted from the fitted

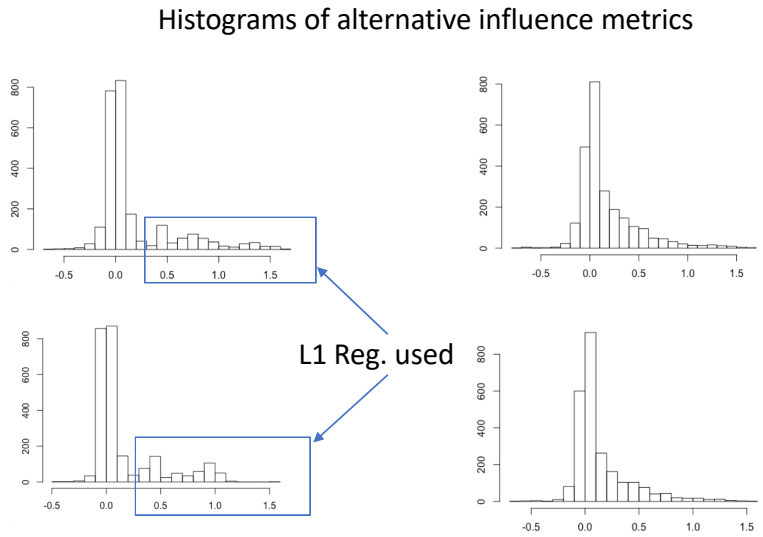

Figure 2-15: Histograms of alternative influence metrics. Dataset #20, 5 hidden skills. Left-column, with regularization. Top row: three repeated inputs. Bottom row: two repeated inputs.

RNN are not acting as a good predictor of whether questions belong to the same hidden skill group. The high AUC of the RNN suggests that this failure is not due to a lack of signal in the underlying data. In fact, the underlying signal can produce perfect clustering even using the simplest of non-RNN methods. Using the correlation between answers to question i and question j as an alternative influence metric between i and j, we obtained perfect clustering in all 40 synthetic datasets when applying a threshold of a correlation of 0.1.

To summarize, after testing all 40 synthetic datasets from the original DKT paper we find the question clustering results to be non-reproducible. For several reasons, the original influence metric proposed does not seem ideally suited to the task, and although alternative influence metrics seemed to have generated a stronger influence signal, they did not produce the expected number of clusters or modular communities, whereas a simple correlation between questions achieved perfect cluster separation. Given the performance of simple correlation as an influence measure, it seems unnecessary, at least for any dataset similar in structure to the synthetic datasets tested, to

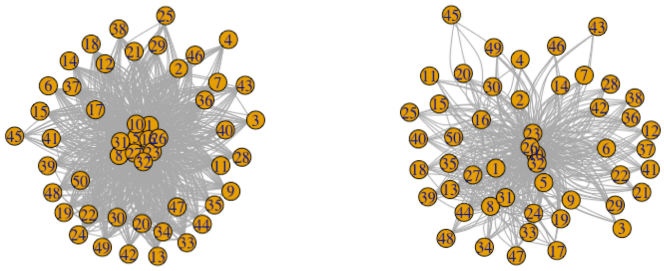

Figure 2-16: Giant connected component with threshold of 0.1 (left) and 0.85 (right) on the alternative influence metric.

attempt to extract question relationships from the fitted RNN output when a much simpler method achieves such strong results. The fact that the RNN produced a high AUC yet fails to establish links between related questions appears contradictory; this surprising result is fully explained in Section 2.5.

2.5 DKT and Unseen Question Sequences

The influence metrics contemplated in the previous section assumes the RNN produces sensible results when any question j and i are treated as question 1 and 2, even though the model was trained on a fixed sequence. This assumption is demonstrably false. Even small perturbations in the sequence of questions in the test set will cause the RNN to lose much or all of its performance. In the training process, the target was always set to the subsequent input. Shifting the target by +/- 1 position caused AUC to fall by 90-100% (using 0.5 as a base for AUC) in the 5-skill datasets. Shifting the target to be two questions ahead caused a 50-60% fall in AUC in the 2-skill datasets, and a fall around 70% when the input itself became the target. Rotating the question sequence in the test set had a less severe effect on AUC, as summarized in Table 2.2. Under rotation, the AUC loss was mainly due to a sharp decline in performance on questions that were preceded by questions that did not precede them in the training set; the predictive performance mostly recovered by the time unfamiliar questions

k (+/-)	1-5	6-12	>13
% AUC lost	4%-6%	7%-12%	13%-16%

Table 2.2: The test set question sequence was rotated by +/- k positions, from k = 1 to 25.

were five or more steps behind the target.

The transmission of historic information within the RNN is governed by the update function $\mathbf{h_t}(\mathbf{x_t}, \mathbf{h_{t-1}}) = \tanh(\mathbf{W_{hx}}\mathbf{x_t} + \mathbf{W_{hh}}\mathbf{h_{t-1}} + \mathbf{b_h})$. Consider the behavior of this function when there is no input $\mathbf{x_t}$ ($\mathbf{x_t} = \mathbf{0}$); let's call such a function $F(\mathbf{h_{t-1}})$. Suppose that question i and j belong to the same skill group, and $i > j$. The RNN is trained to transmit information from question j to question i through exactly i-j recursive applications of $F(\cdot)$; denote such an recursive composite function as $F_{i-j}(\cdot)$. Consequently, the RNN fitting process will seek weights and biases such that, for each k, the function $F_k(\mathbf{h_{t-1}})$ can transmit information between same-skill questions that are exactly k steps apart in the training data. However, for $k' \neq k$, $F_{k'}(\mathbf{h_{t-1}})$ could differ substantially from $F_k(\mathbf{h_{t-1}})$, which may explain why it is easy to fit an RNN that performs well on a fixed sequence but utterly fails under even slight shifting of the target.

When the RNN is trained on a fixed question sequence, it only has to learn weights and biases that produce accurate predictions for the particular target element in the prediction vector at each time step. As a result, the RNN will not be penalized for choosing weights and biases that produce terrible prediction vectors in each time step, other than the one element being tested. The space of solutions that produce good sequences of prediction vectors is much smaller than the space of solutions that produce good predictions for just the target. A good prediction vector will be more likely when the hidden layer implicitly encodes representations of the multi-skill knowledge state without transmitting superfluous information through the hidden layers. It will also be more likely when the hidden layer parameters give $F(\cdot)$ a level of stability (i.e., $F_k(\cdot) \approx F_{k+1}(\cdot)$).

We need a training methodology that preserves much of the RNN's performance even in the face of aggressive question sequence shuffling in the test set. One option

could be to shuffle the training data. For many education data sources, especially paper-based tests, the question sequence is fixed across all students. If learning is presumed to be happening over the course of the test, as is assumed in ITS settings, then randomly scrambling the training data question sequences could create problems. The underlying skill levels at the beginning and end may be different. A model trained using fully scrambled data will fail to estimate the rate of learning. One way to overcome this problem is to partially scramble the data. More specifically, choose the largest integer K such that the expected magnitude of skill change after K questions is low. Then break the ordered questions into $\lceil \frac{|Q|}{K} \rceil$ partitions, and in each batch shuffle the order of questions within the partition (see Figure 2-17).

This partial batch shuffling (PBS) technique preserves the order at the partition level so that the overall learning trend is not shuffled away, and yet it will force the training process to find parameters that enable the RNN to transmit information between related questions without requiring a constant distance between the questions. Partial shuffling aims to push the RNN towards a better representation of the true model, and not simply a model that does well on question sequences similar to those on which it was trained. We evaluated the PBS technique by examining its performance on five aggressive shufflings of the test dataset. The shuffled datasets had 80%, 85%, 90%, 95%, and 100% of their entries randomly shuffled, were rotated by +/- 1 positions, and had the target question shifted 0-4 steps ahead (chosen at random). The RNN was fit using PBS with five partitions, both with and without L1 regularization, and the AUC on the unshuffled test set was also computed for reference. In all cases, using PBS preserved most of the AUC in the face of aggressive shuffling. A small loss of AUC is inevitable when randomly shuffling the test set because it removes the global learning trend that is present in the unshuffled test data. The results for synthetic dataset #20 are presented in Table 2.3 to illustrate the strong performance of PBS.

To get stochastic convergence of the loss function, we had to run the PBS training process for 70 epochs instead of 25. A total of 1,400 distinct batch sequences were sampled during the training process, which is a small percentage, $\frac{1400}{(10!)^5} < 10^{-29}$ to

Figure 2-17: An example of partial shuffling. The original sequence of questions is divided into partitions of five consecutive questions, and in the new sequence the order of the partitions is the same, but the order within each partition is randomly shuffled.

be precise, of the sequences that could have been sampled under PBS. The set of all sequences that could have been sampled under PBS training represents an even smaller fraction of all possible sequences that could have been sampled for the test set: $\frac{(10!)^5}{50!} < 10^{-30}$. We can therefore conclude that PBS training helped the RNN approximate the true model without limitation to the tiny fraction of possible question sequences on which it was trained.

Model	Test AUC	Test AUC, Shuffled
5 skills, shuffled training batches	0.81 (L1 reg.) 0.81 (no L1 reg.)	0.78-0.80 (L1 reg.) 0.76-0.79 (no L1 reg.)
5 skills, normal training batches	0.77 (L1 reg.) 0.83 (no L1 reg.)	0.57-0.61 (L1 reg.) 0.56-0.58 (no L1 reg.)
2 skills, shuffled training batches	0.87 (no L1 reg.)	0.80-0.85 (no L1 reg.)
2 skills, normal training batches	0.88 (L1 reg.) 0.88 (no L1 reg.)	0.56-0.61 (L1 reg.) 0.51-0.61 (no L1 reg.)

Table 2.3: AUC with and without partial batch shuffling while training on the regular test set, with and without L1 regularization. The second column shows the AUC on the unshuffled test set, and the third column shows the range AUC values obtained on five aggressively shuffled test sets.

Interestingly, whereas L1 regularization did not improve AUC when the RNN was fit on static-sequence data and tested on either static or shuffled data, L1 regularization slightly improved AUC on shuffled test data when used in conjunction with PBS, and the magnitude of the benefit seemed to increase with the amount of shuffling in the test set (see Table 2.4). This empirical result raises interesting methodological questions that deserve further investigation, though we offer some thoughts on why

this result occurred. We suspect that regularization's role in enhancing the RNN fitting process is one of refinement rather than discovery. We speculate that regularization in the RNN might only work when the training method, with or without regularization, comes close to approximating the true generative model. If the true generative model is simple (e.g., up to five hidden skills in our case), then regularization may help push the training process from a good model with a few superfluous parameters towards a simpler model that is less prone to over-fitting. We suspect that regularization without PBS will usually be insufficient to approximate the "true model" because there will be too many locally optimal solutions on fixed data sequences that have low parameter L1 norms but that are of very low quality on unseen sequences.

The PBS method, with or without L1 regularization, is a novel general training methodology that could be used for any discrete response time-series governed by a small number of monotonically changing latent variables. The PBS method can be used to digitize fixed-sequence educational datasets, such as national exams, and to make accurate success predictions in adaptive individual tutoring systems when questions are presented in unfamiliar sequences. The PBS method may be helpful in ensuring that there is sufficient variation in the distances between all pairs of questions during training even for datasets that are known to not be in a strict sequence, as the default level of variation in the question sequences may be insufficient. For these reasons, it would make sense for PBS to become a standard training procedure for all future implementations of DKT.

% shuffled:	80%	85%	90%	95%	100%
No L1 Reg.	0.792	0.782	0.776	0.779	0.758
L1 Reg.	0.803	0.792	0.794	0.800	0.786
AUC gain	2.18%	1.99%	3.58%	4.26%	5.61%

Table 2.4: AUC was better with L1 regularizer than without, consistent across different percentages of test set shuffling.

2.6 New Model Framework

The preceding sections discussed many conceptual and computational limitations affecting a wide range of methods for modeling and estimating knowledge. This section develops a substantially new model that aims to overcome some of these limitations, with a view towards interpretability and supporting information exchange between ITS platforms. The most important conceptual difference is that the proposed model treats knowledge/skill mastery as a continuous latent variable with bounded support over [0,1], rather than binary (e.g., BKT, KST, and DCM), or over an unbounded scale with an arbitrary center (e.g., IRT). This section outlines a procedure to estimate latent skills from questions of varying difficulty and derives a closed-form solution for a statistically efficient, consistent, and unbiased estimator in the case of questions that engage one skill at a time. Computationally, the method is suitable for use on low-cost client devices. A general framework to handle complex causality in multi-skill items is proposed, though the computational procedures and statistical power considerations for multi-skill estimation are left as future work.

2.6.1 Model 1: One skill, Many Questions

The model developed in this section deals with the estimation of a single static skill parameter from a set of responses to questions that test that skill, and for which all skills required to succeed on the question can be viewed as prerequisites of the target skill. The new model uses the following notation:

- L: set of learners;
- K: set of pre-defined knowledge units;
- Q: set of questions, indexed by j;
- Q_k: subset of questions in knowledge unit k (we will suppose for now that questions belong to only one knowledge unit);
- Q_ℓ: array of question-response data for learner ℓ (missing values allowed);

- $Q_{k\ell}$: array of question-response data for learner ℓ in knowledge unit k (missing values allowed);
- j: the index for questions.
- c: the index for question classes. A question class is a set of questions that have the exact same difficulty and that belong to the same knowledge unit. Responses to questions from the same class be viewed iid random variables.
- $q_{\ell j}$: learner ℓ's response to question j, encoded as 1 for correct responses and 0 for incorrect (if there are multiple responses from the question class, then just assume there's another index level).
- $R_{k\ell} = \{j | q_{j\ell} = 1 \text{ and } q_j \in Q_k\}$;
- $W_{k\ell} = \{j | q_{j\ell} = 0 \text{ and } q_j \in Q_k\}$.

Note that in this section Q represents data on a question, and is not to be confused with the skill-question Q matrix for KT methods introduced in the literature review.

Likelihood Function

The probability distribution of question-response data can be viewed as a mixture of two distributions: (1) one produced from uninformed guessing, and (2) one produced from a learner who has mastery of the knowledge unit and all strictly necessary prerequisite material. In the latter case, we would still expect some amount of careless error unrelated to the knowledge unit being tested, e.g., by misreading a question. Let $\lambda_{k\ell} \in [0, 1]$ be the parameter that allocates weight between the two distributions. Let G_j be the probability of a lucky guess on question j, and X_j be the probability of a correct response on question j for a learner with full mastery of the knowledge unit and pre-requisite material. Then, for $q_j \in Q_{k\ell}$

$$\mathbb{P}(q_{\ell j} = 1) = (1 - \lambda_{k\ell})G_j + \lambda_{k\ell} X_j \tag{2.29}$$

Note that $(1 - \lambda_{k\ell})$ can be interpreted as the percentage of the time that the learner does not know how to answer the question because of a deficiency in required knowledge.

Now consider the likelihood of $Q_{k\ell}$, on the assumption that the responses are independent, conditioning on $\lambda_{k\ell j}$:

$$\mathcal{L}(\lambda_{k\ell}, G, X|Q_{k\ell}) = \prod_{j|q_{\ell j} \in Q_{k\ell j}} q_{\ell j}\mathbb{P}(q_{\ell j} = 1) + (1 - q_{\ell j})(\mathbb{P}(q_{\ell j} = 0)) \tag{2.30}$$

Homogeneous Known Question Difficulty

Suppose that G and X are known. Further suppose that all of the questions in the knowledge unit are equally difficult and that the probability of success is independent, conditional on $\lambda_{k\ell}$. That is, there is a common G_k and X_k for all $q_j \in Q_k$.

In this case, the maximum likelihood estimate has a unique solution for $\lambda_{k\ell}$, which is given by setting $\hat{Pr}(q_{j\ell} = 1) = \frac{|R_{k\ell}|}{|W_{k\ell}|+|R_{k\ell}|}$ (empirical success rate), where $|R_{k\ell}|$ is the number of right responses in $Q_{k\ell}$ and $|W_{k\ell}|$ is the number of wrong responses. Denote $\hat{Pr}(q_{j\ell})$ as $\hat{p}_{k\ell}$, and we obtain an expression for $\lambda_{k\ell}$ using equation 2.29:

$$\hat{\lambda}_{k\ell} = \begin{cases} 0, & \text{if } \hat{p}_{k\ell} \leq G_j \\ 1, & \text{if } \hat{p}_{k\ell} \geq X_j \\ \dfrac{\hat{p}_{k\ell} - G_k}{X_k - G_k}, & \text{o.w.} \end{cases} \tag{2.31}$$

Heterogeneous Known Question Difficulty

Suppose that the probability of succeeding on a question is still conditionally independent on other questions in the knowledge unit, given $\lambda_{k\ell}$, but that question difficulty varies. The likelihood function given by equation 2.30 can be expanded as:

$$\mathcal{L}(\lambda_{k\ell}|Q_{k\ell}) = \prod_{j \in R_{k\ell}} \left[(1-\lambda_{k\ell})G_j + \lambda_{k\ell}X_j\right] \prod_{j \in W_{k\ell}} \left[1-(1-\lambda_{k\ell})G_j - \lambda_{k\ell}X_j\right]$$

$$\Rightarrow log(\mathcal{L}) = \sum_{j \in R_{k\ell}} log\left((1-\lambda_{k\ell})G_j + \lambda_{k\ell}X_j\right) + \sum_{j \in W_{k\ell}} log\left(1-(1-\lambda_{k\ell})G_j - \lambda_{k\ell}X_j\right)$$

$$\Rightarrow \frac{\partial log(\mathcal{L})}{\partial \lambda_{k\ell}} = \sum_{j \in R_{k\ell}} \frac{X_j - G_j}{(1-\lambda_{k\ell})G_j + \lambda_{k\ell}X_j} + \sum_{j \in W_{k\ell}} \frac{G_j - X_j}{1-(1-\lambda_{k\ell})G_j - \lambda_{k\ell}X_j}$$

Setting this quantity to zero gives us the first order condition for the MLE of $\lambda_{k\ell}$. To solve for λ, we can multiply all terms by the various denominators. To simplify notation, we can write the denominators as $\mathbb{P}(q_{\ell j} = 1)$ for $j \in R_{k\ell}$ and its complement for $j \in W_{k\ell}$, giving

$$\sum_{j \in Q_{k\ell}} \left[(X_j - G_j) \prod_{h \in R_{k\ell} \setminus j} \mathbb{P}(q_{\ell j} = 1) \prod_{h \in W_{k\ell} \setminus j} \left(-(\mathbb{P}(q_{\ell j} = 0)\right)\right] = 0 \tag{2.32}$$

The expression on the left-hand-side of equation 2.32 is a polynomial in $\lambda_{k\ell}$ of degree up to $n-1$, not guaranteed to be concave, where n is the number of observations in $Q_{k\ell}$. Consequently, there may be multiple stationary points in the likelihood function. Since we want the global maximum given by a single variable over the interval [0,1], one could solve for all stationary points algebraically and find the argmax, or numerically approximate the global optimum using a global line search algorithm.

Alternate Method for Estimating $\lambda_{k\ell}$

This section derives a closed-form solution for finding a statistically efficient and consistent estimator for $\lambda_{k\ell}$ from questions of varying difficulty. Suppose unit k has multiple question classes, indexed by c, with distinct difficulty levels (characterized by G_c and X_c). If we happened to only have data from question class c, then the maximum likelihood estimate for $\lambda_{k\ell}$ would be

$$\hat{\lambda}_{k\ell(c)} = \begin{cases} 0, & \text{if } \hat{p}_{c\ell} \leq G_c \\ 1, & \text{if } \hat{p}_{c\ell} \geq X_c \\ \dfrac{\hat{p}_{c\ell} - G_c}{X_c - G_c}, & \text{o.w.} \end{cases} \tag{2.33}$$

where $\hat{p}_{c\ell}$ is the empirical success rate for question class c. The empirical success rate can fall outside the range $p_{c\ell} \in [G_c, X_c]$. In a step to follow, it will be helpful to replace $\hat{p}_{c\ell}$ with a proportion $\tilde{p}_{c\ell}$ that is constrained to be within the range of p:

$$\tilde{p}_{c\ell} = \begin{cases} G_c, & \text{if } \hat{p}_{c\ell} \leq G_c \\ X_c, & \text{if } \hat{p}_{c\ell} \geq X_c \\ \hat{p}_{c\ell}, & \text{o.w.} \end{cases} \tag{2.34}$$

Since $\hat{p}_{c\ell} \xrightarrow{a.s.} p_{c\ell} \quad \forall c$, the estimator $\hat{\lambda}_{k\ell(\cdot)}$ generated by each question class is a statistically consistent estimator for $\lambda_{k\ell}$. Therefore, a weighted average of the $\hat{\lambda}_{k\ell(c)}$ estimators will also be statistically consistent. The data from each question class will not necessarily be equally informative, however. We want a weighting the produces an estimator that is efficient, in the sense that it has the lowest mean squared error, given a set of observations, amongst all statistically consistent estimators.

We can find such an estimator by solving a quadratic constrained optimization problem for the the weights on each question class' estimator to minimize the standard error of the overall estimate. We assume that successes on questions in unit k are independent after conditioning on $\lambda_{k\ell}$.

Let $\sigma^2_{k(c)}$ be the true variance of the estimator $\hat{\lambda}_{k\ell(c)}$, and let $\pi_{c\ell}$ be a random variable representing the observed proportion of successes in question class c after n_c trials, and $p_{c\ell} \in [G_c, X_c]$ be the true underlying probability of success for learner ℓ in question class c. Then the following holds:

$$\min\left(\frac{G_c(1-G_c)}{n_c(X_c-G_c)^2}, \frac{X_c(1-X_c)}{n_c(X_c-G_c)^2}\right) \leq \sigma^2_{k(c)} \leq \text{Var}(\frac{\pi_{c\ell}-G_c}{X_c-G_c}) = \frac{p_{c\ell}(1-p_{c\ell})}{n_c(X_c-G_c)^2} \tag{2.35}$$

with the inequality holding because the support of $\hat{\lambda}_{k\ell(c)}$) is truncated at 0 and 1, whereas the ratio on the LHS can vary more widely since $\pi_{c\ell}$ can fall outside the support of $p_{c\ell}$.

Unfortunately, the true value of $p_{c\ell}$ is unknown, so this particular upper bound cannot be computed. Instead, one could assume maximum variance (as if $p_{c\ell} = 0.5$) to obtain $(0.5)^2$ in the RHS bound's numerator. This looser bound will always hold. Alternatively, one could assume $p_{c\ell} = 0.5w + \pi_{c\ell}(1-w), w \in [0,1]$ and accept that with some reasonably low probability the upper bound will be violated. This probability would decrease with $w \to 1$ and with $\sqrt{n}$.

Without loss of generality regarding the particular choice of consistent estimator $\hat{\sigma}^2_{k(c)}$ used to approximate the true $\sigma^2_{k(c)}$, consider the problem of choosing a weighted average of the estimators $\hat{\lambda}_{k\ell(c)}$ that minimizes the variance of the combined estimate. Assume that the $\hat{\lambda}_{k\ell(c)}$ are independent so that the variance-covariance matrix is diagonal. Denoting the weight on the estimator from class c with θ_c, we express the problem of choosing weights as the following quadratic optimization problem:

$$\begin{aligned} \underset{\theta}{\text{minimize}} \quad & \sum_c \hat{\sigma}^2_{k\ell(c)}\theta_c^2 \\ \text{subject to} \quad & \sum_c \theta_c = 1 \\ & \theta \geq 0. \end{aligned} \tag{2.36}$$

The optimal set of weights when averaging a set of independent consistent estimators to obtain a minimum-variance estimator is given by:

$$\forall i, \quad \theta_i^* = \frac{\prod_{c \neq i} \hat{\sigma}^2_{k\ell(c)}}{\sum_c \prod_{m \neq c} \hat{\sigma}^2_{k\ell(m)}}, \tag{2.37}$$

and by dividing the numerator and denominator by $\prod_c \hat{\sigma}^2_{k\ell(c)}$ we obtain an inverse-variance weighting:

$$\forall i \quad \theta^*_i = \frac{1}{\hat{\sigma}^2_{k\ell(i)}} \left(\sum_c \frac{1}{\hat{\sigma}^2_{k\ell(c)}} \right)^{-1}. \tag{2.38}$$

The inverse variance weighting method is not unique to this application; it is also used, for example, in meta-analysis studies, where different estimates of a treatment effect from different studies are combined to get an overall effect estimate [Sanchez-Meca and Marín-Martínez, 1998]. The optimal value of variance minimization problem 2.36 can easily be shown to be $\left(\sum_c \frac{1}{\hat{\sigma}^2_{k\ell(c)}} \right)^{-1}$:

$$\text{Let } \omega = \sum_c \frac{1}{\hat{\sigma}^2_{k\ell(c)}}; \text{ then } \sum_i \hat{\sigma}^2_{k\ell(i)} \theta^{2*}_i = \sum_i \hat{\sigma}^2_{k\ell(i)} \left(\frac{1}{\hat{\sigma}^2_{k\ell(i)}} \omega^{-1} \right)^2$$
$$= \sum_i \hat{\sigma}^2_{k\ell(i)} \left(\frac{1}{\hat{\sigma}^4_{k\ell(i)}} \omega^{-2} \right) = \omega^{-2} \sum_i \frac{1}{\hat{\sigma}^2_{k\ell(i)}} = \omega^{-2}\omega = \omega^{-1} = \left(\sum_c \frac{1}{\hat{\sigma}^2_{k\ell(c)}} \right)^{-1}.$$

The result 2.38 follows from using a Langrange multiplier for the constraint and solving the first-order condition. It is the same variance-minimizing weighting used in weighted least-squares regression and in meta-analysis studies [Hartung et al., 2008].

The actual efficiency of this estimator depends on the accuracy of the $\hat{\sigma}^2_{k\ell(c)}$. If they are accurate, then θ^* is the weight vector that produces the most efficient estimator for $\lambda_{k\ell}$, and the objective function value from problem 2.36 gives that estimators true variance. If they are inaccurate, however, then the solution is not optimal, and the variance of the chosen estimator is not known but simply approximated by the objective function value. Fortunately, under mild conditions, the estimated variances $\hat{\sigma}^2_{k\ell(c)}$ can generate a tight upper bound on the true variance for the aggregate estimator $\hat{\lambda}_{k\ell}$ induced by θ^*. Such an upper bound allows us to construct a confidence interval around $\hat{\lambda}_{k\ell}$ conservatively. The necessary assumption is that $\hat{\sigma}^2_{k\ell(c)} \geq \sigma^2_{k\ell(c)} \forall c$, (e.g., by shifting the estimate of $p_{c\ell}$ towards 0.5).

Lemma 2. *The optimal value of problem (2.36) using estimated variances that satisfy*

$\hat{\sigma}^2_{k\ell(c)} \geq \sigma^2_{k\ell(c)}\ \forall c$ *is an upper bound to the true variance of the estimator for* $\lambda_{k\ell}$ *induced by the weights* θ^*.

Proof: $\sum_c \hat{\sigma}^2_{k\ell(c)}\theta^{*2}_c \geq \sum_c \sigma^2_{k\ell(c)}\theta^{*2}_c = \text{Var}(\hat{\lambda}_{k\ell})|_{\theta^*}$ by independence of the $\hat{\lambda}_{k\ell(c)}$.

Lemma 3. *The optimal value of problem (2.36) using estimated variances* $\hat{\sigma}^2_{k\ell(c)}$ *is an upper bound to the optimal solution to problem (2.36) using any other variances* $\sigma^2_{k\ell(c)}$ *satisfying* $\sigma^2_{k\ell(c)} \leq \hat{\sigma}^2_{k\ell(c)}\ \forall c$.

Proof: For ease of notation, we will drop the indexing by k and ℓ and will write $\hat{\sigma}^2_{k\ell(c)}$ simply as $\hat{\sigma}^2_c$ and $\sigma^2_{k\ell(c)}$ as σ^2_c.
We want to show the following:

$$\begin{aligned} \max_{\sigma^2} \min_{\theta} \sum_c \sigma^2_c\theta^2_c \quad &\leq \quad \min_{\hat{\theta}} \sum_c \hat{\sigma}^2_c\hat{\theta}^2_c \\ \text{s.t. } \boldsymbol{\sigma}^2 \leq \hat{\boldsymbol{\sigma}}^2 \quad &\qquad \text{s.t. } \sum_c \hat{\theta}_c = 1 \\ \sum_c \theta_c = 1, \boldsymbol{\theta} \geq 0 \quad &\qquad \hat{\boldsymbol{\theta}} \geq 0 \end{aligned} \tag{2.39}$$

We can show this by showing that the closed form optimal solution $\hat{\boldsymbol{\theta}}^*$ to the LHS problem satisfies the KKT conditions for a concave maximization problem equivalent to the RHS problem, and is therefore an optimal solution to the RHS problem.

Recall from problem (2.36) that the optimal solution to the inner minimization problem is given by equation (2.37), which can be re-written as follows:

$$\theta^*_i = C\prod_{c\neq i}\sigma^2_c, \text{ where } C = \frac{1}{\sum_c \prod_{m\neq c}\sigma^2_m}.$$

Problem (2.39) can therefore be re-written as a maximization problem subject to the inner minimization problem's optimality condition (which happens to also satisfy the inner problem's feasibility condition of $\sum_i \theta_i = 1$ as long as $\sigma^2_i > 0\ \forall i$ (if this condition doesn't hold, then the true value of $\lambda_{k\ell}$ would already be known anyway):

$$\begin{aligned}
&\underset{\sigma^2}{\text{Max}} \sum_c \theta_c^{*2}\sigma_c^2 \\
&\text{subject to:} \\
&\boldsymbol{\sigma}^2 \leq \hat{\boldsymbol{\sigma}}^2 \qquad (2.40)\\
&C\prod_{c\neq i}\sigma_c^2 = \theta_i^*, \qquad \forall i = 1, ..., |C| \\
&\boldsymbol{\sigma}^2 \geq 0 \qquad \text{(never binding in practice)}
\end{aligned}$$

Substitute the optimality constraint on θ^* into the objective function to get an maximization problem expressed only in terms of σ^2.

$$\begin{aligned}
&\underset{\sigma^2}{\text{Max}} \sum_i \sigma_i^2 \left(C\prod_{c\neq i}\sigma_c^2\right)^2 \quad = \quad \underset{\sigma^2}{\text{Max}} \sum_i C^2 \left(\prod_c \sigma_c^2\right)\left(\prod_{c\neq i}\sigma_c^2\right) \\
&\text{s.t. } \boldsymbol{\sigma}^2 \leq \hat{\boldsymbol{\sigma}}^2 \qquad\qquad \text{s.t. } \boldsymbol{\sigma}^2 \leq \hat{\boldsymbol{\sigma}}^2 \\
&\boldsymbol{\sigma}^2 \geq 0 \qquad\qquad \boldsymbol{\sigma}^2 \geq 0
\end{aligned}$$

We can simplify objective function:

$$\begin{aligned}
&\Rightarrow \; C^2\left(\prod_c \sigma_c^2\right)\sum_i\prod_{c\neq i}\sigma_c^2 \\
&\quad \text{Since } \sum_i\prod_{c\neq i}\sigma_c^2 = C^{-1}, \text{the objective is: } C\prod_c\sigma_c^2 = \frac{\prod_c \sigma_c^2}{\sum_k \prod_{j\neq k}\sigma_c^2}
\end{aligned}$$

This objective function is concave in $\boldsymbol{\sigma}^2$, which we can treat as our decision vector rather than $\boldsymbol{\sigma}$ since $\boldsymbol{\sigma}$ only appears squared throughout the problem and is non-negative. For ease of notation, let's denote $\boldsymbol{\sigma}^2$ as $\mathbf{x}$.

We can see the function's concavity in $\mathbf{x}$ by examining whether the logarithm of the function is concave. The log of the function is $\sum_c x_c - \log\left(\sum_k \prod_{c\neq k} x_c\right)$. The first term is linear (so both concave and convex). Observe that each product $\prod_{c\neq k} x_c$ is convex. We know that the sum of convex functions is convex, that the logarithm of a convex function is convex, and that the opposite of a convex function is concave,

so we conclude that $-\log\left(\sum_k \prod_{c\neq k} x_c\right)$ is concave. Therefore the objective function is concave.

Next, we examine the KKT conditions.

Consider the problem

$$\begin{aligned}
&\underset{\mathbf{x}}{\text{Max}}\ f(\mathbf{x}) \\
&f(\mathbf{x}) = \frac{\prod_i x_i}{\sum_k \prod_{c\neq k} x_c} \\
&\text{subject to} \\
&\mathbf{g}_1(\mathbf{x}) : \mathbf{x} - \hat{\mathbf{x}} \leq \mathbf{0} \\
&\mathbf{g}_2(\mathbf{x}) : -\mathbf{x} \leq \mathbf{0}.
\end{aligned}$$

The KKT conditions are:

$$\begin{aligned}
&\nabla \mathbf{f}(\mathbf{x}^*) = \nabla \mathbf{g}_1(\mathbf{x}^*)\boldsymbol{\mu}_1 + \nabla \mathbf{g}_2(\mathbf{x}^*)\boldsymbol{\mu}_2 \\
&\boldsymbol{\mu}_1 \circ g_1(\mathbf{x}^*) = \mathbf{0} \text{ (element-wise multiplication)} \\
&\boldsymbol{\mu}_2 \circ g_2(\mathbf{x}^*) = \mathbf{0} \\
&\boldsymbol{\mu}_1 \geq \mathbf{0} \\
&\boldsymbol{\mu}_2 \geq \mathbf{0}
\end{aligned}$$

Since the constrained maximization problem is concave and the inequality constraints are continuously differential convex functions, a point that satisfies the KKT conditions is a globally optimal solution to the problem. Note that $\mathbf{g}_1(\mathbf{x})$ is a matrix whose columns are the gradients of the respective elements of $\mathbf{g}_1(\mathbf{x})$, namely each $x_c - \tilde{x}_c$ differentiated with respect to $\mathbf{x}$. It is easy to see that $\nabla \mathbf{g}_1(\mathbf{x}) = \mathbf{I}$, the identity matrix.

The next step of the proof is to show that $\mathbf{x}^* = \tilde{\mathbf{x}}$ satisfies the KKT conditions. This proposed solution gives $g_1(\tilde{\mathbf{x}}) = 0$, so it satisfies the first set of complementary slackness conditions. Because we have assumed all variances are strictly positive, we have $g_2(\mathbf{x}) : -\tilde{\mathbf{x}} < \mathbf{0}$, so we would have to set $\boldsymbol{\mu}_2 = 0$ to satisfy the second set of complementary slackness constraints. It remains to show that that there exists

$\boldsymbol{\mu}_1 \geq \mathbf{0}$ satisfying $\mathbf{f}(\tilde{\mathbf{x}}) = \nabla \mathbf{g}_1(\mathbf{x})\boldsymbol{\mu}_1 \iff \nabla \mathbf{f}(\tilde{\mathbf{x}}) = \mathbf{I}\boldsymbol{\mu}_1 \iff \nabla \mathbf{f}(\tilde{\mathbf{x}}) = \boldsymbol{\mu}_1$.

Such a $\boldsymbol{\mu}_1$ exists if the smallest possible value for $\nabla_c f(\tilde{\mathbf{x}})$ is non-negative, over any $\tilde{\mathbf{x}} > \mathbf{0}$, $\forall j = 1, ..., |J|$. The elements of $\mathbf{x}$ are interchangeable, so if $\underset{\tilde{\mathbf{x}}>\mathbf{0}}{\text{Min}} \nabla_c f(\tilde{\mathbf{x}}) \geq 0$ for an arbitrary j then the same result holds for all j. By the quotient rule, which produces a positive denominator that can be ignored for the purposes of assessing non-negativity, we have:

$$\nabla_c f(\tilde{\mathbf{x}}) = \frac{\left(\prod_{i \neq c} \tilde{x}_i\right) \sum_k \left(\prod_{i \neq k} \tilde{x}_i\right) - \left(\prod_i \tilde{x}_i\right) \sum_{k \neq c} \left(\prod_{i \neq \{k,j\}} \tilde{x}_i\right)}{(\cdot)^2}.$$

Simplify the numerator by dividing both terms by $\prod_{i \neq c} \tilde{x}_i$ to get

$$\begin{aligned}
&\sum_k \left(\prod_{i \neq k} \tilde{x}_i\right) - \tilde{x}_c \sum_{k \neq c} \left(\prod_{i \neq \{k,j\}} \tilde{x}_i\right) \\
&= \sum_k \left(\prod_{i \neq k} \tilde{x}_i\right) - \sum_{k \neq c} \left(\prod_{i \neq k} \tilde{x}_i\right) \\
&= \prod_{i \neq c} \tilde{x}_i > 0 \text{ since } \tilde{\mathbf{x}} > \mathbf{0} \\
&\Rightarrow \nabla_c f(\tilde{\mathbf{x}}) > 0, \ \forall \tilde{\mathbf{x}} > 0 \\
\Rightarrow \boldsymbol{\mu}_1 = \nabla \mathbf{f}(\tilde{\mathbf{x}}) > 0 &\text{ satisfies KKT conditions}
\end{aligned}$$

Therefore, the worst-case true variance vector $\boldsymbol{\sigma}^2$ that nature can select in LHS of problem (2.39) is $\hat{\boldsymbol{\sigma}}^2$, which proves our bound problem (2.39) is tight. This completes the proof for Lemma 3.

There are several benefits to estimating the variance of a point-estimate for $\lambda_{k\ell}$ in this manner. First, we can use this framework to detect model specification mistakes in a computationally efficient manner using standard hypothesis testing techniques. For example, data from different question classes may point to markedly different

$\lambda_{k\ell}$ values, which would suggest that some of the questions may be testing a skill other than k, or that a linear mixing between guessing and slippage probabilities is too simple a model. Second, we can construct confidence intervals for $\lambda_{k\ell}$ in a computationally efficient manner, without having to rely on bootstrapping. Third, knowing the variance $\hat{\lambda}_{k\ell(c)}$ allows educational software to identify which class of question it should present next to obtain maximum incremental information gain about the knowledge level.

Testing on Simulated Data

Simulated data were generated to compare the speed and accuracy of the closed-form inference procedure for $\lambda_{k\ell}$ to maximum-likelihood estimation. In each iteration of the simulation, there were 50 observations coming from 5 to 10 different question classes, each with a different β_c and η_c values. 1000 simulations were performed.

The MLE estimate was solved in Matlab by taking the maximum of the likelihood function evaluated at all stationary points within [0,1], inclusive of the boundaries. The closed-form solution to the estimator using inverse-variance weights was implemented in R. The average MSE of the estimators was 0.0146 for the MLE method and 0.0139 for the closed-form method. The MLE method in Matlab took approximately 25 times longer than inverse-variance weights in R, which took under ten milliseconds.

2.6.2 Preliminaries for a Continuous and Bounded Multi-Skill Model

The remaining sections discuss several considerations relevant to the construction of a multi-skill model that builds upon section 2.6.1. These sections can be viewed as a preliminary discussion laying the foundation for future work.

One important consideration is the potentially complex nature of the joint distribution of a large number of knowledge state parameters. The $\lambda_{k\ell}$ parameter represents how well learner ℓ knows unit k. More precisely, $1-\lambda_k\ell$ represents the probability that learner ℓ will have to guess an answer in unit k because of a conceptual gap or misun-

derstanding within knowledge unit k, or in one of the skills that is a prerequisite of k. There are good reasons to believe that the distribution for λ will be U-shaped in most populations of learners. Chances are, a randomly selected student has either never learned a topic, or has learned it, with relatively few students being in the middle of the learning process. Whenever this is true, most probability mass will be concentrated around 0 and 1. Intermediary values would occur with greater frequency in an ITS if a human-led or machine-led process attempted to match students to particular questions or topics at their knowledge frontier. Intermediary values for λ will also be more likely for units that are rarely required or revisited in subsequent units, and that are forgotten or become fuzzy without revision. Even in these circumstances, bell-shaped distributions for λ are not plausible priors.

The joint distribution for two skills λ_k and $\lambda_{k'}$ will concentrate mass in the four corners of $[0,1]^2$ if they are independent. If they have a prerequisite relationship, mass will concentrate at [0,0] and [1,1], with some mass around the line segment [0,0] to [1,0] (if the prerequisite is the first coordinate), around the line segment [1,0] to [1,1], but not around [0,1]. Knowing prerequisite relationships may therefore allow us to estimate or bound $\lambda_{k\ell}$ using data reflecting performance on skills other than k.

Let $PR(k)$ denote the set of strict prerequisites for unit k. Let γ_k be a parameter representing the mastery of new knowledge in unit k that is not automatically inherited or known by mastering units in $PR(k)$. If the prerequisite units' λs are independent, then:

$$\lambda_{k\ell} = \gamma_{k\ell} \prod_{i \in PR(k)} \lambda_{i,\ell} \tag{2.41}$$

Alternatively, $\lambda_{k,\ell}$ could be conditionally independent of most of its prerequisites, given a subset of direct "parent" prerequisites. If these parents are independent, we could use the same formula, taking the product over the subset of immediate prerequisites only.

If the prerequisites are not independent (e.g., if they have overlapping prerequisites of their own), then this relationship does not hold, because the right-hand side

will double-count the chance of making a particular conceptual mistake if the root knowledge unit where the mistake lies is inherited by more than one prerequisite. In this case, the correct relationship is between λ and the marginal knowledge mastery parameters γ:

$$\lambda_{k\ell} = \gamma_{k\ell} \prod_{i \in PR(k)} \gamma_{i,\ell} \tag{2.42}$$

Since prerequisite units are likely to have common prerequisites themselves, one would most likely have to use this relationship, which means the γ parameters would be the real quantities of interest. Let $Need(k)$ denote the set of knowledge units that are "strict" descendants of unit k (that is, k is in a chain of strict prerequisites leading to elements in $Need(k)$. Then $\gamma_k \geq \lambda_{k'} \forall k' \in Need(k)$. In other words, the chance of a conceptual error in the "marginal knowledge content" of knowledge unit k should not exceed the chance of a conceptual error in a more advanced unit that requires unit k. This inequality can be viewed as soft analogue to the prerequisite structure used in Knowledge Space Theory.

Although prerequisite structures can be identified by subject matter experts, for large and complex domains it may be necessary to use data-driven causal inference methods to handle the scale. One challenge for data-driven causal inference is the possibility of complex causation. Complex causation can arise with knowledge units or sets of units that are any of the following:

1. not strictly necessary but which raise the expected score in unit k;
2. strictly necessary for learning k but insufficient;
3. sufficient for learning k but not necessary;
4. more than one set of skills sufficient basis for learning k, and the skills within each set are necessary for that set to be sufficient, but are unnecessary on their own.

The first complication is handled naturally through regression. The second re-

quires a conjunctive model, of which there are several contemplated in the literature. The third and fourth scenarios are more complex; they refer to situations in which prerequisites are not strict, in the sense that k can be learned so long as a suitable set of non-strict prerequisites are known, of which there may be many. This setting characterizes most problem-solving tasks have multiple solution pathways. During the literature review, no statistical model capable of recovering such complex causal relationships from data was found. Section 2.8 introduces a model that may be capable of representing complex causality and will be the subject of further work.

2.7 Estimating Mastery Success Rate

The careless error rate $1 - X_c$ can be viewed as the empirical failure rate on questions from class c of learners who have mastered units that require unit k. However, such learners may not share the same asymptotic success rate. An incorrect response to a question in knowledge unit k could be attributed to the following sources:

- A conceptual mistake in the material of knowledge unit k or PR(k). Let's call this "prerequisite error." The probability of such an error is $\lambda_{k\ell}$. In fitting X_c, we can exclude data from learners whose confidence interval (or credible interval) for $\lambda_{k\ell}$ has a lower bound below some threshold (such as 0.95). If we are too strict, we may not have enough data to work with, but if we are not strict enough, then we may introduce identifiability problems.

- A mistake in a skill that is neither in PR(k) nor need(k) but that happens to be relevant to the question (even if unintended by the question creator). Let's call this "other unit error" and let other(k) denote the set of skills not in PR(k) or need(k).

- All other learner-specific sources of error (e.g., error in reasoning, conscientiousness, general problem-solving skills, misreading question, deficiencies in knowledge units or skills exogenous to the γ's). Let's call this "general error."

The probability of "general error" could vary from learner to learner, and it could also vary with learner covariates (e.g., age) and affect.

The parameter X_c can be treated as a random variable whose conditional expectation can be estimated via logistic regression as follows:

$$log\left(\frac{X_{c\ell}}{1-X_{c\ell}}\right) \sim \mu_k + \alpha_\ell + \delta_c + G'_{other(k)}\lambda_{other(k)} + \epsilon,\ \forall \ell | \hat{\lambda}_{k\ell,LB} > \hat{\lambda}_{min}.$$

In this model, α_ℓ is an individual aptitude fixed effect reflecting general reasoning abilities and conscientiousness, μ_k is a subject difficulty fixed-effect. δ_c is a question-class (or question) difficulty fixed-effect (alternatively, we could specify μ_{kj}), $\lambda_{other(k)}$ is a vector of lambda's in the set $other(k)$, $G_{other(k)}$ are the corresponding parameters, and ϵ is the remaining error term. $\hat{\lambda}_{k\ell,LB}$ is a lower bound on the estimate for the skill level, and $\hat{\lambda}_{min}$ is the threshold for including the learner in the dataset of those who are deemed to have "mastered" skill k. If $|other(k)|$ is large, the model may have too many parameters and regularization or variable selection algorithms may be needed to identify a relevant subset within $|other(k)|$ and avoid multicollinearity.

2.8 Model 2: Skills and Questions with Many-to-Many Relationship

This section develops a family of functions for modeling for the probability of success on complex items while representing latent skills over the continuous and bounded domain of $[0,1]^{|}K|$. The family of functions has not been analyzed for identifiability or appropriate parameter fitting procedures, and so the discussion that follows is merely exploratory.

Suppose performance on question class Q_c is conditionally dependent on multiple knowledge units. As with questions that depend on a single knowledge unit, we can model the conditional distribution through the weighting W_c between guessing performance ($W_c = 0$) and mastery performance ($W_c = 1$) on Q_c, only now W_c will depend on multiple knowledge units through a function $S_c : \lambda \rightarrow W_c$.

Suppose that the prerequisite relationships between skills forms a directed acyclic graph (DAG), which we will refer to as a knowledge graph. Let $PA(Q_c)$ represent the set of parent nodes for Q_c in the corresponding knowledge graph.

First consider an additive skill model:

$$S_c(\lambda) = \sum_{k \in PA(Q_c)} w_{ck} \lambda_k^{p_{ck}}, ||w||_1 = 1, w \geq 0, p \geq 0. \tag{2.43}$$

In the additive model, a knowledge unit contributes to mastery independently of the other units, up to some fraction of total performance. They can be unequally weighted, and the marginal improvement in performance can be increasing ($p_{ck} > 1$) or decreasing $\big(p_{ck} \in (0, 1)\big)$ in λ_k. Low values of p imply that sizable probability gains can be achieved with low and modest levels of mastery of the skill, and higher values of p imply that sizable probability gains will not be achieved until a high level of skill mastery is attained. As with the more sophisticated IRT models, this additive model allows for multiple skill dimensions and multiple dimensions of difficulty with respect to those skills, but skills are compensatory, albeit up to a limit.

Conjunctive model (AND gate):

$$S_c(\lambda) = \prod_{k \in PA(Q_c)} \lambda_k^{p_{ck}}, p \geq 0. \tag{2.44}$$

In the conjunctive model, linked knowledge units are not merely helpful, but required. In the educational data mining literature, when knowledge states are modeled with binary components, the conjunctive function produces an AND gate.

Disjoint model (OR gate):

$$S_c(\lambda) = \max_{k \in PA(Q_c)} \left(\lambda_k^{p_{ck}}\right), p \geq 0. \tag{2.45}$$

In this model, any linked knowldge unit is sufficient, none are necessary, performance is determined by the strongest unit.

Complex Causal Sets: Some questions have multiple solution methods that require

different combinations of knowledge. Let $M(PA(Q_c))$ be the set of subsets of $PA(Q_c)$ that represent distinct solution methods to question Q_c. We could suppose that the probability of success is equal to the solution method that has the highest probability of success for the student:

$$S_c(\lambda) = \max_{m \in M} \left(\prod_{k \in m} \lambda_k^{p_{km}} \right), p \geq 0. \tag{2.46}$$

Alternatively, these different solution methods might not be completely disjoint. For example, Q_c might be a question class that is not perfectly homogeneous in its difficulty or knowledge requirements, despite whatever intentions to make it so. Let $f(m) = \left(\prod_{k \in m} \lambda_k^{p_{km}} \right)$ represent method m's contribution to the function S, either in an additive model, or a a partially disjoint model.

Additive model with compound methods:

$$S_c(\lambda) = \sum_{m \in M} w_m f(m), ||w_m||_1 = 1, w \geq 0 \tag{2.47}$$

Partially disjoint model:

$$S_c(\lambda) = ||\mathbf{f}||_N \tag{2.48}$$

In the partially disjoint model, we take the L_N-norm of the vector whose $|M|$ elements are the $f(m)$ functions, instead of taking L_∞ ($\max_{m \in M}$) . A suitably chosen N controls how much weight is given to the largest $f(m)$; a larger N gives more weight to the largest $f(m)$ and $N = 1$ is equivalent to the additive model with compound methods.

Likelihood Function

The likelihood function for the general model can be obtained from the following relationship:

$$\mathbb{P}(q_{\ell j} = 1) = (1 - S_c(\lambda_\ell))G_c + S_c(\lambda_\ell)X_c. \tag{2.49}$$

The likelihood function will be non-convex and may have many local maxima. It remains to be seen whether, in the presence of sufficiently large datasets, gradient descent algorithms, stochastic or otherwise, will tend to arrive at quality local optima that are not prone to over-fitting.

Chapter 3

Network Cost Allocation

3.1 Chapter Abstract

This chapter examines major problems with commonly proposed approaches to network cost allocation games, with a focus on cooperative game theory and its limitations. Such games arise in many contexts, such as determining connection fees for growing electricity networks, micro-grids, fiber-optic cable networks, or even in setting up new municipal water and sewage systems. We consider several challenges arising in scenarios with multi-period network growth and strategic users with reservation prices, and propose an incentive-compatible pricing mechanism that addresses these challenges while ensuring network cost-recovery and avoiding certain user decisions that can lead to significant social welfare losses. Our discussion focuses on electricity grids, but is generalizable to a range of network types.

3.2 Overview

This paper examines a cost allocation problem in which a network is constructed, possibly over time, to connect heterogeneous users to a source node from which services or resources are provided. In practice, this problem arises in determining connection fees for new fibre-optic networks, or for customers joining electricity grids or piped water systems that are still growing, as is the case in many low- and middle-

income countries. Many authors have modeled this problem as a cooperative game and have used solution concepts from cooperative game theory (CGT) to propose cost-allocation policies. We believe these solution concepts to be inappropriate in many cases. In Section 3.3 we summarize the essential terminology and solution concepts from CGT, which we shall reference throughout the paper. In Section 3.4, we discuss several limitations of the modeling frameworks and solution techniques most frequently proposed in the literature. In Section 3.5, we present an alternative pricing mechanism that is dominant strategy incentive compatible, applicable in dynamic settings, computationally tractable, and that recovers the network cost in a manner consistent with common notions of fairness.

3.3 Cooperative Game Theory - General Overview

A cooperative game is defined by a set of players N and a characteristic function $c : 2^N \rightarrow R$, $c(\emptyset) = 0$, that establishes a cost for every subset of N. In cost games (our context), smaller allocations are preferred; in value games, larger allocations are preferred. The characteristic function could be, for example, the cost of constructing a minimum spanning tree or minimum Steiner tree to connect users in subset S to the source node. A game is said to be monotone if for $T \subset S, c(T) \leq c(S)$.

A vector $x \in \mathbb{R}^N$ satisfying $X(N) = c(N)$, where $X(N) \equiv \sum_{i \in N} x_i$, is a cost-allocation for the set N, and if $x_i \leq c(\{i\}) \forall i \in N$ then x is called an imputation. Imputations are cost-allocations that charge the agent less than the cost of serving them if the system were designed solely for that agent's use. The set of imputations is always non-empty [Bkakar et al., 2010]. This notion can be further generalized to coalitions of agents: a game's core is the set of cost-allocation vectors for which $X(S) \leq c(S) \forall S \subset N$, which implies that no subset of agents in the grand coalition (N) is being charged more than the cost of serving that subset exclusively / separately $(c(S))$. The difference $e(S, x) \equiv c(S) - X(S)$ is called the excess of coalition S with respect to allocation x. The excess represents how much better off an allocation x leaves coalition S compared to the cost they would bear if they seceded from the

grand coalition and incurred cost $c(S)$. Thus, the core can also be characterized as the set of imputations that ensure that every coalition has non-negative excess. A related concept is the nucleolus, which is the imputation that maximizes the minimum excess. In other words, the nucleolus is the allocation that gives the best possible outcome to the worst-off coalition. The nucleolus always exists, is unique, and if the core is non-empty, it lies in the core.

In cooperative game theory (CGT), it is typically assumed that players can form coalitions that have the ability to secede from the grand coalition, thereby generating a cost of $c(S)$ to be allocated amongst the members of S. When allocations lie outside the core, then at least one coalition has an incentive to secede since their cost allocation is greater than what they could obtain without the other players. For this reason, many CGT solutions methods focus on maintaining the stability of the grand coalition, and hence seek cost-allocations in the core. When the core is empty, then at least one coalition will have an incentive to secede (negative excess) no matter what cost allocation is devised. In these cases, solution techniques tend to focus on finding the nucleolus, which, by maximizing the excess of the worst-off coalition, gives the allocation that creates the tightest possible bound on the size of any coalition's incentive to secede.

An alternative solution criterion is to seek cost-allocations that satisfy certain axioms one might associate with fairness. Three axioms in particular have been given great importance in the field: 1) Symmetry: if two agents i and j are equivalent—that is, they make equivalent marginal contributions to all subsets $S \subset N$ not containing i or j—then their allocations should be the same; 2) Linearity: if we consider two cooperative games (with identical or different grand coalitions) and view them as one game, with a characteristic function equal to the sum of the original games, then the cost allocation rule should assign to each player the sum of their allocations that would have been assigned in each individual game; 3) Zero Null Player: a player that contributes no marginal cost to any coalition should be charged nothing. Every one-stage cooperative game has a unique cost-allocation vector satisfying these conditions. The allocations are called Shapley values, and an agent's Shapley value (SV) is given

by

$$\phi_i = \sum_{S \subseteq N \setminus \{i\}} \frac{|S|!(n-|S|-1)!}{n!} \left(c(S \cup \{i\}) - c(S) \right), \text{where } n = |N| \tag{3.1}$$

.

An agent's SV takes an average of the agent's marginal contribution to the characteristic function evaluated over all subsets of the grand coalition. The weights are chosen to make all subset sizes equally likely. Alternatively, weights could be chosen to give equal importance to all subsets—this weighting that yields the Banzhaf value, which is related to SV but does not preserve the linearity property [Banzhaf III, 1965].

The SV may or may not lie in the core. However, for convex games, the core is non-empty and contains the SV [Shapley, 1971, Granot et al., 2002, Conitzer and Sandholm, 2004]. A cost game is convex iff $\forall S, T \subset N, c(S \cap T) \leq c(S) + c(T) - c(S \cup T)$. For value games, the inequality is reversed. Equivalently, a cost game is convex if $\forall i \in N$ and all $T \subseteq S \subseteq N \setminus \{i\}$, we have $c(S \cup \{i\}) - c(S) \leq c(T \cup \{i\}) - c(T)$. This definition says that for any monotonically increasing sequence of sets (a sequence satisfying $T_0 \subseteq T_1 \subseteq ... \subseteq T_k \setminus \{i\}$) not containing i, agent i's marginal contribution towards the characteristic function is decreasing. For games with infinite players (or rather, infinitesimal players), the Aumann-Shapley value is a generalization of the SV and shares the same defining properties as SV.

Implementing cooperative game solution methods is intractable in some cases. A naive computation of SVs requires summing a number of terms that grows super-exponentially in the size of the grand coalition, though problems in practice may have special structures enabling polynomial-time computation. Unless games are convex, checking whether an allocation belongs to the core is NP-complete [Conitzer and Sandholm, 2004, Deng and Papadimitriou, 1994].

3.4 Limitations of Common Network Cost Allocation Models and Solution Techniques

The network cost-allocation literature pertaining to electricity grids is split primarily between papers that employ CGT solution concepts, and those that do not consider strategic behavior. We conducted an extensive, albeit not exhaustive, assessment of articles on network cost allocation with relevance to electricity grids or similar infrastructure networks [Ruiz and Contreras, 2007, Junqueira et al., 2007, Molina et al., 2013, Rotoras et al., 2002, Ghayeni and Ghazi, 2011, Gil et al., 2006, Granot et al., 2002, 1996], and found several areas for potential improvement.

3.4.1 Multi-Stage Network Growth

Many large infrastructure networks, including electricity grids, grow over time. Users may be able to choose to join the network now, or at some future time period, and the cost-allocation may depend on the time period, or at least on the set of users who have already joined. Time dynamics add realism and complexity to the problem, but were not considered in the papers reviewed.

3.4.2 Computational Tractability

Computational tractability is an important consideration in the design of any solution technique to the network cost allocation problem, as the number of network nodes in an electricity grids can be quite large. With respect to SVs and checking core membership, computational tractability is not guaranteed in general, but can be achieved in problems with special structures. Most papers that advocated the use of SVs did not address computational tractability. However, two papers reviewed did prove that the nucleolus and SVs can be computed efficiently in the cases of single-period tree networks with homogeneous users [Granot et al., 1996] or heterogeneous users Granot et al. [2002]. The proof technique, however, required that arc costs be determined by the maximum type of all upstream agents using the arc in their path

to the source. This assumption means that arc costs do not need to increase as the number of users increases, or as the total demand flowing over the arc increases, as long as the maximum user type does not increase. Indeed, an arc could connect an infinite number of users for the same cost as an arc that serves one user, as long as the maximum user type is the same in both sub-trees served by those arcs. These arc cost assumptions may be unrealistic in many settings, but it is unclear whether SV computations and checking core membership can be performed efficiently under more complex cost assumptions.

3.4.3 Why Cost Games Are the Wrong Framework

In cooperative cost games, it is often assumed that all players must belong to at least some coalition (even if singleton), and cannot choose to realize a cost of zero by not participating in the game. This may be true in some real contexts if participation is enforced, but it is not true in many real-world scenarios in which agents are free to decline services. Optional participation characterizes many service networks, such as electricity grids in countries where many households are not yet connected, and may decline the opportunity to connect if the connection price is too high.

That agents may choose not to participate implies that they have valuations for the network's services, and may decline to participate if their cost allocation exceeds their valuation. The cost-allocation game, therefore, is actually a cooperative value game, with the following characteristic function for social welfare: $u(S) = v(S) - c(S)$, where $v(S)$ is the sum of agents' valuations for obtaining services from the network serving subset S. The core, SVs, and nucleolus of this value game may be very different from the cost-game. Since valuations are private information in most real-world contexts and would need to be revealed for the core, SV, and nucleolus to be computed, it may not be possible to implement these solution techniques at all, unless they happen to coincide with the payments that would lead to a truth-telling mechanism.

Ultimately, if participation is optional, then SV and other CGT solutions to cooperative cost games may actually be unstable and could yield poor welfare outcomes. For example, consider a network consisting of k users equally spaced along a line,

with identical demand needs and private valuations of $1 + \delta$. Suppose that the cost of each arc is 1. The cost allocation $x = \mathbf{e}$, the unit vector, is stable (no agent or coalition can do better by seceding), and achieves the optimal social welfare ($k\delta$). If we were to use the SV of the cost game, ignoring users' private valuations (which would be unknown to the cost allocator in practice), then the first node's SV would be $\frac{1}{k!}$, and the other first few nodes' values would also be small, since these nodes contribute to the marginal cost only in a very small number of coalitions, whereas the last few nodes will have very high values, since they contribute to the marginal cost in many coalitions. For a small enough δ, the last node will chose not to participate. The second last node will become the last node, and again, for a small enough δ, this node will choose not to participate. By induction, $\exists\delta$ such that only the first node will participate, causing us to lose $\frac{k-1}{k}\%$ of the possible social welfare. Finally, if we suppose that the first node's valuation is in fact 1, not $1 + \delta$, then the first node may participate but obtain no consumer surplus. This example shows that the SV can cause all social welfare to be lost if a cost-game framework is used when agents' participation is optional.

3.4.4 Excessive Concern for Stability

Cooperative game theory's central solution concepts have been developed and applied largely out of a concern for solution stability—that is, minimizing or eliminating incentives for coalitions to secede. However, the focus on stability is misplaced in many application areas in which important constraints on coalition secession exist. There are many contexts in which laws, licenses, regulations, or the rights of the service provider, make it possible to prevent certain coalitions from seceding and still retain service. Specifically, while coalitions may be able to coordinate their participation, or non-participation, and while a characteristic function may indeed specify a cost value for each possible subset of agents who decide to participate, coalitions may be unsuccessful in demanding and receiving service separately from other participating players. For example, supposing a coalition of electricity grid consumers realized its imputation of transmission line charges $X(S)$ had negative

excess and requested to be considered and billed separately from other consumers—that is, for the amount that they would owe if the other consumers did not exist—the transmission company might be within its rights to simply decline the request. In this context, where network services are provided by a company or government responding to agents' decisions to participate in the network or service program, agents can still secede by not participating, but possibly few or no coalitions can secede in the traditional sense of demanding and obtaining an imputation bounded by $c(S)$.

A more striking example in which secession is altogether impossible appears in several papers where a final coalition had already been formed and a cost value realized. Nonetheless, several papers used CGT to devise allocations after the fact [Molina et al., 2013, Junqueira et al., 2007]. However, when the possibility of secession is indeed limited to non-participation, CGT solution concepts are no longer appropriate. Rather than focus on marginal contributions to coalitions that cannot realistically form, or worry about incentivizing secessions that cannot happen in practice, one ought to seek imputations that can lead to good social welfare outcomes.

3.4.5 Questionable Notions of Fairness

In addition to stability, the notion of fairness has played an important role in the development of solutions to cost allocation games. The particular meaning of fairness, however, can be rather narrow, if not problematic. Fairness—regularly equated with SVs in the literature—is often defined in terms of whether an allocation reflects an agent's marginal contribution to various possible coalitions. This particular definition of fairness is mathematically convenient to work with; however, this mathematical convenience seems to have cast this version of fairness as a normative gold-standard. In many real world cooperative cost games, the reasons that certain agents have advantages over others may be arbitrary, due to chance, or to socioeconomic disadvantage. Concern for fairness, from an equity perspective, could entail ignoring or reducing the benefits that might accrue from these (potentially arbitrary) advantages, whereas a Shapley interpretation of fairness deliberately ensures that these advantages are reflected in the allocation. These two interpretations of fairness

are diametrically opposed.

To illustrate, suppose an electricity generator must choose a location for a coal plant at one of two extremes of a line segment, with a set of identical users being equally spaced along the line segment. The choice of either location is arbitrary, yet will confer far lower SV transmission cost allocations to nearby nodes that, ex ante, were identical to the corresponding nodes on the other side of the mid-point. Despite using SV, the fairness of such an allocation is certainly debatable.

3.5 Alternate Approach: "Fair Share" with Rebates

As before, we will consider how to allocate the cost of a network amongst users wishing to connect to a source node. We will make several assumptions that differ from the CGT literature. Specifically, we view the cost allocation problem not as a single-period cost-game but as a multi-period value game with imperfect information. Assume user u obtains a constant discountable payoff x_u in each time period he belongs to the network. A user's payoff is private information, and his reservation price at time t is the net present value of future payoffs, which can be modeled as the value of an annuity or perpetuity, depending on whether a finite time horizon is used. A user may join the network at one of T discrete time periods, or choose not to join by the final period. Joining is an absorbing state. The set of possible histories has size $|T+1|^N$, and the extensive form game has a vector-valued characteristic function that maps between the set of possible histories and the set of costs at each point in time: $c : |T+1|^N \to \mathbb{R}^T$. Coalitions do not have the ability to demand service or provide themselves with service at the cost $c(H_S)$, where H_S is the history of coalition S. The cost function would likely reflect a minimum cost spanning tree or Steiner tree in the first period, with minimum cost extensions to nodes that decide to join in subsequent stages. Arc costs should be increasing in actual or expected future capacity requirements. We assume a common discount rate across all agents.

The challenge is to establish a cost-allocation for each period that stands a good chance of yielding high social welfare even when valuations are unknown. For historical, political, and cultural reasons, setting connection fees in the electricity sector on the basis of private valuations is considered unacceptable [Ortega et al., 2008]. A direct mechanism that asks for user's private valuations is therefore not considered.

A politically acceptable direct mechanism could simply be to ask whether a user's reservation price is above a threshold; this could take the form of a fee proposal. We develop a set of conditions and a specific fee policy for a dominant strategy incentive compatible (DSIC) mechanism of this form. The mechanism essentially renders users myopic: they will accept connection fees below their reservation price and reject offers above, for any discount rate, and no matter what the strategies of other players. When users do not behave myopically, social welfare may fall in two ways. First, users lose utility if they delay joining when it could have been profitable; second, users who delay for strategic reasons will contribute to the network growing more incrementally rather than in large chunks early on. The more incremental the network grows, the further the final network design will be (in typical cases - not necessarily in all cases) from the final minimum spanning tree or minimum Steiner tree, and the higher the overall network cost. These potential social welfare losses are illustrated in examples presented in Section 3.6.

The quality of the proposed fee policy with respect to social welfare vs. other fee schemes may depend on the distribution of private valuations and various network structure idiosyncrasies. This aspect will be explored more thoroughly via simulation and bound derivation in our future work. In the remaining subsections of this paper we will simply present the DSIC mechanism on a tree network, and several extensions for non-tree networks.

3.6 Examples of Cost Allocation Rules Failing

3.6.1 Marginal Cost Pricing

One option for assigning infrastructure costs to users is to charge them the marginal cost of them joining the network. This policy ensures that the full cost of the network is recovered, and that users do not pay for components of the infrastructure that they do not use. However, later joiners may use network components that have already been paid for by others, and so they are effectively subsidized by the early joiners. Situations may also arise where it is not possible for a user to join the network unless others do so simultaneously, and split the cost of shared lines.

Example. Consider the following potential network with two users, each with reservation price of 6, and with line costs of 4, 1 and 1. The users must decide whether to join the network now, or to join one period later. If both users join now, they will each pay 3 to join the network, and enjoy a consumer surplus of $6 - 3 = 3$. If one joins now and the other waits, the first joiner pays 5 and obtains consumer surplus of 1, while the second joiner pays 1 and enjoys a consumer surplus of $\frac{5}{1+\alpha}$. If both delay, they obtain a consumer surplus of $\frac{3}{1+\alpha}$. Unfortunately, delaying is a sub-optimal outcome, but also a Nash Equilibrium.

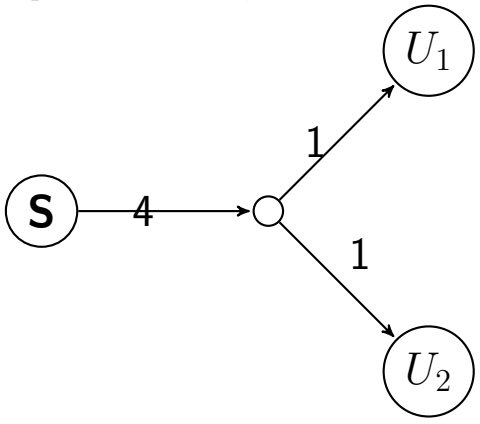

3.6.2 Average Cost Pricing (ACP)

Average cost pricing (ACP) is a common regulatory requirement imposed on natural monopolies to protect consumers. Under ACP, a user pays the average cost of the

system at the time of his or her joining. At a given point in time, new users who impose above average marginal costs of connecting to the system are subsidized by new users who impose below average costs. This reduces fairness and may also price some potential users out of the network who would have otherwise joined if they did not have to subsidize the expensive users. ACP can create incentives to delay joining the network in anticipation of being able to split shared network infrastructure amongst more people.

Example. Consider a potential network similar to Example 1, but with one user's private edge costing 11 instead of 1. Suppose that the reservation prices are 20 and 6 for user 1 and user 2 respectively. Suppose the strategy space consists of joining now, joining in the next period, or not joining at all. The payoff matrix for the two users is shown in Table 3.1.

Table 3.1: Payoff Matrix for Users in Example 2 under Average Cost Pricing

	User 2		
U1 ↓	Now	Later	Never
Now	12, -2	5, $\frac{-2}{1+\alpha}$	**5, 0**
Later	$\frac{12}{1+\alpha}$, 1	$\frac{12}{1+\alpha}$, $\frac{-2}{1+\alpha}$	$\frac{5}{1+\alpha}$, 0
Never	0, 1	0, $\frac{1}{1+\alpha}$	0,0

Observe that the Nash Equilibrium is for user 1 to join now and for user 2 not to join at all, even though user 2 is relatively much less expensive to connect to the system, and his private valuation is sufficient to cover his connection path to the source. This outcome is sub-optimal and unfair.

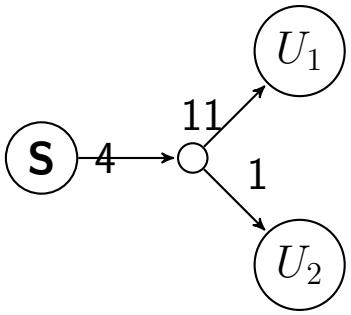

3.6.3 Shapley Values

We can also construct an example for which Shapley Values give the worst possible outcome and average cost pricing (ACP) gives the best. Consider a linear network with identical cost edges c_e. Using Shapley Values, allocation $\phi_i^{SV} = c_e \sum_{j<i} \frac{1}{n-j}$ and under ACP, $\phi_i^{ACP} = c_e$.

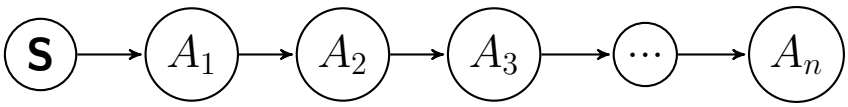

Suppose we let valuation $v_1 = c_e$ and for $i > 1$ let valuations $v_i = c_e H(i) - \epsilon$, where $H(i)$ is the harmonic sum: $\sum_{j=1}^{i} \frac{1}{j}$, and $\epsilon > 0$ is a small number. Observe $\phi_n^{SV} > v_n$ so the n^{th} (final) agent will withdraw from the game. The ϕ^{SV} are recomputed, and now agent $n-1$ faces the same problem, and withdraws from the game. By induction, all agents except the first withdraw, but even it obtains no utility from joining, so the SV causes 100% of the social welfare to be lost. This was not inevitable. The maximum social welfare happens to be V(N) - C(N), and can be achieved by the individually rational allocation ϕ^{ACP} since $\phi_i^{ACP} \geq c_e \; \forall i$.

3.6.4 Proposed Mechanism

General Summary: the pricing policy determines a value ϕ_{ut} for user u at time period t based on the history of who else has already joined the network by time t. ϕ_{ut} is non-increasing in t. The user will pay ϕ_{ut} if he joins at time t, and in each subsequent period will receive a rebate (transfer) $r_{ut'} \equiv \phi_{u,t'-1} - \phi_{u,t'}, r_{ut'} \geq 0 \forall t'$. For the mechanism to be self-funding, we require that the sum of all charges, minus rebates, equals the network cost at each time period. Alternatively, permitting some risk and access to financing, the network service provider may form beliefs about the user types and future joining behavior, and may instead require that the sum of all charges, minus rebates, equals the network cost, on expectation, by the final period.

Necessary and Sufficient Conditions for DSIC of a pricing and rebate proposal: if the pricing and rebate scheme causes the following relation to hold

for any discount rate and for all possible future opponent actions, then the pricing mechanism is dominant strategy incentive compatible.

$$\text{If } t' = inf\{t : \phi_t \leq R_p\} \Rightarrow t' \in Argmax_t NPV(t), R_p \text{ is reservation price} \quad (3.2)$$

In other words, if the mechanism ensures that the first time period in which the fee offered falls below the reservation price belongs to the utility maximizing set of possible time periods for joining the network, then the mechanism is DSIC.

Several sufficient conditions for a DSIC mechanism now become immediately obvious:

$$\text{The sequence } \{\phi_t\} \text{ is non-decreasing} \quad (3.3)$$

$$\text{The sequence } \{\phi_t\} \geq \frac{\phi_{t-1}}{(1+\alpha)} \forall t \text{ (decreases by less than the discount factor)}, \quad (3.4)$$

for when these conditions hold, users who obtain positive utility by joining at some t' will find their utility monotonically decreasing for $t > t'$.

3.6.5 Tree Networks, Users Pay Fraction of Demand Flowing over Arcs

Next we consider a specific formula for ϕ_{ut} and prove that it is DSIC on tree networks. This particular cost-assignment method assigns the entire cost of arcs to current users, even if the network designer has built extra-capacity unused by these users in anticipation of future demand.

By assumption the network has the form of a tree, so the demand d_u must be carried across all edges in the path connecting the source to u. Let P_u be the set of edges connecting the source to u. Let $U(e) = \{u : e \in P_u\}$ be the set of users whose paths contain edge e. Let ρ_{ue} denote user u's share of edge e, which is calculated as u's fraction of the total demand passing through edge e.

$$\rho_{ue} = \frac{d_u}{\sum_{u' \in U(e)} d_{u'}} \tag{3.5}$$

The user's fair share of the cost of e is therefore $\rho_{ue} c_e$. Evidently, the cost of the edge is recovered by charging this price to all users: $\sum_{u \in U(e)} \rho_{ue} = 1 \Rightarrow \sum_{u \in U(e)} \rho_{ue} c_e = c_e$. Since the cost of each edge is recovered in this manner, the total network cost is also recovered (for a single time period).

If we let ϕ_u denote the total charge for user u, then $\phi_u = \sum_e \rho_{ue} c_e$, noting that $\rho_{ue} = 0$ for edges that are not in the path from the source to u. It also holds that $\sum_{u \in U} \phi_u =$ Total System Cost.

Now let's consider this fee structure in a multi-time period setting. Observe that ρ_{ue} can change over time, because new users can enter the network (but old users cannot leave, or at least cannot be refunded their connection fee). Therefore ρ_{uet} is monotonically decreasing in time as new users join and may dilute u's share of usage of edge e. Therefore the time-dependent fair usage connection fee ϕ_{ut} is also monotonically decreasing with time. The rebates provided are as described above.

Proposition 1. For a user that joins at time t_0, the system operator can offer a rebate at each time $t > t_0$ that satisfies $r_{ut} = \phi_{u,t} - \phi_{u,t-1}$. These rebates can be financed from the system operator's profit under the fair-usage pricing scheme.

Proof Under this rebate scheme, the system operator's fee minus rebates to any user u at any time t sum to ϕ_{ut}. It is easy to verify that $\phi_{ut_0} - \sum_{t'=t_0+1}^{t} r_{ut'} = \phi_{ut}$. Since at time t the system operator has ϕ_{ut} from each user, the cost of the system is covered, since $\sum_u \phi_{ut} = \sum_u \sum_e \rho_{uet} c_e = \sum_e \sum_u \rho_{uet} c_e = \sum_e c_e \sum_u \rho_{uet} = \sum_e c_e \times 1 =$ Total Cost , $\forall t$.

Proposition 2. The pricing mechanism proposed (fair usage charge and rebates) is DSIC.

Let PVB(t) denote the present value of the benefits (utility) of joining the network at time t. PVB(0) $= \frac{x_u}{\alpha}$, and $PVB(t) = \frac{x_u}{\alpha(1+\alpha)^t}$. As before, $PVC(t) = \frac{\phi_{ut}}{(1+\alpha)^t} - \sum_{\tau=t+1}^{\infty} \frac{r_{u\tau}}{(1+\alpha)^\tau}$. Denote the net present value of joining at time t as $NPV(t) = PVB(t) - PVC(t)$.

A rational user decides when to join the network by solving the following problem:

$$\max_t \left(PVB(t) - PVC(t)\right)$$
$$= \max_t \left(\frac{x_u}{\alpha(1+\alpha)^t} - \left[\frac{\phi_{ut}}{(1+\alpha)^t} - \sum_{\tau=t+1}^{\infty} \frac{r_{u\tau}}{(1+\alpha)^\tau} \right] \right)$$

If $NPV(t)$ is negative for all finite values of t then the user's best option is to never join ($NPV(\infty) = 0$).

An equivalent formulation of the problem is to compare the net present value (NPV) of joining immediately to joining in some future time period, and choosing t to maximize the difference. When the solution is negative, the user cannot improve his utility by waiting. We shall suppose that the initial price offered, ϕ_{u0} is below the reservation price $\frac{x_u}{\alpha}$, and will see that the user will always accept ϕ_{u0}.

$$\Delta NPV = NPV(t) - NPV(0) = [(PVB(t) - PVB(0)] + [PVC(0) - PVC(t)] \tag{3.6}$$

$$\max_t \left(\frac{x_u}{\alpha} \left(\frac{1}{(1+\alpha)^t} - 1 \right) + \phi_{u0} - \sum_{\tau=1}^{\infty} \frac{r_{u\tau}}{(1+\alpha)^\tau} - \frac{\phi_{ut}}{(1+\alpha)^t} + \sum_{\tau=t+1}^{\infty} \frac{r_{u\tau}}{(1+\alpha)^\tau} \right) \tag{3.7}$$

Using $\phi_{ut} = \phi_{u0} - \sum_{\tau=1}^{t} r_{u\tau}$, we make this substitution, and also cancel out the component of the two summands shared in common ($\tau \geq t+1$), to obtain with some additional simplifications:

$$\max_t \left(\frac{x_u}{\alpha} \left(\frac{1}{(1+\alpha)^t} - 1 \right) + \left(\phi_{u0} - \frac{\phi_{u0}}{(1+\alpha)^t} + \frac{1}{(1+\alpha)^t} \sum_{\tau=1}^{t} r_{u\tau} - \sum_{\tau=1}^{t} \frac{r_{u\tau}}{(1+\alpha)^\tau} \right) \right) \tag{3.8}$$

And with further simplifications:

$$\max_t \left(\left(\phi_{uo} - \frac{x_u}{\alpha} \right) \left(1 - \frac{1}{(1+\alpha)^t} \right) + \left(\frac{1}{(1+\alpha)^t} \sum_{\tau=1}^{t} r_{u\tau} - \sum_{\tau=1}^{t} \frac{r_{u\tau}}{(1+\alpha)^\tau} \right) \right) \quad (3.9)$$

This equation reveals important properties about the household's investment decision under the proposed fee and rebate policy. First, it is important to observe that the difference in the second major set of parentheses, the one that compares different present values of the rebates, is necessarily non-positive, and equals zero if $t = 1$.

Therefore the user will not delay joining if the first term in the expression is also negative, making the whole expression negative for all t. This is equivalent to requiring that, $\phi_{u0} < \frac{x_u}{\alpha}$. Since $\frac{x_u}{\alpha}$ is the present value of the utility stream from joining the grid, it should also be the user's reservation price. So if ϕ_{u0} is below the reservation price (which it is by assumption), then the maximization expression is necessarily negative for any positive value of t, and the user's best decision is to accept the price ϕ_{u0}. Therefore, the rebate mechanism causes rational users to accept all prices below their reservation prices.

In a two-stage game, with only one time-period of possible delay, right-hand terms cancel and users will only base their decision on whether the price offered is less than or equal to their reservation price. However, if $t > 2$, then the right-hand term can become slightly negative, which begs the question of whether users might accept ϕ_{u0} higher than their reservation price, in anticipation of the rebates. It is easy to show that this is not the case. If $\phi_{u0} > R_p$ then the first term is positive, and we can choose $t = 1$ to make the rebate portion equal 0, giving a utility improvement from delaying by one time period. This holds for any feasible set of rebates and any discount rate. Therefore, under the rebate scheme, users will reject prices above their reservation price, making user behavior identical to that of myopic users.

These results hold for any discount rate, and for any future actions of other players, since the fee policy ensures non-negative rebates, and so the policy is DSIC. The proof technique did not actually require a particular structure for ϕ_{ut} and r_{ut} aside from the

conditions outlined under General Summary. Therefore, any fee structure meeting those conditions is DSIC.

3.6.6 Tree Networks, Users Pay Fraction of Arc Capacity

So far we have considered networks with uncapacitated edges. In reality, many networks edges have capacity constraints. Suppose that each edge now has capacity μ_e. A reasonable cost-causality function will charge users only for the fraction of the capacity of each edge that they use, with the cost of excess capacity (CEC) being socialized in some manner. The firm can choose to assign some fraction β of the CEC to existing users upfront, or the firm can finance it through bank loans, and recover the cost through charges to future users who will use the excess capacity. Under this model, the user's charge, not counting CEC, is constant. However, a user's share of the CEC, although decreasing in the number of users (*certeris paribus*), could also increase if the firm decided to make a lot of long term network investments.

In this setting, we can develop a modified fee and rebate structure based on cost causality, and that is DSIC under some conditions. Let P_u represent the set of edges in the path from the source to user u, let U_t denote the users who have joined by time t, and let $TD_t = \sum_{u' \in U_t} d_{u'}$ be the total demand at time t. Let the fee charged to user u if he joins at time t be: $\phi'_{ut} = \sum_{e \in P_u} c_e \frac{d_u}{\mu_e} + \beta_t CEC_t \frac{d_u}{TD_t}$. Define the rebates as before. Note that the first summand is independent of time; the only time-dependent part is the fraction of the CEC assigned to the user. This quantity could increase or could decrease with time. In order to ensure non-negative rebates to make the scheme DSIC (and similar in most respects to the uncapacitated scheme), we must impose the following condition on the firm's choice of β and CEC_t : $\beta_t CEC_t \frac{1}{TD_t} \leq \beta_{t-1} CEC_{t-1} \frac{1}{TD_{t-1}} \forall t$. Without this condition, the user may be charged additional fees of an unknown quantity in the future and we lose the DSIC property.

3.7 Cost Assignment for Networks with Cycles

At the level of household distribution from electricity substations, grids typically take the form of trees. At the level of high-voltage networks connecting substations, there tends to be more cycles in the network so that the grid will continue to work even if a particular line fails. At very high voltage levels (the transmission grid), the network is highly "meshed." Transmission networks are supposed to have N-1 survivability, which means that no service disruptions should occur if any individual transmission line fails.

Determining a user's cost-causality is more straightforward on trees than on networks with cycles, because trees imply a unique path from the power source to the user. We propose two ways to determine cost-causality on networks with cycles.

Method 1: Assign the cost of each of the insurance edges that provide N-1 survivability to those who benefit from them, and socialize the cost of all other edges providing N-k survivability.

Assumptions: arcs are uncapacitated and the network has full N-1 survivability (at least two disjoint paths from the source to each user).

1. Start with the original network, and find the lowest cost path from S to u; add the edges to the set $MAIN_u$
2. Remove the edges in $MAIN_u$ from the network and find the lowest cost path from S to u; add edges to the set INS_u (INS for insurance)
3. Repeat the first two steps for all users
4. Obtain the set of insurance edges allowing N-1 survivability: $\bigcup_{u \in U} INS_u - \bigcup_{u \in U} MAIN_u$

Method 2: Socialize the cost of all N-k edges.

Case 1: Assume edges are uncapacitated. Find the "default" tree. A sensible choice is the minimum cost spanning tree, which is easy to find. The difference between the full set of edges and the tree is the set of "insurance" edges, whose costs can be socialized.

Case 2: Assume edges are capacitated. Let the default network be the lowest cost subset of the pre-existing network that has feasible flows from the source to all users. One can formulate the problem of finding the default network as a mixed integer program (MIP). Denote whether arc e is used with $z_e \in \{0, 1\}$, let $x_e \in R^+$ be the flow over arc e, μ_e be the capacity of arc e, and c_e be the arc's positive construction cost. The MIP to find the default network can be formulated as follows:

$\underset{\mathbf{x},\mathbf{z}}{min}\ c'z$, subject to: $x_e \leq \mu_e z_e$

$Ax = d$ (feasible flows)

We can subtract the default network's edges from the full set to obtain the set of "insurance" edges, the cost of which can be socialized.

3.8 Conclusion

The mechanism by which network costs are allocated to potential users can affect the speed of network growth and the level of social welfare achieved. Most cost allocation strategies advocated in the literature for infrastructure networks, such as electricity transmission or distribution lines, are deeply problematic. Cooperative cost games do not properly reflect key features of such infrastructure networks: optional participation, limited options for secession with service, reservation prices—which imply that the game is actually a value-game with imperfect information, and multi-stage network growth. These problems could mean that solution techniques from CGT are entirely inappropriate, and will yield cost-allocations that are inefficient, inequitable, and vulnerable to strategic gamesmanship by users.

It is possible to design network cost-allocation mechanisms that do not rely on CGT solution concepts using a more realistic modeling framework—one in which users have private valuations, where the network can grow over time, is not restricted to being a tree, and can have capacitated or un-capacitated arcs with costs that are proportional to network size. To avoid social welfare losses associated with strategic user behavior, we developed a set of DSIC cost-allocation mechanisms for a variety of network types. We are optimistic that this family of mechanisms will result in

superior social welfare outcomes, and will test these mechanisms via simulation in our future work.

Chapter 4

Network Cost Allocation, Part II

Introduction

This chapter provides a mathematical analysis of the impacts of certain pricing regulations on single-period network construction. Simple models are developed for which optimal policies can be derived and compared without relying on simulations, which may fail to generalize or reveal underlying mathematical features of interest. In the simple models considered, customers are ordered along a one dimension lattice from the source node. At the end of the chapter, a method is outlined for computing a reasonable lower-bound for the expected value of 2D networks by decomposing the 2D network into 1D segments.

The regulations examined concern rules that aim to protect private reference price information. If a monopolist learns the private reference prices of potential users, it could offer them their reference price to join the network and convert all consumer surplus into profit. Consequently, consumers should not want to reveal information to the monopolist about their reference price. To protect consumers from abuse, a regulator could require that the monopolist simultaneously offer the same joining price to all users to whom offers are made. Moreover, the regulator could require that the offer-acceptance process be binding on both parties so that no party can renege after information about private reference prices is revealed.

By contrast, if the monopolist can be trusted not to use revealed user information

against the users' interests in the future, then the regulator could permit the monopolist to build the network to include the most profitable subset of price-accepting users after the users' responses to the offer are known.

4.1 Summary of Chapter Results

It can be profitable for the monopolist to offer to build the network even when the expected marginal profit per potential user ("drift") is negative and the monopolist is required to serve all potential users who accept the price offer. This counter-intuitive scenario is especially true for smaller networks or when the reward per price-taking user is very high, and a simple expression is derived that characterizes the sets of parameters that make this possible for a 1D lattice.

For scenarios where the monopolist can choose to make binding offers to a subset of users in an area instead of to all of them, an expression for the profit-maximizing number of offers is derived. This profit-maximizing number of offers is shown to be substantially smaller than the break-even number of offers, and may result in massive social welfare losses vs a larger network that merely breaks-even from the monopolist's perspective.

When the monopolist is permitted to build the network up to any subset of users of its choosing after the users signal their willingness to pay the offer price, expected profit will be higher. In particular, when drift is negative, expected profit for the best subset of price-takers on a 1D lattice can be many multiples higher than when the monopolist must connect to all price-takers, and can be approximated by the expected maximum of a random walk or the expected maximum of a Wiener process.

In all sections that follow, reference prices are private but come from some i.i.d. distribution known to the monopolist, $F_r(\cdot)$, $f_r(\cdot)$. Unless otherwise indicated, the profit-maximizing network monopolist proposes a connection fee, the potential user accepts or rejects the fee, and the monopolist connects those who accept and receives the proposed fee. Unlike Chapter 3, the monopolist can incur losses and is limited to offering a fixed price rather than a price-proposal that depends on the decisions of

other potential users.

4.2 Two User Types, Constant Expected Inter-User Extension Cost

Suppose the cost to expand the network between adjacent potential users on the 1D lattice is always 1, and that there are only low and high user types defined by their reference prices: $r_\ell = 0$, and $r_h = 2$. Clearly, if the monopolist offers to build a network at all, it will propose $\pi = r_h = 2$ as the joining price. The profit or loss for building the network up to user n can be modelled as a non-reflected random walk with n steps: profit changes by +1 or -1 depending on the next user's type. Consequently, if the high-type probability, p_h is greater than 0.5 the monopolist will make price offers and the network will be built up to the last high-type. If $p_h = 0.5$ then expected profit is zero and a risk-neutral monopolist will be indifferent. If $p_h < 0.5$ then the monopolist will not offer to build the network when it must make offers to a large number of potential users, but may offer to build for a small number because it only has to build up to the last user who accepts, and the reward from that last user is therefore guaranteed, possibly offsetting the negative expected value of having to build the network up to that point. For example, if there are only two potential users, the expected profit from building the network up to the last user who accepts the price is strictly positive: $1 + p_h$.

To generalize the analysis, let $r_h = \theta$, let V_n be the monopolist's profit if it builds the network up to the n^{th} household (regardless of whether n is high-type or low-type), and let $\bar{V}_n$ be the monopolist's profit if it makes offers to build the network at price $\pi = \theta$ to the first n potential users and then builds up to the last price-taker. Define $S_n \equiv \sum_{i=1}^{n} r_i$. Further, let marginal cost of extending the network be normally distributed, $\mathcal{N}(c, \sigma_c)$, and independent of user type. Then,

$$V_n = S_n - \sum_{i=1}^{n} c_i$$

$$\mathbb{E}[V_n] = \mathbb{E}[S_n] - cn = np_h\theta - cn = n(p_h\theta - c)$$

$$Var[V_n] = Var[S_n - \sum_{i=1}^{n} c_i] = Var[\sum_{i=1}^{n} r_i] + n\sigma_c^2 = n(\sigma_r^2 + \sigma_c^2)$$

$$\Rightarrow Var[V_n] = n(\theta^2(p_h)(1-p_h) + \sigma_c^2)$$

If $p_h\theta \geq c$ then drift is positive and the monopolist makes offers.

It can sometimes be profitable to make offers even when $p_h\theta < c$ because V_n is a lower bound on the actual profit realized in a system containing n potential users, as the monopolist does not have to connect to the final sequence of users if they are all the low-type. The sets of θ, c, p_h, and n that will still lead to monopolist participation are characterized in the following paragraphs.

Let n' denote the index of the last high-type user ($n' \leq n$). Recall that $\bar{V}_n$ denotes the actual profit realized when an offer to n users is made and the network is built up to n', the last high-type. A risk-neutral monopolist's participation constraint is $\mathbb{E}[\bar{V}_n] \geq 0$. We can expand the expected profit expression as follows: $\mathbb{E}[\bar{V}_n] = \sum_{i=1}^{n}[\mathbb{E}[V_{i-1}] + (\theta - c)p_h(1-p_h)^{n-i}]$, letting $V_0 = 0$. Each i represents a possible value of n'. The $(\theta - c)$ term represents the expected amount that we will gain $(\theta - c)$ from user n', the $\mathbb{E}[V_{i-1}]$ represents the expected value of serving the $i-1$ customers behind n', and $p_h(1-p_h)^{n-i}$ is the probability that $n' = i$. Observing that $\sum_{i=1}^{k}(1-p_h)^{k-i}$ is a partial geometric sum, we can use the substitution $\sum_{i=1}^{k}(1-p_h)^{k-i} = \frac{1}{p_h} - \frac{(1-p_h)^k}{p_h}$

and simplify the expression for $\mathbb{E}[\bar{V}_n]$ as follows:

$$\mathbb{E}[\bar{V}_n] = (\theta - c)[1 - (1 - p_h)^n] + p_h \sum_{i=1}^{n} \mathbb{E}[V_{i-1}](1 - p_h)^{n-i}$$

$$= (\theta - c)[1 - (1 - p_h)^n] + p_h \sum_{i=1}^{n} (i-1)(p_h\theta - c)(1 - p_h)^{n-i}$$

$$= (\theta - c)[1 - (1 - p_h)^n] + p_h(p_h\theta - c) \sum_{i=1}^{n} (i-1)(1 - p_h)^{n-i}.$$

We can further simplify the summand in the expression:

$$\sum_{i=1}^{n} (i-1)(1 - p_h)^{n-i} = \sum_{j=0}^{n-2} \sum_{k=0}^{j} (1 - p_h)^k$$

$$= \sum_{j=0}^{n-2} \left[\frac{1}{p_h} - \frac{(1 - p_h)^{j+1}}{p_h} \right] = \frac{n-1}{p_h} - \frac{1 - p_h}{p_h} \sum_{j=0}^{n-2} (1 - p_h)^j$$

$$= \frac{(1 - p_h)^n + np_h - 1}{p_h^2}.$$

Using this result, we can express $\mathbb{E}[\bar{V}_n]$ as follows:

$$\mathbb{E}[\bar{V}_n] = (\theta - c)[1 - (1 - p_h)^n] + p_h(p_h\theta - c) \left[\frac{(1 - p_h)^n + np_h - 1}{p_h^2} \right]$$

$$= (\theta - c)[1 - (1 - p_h)^n] + (p_h\theta - c) \left[\frac{(1 - p_h)^n + np_h - 1}{p_h} \right]$$

$$= \frac{p_h(\theta - c) - p_h(\theta - c)(1 - p_h)^n + (p_h\theta - c)(1 - p_h)^n + (p_h\theta - c)np_h - (p_h\theta - c)}{p_h}$$

$$= \frac{np_h(p_h\theta - c) - (1 - p_h)^n c(1 - p_h) + c - cp_h}{p_h}$$

$$= \frac{np_h(p_h\theta - c) - c(1 - p_h)^{n+1} + c - cp_h}{p_h}$$

$$= n(p_h\theta - c) + \frac{-c(1 - p_h)^{n+1} + c - cp_h}{p_h}$$

$$= n(p_h\theta - c) + \frac{c(p_l - p_l^{n+1})}{p_h} \tag{4.1}$$

The first term in the expression is the expected value of building the network up

to the n-th potential user and collecting θ from the high-types. The second term represents a cost adjustment for actually only having to build the network up to the last high-type. Using this expression for expected profit, the condition $\mathbb{E}[\bar{V}_n] \geq 0$ can be expressed as follows:

$$\mathbb{E}[\bar{V}_n] \geq 0 \iff p_h\theta - c \geq \frac{c[p_\ell^{n+1} - p_\ell]}{np_h} \tag{4.2}$$

The LHS is the drift. Observe that since $p_\ell > p_\ell^{n+1}$ the RHS is always negative and does not contain θ, whereas the LHS is linearly increasing in θ. Therefore, for any values of $p_h > 0$, c, and $n > 0$ there is an interval for which it is profitable to offer to build the network up to the last price-taker despite the presence of negative drift ($p_h\theta - 1 < 0$). The interval is given by $\theta \in (\frac{c[p_\ell^{n+1} - p_\ell]}{np_h^2} + \frac{c}{p_h}, \frac{c}{p_h})$.

The expression can be conditioned on three parameters instead of four simply by scaling the monetary units such that $\theta \equiv 1$, with c expressing some fraction of the joining price, or $c \equiv 1$, with θ expressed as a multiple of the unit extension cost. Figure 4-1 shows how large θ needs to be as a multiple of the unit extension cost c as n and p_h vary. In cases with lower values of n or p, the minimum θ is much smaller than the negative drift threshold of $\frac{c}{p_h}$, as shown in Figure 4-2.

4.3 Making Offers to Subset of Potential Users

If the monopolist is allowed to choose to make offers to $m < n$ potential users instead of all n in the lattice, the expected change in profit by going from m offers to $m + 1$ offers is given by $p_h\theta - c + cp_\ell^{m+1}$. This can be confirmed by solving the difference equation: $\Delta(m) = \mathbb{E}[\bar{V}_{m+1}] - \mathbb{E}[\bar{V}_m]$. The second order difference equation is strictly negative, $c(p_\ell^{m+2} - p_\ell^{m+1})$, which means that as $m \to max(n, \infty)$ the expected profit $\mathbb{E}[\bar{V}_m]$ will cross zero at most once. The optimal number of offers m^* is therefore the largest integer $m \leq n$ satisfying $\Delta(m-1) \geq 0$, which can be expressed as:

$$m \leq \frac{ln(1 - \frac{p_h\theta}{c})}{ln(1 - p_h)}$$

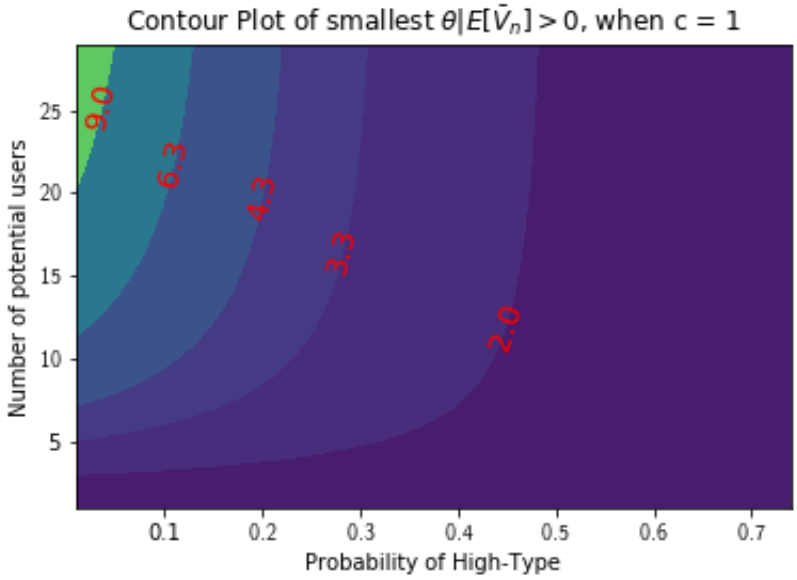

Figure 4-1: Contour plot of the minimum θ, as a multiple of unit extension cost c, that makes offering to build the network up to the last price-taker profitable on expectation.

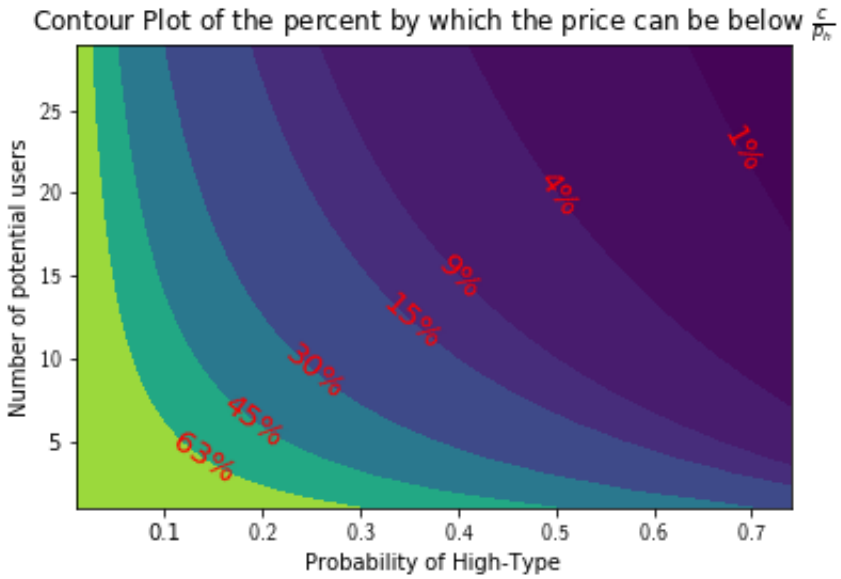

Figure 4-2: Contour plot of the percentage by which θ can fall below the negative drift threshold $\frac{c}{p_h}$ while maintaining positive expected profit.

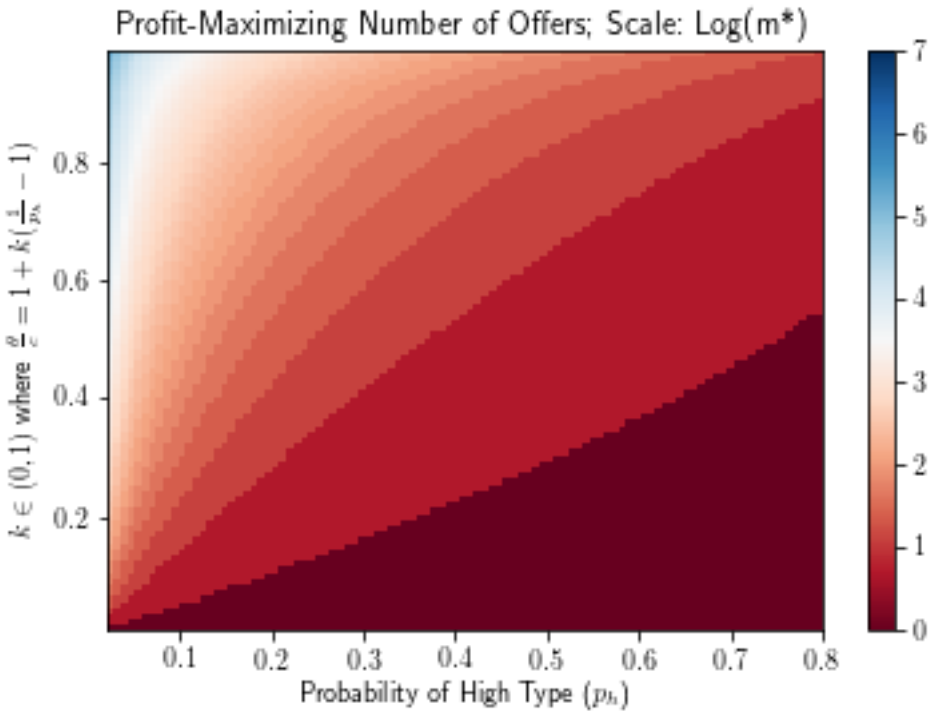

Figure 4-3: A heat map of the natural log of the profit-maximizing number of offers as a function of parameters characterizing the extent the negative drift.

. Figure 4-3 shows that m^* is confined to single digit integers over most of the domain of the plausible values of p_h and θ that have negative drift (scaling units such that $c = 1$ without loss of generality).

To assess whether m^* is small, consider quotient $\frac{m*}{M}$, where $M = \max\{m|\mathbb{E}[\bar{V}_m] \geq 0, m \leq n\}$. M is the largest number of potential users that can be made offers while maintaining positive expected profit, and its range of values is presented in Figure 4-4. The expression for M is cumbersome, involving the Lambert W product-log function, and not particularly informative on its own, so the quotient is numerically computed for a relevant parameter range and presented in Figure 4-5.

If the support for references prices is truly $r_i \in \{r_\ell, r_h\}$, where $r_\ell < c$ and $r_h > c$, then m^* is also the socially optimal solution. However, if the support of r_i contains multiple values above c then m^* may not be socially optimal, the extent and implications of which are discussed in a later section 4.6.

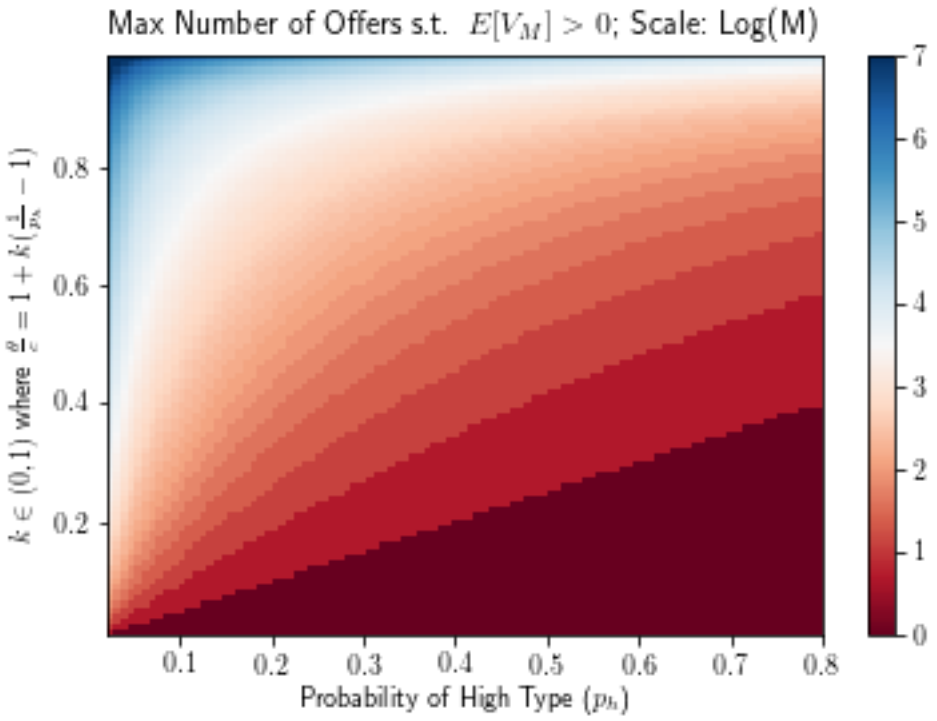

Figure 4-4: A heat map of the natural log of the largest number of offers than can be made while achieving positive expected profit. The axes control the extent of the negative drift.

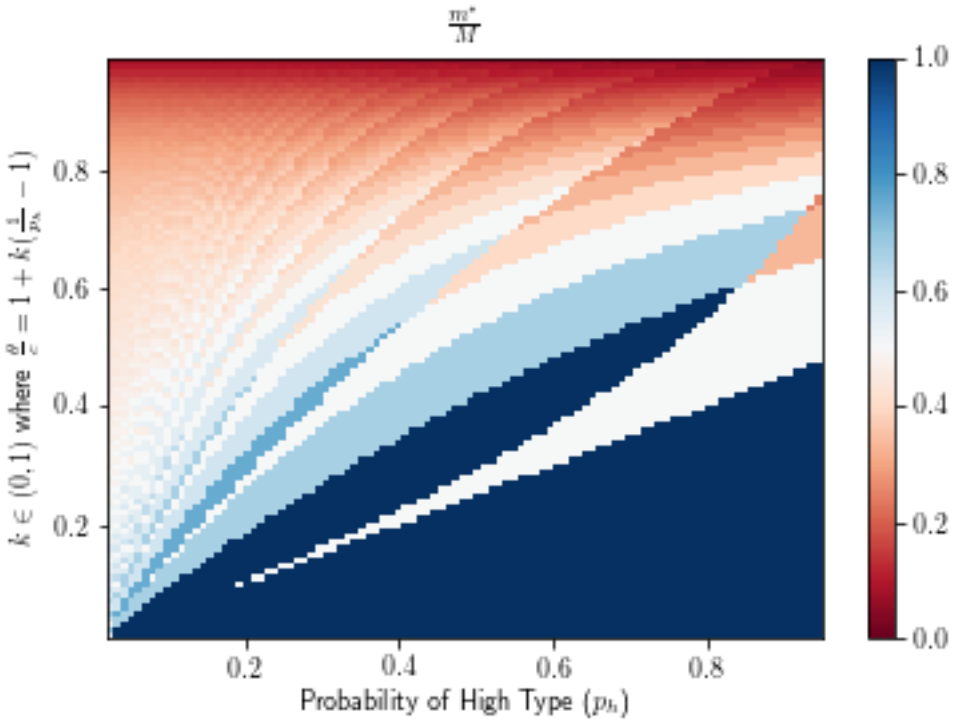

Figure 4-5: Quotient of the profit-maximizing number of offers vs the maximum number of offers that could be made while maintaining positive profit.

4.4 Connect Best Subset After Types Revealed

A monopolist who is given permission to serve any subset of potential users who accept the joining price will realize substantially higher expected profits than when they are required to connect all users who accept. In this setting, the monopolist can decline to build unprofitable networks, and consequently the expected profit is non-negative for all networks no matter what the size or negative drift.

In the case of the simple random walk with negative drift ($\theta = 2, c = 1, \sigma_c = 0, p_h \leq 0.5$, n potential users), the monopolist's expected profit is the expected maximum of the random walk after n steps. Denote this maximum as M_n. The distribution for M_n is given by Mann [2017] and the expectation can be computed in closed form, albeit not by hand. By contrast, the limiting case when $n \to \infty$ is simply $\mathbb{E}[M_\infty] = \frac{p_h}{1-2p_h}, \forall p_h \in (0, 0.5)$ Finch [2018]. To see why, note that the distribution for M_∞ is given by $P(M_\infty \geq k) = (\frac{p_h}{1-p_h})^k$ Shevchenko [2017], Shalop [2017]. Since M_∞ is a positive integer random variable, $\mathbb{E}[M_\infty] = \sum_{k=0}^{\infty} P(M_\infty > k) = \sum_{k=1}^{\infty} (\frac{p_h}{1-p_h})^k = \frac{\frac{p_h}{1-p_h}}{1-\frac{p_h}{1-p_h}} = \frac{p_h}{1-2p_h}$.

The expected profit in the face of negative drift is strictly positive for the finite and infinite random walk when the best subset is chosen. Relaxing the requirement that all customers who accept the joining price be served will enlarge the set of conditions under which a network will be constructed in the first place. However, if high-type users were to believe that revealing their type when no network is built could adversely affect them in the future (e.g., possible price discrimination in a future period if the monopolist does not build the network in the current period), then the offer mechanism could fail to be incentive compatible, and some or all of the gains from picking the optimal subset could be lost, depending on the strategic behavior of potential users.

Figure 4-6 compares $\mathbb{E}[M_\infty]$ to $\mathbb{E}[V_{m^*}]$ for sets of parameters producing a random walk with negative drift: $\theta = 2c, \sigma_c = 0, p_h \in (0, .5)$. The expected profit is expressed units of c. For values of p_h close to 0.5 the drift becomes less negative and the expected profit when subset selection is allowed becomes orders of magnitude larger

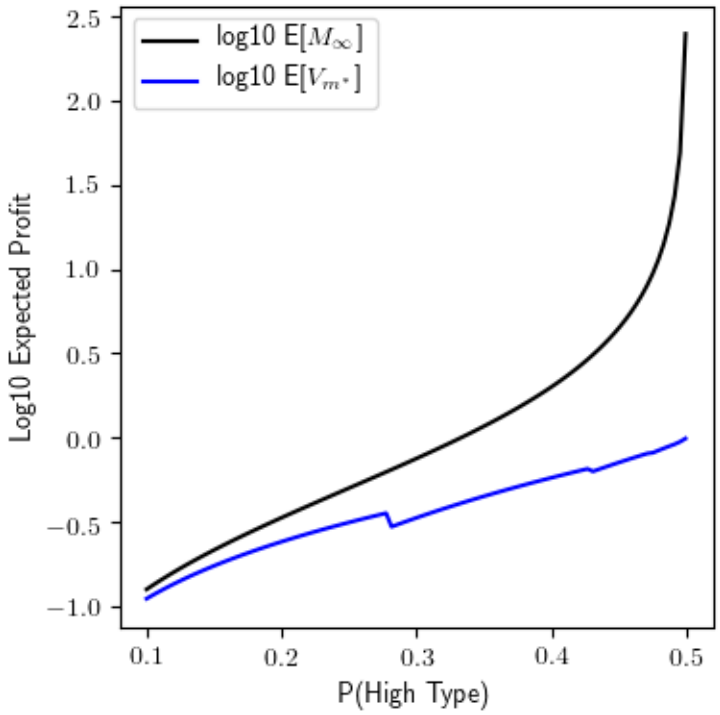

Figure 4-6: Expected profit from optimal subset after types revealed vs expected profit from optimal number of offers with hidden types

than when all price-takers must be served. The magnitude of the advantage for subset selection will be lower if the network is finite and smaller in n, whereas $\mathbb{E}[V_{m^*}]$ will not change as n decreases until $n = m^*$. In the case of sufficiently strong positive drift on a finite network, the optimal subset is unlikely to differ substantially from $m^* = n$. The benefit of allowing subset selection is most pronounced when n is large and when there is negative drift close to 0. These are the circumstances in which it would be most worthwhile to evaluate whether allowing subset selection is worth the risk of exposing users' private reference prices to future monopolist action.

4.5 Multiple User Types

This section outlines how to compute the expected value $\mathbb{E}[V_n(\pi)]$ of offering price π to potential users on a 1D lattice of length n (inclusive of ∞) when users' reference prices have any i.i.d. distribution. A simple solution to maximize $\mathbb{E}[V_\infty(\pi)]$ under

the uniform distribution $r_i \sim U[u, l]$ is presented, and shown to be a lower bound for finite n, whose maximum can be obtained by finding the roots of a $n+1$ degree polynomial or, more practically, through a univariate grid search over $\pi \in [u, l]$.

As before, suppose the monopolist is *required* to connect those who accept the fee proposal. If the reference prices are discrete random variables, the analysis would be similar to the continuous case: the monopolist's choice set would change from $\pi \in [r_\ell, r_h]$ to $\pi \in \{r_1, r_2, ..., r_k, \}$ since no other choice can be optimal (unless non-participation is optimal); for any π between two possible discrete reference price values, profit can increase by setting π to the nearest r_i value above it, as no user's participation decision will change. All users whose reference prices are at or above π become high-type, and the others become low-type. Optimal π^* could be found by enumerating $\mathbb{E}[\bar{V}(\pi)]$ over the finite discrete support.

Indeed, even for continuous distributions of reference prices, by setting a single price for all users we reduce the problem to the previously discussed model with two user types; by offering price π to all users, we can make the following substitutions: $p_h = 1 - F_r(\pi) \equiv F_r^c(\pi)$, and $\theta = \pi$. Making these substitutions in equation 4.1, the monopolist maximizes expected profit by solving:

$$\max_{\pi \in [r_\ell, r_h]} \mathbb{E}[\bar{V}_n(\pi)] = \max_{\pi \in [r_\ell, r_h]} \frac{nF_r^c(\pi)\left[F_r^c(\pi)\pi - c\right] - c\left[F_r(\pi)\right]^{n+1} + c - cF_r^c(\pi)}{F_r^c(\pi)}. \quad (4.3)$$

We can obtain the first-order condition (FOC), noting that $f_r(\cdot)$ is the derivative of the CDF and $-f_r(\cdot)$ is the derivative of its complement for symmetric distributions such as the uniform and normal distribution. To simplify the notation, let f, F, and F^c denote $f_r(p)$, $F_r(p)$, and $F_r^c(p)$, respectively. The following result holds for symmetric distributions:

$$FOC : \frac{\partial \mathbb{E}[\bar{V}_n(\pi)]}{\partial \pi} = 0$$

$$\frac{\partial \mathbb{E}[\bar{V}_n(\pi)]}{\partial \pi} = \frac{\left[2nF^c(-f)\pi + n\left[F^c\right]^2 - cn(-f) - c(n+1)F^n f + cf\right]F^c}{\left[F^c\right]^2} - \frac{(-f)\left[nF^c(F^c\pi - c) - cF^{n+1} + c - cF^c\right]}{\left[F^c\right]^2}$$

$$= -2n\pi f + nF^c + \frac{cnf}{F^c} - \frac{c(n+1)F^n f}{F^c} + \frac{cf}{F^c} + n\pi f - \frac{cnf}{F^c} + cf\frac{(1-F^{n+1})}{\left[F^c\right]^2} - \frac{cf}{F^c}$$

$$= -n\pi f + nF^c - c(n+1)f\frac{F^n}{F^c} + cf\frac{(1-F^{n+1})}{\left[F^c\right]^2} = 0. \quad (4.4)$$

The general FOC is not instructive on its own, and the second-order-condition (SOC) only holds for certain combinations of π, n, c that depend on F. For example, in the case of uniformly distributed reference prices, scaled so that $r_i \sim U[0,1]$, the FOC and SOC are given by the following expressions:

$$\text{U[0,1], FOC: } -c(n(1-\pi)-1)\pi^n + c - n(1-\pi)^2 = 0 \quad (4.5)$$

$$\text{U[0,1], SOC: } -2n - c(n+1)\pi^{n-1}(n + \frac{\pi}{(1-\pi)^2}) + \frac{2c}{(1-\pi)^3}(1-\pi^{n+1}) < 0 \quad (4.6)$$

The FOC is a polynomial of degree $n+1$ that may have multiple real roots. Because of the first term in the SOC, $-2n$, the limit of the LHS of the inequality as $n \to \infty$ is clearly $-\infty$, holding other parameters constant. The profit function is therefore concave for sufficiently large n. It is possible that this concavity would not hold if the optimal $\pi^* \to 1$ as $n \to \infty$, but we can show that this is not the case, and in fact $\pi^* \to 0.5$ as $n \to \infty$. For values of π close to 0.5, n does not need to be large: for $c = 0.25, \pi = 0.5$, strict concavity holds at $n = 1$; for $c = 0.5, \pi = 0.5$, strict concavity holds for $n = 3$. Empirically, the expected profit function when $r_i \sim U[0,1]$ appears unimodal and either concave or quasi-concave in π over a fine mesh and for $n \in \{1,2,3,4,5\}$ and over $c \in \{0.1, 0.2, 0.3, 0.4\}$, as shown in Figure 4-7. Since

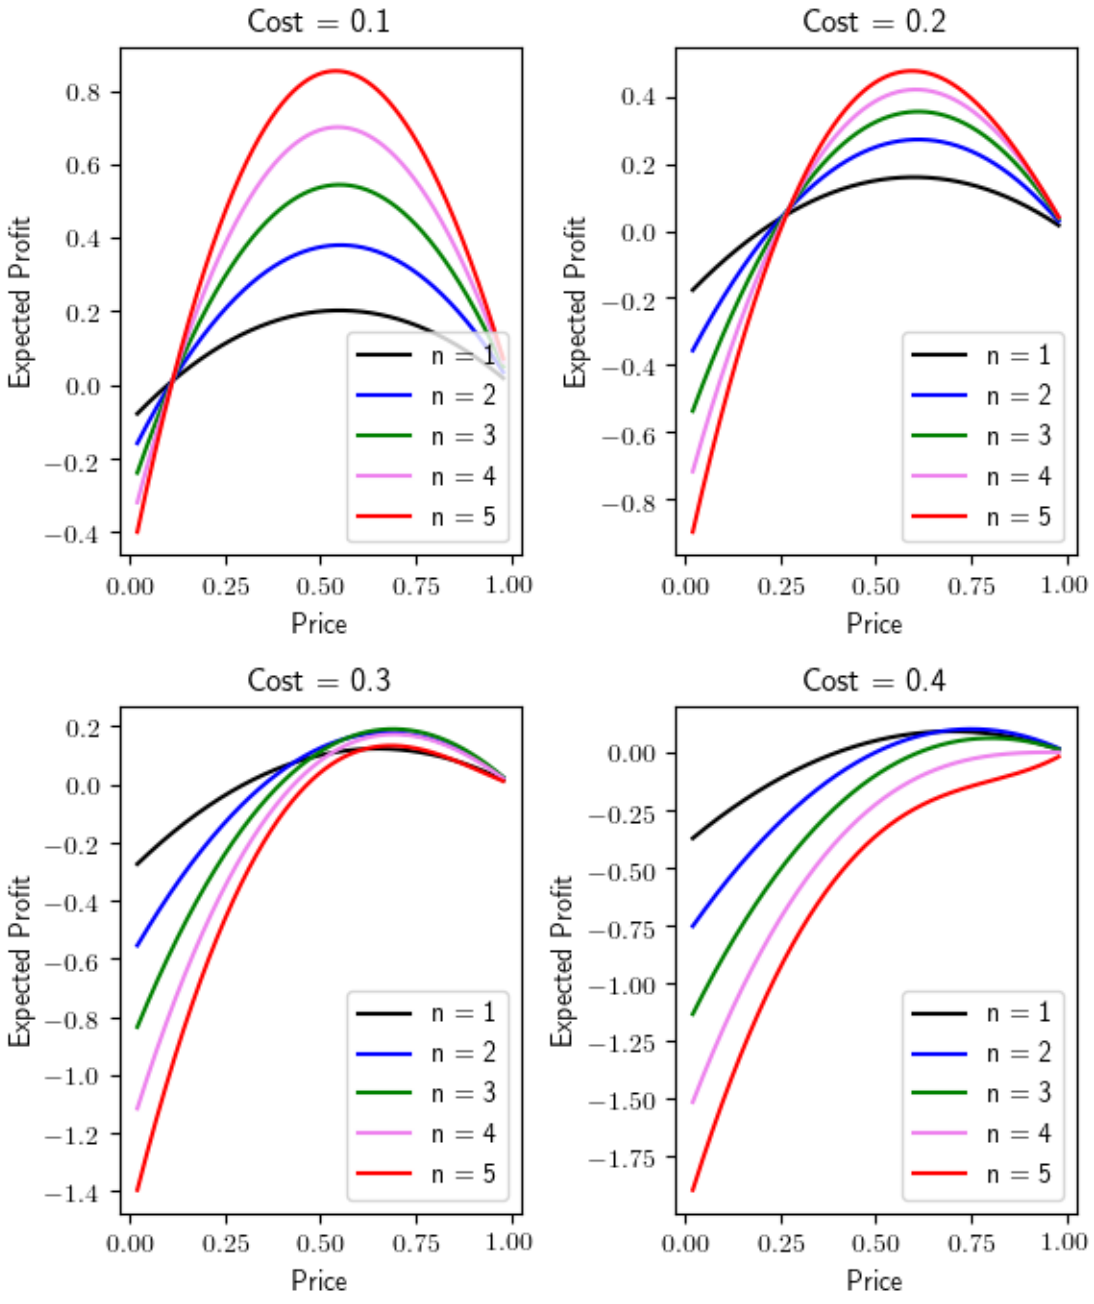

Figure 4-7: Expected profit with respect to price is unimodal and appears concave or quasi-concave for small n and several values of c, where $r_i \sim U[0, 1]$.

expected profit is a function of one bounded variable, π, it may be more practical to approximate the optimal value of π by enumeration over a fine grid for many distributions of interest than to use a gradient-based method when the concavity is not guaranteed.

It is also useful to analyze the limiting case of $n \to \infty$ to establish a lower-bound on the solution for finite n. In the limit, the expected profit per potential user is what matters, so the maximization problem simplifies to $\max_\pi : F^c(\pi)\pi - c$. Suppose the reference prices are uniformly distributed: $r_i \sim U[l, u]$ i.i.d., then the probability that a customer accepts price π is given by

$$p_h = F^c(\pi) = \begin{cases} \frac{u-\pi}{u-l} & \pi \in [l, u] \\ 1 & \pi < l \\ 0 & \pi > u \end{cases} \tag{4.7}$$

The objective function becomes $\max_{\pi \in [l,u]} \frac{-\pi^2 + u\pi}{u-l} - c$, which attains its maximum at $\pi^* = \max\{l, \frac{u}{2}\}$. The optimal offer price for finite n is at least half the maximum of the uniform distribution. We can also obtain a constraint that expresses the maximum unit cost that preserves profitability, given the distribution $U[l, u]$:

$$c \leq \begin{cases} l & \text{if } l > \frac{u}{2} \\ \frac{u^2}{4(u-l)} & \text{if } l < \frac{u}{2} \end{cases} \} \tag{4.8}$$

For the simple case of $l = 0$, $c \leq \frac{u}{4}$, which reveals that for uniformly distributed reference prices the upper price limit has to be at least four times the unit cost. The optimal π for finite n and the expected profit per potential user is larger than the infinite n case. This difference is non-trivial. For $n \leq 10$, the optimal π will be 10%-25% larger across a range of possible unit costs, as shown in Figure 4-8.

For general symmetric unimodal distributions, the FOC of the limiting case is $F^c(\pi^*) = f\pi^*$ and the SOC is $-f' - 2f = 0$. The SOC for the normal distribution

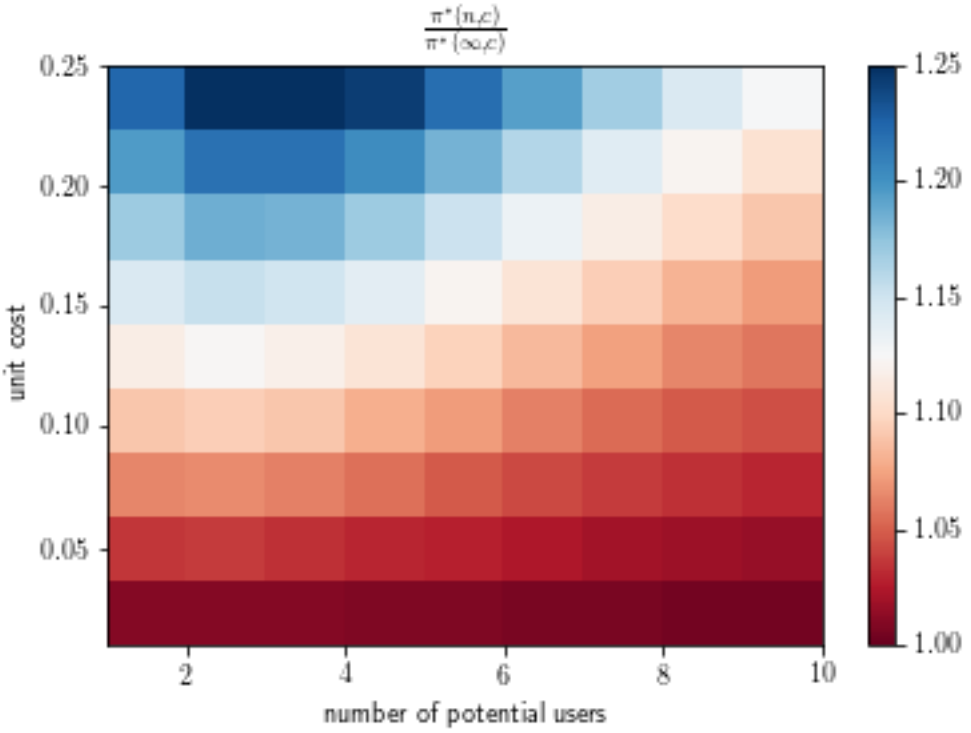

Figure 4-8: Ratio of optimal price for finite n vs. optimal price for infinitely long network. $r_i \sim U[0,1]$.

has a convenient closed form of $-f' - 2f <= 0 \iff \frac{-f'}{f} - 2 < 0 \iff \frac{\pi^* - \mu}{\sigma^2} - 2 < 0 \iff \pi^* < \mu + 2\sigma^2$. This inequality provides a starting point for a univariate grid search or golden section search.

4.5.1 Optimal Individual Prices by Distance

A network monopolist might consider the problem of offering optimal individual prices to each user based on their location on the 1D lattice. To compute optimal individual prices, the monopolist must optimize a non-linear function for expected profit with n price variables. Unfortunately, the expected profit function is not quasi-concave (and therefore not concave either), so stationary points found through a gradient-based method are not guaranteed to be optimal. We are so far unable to offer an elegant, globally-optimal solution. The gradient, Hessian, and proof of non-quasi-concavity for this problem are given in Appendix A. The existence of another solution method that finds the globally optimal set of individual prices is not ruled out.

4.6 Impact on Total Utility of Making Offers to a Subset

Previously, we observed that there can be a large gap between the optimal number of users to whom offers should be made, m^*, and the maximum number to whom offers should be made subject to no losses in expectation, M, when there is negative drift. When there are only two reference prices, all available utility is realized in the form of profit for the monopolist. However, when reference prices are continuous, some of the utility takes the form of consumer surplus from users who's reference prices were above π^*. When a monopolist makes m^* offers instead of n or M offers, this consumer surplus is reduced because fewer users are served, and because the optimal price π^* is slightly higher when the number of offers made is smaller. The difference in total utility between m^* and M offers can be substantial, as will be demonstrated by further examining the uniform case where $r_i \sim U[l, u]$.

The expected consumer surplus for making offers at price π to n users is given by $CS(\pi, n) = n(\frac{u-\pi}{2})(\frac{u-\pi}{u-l})$, which is the expected consumer surplus per user who accepts, times the probability of a user accepting π, times the number of offers made. The total utility for making M offers is therefore given by: $TU(M) = \mathbb{E}[\bar{V}_M(\pi^*_M)] + CS(\pi^*_M, M)$, where π^*_M is the profit maximizing price when M offers are made, where M is the maximum integer such that expected profit, when π is chosen optimally, is non-negative. $TU(m^*)$ is defined similarly for the profit-maximizing number of offers.

As shown in Figure 4-9, the relationship between $TU(M)$ and $TU(m^*)$ depends on the unit cost c and is characterized by four phases. Phase 1 is characterized by non-negative drift, so $M = m^* = n$, which occurs at $c \leq 0.25$ for $r_i \sim U[0, 1]$. Phase 2 is characterised by an abrupt and dramatic loss of total utility when the drift becomes slightly negative. For any level of negative drift, m^* tends towards single digits while M can be orders of magnitude larger. In phase 3, as c increases, the negative drift becomes large enough that M also takes on single digit values and differs only slightly from m^*. The total utility under both offer policies is very low in phase 3, so their relative differences are less consequential. In phase 4, the unit cost causes the drift to be so negative that $m^* = M = 1$, and total utility under both policies will be small and equal for all larger values of c, until c surpasses the upper bound of the support for the reference price distribution, in which case $m^* = M = 0$.

In practice, the true unit cost and reference price distribution may be somewhat uncertain. The analysis presented in Figure 4-9 favors regulatory intervention to cause the monopolist to serve a number closer to the estimated M than the estimated m^*. Utility will be the same in phase 1 and 4, orders of magnitudes higher in phase 2, and very comparable, even if worse, in phase 3, where total utility is low in either case.

4.7 Random Costs and Rewards

The analysis in previous sections relied on several simplifying assumptions to compute expected profit exactly. However, if an approximation of the mean is acceptable, then even when these assumptions are violated it is possible to approximate the mean

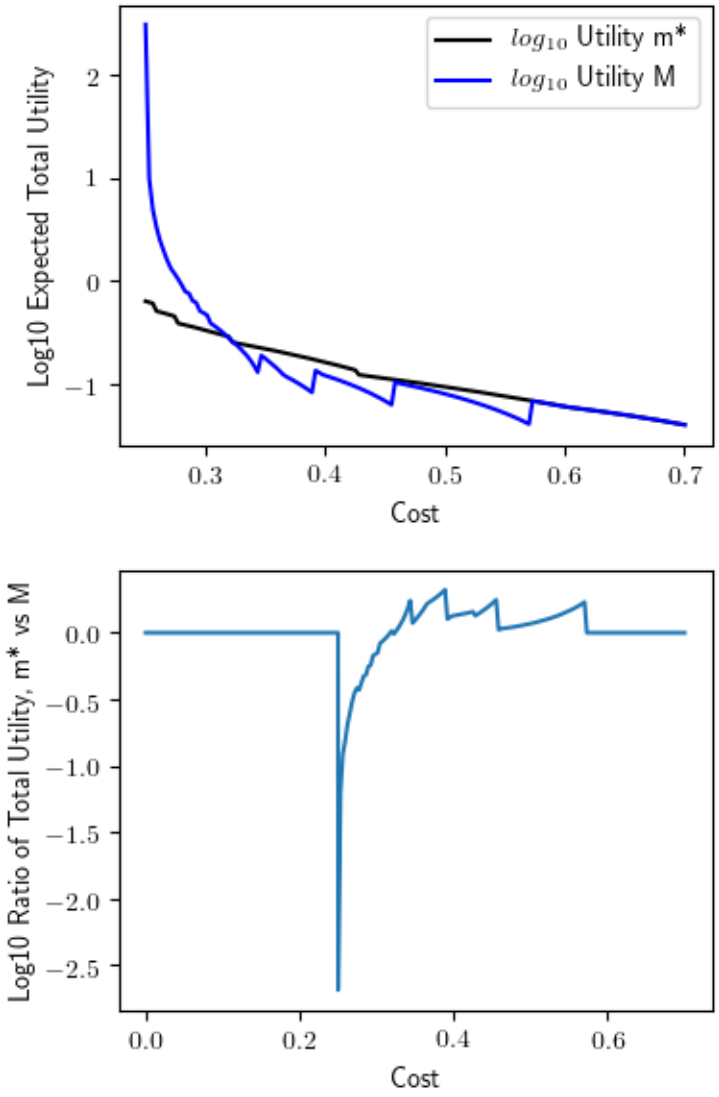

Figure 4-9: Total Utility (log-scaled) vs cost from (0.25,0.7] in top graph and the log of the ratio $\frac{TU(m^*)}{TU(M)}$ vs cost from (0,0.7] on the lower graph. Results are for $r_i \sim U[0,1]$.

with simple closed-form expressions. Previously, we assumed that the revenue for connecting a high-type user is only the offer price π. We also assumed that the expected marginal network extension cost is $\mathcal{N}(c, \sigma_c)$; more strictly, in the section that analyzed the value of picking the optimal subset of price-takers we set $\sigma_c = 0$ and $\pi = 2c$ so that we could treat profit as a discrete random walk and apply known results for its expected maximum.

In most real-world applications, network infrastructure provides users benefits over time and network operators often charge usage-based fees in addition to network connection fees. If all users pay the same usage-based fees and generate the same usage-based costs to the monopolist, then we can continue using the models previously developed by simply redefining π as the net present value of the connection fee plus future profits from usage fees. However, this modification requires the monopolist to know the probability of pricing policy rejection by users. Let $A(\pi)$ be a pricing policy that changes the present value of profit by π per (non-marginal) user joining the network, and by $\pi - c$ for a marginal user joining the network. Denote the probability of policy $A(\cdot)$ being rejected when the marginal profit target is π by $F_A(\pi)$. For any alternative pricing policy B it is not necessary to assume that $F_B(\pi) = F_A(\pi)$. To extend the results of previous sections to the setting with both fixed network connection costs and marginal usage charges, we need only suppose that $F_A(\pi)$ exists and is known for the set of reasonable policies that the monopolist considers using. Explicitly characterizing the optimal policy (i.e., a policy $\Pi() : F_\Pi(\pi) \leq F_A(\pi), \forall A, \forall \pi > 0$) would require specific assumptions about user utility functions, discount rates, risk-aversion, and game-theoretic considerations, all of which may be a worthwhile subject of future research.

The assumption that network joiners will generate identical levels of usage-related profit is rather strong and restrictive. In relaxing this assumption, we will consider two types of variability in the present value profit-contribution of users. First, let γ_i be the component of random variation in user i's contribution that could be learned or estimated, at some cost to the monopolist, before the network construction offer is made, such as whether the user plans to consume a high volume of network services,

is interested in high mark-up premium services, or whether observable physical conditions at or around their house make it cheaper or costlier than average to activate their connection to the network. Second, let ϵ_i be the random variation in user i's profit contribution that cannot be predicted, such as whether a user is disposed to make excessive and unnecessary calls to customer service, or whether hidden physical conditions at or around their house make it cheaper or costlier than average to activate their connection to the network.

In this random-reward setting, $A(\pi)$ is redefined as a policy that improves *expected* present value of profit by π per non-marginal joiner. Let $\mathbb{1}_{A(\pi)}(i)$ be the indicator function for whether user i accepts policy $A(\pi)$. We will also assume $\forall i$ the following: (i) γ_i and ϵ_i have finite variances $\sigma^2_{\gamma,A(\pi)}$ and $\sigma^2_{\epsilon,A(\pi)}$ that may depend on the subset of users who accept policy $A(\pi)$ but are otherwise independent of $\mathbb{1}_{A(\pi)}(i)$; (ii) that $\mathbb{E}[\gamma_i] = \mathbb{E}[\epsilon_i] = 0$; (iii) that each sequence $\{\gamma_i\}$ and $\{\epsilon_i\}$ is i.i.d. internally and independent from the other; (iv) and that any benefit incurred as a result of policy $A()$ causing the population of joiners to be more profitable to serve than the general population is incorporated into the definition of policy $A(\pi)$ delivering π on expectation.

Under these assumptions, we can characterise statistical properties of the random rewards from offering policy $A(\pi)$ to n potential users as follows:

$$\text{Reward from user } i, \text{ given } A(\pi)\text{: } \pi_i = (\pi + \gamma_i + \epsilon_i)\mathbb{1}_{A(\pi)}(i) \quad \forall i$$

$$E\left[\pi_i | A(\pi)\right] = \pi(1 - F_A(\pi)) \Rightarrow E\left[\sum_{i=1}^{n} \pi_i\right] = n\pi(1 - F_A(\pi)) = n\pi F^c_A(\pi)$$

$$Var(\pi_i | A(\pi)) \equiv \sigma^2_{A(\pi)} = \pi^2 F^c_A(\pi) F_A(\pi) + F^c_A(\pi)(\sigma^2_{\gamma,A(\pi)} + \sigma^2_{\epsilon,A(\pi)})$$

$$\Rightarrow Var(\sum_{i=1}^{n} \pi_i | A(\pi)) = n\sigma^2_{A(\pi)} \quad \text{by i.i.d. property.}$$

In the variance calculation, we use the result that for any random variables X, Y such that $X \perp Y, Y \sim Bernoulli(p), \mathbb{E}[X] = 0$, and $\sigma^2_X < \infty$ the variance

of their product is given by $Var(XY) = \sigma_X^2 p$. This can be seen through the following steps: $Var(XY) = \mathbb{E}[(XY)^2] - \mathbb{E}[XY]^2 = \mathbb{E}[X^2Y^2] - \mathbb{E}[X]\mathbb{E}[Y] = \mathbb{E}[X^2]\mathbb{E}[Y^2] - 0\mathbb{E}[Y] = \mathbb{E}[X^2]\mathbb{E}[Y] = \sigma_X^2 p$.

Next, we will relax certain simplifying assumptions on network construction costs before examining the overall stochastic process that describes the combined effect of the random rewards and random costs. Previously, we assumed that the total number of potential users along the lattice is known. This assumption is reasonable for small networks. If the network is large, or if the monopolist needs to evaluate offerings to a very large number of distinct locations, we could also assume that the number of users depends on a counting process $n(d)$, where d is the distance from the source node along the lattice. If we scale the units of d such that the average distance between adjacent potential users is 1, then $\mathbb{E}[n(d)] \approx d$ for sufficiently large d, with equality holding if $n(d)$ is a Poisson counting process with respect to d, analogous to time. For variance, $Var(n(d)) = d\sigma_n^2$, with $\sigma_n^2 = 0$ if $n(d)$ is a Poisson process.

Next, we model the expected cost of building the network to distance d as being linear in d: $\mathbb{E}[C(d)] = cd$, where c is the mean cost per unit of d. As with the random rewards, suppose that deviations $C(d2 - d1) - c(d2 - d1)$ from the expected value of construction cost over any interval can be decomposed into (i) variation that can be learned or approximated at some cost to the monopolist, such as observable terrain conditions, and (ii) sources of variation that cannot be learned in advance, such as hidden terrain conditions. Let $g(\cdot, \omega) : X \to \mathbb{R}\ \forall \omega \in \Omega$ be a piece-wise integrable random function representing how much the conditional expectation of the marginal cost at location x deviates from the overall mean c when the state of nature, ω, is known. More formally, $\mathbb{E}[C(d2 - d1)|g(\cdot, \omega)] = \int_{d1}^{d2}(g(x, \omega) + c)dx$, and we assume that $g(x)$ can be integrated piece-wise. When ω is unknown, we assume that the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ produces definite integrals over $g(x)$ that are random variables satisfying $E\left[\int_{d1}^{d2} g(x)dx\right] = 0\ \forall(d1, d2)$ and $Var\left(\int_{d1}^{d1+\Delta} g(x)dx\right) = \Delta\sigma_g^2\ \forall\Delta > 0$. Let the remaining random noise over the interval, denoted by $\delta(d2-d1)$, be normally distributed with mean of zero, variance proportional to the interval length, and with variance per unit denoted by σ_δ^2. Moreover, if the major statistical

dependencies in cost between segments, if any, are captured by $g(x, \omega)$, then $\delta()$ may come close to satisfying the independent increments property, i.e., the unlearnable random cost perturbation over any interval would be independent of the perturbation observed in past intervals.

Under these assumptions, $\mathbb{E}[C(d)|g(d,\omega)] - C(d)$ can be expected to behave as a Wiener process (a.k.a. Brownian motion process) Cox [1977] with zero drift and standard deviation σ_δ^2 per unit of d. If we cannot condition on $g(d)$, and if $g(d)$ has an independent increments property for sufficiently large increments (allowing for dependencies of small or medium increments), the global behavior of $\mathbb{E}[C(d)] - C(d)$ will approximate that of a Wiener process with zero mean and unit variance $\sigma_g^2 + \sigma_\delta^2$ as $d \to \infty$. The quality of the approximation will be higher if for some reasonably large N we define $\Delta = \frac{d_{max}}{N}$ and have $\int_{d_i}^{d_i+\Delta} g(x)dx \;\dot{\sim}_{iid}\; \mathcal{N}(0, \Delta\sigma_g^2)$, at which point the N independent increments, though discrete, will approach the behavior of a Wiener process, which is known to characterize the limiting behavior of random walks Cox [1977].

Denote with $V(d, A(\pi))$ the stochastic process that characterizes the monopolist's profit as a function of offering to build a network of length d, claiming random rewards π_i from users who accept policy $A(\pi)$, and incurring random costs up to d. If cost perturbations $g(\cdot)$ and $\delta(\cdot)$ are independent from the γ_i and ϵ_i then the following relationships hold:

$$V(d, A(\pi)) = \sum_{i=1}^{n(d)} (\pi + \gamma_i + \epsilon_i) \mathbb{1}_{A(\pi)}(i) - \int_0^d \left[g(x) + c \right] dx - \delta(d). \tag{4.9}$$

Use Wald's Equation $E\left[\sum_{i=1}^{N} X_i\right] = \mathbb{E}[N]\mathbb{E}[X],\ X_i \sim i.i.d., N \in \mathbb{N}$ to get

$$\mathbb{E}[V(d, A(\pi))] = \mathbb{E}[n(d)](\pi F_A^c(\pi) - c) \approx d\left(\pi F_A^c(\pi) - c\right), \tag{4.10}$$

Define drift $\mu_V \equiv \pi F_A^c(\pi) - c,$ and

$$Var(V(d, A(\pi))) = Var\left(\sum_{i=1}^{n(d)} \pi_i | A(\pi)\right) + Var\left(\int_0^d \left[g(x) + c \right] dx\right) + Var(\delta(d))$$

We can simplify each component of the variance as follows:

$$Var(\delta(d)) = d\sigma_\delta^2;$$

$$Var\left(\sum_{i=1}^{n(d)} \pi_i | A(\pi)\right) = E\left[n(d)\right] Var(\pi_i | A(\pi)) + \mathbb{E}[\pi_i | A(\pi)]^2 Var(n(d))$$

$$\approx d\sigma_{A(\pi)}^2 + \pi^2 (F_A^c(\pi))^2 d\sigma_n^2 \ = \ d\left[\sigma_{A(\pi)}^2 + \pi^2 (F_A^c(\pi))^2 \sigma_n^2\right];$$

$$Var\left(\int_0^d \left[g(x) + c \right] dx\right) = d\sigma_g^2.$$

$$\therefore\ Var(V(d, A(\pi))) \approx d\left[\sigma_{A(\pi)}^2 + \pi^2 (F_A^c(\pi))^2 \sigma_n^2 + \sigma_g^2 + \sigma_\delta^2\right]$$

$$= d\left[\pi^2 F_A^c(\pi) F_A(\pi) + F_A^c(\pi)(\sigma_{\gamma, A(\pi)}^2 + \sigma_{\epsilon, A(\pi)}^2) + \pi^2 (F_A^c(\pi))^2 \sigma_n^2 + \sigma_g^2 + \sigma_\delta^2\right]. \tag{4.11}$$

The variance of the stochastic process $V(d, A(\pi))$ is important for several reasons. First, any monopolist that is not risk-neutral need to consider it as part of their investment decision. Second, even for a risk neutral monopolist, it is necessary to know the variance in order to estimate the expected maximum of the process in the case where the monopolist is permitted to choose to connect to a subset of those who accept its offer. Finally, by knowing the variance formula the monopolist can compute the value of revealing the learnable sources of uncertainty and decide whether acquiring such information ahead of time is worth the cost.

4.7.1 Behavior of Risk-Averse Monopolist

Even if the monopolist is risk-averse, it will always prefer larger d to smaller d in a positive drift setting since expected profit increases linearly in d while the standard deviation increases in the square-root of d. The optimal π^*_{RA} for a risk-averse monopolist will satisfy $Var(V(d, A(\pi^*_{RA}))) \leq Var(V(d, A(\pi^*)))$; by differentiating equation 4.11 with respect to π, we can determine whether a risk-averse monopolist will be inclined to lower or raise the joining price from the risk-neutral optimal level. For any differentiable symmetric distribution (i.e., $\frac{\partial F^c}{\partial \pi} = -f$), we obtain the following:

$$\frac{\partial Var(V(d, A(\pi)))}{\partial \pi} = d\left[2\pi F^c F - \pi^2 f F - f(\sigma^2_{\gamma,A(\pi)} + \sigma^2_{\epsilon,A(\pi)})\right] + d\left[F^c\left(\frac{\partial \sigma^2_{\gamma,A(\pi)}}{\partial \pi} + \frac{\partial \sigma^2_{\epsilon,A(\pi)}}{\partial \pi}\right) + 2\pi(F^c)^2\sigma_n^2 - (1 - 2\sigma_n^2)\pi^2 f F^c\right].$$

If $\frac{\partial Var(V(d,A(\pi)))}{\partial \pi} \geq 0$ at $\pi = \pi^*$ then $\pi^*_{RA} \leq \pi^*$. This partial derivative has both positive and negative terms, so little can be said in general about whether $\pi^*_{RA} \leq \pi^*$ without first specifying the distribution $F(\cdot)$. The partial derivative is decreasing in both $\sigma^2_{\gamma,A(\pi)}$ and $\sigma^2_{\epsilon,A(\pi)}$ at rate $-f(\cdot)$, which could cause the partial derivative to turn negative when reward-related variance is extremely high, leading to $\pi^*_{RA} > \pi^*$ for sufficiently large $\sigma^2_{\gamma,A(\pi)}$ or $\sigma^2_{\epsilon,A(\pi)}$. An intuitive interpretation for this scenario is that charging a higher price will generate a higher fixed profit margin per customer while reducing the number of customers such that total revenue may be only slightly lower, but variance in profit from serving users, which is proportional to the number of users, will be reduced by serving fewer users. By contrast, when the partial derivative is positive, the main consideration for the risk-averse monopolist is that revenue generation will be more variable when trying to earn high margins per customer with a lower probability of acceptance per customer.

The partial derivatives $\frac{\partial \sigma^2_{\gamma,A(\pi)}}{\partial \pi}$ and $\frac{\partial \sigma^2_{\epsilon,A(\pi)}}{\partial \pi}$ could be either positive or negative, as

a higher price could lower variance by producing a more homogeneous set of users, but if variation per user scales somewhat proportionally with their usage-based fees, then per-user variation amongst higher-type users might actually be higher than the general population. The effect of σ_n^2 is also ambiguous. If $F^c > \pi f$ then the contribution of the σ_n^2 term to the partial derivative will be positive, in which case higher variation in the user count per unit of distance will nudge a risk-averse monopolist towards a lower π^*_{RA}; the opposite holds true when $F^c < \pi f$.

4.7.2 Connect Optimal Subset - Value of Information

Understanding the variance of $V(d, A(\pi))$ also matters when the monopolist is allowed to build the network up to some $d' \leq d$ after price-taking information, and potentially other information, is revealed. It is not possible to identify $d^* = \text{argmax}_{0,d} V(d, A(\pi^*))$ because the random cost and reward shocks, ϵ_i and $\delta(\cdot)$, cannot be revealed in advance. However, if the monopolist learns how many potential users accept its offer at each value of d, i.e., it learns the function $n(\cdot)$, then it will choose $d' = \text{argmax}_{0,d} \mathbb{E}[V(d, A(\pi^*)|n(\cdot)]$. The monopolist could also make some expenditure to reveal the sequence of learnable reward shocks $\{\gamma_i\}$, and/or to reveal the learnable cost perturbation function $g(\cdot)$. If such information is revealed, then it will choose $d' = \text{argmax}_{0,d} E\left[V(d, A(\pi^*)|n(\cdot), \{\gamma_i\}, g(\cdot)\right]$.

The stochastic processes $\mathbb{E}[V(d, A(\pi^*)|n(\cdot)]$ and $E\left[V(d, A(\pi^*)|n(\cdot), \{\gamma_i\}, g(\cdot)\right]$ are not Wiener processes. However, because of the central limit theorem, and so long as the independent increments property holds for sufficiently large segments of d, their behavior as $d \to \infty$ will asymptotically approach that of a Wiener process Cox [1977]. If distributions for the random rewards and random costs do not severely depart from normality and independence, then even for smaller d a Wiener process might serve as a reasonable approximation, or at least provide a bound for the expected maximum. For example, in the event of positive correlations within the processes $n(d), \{\gamma_i\}, \{\epsilon_i\}, g(\cdot)$, and $\delta(\cdot)$, let alone correlations between them, a Wiener process approximation will underestimate the true variance of $\mathbb{E}[V(d, A(\pi^*)|n(\cdot)]$ and $E\left[V(d, A(\pi^*)|n(\cdot), \{\gamma_i\}, g(\cdot)\right]$. As a result the expected maximum of the Wiener pro-

cess will serve as a lower bound on the true expected maximum, and it will serve as an upper bound in the event of negative correlations in the increments or between the sub-processes.

A Wiener process with zero-drift and variance of 1 per unit of time has an expected maximum between 0 and t given by $\sqrt{\frac{2t}{\pi}}$ Shreve [2008], where π in this case represents the number pi and will be typed numerically henceforth to distinguish it from the decision π. Our processes plays out in units of distance instead of time. The unit variance for $\mathbb{E}[V(d, A(\pi))|n(\cdot)]$ can be extracted from equation 4.11 by dropping the terms that include variance for unrevealed information (i.e., dropping $\sigma^2_{\gamma,A(\pi)}, \sigma^2_{\epsilon,A(\pi)}, \sigma^2_g$ and σ^2_δ). We therefore obtain the following approximation for the expected maximum when drift is zero:

$$\mathbb{E}\left[\max_{d' \leq d}(\mathbb{E}[V(d', A(\pi))|n(\cdot)]\right] \approx \sqrt{\frac{d\left[\pi^2(F^c F + (\sigma_n F^c)^2)\right]}{3.14159...}}. \tag{4.12}$$

The unit variance for $E\left[V(d, A(\pi^*)|n(\cdot), \{\gamma_i\}, g(\cdot)\right]$ can be extracted from equation 4.11 by dropping $\sigma^2_{\epsilon,A(\pi^*)}$ and σ^2_δ, thereby giving the following expected maximum when drift is zero:

$$\mathbb{E}\left[\max_{d' \leq d}(\mathbb{E}[V(d', A(\pi))|n(\cdot), \{\gamma_i\}, g(\cdot)]\right] \approx \sqrt{\frac{d\left[\pi^2(F^c F + (\sigma_n^2 F^c)^2) + \sigma^2_{\gamma,A(\pi)} F^c + \sigma^2_g\right]}{3.14159...}} \tag{4.13}$$

The difference between equation 4.13 and equation 4.12 gives the expected value of learning $\{\gamma_i\}$ and $g(\cdot)$ in the zero-drift setting when subset selection is allowed. The value of learning just $\{\gamma_i\}$ or just $g(\cdot)$ can be computed in a similar manner by dropping the terms σ^2_g or $\sigma^2_{\gamma,A(\pi)}F^c$, respectively, from equation 4.13. These equations give the monopolist a way of establishing a limit on how much it should spend to reveal $\{\gamma_i\}$ and/or $g(\cdot)$. It is worth noting the value of information grows proportional to the square root of d. If the cost of acquiring information grows proportional to d, then $\exists \tilde{d} < \infty$ at which it is no longer profitable to acquire further information.

When a Wiener process has negative drift $\mu < 0$ and standard deviation per unit of time σ, the expected maximum of an infinitely long process is given by $\frac{\sigma}{2|\mu|}$ Resnick [1992]. The expected maximum of the infinite-time process with negative drift is therefore:

$$\mathbb{E}\left[\max_{d<\infty}(\mathbb{E}[V(d, A(\pi))|n(d)])\right] \approx \frac{\pi\sqrt{F^c F + (\sigma_n F^c)^2}}{2|(\pi F_A^c(\pi) - c)|}. \tag{4.14}$$

This expected maximum provides a benchmark and upper bound on the expected maximum under finite d with negative drift. The expected maximum with learned information and under negative drift is as follows:

$$\mathbb{E}\left[\max_{d<\infty}(\mathbb{E}[V(d', A(\pi))|n(\cdot), \{\gamma_i\}, g(\cdot)]\right] \approx \frac{\sqrt{\pi^2(F^c F + (\sigma_n^2 F^c)^2) + \sigma^2_{\gamma,A(\pi)} F^c + \sigma_g^2}}{2|(\pi F_A^c(\pi) - c)|}. \tag{4.15}$$

The expected value of learning $\{\gamma_i\}$ and $g(\cdot)$ is given by the difference between equation 4.15 and equation 4.14. For finite d, however, this difference between the two infinite cases will overestimate the value of learning. The difference is still useful insofar as if the cost of learning $\{\gamma_i\}$ and $g(\cdot)$ up to some finite d exceeds the expected value of learning in the infinite case, then it is not worthwhile to pay to learn $\{\gamma_i\}$ and $g(\cdot)$ up to point d. The expected maximum in the cases of positive or negative drift and finite d are harder to compute and are left out of scope.

4.7.3 Value of Information When Serving All Price-Takers

Returning to the setting in which the monopolist is required to connect all potential users that accept policy $A(\pi)$, we can compute the expected value of learning $g(\cdot)$ and $\{\gamma_j\}$ before deciding whether to make an offer. In this case $\{\gamma_j\}$ has to be indexed over all users within distance d rather than indexing over the price-takers since the price-taking subset has not been revealed yet. This modification requires that $\gamma \perp A(\pi)$; otherwise, it would be necessary to model the effect of learning $\{\gamma_j\}$ on probability $A(\pi)$'s acceptance by individual users. The policy that maximizes

expected profit is to offer to build the network whenever the conditional expectation of profit is non-negative. The event that the conditional expectation is non-negative is given by $\{d(\pi F_A^c(\pi) - c) + \sum_j \gamma_j - \int_0^d g(x)dx \geq 0\}$; denote this event as Y. The expected profit under this policy is $\mathbb{E}[V(d, A(\pi))|Y]\mathbb{P}(Y)$. The value of learning is the difference $\mathbb{E}[V(d, A(\pi))|Y]\mathbb{P}(Y) - \mathbb{E}[V(d, A(\pi))]$. Expanding the second term using the law of total expectation, the difference simplifies to $\mathbb{E}[V(d, A(\pi))|Y^c]\mathbb{P}(Y^c) = \int_{d(F_A^c(\pi)-c)}^{\infty} t\mathbb{P}(\int_0^d g(x)dx - \sum_j \gamma_j = t)dt$. The solvability of this convolution depends on the distributions for $\sum_j \gamma$ and $\int_0^d g(x)dx$. If $\sum_j \gamma$ and $\int_0^d g(x)dx$ are approximately normal and independent for sufficiently large d, then we can treat their sum $\mathcal{S}$ as also being approximately normal with mean drift $\mu_{\mathcal{S}} = \mu_\gamma - \mu_g = 0$ and unit variance $\sigma_{\mathcal{S}}^2 = (\sigma_g^2 + \sigma_\gamma^2)$. This assumption produces an expected value of information that is straightforward to compute:

$$\mathbb{E}[V(d, A(\pi))|Y^c]\mathbb{P}(Y^c) \approx \frac{1}{\sqrt{\tau d\sigma_{\mathcal{S}}^2}} \int_{d\mu_V}^{\infty} (d\mu_V - x)e^{-\frac{x^2}{2d\sigma_{\mathcal{S}}^2}} \partial x. \tag{4.16}$$

4.8 Generalizations to Trees

The one-dimensional lattice assumption is, admittedly, a significant simplification of the general 2D setting in which customers could have arbitrarily distributed locations. The general 2D setting, however, does not yield closed-form optimal policies. Nonetheless, one-dimensional results are relevant to 2D settings. Computing the expected value of making a network construction offer in the 2D setting involves taking the expectation over all minimum spanning trees, the number of which grows exponentially in n. Fortunately, a tractable lower bound on the expected value can be computed using $N < n$ subtrees of the minimum spanning tree that connects all n potential users. This lower bound using subtrees can be computed efficiently by decomposing the tree into 1D lattices and applying techniques developed earlier in this section.

Define $\mathcal{T}_{\mathcal{M}}$ as the minimum spanning tree that connects a potential user subset $\mathcal{M}$

of price-takers to the source node. Let $\mathcal{T}$ be the minimum spanning tree connecting all n potential users to the source node, $\mathcal{S}_{\mathcal{M}}$ as the subtree of $\mathcal{T}$ that connects the subset $\mathcal{M}$ to the source node. Extending the notation introduced for the two-type random walk model, let $V(\mathcal{T}', \pi)$ be the profit of building tree $\mathcal{T}'$ and collecting revenue π from all users for whom $pi \leq r_i$, and let $\bar{V}(\mathcal{T}', \pi)$ be the profit for making offers to the users in tree $\mathcal{T}'$ and building a subtree of $\mathcal{T}'$ only up to those who accept price π.

For any set of price-takers $\mathcal{M}$ it is easy to see that $V(\mathcal{T}_{\mathcal{M}}, \pi) \geq V(\mathcal{S}_{\mathcal{M}}, \pi)$ because the two trees have the same revenue, $\pi|\mathcal{M}|$, but by definition $\mathcal{T}_{\mathcal{M}}$ is a minimum spanning tree, so its cost is less than or equal to $\mathcal{S}_{\mathcal{M}}$. To the extent that pre-existing infrastructure routes (i.e., roads) or geographic/legal constraints (e.g., mountains, zoning) limit the opportunity to profitably deviate from using a subtree of $\mathcal{T}$, this lower-bound will be tighter.

Next, we outline how to compute $\mathbb{E}[\bar{V}(\mathcal{T}, \pi)]$ by decomposing the tree into 1D lattices. Define $N(i)$ as the set of nodes that are downstream from node i; more precisely, downstream nodes refers to nodes for which the shortest path to the source passes through i. Likewise, define $N(I)$ as the set of nodes that are downstream from I, where I represents any chain of nodes in $\mathcal{T}$ satisfying the following conditions: (i) all nodes that are not the first or last node have degree 2; (ii) the first node is the last node if it has degree ≥ 3; (iii) the first node is the last node if it is the source node and has degree ≥ 2; (iv) unless the last node is a terminal node of degree 1 or the source node with degree ≥ 2, it must have degree ≥ 3. Let's denote with A the first such chain emanating from the source node, and denote the rest of the tree as A^c (the complement of A), and the disjoint branches of A^c as $B, C, ...$ Without loss of generality, suppose that the final node of A has degree 3 and that A^c comprises of just two branches, B and C, which themselves may split further. Let α denote the drift in profit (e.g., $F^c(\pi)\pi - c$). We can now begin writing the expected value of offering price π to the potential users in $\mathcal{T}$ using the law of total expectation over two events: (i) when the random realization of reference prices produces at least one

price-taker in A^c, denoted as $P(A^c)$, and (ii) when it produces no price-takers in A^c.

$$\begin{aligned}
E\left[\bar{V}(\mathcal{T})\right] &= E\left[\bar{V}(\mathcal{T})|A^c \cap \mathcal{M} = \emptyset\right](1 - P(A^c)) + E\left[\bar{V}(\mathcal{T})|A^c \cap \mathcal{M} \neq \emptyset\right] P(A^c) \\
&= E\left[\bar{V}(A)\right](1 - P(A^c)) + \left[E\left[V(A)\right] + E\left[\bar{V}(A^c)|A^c \cap \mathcal{M} \neq \emptyset\right]\right] P(A^c) \\
&= \mathbb{E}[\bar{V}_{|A|}]F(\pi)^{|N(A)|} + \alpha|A|(1 - F(\pi)^{|N(A)|}) + \frac{\mathbb{E}[\bar{V}(A^c)]}{P(A^c)}P(A^c) \\
&= \mathbb{E}[\bar{V}_{|A|}]F(\pi)^{|N(A)|} + \alpha|A|(1 - F(\pi)^{|N(A)|}) + \mathbb{E}[\bar{V}(A^c)].
\end{aligned}$$

So far, we can compute the first two terms. For the last term, we can decompose it into $\mathbb{E}[\bar{V}(A^c, \pi)] = \mathbb{E}[\bar{V}(B, \pi)] + \mathbb{E}[\bar{V}(C, \pi)]$, and each of these expectations can be decomposed in the same manner as $E\left[\bar{V}(\mathcal{T}, \pi)\right]$. We can iterate this decomposition process until we reach all the terminal branches. To clarify, a terminal branch is a terminal node plus all nodes along its path to the source before reaching a node of degree ≥ 3. The expected value of terminal branches can be computed as 1D lattices without further decomposition. Letting $\mathcal{TB}$ denote the set of terminal branches, the equation for expected profit when constrained to building subtrees of $\mathcal{T}$ becomes

$$E\left[\bar{V}(\mathcal{T}, \pi)\right] = \sum_{I \in \mathcal{T} \setminus \mathcal{TB}} \left[E\left[\bar{V}_{|I|}(\pi)\right] F(\pi)^{|N(I)|} + \alpha|I|(1 - F(\pi)^{|N(I)|}\right] + \sum_{I \in \mathcal{TB}} E\left[\bar{V}_{|I|}(\pi)\right]$$

4.9 Conclusion

The main goal of this chapter was to identify high-level mathematical properties of the network monopolist's decision problem that are relevant to strategic business decisions and regulation. We established that it can be profitable for the monopolist to offer to build a network even in the presence of negative drift, more so if the monopolist is allowed to choose the optimal number of potential users to receive offers, and especially so if the monopolist can serve only the optimal subset amongst those who accept the offer. Serving an optimal subset allow expected profit to exist in the face of much greater levels of negative drift. It also leads to profit levels that are orders of magnitude higher than when all price-takers must be served. As a result, regulators who wish to see private sector development of microgrids in rural

and remote areas with low density should view this as a strong argument in favor of allowing monopolists to choose to serve only a subset of price-takers.

When the cost and reward levels produce a very slight negative profit drift, then the profit maximizing number of potential users to receive offers can be far fewer than the break-even number, and the break-even number will produce much higher total utility as a result of the greater consumer surplus. In this setting, it may make sense for the infrastructure monopoly to be a public utility, or for public intervention to incentivize a private monopolist to make offers to the break-even number of users.

For the most general 1D model with continuous random costs and rewards, we derived formulas that can be used to estimate the value of acquiring information about the terrain or customer demand to reduce uncertainty about profitability. We believe that the model is sufficiently general that it could serve as a rule-of-thumb for deciding how much to invest in market research and planning when the profitability of a long 1D segment of network infrastructure is uncertain.

The quality of the lower bound on 2D trees presented in Section 4.8 remains to be tested empirically. To test the bound, as well as the multi-time period policies in Chapter 3, we have build a 2D location simulator, and have also acquired building location data for two dozen rural towns and villages in several low-income countries from www.openstreetmap.org. Simulation analysis using these datasets will be undertaken as future work to examine the strengths and weaknesses of our methods under a wide range of 2D location distributions and demand distributions.

Chapter 5

Strategic Procurement of Counterfeits by Retailers

5.1 Introduction

Counterfeiting is a global problem with enormous consequences for businesses, consumer safety, and public health. The global counterfeit market is estimated to have been $1.2 trillion USD in 2017 and is growing [Research and Markets, 2017]. The danger posed by counterfeit pharmaceutical products and medical devices is especially acute, as consumers' lives can be placed in jeopardy, and drug-resistant strains can emerge and be transmitted to others when patients unknowingly take diluted antibiotics. The global annual death toll from counterfeit malaria and tuberculosis drugs alone is estimated to be as high as 700,000 [Karunamoorthi, 2014, IPN, 2011]. The World Health Organization estimates that roughly 10% of medicines sold in developing countries are fake or of poor quality [World Health Organization, 2017], though the percentages for certain products, antibiotics and anti-malaria drugs in particular, is believed to exceed 50% in certain countries in sub-Saharan Africa and Southeast Asia [Ambroise-Thomas, 2012].

Firms and governments have a range of potential tools to combat counterfeiting, such as track-and-trace technologies, single-use product authentication codes, expensive and elaborate packaging, heightened manufacturer regulation, investments in law

enforcement, and consumer awareness campaigns to encourage increased consumer scrutiny and reporting. The potential impact of these measures is hard to estimate *a priori*, the impact may depend greatly on various market-specific and product-specific circumstances that cannot be experimentally manipulated, and the upfront costs can be high enough to discourage policy experimentation. In the absence of hard empirical evidence, mathematical models may offer insight and intuition that could help inform anti-counterfeiting policy development for firms and governments with respect to at least some of these possible countermeasures.

This chapter places particular emphasis on analyzing the profit-maximizing decisions of hypothetical retailers who knowingly procure and sell counterfeits as some fraction of their inventory. We develop three game theoretic models of varying complexity to examine this problem. In the first two models, retailers face a known demand function that is increasing in the proportion of their inventory that is authentic; the major results for these models are given mathematically as a function of the demand curve's parameters. In the third model, we consider the possibility of consumers having the option to use a product verification technology, such as a single-use scratch code texted to a company database, and companies being able to reward checking behavior. The third model involves more complex demand functions, requiring numerical computation of equilibrium counterfeiting levels on specific simulated instances rather than closed-form solutions.

5.2 Literature Review

The game-theoretic literature on counterfeiting has tended to focus on decisions by the manufacturer with respect to quality, price, or branded packaging effort, given anticipated responses by counterfeiters with respect to quality and/or effort to conceal the counterfeits, as well as the probability of counterfeit detection by consumers of varying vigilance [Cho et al., 2015, Quercioli and Smith, 2015]. Some attention has been given to the potential behavior of retailers in the counterfeiting supply chain, such as by cooperating with counterfeiters to introduce fake products into official

supply chains [Green and Smith, 2002, Liu et al., 2005], and by selectively selling counterfeits to certain customer types [Fitzpatrick, 2015].

The paper analyzing a setting most similar to the one considered in this chapter is by Cho et al. Cho et al. [2015]. The authors consider a distribution of proactive and non-proactive consumers (those who check for and may successfully detect deceptive counterfeits vs. those who do not), and the possibility that retailers may strategically vary the percentage of counterfeits that they sell in response to consumer behavior. Their paper differs from this chapter in several respects: first, they examine both deceptive and non-deceptive counterfeits (e.g., flea-market designer handbags), whereas our analysis is limited to deceptive counterfeits; second, they treat product quality as a key decision variable for the authentic manufacturer, whereas we consider quality to be exogenous; third, their analysis relies on a very particular demand function, whereas we aim to provide more general results for a range of plausible demand function shapes. Our focus on deceptive counterfeits and fixed-quality products is appropriate for the pharmaceutical sector.

The analysis by Cho et al. suggested that lowering price or raising quality may be effective strategies against non-deceptive counterfeits (or at least an improvement over the profit-maximizing prize and quality level in the absence of counterfeits), but are likely to be ineffective against deceptive counterfeits. We therefore aim to provide deeper analysis of deceptive counterfeits. Their paper uses a particular model of consumer demand called the vertical differentiation model, which is also used in other recent counterfeiting studies [Qian, 2008, Zhang et al., 2012]. The vertical differentiation model defines consumer utility as $X_i z_i - p$, where $X_i \sim U[0,1]$ is a parameter for consumer i's willingness to pay for product quality, z_i parameterizes consumer i's perception of product quality, and p is the retail price. This model produces the demand function $D(p, z_i) = \max(0, 1 - \frac{p}{z_i})$, which is linear decreasing in p up to $p = z_i$, and for any fixed p is concavely increasing towards an asymptote of 1 over $z_i \in [p, \infty)$.

The assumption that $X_i \sim U[0,1]$ is fairly strong and unlikely to generalize. Although that assumption makes the model easier to analyze with respect to price,

the model is not as mathematically friendly when quality becomes the key variable of interest, especially since perceived quality is a latent cognitive variable whose quantitative relationship with actual counterfeiting levels, the quality of counterfeits, and consumer awareness of counterfeits will be hard to characterize. To overcome this limitation, we develop two alternative models in which demand is an increasing function of the percentage of product that is authentic rather than as a function of quality perception. We define the function by parameters that allow its shape to vary much more widely than the vertical differentiation model, so as to encompass a wide range of possible relationships between counterfeit prevalence and demand. Despite the greater flexibility in the shape of demand, the functional form we deploy leads to closed-form solutions for the strategic decisions of retailers under a range of interesting scenarios.

5.3 Single Retailer, Demand Functions Without Inflection Points

In this section we model aggregate demand with respect to the percentage of products that are real (denoted by θ) with the following functional form:

$$D(\theta) = [\max(0, d(m\theta + b))]^c, \text{ with } m, d > 0, d = (m + b)^{-c}.$$

This family of functions includes functions that are strictly convex ($c > 1$), strictly concave ($0 < c < 1$), and linear ($c = 1$). No particular parameter value need be assumed; the approach in this section is to posit that demand falls as the percentage of authentic product decreases. Even with deceptive counterfeits, this could happen when repeat customers discover counterfeits and switch retailers, through reputation damage, or from customers inferring that the authentic product is of poor quality and reducing their demand. We do not assume a specific mechanism or timeline for consumers learning or reacting to θ, nor is a particular consumer utility model or θ

detection model required. There are many reasonable possibilities with respect to these mechanisms, but we focus on using a family of demand functions that could approximate a wide range of aggregate behavior from such mechanisms rather than speculate on or model these types of mechanisms directly.

Sigmoidal demand, which some may consider to be more realistic, is considered in Section 5.5. The functions can be scaled by $d = (m + b)^{-c}$ to normalize such that $D(\theta = 1) = 1$, representing 100% of the possible demand at the prevailing market-price level; alternatively (and equivalently), we can require that $m+b = 1$ and simply let d=1, which also enforces $D(\theta = 1) = 1$:

$$D(\theta) = [\max(0, m\theta + b)]^c, \text{ with } m > 0 \text{ and } m + b = 1. \tag{5.1}$$

To reduce the number of parameters referenced we will use the formulation without d. The parameter m can be thought of as a slope-like parameter, and $\max(0, b)^c$ as the intercept (i.e., demand when everything is counterfeit). Since $\{\theta | m\theta + b < 0\}$ cannot be optimal, we can limit the domain to $\theta \in [\max(0, \frac{-b}{m}), 1]$ and write $D(\theta) = (m\theta+b)^c$ over this restricted domain. As shown in Figure 5-1, this family of functions include the positive components of all increasing lines passing through (1,1), as well as all transformations of these line components by a positive exponent to create concavity or convexity.

5.3.1 Retailer Profit Maximization

Instead of parameterizing the retail and wholesale prices, we will instead parameterize the retailer's per-unit profit margin for real products as π_r and for fake products as π_f, where $\pi_f > \pi_r$. If monetary units are scaled such that $\pi_r = 1$, then $D(1)\pi_r = 1$ represents the retailer profit, Π_r, when only real goods are sold. The models developed in this section assume that π_f is constant over all θ, whereas in reality lower levels of θ may bring about regulatory or law enforcement risks to the retailer. Since these risks may vary arbitrarily across jurisdictions, and may be very low in some countries, the models presented in this section omit these risks from the retailer's profit function

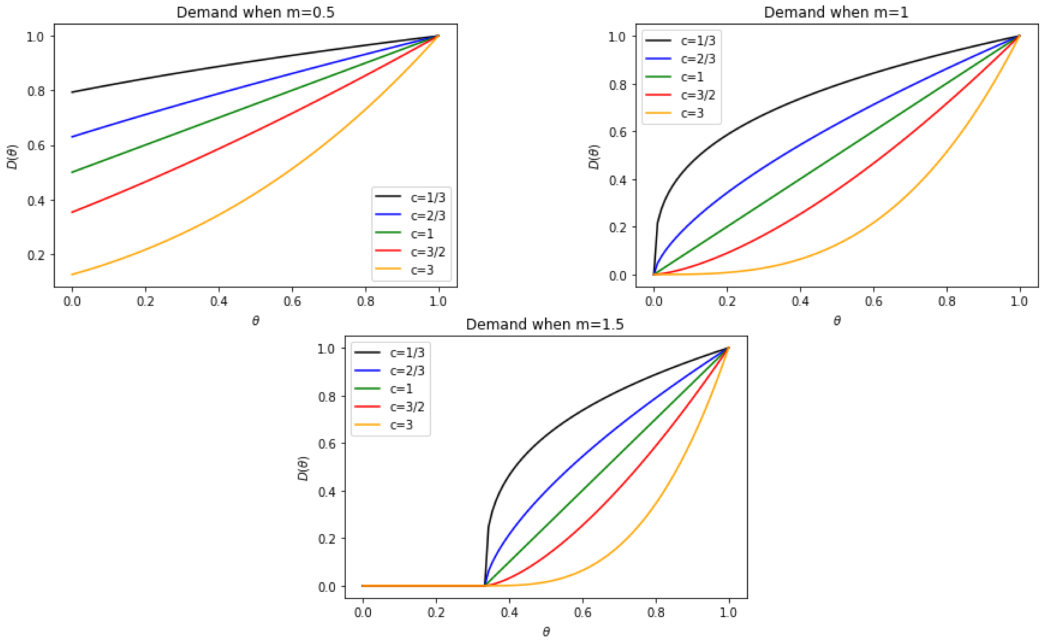

Figure 5-1: Examples to illustrate the flexibility of the family of functions for modeling demand vs percent authentic.

and simply focus on the retailer's hypothetical decision in the absence of penalties from law enforcement. As a result, the equilibrium θ^* described in the following pages can be viewed as a lower bound on the θ that would be obtained if the retailer faces convex costs increasing in the proportion of counterfeits sold.

The retailer's profit function is concave over all θ when $c \leq 1$. The retailer's profit-maximizing θ^* can therefore be set to the nearest boundary if the unconstrained first-order condition's (FOC) solution to the retailer's profit problem falls outside this domain when $c \leq 1$. When $c > 1$, concavity of profit is not guaranteed over the unconstrained domain and the FOC result should be used only if it is above the lower bound $\theta \geq \max(0, \frac{-b}{m})$.

The profit-maximizing percentage of authentic product for the retail to procure can be found by taking the derivative with respect to θ of retailer profit Π_r:

$$\Pi_r = D(\theta)[\pi_r\theta + \pi_f(1-\theta)].$$

$$\Rightarrow \frac{d\Pi_r}{d\theta} = D'(\theta)[\pi_r\theta + \pi_f(1-\theta)] + D(\theta)[\pi_r - \pi_f]$$

Evaluate for the particular form of $D(\theta)$ over allowable domain and solve FOC:

$$cm(m\theta+b)^{c-1}\theta(\pi_r-\pi_f) + cm(m\theta+b)^{c-1}\pi_f + (m\theta+b)^c(\pi_r-\pi_f) = 0$$

Simplify by dividing by $(m\theta+b)^{c-1}$

$$\Rightarrow cm(\pi_r-\pi_f)\theta + cm\pi_f + (m\theta+b)(\pi_r-\pi_f) = 0$$

$$\Rightarrow m(\pi_r-\pi_f)(c+1)\theta = -b(\pi_r-\pi_f) - cm\pi_f$$

$$\Rightarrow \theta^* = \min\left(1, \max\left(0, \frac{-b}{m}, \frac{-b(\pi_r-\pi_f) - cm\pi_f}{m(c+1)(\pi_r-\pi_f)}\right)\right)$$

$$\Rightarrow \theta^* = \min\left(1, \max\left(0, \frac{-b}{m}, \frac{c\pi_f}{(c+1)(\pi_f-\pi_r)} - \frac{b}{m(c+1)}\right)\right)$$

$$\Rightarrow \theta^* = \min\left(1, \max\left(0, \frac{m-1}{m}, \frac{c\pi_f}{(c+1)(\pi_f-\pi_r)} - \frac{(1-m)}{m(c+1)}\right)\right) \tag{5.2}$$

The optimal authenticity level θ^* is therefore unique, and it is straightforward to show that any interior solution θ^* is decreasing in b, increasing in m, increasing as π_r approaches π_f from below, increasing in c, and decreasing in π_f. These directional results are true for all functions permitted by the functional form of demand described above, and do not require convexity or concavity. We can extract some useful insights from the partial derivatives. The dependence of θ^* on c, π_f and m is illustrated in the collection of heatmaps presented in Figure 5-2, with redder areas indicating higher counterfeiting.

In all heatmaps, the darkest blue area indicates parameter combinations for which it is most profitable to sell only authentic products, whereas the darkest red represents combinations for which only counterfeits are sold. For $m = 0.333$ and $m = 0.667$ and $c \leq 1$ we can see the presence of fairly rapid phase transition between all-real to all

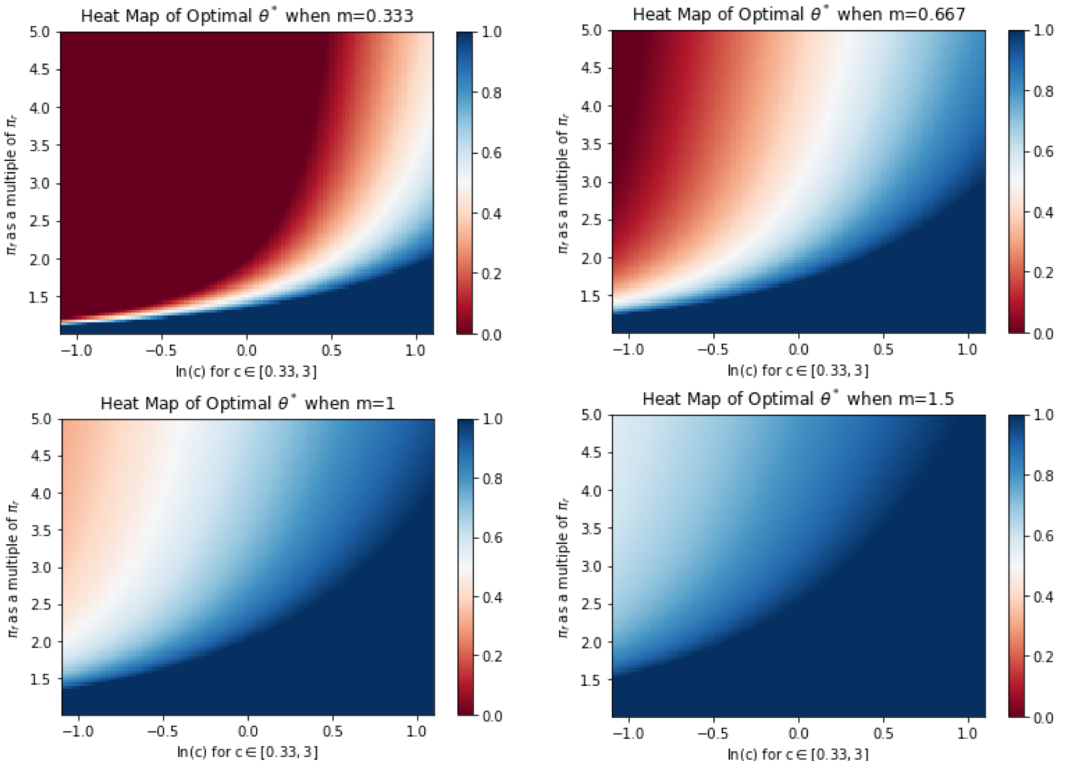

Figure 5-2: Heat maps of the profit-maximizing percentage of real product with different values of m. The vertical axis is the profit margin on fake goods as a multiple of the profit margin on real goods. The horizontal axis is the convexity parameter c on a natural log scale.

or mostly fake products as π_f increases. This result is of practical significance. It suggests that if consumer demand is not especially responsive to increased counterfeit prevalence, and if raising consumer awareness or aversion to fakes is expensive, then hurting the profit margins on counterfeits by attacking the counterfeit supply chain or driving up law enforcement penalties may have a huge effect if π_f is close to the phase transition threshold. By contrast, if π_f is far from the phase transition threshold then law enforcement efforts to drive down π_f may fail to reduce counterfeits, and public efforts may be better spent on campaigns to shift the shape of the demand curve (i.e., higher m and/or c).

The partial derivatives below hold only if $\theta^* \in (\frac{m-1}{m}, 1)$; otherwise, the optimal solution is at the boundary and we have either full demand or no demand, and the partial derivatives would instead be zero (unless the unconstrained optimal solution happens to fall exactly on the boundary).

$$\frac{\partial \theta^*}{\partial m} = \frac{1}{m^2(c+1)} \tag{5.3}$$

$$\frac{\partial \theta^*}{\partial b} = \frac{-1}{(1-b)^2(c+1)} \tag{5.4}$$

$$\frac{\partial \theta^*}{\partial \pi_r} = \frac{c}{(c+1)} \frac{\pi_f}{(\pi_f - \pi_r)^2} \tag{5.5}$$

$$\frac{\partial \theta^*}{\partial \pi_f} = \frac{c}{(c+1)} \frac{-\pi_r}{(\pi_f - \pi_r)^2} \tag{5.6}$$

$$\frac{\partial \theta^*}{\partial c} = \frac{1}{(c+1)^2} \left[\frac{\pi_f}{\pi_f - \pi_r} + \frac{b}{m} \right] \tag{5.7}$$

The partial derivatives of an interior solution θ^* with respect to the quasi-slope and quasi-intercept parameters, m and b, do not depend on the profit margins. Likewise, the impact of changes in π_r and π_f on an interior solution θ^* depends only on the difference between the two profit margins and the convexity/concavity parameter c, and not on m or b, though m and b can affect whether an interior solution is realized for given values of π_r, π_f and c. The more concave the demand function, the less sensitive it is to changes in the profit margin spread. More importantly, because the

square of the difference in profit margins appears in the denominator, an interior solution θ^* will increase at an increasingly sharp rate as π_f approaches π_r, whereas a similar magnitude change in π_f will have far less impact on an interior solution θ^* when the profit margin spread is large.

5.3.2 Minimium Authenticity Level

Under this model, the minimum percentage of authentic product that a retailer should sell, no matter how profitable it is to sell counterfeits, can be found by evaluating $\lim_{\pi_f \to \infty} \theta^*(\pi_f)$, which gives

$$\inf_{\pi_f} \theta^*(\pi_f) = \frac{c}{c+1} - \frac{b}{m(c+1)} = \frac{c}{c+1} - \frac{(1-m)}{m(c+1)}. \tag{5.8}$$

The intuition for the existence of such a quantity is that even when the profit on fakes is enormous, the retailer may have to still sell some authentic product to induce the consumer to take the risk of buying a product that could be fake. If the quasi-intercept parameter b is positive, then such a demand function would imply that some demand will remain even if all products are counterfeit. In such a situation, if b is too large then the minimum θ^* can be 0. More precisely, the minimum θ^* is strictly positive as long as any of the following (equivalent) conditions hold: $c > \frac{b}{1-b} \iff c > \frac{1-m}{m} \iff b < \frac{c}{c+1} \iff m < \frac{1}{c+1}$. Figure 5-3 displays through a heatmap the values taken by $\inf_{\pi_f} \theta^*(\pi_f)$ over $(c, m) \in [0.2, 5]^2$.

5.3.3 Minimum Profit Margin for Counterfeiting

We can find the threshold $\bar{\pi}_f$ at or below which no counterfeiting should occur by setting the equation for θ^* greater than or equal to 1 and isolating π_f on one side of the inequality. For ease of interpretation we can let $\pi_r = 1$ so that π_f's units are in multiples of π_r.

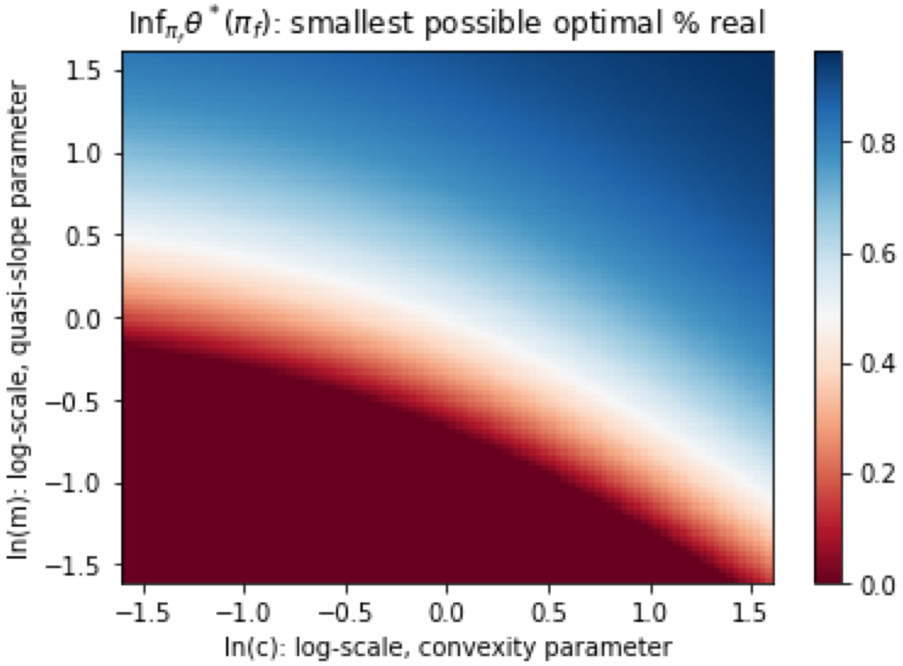

Figure 5-3: Minimum % real as a function of m and c, which are plotted on a logarithmic scale.

$$\theta^*(\pi_f) \geq 1, \ \forall \ \pi_f \leq cm\pi_r + \pi_r \tag{5.9}$$

From Equation 5.9 and Figure 5-4 we see that demand curves that are convex and/or have a shallower slope parameter tend to enable counterfeiting at lower values of π_f, whereas a steeper slope parameter and increased convexity multiplicatively interact to produce a higher required profit margin for counterfeits. When data on counterfeit prevalence do not exist for particular products or markets, estimates on the profitability of fakes and the likely shape of the demand curve could provide manufacturers and law enforcement with preliminary guidance on how much that market is at risk for being targeted by counterfeiters. It would be interesting to see whether the prevalence of deceptive counterfeits in the real world is consistent with this model, or if it is simply an artefact of idiosyncratic or overly stylized mathematical assumptions. We are hopeful that it is the former, and that the result has some

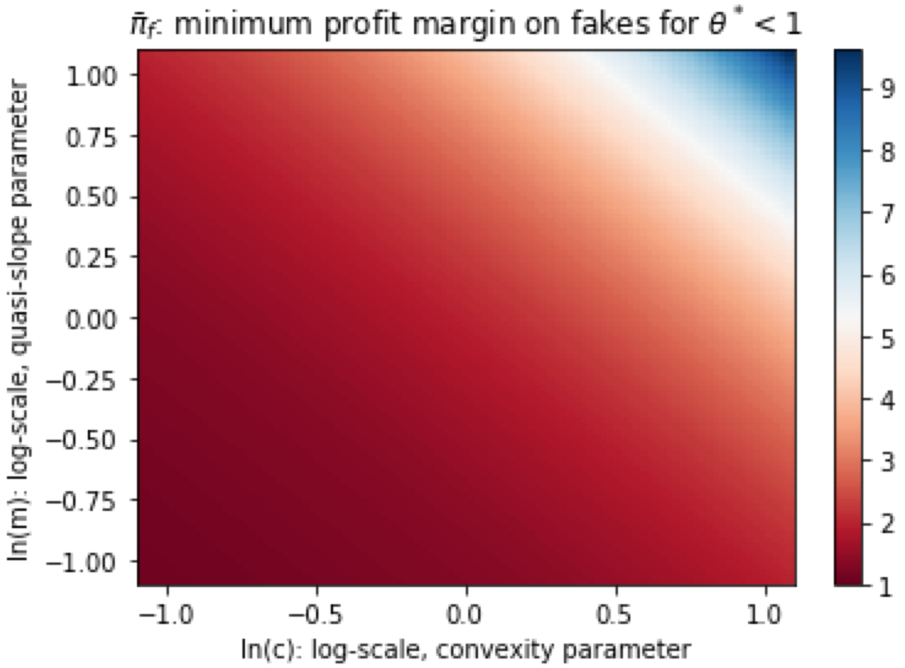

Figure 5-4: Minimum profit margin for fakes, at or below which it is not profitable for the retailer to procure counterfeits.

intuitive appeal. For example, it would make sense that for high-end jewelers and art auctioneers the discovery of even a small proportion of counterfeits could have disastrous consequences for future demand (high m and c). That steep fall in demand could be sufficient to deter strategic procurement of counterfeits even if the profit margin on fakes is multiples higher than on an authentic piece.

5.4 Multiple Retailers, Shared Reputation

So far we have characterized the profit-maximizing counterfeit procurement for a single retailer and examined several implications for law enforcement and manufacturers. Fortunately, the mathematical structure of the solution is amenable to deeper game-theoretic analysis. Of particular interest, especially to franchises or firms with decentralized procurement, is the question of how a shared reputation may affect counterfeit procurement levels. We examine several types of shared reputation and find that a shared reputation is subject to a free-rider problem, leading to higher levels

of counterfeiting than when each retailer's demand function is solely a function of its own decisions. Suppose n retailers share a reputation; that is, suppose $\theta = \sum_i^n \lambda_i \theta_i$, where $\sum_i^n \lambda_i = 1, \lambda_i \geq 0$. This situation could represent, for example, a franchise that allows franchisees some discretion over procurement.

5.4.1 One Strategic and One Myopic Retailer

Let's first consider the case of $n = 2$, so $\theta = \lambda_1\theta_1 + \lambda_2\theta_2 = \lambda_1\theta_1 + (1 - \lambda_1)\theta_2$. This means the retailers face the following demand functions:

$$D_1(\theta_1) = [\max(0, (m(\lambda_1\theta_1 + \lambda_2\theta_2) + b))]^c \tag{5.10}$$

$$D_2(\theta_2) = [\max(0, (m(\lambda_1\theta_1 + \lambda_2\theta_2) + b))]^c. \tag{5.11}$$

We can view these functions as being the same as the original demand function, but with a parameter transformation of m and b such that $m_1 = m\lambda_1$, $m_2 = m(1-\lambda_1)$, $b_1 = m(1 - \lambda_1)\theta_2 + b$, and $b_2 = m\lambda_1\theta_1 + b$.

Observing that $m_1, m_2 < m$ and $b_1, b_2 > b$, one can conclude that a shared reputation will cause equilibrium authenticity levels to fall. We will first examine the case where only one retailer realizes that there is a shared reputation. Later, we will assume that both retailers are aware of the shared reputation. Supposing that Retailer 1's strategy is myopic (non-strategic), we have $\theta_1^* = \theta^*$. Retailer 2's best response is

$$\theta_2^* = \min\left[1, \max\left[0, \frac{c\pi_f}{(c+1)(\pi_f - \pi_r)} - \frac{b + m\lambda_1\theta_1^*}{(1-\lambda_1)m(c+1)}\right]\right], \tag{5.12}$$

and the partial derivative of $\theta_2^* \in (0, 1)$ with respect to λ_1 (or λ_2) is as follows:

$$\frac{\partial\theta_2^*}{\partial\lambda_2} = \frac{m(\theta_1^* - 1) + 1}{\lambda_2^2 m(c+1)} \iff \frac{\partial\theta_2^*}{\partial\lambda_1} = -\frac{m(\theta_1^* - 1) + 1}{(1-\lambda_1)^2 m(c+1)}. \tag{5.13}$$

It is straightforward to show that $\frac{\partial \theta_2^*}{\partial \lambda_2}$ is non-negative for all positive values of m, as expected. The rate of the fall in the optimal authenticity level gets increasingly severe as the retailer's influence on the overall θ falls. When we set Retailer 2's best response to myopic Retailer 1 to be 0, we can establish an inequality that gives the threshold value of λ_1 (equivalently, the minimum value of λ_2) above which $\theta_2^* = 0$. That is, if the myopic retailer is responsible for too much of the shared reputation, the strategic retailer may end up only procuring counterfeits:

$$\text{Solve for } \lambda_1 : \theta_2^* = \frac{c\pi_f}{(c+1)(\pi_f - \pi_r)} - \frac{b + m\lambda_1\theta_1^*}{(1-\lambda_1)m(c+1)} \leq 0 \quad (5.14)$$

$$\Rightarrow \text{Retailer 2 all-counterfeit threshold} : \lambda_1 \geq \frac{c\pi_f m - (1-m)(\pi_f - \pi_r)}{c\pi_f m + m\theta_1^*(\pi_f - \pi_r)}. \quad (5.15)$$

For all allowable values of m, π_f, π_r and c this threshold is, indeed less than 1. Note that the θ_1^* also depends on these parameters and would need to be recomputed whenever the threshold is calculated using new parameters. The shape of the strategic retailer's best response, as a function of the myopic retailer's contribution to the strategic retailer's reputation, is shown in Figure 5-5 for two representative scenarios.

5.4.2 Two Strategic Retailers, Shared Reputation

Next, consider the case where two identical retailers are aware that they share a reputation according to $\theta = \lambda_1\theta_1 + \lambda_2\theta_2$, where $\lambda_1 + \lambda_2 = 1$ and $\lambda_i \geq 0$. What is the Nash Equilibrium? Let $BR_1(\theta_2)$ denote the best response of Retailer 1 to the choice of Retailer 2, and vice-versa for $BR_2(\theta_1)$. The best response functions follows from Equation 5.12:

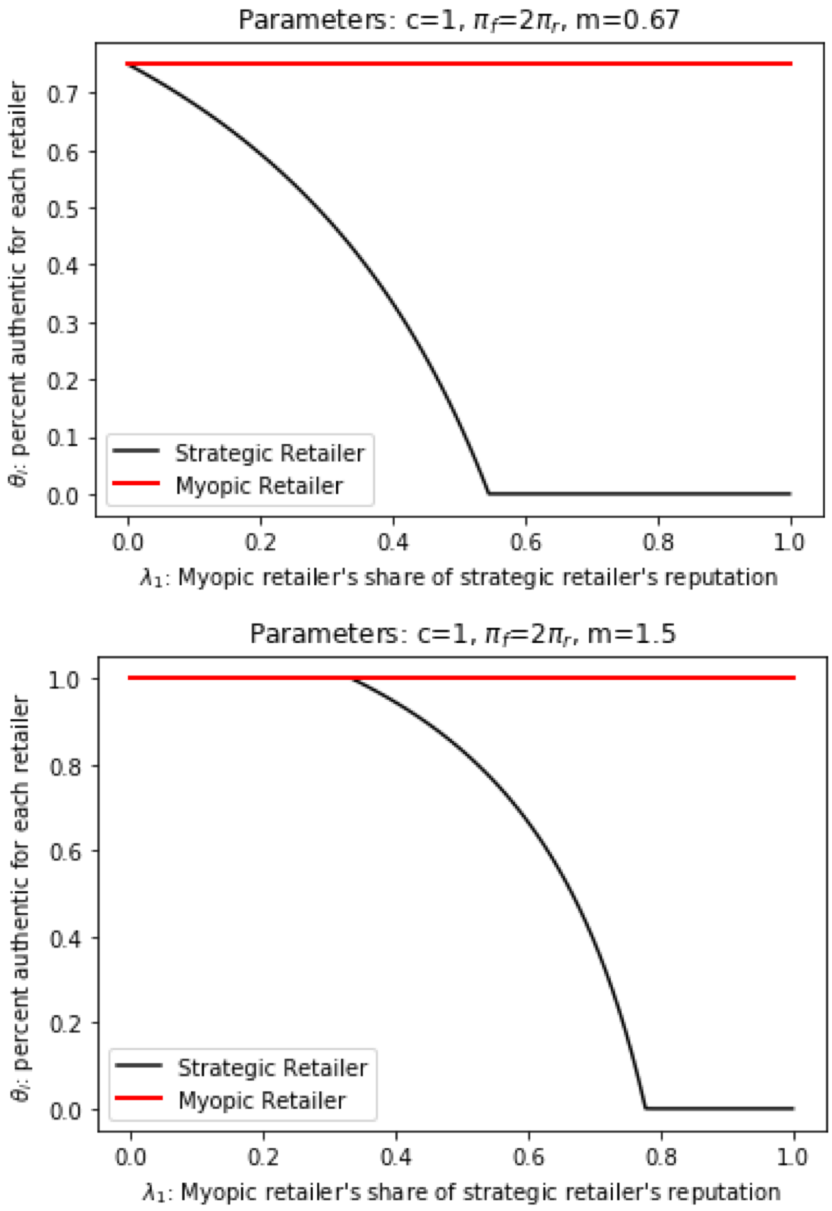

Figure 5-5: The equilibrium θ of a strategic retailer and a myopic retailer as a function of how strongly the myopic retailer's decision θ affects the strategic retailer's reputation.

$$BR_1(\theta_2) = \min\left[1, \max\left[0, \frac{c\pi_f}{(c+1)(\pi_f - \pi_r)} - \frac{b + m(1-\lambda_1)\theta_2}{\lambda_1 m(c+1)}\right]\right] \tag{5.16}$$

$$BR_2(\theta_1) = \min\left[1, \max\left[0, \frac{c\pi_f}{(c+1)(\pi_f - \pi_r)} - \frac{b + m\lambda_1\theta_1}{(1-\lambda_1) m(c+1)}\right]\right] \tag{5.17}$$

When $BR_{-i}(\theta_i) \in (0,1)$, the best response function is linear with respect to the opponent's strategy; otherwise, it is rounded to the boundary. Figure 5-6 illustrates the best-response functions in two settings with linear demand and equally shared reputations. The Nash Equilibrium is the point of intersection, which, in the two examples showin in Figure 5-6, involves more counterfeiting than in the individual reputation scenario.

Denote the slope of the line for the interior solution, if applicable, of these best response functions as α_i (note that it is not positive) and the hypothetical intercept of that line as β_i:

$$\alpha_i = -\frac{(1-\lambda_i)}{\lambda_i(c+1)},$$
$$\beta_i = \frac{c\pi_f}{(\pi_f - \pi_r)(c+1)} - \frac{b}{\lambda_i m(c+1)}.$$

Using the symmetry between the retailers, we can now solve for the Nash Equilibrium by solving $\theta^*_{i,NE} = BR_i(\theta^*_{-i,NE}) = BR_i(BR_{-i}(\theta^*_{i,NE}))$:

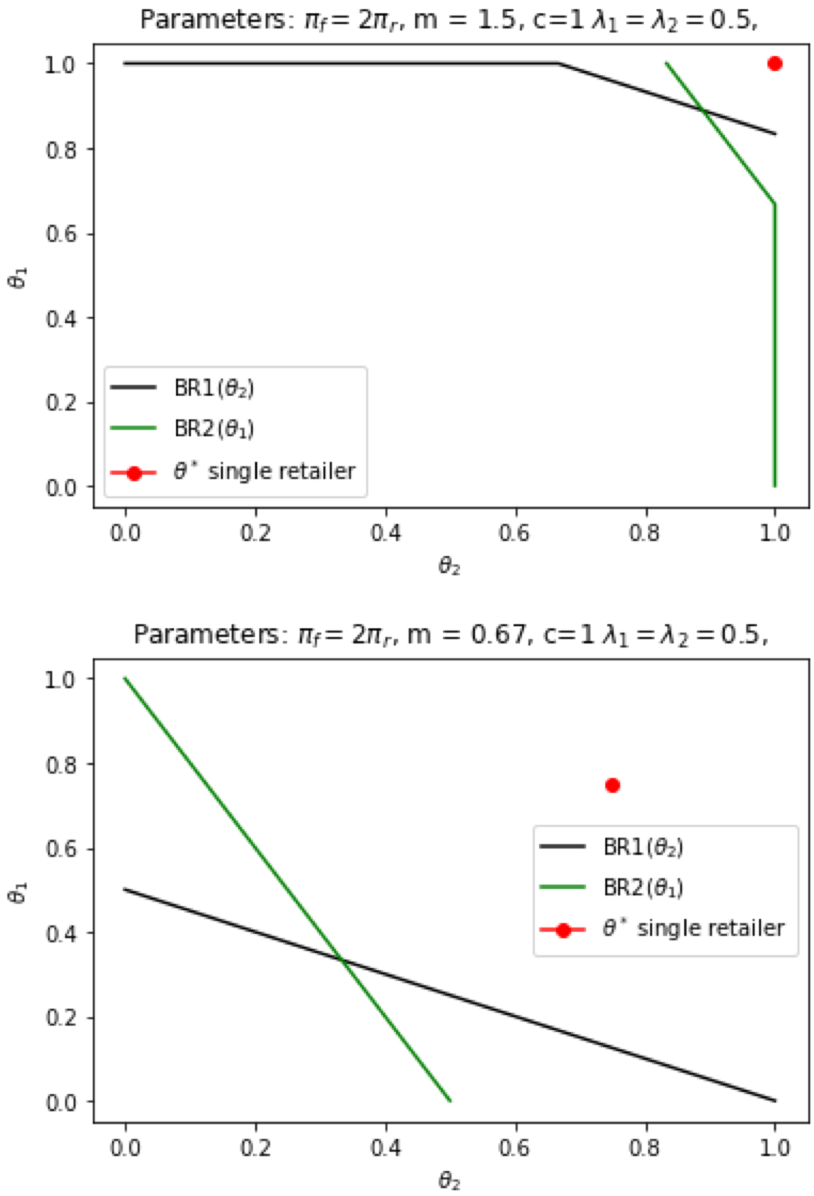

Figure 5-6: Best response functions for two retailers whose decisions equally contribute to a common reputation.

$$\text{for both retailers, } BR_i(\theta^*_{-i}) = \min\left[1, \max\left[0, \alpha_i\theta^*_{-i} + \beta_i\right]\right] \tag{5.18}$$

$$\Rightarrow \theta^*_{i,NE} = \min\left[1, \max\left[0, \alpha_i \min\left[1, \max\left[0, \alpha_{-i}\theta^*_{i,NE} + \beta_{-i}\right]\right] + \beta_i\right]\right] \tag{5.19}$$

$$\Rightarrow \theta^*_{i,NE} = \begin{cases} \min\left[1, \max\left[0, \frac{\alpha_i\beta_{-i}+\beta_i}{1-\alpha_i\alpha_{-i}}\right]\right] & \text{if } \theta^*_{i,NE} \in \left(\frac{1-\beta_{-i}}{\alpha_{-i}}, \frac{-\beta_{-i}}{\alpha_{-i}}\right) \\ \min\left[1, \max\left[0, \beta_i\right]\right] & \text{if } \theta^*_{i,NE} \geq \frac{-\beta_{-i}}{\alpha_{-i}} \\ \min\left[1, \max\left[0, \alpha_i + \beta_i\right]\right] & \text{if } \theta^*_{i,NE} \leq \frac{1-\beta_{-i}}{\alpha_{-i}} \end{cases} \tag{5.20}$$

Although the expressions are cumbersome, we present the expanded forms for the three cases ($\theta^*_{i,NE}$ interior to $\left(\frac{-\beta_{-i}}{\alpha_{-i}}, \frac{1-\beta_{-i}}{\alpha_{-i}}\right)$, above the interior, and below the interior) for the sake of completeness, though readers interested in replication may find it easier and less error-prone to numerically find the fixed point solving Equation 5.19 for plotting and computational analysis.

$$\theta^*_{i,NE} = \begin{cases} \min\left[1, \max\left[0, \frac{\frac{(1-\lambda_i)}{\lambda_i(c+1)^2}\left(\frac{b}{\lambda_{-i}m} - \frac{c\pi_f}{(\pi_f-\pi_r)}\right) + \frac{c\pi_f}{(\pi_f-\pi_r)(c+1)} - \frac{b}{\lambda_i m(c+1)}}{1-\left(\frac{(1-\lambda_i)}{\lambda_i(c+1)}\right)\left(\frac{(1-\lambda_{-i})}{\lambda_{-i}(c+1)}\right)}\right]\right] \\ \min\left[1, \max\left[0, \frac{c\pi_f}{(\pi_f-\pi_r)(c+1)} - \frac{b}{\lambda_i m(c+1)}\right]\right] \\ \min\left[1, \max\left[0, -\frac{(1-\lambda_i)}{\lambda_i(c+1)} + \frac{c\pi_f}{(\pi_f-\pi_r)(c+1)} - \frac{b}{\lambda_i m(c+1)}\right]\right] \end{cases} \tag{5.21}$$

As shown in the second plot, the increase in the equilibrium proportion of counterfeits can be dramatic when reputations are shared. The common reputation need not be shared equally. Figure 5-7 shows, over four scenarios, the Nash Equilibrium decisions $\theta^*_{i,NE}$ as Retailer 1's proportion of the shared reputation varies from 0 to 1; implicitly, Retailer 2's proportion varies from 1 to 0. For reference, the optimum real goods procurement level for a single retailer (i.e., retailers with independent rep-

utations) is plotted in red, and is greater than or equal to the market equilibrium in all cases, as expected. The average level of real goods for the market is plotted in green when each retailer's market share is made to equal its influence on the shared reputation, and in blue when the retailers' market shares are each 50% despite the unequal share of reputation. In the latter case, the market prevalence of counterfeits is strictly higher than when each retailer's reputation is independent.

The most important practical insight from Figure 5-7 is the extent to which a retailer's procurement of counterfeits can surge when their share of reputation begins to fall from 0.5, even as the market counterfeit level decreases as a result of the better procurement by the main reputation bearer. For a franchise or retail chain, this would suggest that procurement managers at flagship stores will have less of an incentive to procure counterfeits than secondary stores, and also suggests that variation in counterfeit procurement rates across different retailers should be expected. In the lower-left graph, it is interesting to observe that procurement of authentic goods collapses entirely when reputation is equally shared, whereas in the top-right graph the highest market procurement of real goods occurs over an interval centered on $\lambda_1 = 0.5$ (and also the edge case of $\lambda_1 = 0$ or 1). In the latter case, the retailers are so strongly committed to full authentic procurement that they will not introduce any counterfeits until their proportion of reputation falls below 0.43 (approximately), and will begin procuring real products once their share of reputation exceeds 0.17 (approximately). Interestingly, as shown in the lower-right graph in Figure 5-7, the reputation bearer slightly increases their procurement of real goods as their proportion of reputation falls from 1 to 0.7 (approximately); the retailer does so in order to offset the demand-depressing effect of reputation influence shifting to the retailer who initially procures only counterfeits. The gains from unilaterally improving the shared reputation depend on the shape of the demand curve (stronger with concavity and a higher quasi-slope parameter). It appears these gains were not sufficiently strong to induce similar behavior in the other three settings.

Next, we will examine more flexible notions of reputation sharing, and will then compare the equilibrium outcomes across a few scenarios to investigate the persistence

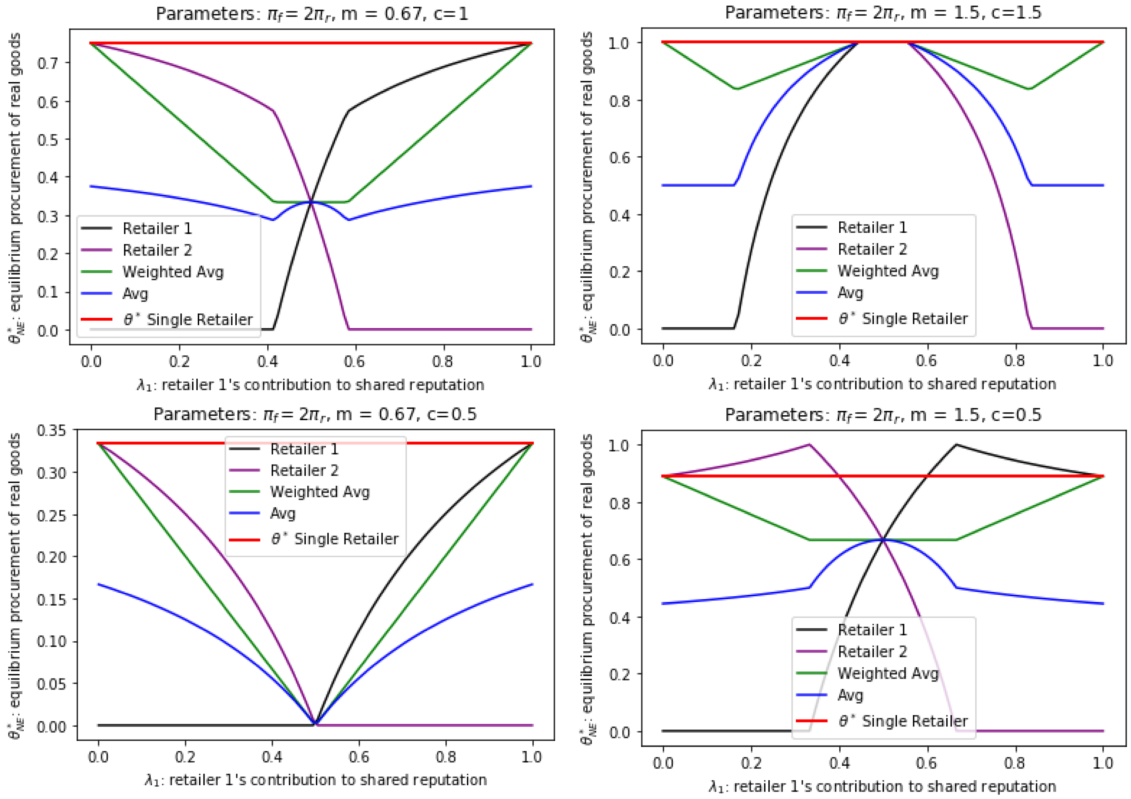

Figure 5-7: Nash Equilibrium decisions $\theta^*_{i,NE}$ as Retailer 1's contribution to a fully shared reputation varies from 0 to 1. Weighted average (green) is the overall market level of authentic goods weighted by share of reputation. Average (blue) is overall market level of authentic goods with 50% contribution by each retailer.

and severity of the effect of reputation sharing.

5.4.3 Multiple Retailers, Partially Shared Reputation

Now consider a shared reputation model in which, for each retailer, λ percent of the retailer's reputation comes from its own actions, and $(1-\lambda)$ percent comes from the average of other retailers' actions. This arrangement allows for all retailers to have more than 50% influence over the perceived θ ascribed to them by consumers. Treating all retailers as identical, we can solve the Nash Equilibrium by substituting θ_i for θ_{-i} in Equation 5.12 and solving $\theta_i^* = BR_i(\theta_i^*)$, since $\theta_i^* = \theta_{-i}^*$ by symmetry. We therefore want to solve $\theta = \min\left[1, \max\left[0, \frac{c\pi_f}{(c+1)(\pi_f-\pi_r)} - \frac{b+m(1-\lambda)\theta}{\lambda m(c+1)}\right]\right]$. When the equation is satisfied for some θ such that the inner expression is interior to (0,1), we can directly solve for $\theta_{i,NE}^*$ as $\theta = \frac{c\pi_f\lambda m - b(\pi_f-\pi_r)}{m(c\lambda+1)(\pi_f-\pi_r)}$. In order for this θ to produce an interior solution, it must fall within the following domain: $\theta \in \left(\frac{c\pi_f\lambda m}{(\pi_f-\pi_r)((1-\lambda)m+b)} - \frac{\lambda m(c+1)}{b+m(1-\lambda)}, \frac{c\pi_f\lambda m}{(\pi_f-\pi_r)((1-\lambda)m+b)}\right)$. These are the boundaries of the interior, assuming that $b + m(1-\lambda)\theta_i > 0$. If it happens to be negative, then the boundaries of the interior are reversed. To simplify notation, denote the lower boundary of the interior as L and the upper boundary as U. We can formally characterize the Nash Equilibrium as follows:

$$\theta_{i,NE}^* = \begin{cases} \min\left[1, \max\left[0, \frac{c\pi_f\lambda m - b(\pi_f-\pi_r)}{m(c\lambda+1)(\pi_f-\pi_r)}\right]\right] & \text{if } \theta_{i,NE}^* \in (L,U) \\ 0 & \text{if } \theta_{i,NE}^* \leq L \\ 1 & \text{if } \theta_{i,NE}^* \geq U \end{cases} \tag{5.22}$$

As shown in Figure 5-8, the retailer in a market with a partially shared reputation procures more authentic goods than when other retailers are myopic. Compared to a market with a fully shared reputation, the retailer will procure fewer authentic goods than the reputation-bearer, but more than the reputation free-rider. Having multiple retailers in the partially shared reputation setting produces the same equilibrium as

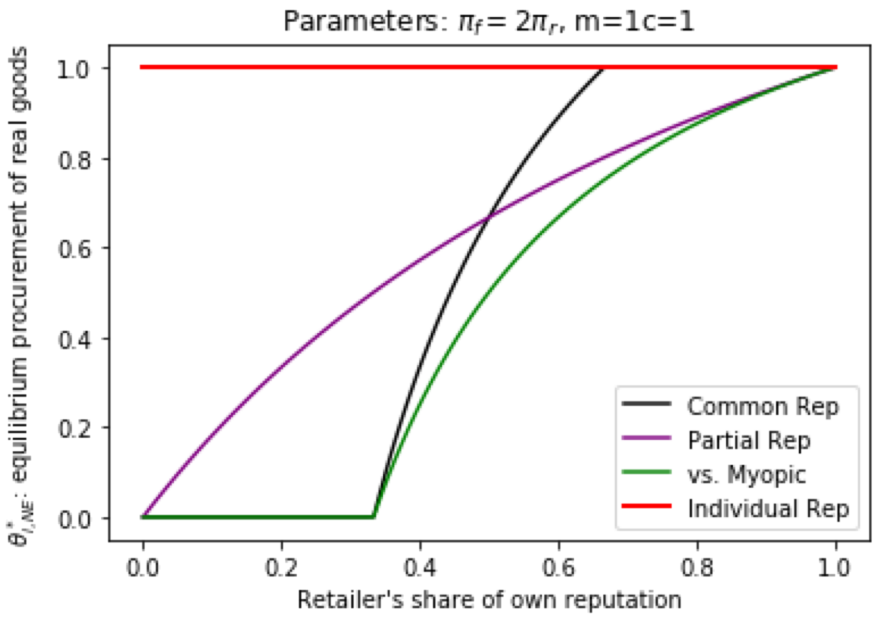

Figure 5-8: A retailer's authentic procurement level as a function of λ_i. The partially shared reputation is shown in purple. The fully shared reputation is shown in black. The decision taken when the other retailer(s) are myopic is shown in green.

the two-retailer partially shared reputation setting as long as the parameter λ that weights the importance of a retailer's own actions on their reputation is unchanged.

5.4.4 Multiple Retailers, Fully Shared Reputation

The case of multiple retailers who fully share a common reputation can be represented by $\theta = \sum_{i=1}^{k} \lambda_i \theta_i$, $\lambda \geq 0$ and $||\lambda||_1 = 1$. The best response functions have roughly the same form as in the case of two retailers. We need only a slight change in notation, letting θ_{-i} denote the weighted average of the actions of the other retailers: $\frac{1}{1-\lambda_i} \sum_{j \neq i} \lambda_j \theta_j$. The best response function for any retailer (indexed by i) is therefore given by

$$BR_i(\theta_{-i}) = \min \left[1, \max \left[0, \frac{c\pi_f}{(c+1)(\pi_f - \pi_r)} - \frac{b + m(1-\lambda_i)\theta_{-i}}{\lambda_i m(c+1)} \right] \right].$$

Let's first consider the case where $\lambda_i = \frac{1}{k}$ for $i = 1, ..., k$.

In fact, if $m < 1$, then $\forall\ K \geq \frac{c\pi_f m(c+1)}{(c+1)(\pi_f - \pi_r)(1-m)}$ then the Nash Equilibrium is $\theta^*_{NE} = 0$ represents a market collapse in the sense that no authentic products will be procured. Conceptually, a market collapse equilibrium happens if all retailers choose $\theta_i = 0$ and the x-axis intercept for the demand function doesn't become positive until a value of θ that no-one can reach unilaterally by changing their own θ_i. Such a scenario may characterize certain retail markets where it is generally understood by consumers that goods presented as branded luxury goods are knock-offs, such as flea markets and open-air markets, and in such settings a vendor would have no incentive to procure authentic luxury goods (even if some consumers are willing to pay full price for them) because no consumers would believe the goods to be authentic, given the high prevalence of fakes in the market.

5.5 Single Retailer, Logistic Demand Function

It may be more realistic, though not necessarily more tractable, to model aggregate demand with respect to θ using an analytical function that is zero or close to zero when θ is close to zero, and one or close to one when θ approaches 1, and whose derivative is close to zero at these two extremes. Such a function would be reasonable if the effect of a small amount of counterfeits on demand is initially small, perhaps because small amounts of counterfeits for certain types of products are not scandalous and go undetected or unreported by the media. Likewise, if mostly counterfeit goods are sold, it is conceivable that consumer confidence in the product will make a slow recovery as the counterfeit level goes from being extremely terrible to slightly less terrible.

Sigmoidal functions have these properties, and the most natural one, which we will consider, is the logistic function:

$$D(\theta) = \frac{1}{1 + e^{-k(\theta - x_0)}}. \tag{5.23}$$

The parameter k controls the scale of the function (how stretched or compressed the S-shape is over the interval [0,1]) and x_0 is its center. The function's shape over [0,1] could be made strictly concave or strictly convex by adjusting the center x_0 and re-scaling the vertical axis, and is therefore quite flexible in the variety of relationships it can model. This function has additional intuitive appeal. The logistic function's derivative is elegant:

$$kD(\theta)(1 - D(\theta). \tag{5.24}$$

The factor $(1-D(\theta))$ represents the fraction of customers who have yet to buy, and it would make sense that the fewer customers remaining to be converted, the slower the growth in demand. Likewise, if $D(\theta)$ is low, because θ is low, then the rate of converting remaining customers may be low for incremental improvements in θ. However, it is not necessarily true that the conversion rate of potential customers is monotonic in θ; in practice, it could peak around some "good enough" level that is less than 1, around which most consumers' decision threshold may be concentrated. Notwithstanding this limitation, the logistic function has most of the important properties needed to model aggregate demand.

Although $D(\theta)$ only evaluates to 0 and 1 in the limit, we can still use this functional through the normalized function: $D_N(\theta) = \frac{D(\theta)-D(0)}{D(1)-D(0)}$, which will result in $D_N(1) = 1$ and $D_N(0) = 0\ \forall k, x_0$. Moreover, since this normalization is an affine transformation, we have $D'_N(\theta) = \frac{D'(\theta)}{D(1)-D(0)}$.

5.5.1 Retailer Profit Maximization

We model and solve the retailer's profit maximization problem in the same manner as in Section 5.3.1, but using the logistic demand function:

$$\Pi_r = D(\theta)[\theta\pi_r + (1-\theta)\pi_f]$$

$$\text{let } \pi_r = 1$$

$$= D(\theta)[\theta + (1-\theta)\pi_f]$$

$$\text{Take the derivative w.r.t. } \theta$$

$$\frac{d\Pi_r}{d\theta} = D'(\theta)[\theta + (1-\theta)\pi_f] + D(\theta)[1-\pi_f] \tag{5.25}$$

The First Order Condition is given by:

$$D'(\theta)[\theta + (1-\theta)\pi_f] + D(\theta)[1-\pi_f] = 0, \text{ and by Equation 5.23}$$

$$\Rightarrow kD(\theta)(1-D(\theta))[\theta + (1-\theta)\pi_f] + D(\theta)[1-\pi_f] = 0.$$

$$\text{Simplify by dividing both sides by } D(\theta) \text{ (strictly positive)}$$

$$\text{and express } (1-D(\theta)) \text{ in full form:}$$

$$\Rightarrow k\left(\frac{e^{-k(\theta-x_0)}}{1+e^{-k(\theta-x_0)}}\right)[\theta + (1-\theta)\pi_f] + [1-\pi_f] = 0$$

$$\Rightarrow k\left(\frac{1}{e^{k(\theta-x_0)}+1}\right)[\theta + (1-\theta)\pi_f] + [1-\pi_f] = 0. \tag{5.26}$$

Using this equation, one can quickly compute optimal unconstrained θ^* numerically, given k, x_0, and π_f. However, the expression for θ^* involves the analytic continuation of the product-log function, $W_n(\cdot), n \in \mathbb{Z}$, and has no closed form[1]:

$$\theta^* = \frac{-(\pi_f-1)W_n\left(e^{\left(\frac{kx_0+\pi_f(-x_0k+k-1)+1}{(\pi_f-1)}\right)}\right) + \pi_f(k-1)+1}{(\pi_f-1)k}. \tag{5.27}$$

As a result of the expression θ^* being unwieldy, the closed-form expressions and analyses obtained in the preceding sections that required θ^* do not have analogs when using logistic demand. While such quantities could be examined computationally and

[1]Expression obtained through Wolfram-Alpha using Pro Premium account

plotted for specific combinations of π_f, k and x_0, the ability to generate insights that are independent of specific parameters is somewhat diminished. This loss highlights the general importance of balancing realism and tractability when choosing functions and models.

However, some of the analyses from the preceding sections can be reproduced. As before, there is a minimum profit-margin for counterfeits (as a multiple of π_r), below which it is optimal to procure only authentic products. Given some k and x_0, we solve for the value of π_f that causes $\theta^* = 1$ to be a solution to equation 5.26. Denote this minimum profit level for fakes as $\pi_f^{min}(k, x_0)$. We can compute this quantity by letting $\theta = 1$ in equation 5.26:

$$\left(\frac{k}{1+e^{k(1-x_0)}}\right)[1+(1-1)\pi_f]+[1-\pi_f]=0$$

$$\left(\frac{k}{1+e^{k(1-x_0)}}\right)=\pi_f$$

$$\therefore \pi_f^{min}(k,x_0)=\frac{k}{1+e^{k(1-x_0)}}+1 \tag{5.28}$$

This derivation of $\pi_f^{min}(k, x_0)$ depends on the assumption that $\frac{d\theta^*}{d\pi_f} < 0$, which held true in the preceding sections. This assumption held true for all logistic demand parameter combinations that were checked manually but has not been proven analytically. It is straightforward to show that the equivalent condition for a normalized logistic demand function $D_N(\theta) = \frac{D(\theta)-D(0)}{D(1)-D(0)}$ is given by

$$\pi_{f,norm}^{min}(k,x_0)=\left(\frac{1}{D(1)-D(0)}\right)k\frac{e^{-k(1-x_0)}}{(1+e^{-k(1-x_0)})^2}+1. \tag{5.29}$$

Proof:

$$\frac{d\Pi_r}{d\theta} = D'_N(\theta)[\theta + (1-\theta)\pi_f] + D_N(\theta)[1-\pi_f]$$

$$\text{Evaluate FOC at } \theta^* = 1, \text{ use } D'_N(\theta) = \frac{D'(\theta)}{D(1)-D(0)} \text{ and } D_N(1) = 1$$

$$\Rightarrow \frac{D'(1)}{D(1)-D(0)}(1) + D_N(1)(1-\pi_f) = 0$$

$$\Rightarrow \frac{D'(1)}{D(1)-D(0)} + (1)(1-\pi_f) = 0$$

$$\Rightarrow \frac{D'(1)}{D(1)-D(0)} + 1 = \pi_f$$

$$\text{Expand } D'(1) = kD(1)(1-D(1)) = k\frac{1}{1+e^{-k(1-x_0)}}\frac{e^{-k(1-x_0)}}{1+e^{-k(1-x_0)}}$$

$$\Rightarrow \pi_f^{min}(k, x_0) = \left(\frac{1}{D(1)-D(0)}\right) k\frac{e^{-k(1-x_0)}}{(1+e^{-k(1-x_0)})^2} + 1$$

The logistic demand curve does not lend itself to producing closed form expressions for the other quantities of interest, though they could be estimated and plotted through numerical computation and simulation over a range of parameter instances. However, if numerical and simulation methods are to be used in the first place, then it may be worthwhile forgoing the logistic demand curve in favor of a more complex demand model that can reflect a variety of assumptions about consumer utility. Section 5.6 takes this approach in order to model the potential impact of consumer-level product verification technology and financial incentives for verification.

5.6 Product Verification Services

One strategy in the fight against counterfeits is to make it easier for consumers to identify authentic products. Some companies use fancy packaging designs with holographic foil to raise the cost and difficulty of mimicry, though this does not prevent the packaging from being copied or imitated. Even when such copies are imperfect, many consumers may not know what the correct package was supposed to look like, and even if they do, the imperfections may go undetected. In low-income countries, even major typographical or spelling errors may be hard to recognize for a large frac-

Figure 5-9: Sproxil Mobile Authentication Service. Photo from Fight the Fakes [2015].

tion of the population. Consumer awareness of counterfeits will also vary by product and location.

To overcome this challenge, a company called Sproxil, headquartered in Cambridge MA, developed a consumer product verification service for manufacturers called Mobile Authentication Service (MAS). MAS involves adding a covert single-use scratch code, akin to the serial number, to the product packaging. The consumer texts this message to a MAS phone number, the number is checked against a database, and the consumer receives confirmation, or a follow-up message or call if the number does not match the database (see Figure 5-9). Optionally, the manufacturer using MAS has the option to give the consumer financial incentives, loyalty points, mobile phone credit, etc., every time they authenticate a real product.

The point at which consumers can access the scratch code may vary by retailer practice. Some retailers may allow potential customers to authenticate the product before purchase if the code is accessible on the exterior packaging, in which case the retailer bears the risk of the product being fake. Other retailers may promise a refund if the product turns out to be fake, or may be obliged if the payment was

made through a credit card company with strong consumer protection policies, but may insist that the consumer conduct the verification on-site just after purchase. In other cases, the consumer may have to check the product at home, and might not receive hassle-free redress from the retailer if the product turns out to be fake. In this case, the retailer might keep its profit from the transaction, but could still lose the customer's future business. The customer is out of pocket for the price of the item, but at least knows the authenticity of the product and can avoid consuming it if doing so carries risk of harm.

The possibility of product authentication should be expected to greatly affect the decisions of retailers regarding the level of willful procurement of black market and grey market products that are, or likely to be, counterfeit. The problem of modeling aggregate consumer decisions is now more complex, as it depends on both the prevalence of counterfeits, whether MAS is available, and whether incentives to use MAS are offered. The procurement decision of retailers is also more complex. It must take into account not only aggregate demand, but also the propensity of consumers to check products, which could involve various direct and indirect penalties when they are caught selling counterfeit products. Finally, this situation now involves the manufacturer, who must decide whether to adopt MAS, and how much to incentivize product verification, if at all, based on the retailer's procurement decision, and aggregate consumer demand and checking behavior. To handle the added complexity of the model, we take a numerical approach to computing equilibrium decisions for each type of agent and use graphics from a range of scenarios to communicate the patterns and findings of interest.

5.6.1 Multi-Level Sequential Decision Model

In this section we present the notation and mathematical structure for the sequential decision game played by the manufacturer, M, the retailer , R, and consumers, C. The manufacturer decides on an average expenditure x_m per product sold so as to incentivize the consumers to use the mobile authentication service. We let $x_m \in [0, \pi_m]$, where π_m is the manufacturer's marginal profit per product sold. We will

assume that π_m is independent of the sales quantity.

As in the previous sections, the retailer has inventory sourced from the manufacturer and some from illegitimate (i.e., counterfeit) sources. The retailer decides on a fraction $\theta \in [0, 1]$ of its inventory that will be "authentic," or procured from the manufacturer. We will the complement, $\theta^c = 1 - \theta$, to represent the percent of goods procured from counterfeiters. The retailer sells authentic and fake products for the same price (p). For now, we will treat the price as exogenous to the decision problem, which may be reasonable if the retailer must stay close to the market price or MRSP to avoid transmitting a signal that it is covertly selling cheap counterfeits.

We do not consider lead time or inventory holding costs for the retailer. Although retailers might sell counterfeits to different customer types at different rates in practice, we will assume that the retailer provides all customer types with an authentic product with probability θ, and a counterfeit with probability θ^c.

Previously, we modeled consumer demand with a simple family of differentiable functions that are increasing in θ, and the shape of the curve was given by mathematical parameters that did not have a strong practical interpretation tying it to a particular consumer utility model. This high-level approach is harder to justify in the more complex setting involving MAS. Instead, we propose a decision model in which each consumer $i \in C$ makes decisions based on three individual parameters and three decisions set by other agents:

1. β_i: their expected utility for consuming the real product;
2. η_i: their expected utility for consuming a counterfeit;
3. ϵ_i: the consumer's cost to verify the product's authenticity using MAS;
4. p, θ, and x_m.

In our setting, the parameter η_i will take negative values to represent the serious negative consequences or possibility of death from using counterfeit medicine, medical devices, safety equipment, etc. The parameter ϵ is meant to capture all direct, indirect, and subjective costs. It can be thought of as the minimum financial incentive

needed to get the consumer's attention and induce them, with high probability, to take the time to go through the MAS steps even if there was no risk of counterfeits being present.

The population of consumers is assumed to be heterogeneous. In our numerical experiments, we treat the three individual parameters as varying according to a few different multivariate normal distributions. Based on these parameters, each individual consumer will decide whether or not to buy the product at price p:

$$D_i = \begin{cases} 1 & \text{if customer i decides to buy} \\ 0 & \text{if customer i decides to not buy} \end{cases}$$

In a best-case scenario for consumers and manufacturers, the consumer will be able to verify the product before purchase. If the individual buys the product, then they are faced with a second decision of whether or not to use the authentication service. This decision will be made before consumption of the product, so that if consumer i checks and finds that their medicine is counterfeit, they do not consume it. To take a conservative approach to estimating the benefit of MAS, we consider the case where the consumer is only able to verify the product after purchase. We do not assume that the consumer will be able to get a refund if a fake product is detected. Since there could be an unpredictable and highly variable amount of effort required to get a refund, this assumption may be somewhat realistic where consumer protection norms are weak and credit card usage is low. If refunds after discovering counterfeits is possible with low effort by the consumer, then both product demand and the usage of MAS would be even higher than implied by this section. The individual consumer's checking decision is denoted by

$$X_{c,i} = \begin{cases} 1 & \text{if customer i decides to check} \\ 0 & \text{if customer i decides to not check.} \end{cases}$$

The order in which the game proceeds is as follows:

1. The manufacturer chooses an incentive level x_m for MAS usage, which becomes

public information.

2. The retailer chooses a fraction θ of its inventory to be authentic. Although the choice θ is private information, we will suppose that the manufacturer knows the retailer's utility function (crucially, profit margin difference between real and fake goods) can therefore compute the retailer's best response to consumer aggregate demand.

3. A consumer chooses whether or not to buy; if they buy, then they make a second decision to authenticate or not. Each consumer could be viewed as having a monotonic belief function $\hat{\theta}_i = B_i(\theta)$ that generates their estimate of product authenticity levels (or product quality levels) based on news, awareness campaigns, their personal experience with a product, word of mouth, etc. For now, we will let $\hat{\theta}_i = \theta$. We assume that the retailer and manufacturer can estimate the aggregate best response of consumers to x_m and θ; alternatively, if the retailer's profit function is concave with respect to aggregate $D(\theta, x_m)$, then they can eventually arrive at a profit-maximizing θ^* through incremental profit-improving steps.

Aggregate Demand and Checking

The sum of all customers' buying and checking decisions are represented at the "aggregate demand level", D $\in [0, D_{max}]$. We re-scale this quantity by setting $D_{max} = 1$ so that it represents the fraction of the market that buys, and $D \in [0, 1]$. The fraction of the market that both buys and verifies using MAS is defined as $X_c \in [0, 1]$. The aggregate consumer best response functions $D(\theta, x_m)$ and $X_c(\theta, x_m)$ depend on the distribution of consumer parameters that are part of the consumer utility model.

Manufacturer Utility Function

The manufacturer's utility function is defined as

$$U_M(x_m, \theta, X_c) = \pi_m D\theta - x_m \theta X_c = \theta(D\pi_m - x_m X_c). \tag{5.30}$$

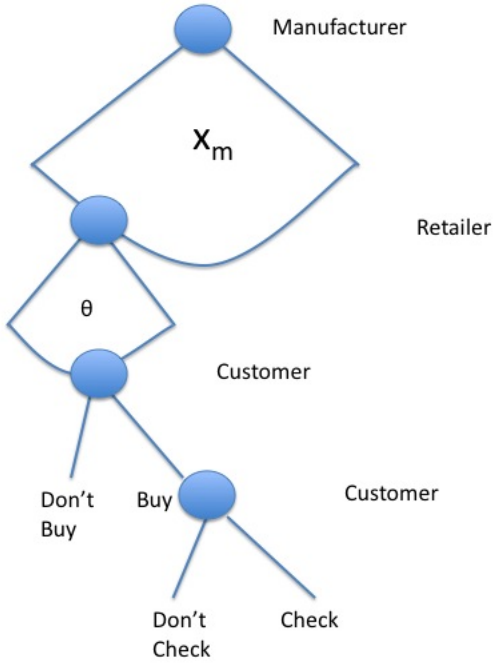

Figure 5-10: Sequential Game Structure

For now, we assume the manufacturer receives no negative utility if the consumer tests a product and finds a counterfeit targeting that manufacturer's brand.

Retailer Utility Function

The retailer's utility is their expected profit from real and fake sales, minus some penalty that is given as a function, $\alpha(\cdot)$, of the number of fakes detected. We assume that $\alpha(\cdot)$ is increasing in the number of fakes detected (e.g., risk of manufacturer punishment, lost consumer loyalty, law enforcement, etc):

$$U_R(\theta, X_c) = \theta D\pi_r + (\theta^c)D\pi_f - \alpha(D, X_c, \theta). \tag{5.31}$$

As before, π_r is the retailer's unit profit on real products, π_f is its profit on fake products, and $\pi_r < \pi_f$. The α penalty function for detected fakes is assumed to be positive and strictly increasing. The quantity $\theta^c X_c$ is the number of detected fakes. The penalty should have a component that is linear in the number of fakes detected, as this can represent the expected value of customers attempting to return

merchandise and losing future sales from those customers, with some probability. The penalty should also have a strictly convex component that increases in the proportion of the retailer's sales that are discovered to be counterfeit. This component can represent reputation damage and the growing risk of penalties from regulation or law enforcement. One possible penalty function could take the following form:

$$\alpha(D, X_c, \theta) = aDX_c(1-\theta) + b(cX_c(1-\theta))^d, \tag{5.32}$$

where a, b, c, d are positive constants, and $d > 1$. Note: b induces a vertical transformation in the convex component of the penalty, while c induces a horizontal transformation with respect to the proportion of sales that are detected as fake.

Customer Utility

Recall:

- β_i: the benefit for person i of consuming an authentic product;
- η_i: the utility of consuming a fake product (negative);
- ϵ_i the magnitude of cost (direct costs and "effort" cost) of going through the product verification process;
- D_i and $X_{c,i}$: indicator variables for whether person i buys and checks the product, respectively, with $X_{c,i} \leq D_i$;
- p the market price of the product (assume exogenous).

For now, assume that consumers can only test a product's authenticity after purchase, and does not get reimbursed by anyone if a fake is detected.

Consumer i's utility is given as follows:

$$U_{c,i}(D_i, X_{c,i}, \theta, x_m) = D_i\left[\theta(\beta_i + X_{c,i}(x_m - \epsilon_i)) + \theta^c((1-X_{c,i})\eta_i - X_{c,i}\epsilon_i) - p\right]. \tag{5.33}$$

It is more useful, however, to model purchasing and verifying separately, and we can do so since these decisions are made sequentially. First, consider the utility that depends on $X_{c,i}$, assuming a purchase has already been made:

$$U^{v}_{c,i}(X_{c,i}, \theta, x_m) = \theta \left[\beta_i + X_{c,i}(x_m - \epsilon_i)\right] + \theta^c \left[(1 - X_{c,i})\eta_i - X_{c,i}\epsilon_i\right]. \tag{5.34}$$

Let $U^{v*}_{c,i}(X_{c,i}, \theta, x_m)$ be the max of this function over $X_{c,i}$. Then $D_i = 1$ iff $U^{v*}_{c,i}(X_{c,i}, \theta, x_m) \geq p$, and $X_{c,i} = D_i = 0$ otherwise.

Aggregate Demand Function - Specifics

Having defined the consumer's utility function, we can now begin to characterize the aggregate best response functions $D(\theta, x_m)$ and $X_c(\theta, x_m)$. Assuming customers wish to maximize their utility, customer i purchases if and only if

$$\max_{X_c \in \{0,1\}} \theta(\beta_i + X_c(x_m - \epsilon_i)) + (1 - \theta)((1 - X_c)\eta - X_c\epsilon_i) - p \geq 0$$

$$\iff$$

$$\max\{\theta\beta_i + (1 - \theta)\eta_i,\ \theta\beta_i + \theta x_m - \epsilon_i\} \geq p \tag{5.35}$$

$$\iff$$

$$\min\left(\frac{p - \eta_i}{\beta_i - \eta_i}, \frac{p + \epsilon_i}{\beta_i + x_m}\right) \leq \theta. \tag{5.36}$$

Expression 5.36 shows that $D(\cdot, x_m)$ is a right-continuous, nondecreasing, function of θ, which is to say that the utility has the intended property that more consumers will buy the product as the percentage of authentic goods increases. Denote A as the event that the first argument of Expression 5.35 is greater than or equal to p, and B as the event that the second argument is greater than or equal to p. Then $Pr(D_i = 1|\theta, x_m) = Pr(A \cup B|\theta, x_m)$.

If we assume that the consumer parameters are drawn from a distribution that is the same for all consumers, then $D(\theta, x_m) = Pr(D_i = 1|\theta, x_m)\ \forall i \in C$ and we can

stop indexing β, η, ϵ by consumer. We can therefore express $D(\theta, x_m)$ as

$$D(\theta, x_m) = 1 - Pr\left(A^c \cap B^c\right) = 1 - Pr(\theta\beta + (1-\theta)\eta \leq p \ \cap \ \theta\beta + \theta x_m - \epsilon \leq p\right). \tag{5.37}$$

Observe that for any fixed β, we can define the set of ϵ and η for which condition A^c and B^c hold:

$$\epsilon \geq \underline{\epsilon}(\theta, x_m, \beta, p) = \theta\beta + \theta x_m - p \tag{5.38}$$

$$\eta \leq \bar{\eta}(\theta, \beta, p) = \frac{p - \theta\beta}{1-\theta}, \quad \text{for } \theta < 1 \tag{5.39}$$

Note that for $\theta = 1$ the condition A^c will hold for any η if $\beta \geq p$, and for no η otherwise. With a domain now specified for ϵ and η, we can express $D(\theta, x_m)$ with the following integral:

$$D(\theta, x_m) = 1 - \int_{-\infty}^{\infty} \int_{-\infty}^{\underline{\eta}(\theta,\beta,p)} \int_{\bar{\epsilon}(\theta,x_m,\beta,p)}^{\infty} f(\epsilon, \eta, \beta) d\epsilon d\eta d\beta \tag{5.40}$$

where $f(\epsilon, \eta, \beta)$ is the joint density function for these parameters in the consumer population.

Aggregate Checking Function - Specifics

The aggregate checking function is more complex than the aggregate demand function. Recall that $X_c(\theta, x_m)$ is not the percentage of buyers who check, but the percentage of potential consumers in the market who buy and check. It is easier to first consider the decision of those who have already made a purchase decision. Those consumers will check the product if and only if

$$(1-\theta)\eta \leq \theta x_m - \epsilon \iff \begin{pmatrix} \theta \geq \frac{\eta+\epsilon}{\eta+x_m} & \text{if } \eta + x_m > 0 \\ \eta - \epsilon \geq 0 & \text{if } \eta + x_m = 0 \\ \theta \leq \frac{\eta+\epsilon}{\eta+x_m} & \text{if } \eta + x_m < 0 \end{pmatrix}. \tag{5.41}$$

Expression 5.41 shows two distinct motivations for buyer checking. One group of buyers will check when the product authenticity is sufficiently high because they are seeking the reward from x_m. Another group will only check when the authenticity is sufficiently low, to avoid disutility η. Which group a customer belongs to depends on the customer's η compared to the manufacturer's offered incentive x_m.

Given a joint distribution for β, η, and ϵ, the function $X_c(\theta, x_m)$ can be expressed as the probability of the following event:

$$\{\text{Buy and Check}\} \iff$$
$$\{\epsilon \leq \theta x_m - \eta(1-\theta)\} \text{ AND } \{\{\frac{p-\eta}{\beta-\eta} \leq \theta\} \text{ OR } \{\frac{p+\epsilon}{\beta+x_m} \leq \theta\}\}. \tag{5.42}$$

We can compute the probability of the event of buying and checking as the probability of both the checking condition and the first buying condition being true, plus the probability of both the checking condition and second buying condition being true whilst the first buying condition is false. Mathematically, this is expressed as follows:

$$Pr(\text{Buy and Check}) = \int_{-\infty}^{0} \int_{\frac{p-\eta(1-\theta)}{\theta}}^{\infty} \int_{0}^{\theta x_m - \eta(1-\theta)} f(\eta, \beta, \epsilon) \partial\eta \partial\beta \partial\epsilon + \tag{5.43}$$
$$\int_{-\infty}^{0} \int_{0}^{\frac{p-\eta(1-\theta)}{\theta}} \int_{0}^{\min\{\theta x_m - \eta(1-\theta),\ \theta\beta + \theta x_m - p\}} f(\eta, \beta, \epsilon) \partial\eta \partial\beta \partial\epsilon \tag{5.44}$$

For complicated distributions, this integral could simply be computed using Monte Carlo simulation to sample from the distribution and compute the proportion of a large number of samples that satisfy condition 5.42. The function $X_c(\theta, x_m)$ can be numerically approximated by evaluating this integral over a grid of values for θ and

x_m. The percentage of buyers who check, as opposed to the percentage of consumers who will buy and check, can be computed as follows:

$$Pr(\text{Check}|\text{Buy}) = \frac{X_c(\theta, x_m)}{D(\theta, x_m}). \tag{5.45}$$

The shape of the function represented by Equation 5.45, the percentage of buyers who check, is not amenable to algebraically-driven conceptualization or algebraic analysis of its partial derivatives. We present graphs of the function for a few different multivariate normal distributions in Section 5.7.

Aggregate Checking Function - Special Case

Before proceeding to analyze the retailer's decision, we consider a special case for which the percent of buyers who check is more amenable to analysis. In this case, we suppose that all consumers are given the product for free (e.g., medical or pharmaceutical products funded by insurance or government), and they either exert effort to check the product or consume it without checking. Let $H(\theta, x_m)$ denote the probability of checking under this setup.

Let $\eta_i \sim N(\mu_\eta, \sigma_\eta^2)$ and $\epsilon_i \sim N(\mu_\epsilon, \sigma_\epsilon^2) \quad \forall i$, and suppose $cov(\eta, \epsilon)$ is known. The checking condition can be expressed as $\theta x_m > \epsilon_i - \eta_i(1-\theta)$. The left-side of this inequality is the product of the decision of the manufacturer and retailer. The right side is a linear combination of two normally distributed random variables. We can therefore represent the probability of checking using a cumulative distribution function: $H(\theta, x_m) = \Phi_{Z_\theta}(\theta x_m)$, where Φ_{Z_θ} denotes the CDF of the random variable $Z_\theta = \epsilon - \eta(1-\theta)$. The CDF is the following normal distrubtion: $Z_\theta \sim N(\mu(\theta), \sigma^2(\theta))$, where $\mu(\theta) = (\theta - 1)\bar{\eta} + \bar{\epsilon}$, and $\sigma^2(\theta) = (1-\theta)^2 Var(\eta) + Var(\epsilon) + 2(1-\theta)cov(\eta, \epsilon))$.

We can obtain closed-form expressions for the partial derivatives of the checking proportion with respect to the incentive x_m and the authenticity rate θ:

$$\forall x_m \in [0,1] \quad \frac{\partial H(\theta, x_m)}{\partial \theta} = \frac{\partial \Phi_{Z_\theta}(\theta x_m)}{\partial \theta}$$
$$= \frac{1}{\sqrt{\pi}} \exp\left(-\frac{(x_m\theta - \mu(\theta))^2}{2\sigma^2(\theta)}\right) \cdot \frac{\partial}{\partial \theta}\left[\frac{x_m\theta - \mu(\theta)}{\sqrt{2\sigma^2(\theta)}}\right] \tag{5.46}$$

$$\forall \theta \in [0,1] \quad \frac{\partial H(\theta, x_m)}{\partial x_m} = \frac{\partial \Phi_{Z_\theta}(\theta x_m)}{\partial x_m} = \frac{\partial}{\partial x}\Phi_{Z_\theta}(x)|_{x=x_m\theta} \cdot \frac{\partial}{\partial x_m}[x_m\theta]$$
$$= \frac{\theta}{\sqrt{2\Pi\sigma^2(\theta)}} \exp\left(-\frac{(x_m\theta - \mu(\theta))^2}{2\sigma^2(\theta)}\right) \tag{5.47}$$

5.6.2 Retailer Best Response

We have now derived two functions representing the aggregated best responses of heterogeneous consumers. To solve the game's sub-game perfect equilibrium, we must determine the retailer's best response (BR), and then the manufacturer's. First, consider the retailer's BR if we assume D and X_c are inelastic with respect to θ, and that the penalty function has the form $\alpha(D, X_c, \theta) = aDX_c(1-\theta) + b(cX_c(1-\theta))^d$. Penalty functions of this form have a linear component, representing the expected loss per customer when a customer detects a fake (e.g., reduction in future purchases from that customer), and the second term makes the penalty strictly convex (increasing) to reflect more severe reputation damage, regulatory risk, supplier sanctions, or criminal prosecution.

Under this model, optimal θ^* for an interior solution is given by the first-order condition:

$$\theta^* = 1 - \left(\frac{D(\pi_f - \pi_r) - aDX_c}{bd(cX_c)^d}\right)^{\frac{1}{d-1}} \tag{5.48}$$

The D and X_c may be inelastic with respect to θ over very short time intervals since there will be a delay between higher counterfeit prevalence and consumer per-

ceptions shifting, but over long time periods D and X_c depend on θ. The retailer must choose long-term θ knowing that the customers will eventually respond to *theta*. Consequently, it would be useful to solve the FOC using $D(\theta, x_m)$ and $X_c(\theta, x_m)$ instead of fixed D and X_c. To do this, we would need ∇D and ∇X_c.

To compute $\frac{\partial D(\theta, x_m)}{\partial \theta}$, we apply the Leibnitz Rule since θ appears in the lower bound of the inner integrals:

$$\frac{d}{d\theta}D(\theta) = \int_{-\infty}^{\infty}\int_{\underline{\eta}(\theta)}^{\infty} f(\underline{\epsilon}(\theta), \eta, \beta)\underline{\epsilon}'(\theta)d\eta d\beta + \int_{-\infty}^{\infty}\int_{\underline{\epsilon}(\theta)}^{\infty} f(\epsilon, \underline{\eta}(\theta), \beta)d\epsilon \cdot \underline{\eta}'(\theta)d\beta \quad (5.49)$$

Given the complexity of the expression for X_c, we do not consider it practical to produce an expression for ∇X_c, and recommend solving the retailer's best response through numerical approximations. Alternatively, one could use the partial derivatives of $H(\theta, x_m)$, given in equations 5.46 and 5.47.

5.7 Results of Numerical Simulations

We found it most practical to solve the sub-game perfect equilibrium using brute-force numerical techniques. The aggregate demand and aggregate checking best response functions were computed by performing numerical integration. The functions were evaluated over a mesh grid of $[\theta, x_m] \in [0, 1]^2$ and the resulting two 2D-arrays were treated as an approximate best response surface. The retailer's profit, given $D(\theta, x_m)$ and $X_c(\theta, x_m)$, was then computed over the same mesh-grid. The retailer's penalty function was convex and took the form given by Equation 5.32. For each value of x_m, the θ that maximized retailer profit was treated as the retailer's best-response. The manufacturer's profit was then computed as a function of x_m, which determined the value of θ, D, and X_c.

We present results for three different distributions of β, η, and ϵ. In case 1, $\bar{\beta} \approx |\bar{\eta}| \approx \bar{\epsilon}$, which means that the average perceived harm from consuming a counterfeit is about the same as the average perceived benefit and the average incentive

needed to induce checking (in the absence of any harm from counterfeits). In case 2, $|\bar{\eta}| > \bar{\beta} > \bar{\epsilon}$, which is meant to representative of tuberculosis (TB) medication. For TB, the patient must go through six months of uninterrupted treatment. Missing doses (i.e., by consuming counterfeit antibiotics) can result in the patient needing to restart treatment all over again, or the development of drug-resistant TB strains. The benefit of consuming the medication should generally be perceived to be higher than the effort to check the medication package. In case 3, $\bar{\epsilon} > \bar{\beta} > |\bar{\eta}|$, which represents products with lower perceived value and with relatively lower stakes from consuming counterfeits from time to time. Blood pressure medication might fall into this category, as it helps with the maintenance of a long-term chronic condition that is not immediately life-threatening, and missing a dosage here and there is not as devastating as it is in the case of TB. The inequalities represent the means of each distribution; individual values of each parameter still vary randomly from one consumer to the next.

The shape of the demand curve, with and without MAS, is presented for these three cases in Figures 5-11 to 5-16. Note that the vertical axes in these figures do not have identical ranges. In cases 1 and 2, the enabling product verification raises demand substantially even without incentives for checking. Incentives for checking raise demand in all three cases, though in case 2 demand is already close to 100% when checking is enabled.

The aggregate checking functions for the three cases are presented in Figures 5-17, 5-18, and 5-19. Checking increases sharply with both incentives and the introduction of counterfeits in case 2 (i.e., products like TB meds) but depends mainly on incentives in case 3 (i.e., products like BP meds).

5.7.1 Retailer Profit

The penalty faced by the retailer depends on the number of times they are caught selling fake products. Figures 5-20, 5-21, and 5-22 show the shape of the retailer penalty function in the three cases, with no incentive to check (green line) and with an incentive equal to 20% of the manufacturer's profit per product. The y-axis scale

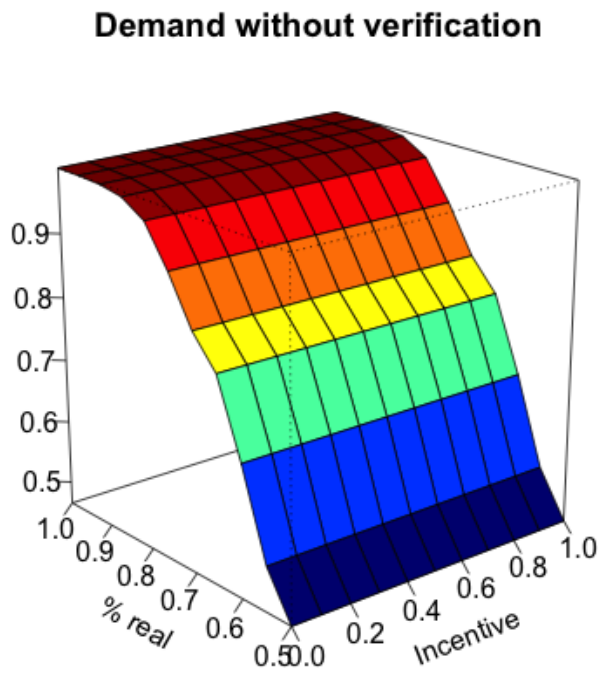

Figure 5-11: Case 1: demand without verification. $\bar{\beta} \approx |\bar{\eta}| \approx \bar{\epsilon}$

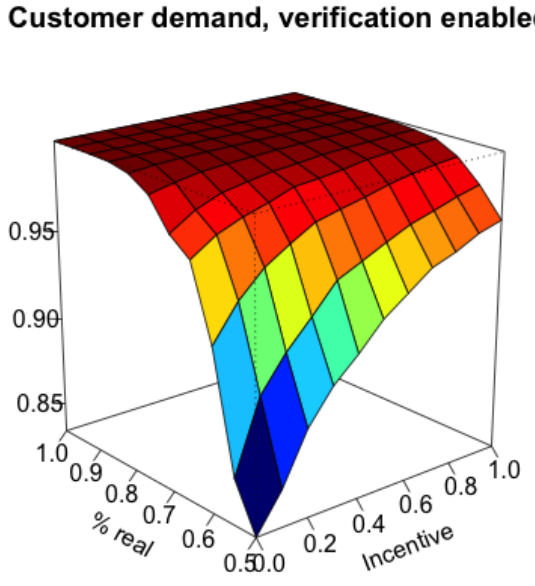

Figure 5-12: Case 1: demand with verification enabled. $\bar{\beta} \approx |\bar{\eta}| \approx \bar{\epsilon}$

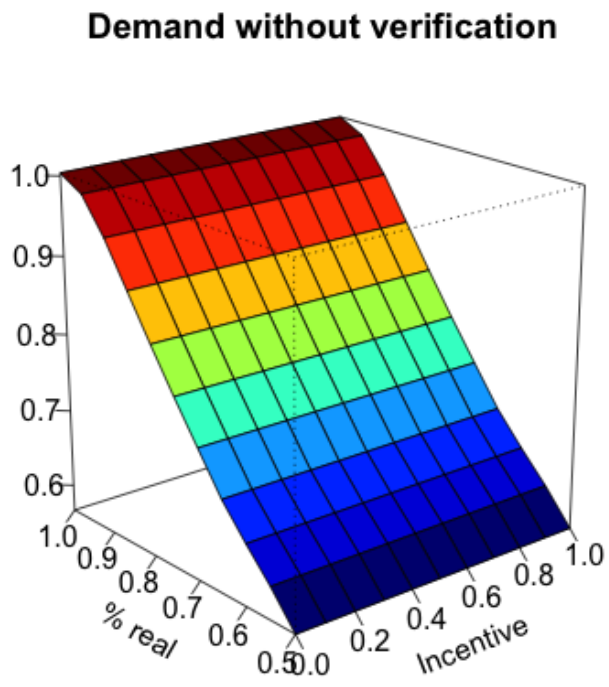

Figure 5-13: Case 2: demand without verification. $|\bar{\eta}| > \bar{\beta} > \bar{\epsilon}$

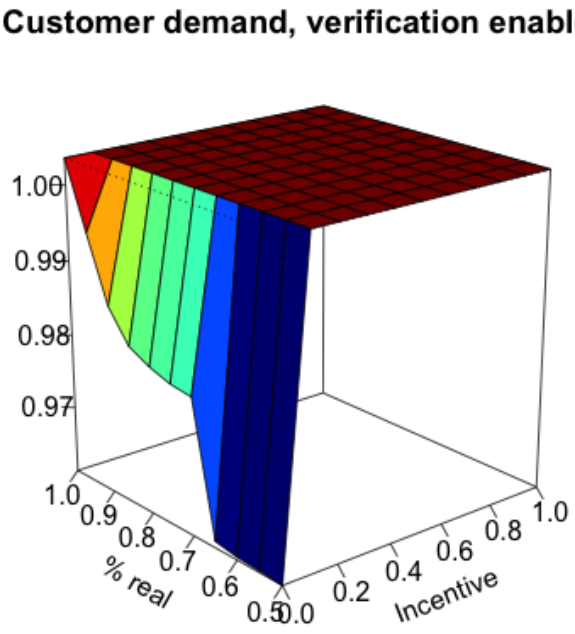

Figure 5-14: Case 2: demand with verification enabled. $|\bar{\eta}| > \bar{\beta} > \bar{\epsilon}$

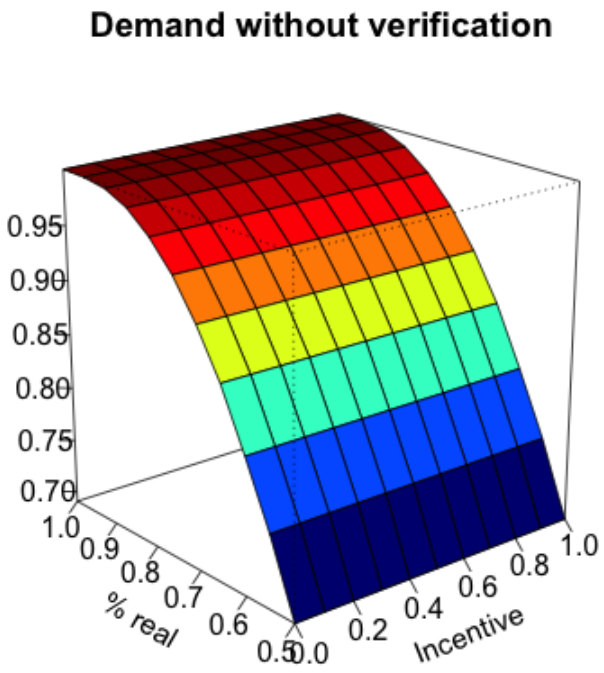

Figure 5-15: Case 3: demand without verification. $\bar{\epsilon} > \bar{\beta} > |\bar{\eta}|$

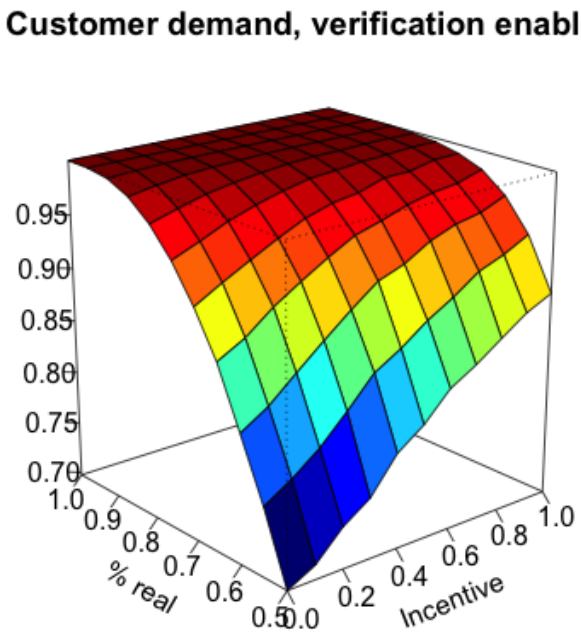

Figure 5-16: Case 3: demand with verification enabled. $\bar{\epsilon} > \bar{\beta} > |\bar{\eta}|$

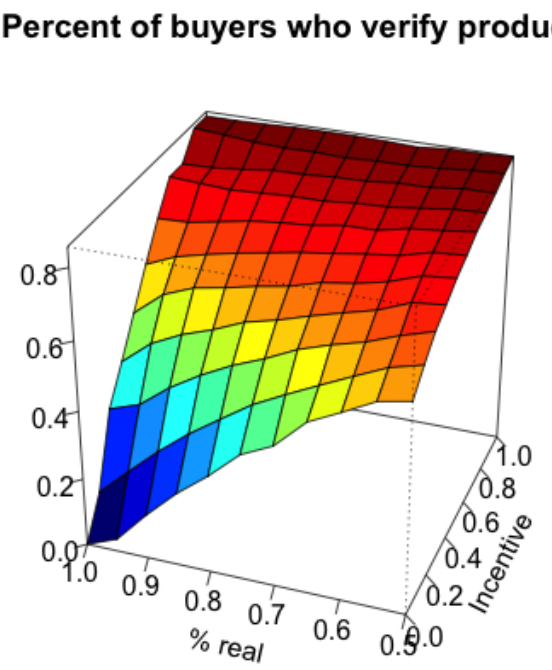

Figure 5-17: Case 1: percent of buyers who verify product. $\bar{\beta} \approx |\bar{\eta}| \approx \bar{\epsilon}$

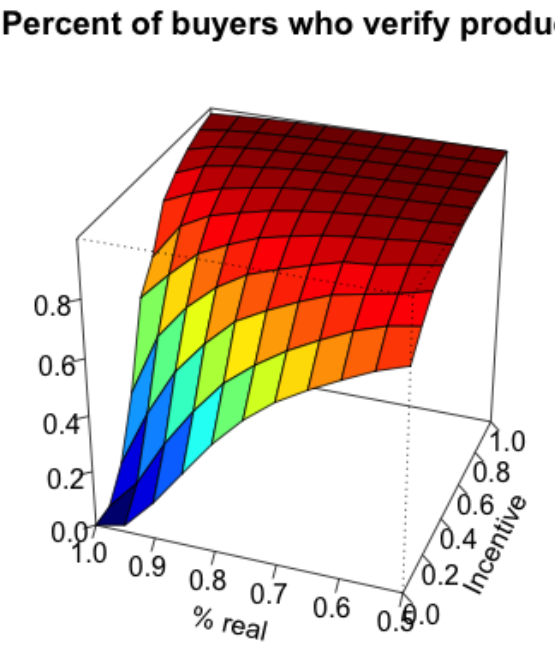

Figure 5-18: Case 2: percent of buyers who verify product. $|\bar{\eta}| > \bar{\beta} > \bar{\epsilon}$

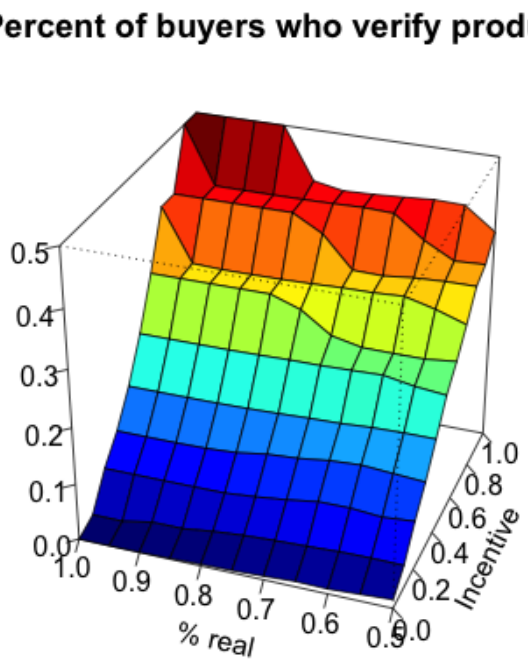

Figure 5-19: Case 3: percent of buyers who verify product. $\bar{\epsilon} > \bar{\beta} > |\bar{\eta}|$

showing the magnitude of the retailer's penalty is set such that 1 is the magnitude of the retailer's profit when it sells only authentic products. Note that the axis range is not the same in the three figures. The penalty for retailers is much higher in cases 1 and 2 because of high checking activity by customers, and an order of magnitude lower in case 3 even with the 20% checking incentive.

Under these penalty functions, the we can plot the retailer's profit as a function of θ in each scenario, and under varying incentive levels. Note that the dots are retailer's profit when no verification technology is enabled. As shown in Figures 5-23, 5-24, and 5-25, the profit functions are quasi-concave, so an unsophisticated profit-maximizing retailer could eventually arrive at the global optimum through small incremental greedy adjustments to θ. The lack of smoothness and of strict concavity in the curves is a result of the numerical approximations. We can observe that for case 2 (products like TB medication), the retailer's optimal θ increases substantially when verification is introduced. The increase is good but more modest in case 1, and small for case 3 until large incentives are introduced. The model suggests that a company considering adopting MAS but unwilling to take the risk of paying for an

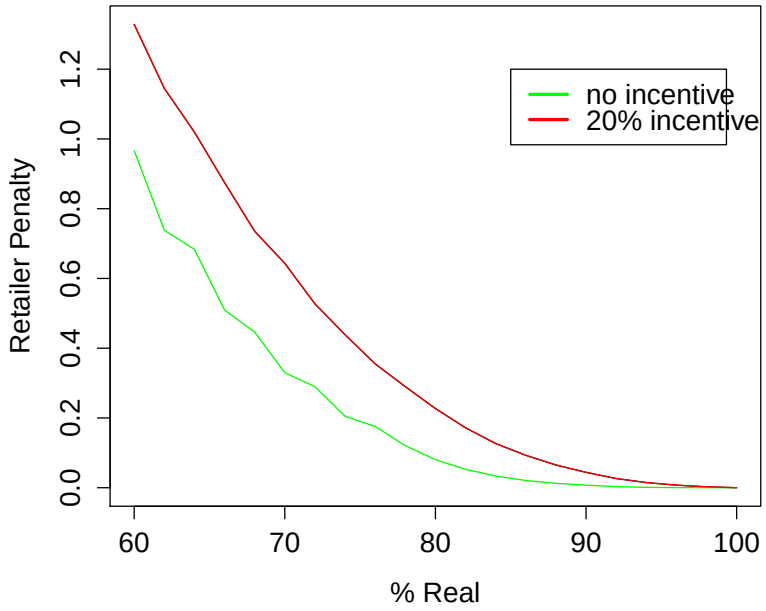

Figure 5-20: Case 1: retailer penalty for selling counterfeits as a function of θ.

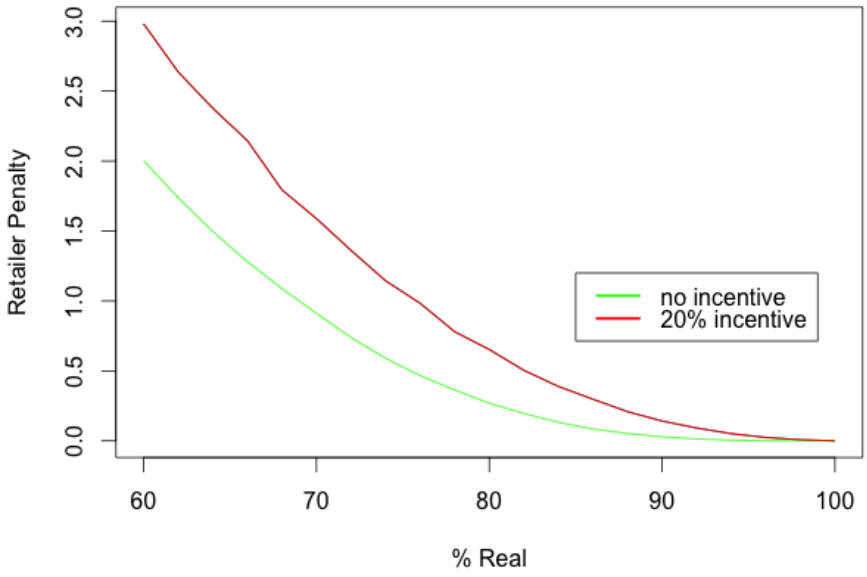

Figure 5-21: Case 2: retailer penalty for selling counterfeits as a function of θ.

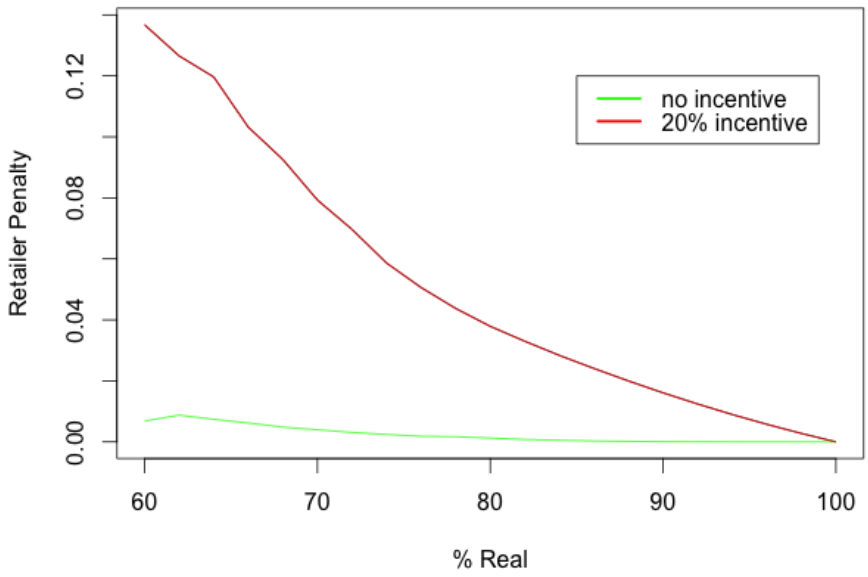

Figure 5-22: Case 3: retailer penalty for selling counterfeits as a function of θ.

incentives program may get the best initial results on products that consumers have naturally high motivations to check, such as TB medications.

To illustrate the effect of much steeper penalties or much higher profit margins on counterfeits on retailer profits, we have generated plots for each case covering these two scenarios plus the base case (Figures 5-26, 5-27, and 5-28). In each graph, the base case involves no financial incentive for product verification. In case 2, the retailer's decision is not sensitive to higher margins on counterfeits, presumably because of the demand-suppressing effect of higher counterfeit levels when η is high. In case 2, however, retailer profit is very sensitive to higher regulatory penalties, and the profit-maximizing authenticity rises substantially. In case 1, the base scenario and high-regulatory penalty scenario have flat profit curves around the optimum, so the set of strategies that retailers may perceive to be optimal might range along the flat portion of the curve. Stronger penalties unequivocally make some of the lower values of θ that appear near-optimal in the base scenario very unprofitable, so we might expect stronger penalties to also have good effects for products characterized

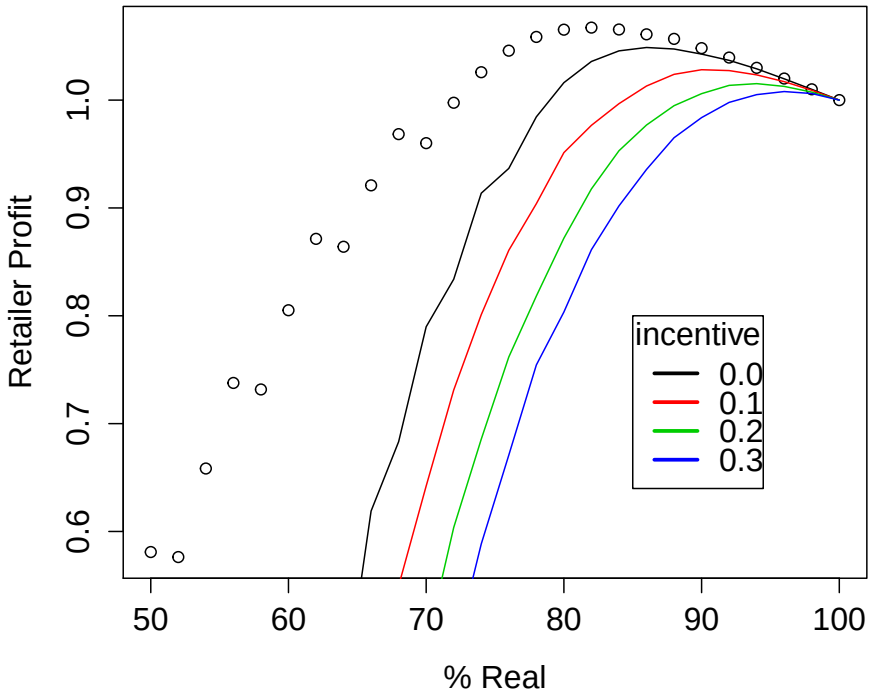

Figure 5-23: Case 1: retailer profit as a function of θ.

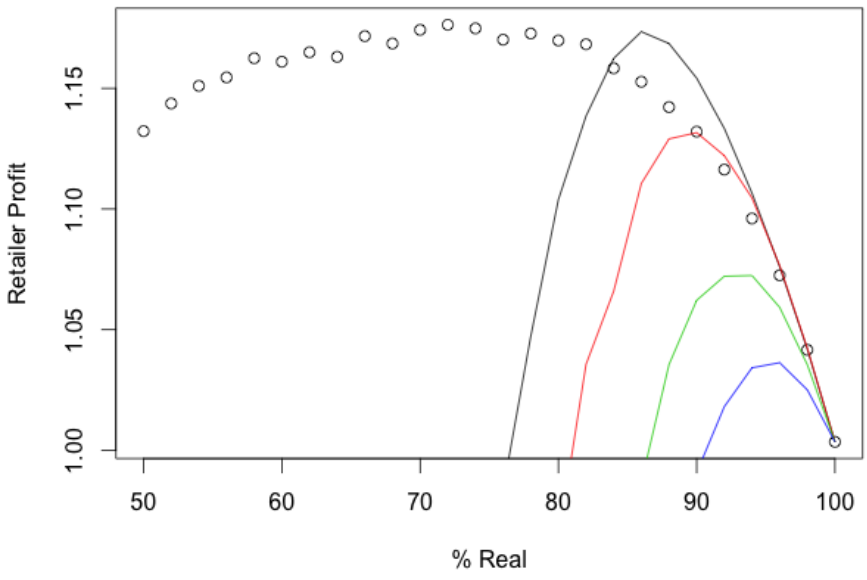

Figure 5-24: Case 2: retailer profit as a function of θ.

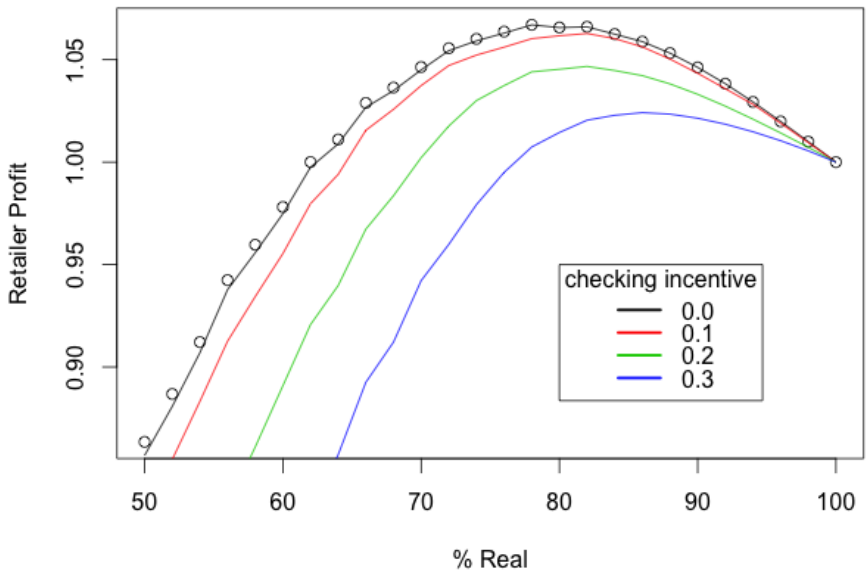

Figure 5-25: Case 3: retailer profit as a function of θ.

by case 1. Higher regulatory penalties had no effect in case 3. Overall, this suggests that law-enforcement approaches targeting and steeply penalizing retailers caught selling counterfeits will have more effect in markets for products that consumers are reasonably likely to check. By contrast, in cases 1 and 3, higher profit margins on counterfeits greatly boosted overall retailer profit and counterfeit prevalence. In these cases, law-enforcement strategies targeting the upstream the suppliers/manufacturers of counterfeits and their distribution chains may be more effective at reducing counterfeit procurement, at least to the extent that such activity can drive up the cost of counterfeits to retailers and thereby reduce π_f.

From the perspective of the manufacturer, what is of particular interest is the retailer's best response θ given the manufacturer's choice of x_m. In our simulations, θ improves in all scenarios as x_m, with one curious exception around $x_m = .6$ for case 3 when profit margins on fakes were four times greater than real products. One possible explanation is that the checking incentive is increasing the number of consumers willing to buy, but that these marginal consumers are driven more by the

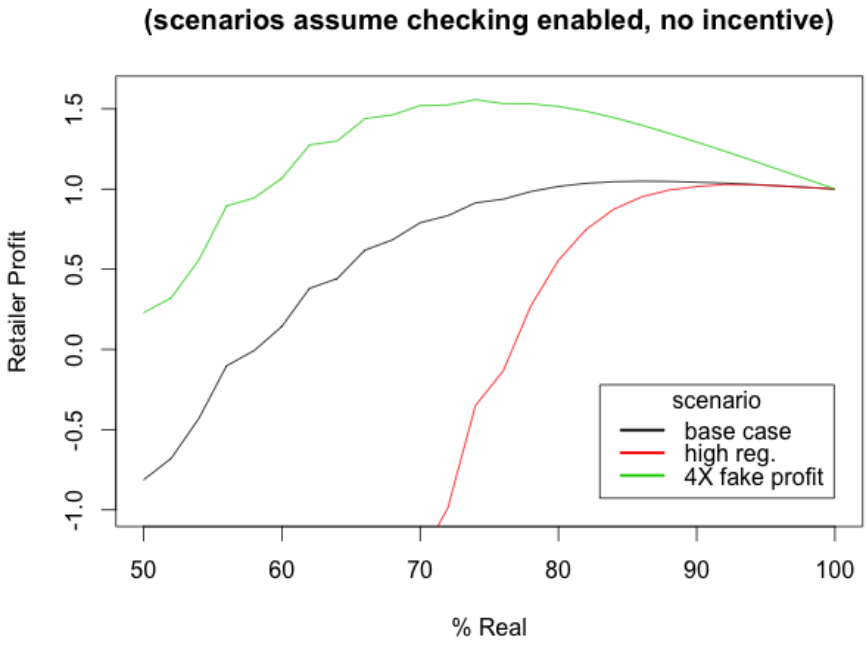

Figure 5-26: Case 1: retailer profit as a function of θ, special scenarios.

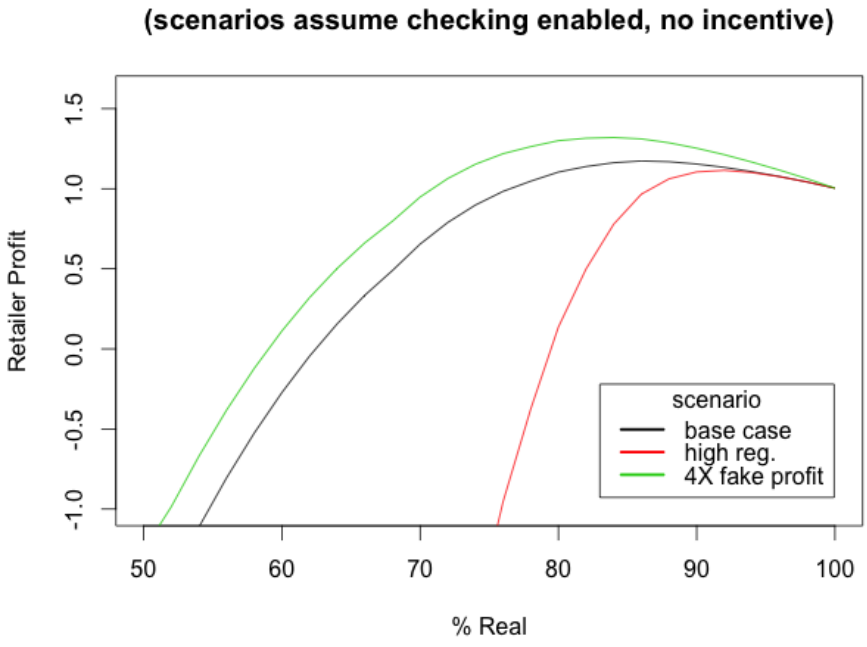

Figure 5-27: Case 2: retailer profit as a function of θ, special scenarios.

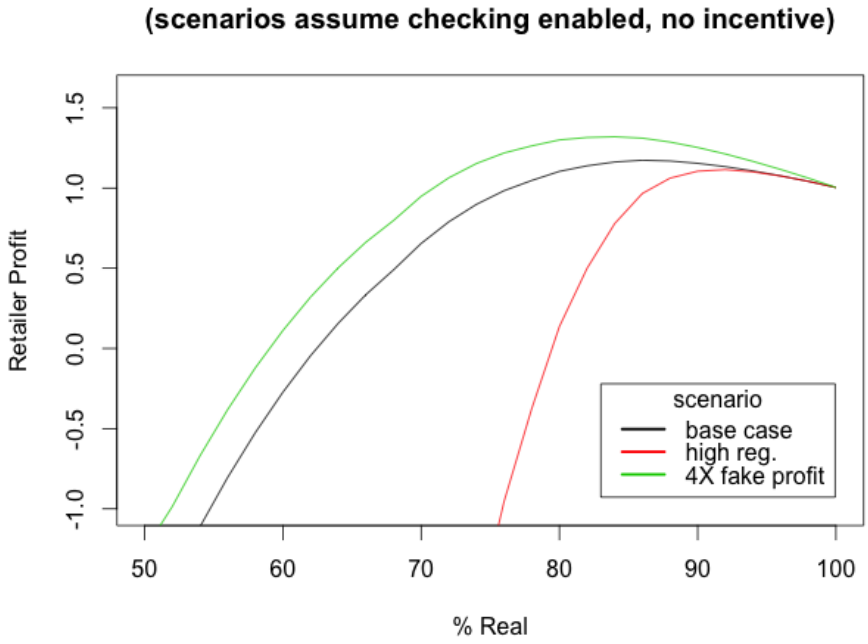

Figure 5-28: Case 3: retailer profit as a function of θ, special scenarios.

lottery-effect than a reduction in the risk from η, and so they may be more willing to participate in a market with counterfeits than the base consumers who purchased even without incentives. This may not be the mechanism at play, and is offered simply to explain that there is not a mathematical requirement that the best-response curve be monotone increasing in x_m.

5.7.2 Manufacturer Profit

We plot the manufacturer's profit vs. the manufacturer's chosen incentive level for consumer product verification. Profit is scaled so that 1 represents the maximum profit attainable if retailers procure no counterfeits. Figures 5-29, 5-30, and 5-31 show the results for our three cases, including the high-penalty and high π_f scenarios. For cases 1 and 2, an incentive in the range of 0% to 20% seems to have a slightly positive or neutral effect on manufacturer profit. Given the impact on public health and consumer welfare, a case could be made for small government subsidization or incentives for manufacturers who adopt product verification technologies; these sub-

sidies could be made temporary in order to merely reduce the risk of incentivizing MAS usage by consumers, as MAS incentives may be profit-enhancing in certain markets and product lines. Interestingly, in the case 3 when π_f is high, the manufacturer can afford to give away very large incentives to customers who check because a large proportion of potential customers will be more willing to buy the product when counterfeit rates are lowered by those who do check, but they themselves will not be bothered to check the product and claim the incentives.

In all cases, the manufacturer's profit from adopting verification technology, even without incentives, was larger (by 9%-19%) than without the technology (see Figures 5-32, 5-33, and 5-34). These figures are before factoring in the cost of the technology itself, which obviously needs to be weighed against the benefits. That said, the cost per product of a covert scratch code and maintaining database that can be checked by mobile phone SMS is not very high. The plots suggest that a manufacturer interested in adopting MAS might have a very large margin of error in their choice of x_m, as many values of x_m will still leave them better off than if they had not adopted MAS in the first place. The increase in the manufacturer's profit by adopting MAS comes both from increased consumer demand and recapturing market share from counterfeiters.

These graphs and the others presented in this chapter are meant to illustrate the range, variability, and consistency of possible outcomes for different product classes and markets. They are not meant to suggest a particular precise outcome for specific products or countries, as doing so would require much more empirical work to estimate product-specific and market-specific parameters. The parameters and code used in our numerical simulations will be posted to GitHub by October 30, 2019 for those wishing to examine our parameter assumptions, replicate the results, or test other scenarios https://github.com/MichaelBeeler?tab=projects. We chose not to list in the body of the chapter the particular parameters used to avoid the impression that we are giving particular results for specific products or market types. The parameters were chosen to test a range of circumstances so that consistent model behavior would most likely not be an idiosyncratic artefact of arbitrary model assumptions. We aimed to show the scope of possibilities and offer a starting point for further research.

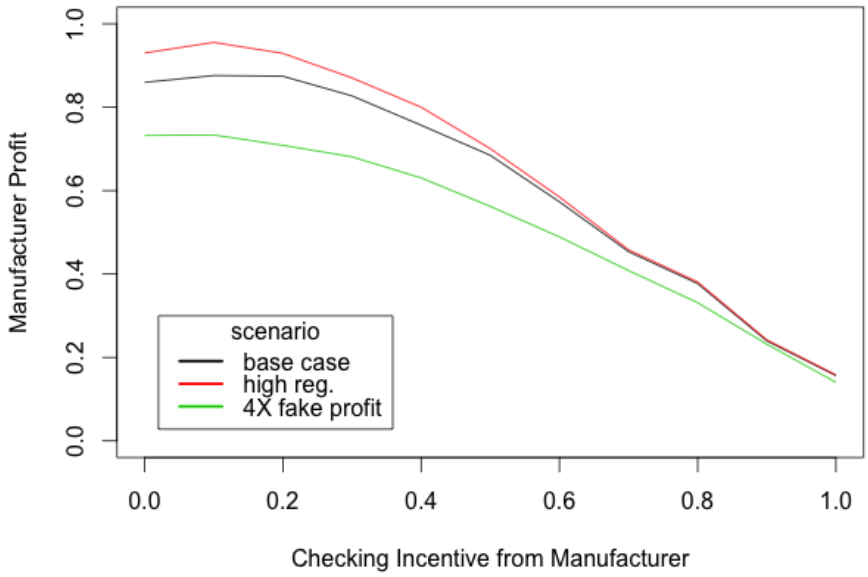

Figure 5-29: Case 1: manufacturer profit vs incentive level for checking.

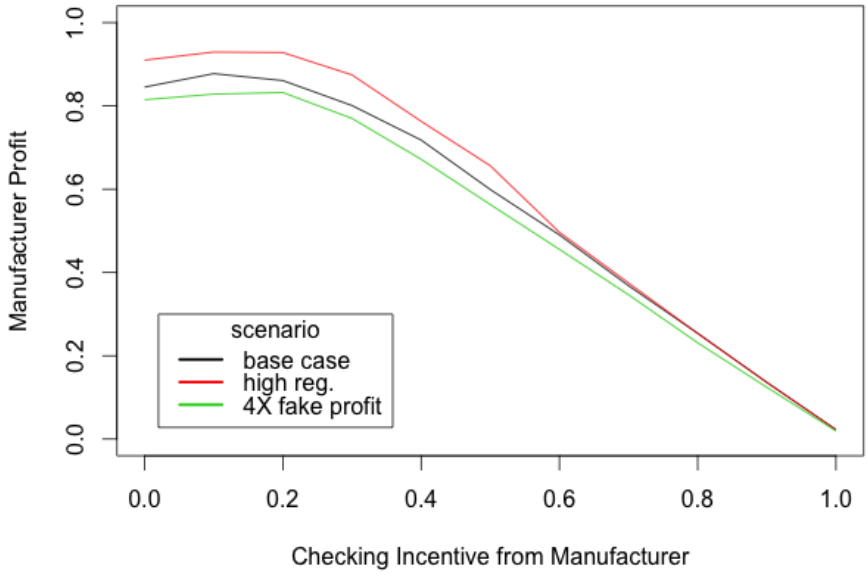

Figure 5-30: Case 2: manufacturer profit vs incentive level for checking.

Figure 5-31: Case 3: manufacturer profit vs incentive level for checking.

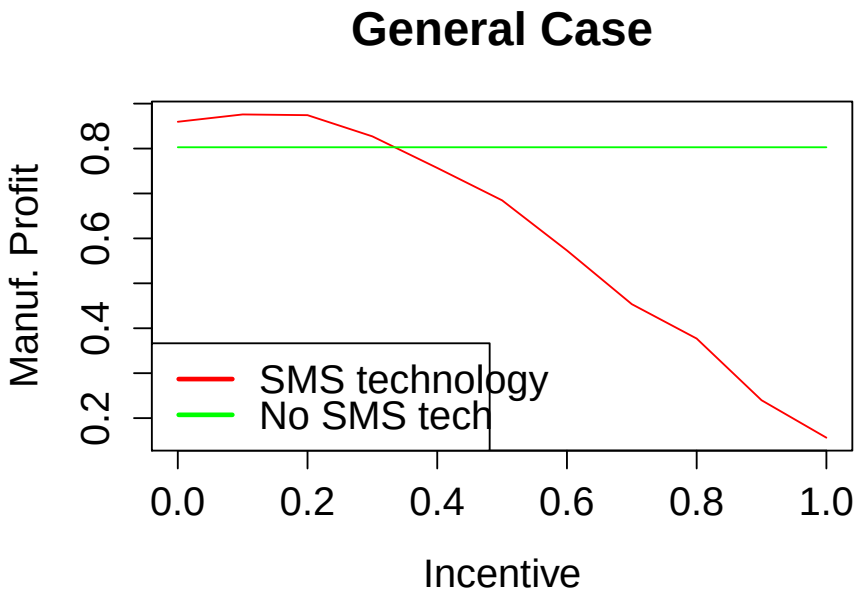

Figure 5-32: Case 1: manufacturer profit vs incentive level for checking. Horizontal line is profit without verification technology.

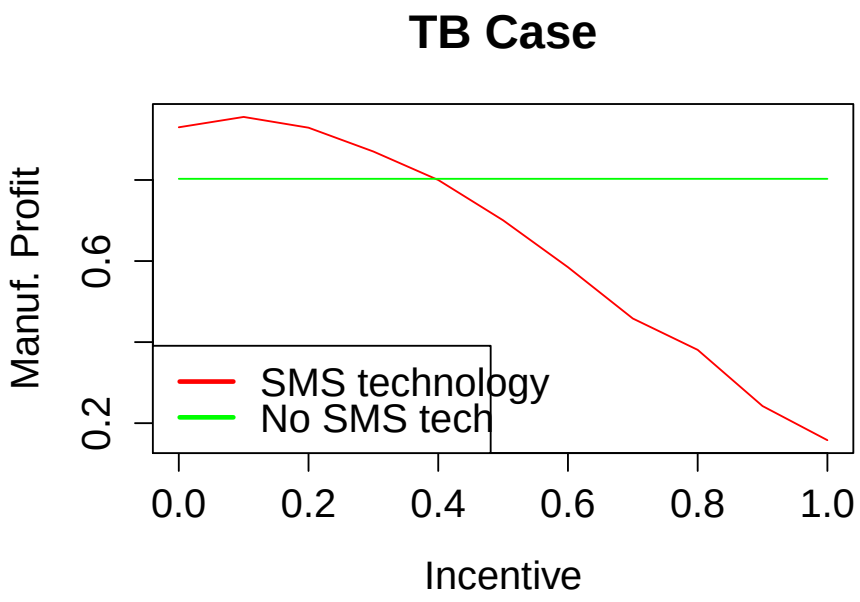

Figure 5-33: Case 2: manufacturer profit vs incentive level for checking. Horizontal line is profit without verification technology.

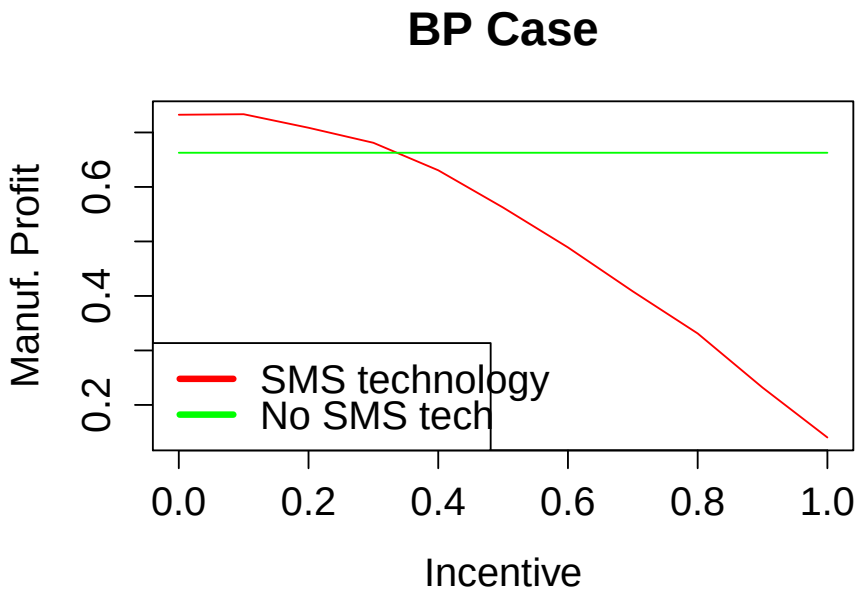

Figure 5-34: Case 3: manufacturer profit vs incentive level for checking. Horizontal line is profit without verification technology.

5.8 Conclusion

The models presented in this chapter provide a range of analytical and simulation tools for understanding how governments and firms could reduce procurement and sale of counterfeit products to consumers. Both the analytical and simulation sections' models suggest that the equilibrium prevalence of counterfeits depends greatly on the sensitivity of consumer demand to counterfeits. Our models also suggest that the relative effectiveness of law enforcement / discovery action against retailers vs action against manufacturers/suppliers of counterfeits may vary greatly depending on the profit-margins on counterfeits and whether consumers are naturally motivated to check for counterfeits and punish or avoid retailers who sell them fake goods.

The simulation model is the first quantitative study of the potential impact of MAS on the profit of brands targeted by counterfeiters. The evidence, coming from simulated agents in a complex game theoretic model, is necessarily inconclusive; it should be viewed as suggestive and preliminary. Our simulations found that MAS was profit-enhancing across all simulated scenarios. We also found that large incentives for checking could be offered before profit would fall to pre-MAS levels. Given the high stakes for consumers and public health and manageable costs (if not outright profit improvement) to manufacturers/brands, governments may want to consider mandating and/or subsidizing MAS combined with programs to incentivize consumer checking.

Chapter 6

Conclusion and Future Work

This thesis examined three problems of broad public interest in a manner that oriented the lines of inquiry and model development around concerns and challenges of particular relevance in low-income countries. In each chapter, the goal was not to present an idiosyncratic solution to a particular problem, but to develop general conceptual frameworks or methodological improvements to aid decision-making and planning in areas that have received less attention from the operations research and management science profession. While the work presented makes progress towards this goal, there is certainly much more to be done on each problem. Whereas the individual chapter conclusions summarized the main contributions of their respective chapters, this conclusion will discuss the most promising directions for future work.

The education technology sector is undergoing rapid growth. The amount and quality of data available to educational software platforms will generate more statistical power than has been traditionally available; this trend could support more sophisticated diagnostic models, including models for complex causality, as contemplated in Chapter 2, Section 2.8. Such models raise interesting methodological questions that could be answered using synthetic data if sufficiently large datasets are not publicly available: what algorithms are best suited to fit complex causal models with non-convex loss functions? How much data is needed, and how varied does the data need to be, to recover the functional form? Given the correct functional form, how much data is needed to precisely estimate the model parameters?

Another direction for future work would be to identify other applications of the partial shuffling training method developed for deep knowledge tracing. The shuffling method is an example of "data augmentation," whereby new samples were generated stochastically from slightly altered existing samples, but in a way that seeks to preserve underlying properties of the data sequence. In this case, the underlying property was gradual learning, which is why we shuffled within contiguous partitions. In other time-series problems, there will be different, and perhaps stricter, constraints on how data can be shuffled without violating causal principles or producing nonsensical training samples. However, to the extent that such data augmentation is possible, it can help the training process find a more robust model that will perform better on unseen sequences. It would seem this method would be more likely to succeed on quantitative time series data generated from either a stationary process or one that evolves slowly. It would seem less likely to succeed on many natural language processing applications, since slight permutations in word order can change meaning or produce incoherencies.

The chapters on network infrastructure cost allocation and pricing are admittedly theoretical and abstract. To increase the relevance of this work to practitioners, the general insights and high-level recommendations ought to be tested through extensive policy simulation using the geographic coordinates of buildings from a large sample of villages and towns that lack electricity. We have sourced 40 such datasets from http://OpenStreetMap.org, including towns from various countries in sub-Saharan Africa and South Asia, which we will use to evaluate various cost allocation policies in our future work. If that simulation-based policy evaluation proves fruitful, then it may make sense to incorporate into the simulation system-specific costs at a more granular level, such as the cost of the initial power generation infrastructure, transmission losses, etc., and to test these policies against actual pricing practices of surveyed microgrid operators as a benchmark.

The anti-counterfeiting chapter provided a flexible framework for modeling the equilibrium counterfeit procurement decisions of strategic, profit-maximizing retailers in a variety of settings. The parameters that define the consumer's utility model are

amenable to empirical estimation through observed behavior and/or surveys, though the results of survey instruments may be sensitive to how losses and gains are framed. Given the severity of the problem of pharmaceutical counterfeiting, empirical work to get parameters for the consumer utility model and for the retailer penalty function in different markets may be well worth the cost. Estimates of θ obtained by sampling various product lines and a variety of retailers could be compared with predicted θ as a means of model calibration and validation. Even if the predicted θ has a high error, the directional impact of various covariates and reputation-sharing conditions can still be tested through such empirical research. It would also be important to establish how variable θ is amongst retailers in the same market, as the behavior of dishonest retailers may depend greatly on the prevalence of honest or myopic retailers.

Future research should also consider several angles that were briefly mentioned in the anti-counterfeiting chapter. How informed are consumers about counterfeits and how widely do consumer beliefs vary? Does increased consumer awareness of counterfeits risk hurting demand in a way that could create incentives for brands to conceal the extent of the problem? Could consumers' subjective perceptions of disutility from counterfeits be influenced by mobile authentication services simply by being reminded of the problem around the time of a purchase decision? With trillions of dollars and millions of lives at stake each year, these and other questions whose answers could improve responses to global counterfeiting deserve further attention.

Appendix A

Optimal Individual Prices for Potential Network Users

With reference to Chapter 4, Section 4.5.1, the network monopolist could face the problem of assigning different prices to different users based on their distance from the source node on the 1D lattice. This appendix presents the expected profit maximization function, its gradient, and Hessian. The function is shown not to be quasi-concave. As a result, a stationary point reached through gradient-based methods will not necessarily be a global optimum.

Define π_i as the price offered to the customer i positions from the source node. Define $V^*_{n,\text{diff},F}$ as the optimal profit when there are n potential users who can be offered different prices, and reference prices follow distribution F. The problem of choosing optimal π^*_i for all users to maximize expected profit is neither concave nor pseudo-concave, even for the simple uniform case where $r_i \sim U[0,1]$. Scale the currency units so that $\sup\{r\} = 1$, causing c to the marginal network extension cost as a fraction of the largest possible reference price.

$$V^*_{n,\text{diff},F} \equiv \max_{\pi_1,\ldots,\pi_n,} (E\left[\sum_{i=1}^{n} \pi_i \mathbb{1}_{\pi_i \leq r_i} - c \max_{j:\pi_j \leq r_j} j\right], 0)$$

$$\iff \max_{\pi_1,\ldots,\pi_n} \left(\left(\sum_{i=1}^{n} \pi_i F^c(\pi_i) - c \sum_{i=1}^{n-1} i F^c(\pi_i) \prod_{j=i+1}^{n} F(\pi_j) - cn F^c(\pi_n) \right), 0 \right)$$

For $r_i \sim U[0,1], F(\pi_i) = \pi_i$ and $F^c(\pi_i) = 1 - \pi_i,$ for $\pi_i \in [0,1],$ hence

$$V^*_{n,\text{diff},U} = \max_{\pi_1,\ldots,\pi_n \in [0,1]} \left(\sum_{i=1}^{n} \pi_i (1-\pi_i) - c \sum_{i=1}^{(n-1)} i(1-\pi_i) \prod_{j=i+1}^{n} \pi_j - cn(1-\pi_n) \right),$$

noting that 0 can be attained by $\pi = 1$.

Unfortunately, $V^*_{n,\text{diff},F}$ is not a concave function; moreover, the term in it's middle summand, $-i(1-\pi_i)\prod_{j=i+1}^{n}\pi_j$ in $-c\sum_{i=1}^{(n-1)} i(1-\pi_i)\prod_{j=i+1}^{n}\pi_j$, is not even a quasi-concave function. The lack of quasi-concavity, which implies a lack of concavity, can be shown by example. The following two pricing policies can be used for verification by the reader: $\boldsymbol{\pi_1} = (0.5, 1, 0.5, 1, 0.5, 1, 0.5, ...)$ and $\boldsymbol{\pi_1} = (1, 0.5, 1, 0.5, 1, 0.5, 1, ...)$ compared to their average $\bar{\boldsymbol{\pi}} = (0.75, 0.75, 0.75, ...)$. As a result, a stationary point found through gradient-based methods cannot assumed to be a maximum, let alone a global maximum. For reference to a reader who may nonetheless wish to try gradient-

based methods, the elements of the gradient and Hessian are provided below:

$$\forall i \in \{1, ..., n-1\} \text{ we have } \nabla_i V^*_{n,\text{diff},U} = -2\pi_i + 1 + ci \prod_{j>i} \pi_j - c \sum_{j<i} j(1-\pi_j) \prod_{\substack{k>j \\ k\neq n}} \pi_k$$

$$\text{and } \nabla_{i=n} V^*_{n,\text{diff},U} = -2\pi_i + 1 + cn - c \sum_{j<n} j(1-\pi_j) \prod_{\substack{k>j \\ k\neq n}} \pi_k$$

$$\Rightarrow \forall i < j, \nabla_{ij} V^*_{n,\text{diff},U} = ci \prod_{\substack{k>i \\ k\neq j}} \pi_k - c \sum_{k<i} k(1-\pi_k) \prod_{\substack{\ell>k \\ \ell\neq i \\ \ell\neq j}} \pi_\ell; \text{ and } \forall i, \nabla_{ii} V^*_{n,\text{diff},U} = -2.$$

Note: the Hessian matrix is symmetric with diagonal entries of -2.

Bibliography

P. Ambroise-Thomas. The tragedy caused by fake antimalarial drugs. *Mediterranean Journal of Hematology and Infectious Diseases*, 4(1), 2012.

E. B. Anderson. The numerical solution of a set of conditional estimation equations. *Journal of the Royal Statistical Society, Series B*, 34:42–54, 1972.

P. Anselmi, E. Robusto, L. Stefanutti, and D. de Chuisole. An upgrading procedure for adaptive assessment of knowledge. *Psychometrica*, 81(2):461–482, 2016.

F. B. Baker and S.-H. Kim. *Item Response Theory: Parameter Estimation Techniques, Second Edition*. CRC Press, Boca Raton, 2004.

J. Banzhaf III. Weighted voting does not work: a mathematical analysis. *Rugters Law Review*, 19:317–343, 1965.

M. A. Barton and F. M. Lord. An upper asymptote for the three-parameter logistic item-response model. *Educational Testing Service Research Bulletin*, 81, 1981.

M. Beeler. Interview with aleks scientist at edm conference, buffalo, ny. Personal Communication, 2018.

A. Birnbaum. Efficient design and use of tests of a mental ability for various decision-making problems. series report no. 58-16. Technical report, Randolph Air Force Base, TX: USAF School of Aviation Medicine, 1957.

A. Birnbaum. *Some latent trait models and their use in inferring an examinees's ability*, pages 397–479. Addison-Wesley, Reading, NA, 1968.

R. Bkakar, S. V.S., N. Padhy, , and H. Gupta. Probabilistic game approaches for network cost allocation. *IEEE Transactions on Power Systems*, 25(1):51–58, 2 2010.

B. Bloom. The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational Researcher*, 13:4–16, 1984.

R. D. Bock and M. Aitkin. Marginal maximum likelihood estimation of item parameters: Application of an em algorithm. *Psychometrika*, 46(4):443–459, Dec 1981. ISSN 1860-0980. doi: 10.1007/BF02293801.

H. Cen, K. Koedinger, and B. Junker. Learning factors analysis – a general method for cognitive model evaluation and improvement. In M. Ikeda, K. D. Ashley, and T.-W. Chan, editors, *Intelligent Tutoring Systems*, pages 164–175. Springer Berlin Heidelberg, 2006. ISBN 978-3-540-35160-3.

H. Cen, K. Koedinger, and B. Junker. Comparing two irt models for conjunctive skills. *ITS 2008. LNCS*, 5091:796–798, 2008.

M. Chi, K. Koedinger, G. Gordon, P. Jordan, and K. VanLehn. Instructional factors analysis: A cognitive model for multiple instructional interventions. In M. Pechenizkiy, T. Calders, C. Conati, S. Ventura, C. Romero, and J. Stamper, editors, *Proceedings of the 4th Intl. Conf. on Educational Data Mining (EDM)*, pages 2984–2991, july 2011.

S.-H. Cho, X. Fang, and S. Tayur. Combating strategic counterfeiters in licit and illicit supply chains. *Manufacturing & Service Operations Management*, 17(3), 2015.

P. A. Cohen, J. A. Kulik, and C. L. C. Kulik. Educational outcomes of tutoring: a meta-analysis of findings. *American Educational Research Journal*, 19:237–248, 1982.

V. Conitzer and T. Sandholm. Computing shapley values, manipulation value division schemes, and checking core membership in multi-issue domains. In *Proc. of the National Conference on Artificial Intelligence*, AAAI'04. AAAI Press, July 2004. ISBN 0-262-51183-5. [Accessed 27-June-2019].

A. T. Corbett and J. R. Anderson. Knowledge tracing: Modeling the acquisition of procedural knowledge. *User modeling and user-adapted interaction*, 4:253–278, 1994.

D. Cox. *The Theory of Stochastic Processes*, chapter 5.2, page 203. Routledge, 1st edition edition, 1977.

J. de la Torre. The generalized dina model framework. *Psychometrika*, 76:179–199, 04 2011. doi: 10.1007/s11336-011-9207-7.

X. Deng and C. H. Papadimitriou. On the complexity of cooperative solution concepts. *Mathematics of Operations Research*, 19(2):257–266, 1994. ISSN 0364765X, 15265471. [Accessed 27-June-2019].

J.-P. Doignon and J.-C. Falmagne. Spaces for the assessment of knowledge. *International Journal of Man-Machine Studies*, 23:175–196, 1985.

J.-P. Doignon and J.-C. Falmagne. *Knowledge Spaces*. Springer-Verlag, 1999.

S. Doroudi and E. Brunskill. The misidentified identifiability problem of bayesian knowledge tracing. In X. Hu, T. Barnes, A. Hershkovitz, and L. Paquette, editors, *Proceedings of the 10th International Conference on Educational Data Mining*, pages 143–149. International Educational Data Mining Society, July 2016.

S. Embretson. A multicomponent latent trait model for diagnosis. *Psychometrika*, 78 (1):14–36, 2013. doi: 10.1007/s11336-012-9296-y.

S. Embretson and X. Yang. The multicomponent latent trait model for diagnosis: Applications to heterogeneous test domain. *Applied Psychological Measurement*, 39(1):16–30, 2015.

S. E. Embretson. A general multicomponent latent trait model for response processes. *Psychometrika*, 49:175–186, 1984.

J. Falmagne. A latent trait theory via stochastic learning theory for a knowledge space. *Psychometrika*, 54:283–303, 1989.

J. Falmagne and J.-P. Doignon. A class of stochastic procedures for the assessment of knowledge. *British Journal of Mathematical and Statistical Psychology*, 41:1–23, 1988a.

J. Falmagne and J.-P. Doignon. A markovian procedure for assessing the state of a system. *Journal of Mathematical Psychology*, 32:232–258, 1988b.

J.-C. Falmagne, M. Koppen, M. Villano, J. Doignon, and L. Johannesen. Introduction to knowledge spaces: How to build test and search them. *Psychological Review*, 97: 201–224, 1990a.

J.-C. Falmagne, M. Koppen, M. Villano, J.-P. Doignon, and L. Johannesen. Introduction to knowledge spaces: How to build, test, and search them. *Psychological review*, 97(2):201–224, 1990b.

Fight the Fakes. Mobile Verification of Medicines comes to Saudi Arabia. http://fightthefakes.org/updates/mobile-verification-of-medicines-comes-to-saudi-arabia/, 2015.

S. Finch. How far might we walk at random?, 2018. https://arxiv.org/pdf/1802.04615.pdf.

G. H. Fischer. Conditional maximum-likelihood estimations of item parameters for a linear logistic test model. Research Bulletin 9, University of Vienna, Psychological Institute, 1972.

A. Fitzpatrick. Do informed consumers reduce the price and prevalence of counterfeit drugs? evidence from the antimalarial market. http://www-personal.umich.edu/~fitza/Fitzpatrick_jmp.pdf, 2015.

M. Ghayeni and R. Ghazi. Transmission network cost allocation with nodal pricing approach based on ramsey pricing concept. *IET Generation, Transmission & Distribution*, 5(3):4568–4577, 2 2011. [Accessed 27-June-2019].

H. Gil, F. Galiana, , and E. da Silva. Nodal price control: a mechanism for transmission network cost allocation. *IEEE Transactions on Power Systems*, 21(1):3–10, 2 2006. [Accessed 27-June-2019].

D. Granot, M. Maschler, G. Owen, and W. R. Zhu. The kernel/nucleolus of a standard tree game. *Int. J. Game Theory*, 25(2):219–244, Jan. 1996. ISSN 0020-7276. doi: 10.1007/BF01247104. [Accessed 27-June-2019].

D. Granot, J. Kuipers, and C. S. Cost allocation for a tree network with heterogeneous customers. *International Journal of Game Theory*, 27(4):647–661, 2002. [Accessed 27-June-2019].

A. Graves, A. rahman Mohamed, and G. Hinton. Speech recognition with deep recurrent neural networks. In *Proceedings of 2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, page 6645–6649, Vancouver, Canada, 2013.

R. T. Green and T. Smith. Executive insights: countering brand counterfeiters. *Journal of International Marketing*, 10(4):89–106, 2002.

J. Hartung, G. Knapp, and B. K. Sinha. *Statistical Meta-Analysis with Applications*. John Wiley & Sons, 2008.

J. Heller, L. Stefanutti, P. Anselmi, and E. Robusto. On the link between cognitive diagnostic models and knowledge space theory. *Psychometrika*, 80(4):995–1019, Dec. 2015.

R. A. Henson, J. L. Templin, and J. T. Willse. Defining a family of cognitive diagnostic models using log-linear models with latent variables. *Psychometrika*, 74(2):191–210, June 2009.

C. Hockemeyer. Bibliography on knowledge spaces. http://kst.hockemeyer.at/kst-bib.html, 2018. [Accessed 2018-August-31].

IPN. IPN Press Release: Fake drugs kill over 700,000 people every year - new report. https://archive.is/ipW8i, 2011.

M. Junqueira, L. da Costa, L. Barroso, G. Oliveira, L. Thome, and M. Pereira. An aumann-shapley approach to allocate transmission service cost among network users in electricity markets. *IEEE Transactions on Power Systems*, 22(4):1532–1546, 11 2007. [Accessed 27-June-2019].

K. Karunamoorthi. The counterfeit anti-malarial is a crime against humanity: a systematic review of the scientific evidence. *Malaria Journal*, 13, 2014.

M. Khajah, R. V. Lindsey, and M. C. Mozer. How deep is knowledge tracing? In *Proc. of the 9th Conference on Educational Data Mining EDM 2016 (Raleigh, NC, USA)*, volume abs/1604.02416, pages 94–101. CoRR (Computing Research Repository), 2016. [accessed 10-June-2019].

D. P. Kingma and J. L. Ba. Adam: a method for stochastic optimization. In *The International Conference on Learning Representations (San Diego, USA)*, 2015.

M. Koppen. Extracting human expertise for constructing knowledge spaces: An algorithm. *Journal of Mathematical Psychology*, 37:1–20, 1993.

M. Koppen and J.-P. Doignon. How to build a knowledge space by querying an expert. *Journal of Mathematical Psychology*, 34:311–331, 1990.

W.-W. Liao, R.-G. Ho, Y. C. Yen, and H. C. Cheng. The four-parameter logistic item response theory model as a robust method of estimating ability despite aberrant responses. *Social Behavior and Personality: an international journal*, 40, 11 2012. doi: 10.2224/sbp.2012.40.10.1679.

K. Liu, J. Li, Y. Wu, and K. K. Lai. Analysis of monitoring and limiting of commercial cheating: a newsvendor model. *Journal of the Operational Research Society*, 56(7): 844–854, 2005.

J. Mann. Distribution of the maximum of a (infinite) random walk. Mathematics Stack Exchange, 2017. https://math.stackexchange.com/q/2148246 (version: 2017-02-17), author page: (https://math.stackexchange.com/users/126754/jonmann).

T. Mikolov, M. Karafiat, L. Burget, J. H. Cernock, and S. Khudanpur. Recurrent neural network based language model. In *INTERSPEECH 2010, 11th Annual Conference of the International Speech Communication Association*, pages 1045–1048, Makuhari, Japan, 2010.

S. Minn, Y. Yu, M. C. Desmarais, F. Zhu, and J.-J. Vie. Deep knowledge tracing and dynamic student classification for knowledge tracing. https://arxiv.org/abs/1809.08713, 2018. [Accessed 10-June-2019].

Y. Molina, O. Saavedra, and H. Amaris. Transmission network cost allocation based on circuit theory and the aumann-shapley method. *IEEE Transactions on Power Systems*, 28(4):4568–4577, 11 2013. [Accessed 27-June-2019].

M. P. R. Ortega, J. I. Pérez-Arriaga, J. R. Abbad, and J. P. González. Distribution network tariffs: A closed question? *Energy Policy*, 36(5):1712 – 1725, 2008. ISSN 0301-4215. doi: https://doi.org/10.1016/j.enpol.2008.01.025. [Accessed 27-June-2019].

Z. A. Pardos and N. T. Heffernan. Modeling individualization in a bayesian networks implementation of knowledge tracing. In *Proceedings of the 18th International Conference on User Modeling, Adaptation, and Personalization*, pages 255–266, june 2010. doi: 10.1007/978-3-642-13470-8_24.

C. Piech, J. Spencer, J. Huang, S. Ganguli, M. Sahami, L. J. Guibas, and J. Sohl-Dickstein. Deep knowledge tracing. *CoRR*, abs/1506.05908, 2015.

Y. Qian. Impacts of entry by counterfeiters. *Quarterly Journal of Economics*, 123 (4):1577–1609, 2008.

E. Quercioli and L. Smith. The economics of counterfeiting. *Econometrica*, 83(3): 1211–1236, 2015.

G. Rasch. Probabilistic models for some intelligence and attainment tests. Technical report, Copenhagen: Danish Institute for Educational Research, 1960.

M. D. Reckase and R. L. McKinley. The discriminating power of items that measure more than one dimension. *Applied Psychological Measurement*, 15(5):361–373, 1991. doi: 10.1177/014662169101500407.

Research and Markets. Global Brand Counterfeiting Report, 2018. https://www.researchandmarkets.com/research/7j7l2n/global_brand?w=4, 2017.

S. Resnick. *Adventures in Stochastic Processes*, chapter 6.8. Birkhauser Boston Inc., 1992.

R. Rotoras, T. Lefevre, and R. Pacudan. Marginal transmission pricing and supplemental cost allocation method: A case of philippines. *Electric Power Systems Research*, 63(3):213 – 227, 2002. ISSN 0378-7796. doi: https://doi.org/10.1016/S0378-7796(02)00119-0. [Accessed 27-June-2019].

P. Ruiz and J. Contreras. An effective transmission network expansion cost allocation based on game theory. *IEEE Transactions on Power Systems*, 22(1):136–144, 2 2007. [Accessed 27-June-2019].

J. Sanchez-Meca and F. Marín-Martínez. Weighting by inverse variance or by sample size in meta-analysis: A simulation study. *Educational and Psychological Measurement*, 58(2):211–220, 1998.

M. Schrepp. A generalization of knowledge space theory to problems with more than two answer alternatives. *Journal of Mathematical Psychology*, 41:237–243, 1997.

M. Schrepp. Extracting knowledge structures from observed data. *British Journal of Mathematical and Statistical Psychology*, 52(2):213–224, 1999.

M. Schrepp. About the connection between knowledge structures and latent class models. *Methodology*, 1(3):93–103, 2004.

Shalop. Distribution of the maximum of a (infinite) random walk. Mathematics Stack Exchange, 2017. https://math.stackexchange.com/q/2148285 (version: 2017-02-18), author page: https://math.stackexchange.com/users/224467/shalop.

L. Shapley. Cores of convex games. *International Journal of Game Theory*, 1(1): 11–26, 1971. [Accessed 27-June-2019].

G. Shevchenko. Distribution of the maximum of a (infinite) random walk. Mathematics Stack Exchange, 2017. https://math.stackexchange.com/q/2148319 (version: 2017-02-17; username: zhoraster).

S. E. Shreve. *Stochastic Calculus for Finance II: Continuous Time Models*, page 114. Springer, 2008.

UNESCO Institute of Statistics. The world needs almost 69 million new teachers to reach the 2030 education goals, 2016. http://uis.unesco.org/sites/default/files/documents/fs39-the-world-needs-almost-69-million-new-teachers-to-reach-the-2030-education-goals-2016-en.pdf.

S. E. Whitely. Multicomponent latent trait models for ability tests. *Psychometrika*, 45:479–494, 1980.

K. H. Wilson, Y. Karklin, B. Han, and C. Ekanadham. Back to the basics: Bayesian extensions of irt outperform neural networks for proficiency estimation. In *9th International Conference on Educational Data Mining*, June 2016.

World Health Organization. 1 in 10 medical products in developing countries is substandard or falsified. https://www.who.int/news-room/detail/28-11-2017-1-in-10-medical-products-in-developing-countries-is-substandard-or-falsified, 2017.

Y.-C. Yen, R.-G. Ho, W.-W. Laio, L.-J. Chen, and C.-C. Kuo. An empirical evaluation of the slip correction in the four parameter logistic models with computerized adaptive testing. *Applied Psychological Measurement*, 36(2):75–87, 2012. doi: 10.1177/0146621611432862.

C.-K. Yeung and D.-Y. Yeung. Addressing two problems in deep knowledge tracing via prediction-consistent regularization. https://arxiv.org/pdf/1806.02180.pdf, 2018. [Accessed 10-June-2019].

M. V. Yudelson, K. R. Koedinger, and G. J. Gordon. Individualized bayesian knowledge tracing models. In H. C. Lane, K. Yacef, J. Mostow, and P. Pavlik, editors, *Proceedings of the 16th international conference on Artificial intelligence in education*, pages 171–180. Springer-Verlag Berlin, July 2013.

J. Zhang, L. J. Hong, and R. Q. Zhang. Fighting strategies in a market with counterfeits. *Annals of Operations Research*, 192(1):49–66, 2012.