Information Procurement and Delivery: Robustness in Prediction Markets and Network Routing

by

Stanko B. Dimitrov

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Industrial and Operations Engineering)
in The University of Michigan
2010

Doctoral Committee:
Associate Professor Marina A. Epelman, Co-Chair
Assistant Professor Rahul Sami, Co-Chair
Professor Michael P. Wellman
Associate Professor Jussi Samuli Keppo
Adjunct Assistant Research Scientist Dushyant Sharma

ACKNOWLEDGEMENTS

I would like to thank my family and friends for their support. I thank my academic advisers, Marina Epelman and Rahul Sami, for their continuous support and encouragement. I appreciate my committee members: Jussi Keppo, Dushyant Sharma, and Michael Wellman, for taking the time to give me their input on my work.

The work presented in this dissertation is based on joint work with Marina Epelman, Rahul Sami, and Dushyant Sharma.

I thank Z. Morley Mao, Xu Chen, and Yin Zhang for their helpful discussions and feedback for Chapter II. I am grateful to Yiling Chen, P.J. Healy, John Ledyard, and David Pennock for useful feedback on prediction markets.

I would like to thank all of the various sources of funding I received during my graduate studies: NSF IGERT grant 0114368, NSF grant CCF-0728768, and 2009 Spring/Summer Research Grants Program at the University of Michigan.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

CHAPTER I

Introduction

Information is one of the most valuable commodities in business, as the products of many companies today are simply bits of information. The business models of these companies require explosive amounts of data collection and storage. Useful information must be extracted from these vast amounts of data, and methods for transmitting the data in a manageable way need to be developed. In this dissertation, we propose methods for information delivery and procurement in order to address these important issues.

In this introduction, we describe three problems arising from uncertainty in information delivery and procurement systems. Specifically, we describe one problem in network routing and two problems in prediction markets. We outline the contributions of this dissertation, addressing each problem through the design and analysis of robust algorithms and protocols that account for uncertainty.

- Problem 1: **Finding optimal routing polices in data networks that account for router based active congestion control.** Most traffic on current backbone networks is transmitted, end-to-end, using Transmission Control Protocol (TCP). TCP inherently incorporates congestion control by resending lost packets from the sender to the receiver. Due to its inherent congestion control methods, TCP is too slow to accommodate services such as Internet Protocol telephony, Internet Protocol television, and Instant Messaging, all growing in popularity. Such services use protocols like the User Datagram Protocol (UDP) that do not have built-in congestion control. In order to manage throughput across a network, congestion control must come from network components such as the routers. Classically, router based congestion control has been addressed by the Random Early Drop (RED) [34] algorithm. However, such routing protocols do not account for both routing and congestion control at the same time, severely limiting the network throughput. We propose a mathematical model to generate improved routing policies, while also taking into account congestion control. The proposed model is not only easily extended to incorporate RED, but can also take into account demand uncertainty when generating routing policies. According to our computational experiments, the resulting polices are at least 20% better than those currently used in a real world network [2].

- Problem 2: **Analyzing and alleviating the impact of non-myopic actions of risk neutral traders in prediction markets.** Prediction markets are an information aggregation tool in which participants trade on the outcome of a future event. Some forms of subsidized prediction markets have been proven to accurately aggregate traders' beliefs [42, 63]. Subsidized prediction markets have grown in popularity, and a number of Fortune 500 companies are using them to aid in decision making [24]. When introduced, subsidized prediction markets were only proven to accurately aggregate the beliefs of risk neutral traders that do not take into account future payoffs, *i.e.*, are myopic. However, we show that if traders are allowed to take into account future payoffs and have complementary information (meaning that a trader can earn a greater profit from knowing the other traders' private information in addition to her own), the incentive to not fully reveal information exists. We design a prediction market that adjusts non-myopic traders' incentives, leading to all traders revealing their true beliefs.

- Problem 3: **Analyzing and alleviating the impact of risk averse traders in prediction markets.** In practice people are not risk neutral, but tend to be risk averse. Therefore, accurately aggregating the beliefs of risk averse traders is key to creating practical prediction markets. We show that current prediction market mechanisms do not accurately aggregate the beliefs of risk averse traders, and propose one prediction market mechanism that does. Unfortunately, the reward distributed by this mechanism decreases exponentially with the number of traders. We prove that this exponential decrease is unavoidable by showing that for any subsidized prediction market mechanism to aggregate the beliefs of risk averse traders, it must exponentially decrease trader rewards.

1.1 Contributions

The contributions of this dissertation are:

Problem 1: Finding routing policies that maximize the amount of information received at destinations subject to a standard active congestion control method. For the network routing problem, we propose an optimization model that incorporates active congestion control in a multi-commodity network. We show that in general the problem is $\mathcal{NP}$-hard. However, using a robust instance of the problem, we show that a routing policy generated by applying a standard nonlinear programming optimization software to the robust instance outperforms the routing polices currently used in a real world network.

Problem 2: What happens when traders in a prediction market take into account future payoffs? For this problem, we show that when traders have complementary information, they have an incentive to bluff, meaning that they will not fully reveal their information when they trade in a market with an unlimited number of trades. We then propose a new prediction market to curb these incentives, resulting in traders fully revealing their information as the number of trades increases.

Problem 3: What happens when traders are risk averse? In studying this problem, we first characterize the desired properties all subsidized prediction market mechanisms should satisfy. Second, we observe that the reward distributed to traders with arbitrary risk averse preferences must be non-negative (this is not the case for current subsidized prediction market mechanisms). Third, we propose a prediction market mechanism that satisfies the desired properties. However, this mechanism exponentially reduces the rewards distributed to traders as the number of traders increases. Finally, we show that for any prediction market to satisfy the desired properties in the presence of traders with arbitrary risk averse preferences, the reward must decrease exponentially with the number of traders.

1.2 Outline of the Dissertation

In Chapter II, joint work with Marina Epelman and Dushyant Sharma, we describe the routing model used to generate routing polices while taking into account current active congestion control methods, and apply the model to the Abilene network. Further, we show that a routing policy that takes into account natural demand fluctuation performs better than the currently deployed routing policies. In Chapter III we introduce prediction markets, some of the related work, and mention the overlap that exists between Chapters IV and V. In Chapter IV, joint work with Rahul Sami, we show that under certain information settings, prediction markets using the log market scoring rule do not reach a market equilibrium because traders bluff, *i.e.*, never fully reveal their true beliefs. In order to address the issue of bluffing, we propose a discounted log market scoring rule and show that this market converges to the truthful prediction at an exponential rate as the number of trades in the market increases. In Chapter V, joint work with Marina Epelman and Rahul Sami, we present a summary of desirable properties of subsidized prediction market mechanisms, and propose a mechanism that possesses these properties even when risk averse players are present. In the same chapter we show that all prediction markets that satisfy the desirable properties must exponentially decrease the reward distributed to players as the number of players increases. In Chapter VI, we leave the reader with concluding remarks and future work.

The work presented in this dissertation is based on joint work with coauthors indicated above. As such, each chapter is intended to be a self contained exploration of the problem addressed. This means that the notation may differ across chapters and there is some overlap across introductory sections. In particular, Chapter III introduces prediction markets, discusses some related work, and introduces some of the notation used in Chapters IV and V. However, Chapters IV and V address two different problems in prediction markets. Therefore, the introductions of Chapters IV and V do refer to some of the same background material, and the notation introduced in Chapter III is modified in the subsequent chapters to appropriately present the results discussed in each of the chapters.

CHAPTER II

New Models of Network Routing under Active Congestion Control

This chapter is based on joint work with Marina Epelman and Dushyant Sharma. This work is submitted for publication at the time of this writing.

2.1 Introduction

In this chapter we consider network routing under congestion control. We focus on *active* congestion control. We say that a congestion control method is active if the amount of flow sent into a network component, such as an arc, is a function of the network status; for example, in computer networks, some routers are designed for congestion control and as information is passed through the routers some packets may be dropped. (Contrast this with a congestion control method that preserves the amount of flow on every component of the network; for instance, a road network, in which cars travel slower on a congested highway, but remain in the network until they reach their destinations.) Though the proposed model and techniques easily extend to most networks with congestion control, we focus on computer networks with active congestion control. In the remainder of this section we first present background required to understand the current routing policies and congestion control techniques used in computer networks, and then describe the question we address in this chapter.

In Section 2.2 we present our mathematical model for determining routing policies in a network with a particular type of congestion control and show, in Section 2.3, that it is $\mathcal{NP}$-hard. In Section 2.4 we present a real world computer network, and discuss the performance of different routing policies generated using locally optimal solutions of the model in Section 2.5. Finally, in Section 2.6 we propose a robust formulation of the model and present its performance relative to the routing policy currently used.

2.1.1 Model Idea

We design a routing policy in a computer network using a generalization of a multi-commodity network flow model described, for instance, by Ahuja *et al.* [4]. In our model, every origin and destination (OD) pair in the network constitute a commodity, and every commodity has a fixed amount of demand that is sent from the origin node to the destination node. The main difference from the generalized multi-commodity network flow model, as we discuss later, is that the amount of flow of each commodity received at the head of an arc is a function of the total flow on that arc. After introducing some background, we will revisit the model and explain it, and its relation to previous work, in greater detail.

2.1.2 Computer Network Background

In this section we discuss some of the current network protocols used in computer networks and point out which ones will be captured by the proposed mathematical model. In a telecommunications network, one has to make two decisions: one is what path(s) the information is going use from the source to the destination, using a routing protocol, and the other is what type of flow management will be used to improve quality of service in the network, using a network end-to-end protocol. In particular, we will describe Open Shortest Path First (OSPF) and Multiprotocol Label Switching (MPLS) routing protocols, and the implications of each one for finding routing policies used in the network. We will then describe Random Early Detection (RED), an active congestion control method currently available in computer networks, and the implications of taking RED into account when generating an optimal routing policy.

As a notational issue, it is important to note that computer networks transmit information in packets, or discrete blocks of information. However, in a network flow model each commodity is thought to be continuous and as such we can think of each commodity's flow in the proposed model as the rate of packets for that commodity in the computer network.

Network Protocols

Computer networks tend to use a variant of one of two end-to-end transmission protocols. One is the transmission control protocol (TCP) and the other is the user datagram protocol (UDP).

A packet sent using TCP is acknowledged by the receiver and, if a sender does not receive the acknowledgment in a given timeframe, the packet is re-sent after waiting an exponentially increasing amount of time. As one would expect, TCP inherently slows down the throughput due to the constant acknowledgments that are sent back to the senders. A network using TCP, by definition of the protocol, has built-in active congestion control in that it modifies the transmission rate, *i.e.*, the effective demand of each commodity, in response to congestion in the network.

UDP is a protocol commonly used to transmit Voice Over Internet Protocol (VoIP), Internet Protocol Television (IPTV), and Instant Messaging (IM) services traffic. It does not have built-in acknowledgment for every received packet, and all packets are sent once, with no guarantee of being received. If UDP is used without any congestion control, severe congestion could occur in the network. Moreover, a network using UDP without congestion control is susceptible to a denial of service attack during which an attacker floods the network with UDP packets. With the growing prominence of VoIP, IPTV and IM traffic on the Internet, examining the impact of active congestion control on computer networks with UDP traffic is of interest.

Routing Protocols

Open Shortest Path First (OSPF)

The Open Shortest Path First routing protocol, OSPF, was first proposed in 1989 and has been modified four times since the original request for comments (RFC) was posted. The RFC for the current version of OSPF was posted in April of 1998 [61]. OSPF is a routing policy used in intranet networks, *i.e.* Autonomous Systems, which are networks administered by a single organization. Abstractly OSPF can be described as routing demand along a single path for every origin and destination (OD) pair in the network, namely, the shortest path between those origin and destination nodes. As a rule of thumb, Cisco Systems recommends the arc weights used in the shortest path calculation to be set to the inverse of the arc capacities [70]. In practice, since arc capacities seldomly change, the paths between nodes are updated rather infrequently. As most networks, specifically the ones we examine in Section 2.4, are still using OSPF, we will use network performance under OSPF routing as a benchmark for assessing network performance under Multiprotocol Label Switching, which we will describe next.

Multiprotocol Label Switching (MPLS)

Multiprotocol Label Switching, MPLS, was proposed in January of 2001 [55].

MPLS differs from OSPF in that each OD pair in the network has multiple paths, which may or may not be disjoint, simultaneously able to carry positive flow from the origin to the destination.

When a packet first enters the network, the incoming router looks at the destination of the packet and chooses which of the possible paths it should follow to its destination. It then assigns the appropriate label to the packet and forwards it to the first node in the determined path.

Associated with each intermediate router in the network is an *MPLS routing table.* Each row of the table contains information that determines, given the neighbor the packet came from and its current label, the next node on the packet's path to its destination, and a new label to attach to the packet (the new label can be interpreted by the next router in the same manner). Once a packet is received by a router, the router removes the current label of the packet, identifies the appropriate row of its MPLS table, and attaches a new label and forwards the packet accordingly.

The proposed research is to find a routing policy which is feasible in a network using MPLS while accounting for both active congestion control and demand uncertainty. Though currently OSPF is used in intranet networks, most networks are beginning to port to MPLS. Therefore, now is the time to address the issue of finding good MPLS routing policies for networks facing congestion control and demand uncertainty.

Random Early Detection (RED)

In a computer network, whenever a packet, or a datagram of information, is forwarded from one router to another, that packet must be examined by the forwarding router. If two or more packets need to be examined, then all packets not being examined are placed in a queue of finite capacity, say u, and serviced in a first in first out (FIFO) manner. However, if the rate of incoming packets is greater than the router service rate, the queue will reach capacity u and no incoming packets will be enqueued. To address the issue of the resulting starvation, the Random Early Detection, RED, congestion avoidance mechanism was introduced [34]. The mechanism separates the router queue capacity into three regions, characterized by parameters $0 \leq \beta \leq \gamma \leq u$. In the first region, between an empty queue and a queue length of β, all incoming packets are enqueued to wait for service. In the second region, with queue length between β and γ, the probability that a packet is enqueued is determined by a decreasing linear function with a slope of $-\alpha$, for $\alpha > 0$. Finally, in the third region, with queue length between γ and u, none of the incoming packets are enqueued into the router queue. Since γ determines the effective capacity of the queue, without loss of generality we will let $\gamma = u$ for the remainder of this chapter.

2.1.3 Gain Functions

Though our work is motivated by an application to computer network routing, the nominal problem being addressed is an extension of the generalized network flow problem, as described by Ahuja *et al.* [4, Chapter 15]. In a single commodity setting, we denote by x_{ij} the amount of flow sent from node i to node j on arc (i, j) and by y_{ij} the amount of flow received at node j from node i on arc (i, j). In the classic network flow model we would have $y_{ij} = x_{ij}$. In a generalized flow setting, however, we have $y_{ij} = x_{ij}\mu_{ij}$, where $\mu_{ij} > 0$ represents the proportional loss or gain of flow on arc (i, j). In a network with congestion control imposed by the RED algorithm, we have $y_{ij} = x_{ij}f_{ij}(x_{ij})$. In this setting $f_{ij}(x_{ij})$ represents the loss, due to congestion at router at node i, that takes place on arc (i, j). Since $f_{ij}(x_{ij})$ is a function of x_{ij}, this new model is a further generalization of the generalized network flow model.

As noted in Section 2.1.2, we can define a function $g(t)$, the probability of a packet being enqueued by the router as a function of the queue length t, as

$$g(t) = \begin{cases} 1 & 0 \le t \le \beta, \\ 1 - \alpha(t - \beta) & \beta \le t \le u, \\ 0 & u \le t. \end{cases}$$

We chose $\alpha = \frac{1}{u-\beta}$ to guarantee continuity of $g(\cdot)$. Note also that with this choice of α the effective capacity of the arc is u, as desired. As defined, $g(t)$ determines the loss on an arc as a function of y_{ij}, because every enqueued packed will be serviced. Therefore, the $f(t)$ function satisfies $f(t) = g(tf(t))$. When $g(t)$ is defined as above, the resulting $f(t)$ function is:

$$f(t) = \begin{cases} 1 & 0 \le t \le \beta, \\ \frac{1+\alpha\beta}{1+\alpha t} & \beta \le t. \end{cases} \tag{2.1}$$

We denote by $h(t) = tf(t)$ the amount of flow received as a function of t. With $f(t)$ given by (2.1),

$$\begin{aligned} h(t) &= tf(t) \\ &= \begin{cases} t & 0 \le t \le \beta \\ \frac{1+\alpha\beta}{\alpha} - \frac{1+\alpha\beta}{\alpha(1+\alpha t)} & \beta < t, \end{cases} \end{aligned}$$

and with α as above, $h(t) \to u$ as $t \to \infty$, ensuring that the capacity is not exceeded. Observe also that $h(t)$ is non-decreasing and concave.

In a multi-commodity setting that we will consider in this chapter, the flow conservation relationship becomes $y^k_{ij} = x^k_{ij} f\left(\sum_l x^l_{ij}\right)$ for each arc (i, j) and each commodity k.

Shigeno [69] studied the problem of finding optimal routing polices for a single commodity using congestion control. He refers to the function $f(t)$ above as a *gain* *function*, and defines a *concave gain function* to be a gain function $f(t)$ such that $h(t) = tf(t)$ is concave and non-decreasing. Note that the RED congestion function $f(t)$ given by (2.1) is less than or equal to 1 for all t, and thus represents a *loss* of flow on an arc. However, to stay consistent with Shigeno we refer to $f(t)$ as a gain function.

2.1.4 Proposed Problem

In this chapter we propose a mathematical model for Internet routing under active congestion control. We propose a continuous flow model (rather than a packet burst model) that is motivated by computer networks using UDP with RED deployed in the network. Though the proposed model is $\mathcal{NP}$-hard, we show that routing policies obtained by applying a nonlinear programming solver to the model improve network performance over existing routing policies for a computer network.

2.1.5 Previous Research

Shigeno [69] introduced the concept of a *concave gain function*, and associated a concave gain function with every arc in a network. He showed that in a single commodity network flow problem with concave gains, a routing policy maximizing the total flow received at the destination node can be found in polynomial time. As our problem is a multi-commodity flow problem with concave gains, it helps to think of the proposed problem as the multi-commodity flow generalization of Shigeno's work.

Several studies besides Shigeno have looked at generating routing policies in networks while taking into account congestion. For example, Sheffi [67] addresses the issue of roadway congestion in optimal traffic selection. In his models he proposes a convex function representing the travel time on a roadway (*i.e.* an arc of a network).

As the number of users of a roadway segment increases, so does their travel time on that segment, until the capacity of the roadway is reached. The users of the roadway are assumed to minimize their total travel time. This model is similar to the one of interest; however, all passengers on a roadway remain on a roadway, while in our model we can remove users. Studies addressing demand uncertainty using a robust routing scheme will be discussed in detail in Section 2.6.

2.2 Multi-Commodity Network Flow Problem with Nonlinear Gains

We begin by defining the Multi-Commodity Network Flow problem with Nonlinear Gains (MCFPNG). Let $G = (N, A)$ be a directed graph with node set N and arc set A, and $(o^k, d^k) \in N \times N$ for $k = 1, \ldots, K$ be origin-destination node pairs for K commodities. We consider the arcs to have infinite capacity, and let $f_a : \mathbb{R}_+ \to \mathbb{R}_+$ for $a \in A$ be *gain functions* associated with each of the arcs. Let s^k be the supply of commodity k originating at its origin o^k, and c^k be the value per unit of this commodity delivered to destination d^k. We will use the following notation:

- $x_i^k \equiv$ amount of flow of commodity k present at node $i \in N$;
- $\alpha_{ij}^k \equiv$ the fraction of commodity k present at node i sent to node j, on arc $(i, j) \in A$;
- $\delta^+(v)$ and $\delta^-(v)$ denote the sets of nodes in N that are end points of arcs coming out of node v and coming into node v, respectively.

There are several possible ways of defining the MCFPNG; our version is, essentially, a weighted maximum flow problem with multiple commodities and (nonlinear) gain functions on each arc.

Specifically, we define (MCFPNG) as:

$$\text{(2.2)} \quad \max_{\alpha, x} \quad \sum_{k=1}^{K} c^k x^k_{d^k}$$

$$\text{(2.3)} \quad \text{s.t.} \quad x^k_{o^k} = s^k \qquad k = 1, \ldots, K$$

$$\text{(2.4)} \quad \sum_{(i,j) \in A} \alpha^k_{ij} = 1 \qquad k = 1, \ldots, K,\ i \in N : i \neq d^k$$

$$\text{(2.5)} \quad \sum_{(i,j) \in A} \alpha^k_{ij} = 0 \qquad k = 1, \ldots, K,\ i = d^k$$

$$\text{(2.6)} \quad \sum_{i \in \delta^-(j)} \alpha^k_{ij} x^k_i f_{ij} \left(\sum_{l=1}^{K} \alpha^l_{ij} x^l_i \right) - x^k_j = 0 \qquad k = 1, \ldots, K,\ j \in N : j \neq d^k$$

$$\text{(2.7)} \quad \alpha^k_{ij} \geq 0 \qquad (i,j) \in A,\ k = 1, \ldots, K.$$

Here, the objective function (2.2) maximizes the weighted sum of flows of each commodity delivered to the destination nodes, while constraints (2.3) indicate the supply of each commodity. Constraints (2.4) and (2.5), together with (2.7), ensure that the entire amount of commodity k available at node i is routed along the edges emanating from i, with the exception of the destination node for that commodity. Constraints (2.6) calculate the amount of commodity k available at node j by tracking the flow of that commodity routed, and lost, on each of the arcs coming into node j.

2.3 Complexity of MCNFCG

This section is dedicated to proving the following

Theorem II.1. *MCFPNG is $\mathcal{NP}$-hard.*

The proof is done by reduction from the Set Cover problem:

Definition II.2. *Set Cover* in minimization form is defined as follows: given

- set $U = \{e_1, e_2, \ldots, e_m\}$ and
- collection of subsets $S_j \subseteq U, \quad j \in \{1, 2, \ldots, n\}$,

find a minimum set cover, $i.e.$set $J \subseteq \{1, 2, \ldots, n\}$ such that $\cup_{j \in J} S_j = U$, of minimum cardinality.

Given a Set Cover instance, we construct the following directed graph G, as a 4 layer network:

Layer 1 consists of one node for every element, $e_i \in U,\ i \in \{1, \ldots, m\}$,

Layer 2 consists of one node for every subset, $S_j,\ j \in \{1, \ldots, n\}$,

Layer 3 consists of one node, I,

Layer 4 consists of one node, t.

The layers are connected in the following manner:

- For every i and j such that $e_i \in S_j$ there is a directed arc (e_i, S_j),
- For every j there is a directed arc (S_j, I),
- There is a directed arc (I, t).

For example, consider the following instance of set cover:

$$\begin{aligned} U &= \{e_1, e_2, e_3, e_4, e_5\} \\ S_1 &= \{e_1, e_2, e_3\} \\ S_2 &= \{e_1, e_2, e_3, e_4\} \\ S_3 &= \{e_1, e_4\} \\ S_4 &= \{e_1, e_5\}. \end{aligned}$$

The corresponding directed graph constructed as described above is depicted in Figure 2.1.

To define an instance of MCFPNG, we define the following $m + 1$ commodities:

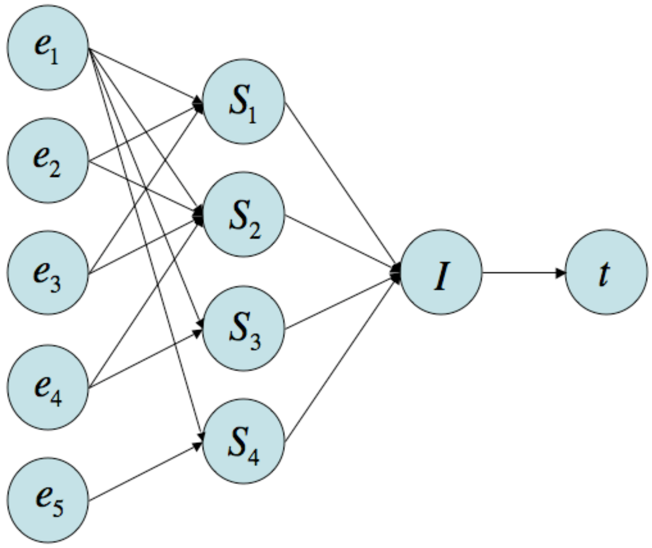

Figure 2.1: Example of graph construction

Commodity 0 from origin node I to destination node t, $i.e.o^0 = I$ and $d^0 = t$, with $s^0 = 1$ and $c^0 = 1$;

Commodity i, for $i = 1, \ldots, m$ from origin node e_i to destination node t, $i.e.o^i = e_i$ and $d^i = t$, with $s^i = 1$ and $c^i = 0$.

Finally, the gain function of each arc is

$$f_a(t) = f(t) = \begin{cases} 1 & t < 1 \\ \frac{1}{t} & 1 \leq t \end{cases}$$

for all $a \in A$, fitting the definition of a concave gain function. Though we can consider the capacities of the arcs to be infinite, notice that the form of the gain function $f(t)$ above implies that the effective capacity of each arc is 1.

Note that the there were $O(nm)$ steps required to transform the instance of set cover into an instance of MCFPNG.

Given an instance of Set Cover, the corresponding MCFPNG instance resulting from this transformation is:

$$
\begin{array}{lll}
(2.8) & \max\limits_{\alpha, x} \; x_t^0 & \\
(2.9) & \text{s.t.} \;\; x_I^0 = 1 & \\
(2.10) & x_{e_k}^k = 1 & k = 1, \dots, m \\
(2.11) & \sum\limits_{(i,j) \in A} \alpha_{ij}^k = 1 & k = 0, \dots, m, \; i \in N : i \neq d^k \\
(2.12) & \sum\limits_{(i,j) \in A} \alpha_{ij}^k = 0 & k = 0, \dots, m, \; i = d^k \\
(2.13) & \sum\limits_{i \in \delta^-(j)} \alpha_{ij}^k x_i^k f_{ij} \left(\sum\limits_{l=0}^{m} \alpha_{ij}^l x_i^l \right) - x_j^k = 0 & k = 0, \dots, m, \; j \in N : j \neq d^k \\
(2.14) & \alpha_{ij}^k \geq 0 & (i,j) \in A, \; k = 0, \dots, m.
\end{array}
$$

Below we explore some of the properties of optimal solutions to this problem. For convenience, we will use the following notation in the rest of this section:

$$
X_v = \sum_{i=1}^{m} x_v^i, \; v \in N, \tag{2.15}
$$

$i.e. X_v$ denotes the total amount of flow of commodities 1 through m present at node v.

Proposition II.3. *A feasible solution* (α, x) *of* (2.8) – (2.14) *is optimal if and only if it minimizes* X_I.

Proof. Recall that

$$
X_I = \sum_{i=1}^{m} x_I^i.
$$

For any feasible solution (α, x), $\alpha_{It}^k = 1 \; \forall k$ due to (2.11), as there is only one arc out of I. In addition, $x_I^0 = s^0 = 1$, and so $X_I + x_I^0 \geq 1$, implying that $f(X_I + x_I^0) = \frac{1}{X_I + x_I^0}$. Moreover, $x_t^0 = x_I^0 \cdot f(X_I + x_I^0) = \frac{1}{1 + X_I}$. Therefore, as the objective of (2.8) is to maximize x_t^0, a feasible solution is optimal if and only if it minimizes X_I. □

Proposition II.4. *There exists an optimal solution (α, x) which has integral values (0 or 1) of variables α and x associated with all arcs from layer 1 to layer 2.*

Proof. If in a feasible solution the values of α are integral (0 or 1) for all arcs from layer 1 to layer 2, then, according to (2.11), this solution has positive flow on exactly one arc (e_i, S_j) for each $i = 1, \ldots, m$. Since the supply at each node e_i is $s^i = 1$, the amount routed on all arcs from layer 1 to layer 2 is 0 or 1. Thus, we only need to show existence of a solution with integral values of α associated with all such arcs.

Let (α, x) be an optimal solution, and suppose that there is a node e_i in layer 1 such that the α values at this node to layer 2 split commodity i between two or more arcs. Without loss of generality, let the corresponding nodes in layer 2 be S_1 and S_2, $i.e. \alpha_{e_i,S_1} > 0$ and $\alpha_{e_i,S_2} > 0$.

Recall that the total flow of commodities $1, \ldots, m$ present at node I is equal to $X_I = \sum_{j=1}^{n} X_{S_j} f(X_{S_j})$, by construction of G. Without loss of generality, assume $X_{S_1} \geq X_{S_2}$.

Consider solution $(\tilde{\alpha}, \tilde{x})$ that is obtained from solution (α, x) by moving the flow of commodity i from arc (e_i, S_2) to arc (e_i, S_1), $i.e. \tilde{\alpha}^i_{e_i,S_2} = 0$ and $\tilde{\alpha}^i_{e_i,S_1} = \alpha^i_{e_i,S_1} + \alpha^i_{e_i,S_2}$, while all other α values at layer 1 nodes remains the same. It is easy to verify that solution $(\tilde{\alpha}, \tilde{x})$ is feasible. We will show that

$$\sum_{j=1}^{n} \tilde{X}_{S_j} f(\tilde{X}_{S_j}) \leq \sum_{j=1}^{n} X_{S_j} f(X_{S_j}),$$

and thus, in view of Proposition II.3, $(\tilde{\alpha}, \tilde{x})$ is optimal. In fact, it only needs to be shown that

$$\tilde{X}_{S_1} f(\tilde{X}_{S_1}) + \tilde{X}_{S_2,} f(\tilde{X}_{S_2}) \leq X_{S_1} f(X_{S_1}) + X_{S_2} f(X_{S_2}),$$

since the other values remain unchanged.

Due to the supply constraint at node e_i, $x^i\alpha^i_{(e_i,S_1)} + x^i\alpha^i_{(e_i,S_2)} = x^i\tilde{\alpha}^i_{(e_i,S_1)} + x^i\tilde{\alpha}^i_{(e_i,S_2)} \leq 1$. Therefore, by the definition of $f(\cdot)$, $x^i_{S_1} = \alpha^i_{(e_i,S_1)}x^i_{e_i}$, $x^i_{S_2} = \alpha^i_{(e_i,S_2)}x^i_{e_i}$, $\tilde{x}^i_{S_1} = \tilde{x}^i_{e_i} = x^i_{e_i}(\alpha^i_{(e_i,S_1)} + \alpha^i_{(e_i,S_2)})$ and $\tilde{x}^i_{S_2} = x^i_{e_i}\tilde{\alpha}^i_{(e_i,S_2)} = 0$.

Consider the following two cases:

Case 1: $X_{S_2} < 1$**.** In this case $X_{S_2}f(X_{S_2}) = X_{(S_2,I)}$ and $\tilde{X}_{S_2}f(\tilde{X}_{S_2}) = \tilde{X}_{S_2} = X_{S_2} - x^i_{S_2}$. Furthermore, $\tilde{X}_{S_1}f(\tilde{X}_{S_1}) = (X_{S_1} + x^i_{S_2}) \cdot f(X_{S_1} + x^i_{S_2})$. Notice that the function $f(t)$ satisfies

$$(t + \Delta)\, f(t + \Delta) \leq t\, f(t) + \Delta \text{ for all } t \geq 0 \text{ and } \Delta \geq 0,$$

and so

$$\begin{aligned}
\tilde{X}_{S_1}f(\tilde{X}_{S_1}) + \tilde{X}_{S_2,}f(\tilde{X}_{S_2}) &= X_{S_2} - x^i_{S_2} + (X_{S_1} + x^i_{S_2}) \cdot f(X_{S_1} + x^i_{S_2}) \\
&\leq X_{S_2} - x^i_{S_2} + X_{S_1}f(X_{S_1}) + x^i_{S_2} \\
&= X_{S_1}f(X_{S_1}) + X_{S_2}f(X_{S_2}),
\end{aligned}$$

as desired.

Case 2: $X_{S_1} \geq X_{S_2} \geq 1$**.** Notice that in this case $X_{S_1}f(X_{S_1}) + X_{S_2}f(X_{S_2}) = 2$. In the new solution, $\tilde{X}_{(S_1,I)} \geq X_{(S_1,I)} \geq 1$, and hence $\tilde{X}_{S_1}f(\tilde{X}_{S_1}) = 1$. On the other hand, due to the form of the gain function, $\tilde{X}_{S_2}f(\tilde{X}_{S_2}) \leq 1$, and hence the total amount of commodities 1 though m arriving at node I will not increase.

Modifications above removed the flow on one of the arcs between layers 1 and 2 without loss of optimality. Applying this procedure repeatedly will generate an optimal solution with positive flow on exactly one arc (e_i, S_j) for each $i = 1, \ldots, m$.

Finally, notice that each step of the above modification procedure takes a constant amount of time to execute, and needs to be applied at most mn times. $\square$

The procedure outlined in the above proof can be applied to any feasible solution (α, x), producing a feasible solution $(\tilde{\alpha}, \tilde{x})$ with integral values on arcs from layer 1 to layer 2 and with the same or better objective function value in at most mn steps.

Proposition II.5. *A feasible solution of* (2.8) – (2.14) *with integral flow (0 or 1) on each arc from layer 1 to layer 2 is optimal if and only if it minimizes the total number of arcs from layer 2 to layer 3 with nonzero flow.*

Proof. In any feasible solution (α, x) with integral flow (0 or 1) on each arc from layer 1 to layer 2, $X_{S_j} = \sum_{i=1}^{m} x^i \alpha^i_{(e_i,S_j)} \in \mathbb{Z}_+$ for all $j = 1, \ldots, n$. Therefore, for any j,

$$X_{S_j} f(X_{S_j}) = \begin{cases} 0 & X_{S_j} = 0, \\ 1 & X_{S_j} > 0, \end{cases}$$

and hence

$$X_I = \sum_{j=1}^{n} X_{S_j} f(X_{S_j}) = \sum_{j=1}^{n} \mathbb{I}(X_{S_j} > 0),$$

which, in view of Proposition II.3, implies the conclusions of the proposition. □

Proposition II.6. *Suppose* (α, x) *is an optimal solution of* (2.8) – (2.14) *with integral flow (0 or 1) on each arc from layer 1 to layer 2. Then the set* $J = \{j : X_{(S_j,I)} > 0\}$ *is an optimal solution to the corresponding instance of the minimum set cover problem.*

Proof. First notice that any feasible solution of MCFPNG (2.8) – (2.14) with integral flow on each arc from layer 1 to layer 2 corresponds to a set cover J constructed as follows: $j \in J$ if and only if $\alpha^i_{(e_i,S_j)} = 1$ for some e_i, or equivalently, $X_{S_j} > 0$ (and integral). Supply constraints at e_i, $i = 1, \ldots, m$, imply that J is, indeed, a cover, and we have

$$X_I = \sum_{j=1}^{n} X_{S_j} f(X_{S_j}) = \sum_{j=1}^{n} \mathbb{I}(X_{S_j} > 0) = |J|.$$

Conversely, every minimal set cover J (i.e., one that does not contain a set cover of smaller cardinality) can be represented by a feasible solution of the corresponding MCFPNG with integral flow on each arc from layer 1 to layer 2 as follows: For each $i = 1, \ldots, m$, pick any $j \in J$ such that $e_i \in S_j$ and set $\alpha^i_{(e_i,S_j)} = 1$, *i.e.* direct 1 unit of flow along the arc (e_i, S_j) (since J is a set cover, such j can always be found). Assign flows on arcs out of layer 2 and layer 3 nodes accordingly to satisfy flow gain constraints for the arcs and flow balance constraints for the nodes. Again, we have

$$|J| = \sum_{j=1}^{n} \mathbb{I}(X_{S_j} > 0) = \sum_{j=1}^{n} X_{S_j} f(X_{S_j}) = X_I.$$

Thus, by Proposition II.5, the minimum set cover can be obtained by finding an optimal solution to MCNFNG (2.8) – (2.14) with integer flows on arcs from layer 1 to layer 2. □

To complete the reduction, in view of Proposition II.6, one only needs to recall that an arbitrary optimal solution to (2.8) – (2.14) can be modified in at most mn steps into one with integral flows from layer 1 nodes to layer 2 nodes.

We have thus established that the MCFPNG as modeled by (2.2) – (2.7) is $\mathcal{NP}$-hard. Nonetheless, a nonlinear programming solver can be used to successfully find locally optimal solutions to instances of the MCFPNG. It is important to note that the complexity proof hinged on the fact all of the flow present at any non-destination node must be sent, constraint (2.11), and that there were no self loops in the routing of the commodities. If either of these two conditions were not satisfied then the resulting problem is no longer difficult to solve. For example, if we relax (2.11) ($\sum_{(i,j)\in A} \alpha^k_{ij} \leq 1, \quad \text{for } k = 0, \ldots, m, \; i \in N : i \neq d^k$), then all of the commodities of value zero in the construction would remain at their destinations leading to a trivial solution. Similarly, if we had self-loops, all of the zero value commodities would be sent on the self loops effectively never interfering with commodity of value. The remainder of the chapter is dedicated to numerical experiments with this model.

2.4 Abilene Network

As a testbed for numerical experiments in the following sections we used the Abilene Network depicted in Figure 2.2. The Abilene Network is the backbone network of the Internet2 community. The Internet2 community is a not for profit consortium of universities, companies, and government agencies that develops and deploys advanced network applications critical to the progress of the Internet. We have obtained network usage data in the Abilene Network in the form of 24 weekly data sets during a 6-month period of time in 2004 (March 1, 2004 to September 4, 2004) [1]. As the usage data does not span the entire 6-month period, we do not present the data instances in absolute terms (e.g. 10 AM August 12), but instead use relative terms (e.g. instance 2783). Due to the time the data was collected, for the remainder of this chapter we discuss the Abilene Network as it existed in 2004, during which time OSPF was used to deliver traffic between every origin and destination pair, and no congestion control protocols were implemented.

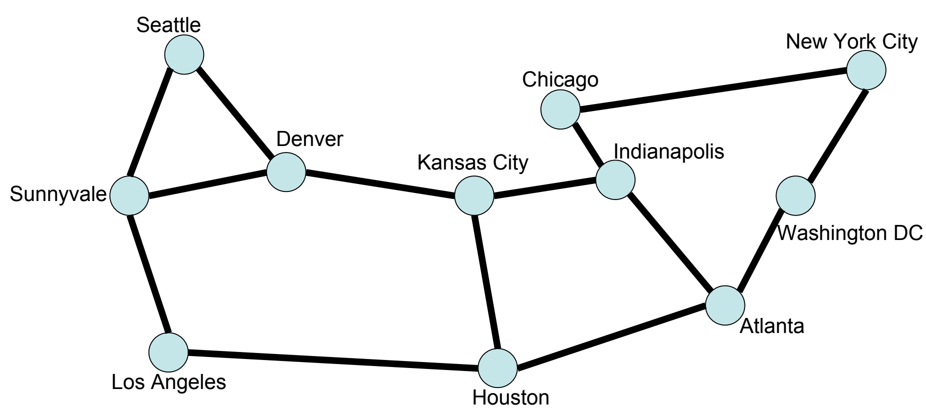

Figure 2.2: The Abilene Network

2.5 Numerical Experiments

In the remainder of the chapter we will present results of our numerical studies. In this section, we first describe the specifics of how we applied our model to the Abilene network, including a discussion of the objective function used, and present the first set of numerical results.

The Abilene network usage data, provided at [1], consists of the demands between every origin and destination pair in the 11-node network over five minute intervals. It should be pointed out that the data collected consist of arc flows during the specified period of time (*e.g.* 5 minutes). This arc flow data is then converted into estimates of demand between every OD pair with a method such as the one discussed by Roughan *et al.* [64]. The method used is not the topic of this chapter, but it is important to note that the demand data provided at [1] is only an estimate of the true demand.

For our numerical experiments we encoded the model (2.2) – (2.7) in AMPL [35], and used the data captured from the Abilene Network to define a family of data files, one for each time interval (we aggregated the data into hour-long intervals). We used SNOPT [38], a nonlinear programming solver, to solve the resulting problem instances. In preliminary experiments with MCFPNG models, we found SNOPT to perform better than several other popular nonlinear solvers, possibly due to the high level of nonlinearity of the constraints. SNOPT uses a sequential quadratic programming algorithm to find locally optimal solutions, and has been successfully used in solving nonlinear mathematical programs [32, 40]. We approximated the (non-differentiable) RED function outlined in Section 2.1.3 by

$$f(t) = \frac{1}{t+1}. \qquad (2.16)$$

Note that this gain function implies that every arc in the network has effective capacity of 1. Correspondingly, we scaled the demand data provided at [1] by a constant factor, in part so that no commodity had a demand greater than 1. (The scaling factor is intended to calibrate the demands in the network to be consistent with arcs having capacity of 1 unit, while maintaining the *relative* demand levels between the OD pairs, and to be sufficiently high to justify deployment of congestion control. Thus, the value of the scaling factor was arrived at, to some extent, by trial and error. It is worthwhile pointing out that, after scaling, the overall level of demand in the network, as compared to the capacity, is quite high. In particular, under routing policies considered below, only a small fraction of some of the commodities reaches the destination.)

2.5.1 Role of Objective Function

In (2.2) – (2.7) we presented the model with an objective function that maximizes the weighted total of commodities received at all the destinations:

$$\max_{\alpha,\, x} \sum_{k=1}^{K} c^k x^k_{d^k}. \tag{2.17}$$

Alternatively, we may choose to maximize the weighted total *fractions* of commodities received at all of the destinations:

$$\max_{\alpha,\, x} \sum_{k=1}^{K} c^k \frac{x^k_{d^k}}{s^k}. \tag{2.18}$$

This objective function is appropriate in models motivated by the use of congestion control in UDP networks, common in VoIP and IPTV applications, in which the fraction of commodity delivered reflects the quality of service for that commodity. Moreover, using objective (2.18) provides an incentive for commodities with smaller demands to be routed to their destinations.

Unfortunately, optimizing with respect to either one of these objective functions may still lead to starvation, *i.e.*a situation in which some commodities are ignored and not routed to their destination because of their lower relative values. This can be circumvented, for instance, by adding constraints assuring a minimum performance guarantee for all of the commodities. (For example, constraints can stipulate that at least $z\%$ of every commodity must the received at the destination node.)

However, in a network with congestion control, imposing a minimum performance guarantee could lead to an infeasible problem instance. An alternative approach to avoiding starvation is to utilize a "max-min" objective function. For example, objective function (2.17) can be modified as:

$$\max_{\alpha,\, x} \min_{k\in\{1,\ldots,K\}} c^k x^k_{d^k}. \tag{2.19}$$

Similarly, we can define the max-min modification of objective function (2.18). The issue with either of these formulations is that the resulting routing policies may be hindered by a commodity that simply has very little demand to begin with. An optimal routing policy will optimize for the commodity with lowest demand and may actually perform much worse, by any other metric, than one obtained with the cumulative objective function approach.

In the numerical experiments discussed in this chapter, we used objective function (2.18), as our work is motivated by the use of congestion control in UDP networks. However, the above considerations should be taken into account, and alternative objective functions should be considered, when applying similar optimization models to determine routing policies in networks with specific cost and quality of service considerations.

2.5.2 Numerical Results

In the first set of numerical experiments, we considered the first week of usage data (aggregated into hour-long increments). For each of the resulting 168 data instances, we compared the performance of the current routing policy, OSPF, and the MPLS routing policy optimized for that data instance, when each is subject to congestion control. Recall that performance comparisons were done with respect to the objective function in the form (2.18); we varied the weight coefficients c^k to arrive at three different sets of problem instances.

It should be pointed out that, since the feasible region of MCFPNG is not convex, there might be multiple local, but not global, solutions to (2.2) – (2.7). Thus, the routing policy found by the solver may not be truly optimal for the problem. Moreover, which (local) optimum is found by the solver is dependent on the initial routing policy used as a starting point of the optimization algorithm. We chose to initialize the solver with the robust routing policy, discussed in Section 2.6, when finding the proposed MPLS routing policy. The reasons for this choice will become clear in the following section. In reporting our computational results, we nonetheless refer to the solutions found by the solver as optimal, for simplicity of presentation.

In Figures 2.3–2.6, we compare the performance of the OSPF and MPLS routing policies with three different objective functions. The objective functions differ by having different values of commodity weights, c^k. As we do not have precise information regarding the values of the commodities in the network we consider three different choices of the weights, to see if there is a performance change amongst the resulting problem instances. The key takeaway from Figures 2.3–2.6 is that, regardless of the commodity weights used the MPLS routing policy, optimized for every demand instance, performs better than the OSPF routing policy. Each objective function is the result of different valuations of commodities in the network; to provide a framework for comparison between results reported in Figures 2.3, 2.4 and 2.5, we note that node 1 (Sunnyvale) provides roughly 4% of the overall demand in the network, while over 23% of the overall demand originates in node 10 (Washington, DC). Note that the MPLS routing policy, when optimized for each demand instance, performs much better than the OSPF routing policy in every instance and for every objective function used. As shown in Figure 2.6, the improvement over the OSPF routing policy is at least 27% when the objective function values all commodities equally; similar improvements are obtained for other objective functions as well. To summarize, a routing policy that (i) takes into account the actual demand, and (ii) allows routing of each commodity on multiple paths performs much better than a policy, such as OSPF, that does not.

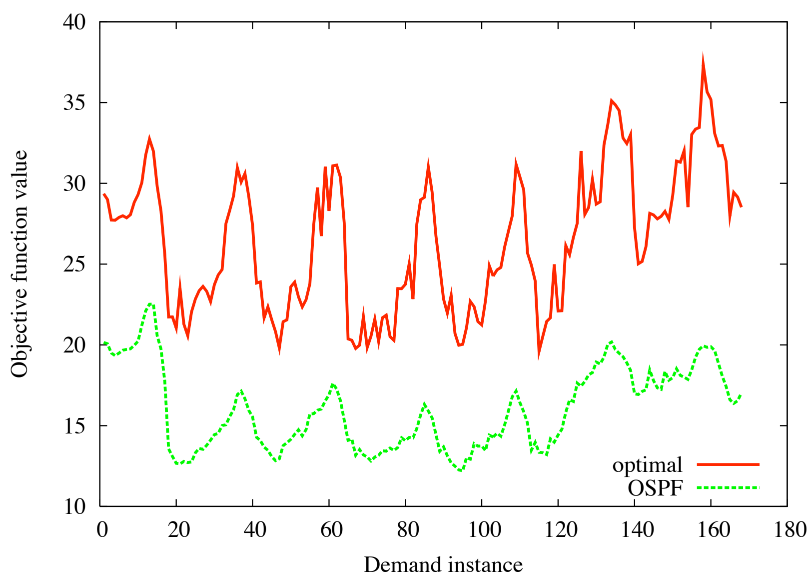

Figure 2.3: Commodities from node 1 are 5 times more valuable than all other commodities

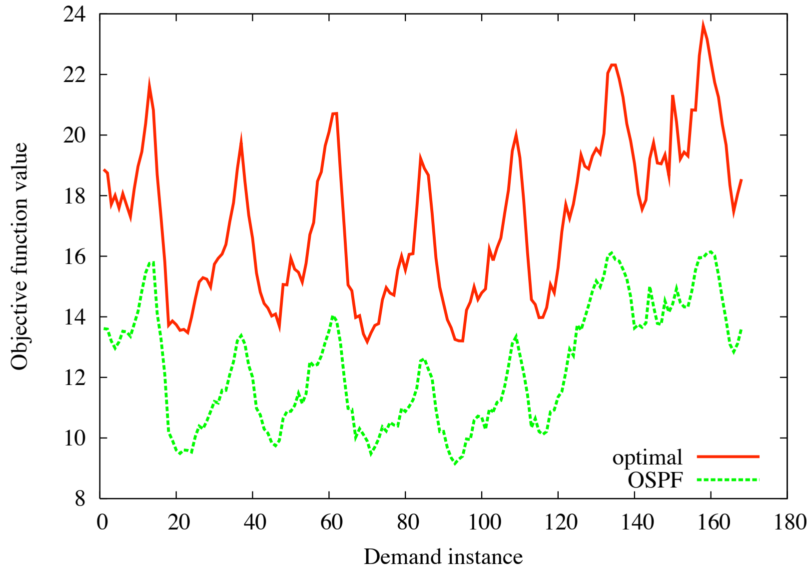

Figure 2.4: Commodities from node 10 are 5 times more valuable than all other commodities

2.6 Robust MPLS Routing Policies

In this section we present a robust reformulation of (2.2) – (2.7) and study empirical performance of the resulting routing policy.

In Section 2.5.2 we compared the performance of the OSPF routing policy to the MPLS routing policy that is optimized for every demand instance. As one would expect, the latter outperformed the former in every instance. In practice, however, demand in a network fluctuates continuously and is not known in advance; thus, re-optimizing the routing policy for every short time period would not be feasible. Therefore, in this section we will use ideas of robust optimization to propose a robust counterpart of MCFPNG and find one MPLS routing policy which performs well for a variety of demand instances observed in the network.

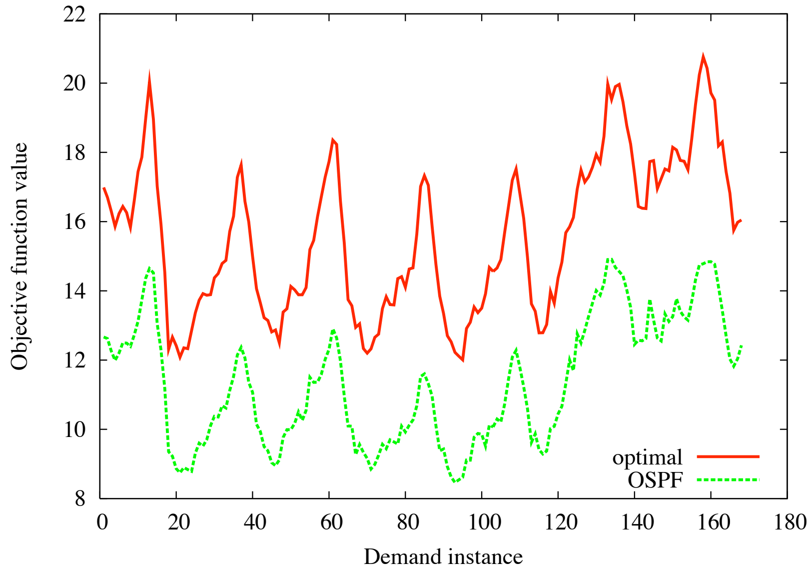

Figure 2.5: All commodities equal value

2.6.1 Robust Reformulation

Modern robust optimization was simultaneously introduced by Ben-Tal and Nemirovski [73] and El Ghaoui *et al.* [37]. Robust optimization is a mathematical programming modeling technique used when the problem data is not known exactly, but instead it is known (or assumed) that any data instance from an *uncertainty set* can be the problem data. The objective is to find a solution that is optimal among the solutions feasible under *all* possible data realizations. Using this approach Ben-Tal and Nermirovski [72] and El-Ghaoui and Lebret [36] pose and solve problems in robust truss topology design and robust least-square optimization, respectively. Recently robust optimization has attracted a lot of attention and has been considered for portfolio selection problems [39], integer programming and network flow problems [12], supply chain management [13], inventory theory [14], radiation treatment

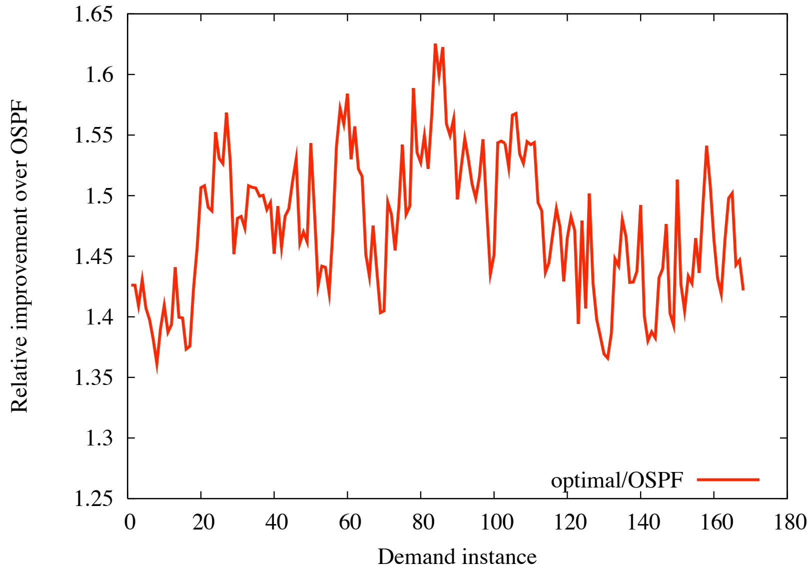

Figure 2.6: All commodities are of equal value; fraction improvement

planning [25], *etc.*

Robust optimization and related approaches have also been applied to computer routing and network flow problems [10, 11, 8, 68, 18]. For example, Applegate and Cohn [8] look at minimizing the maximum link utilization over a set of feasible demand realization in a network using MPLS. Chekuri [18] provides a survey paper of the work to date on using robust optimization to create routing policies. Chekuri characterizes the work of Applegate and Cohn as working on *oblivious routing*, meaning that they design a routing policy that is used for a set of possible demand realizations without knowing the exact realization. The objective of an oblivious routing policy, as described by Chekuri, is to minimize congestion. This is a valid objective function to consider for a Virtual Private Network (VPN) setting, where there is a strict limit on the amount of flow that can pass through any one point in the network. The main distinguishing factor among works on oblivious routing is the structure of the uncertainty set, *i.e.*the set of possible demand realizations, considered. For example, [8] considers a discrete set of demand realizations, while [11] is the first to consider a polyhedral set of demand realizations.

As far as we know, none of the current work in oblivious routing takes into account active congestion control. Though minimizing congestion seems like a good way to accomplish a level of congestion control, it may not capture the tradeoffs that need to take place between different commodities in the network. Moreover, since active congestion control is accomplished, in part, through packet loss, it is not clear whether the routing policies obtained from models that explicitly take active congestion control into account will be the same as those obtained by simply minimizing congestion without modeling congestion control.

A robust counterpart of the problem (2.2) – (2.7) takes into account multiple demand instances ($m \in \{1, \dots, M\}$), and finds the routing policy maximizing the minimum performance over all of these demand instances, mathematically written as:

(2.20)

$$
\begin{aligned}
\text{(RMCFPNG)} \quad \max_{\alpha,\, x} \; \min_{m} \quad & \sum_{k=1}^{K} c^k x^k_{d^k,m} \\
\text{s.t.} \quad & x^k_{o^k,m} = s^k_m && k = 1, \dots, K,\ m = 1, \dots, M \\
& \sum_{(i,j)\in A} \alpha^k_{ij} = 1 && k = 1, \dots, K,\ i \in N : i \neq d^k \\
& \sum_{(i,j)\in A} \alpha^k_{ij} = 0 && k = 1, \dots, K,\ i = d^k \\
& \sum_{i\in\delta^-(j)} \alpha^k_{ij} x^k_{i,m} f_{ij}\left(\sum_{l=1}^{K} \alpha^l_{ij} x^l_{i,m}\right) - x^k_{j,m} = 0, \\
& \qquad k = 1, \dots, K,\ j \in N : j \neq d^k,\ m = 1, \dots, M \\
& \alpha^k_{ij} \geq 0 && (i,j) \in A,\ k = 1, \dots, K.
\end{aligned}
$$

Here, s_m^k is the demand for commodity k in data instance m, and $x_{i,m}^k$ is amount of flow of commodity k present at node $i \in N$ in instance m (note that the values of α's remain the same for all data instances, and thus specify a routing *policy*).

2.6.2 Numerical Results

Robust MPLS Routing policy: First Week

To formulate the robust counterpart of MCFPNG as given by (2.20), we need to specify the uncertainty set, *i.e.*the collection of demand instances to be considered. In considering the data for the Abilene network, we observed that the demand fluctuated following, to a large extent, a daily pattern. Roughly speaking, demand during the day was significantly higher than during early morning and night hours, which is to be expected. (This pattern was less pronounced during the weekends, but was still present.) Based on this observation, we constructed an uncertainty set with three demand instances as follows. We considered the demand data for the first week of usage, separated each of the 7 days into three eight-hour intervals (morning, day, and night), and averaged the demand over these eight-hour intervals across the week. Thus, the three demand instances included in the uncertainty set reflect average hourly demand during mornings, days and nights during the first week of usage.

The resulting robust problem (2.20) has $M = 3$, and can be solved using MINOS. We then compared the performance of the resulting routing policy on the 168 demand instances considered in section 2.5.2 (recall that each of these instances corresponds to the demand during an hour-long interval during the first week). The results are summarized in Figures 2.7– 2.10. In each plot of this figure, the robust MPLS routing policy is the (possibly local) solution of the single problem (2.20), while OSPF and optimal MPLS policies are the same as were found in section 2.5.2. In particular, each optimal MPLS routing policy has been optimized for the demand instance at hand, and thus changes every hour.

As we mentioned in section 2.5.2, each optimal MPLS policy was obtained by initializing the solver at the robust MPLS policy found by solving (2.20), thus assessing the improvement to the robust policy that can be made by modifying it to suit a particular demand instance. As demonstrated in Figures 2.7– 2.10, it appears that fine-tuning the policy to a specific demand instance makes for very little improvement over the robust routing policy. In particular, Figure 2.10 compares, for each data instance, the improvement over the OSPF routing policy realized by the robust MPLS policy with the improvement over OSPF realized by the MPLS policies optimal for each instance. As expected, the improvement achieved by the robust policy is never better, but the plots are fairly similar: the robust policy improves performance by at least 27% in each instance, compared to the improvement of at least 36% realized by the optimal policies.

To summarize, simply by taking into account the natural demand fluctuations in a network, we propose a single, robust, MPLS routing policy which much better performance than OSPF, and comparable performance with each of the optimal MPLS policies.

Robust MPLS Routing policy: 24 Weeks

In Section 2.6.2 we proposed a robust MPLS routing policy, obtained by considering an uncertainty set considering of three data instances capturing morning, day and night demand patters. Recall that we constructed these three instances by averaging demands in the corresponding 8-hour periods during the first of the 24 weeks of data available to us. Since at first glance the demand follows a weekly, as well as daily, pattern, a natural next step is to assess how the robust policy of Section 2.6.2

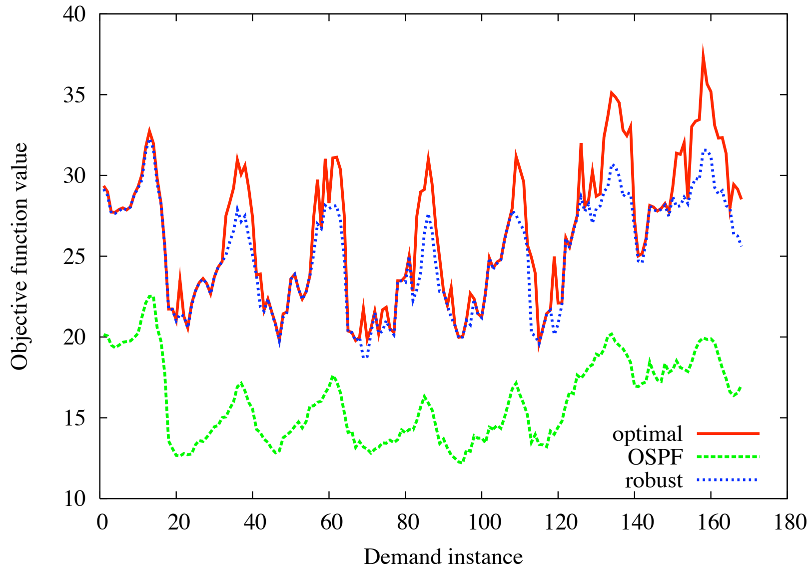

Figure 2.7: Commodities from node 1 are 5 times more valuable than all other commodities

would perform in the following weeks. This is the subject of this section. Here, we limit our discussion to the case where all commodities are of equal value, but the results discussed carry over to the other cases as well. At this point it is important to note that 4 days of the 24 weekly datasets obtained overlap due to the way the data was partitioned at the time it was collected. This means that for the plots in Figures 2.11–2.14, 96 data instances (2.4% of all the instances) appear twice every time a robust routing policy is evaluated.

The plot in Figure 2.11 presents the performance of the robust routing policy obtained in Section 2.6.2 (based on week 1 data) relative to the performance of the OSPF routing policy for each hour-long period in the 24-week period. This is done, as before, by plotting the ratio of the objective function values of the robust policy and the OSPF policy for each demand instance. If the ratio is greater than one, the robust policy performs better than OSFP for that demand instance, as is the case

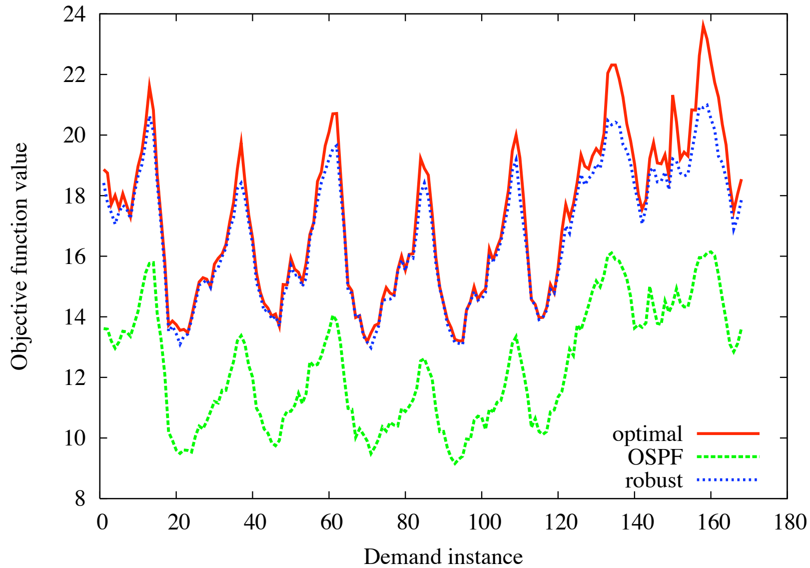

Figure 2.8: Commodities from node 10 are 5 times more valuable than all other commodities

for 86% of the instances.

A closer study of Figure 2.11 reveals that the robust policy based on week 1 data performs better than the OSPF routing policy on all instances except for weeks 3, 4 and 5, which correspond to the period from March 15 until April 5 of 2004. Examining the data, we noted that these three weeks not only have the greatest demand over all twenty-four weeks, but also have somewhat different demand patterns than week 1. For example, 35% of the total demand in week 4 originated from a node that provided only 5% of the total demand in week 1. As the robust policy is catered to the week 1 demand patterns, it is no surprise that it did not perform well when the demand pattern was significantly different. (Unfortunately, we were unable to identify a cause or an explanation of such uncharacteristic demand patterns during this time period.)

Based on the above observations, we computed a different robust routing policy,

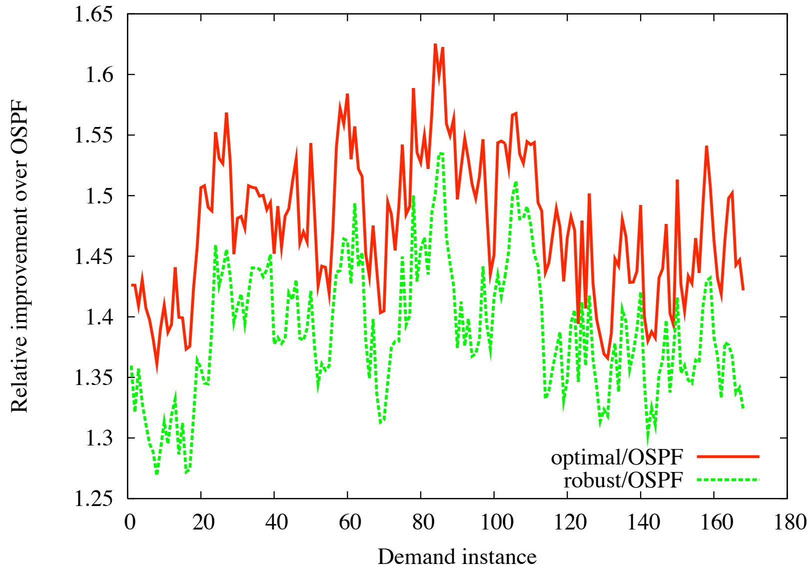

Figure 2.9: All commodities equal value

basing it on week 4 data, and plotted its performance (relative to that of the OSPF policy) in Figure 2.12. Notice that this robust policy performs well (i.e., better than OSPF) *only* during weeks 3, 4 and 5, further confirming that this behavior is due to the fact the demand pattern differs between these weeks and the rest of the time period considered.

We still hypothesized that, aside from anomalous behavior in weeks 3, 4 and 5, overall demand remains stable from week to week due to the self-similar nature of demand patterns [27, 51]. Indeed, the first robust policy continued to perform well in week 6 and beyond. We also looked at the performance of the robust routing policy based on week 12 data in Figure 2.13. This policy performs better than the other two routing policies, and seems to perform at least as well as the OSFP routing policy on most of the data instances. There is still, however, a drop in performance in weeks 3, 4 and 5.

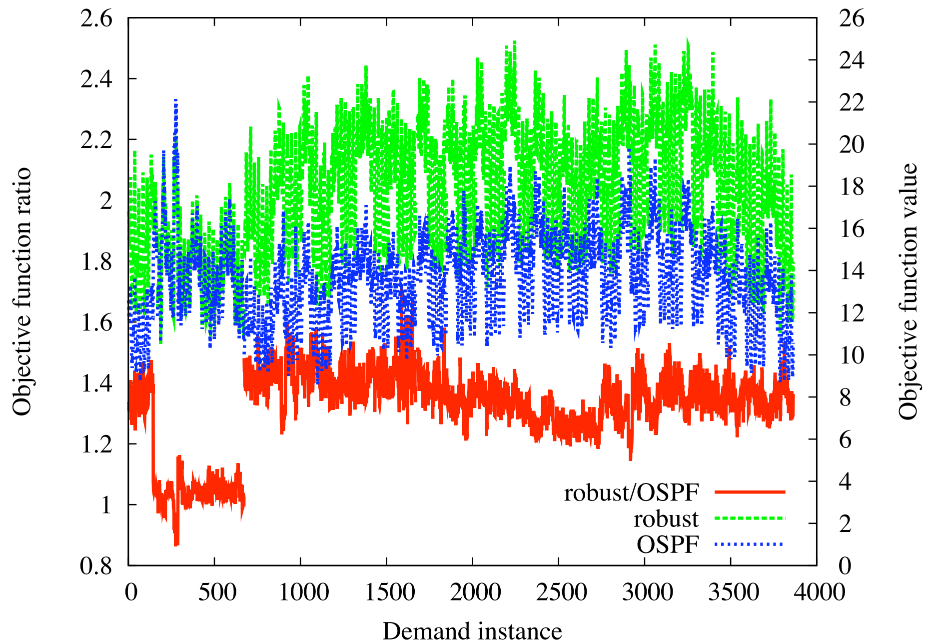

Figure 2.10: All commodities equal value, fraction improvements

Finally, Figure 2.14 shows that the objective function values of the robust routing policy based on week 12 data and the OSPF policy over time follow the same trend, suggesting that, if the ratio of the objectives is greater than one, it is due to the improvement achieved by the robust policy over OSPF, not deterioration of OSPF performance (the same is observed with the other robust routing policies). Figure 2.14 also suggests that a sharp reduction in the performance of a currently implemented robust policy could serve as an indication that the demand pattern is changing and the robust routing policy needs to be reevaluated using more current demand data.

2.7 Conclusion and Future Work

We introduced a way to model congestion control in networks facing congestion via the multi-commodity network flow problem with nonlinear gains (MCFPNG), and proved that the resulting model is $\mathcal{NP}$-hard. We applied the model to the

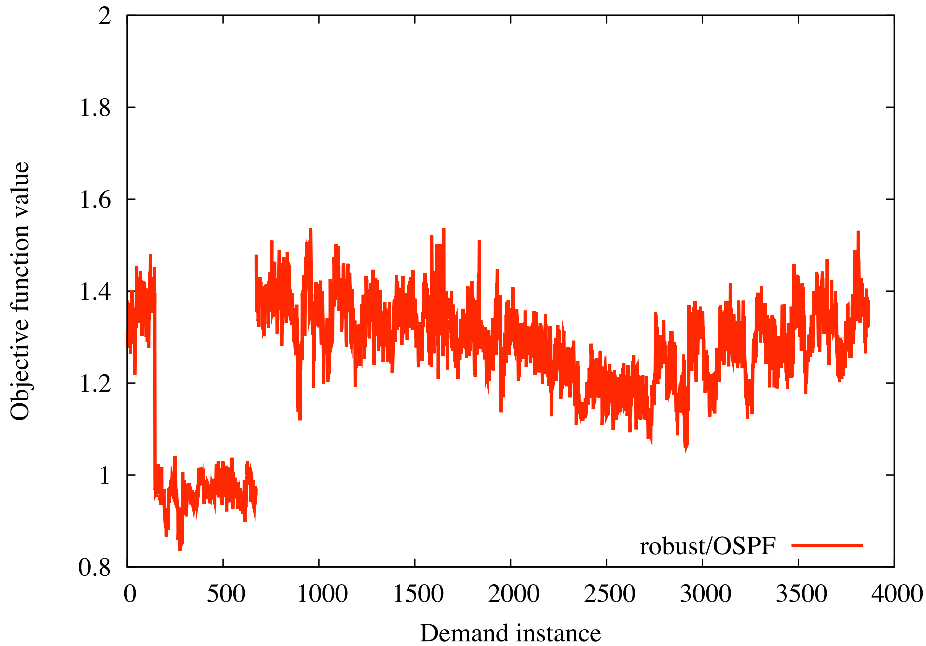

Figure 2.11: Relative performance of the robust routing policy based on week 1 data and the OSFP policy

Abilene network and showed that a better routing policy can be developed by taking into account demand fluctuations — either by designing optimal routing policies for the demand at hand, or, when the above is not possible or desirable, designing robust routing policies.

As presented, our results only give relative performance guarantees by showing empirically that the robust routing polices relatively outperform the OSPF routing policy currently used in the Abilene network. In the future we would like to give absolute performance guarantees on any routing policy we generate by using an approximation algorithm to solve our model.

As a further avenue of research, we would like to consider applications of our model in other areas of network routing. For instance, a modified version our model could be used to find the minimum power levels for wireless routers so as to meet a predetermined performance guarantee. For example, we would like the examine a

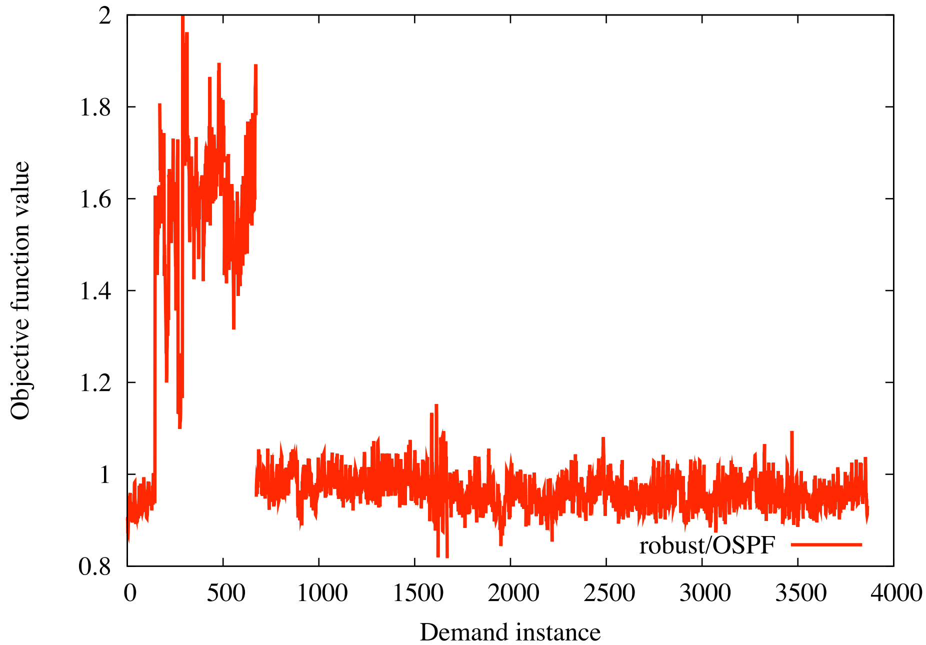

Figure 2.12: Relative performance of the robust routing policy based on week 4 data and the OSFP policy

problem such as: what is the minimum power need in a wireless network such that no more than 5% of packets in the network are lost? We believe that we can extend the model we presented to find a near optimal solution to these types of problems.

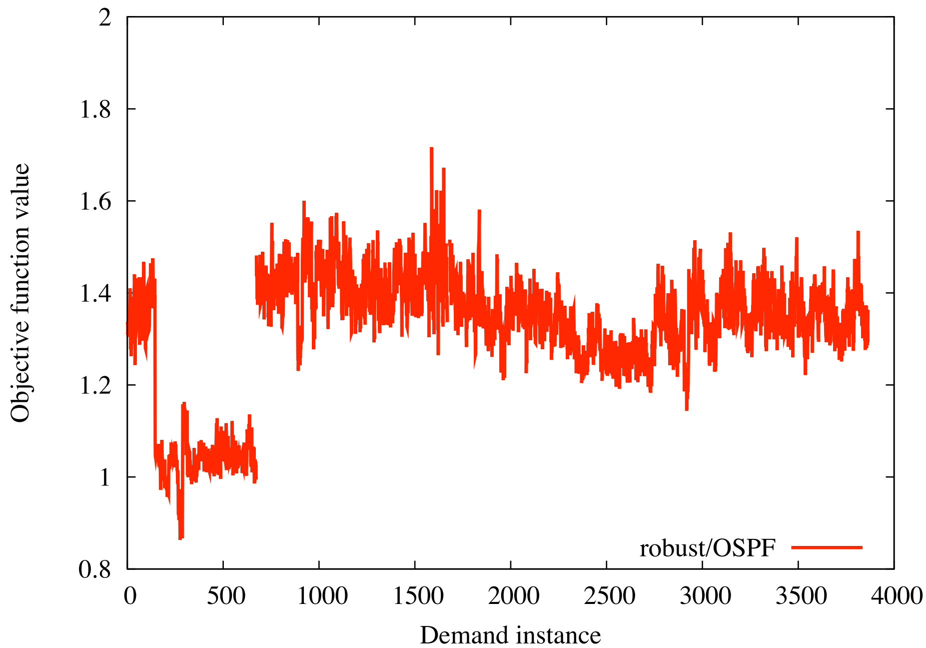

Figure 2.13: Relative performance of the robust routing policy based on week 12 data and the OSFP policy

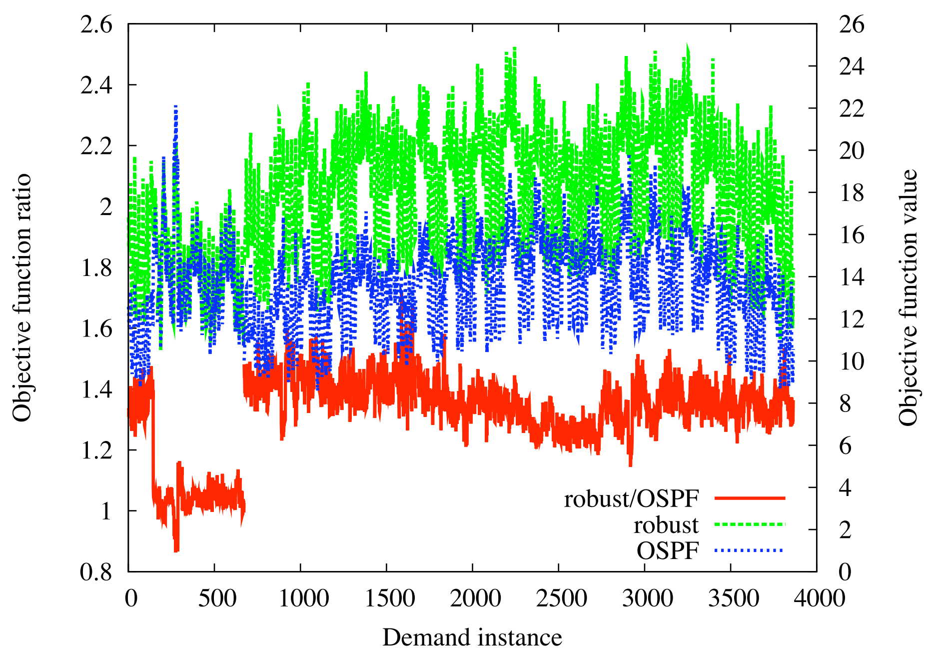

Figure 2.14: Left axis: Relative performance of the robust routing policy based on week 12 data and the OSFP policy; Right axis: Objective function values of the two policies

CHAPTER III

Introduction to Prediction Markets

In this chapter we introduce prediction markets in greater detail, and present some of the notation we will use in Chapters IV and V.

3.1 Prediction Markets Overview and Notation

A prediction market is an information aggregation tool used to elicit and aggregate participants' beliefs regarding the outcome of a future event. Prediction markets are used to predict future political and social events in markets such as InTrade [47], Iowa Electronic Market [46], and Hollywood Stock Exchange [45]. Prediction markets are also used in corporate settings. For example, Microsoft used a prediction market to determine if an internal product would meet its delivery date [53]. Through this market, the director learned that the project was behind schedule. Best Buy is also using prediction markets to assess the success of new products and ideas [31]. One of their prediction markets accurately forecasted the sales of a new laptop service package offered by the company.

Prediction markets can be unsubsidized or subsidized. The InTrade market is an unsubsidized prediction market in which traders use real money to trade. The Microsoft prediction market is a subsidized prediction market in which each trader is initially endowed with \$50. We will discuss why a corporation would be interested in subsidizing prediction markets later on in this chapter.

The prediction markets above aggregate participants' beliefs by incentivizing them to reveal their private beliefs about the event. For example, assume there are n traders in the market and participant i has private information, s_i, pertaining to the outcome of a future event ω. For ease of exposition, we assume the future event is binary, $\omega = 1$ or $\omega = 0$. An individual running a market contingent on ω is interested in eliciting the probability of ω given all of the traders' private beliefs, $P\{\omega|s_1, s_2, \ldots, s_n\}$. Prior to any trades in the market, each trader i has a private belief on ω, denoted $P\{\omega|s_i\}$. With every trade, all traders may update their private beliefs based on trades previously executed in the market. Trader i, after observing all of the first k trades, has a new updated private belief, $P\{\omega|r_1, r_2, \ldots r_{i-1}, s_i, r_{i+1}, \ldots r_k\}$, that reflects her assessment of other traders' private information and the executed trades (r_j is the trade of trader j).

In practice, traders in prediction markets trade securities whose ultimate value is contingent on the outcome of future events. For example, if a trader buys a "*Yes*" contract for a security contingent on ω, then she stands to earn \$1 if $\omega = 1$ (the event occurs) and \$0 if $\omega = 0$ (if the event does not occur). Therefore, assuming traders are risk neutral, the current market price in real world markets is interpreted as the market consensus on the event taking place [79].

Though every trader initially has her private belief $P\{\omega|s_i\}$, this belief, or any subsequently updated belief after trading, may not be entirely revealed when trading in a prediction market due to the trader's risk preferences. If this situation arises, then prediction markets may not be accurately aggregating all of the traders' beliefs. We will discuss this issue along with different forms of prediction markets in the next section.

3.2 Types of Prediction Markets

Prediction markets can be either unsubsidized or subsidized. There are reasons to use one over the other, such as using subsidized prediction markets to guarantee market liquidity. On the other hand, an unsubsidized prediction market does not require any subsidy for information to be aggregated. We discuss the differences in greater detail below.

3.2.1 Unsubsidized Prediction Markets

Unsubsidized prediction markets are similar to financial markets where orders to buy and sell securities are matched in a Continuous Double Auction (CDA). In a CDA prediction market, an agent willing to sell a stock sets an *ask* price and an agent willing to buy a stock sets a *bid* price. If the ask price is less than or equal to the bid price, then the two orders are matched and the orders are cleared. In these markets trade continues until no trader would like to trade. At this point, the market is said to reach equilibrium and the market price is said to be the equilibrium market price. However, "no-trade" theorems can be used to characterize such equilibria (assuming traders are maximizing their utility from participating in the CDA) [41, 49, 54, 65, 74, 75] (see the survey paper by Sent [66] for a detailed overview).

To summarize these "no-trade" results, if all of the agents are rational and have the same risk averse preferences, then an agent willing to trade is informing the other traders that she has valuable information and will be making a profit at their expense. Therefore, those trades would not be willing to trade. This means that a trader cannot profit from gathering more information in the market. This may seem like an artificial construct, as trade occurs daily in financial markets, but these results have been show in laboratory experiments by Angrisani *et al.* [7]. In the laboratory experiment, the authors show that participants, with the same risk preferences, initially trade, but as the number of trading rounds increases, the participants approach a no-trade situation.

Both theoretically and in the experiment by Angrisani *et al.* traders need to have the same risk preferences for "no-trade" results to hold. In practice, traders have heterogeneous risk preferences. Preference heterogeneity may explain why financial markets have such a high volume of trade. As we will see in the next section, having traders with different risk preferences, within CDA prediction markets, does not guarantee traders' beliefs will be accurately aggregated.

CDA prediction markets also run into the issue of elicitation, *i.e.*, traders may not have an opportunity to reveal their information in the market. Consider the case where only one trader is present in a CDA prediction market. Because this trader is the only one present, there is no possibility for her to trade with anyone else in the market. Therefore, any information this trader has is not captured by the market.

Risk Aversion

In CDA prediction markets, if more than one trader is present in the market, and all are willing to trade, the reported beliefs on the contingent event might not be their true beliefs. In fact, risk averse traders will report their risk neutral probability estimates of the event, as discussed by Kadane and Winkler [48]. When a trader is risk averse, her risk preference and belief become intertwined. For example, if a trader has private belief of $P\{\omega|s_i\}$ on ω, and a utility function $u(\cdot)$, then, as described by Nau [56], she would report $\hat{P}\{\omega|s_i\}$, where $\hat{P}\{\omega|s_i\} \propto P\{\omega|s_i\}u'(\omega|s_i)$, with $u'(\omega|s_i)$ being her marginal utility if the event occurs conditional on her information s_i. The risk neutral probability of a trader is the probability an observer would assign to the trader's beliefs given her trading behavior and assuming the trader were risk neutral.

Further, it can be shown that even though price is converging in a market, the market may only be converging to a consensus in risk neutral probabilities and not the true probabilities [57]. This result has the implication that the revealed consensus in a CDA prediction market is not the consensus that would be reached if all of the traders traded according to their private beliefs, $P\{\omega|s_i\}$. If the reached consensus is in terms of the risk neutral probabilities and the risk references are unknown, then we cannot calculate the true probability from the risk neutral probabilities.

Pennock [62] offers a detailed exploration of the behavior of traders with different risk preferences in a CDA setting. He points out that, although traders are risk averse, they will trade with respect to their true probabilities. This means that, if the current market consensus is less than their true probability, a trader, regardless of risk preferences, would have a positive demand for this security. However, the risk preferences determine the extent of each trader's demand. From this result it follows that *if the traders all have the same initial beliefs on the traded events*, then the market will converge to the truthful probability. He then goes on to characterize the equilibrium of a market if all of the agents have constant absolute risk aversion or generalized logarithmic utility. These results allow the market equilibrium to be interpreted correctly by a decision maker.

Prediction markets are used for information aggregation in the corporate setting [24]. Since CDA prediction markets may not be incentivizing trade due to the "no-trade" theorems, and have issues with eliciting traders' beliefs when there are few of them trading, most corporate predictions markets are subsidized. In the next section we review a few forms of subsidized prediction markets, and discuss in detail market scoring rule prediction markets.

3.2.2 Subsidized Prediction Markets

For someone interested in eliciting the beliefs of all participants in a prediction market regardless of the number of traders present, a way to ensure this elicitation is to subsidize the market via a market maker. In a prediction market using a market maker, all traders effectively trade with the market maker so that market liquidity is guaranteed. Subsidizing the market ensures that a no-trade situation does not arise. However, if participants are risk averse, traders will still only report their risk neutral probabilities.

A simple market maker is one that accepts all bid and ask offers in a CDA setting. The issue with such an approach is that the amount of subsidies introduced into the market may not be bounded.

Another market maker prediction market is the dynamic pari-mutuel market [63]. The dynamic pari-mutuel market is based on traditional pari-mutuel markets used in horse racing, but allows for early sale at a dynamically changing price in order to encourage early trade by informed traders. Standard pari-mutuel markets distribute the winnings proportionally to the total wager of the winning horse. For example, if Alice places \$10 on the winning horse and Bob places \$5 on the winning horse, then Alice receives $\frac{2}{3}$ of the total winnings and Bob receives $\frac{1}{3}$. In dynamic pari-mutuel markets, the proportion of winnings is also dependent on when the bet was made; in the example above, if Bob made the bet on the winning horse (the event that occurred) before Alice then he would receive more than $\frac{1}{3}$ of the total winnings.

A type of subsidized prediction markets is market scoring rule (MSR) prediction markets, introduced by Hanson [42]. (Proper) Market scoring rules are derived from (proper) scoring rules. The notion of scoring rules was introduced by Brier [16], in the form of the quadratic scoring rule (which is proper), to measure the accuracy of weather forecasters. *Proper* scoring rules provide a way to reward or evaluate a forecaster in a way that motivates honest probabilistic forecasts. Other proper scoring rules, like the logarithmic scoring rule, were later introduced again to assess the quality of weather reports [76, 77]. Most of the early work on scoring rules assumed that weathermen were expected score maximizers, *i.e.*, risk neutral. If forecasters are risk averse then, as described by Kadane and Winkler [48], the forecasters will report their risk neutral probabilities just like they would in a CDA prediction market. We revisit this discussion in Chapter V when we analyze the impact of risk averse traders in subsidized prediction markets.

Hanson [42] showed that if a trader is risk neutral and does not take into account future payoffs, then she will report her true belief in a MSR prediction market, and the amount of market subsidy in the market is bounded regardless of the number of traders in the market. Moreover, as MSR markets are subsidized, these results hold whether there is a large or a small number of traders present. Unfortunately, this result hinges on two key assumptions:

1. Traders do not take into account future payoffs (in this setting we can think of this as traders only trading once in a MSR market).

2. Traders are risk neutral.

Both of these assumptions do not hold in practice. Traders may trade any number of times in a prediction market, thus giving them the opportunity to take into account future payoffs, and traders are inherently risk averse. Therefore, in Chapter IV we analyze the effect on MSR prediction markets of allowing traders to take into account future payoffs, and in Chapter V we analyze the impact of risk averse traders in prediction markets.

MSR prediction markets may be presented as financial markets with traders buying and selling securities with the market maker. However, underlying this interface is a price function (a price function determines the cost to buy and sell securities in the market given the current market state) derived from scoring rules. Therefore, it helps to think of these markets as traders simply reporting, in turn, to the market maker a probability, their assessment on the probability of the event taking place, and being scored with respect to this report. We denote by r_t the report of the tth trader in the market. For example, using the log proper MSR, if the tth trade in the market, made by trader A, consists of a report r_t when the previous report was r_{t-1}, trader A receives a reward that is an affine function of $\log(r_t) - \log(r_{t-1})$ if the event occurs and a reward that is an affine function of $\log(1-r_t) - \log(1-r_{t-1})$ if the event does not occur. For example, consider a risk neutral trader trading only once in a market using the log MSR. She observes the current market position, the last made report to the market maker, to be 0.3 and has an updated belief of 0.8 of the event occurring. As this situation satisfies all of the conditions defined by Hanson [42], this trader will report her true belief. If the event occurs, she stands to make a profit of $\log 0.8 - \log 0.3 = \log \frac{0.8}{0.3}$, and if the event does not occur, she stands to make a loss of $\log 0.2 - \log 0.7 = \log \frac{0.2}{0.7}$, since making a report of 0.8 of the event occurring is equivalent to making a report of 0.2 of the event not occurring.

3.3 Dynamics and Model

For the remainder of the dissertation we will only consider subsidized prediction markets. In Chapter IV we consider market scoring rule prediction markets, and in Chapter V we consider all subsidized prediction markets. Before we present our results we will discuss how we model subsidized prediction markets.

As previously mentioned, the purpose of prediction markets is to aggregate traders' beliefs on the outcome of an event. Mathematically, we are interested in finding $P\{\omega|s_1, s_2, \ldots, s_n\}$, where s_i is the signal, or private information, of trader i and ω is the event traded in the market. Initially we assume that every trader has a common prior on the joint probability of the event and all of the other traders' signals, *i.e.*, $P\{\omega, s_1, s_2, \ldots, s_n\}$, $P\{\omega, s_i\}$, and $P\{s_i\}$ $\forall i \in \{1, \ldots, n\}$. Every trader observes her private signal and forms a belief on the outcome of the event conditional on that signal, *i.e.*, the belief of trader i is $P\{\omega|s_i\}$.

Without loss of generality we assume that traders trade in turn (trader 1 trades first, then trader 2, *etc.*). We do not assume the traders necessarily know the identity of traders 1, 2, and so forth. Once trader 1 trades, all of the other traders update their beliefs conditional on her trade. For example, assume that trader 1 is myopic and risk neutral. This means she would report $r_1 = P\{\omega|s_1\}$, as every other trader knows $P\{s_1\}$ and $P\{\omega, s_1\}$, the other traders can deduce trader 1's signal realization. Even if the first trader does not honestly report her belief and instead reports a different value of r_1, any trade she makes reveals some information on her signal realization. Therefore, all other agents will update their belief on ω in a Bayesian manner given r_1. With these updated beliefs, the second trader will trade and all of the other traders, including the first, will update their beliefs. This continues until no trader trades and the market is said to reach its equilibrium price.

In this dissertation we consider two slightly different variants of the dynamics above. In Chapter IV we consider a setting where every trader knows who the other traders are. That is, if Alice is trader 1, Bob is trader 2, and Carol is trader 3, then when trader 1 trades, all of the traders know Alice made a trade. Similarly if trader 3 trades, then all of the traders know Carol made a trade. On the other hand, in

Chapter V, we assume the traders cannot tell who is making the actual report. The main difference between these two settings is that if a trader were to trade repeatedly, in the first setting (Chapter IV) traders can tell that the same trader made multiple trades, and in the second setting (Chapter V) they cannot.

We can now introduce the notion of an *information structure*, defined in greater detail in Section 5.2. An information structure describes the truthful prediction ($P\{\omega|s_1, s_2, \ldots, s_n\}$) for all possible signal realizations for n traders, assuming each trader trades in the market once. For example, for $n = 3$ and for every trader having an equal probability of observing one of two signals (*left, right*), one possible information structure is seen in Figure 3.1. In the figure, if trader 1 observed a *right* signal and trader 2 observed a *right* signal, then the truthful prediction given those observations would be 0.7. At that point, if the third trader observed a *left* signal, then the truthful prediction would be 0.65. Had the signals observed by the three traders been different, a different sequence of probability updates would result — all possible outcomes are presented in the information structure.

The notation introduced in this chapter is common among the following two chapters. In every chapter we will expand and simplify this notation as needed so as to best convey the results.

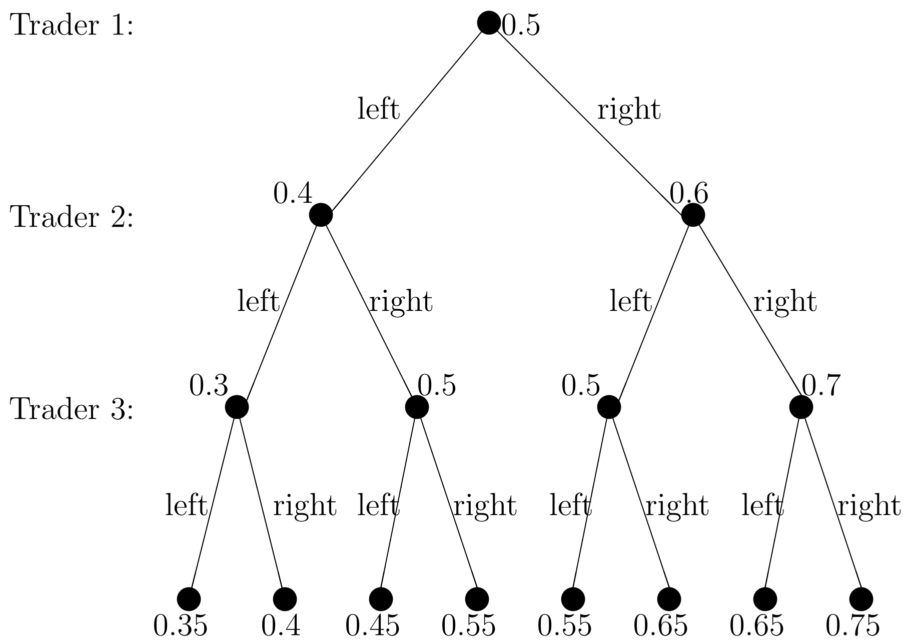

Figure 3.1: An example of an information structure

CHAPTER IV

Non-myopic Strategies in Prediction Markets

This chapter is based on joint work with Rahul Sami; parts were reported in [28, 29].

4.1 Introduction

As discussed in Chapter III, prediction markets may be used to aid in decision making by aggregating the beliefs of participants on the outcome of a future event. Ideally, the market participants truthfully reveal their beliefs, and all of their private information is captured by the market. In this chapter, we characterize a class of situations where risk neutral traders have an incentive to deviate from the truthful strategy is a market scoring rule prediction market.

The successful aggregation of information through prediction markets thus relies critically on traders adjusting their beliefs in response to other traders' trades. However, this responsiveness can also have a drawback in the operation of the market: A trader may attempt to first mislead other traders about the value of the security, and then exploit their inaccurate information in later trades. Awareness of, and reaction to, this problem can lead traders to be overly cautious about making inferences from market prices, thus weakening the aggregative powers of the market. As a result, prediction markets have always had to grapple with this *perceived* threat of manipulation, even when actual manipulation is absent. It would be very useful to have a characterization of market situations in which such manipulation is possible (or impossible); due to the strategic complexity of traditional double-auction markets, such characterizations have been difficult to achieve.

With the recent rapid growth of markets designed primarily for information aggregation, researchers have developed new market designs that are tailored to incentivize informed agents to trade and to reveal their private information in a timely manner. Hanson's market scoring rule [42] is an innovative tradable security; it is based on the idea of a *proper scoring rule* [16]. Pennock's dynamic pari-mutuel market [63] is another new market design that is based on the traditional pari-mutuel market form used in horse racing, but allows for early sale at a dynamically changing price in order to encourage early trade by informed traders. Apart from their other advantages, these new market forms are promising for another reason: As one side of each individual trade is held by an automated market maker with a predetermined (and fairly simple) strategy, these market forms are much more amenable to formal analysis. For the market scoring rules, it has been proven that honest revelation of private information is myopically optimal [42]. A similar (although slightly weaker) characterization of myopically optimal strategies in dynamic pari-mutuel markets is reported by Nikolova and Sami [58]. However, much of the concern about manipulation in prediction markets is based on *non-myopic strategies*: strategies in which the attacker sacrifices some profit early in order to mislead other traders, and then later exploits erroneous trades by other traders, thereby gaining an overall profit. As yet, very little is known theoretically about the existence and characterization of manipulative non-myopic strategies in these markets.

Our Results

In this chapter, we study trading strategies in the logarithmic market scoring rule prediction market. We model a general Bayesian framework in which traders receive information signals relevant to the event to be predicted, and trade in the prediction market to maximize their expected payoffs. Our model captures the fact that traders learn from prior trades as well as their own signals. This way, the market itself is represented as an extensive form game played between partially informed traders. The logarithmic market scoring rule allows the traders' moves and profits to be connected to the information-theoretic notion of *entropy*. Our analysis builds on this connection, and we show that it allows meaningful analysis of the informativeness of market prices.

We show that, if traders' initial signals are independent, it is generically true that the myopically optimal strategy of trading honestly is *not* an equilibrium of this extensive-form game. In other words, if a trader believes that future traders will believe that she is playing myopically, she can profit by dishonest trading. Thus, we demonstrate that strategies that involve deception of future traders are a real possibility under a wide range of information conditions.

We propose a simple scheme, the *discounted market scoring rule*, in which traders' payoffs for market transactions are explicitly discounted over time. This reduces the potential gain from correcting a misled trader, thereby reducing the threat of deceptive, non-myopic strategies. We analyze the market game in the presence of discounting, and show that, although non-myopic trading might still be profitable, the market converges to the optimum value (the probability of the event given the traders' signal realizations) in a very strong sense: In any equilibrium, the relative entropy of the actual market price with respect to the optimal market price decreases exponentially over time, at a rate that can be lower-bounded in terms of the discount factor. For a market operator who is running a prediction market to aggregate all known information about a particular event, this provides a way to quantify and limit the uncertainty in price accuracy due to non-myopic bluffing strategies. Our analysis also reveals conditions under which the myopic strategy is in fact the only equilibrium strategy.

Related Work

There have been several field and experimental studies of manipulation in prediction markets. Strumpf and Rhode [71] conducted experiments on manipulating prices in the Iowa Electronic Market. Hanson *et al.* [44] experimentally study whether agents with an incentive to manipulate prices can influence the trading price of a security. They found that other agents who were aware of potential manipulation adjusted for this possibility, thus limiting the effects of the manipulation attempts.

There is rich literature on manipulation in financial markets, which are closely related. This literature has studied manipulation based on releasing false information (perhaps through trades in other markets), as well as manipulation that only requires strategic manipulation in a single market; the latter form of manipulation is closely related to our study here. Allen and Gale [6] describe a model in which a manipulative trader can make a deceptive trade in early trading rounds, and then profit in later rounds, even though the other traders are aware of the possibility of deception and act rationally. They use a stylized model of a multi-period market; in contrast, we seek to exactly model a market scoring rule model. Apart from other advantages of detailed modeling, this allows us to construct simpler examples of manipulative scenarios: The model in [6] needs to assume traders with different risk attitudes to get around no-trade results, which is rendered unnecessary by the inherent subsidy in the market scoring. Our model requires only risk neutral traders, and exactly captures the market scoring rule prediction markets. We refer readers to the paper by Chakraborty and Yilmaz [17] for references to other research on manipulation in financial markets.

Feigenbaum *et al.* [33] also study prediction markets in which the information aggregation is sometimes slow, and sometimes fails altogether. In their setting, the aggregation problems arise from a completely different source: The traders are nonstrategic, but extracting individual traders' information from the market price is difficult. Here, we study scenarios in which extracting information from prices would be easy if traders were not strategic; the complexity arises solely from the use of non-myopic strategies.

Nikolova and Sami [58] present an instance in which myopic strategies are not optimal in an extensive-form game based on the market, and suggest (but do not analyze) using a form of discounting to reduce manipulative possibilities in a prediction market. We draw on a generalization of this instance as the starting point of our analysis. Plott *et al.* [9] also proposed a form of discounting in an experimental pari-mutuel market, and showed that it promoted early trades. Unlike the pari-mutuel market, the market scoring rule has an inherent subsidy, so it was not obvious that discounting would have strategic benefits in our setting as well.

Our work is most closely related to independent work by Chen *et al.* [22, 23]. Chen *et al.* study non-myopic strategies in prediction markets; their initial results [23] were reported at the same time as the preliminary version of our results [28]. They study a similar Bayesian model of a market scoring rule market, with an information structure that differs in one key aspect from ours. Our nonexistence of myopic equilibria results assume that traders' signals are independently generated, and that different combinations of the signals lead to different expectations of the event occurring. Chen *et al.* model signals as *conditionally independent*, conditioned on the eventual truth of the event under consideration. This difference in models leads to opposite results: they show that, under their model, following the myopic strategy is an equilibrium strategy. This difference can be represented by looking at the two settings in terms of complementary and substitutable information. In the conditionally independent setting traders' signals are substitutable. In our model, the signals are complementary. When traders have substitutable signals, it becomes a race between traders as to who will earn most of the profit available in the market by truthfully revealing their information. However, when the signals are complementary, a trader stands to earn a larger reward if she knows the other trader's signal. Therefore, the first trader will bluff in hopes of having the second trader reveal some information about their signal and thus the first trader stands to gain a larger profit than by being truthful during the first trade. Both of these conditions can occur in practice. For example, if traders are trying to predict the outcome of an election where every trader is a voter, then a trader, knowing all of the other traders' information, will earn all available profit in the market. However, if traders instead are each asked the probability oil will be found in a particular well, then the traders might have substitutable information. Further, Chen *et al.* [22] construct an example three-round market in which the conditional independence condition does not hold, and show that it admits an equilibrium strategy in which the first trader bluffs with some nonzero probability.

Börgers *et al.* [15] study when signals are substitutes and complements in a general setting. Our analysis and convergence result suggests that prediction markets are one domain where this distinction is of practical importance.

Structure of the Chapter

The rest of this chapter is structured as follows: In Section 4.2 we describe the 2-player model we use to highlight deception threats, and we introduce some equilibrium concepts. In Section 4.3, we formally analyze the simple 2-player model and show that there exists no finite equilibrium in this setting. In Section 4.4 we generalize the 2-player model to any finite number of players and signals and extend the result to other scoring rules. In Section 4.5 we show that a simple discounted market scoring rule reduces the opportunity for non-myopic strategies, and the market price converges to the optimal price at a rate bounded in terms of the discounting parameter. In Section 4.6 we discuss how our results may be generalized and used to gain insight about more complex markets. We draw parallels with classical bargaining theory, and sketch directions for future research.

4.2 A Simple Model: Two Players and Two Signals Each

In this section, we describe a model of an extremely simple prediction market setting. The setting is as follows: A prediction market is designed to predict a future event ω, by trading in a security F based on ω. Two players, P1 and P2, are each endowed with some private information about ω. We assume the simplest possible case, in which P1 and P2 each have a single bit of information (s_1, s_2 respectively) relevant to ω. (Equivalently, they each receive a signal that can take on two possible values: 0 or 1). Further, we assume that the traders are risk neutral, and share a common (and accurate) prior probability distribution over P1, P2, and ω.

The prior probability distribution can be completely specified by specifying the prior probabilities of the signals and the conditional probability of ω given each combination of signals. We assume that the two signals are independent. Further, we assume for simplicity that s_1 is 0 or 1 with equal probability. The probability that s_2 is 1 is given by a parameter $0 < \kappa < 1$. Thus, the model can be fully specified by specifying four probabilities $p_{00} = P\{\omega | s_1 = 0, s_2 = 0\}$ (or $P\{\omega|00\}$ for short), $p_{11} = P\{\omega|11\}$, $p_{01} = P\{\omega|01\}$ and $p_{10} = P\{\omega|10\}$. We study the behavior of the market for different values of the parameters $p_{00}, p_{11}, p_{01}, p_{10}$ and κ. Note that p_{ij} may be thought of as the probability of ω given signals i and j. We summarize the probability of ω given the players' signals in Table 4.1:

Table 4.1: Probability Realizations

Probability	Signals
p_{00}	00
p_{01}	01
p_{10}	10
p_{11}	11

We assume that the trade in security F is conducted using a market scoring rule [42]. Players make a sequence of market moves; in each move, the player reports a probability r_i. At the end, when the event ω is revealed, the move earns a player a net score $\text{Score}(\omega, r_i) - \text{Score}(\omega, r_{i-1})$, where S is some proper scoring rule. In this chapter, we assume the *logarithmic scoring rule*. The market maker seeds the market with a value p_s which is irrelevant to our analysis. We consider a sequence of alternating moves in which P1 moves first, P2 moves next, P1 potentially moves again, and so on.

In a market using the logarithmic scoring rule, the score of any one move is a constant multiple of $\log r_i - \log r_{i-1}$ if ω occurs. Without loss of generality we assume that the constant multiple is 1 for our analysis of the market scoring rule. In Section 4.5, we propose a scheme in which the constant multiplier changes over time.

4.2.1 Myopic Behavior

We now analyze the price dynamics if each trader followed her myopically optimal strategy. There are two additional probabilities ϕ_0 and ϕ_1 that arise in the analysis of the myopic behavior because of P1's uncertainty of P2's signal. Suppose P1 saw $s_1 = 1$. She would then condition her prior on this information, resulting in a posterior in which she ascribes probability κ to the possibility that the optimal probability (the probability of ω given the traders' signals) is p_{11}, and probability $1-\kappa$ to the possibility that the optimal probability is p_{10}. In the balance, her belief about the likelihood of ω would be in between that implied by p_{11} and that implied by p_{10}. Therefore, her optimal myopic strategy if she observed $s_1 = 1$ would be to report probability $\phi_1 = \kappa p_{11} + (1-\kappa)p_{10}$, or simply the expected optimal probability conditioned on her seeing 1 as her signal. Likewise, if P1 saw $s_1 = 0$, she would move to a point ϕ_0 defined in terms of p_{01} and p_{00}.

After P1's move P2 cannot directly see s_1, but she can infer what P1's myopic actions would have been in each case. We assume that we are in the non-degenerate case in which $\phi_0 \neq \phi_1$; this allows us to focus on strategic threats instead of difficulties in extracting signals from the price. Then, P2 can infer the value of s_1; combining this with the value of s_2 that P2 observed, she can calculate the best possible estimate of the conditional probability of ω. Due to the myopic strategyproof properties of the market scoring rules, P2 would move to p_{00}, p_{01}, p_{10}, or p_{11}. Subsequently, neither player would have an incentive to move again. Thus, if players followed their myopic strategies, the market would perform remarkably well: All information would be aggregated optimally in just two trades. Further, in general, both players would make a profit in expectation in this market.

4.2.2 Non-myopic Behavior and Bluffing

Now, suppose that the players were not restricted to myopic behavior. Specifically, a player may deviate from the myopic strategy to *exploit the other player's reaction*, and make a greater total profit through subsequent moves. Consider the ways in which P1 can deviate from her original myopic strategy. We restrict our attention to strategies in which P1 moves to either ϕ_1 or ϕ_0 in the first round. These are the two positions that P2 is expecting to see the market in, and thus we can reason about the reaction that P2 would make to the move; this would be difficult if the move was to a different point.

Thus, we are interested in the following kind of bluffing strategies for P1: Suppose P1 sees $s_1 = 1$. She could move to ϕ_0 in the first round, instead of her myopically optimal strategy of moving to ϕ_1. Now, if P2 is expecting myopic behavior, she would incorrectly infer that $s_1 = 0$, and correspondingly report the wrong probability: p_{00} instead of p_{10}, or p_{01} instead of p_{11}. P1 can see the reported probability by P2, and make a subsequent correcting move: $p_{00} \rightarrow p_{10}$ or $p_{01} \rightarrow p_{11}$ respectively. P1's incentive to bluff is determined by the profitability of this bluffing strategy relative to the honest (myopic) strategy. If the myopic strategy is superior to bluffing, for both values of s_1, P1 would follow this strategy. Then, P2 would have no reason to bluff (because P1 would not move again). Thus, checking if the myopic strategy is an equilibrium is equivalent to checking if P1's expected profit from bluffing is less than her expected profit from the myopic strategy, assuming that P2 will be myopic.

Suppose that the bluffing strategy has a strictly higher profit than the myopic strategy for player P1. It follows that P1 will bluff with some probability ζ. Note that P2 can analyze P1's profit in different scenarios, and thus, can infer that P1 would not necessarily be truthful. Now, we characterize equilibria in which the bluffing probability ζ is known to P2, who takes it into account and reacts accordingly. It must be that $0 < s < 1$, because otherwise P2 would know s_1 with certainty. Now, from P2's point of view, the market looks very similar to the market we just analyzed for P1: She sees s_2, and assigns some probability ζ to $s_1 = 1$. The myopic optimal response for P2 *taking into account the probability that P1 is bluffing* can be determined: it is a function of ζ, s_2 and the position (ϕ_1 or ϕ_0) that P1 left the market. Next, we can repeat the analysis from P2's point of view, and determine if the myopic response is optimal for P2, or if she too would rather bluff with some probability. The analysis follows exactly as done for P1, except that the role of s_1 and s_2 are interchanged, or equivalently, the probabilities p_{01} and p_{10} are interchanged.

Next we show that with an informativeness condition for all points of p_{00}, p_{11}, p_{01}, and p_{10} the myopic strategy is not an equilibrium strategy for the two-player, two-valued signal setting and in general for the n-player m-valued signal setting. Using the informativeness condition we show that no finite equilibrium exists in both of the settings for the logarithmic market scoring rule.

4.2.3 Equilibrium Concepts

The prediction market model we have described is an extensive-form game between two players with common prior probabilities but asymmetric information signals. Specifying a plausible play of the game involves specifying not just the moves that players make for different information signals, but also the beliefs that they have at each node of the game tree.

Informally, an *assessment* $A_i = (\sigma_i, \mu_i)$ for a player i consists of a *strategy* σ_i and a *belief system* μ_i. The strategy dictates what move the player will make at each node in the game tree at which she has to move. We allow for strategies to be (behaviorally) mixed; indeed, a bluffing equilibrium must involve mixed strategies.

To avoid technical measurability issues, we make the mild assumption that a player's strategy can randomize over only a finite set of actions at each node. The belief system component of an assessment specifies what a player believes at each node of the game tree. In our setting, the only relevant information a trader lacks is the value of the other trader's information signal. Thus, the belief at a node consists of an assignment of a value to the probability that the other player received a '1' signal, contingent on reaching the node.

An assessment profile (A_1, A_2), consisting of an assessment for each player, is a *weak Perfect Bayesian Equilibrium iff*, for each player, the strategies are sequentially rational given their beliefs *and* their beliefs at any node that is reached with nonzero probability are consistent with updating their prior beliefs using Bayes' rule, given the strategies. This is a relatively weak notion of equilibrium for this class of games; frequently, the refined concepts of Perfect Bayesian Equilibrium or sequential equilibrium are used. Our results involve proving the nonexistence of myopic weak PBEs, and characterizing the set of all weak PBEs. They thus hold *a fortiori* for refinements of the weak PBE concept, including those mentioned above. For a formal definition of the equilibrium concept, we refer the reader to the book by Mas-Colell *et al.* [52].

Given the strategy components of a weak PBE profile, the belief systems of the players are completely defined at every node on the equilibrium path (*i.e.*, every node that is reached with positive probability). In the remainder of this chapter, we will not consider players' beliefs off the equilibrium path. Thus, we will abuse notation slightly by simply referring to an "equilibrium strategy profile", leaving the beliefs implicit.

4.3 Analysis of the Simple Model

Building on the intuition developed in Section 4.2.2, we now consider an analytical proof to show that player 1 has an incentive to bluff. It turns out to be easiest to analyze the logarithmic market scoring rule in this case, because we can reduce it to a standard result on information-theoretic (Shannon) entropy. At the end of the section, we discuss why we expect this result to hold for other scoring rules.

4.3.1 Generic Bluffing

In this subsection, we show that the myopic strategy profile is generically not a weak Perfect Bayesian equilibrium.

In order to show that there is a *strictly* profitable deviation from the myopic strategy profile, we first need to exclude certain degenerate cases. In particular, we restrict our attention to instances that satisfy the following *general informativeness condition*:

Definition IV.1. An instance of the prediction market with n players satisfies the *general informativeness condition* if there is no vector of signals for any $n-1$ players that makes the n^{th} player's signals reveal no distinguishing information about the optimal probability. Formally, for $n = 2$, the following property must be true: $\forall i, \bar{i}, j, \bar{j}$ such that $i \neq \bar{i}, j \neq \bar{j}$: $p_{ij} \neq p_{i\bar{j}}$ and $p_{ij} \neq p_{\bar{i}j}$. For $n > 2$, using the notation of Section 4.4, we must have $\forall \mathbf{i}\ \forall \bar{j} \neq j, p(j, \mathbf{i}) \neq p(\bar{j}, \mathbf{i})$,where $j, \bar{j}$ are two possible signals for any one player, and $\mathbf{i}$ is a vector of signals for the other $n-1$ players.

Consider a game of two players with each player seeing one of two signals. The optimal probability if player 1 sees signal i and player 2 see signal j is p_{ij}. As before we assume that player 2 has a probability of κ of seeing a one. As player 1 is playing first, her honest belief, conditioned on her signal, would be given by:

$$\begin{aligned}\phi_1 &= \kappa p_{11} + (1-\kappa)p_{10} \quad \text{expectation if P1 sees } s_1 = 1\\ \phi_0 &= \kappa p_{01} + (1-\kappa)p_{00} \quad \text{expectation if P1 sees } s_1 = 0\end{aligned}$$

The probabilities ϕ_1 and ϕ_0 determine the optimal myopic moves for player 1.

We first show that, if an equilibrium profile involves deterministic strategies, it must be the myopic strategy profile:

Lemma IV.2. *Consider any equilibrium strategy profile. If player 1 has a deterministic strategy of playing ϕ_u when she receives a 1 signal and $\phi_v \neq \phi_u$ when she receives a 0, then $\phi_u = \phi_1$ and $\phi_v = \phi_0$.*

Proof. Assume that $\phi_u \neq \phi_1$. Whenever player 1 plays ϕ_u player 2 will deduce that player 1 has observed a 1; then, player 2 will capture all remaining surplus. Player 1 thus gets at most the profit she earns from her first move. However, by definition of myopic optimality, ϕ_1 would yield a higher profit to player 1 in the first round. Therefore, player 1 has a profitable deviation in expected payoff from ϕ_u to ϕ_1. A similar argument holds for ϕ_v and ϕ_0. □

Consider the situation in which player 1 observes a 1 as her signal. We will want to compare the expected profits that player 1 could earn through different first-round moves. Assume that the market starts at an arbitrary point p_s. The market scoring rule payoffs are additive, in the sense that the total payoff for two consecutive moves is exactly the payoff of moving from the starting point to the end point of the second move. Now, in order to estimate the relative profitability of bluffing, we can think of the bluffing strategy as a move from p_s to the honest position ϕ_1 followed by a move from ϕ_1 to ϕ_0. Therefore, when comparing the two strategies the initial move p_s to ϕ_1 cancels out. In order to eliminate the irrelevant p_s from our comparison, we treat the myopic move as if it had profit 0, and analyze the incremental profit or loss of the move from ϕ_1 to ϕ_0.

We now express the expected profits in terms of information-theoretic entropy. To this end, we observe that the following two expressions are equivalent:

$S(p_i, p_j)$ The expected log score from moving from position p_i to p_j, with p_j being the true belief:

$$S(p_i, p_j) = p_j \log \frac{p_j}{p_i} + (1 - p_j) \log \frac{1 - p_j}{1 - p_i}$$

$D(p_j||p_i)$ The relative entropy of two probability mass functions $p(x)$ and $q(x)$ is defined in [26] as:

$$D(p||q) = \sum p(x) \log \frac{p(x)}{q(x)}$$

With $p(x) = \begin{cases} p_j & x = \text{Y} \\ 1 - p_j & x = \text{N} \end{cases}$ and $q(x) = \begin{cases} p_i & x = \text{Y} \\ 1 - p_i & x = \text{N} \end{cases}$.

Lemma IV.3. *Assume that Player* 2 *expects Player* 1 *to play honestly, and reacts accordingly. If Player 1 observes* $s_1 = 1$ *and bluffs, her expected increase in score (over following the myopic strategy) is* $qD(p_{11}||p_{01}) + (1 - q)D(p_{10}||p_{00}) - D(\phi_1||\phi_0)$.

Proof. We analyze the change in Player 1's score due to her two moves (deviation and subsequent correction) separately. Given $P1$'s information ($s_1 = 1$), the expected probability of the event happening is ϕ_1. Thus, the expected deviation move score for player 1 is:

$$\begin{aligned} S(\phi_1, \phi_0) &= \phi_1 \ln \tfrac{\phi_0}{\phi_1} + (1 - \phi_1) \ln \tfrac{1-\phi_0}{1-\phi_1} \\ &= -D(\phi_1||\phi_0) \end{aligned}$$

As player 2 has a probability of κ of seeing a one, player 1 will have to have a corrective step from p_{01} to p_{11} with probability κ. Similarly with probability $1 - \kappa$

player 1 will make her corrective step from p_{00} to p_{10}. Therefore the expected score from the corrective step is $\kappa S(p_{01}, p_{11}) + (1-\kappa)S(p_{00}, p_{10})$.

$$\begin{aligned} S(p_{01}, p_{11}) &= p_{11} \ln \tfrac{p_{11}}{p_{01}} + (1-p_{11}) \ln \tfrac{1-p_{11}}{1-p_{01}} \\ &= D(p_{11}||p_{01}) \end{aligned}$$

$$\begin{aligned} S(p_{00}, p_{10}) &= p_{10} \ln \tfrac{p_{10}}{p_{00}} + (1-p_{10}) \ln \tfrac{1-p_{10}}{1-p_{00}} \\ &= D(p_{10}||p_{00}) \end{aligned}$$

□

Theorem IV.4. *Suppose two players are trading in a market with alternate moves; without loss of generality, suppose $P1$ makes the first move. Suppose that the general informativeness condition holds.*

Then, there is no weak PBE strategy profile in which $P1$ always moves to some ϕ_u in the first round when she sees $s_1 = 1$, and $P1$ always moves to $\phi_v \neq \phi_u$ in the first round when she sees $s_1 = 0$.

Proof. Let (σ_1, σ_2) be a weak PBE equilibrium strategy. For contradiction, suppose that σ_1 requires $P1$ to follow the myopic strategy in the first round. By lemma IV.2, $P1$ must move to ϕ_1 when $s_1 = 1$ and ϕ_0 when $s_1 = 0$. Now, in equilibrium, $P2$ will take this into account, and will therefore know both bits after the first round. She can capture all the remaining surplus by moving to $p_{00}, p_{11}, p_{01}, p_{10}$ depending on s_2 and the inferred value of s_1. Thus, in any equilibrium, she will eventually move to the optimal point. Now, consider a deviation from this strategy in which $P1$ bluffs in the first round and corrects P2's final move at the end. When $s_1 = 1$, Lemma IV.3 shows that the expected additional score increase if $P1$ bluffed is given by:

$$\kappa D(p_{11}||p_{01}) + (1-\kappa)D(p_{10}||p_{00}) - D(\phi_1||\phi_0)$$

Now, from a well-known convexity property of the relative entropy [26, pp.30], we have:

$$\kappa D(p_{11}||p_{01}) + (1-\kappa)D(p_{10}||p_{00}) \geq$$

(4.1) $$D(\kappa p_{11} + (1-\kappa)p_{10}||\kappa p_{01} + (1-\kappa)p_{00})$$

$$\Rightarrow \quad \kappa D(p_{11}||p_{01}) + (1-\kappa)D(p_{10}||p_{00}) \geq$$

(4.2) $$D(\phi_1||\phi_0)$$

$$\Rightarrow \quad \kappa D(p_{11}||p_{01}) + (1-\kappa)D(p_{10}||p_{00}) - D(\phi_1||\phi_0) \geq$$

(4.3) $$0$$

Inequality (4.2) follows from the definition of ϕ_1 and ϕ_0. Thus, inequality (4.3) implies that bluffing is always at least as profitable as behaving myopically by lemma IV.3. Moreover, inequality (4.3) is strict when $\kappa \neq 0, 1$ and $p_{10} \neq p_{00}, p_{11}$; this follows directly from the log sum inequality [26]. Thus, bluffing will be a strictly profitable deviation under the conditions of the theorem, and hence the myopic strategy for $P1$ cannot be part of an equilibrium profile. □

We observe that the general informativeness condition we assumed is sufficient but not necessary.

4.3.2 Nonexistence of Finite Equilibrium

Theorem IV.4 and Lemma IV.2 show that there is no equilibrium in which player 1 follows a deterministic strategy that is dependent on her signal. If there was such an equilibrium, then, in equilibrium, player 2 would infer player 1's bit and move to the optimal point.

Now, it follows that there is no weak PBE equilibrium strategy profile for the extended trading game, under the same assumptions, that satisfies the condition that the market is in the optimal state with certainty after some finite number, k, of rounds.

Theorem IV.5. *Under the general informativeness condition there is no weak PBE strategy profile in which the market is certain to be in the optimal state after k rounds, for any finite k.*

Proof. Suppose the general informativeness condition is met. By the log sum inequality, no matter the priors, *i.e.* the players beliefs of the other player's signals, it is always profitable for the player currently playing in the market to bluff.

From theorem IV.4, there are two cases that could arise in equilibrium.

Case (i): $P1$ plays some strategy p_u with certainty regardless of the value of s_1. In this case, $P2$ has learned nothing about $P1$'s bit, and thus, the conditions of the theorem always hold after 1 round.

Case (ii): $P1$ plays a mixed strategy for at least one value of s_1. Now, we claim that there is a position ϕ_u such that P1 moved to ϕ_u with nonzero probability ξ when $s_1 = 0$ *and* with nonzero probability ξ' when $s_1 = 1$. Suppose there was no such ϕ_u. Then, the support of $P1$'s first-round strategy for different values of s_1 would be completely disjoint, and thus, $P2$ could infer $P1$'s bit exactly. Thus, $P1$ would effectively have a deterministic strategy; a simple extension of Lemma IV.3 shows that the myopic strategy would be as good as this.

Observe that ϕ_u is played with strictly positive probability. Further, conditioning on $P1$ moving to ϕ_u in the first round, $P2$ would assign some probability $\hat{\kappa} \neq 0, 1$ to $s_1 = 1$. $P1$'s beliefs about s_2 haven't changed at all after the first round. Thus, the conditions of theorem IV.4 still hold after the first round, conditional on ϕ_u being played.

Repeating this argument for each of the first $k - 1$ rounds, and conditioning on one of the strategies in the support of each round, shows that the conditions of the theorem still hold with some nonzero (albeit small) probability. Thus, the price cannot converge with certainty after k rounds. $\square$

4.4 Generalizing the Results

We now move to a setting with n players and m signals each, for arbitrary n and m. We will use the following notation:

$\mathcal{M}$ The set of all players.

$\mathbf{i} \in \{0, n-1\}^{n-1}$ is a vector of the signals for all players *other than* player i.

$(j, \mathbf{i}) \in \{0, n-1\}^n$ is a vector of the signals for all players; j denotes player i's signal.

κ_j^k prior probability that player k sees signal j

$\kappa_{\mathbf{i}} = \prod_{k \in \mathcal{M}/\{1\}} \kappa_{\mathbf{i}_k}^k$ is the probability of players 2 through n of seeing the signals specified in $\mathbf{i}$.

$p_{(j,\mathbf{i})}$ is the optimal prediction with signal vector $(j, \mathbf{i})$

In the following scenario, assume that all players other than player 1 are behaving myopically, and player k moves kth in the market. Below we determine if player 1 has an incentive to deviate from the myopic strategy. As player 1 is playing first we define the following myopic optimal moves:

$$\phi_j = \sum_{\mathbf{i}} \kappa_{\mathbf{i}} p_{(j,\mathbf{i})} \qquad \text{if she sees signal } j$$

For the following assume that player 1 observes j as her signal, but is contemplating pretending to have $\bar{j}$ instead of following the myopic strategy.

Claim IV.6. *Player 1 has an expected score increase of* $\sum_{\mathbf{i}} \kappa_{\mathbf{i}} D(p_{(j,\mathbf{i})}||p_{(\bar{j},\mathbf{i})}) - D(\phi_j||\phi_{\bar{j}})$ *if she bluffs and then corrects the position after others have played myopically.*

Proof. The change in expected score due to the initial deviation move is equivalent to moving from ϕ_j to $\phi_{\bar{j}}$:

$$\begin{aligned} S(\phi_j, \phi_{\bar{j}}) &= \phi_j \ln \tfrac{\phi_{\bar{j}}}{\phi_j} + (1-\phi_j) \ln \tfrac{1-\phi_{\bar{j}}}{1-\phi_j} \\ &= -D(\phi_j||\phi_{\bar{j}}) \end{aligned}$$

As each player, k, has a probability of κ_j^k of seeing j as her signal and is behaving myopically it means that player 1 will move from $p_{(\bar{j},\mathbf{i})}$ to $p_{(j,\mathbf{i})}$ with probability $\kappa_{\mathbf{i}}$. Summing over all possible $\mathbf{i}$ we have the expected profit of the corrective step being $\sum_{\mathbf{i}} \kappa_{\mathbf{i}} S(p_{(\bar{j},\mathbf{i})}, p_{(j,\mathbf{i})})$
For any given $\mathbf{i}$:

$$\begin{aligned} S(p_{(\bar{j},\mathbf{i})}, p_{(j,\mathbf{i})}) &= p_{(j,\mathbf{i})} \ln \tfrac{p_{(j,\mathbf{i})}}{p_{(\bar{j},\mathbf{i})}} + (1-p_{(j,\mathbf{i})}) \ln \tfrac{1-p_{(j,\mathbf{i})}}{1-p_{(\bar{j},\mathbf{i})}} \\ &= D(p_{(j,\mathbf{i})}||p_{(\bar{j},\mathbf{i})}) \end{aligned}$$

Thus, the expected profit obtained in the corrective step is $\sum_{\mathbf{i}} \kappa_{\mathbf{i}} D(p_{(j,\mathbf{i})}||p_{(\bar{j},\mathbf{i})})$ □

Theorem IV.7. *The honest strategy profile is not a weak PBE equilibrium.*

Proof. Player 1 will deviate from the myopic strategy if her expected score of deviating is greater than her expected score of behaving myopically. In particular:

$$\begin{aligned} &\sum_{\mathbf{i}} \kappa_{\mathbf{i}} D(p_{(j,\mathbf{i})}||p_{(\bar{j},\mathbf{i})}) - D(\phi_j||\phi_{\bar{j}}) > 0 \\ \iff &\sum_{\mathbf{i}} \kappa_{\mathbf{i}} D(p_{(j,\mathbf{i})}||p_{(\bar{j},\mathbf{i})}) > D(\phi_j||\phi_{\bar{j}}) \\ \iff &\sum_{\mathbf{i}} \kappa_{\mathbf{i}} D(p_{(j,\mathbf{i})}||p_{(\bar{j},\mathbf{i})}) > D(\sum_{\mathbf{i}} \kappa_{\mathbf{i}} p_{(j,\mathbf{i})}|| \sum_{\mathbf{i}} \kappa_{\mathbf{i}} p_{(\bar{j},\mathbf{i})}) \end{aligned} \tag{4.4}$$

We use the following standard result on the convexity of the relative entropy, for any set of convex multipliers $\{\lambda_i\}$:

$$D(\sum_i \lambda_i p_i || \sum_i \lambda_i q_i) \leq \sum_i \lambda_i D(p_i || q_i) \tag{4.5}$$

Further, we note that equality can hold only if either all p_i with $\lambda_i > 0$ are identical, or $p_i = q_i$ for at least one pair. (This is implicit in [26, pp. 29-30].)

By the general informativeness condition, $p_{(j,\mathbf{i})} \neq p_{(\bar{j},\mathbf{i})}$. This condition also implies that, if $\bar{\mathbf{i}}$ is constructed by changing any one component of $\mathbf{i}$, we have $p_{(j,\bar{\mathbf{i}})} \neq p_{(j,\mathbf{i})}$. As $\bar{\mathbf{i}}$ can occur with positive probability, the inequality (4.4) holds strictly.

□

4.4.1 Nonexistence of Finite Equilibrium

We can extend this result to show nonexistence of finite informative equilibria. For contradiction, consider any weak PBE strategy profile σ that always results in the optimal market prediction after some number t of rounds in a potentially infinite game. By theorem IV.7, player 1 cannot play myopically in her first move under σ; otherwise, she would have a profitable deviation. Thus, after some move that player 1 makes with positive probability under σ, there must be at least two feasible values of her signal that could have led to that move. We are now left with a reduced game with a different order of play, perhaps a smaller set of signals that player 1 could have, and perhaps different values of some κ^1_j. σ must be consistent with a weak PBE of this reduced game. However, the independence and general informativeness conditions still hold. Thus, we can apply theorem IV.7 again. Inductively, we cannot have convergence by any finite t.

Intuitively if a player has some private information that may be revealed by her signal, she has an incentive to not fully divulge this information when she may play multiple times in the market. The general informativeness condition simply guarantees that no matter the signal distribution, a player will always have some private information revealed by her signal.

4.4.2 Implications for Other Scoring Rules

We believe that these results are not artifacts of the logarithmic scoring rule in particular. The result above relies on the strict convexity of the Kullback-Leibler divergence, the divergence associated with the log market scoring rule. The results above holds for all scoring rules with strictly convex divergence functions. We believe that many, if not all, other scoring rules have an associated convex divergence function. In particular, there is a strictly convex divergence function associated with the quadratic market scoring rule as well. This means that a similar analysis holds for the quadratic scoring rule.

4.5 Discounting and Entropy Reduction

In this section we propose a discounted market scoring rule, characterize equilibria in this market in terms of entropy, and use this to show that the market converges to the optimal price in any equilibrium.

4.5.1 The Discounted Market Scoring Rule

One way to address the incentives traders have to bluff in a market using the log market scoring rule is to reduce future payoffs using a discount parameter, perhaps resulting in an incentive to play the myopic strategy. Using this intuition, we propose the discounted log-MSR market.

Let $\delta \in (0, 1)$ be a discount parameter. The **δ-discounted market scoring rule** is a market microstructure in which traders update the predicted probability of the event under consideration happening, just as they would in the regular market scoring rule. However, the value (positive or negative) of trade is discounted over time. For simplicity, we assume a strict alternating sequence of trades. Suppose a trader moves the prediction from p to q in his ith move, and the event is later observed to happen. The trader would then be given a payoff of $\delta^{i-1}(\log q - \log p)$. On the other hand, if the event did not happen, and the player moves the prediction from p to q in his ith move he would earn a payoff of $\delta^{i-1}(\log(1-q) - \log(1-p))$.

Using the definition of a market scoring rule price function as defined in [43], we note that the price function in the δ-discounted market scoring rule is (4.6).

$$\frac{\exp((h_i - a_i)/(b \cdot \delta^t))}{\sum_j \exp((h_j - a_j)/(b \cdot \delta^t))} \tag{4.6}$$

h_j is the amount of shares of security j sold so far. b and a_j are constants defined by the non-discounted market scoring rule.

Clearly, the myopic strategic properties of the market scoring rule are retained in the discounted form. We will show that the discounted form can have better non-myopic strategic properties.

We present a multiplicative form of discounting. However, there may be other forms of discounting that could be used. One alternative would be to charge each traders a fixed coupon price to trade in the mechanism. In particular, before any trader trades they pay a fixed cost to the market maker. This approach may actually disincentivize some traders from participating in this market if their expected profit from participating without the fixed cost is less than the participation cost. Another approach to discounting is to fix an ending time of the security. For example, the market will be closed after 100 trades and the traders will receive their payoff once the outcome of the event is observed. This approach has the advantage of each trader revealing their true beliefs with their last trade in the market. However, closing the event early eliminates the opportunity for traders to gather information past 100 trades and have that information aggregated by the market. As the δ-discounted market scoring rule does not suffer from these drawbacks, we believe it is the way to discount.

4.5.2 Convergence and Entropy Reduction

The discounted log-MSR may admit equilibria in which players play non-myopic strategies, *i.e.*, they bluff with some probability. We want to show that, in any weak-PBE profile σ, the market price will converge towards the optimal value for the particular realized set of information signals. In other words, although complete aggregation of information may not happen in two rounds, it does surely happen in the long run.

We now present a natural metric $\mathcal{D}^i$ that quantifies the degree of aggregation in the prediction market after any number of trades of strategy profile σ: $\mathcal{D}^i$ is the expectation, over all possible signal realizations *and* the randomization of moves as dictated by σ, of the relative entropy between the optimal price (given the realization of the signals) and the actual price after i rounds.

Formally, for a strategy profile σ, and a number of rounds i, a *signal node* π consists of a realization of the signals of the two players, and a sequence of i trades in the market. The aggregative effect of the strategy profile σ is summarized by the collection of signal nodes π that can be reached, the market price $r_i(\pi)$ after the last trade in π, and the associated ex-ante probability $P\{\pi\}$ of reaching each such signal node. Now, we define

$$\begin{aligned}\mathcal{D}^i(\sigma) &= E_\sigma[D(r^*||r_i)] \\ &= \sum_{\pi:\pi\in\sigma} P\{\pi\} D(r^*(\pi)||r_i(\pi)) \ ,\end{aligned}$$

where $r^*(\pi)$ denotes the optimal trading price given the realization of signals in

π. Hereafter, we abuse notation by merely writing $\mathcal{D}^i$ for $\mathcal{D}^i(\sigma)$; the profile under consideration is obvious from the context. For $i = 0$, the signal nodes π correspond to different realizations of the signals.

If $\mathcal{D}^i = 0$, it implies that the market will always have reached its optimal price for the realized signals by the ith rounds. If $\mathcal{D}^i > 0$, it indicates that, with positive probability, the market has not yet reached the optimal price. $\mathcal{D}^i$ is always nonnegative, because the relative entropy is always nonnegative.

We now show that, in addition to measuring the distance from full aggregation, $\mathcal{D}^i$ also enables interesting strategic analysis. The key result is that $\mathcal{D}^i$ can be related to the expected payoff of the ith round move in the *non-discounted* (standard) log-MSR:

Lemma IV.8. *Let M^i denote the expected profit (over all signal nodes π) of the ith trade under profile σ. Then, $M^i = \mathcal{D}^{i-1} - \mathcal{D}^i$*

Proof.

$$\begin{aligned} M^i &= \sum_{\pi:\pi\in\sigma} P\{\pi\}[r^*(\pi)[\log r_i - \log r_{i-1}] \\ &\qquad +(1-r^*(\pi))[\log(1-r_i) - \log(1-r_{i-1})] \\ &= \sum_{\pi:\pi\in\sigma} P\{\pi\}[D(r^*||r_{i-1}) - D(r^*||r_i)] \\ &= \mathcal{D}^{i-1} - \mathcal{D}^i \end{aligned}$$

This first equality holds by the definition of M^i and the second by the definition of relative entropy. □

This suggests another interpretation for $\mathcal{D}^i$: $\mathcal{D}^i$ represents the expected value of the potential profit left for trades *after* the ith trade.

Given the definition of M^i, we can now define $\tilde{M}^i$ as the expected profit of the ith trade in the discounted log MSR. We have assumed discounting after every even trade, *i.e.*, after every trade by player 2.

$$\tilde{M}^i = \begin{cases} \delta^k(\mathcal{D}^{2k} - \mathcal{D}^{2k+1}) & \forall i = 2k \\ \delta^k(\mathcal{D}^{2k+1} - \mathcal{D}^{2(k+1)}) & \forall i = 2k+1 \end{cases} \tag{4.7}$$

Using the definition of $\tilde{M}^i$ we write the total profit for player 1 in the δ-discounted MSR as:

$$\begin{aligned} \text{P1 payoff} &= \sum_{i:i=2k,k\in\mathbb{Z}^+} \tilde{M}^i \\ &= \mathcal{D}^0 - \mathcal{D}^1 + \delta(\mathcal{D}^2 - \mathcal{D}^3) + \ldots \end{aligned}$$

We can similarly rewrite player 2's expected payoff as:

$$\begin{aligned} \text{P2 payoff} &= \sum_{i:i=2k,k\in\mathbb{Z}^+} \tilde{M}^{i+1} \\ &= \mathcal{D}^1 - \mathcal{D}^2 + \delta(\mathcal{D}^3 - \mathcal{D}^4) + \ldots \end{aligned}$$

We reiterate that the definition of $\mathcal{D}^i$ is *not* dependent on δ; it is a measure of the informational distance between the prices after i trades in profile σ and the optimal prices. Of course, the stability of a given strategy profile σ may change with δ.

4.5.3 Bounding Relative Entropy

In this section, we bound the value of $\mathcal{D}^n$ for large n, in any weak PBE. Recall that an instance of the two-person market game consists of a set of optimal points for different signal realizations $\{p_{00}, p_{01}, p_{10}, p_{11}\}$ and prior beliefs κ_1 and κ_2 about the probability that player 1 (and player 2, respectively) will receive the 1 signal. The optimal points remain unchanged through the entire course of trade, but the players' beliefs about each other's signal distribution changes as trade proceeds. Thus, it is useful to separate the instance description into a *configuration* $\{p_{00}, p_{01}, p_{10}, p_{11}\}$ and a *signal distribution* (κ_1, κ_2).

We will express our convergence bound in terms of an invariant of the market configuration, the *complementarity coefficient*, which we define below. Fix a particular configuration of the optimal points. Now, any pair of probabilities (κ_1, κ_2) determines an instance of the game. Let $\pi_M^1(\kappa_1, \kappa_2)$ denote the expected profit of player 1 if she traded first, and followed the myopically optimal (*i.e.*, honest) trading strategy in the first round of trade. Similarly, let $\pi_M^2(\kappa_1, \kappa_2)$ denote the expected profit of player 2 *if player 2 had traded first*, and followed the myopically optimal (*i.e.*, honest) trading strategy in the first round of trade. Let $\mathcal{D}^0(\kappa_1, \kappa_2)$ denote the initial profit potential of this instance: the expected *total* profit in moving from p_s to the optimal point. In other words, $\mathcal{D}^0(\kappa_1, \kappa_2)$ represents the sum of both players' expected profits if they both followed the myopic strategy. Now, define the complementarity coefficient $C(\kappa_1, \kappa_2) = \frac{\pi_M^1(\kappa_1,\kappa_2)+\pi_M^2(\kappa_1,\kappa_2)}{\mathcal{D}^0(\kappa_1,\kappa_2)}$. The reason for using the term 'complementarity' comes from the following observation. Under the myopic strategy profile, whichever order the players trade, their total profit will be $\mathcal{D}^0(\kappa_1, \kappa_2)$. If this is greater than the sum of what they could each have earned playing first, *i.e.*, $\pi_M^1(\kappa_1, \kappa_2) + \pi_M^2(\kappa_1, \kappa_2)$, then this indicates that their individual bits of information have increasing marginal value, *i.e.* they are complements. The *lower* $C(\kappa_1, \kappa_2)$ is, the *greater* the complementarity of information. Finally, let the *complementarity bound* C of the configuration be the minimum, over all values of (κ_1, κ_2), of $C(\kappa_1, \kappa_2)$.

Under our independence assumption $C(\kappa_1, \kappa_2)$ is always less than 1; indeed, if it were greater than 1, the myopic strategy profile would be an equilibrium. We do not yet have a good characterization of the complementarity coefficient. However, based on sample configurations, we have observed that it is nontrivial (not always 0), and often quite close to 1. Note that if $C = 0$, the myopic strategy may not involve any movement by either player, and thus, we could have lack of information aggregation even with the myopic strategy. We exclude such degenerate cases, and assume that $C > 0$.

Now, fix a discount parameter δ, and a particular instance (including probabilities (κ_1, κ_2)) of the two-person market game induced by the δ-discounted log-MSR. Let σ be any weak PBE of this market game.

Without loss of generality we assume that player 2 moves second in the market. In the first round, player 1 will follow some (perhaps mixed) strategy σ_1 dictated by σ. Under the equilibrium strategy profile σ, player 2 will revise her beliefs consistent with the profile σ and the realized move. Let $\pi^2_M|\sigma_1$ denote the expected profit of player 2 if she played her myopically optimal strategy conditioned on her revised beliefs. From theorem IV.9 we note that $\pi^2_M|\sigma_1 \geq \pi^2_M$.

Theorem IV.9. $\pi^2_M|\sigma_1 \geq \pi^2_M$

Proof. Consider any strategy σ for player 1 such that setting $\sigma_1 = \sigma$ minimizes $\pi^2_M|\sigma_1$. We will show that σ must involve player 1 not moving the market price at all.

We will first argue that σ has support on a single point only. If not, then σ would have support over a set of points: at least two points A, B, and perhaps a set of other points R. In this case, we show that we can construct a strategy σ' that reduces the objective function by "mixing" points A and B. For simplicity, we assume that $\kappa_1 = 0.5$, *i.e.*, that player 1 has equal chance of seeing either a 1 or a 0. Any other value of κ_1 can easily be substituted into the proof below, but it would slightly clutter the notation. Define u_A and u_B as the probability (under σ) that player 1 will play A and B given she saw 1 as her signal. Similarly we define v_A and v_B as the probability player 1 will play A and B given she saw 0 as her signal. Let p_A be the probability player 1 plays A and similarly p_b be the probability player 1 plays B. Without loss of generality let $p_B < p_A$. Define $\alpha = P\{s_1 = 1|A\} = \frac{u_A}{p_A}$ and

$\beta = P\{s_1 = 1|B\} = \frac{u_B}{p_B}$. With these definitions we can define the myopic move of player 2 given that market is at A and $s_2 = 1$ as $\phi_1^\alpha = \alpha p_{11} + (1-\alpha)p_{01}$; likewise, if $s_2 = 0$ as $\phi_0^\alpha = \alpha p_{10} + (1-\alpha)p_{00}$. Similarly ϕ_0^β and ϕ_1^β are defined as the myopic moves of player 2 given the market is at B.

Now, let C = (A+B)/2 be the midpoint of A and B, and consider a new strategy σ' over points A, C, and the same set of remaining points R. Under σ' she will mix over A and C with probability $p_A - p_B$ and $2p_B$ respectively. As before we can define $\gamma = P\{s_1 = 1|C\} = \frac{\frac{p_B}{p_A}u_A + u_B}{2p_B} = \frac{1}{2}\frac{u_A}{p_A} + \frac{1}{2}\frac{u_B}{p_b} = \frac{\alpha+\beta}{2}$. We can now define the myopic move of player 2 given the market is at C and $s_2 = 1$ as $\phi_1^\gamma = \gamma p_{11} + (1-\gamma)p_{01}$. Likewise if $s_2 = 0$ and the market is at C the myopic move of player 2 is defined as $\phi_0^\gamma = \gamma p_{10} + (1-\gamma)p_{00}$.

We now characterize $\pi_M^2|\sigma$ as:

$$
\begin{aligned}
\pi_M^2|\sigma \;\; = \;\; & 0.5(p_A - p_B)[qD(\phi_1^\alpha||q\phi_1^\alpha + (1-q)\phi_0^\alpha) \\
& +(1-q)D(\phi_0^\alpha||q\phi_1^\alpha + (1-q)\phi_0^\alpha)] \\
& +0.5p_B[qD(\phi_1^\alpha||q\phi_1^\alpha + (1-q)\phi_0^\alpha) \\
& +(1-q)D(\phi_0^\alpha||q\phi_1^\alpha + (1-q)\phi_0^\alpha)] \\
& +0.5p_B[qD(\phi_1^\beta||q\phi_1^\beta + (1-q)\phi_0^\beta) \\
& +(1-q)D(\phi_0^\beta||q\phi_1^\beta + (1-q)\phi_0^\beta)] \\
& + \text{remaining profit over } R
\end{aligned}
$$

We also characterize $\pi_M^2|\sigma'$ as follows, writing p_A as $(p_A - p_B) + p_B$ to facilitate comparison with σ':

$$\begin{aligned}\pi_M^2|\sigma' &= 0.5(p_A - p_B)[qD(\phi_1^\alpha||q\phi_1^\alpha + (1-q)\phi_0^\alpha)\\ &+(1-q)D(\phi_0^\alpha||q\phi_1^\alpha + (1-q)\phi_0^\alpha)]\\ &+0.5 \cdot 2p_B[qD(\phi_1^\gamma||q\phi_1^\gamma + (1-q)\phi_0^\gamma)\\ &+(1-q)D(\phi_0^\gamma||q\phi_1^\gamma + (1-q)\phi_0^\gamma)]\\ &+ \text{ remaining profit over } R\end{aligned}$$

From the definitions of the myopic moves given the market states, note that $\phi_1^\gamma =$

$\frac{\phi_1^\alpha+\phi_1^\beta}{2}$ and $\phi_0^\gamma = \frac{\phi_0^\alpha+\phi_0^\beta}{2}$. This means that $\pi_M^2|\sigma'$ can be bounded as :

$$\begin{aligned}
\pi_M^2|\sigma' &= 0.5(p_A - p_B)[qD(\phi_1^\alpha||q\phi_1^\alpha + (1-q)\phi_0^\alpha) \\
&\quad +(1-q)D(\phi_0^\alpha||q\phi_1^\alpha + (1-q)\phi_0^\alpha)] \\
&\quad +0.5 \cdot 2p_B[qD(\phi_1^\gamma||q\phi_1^\gamma + (1-q)\phi_0^\gamma) \\
&\quad +(1-q)D(\phi_0^\gamma||q\phi_1^\gamma + (1-q)\phi_0^\gamma)] \\
&\quad + \text{remaining profit over } R \\
&= 0.5(p_A - p_B)[qD(\phi_1^\alpha||q\phi_1^\alpha + (1-q)\phi_0^\alpha) \\
&\quad +(1-q)D(\phi_0^\alpha||q\phi_1^\alpha + (1-q)\phi_0^\alpha)] \\
&\quad +0.5 \cdot 2p_B[qD(\tfrac{\phi_1^\alpha+\phi_1^\beta}{2}||q\tfrac{\phi_1^\alpha+\phi_1^\beta}{2} + (1-q)\tfrac{\phi_0^\alpha+\phi_0^\beta}{2}) \\
&\quad +(1-q)D(\tfrac{\phi_0^\alpha+\phi_0^\beta}{2}||q\tfrac{\phi_1^\alpha+\phi_1^\beta}{2} + (1-q)\tfrac{\phi_0^\alpha+\phi_0^\beta}{2})] \\
&\quad + \text{remaining profit over } R \\
&< 0.5(p_A - p_B)[qD(\phi_1^\alpha||q\phi_1^\alpha + (1-q)\phi_0^\alpha) \\
&\quad +(1-q)D(\phi_0^\alpha||q\phi_1^\alpha + (1-q)\phi_0^\alpha)] \\
&\quad +0.5p_B[qD(\phi_1^\alpha||q\phi_1^\alpha + (1-q)\phi_0^\alpha) \\
&\quad +(1-q)D(\phi_0^\alpha||q\phi_1^\alpha + (1-q)\phi_0^\alpha)] \\
&\quad +0.5p_B[qD(\phi_1^\beta||q\phi_1^\beta + (1-q)\phi_0^\beta) \\
&\quad +(1-q)D(\phi_0^\beta||q\phi_1^\beta + (1-q)\phi_0^\beta)] \\
&\quad + \text{remaining profit over } R \\
&= \pi_M^2|\sigma
\end{aligned}$$

The last inequality follows from the strict convexity of relative entropy under the general informativeness condition.

Therefore, for any strategy σ with two or more points in its support, there always exists a strategy σ' such that $\pi_M^2|\sigma' < \pi_M^2|\sigma$. This means that for any strategy, σ, for player 1 that minimized $\pi_M^2|\sigma$ the strategy must have only one point in its support. Thus, the strategy does not reveal any information about player 1's bit to player 2.

Suppose that the point in the support, p_i, is such that $p_i \neq p_s$. This again contradicts the fact that σ minimizes $\pi^2_M|\sigma$, as player 2 will always make a positive payoff in expectation if she moves from p_i to p_s; thus, she would have a larger payoff overall if player 1 left the market at p_i instead of p_s. Therefore the strategy that minimizes the expected payoff of player 2 is for player 1 to report p_s. However, player 1 reporting p_s is equivalent to her not trading at all in the first round. Therefore we have shown that $\pi^2_M|\sigma \geq \pi^2_M$. □

(where, for clarity, we have suppressed the dependence on (κ_1, κ_2)). Intuitively the result of the theorem holds, as any move by player 1 reveals some information to player 2 in equilibrium. Any such information would be used by player 2 to reduce her uncertainty on the observed bit of player 1. Due to the complementarity of signals, this results in a higher expected profit for player 2 than if she had no information at all. The latter is equivalent to the situation in which player 2 moves first.

Recall that after the second round, the total expected payoff of both players is at most $\delta\mathcal{D}^2$. We also know that the total expected payoff of player 1 in equilibrium is at least π^1_M; if not, a simple deviation to the myopic strategy would be beneficial. By theorem IV.9, the total expected payoff of player 2 in equilibrium is also at least π^2_M. This means that the total payoff of the first two rounds in the market is at least $\pi^1_M + \pi^2_M - \delta\mathcal{D}^2$. Therefore we can bound $\mathcal{D}^2$ as:

$$\mathcal{D}^2 \leq \mathcal{D}^0 - [\pi^1_M + \pi^2_M - \delta\mathcal{D}^2]$$

This argument generalizes to any even number of rounds, by looking at the total expected profits within the first $2k$ moves, and the remaining profit potential $\delta^k\mathcal{D}^{2k}$:

$$
\begin{aligned}
(4.8) \qquad\qquad \mathcal{D}^{2k} &\leq \mathcal{D}^0 - [\pi_M^1 + \pi_M^2 - \delta^k \mathcal{D}^{2k}] \\
\iff (1-\delta^k)\mathcal{D}^{2k} &\leq \mathcal{D}^0 - [\pi_M^1 + \pi_M^2] \\
&= \mathcal{D}^0 \left(1 - \frac{\pi_M^1 + \pi_M^2}{\mathcal{D}^0}\right) \\
(4.9) \qquad\qquad \iff \mathcal{D}^{2k} &\leq \frac{\mathcal{D}^0\left(1 - C(\kappa_1, \kappa_2)\right)}{1-\delta^k}
\end{aligned}
$$

Note that for any configuration, by definition, $C(\kappa_1, \kappa_2) \geq C$. Thus, we can rewrite inequality (4.9) as:

$$
\mathcal{D}^{2k} \leq \frac{\mathcal{D}^0(1-C)}{1-\delta^k} \tag{4.10}
$$

From inequality (4.10) we see that the bound on $\mathcal{D}^{2k}$ depends only on δ, as $\mathcal{D}^0$ and C are both constants for any configuration of p_{ij}. Now, consider the remainder of the game after k rounds. After any particular sequence of moves that occurs with positive probability, the players would have updated their beliefs about the other player's bit. Thus, the players are left to play a slightly different instance of the 2-player market game, and for smaller stakes. But the configuration of the optimal points stays the same. Therefore, the equilibrium profile σ will also induce an equilibrium profile on the instance of the game after $2k$ rounds. Now, we can repeat this argument to bound $\mathcal{D}^{4k}$ in terms of $\mathcal{D}^{2k}$, etc.

In this way, we rewrite inequality (4.10) in terms of δ and for a round $n = 2km$. We set k such that $\delta^{k/2} = C$, *i.e.* $k = \frac{2\log C}{\log \delta}$. Using this value of k and a value

$m = \frac{n}{2k}$ we note that:

$$
\begin{aligned}
\mathcal{D}^n &\leq \mathcal{D}^0 \left(\tfrac{1-C}{1-\delta^k}\right)^m \\
&= \mathcal{D}^0 \left(\tfrac{1-C}{(1-\delta^{k/2})(1+\delta^{k/2})}\right)^m \\
&= \mathcal{D}^0 (1+\delta^{k/2})^{-m} = \mathcal{D}^0 (1+C)^{-\frac{n}{2k}} \\
&= \mathcal{D}^0 (1+C)^{-\frac{n \log \delta}{4 \log C}} = \mathcal{D}^0 \delta^{n^{\frac{-\log(1+C)}{4\log C}}}
\end{aligned} \tag{4.11}
$$

Note that $\frac{-\log(1+C)}{4\log C}$ depends only on complementarity coefficient of the market configuration. Moreover, the value of $\frac{-\log(1+C)}{4\log C} > 0$ as $C < 1$ from the independence condition. Therefore, inequality (4.11) shows that the relative entropy of the prices with respect to the optimal prices reduces exponentially over time.

Further, the mechanism designer can reduce the value of δ to speed up the convergence to optimum in any weak PBE. One caveat: rapid discounting results in rapidly reducing available profits, and thus, may dissuade traders from participating in the market.

4.6 Discussion and Future Work

In this chapter, we analyze non-myopic strategies in a two-player prediction market setting. We find that the myopic strategies are generically not in equilibrium when non-myopic strategies are admitted, under our independent signals assumption. In a real market, there may be other reasons why players prefer the myopically optimal strategies: In particular, they are much simpler to play, and more robust, and the potential gains from bluffing are often very small. Thus, our results are not in any way meant to imply that market scoring rules are not a useful microstructure for organizing a market. Instead, we believe that the analysis suggested here will be useful in clarifying when markets might be especially susceptible to long-range manipulative strategies. The contrast between our results and the results of Chen *et al.* [22] are especially intriguing. One exciting direction for future research is to fully characterize the class of information structures on which myopic strategies are in equilibrium, and more importantly, are the *only* equilibrium.

We use a simple modification to the market scoring rule, which includes a form of discounting, to ameliorate this potential problem. This allowed us to prove a bound on the rate at which the error of the market, as measured by the relative entropy between perfect aggregation and the actual price distribution, reduces exponentially over time. The exponent depends on the "complementarity coefficient" of the instance. One important direction for future work is to characterize or bound this function; this will lead to a more complete understanding of the convergence rate.

The need for discounting shows a connection to bargaining settings, in which players bargain over how to divide a surplus they can jointly create. In a prediction market, informed players can extract a subsidy from the market maker; moreover, players can pool their information together to make sharper predictions than either could alone, and thus extract an even larger subsidy. They might engage in bluffing strategies to bargain over how this subsidy is divided. Explicit discounting can make this bargaining more efficient.

CHAPTER V

Subsidized Prediction Markets for Risk Averse Traders

This chapter is based on joint work with Marina Epelman and Rahul Sami; parts were reported in [30].

5.1 Introduction and Related Work

As discussed in Chapter III, Hanson [42] has shown that, for risk neutral agents who are myopic (*i.e.*, do not account for the effect of their trades on other traders), it is optimal for each trader to reveal her true beliefs of the outcome of the traded event in a market scoring rule market. This results leaves two questions:

1. What happens when agents take into account future payoffs?

2. What happens when agents are not risk neutral?

The first question was partially addressed in Chapter IV and results are further expanded on in Chen *et al.* [20] by showing that if agents have complementary information, a mixed strategy of bluffing with a certain probability is an equilibrium strategy, and if agents have substitutable information, then the truthful strategy is an equilibrium strategy.

In this chapter, we tackle the second question: *What happens when agents are not risk neutral?* In practice, most people are better modeled as being risk averse in their decision making. We model traders as expected-utility maximizers with an arbitrary weakly monotone and concave utility functions that captures their risk aversion. Current prediction market mechanisms, like the market scoring rule or the dynamic pari-mutuel market [63], do not always give appropriate incentives to risk averse traders. For example, a sufficiently risk averse informed trader, who knows that an event will occur with 80% probability even though the current price is 50%, may not want to push up the price in a market scoring rule market because of the 20% chance of making a loss. This observation suggests that subsidized prediction markets using the current mechanisms may converge to a non-truthful[1] price in a sequential equilibrium. As prediction markets are used in the corporate setting for forecasting future events [24], a decision maker running such a market might make the wrong decision as a result of a non-truthful equilibrium due to the presence of risk averse traders.

If traders have known risk aversion, the scoring rules could be adjusted to compensate, retaining the original incentive properties. In this chapter, we focus on the more common setting, in which traders have *unknown* risk aversion, and study whether it is possible to modify the market mechanism to guarantee myopic honesty while preserving other desirable properties of prediction markets. We first list a set of properties that any prediction market-like mechanism should satisfy: (1) myopic strategyproofness; (2) sequential trade, giving traders the opportunity to update beliefs; (3) a variant of *sybilproofness*, capturing the idea that trading under multiple identities does not yield any direct advantage; and (4) boundedness of the expected subsidy.

[1]As discussed in Chapter III, risk averse agents may participate, but will report their risk neutral probabilities. Since risk neutral probabilities are reported, the market will converge to an equilibrium in the risk neutral probabilities. In this situation we say that the market is converging to a non-truthful equilibrium.

We propose one mechanism that satisfies all of these properties, even in the presence of traders with unknown risk averse preferences. The key building block of our result is a technique, developed by Allen [5], of scoring forecasters by varying their probability of winning a fixed reward. We note that the proposed mechanism has the undesirable property of reducing the expected reward exponentially with the number of agents.

We then establish that such exponentially decreasing rewards are unavoidable for any mechanism satisfying all the properties listed above. It is natural for traders to have lower expected rewards in some contexts, such as when they have no private information; to exclude trivial examples of decreasing rewards, we normalize all rewards by a measure of the intrinsic informativeness of a trader's private information. We show that exponential decrease in the *normalized* expected reward is necessary for any mechanism that satisfies the properties we propose in the presence of arbitrarily risk averse agents. The consequence of this result is that, in any such mechanism, even a trader with a significant amount of private information might find that she can earn a very minute amount if the number of traders is large.

Related Work

In this chapter we only consider subsidized prediction markets. Therefore, in this section we look specifically at risk aversion in scoring rules and subsidized prediction markets. Prior work on risk aversion in unsubsidized prediction markets is summarized in Chapter III. Hanson [42] introduced the concept of a market scoring rule, a form of subsidized prediction market, and proved a myopic strategyproofness property for risk neutral traders, as well as a bound on the total subsidy. Pennock [63] introduced another mechanism, the dynamic pari-mutuel market, for a subsidized prediction market. Both these mechanisms introduce some of the properties in Section 5.2. However, both mechanisms assume risk neutrality of the traders, which is not assumed in this chapter. Lambert *et al.* [50] introduced a self-financed wagering mechanism, *i.e.*, one that is not subsidized, and introduced properties that such a mechanism must satisfy and showed a class of mechanisms that satisfy these properties. The authors assumed risk neutral traders and an absence of subsidy in the mechanism. Chen and Pennock [21], and later generalized by Agrawal *et al.* [3], considered a risk averse market maker in a subsidized market and showed that the market maker has bounded subsidy in most forms of risk aversion. This line of work is useful to show that even a risk averse market maker would be interested in subsidizing a market; however, unlike this chapter, the incentive consequence of risk averse *traders* were not addressed.

As discussed in Chapter III, the market scoring rule prediction market mechanism is based on scoring rules, which originally were introduced to measure the accuracy of weather forecasters. In order to address risk aversion in subsidized prediction markets, we look at previous work on risk aversion in scoring rules. Winkler and Murphy [78] showed that, if forecasters have a *known* risk type, scoring rules can be transformed to recapture the honest reporting property. Chen *et al.* [19] and Offerman *et al.* [59] provided one approach to handling forecasters with unknown risk type: they proposed first figuring out every participant's risk type by asking them a series of questions, and then calibrating their future reports using this data. This mechanism may work for a prediction market mechanism if the group of traders can be pre-screened. However, this may not be the case, and ideally we would like to have an "online" mechanism that can handle traders regardless of their risk type without any calibration. Allen [5] proposed one such "online" scoring rule for traders with arbitrary risk type. Our mechanism is based on Allen's result, and we discuss this idea in Section 5.3.

5.2 Model, Notation, and Definitions

In this section and Section 5.4 we consider a class of mechanisms defined by a set of properties each mechanism in the class satisfies. These properties are satisfied by most subsidized prediction markets described in literature. We do not claim that the outlined properties are sufficient to completely characterize the space of prediction market mechanisms; rather, they identify a class of broad market-like mechanisms. We first describe our basic model of the information and interaction setting in which the mechanism operates, and then list the properties that the mechanisms we study should satisfy.

We consider a class of mechanisms designed to aggregate information from a set of *agents* (or *traders*) in order to forecast the outcome of a future event ω. Each agent i receives a private information signal, s_i, relevant to the outcome of the event; we assume s_i is binary, as is ω.

As before, in the market-like mechanisms that we consider, the agents sequentially interact with the mechanism through trades or *reports*, each of which expresses a predicted probability of the event. Reports are public, and other agents can update their beliefs based on the observed history of reports. We use $r_k \in [0, 1]$ to denote the kth report made in the market, and let $\mu_k = (r_1, \cdots, r_{k-1})$ denote the history up to the start of the kth trade. r_k can thus depend on μ_k as well as any private information available to the trader making the report. We let n denote the total number of trades in the market.

The agents provide names while making reports, but in this chapter we consider a setting in which the identity of the agents making the reports cannot be verified, and that the total number of participating agents is not known. We assume that no agent is identified more than once in a sequence of trades. An agent may, however, masquerade as multiple agents, which will be important in our consideration of sybilproofness. This requirement to treat each trade as if it were from a separate agent is natural for a market setting.

Once the true outcome of the event is realized, $\omega = 1$ if the event occurs and $\omega = 0$ if the event does not occur, the mechanism determines the reward for every agent. The reward for agent i, $\rho(r_i, \mu_i, n, \omega)$, is a function of the report of the agent, the market state at the time the report has been made, the total number of agents participating in the mechanism, and the outcome of the event. We allow the mechanism to randomize the distribution of the rewards, and we propose one such mechanism in Section 5.3. We assume that the reward does not depend on any reports made in the future. This is a nontrivial technical assumption that enables us to simplify the analysis of agents' myopic strategies, as agents can make decisions based on their current beliefs about the outcome, without forming beliefs about future agents' signals and strategies. This assumption is satisfied by most securities markets as well as market scoring rule markets, but is not necessarily true for pari-mutuel markets.

Every agent i values the reward distributed to her by the mechanism according to her value function $V_i(\cdot)$, where $V_i(\cdot)$ is a weakly monotone increasing concave function. We make the normalizing assumption that $V_i(0) = 0$. In order to make her report, the agent maximizes her expected value $E_{\omega \sim p_i} V_i(\rho(r, \mu_i, n, \omega))$ over possible reports $r \in [0, 1]$. The expectation above is taken over the outcome of the event (with respect to the agent's true belief p_i, as emphasized by the notation), as well as any randomization of the mechanism over the rewards. Though there may be other sources of uncertainty in the mechanism, we do not consider them in our model.

We identify the properties mechanisms should satisfy by examining the literature in prediction market design and other wagering mechanisms. In recent literature there is work on characterizing properties of subsidized prediction markets. Hanson [42], in introducing the market scoring rules required a subsidized prediction market be *myopically strategy proof*, meaning that an agent not considering future reward will maximize her payoff by reporting her true belief on the outcome of the event. Similarly, he required that the subsidized prediction market have *bounded market subsidy*. The same properties also hold in the dynamic pari-mutuel market introduced by Pennock [63]. Due to the fact that both Hanson's and Pennock's mechanisms were subsidized, both had *guaranteed liquidity* by having a market maker that is always willing to trade with an agent using a predetermined price function. Prediction markets provide *anonymity*, *i.e.*, the reward given due to a report is independent of who made the report. Finally, prediction markets are *sybilproof*, meaning that an agent reporting once with some information is no better off reporting twice in the market (using different names, commonly referred to as *sybils*) with the exact same information. We note that though anonymity and sybilproofness were not explicitly stated by Hanson or Pennock, they still hold in their proposed mechanisms and were explicitly defined by Lambert *et al.* [50]. The definition of sybilproofness we give is a relaxation of that presented by Lambert *et al.* in that we require the agent being no better off reporting twice, while they require the agent receive the exact same payoff regardless of the agent reporting once or twice.

Using the notation established above, we formally define the desired properties:

P1: Myopically Strategyproof: Let p_i be the true belief of the agent making the ith trade, this belief being determined by the market history up to trade i, μ_i, and her signal, s_i. If this agent trades only once in a market, her reported belief, which maximizes her expected utility in a myopically strategy proof mechanism, will be her true belief,. Equivalently, if an agent does not take future payoff into consideration, she will report her true belief. Mathematically, for any n, $i \in \{1, \ldots, n\}$ and μ_i,

$$\begin{aligned} p_i &= \arg\max_{r \in [0,1]} p_i V_i(\rho(r, \mu_i, n, \omega = 1)) + (1 - p_i) V_i(\rho(r, \mu_i, n, \omega = 0)) \\ &= \arg\max_{r \in [0,1]} E_{\omega \sim p_i} V_i(\rho(r, \mu_i, n, \omega)). \end{aligned} \tag{5.1}$$

(As mentioned above, we use the notation $E_{\omega \sim p_i}$ to represent the expectation over the outcome of the event ω with respect to beliefs p_i, as well as, if appropriate, any randomization of the mechanism over the rewards.) Further, we also require that the expected value an agent gets as a result of such report, which we refer to as an *honest trade*, should be non-negative. In other words, our notion of myopic strategyproofness includes a standard individual rationality condition.

P2: Sybilproofness: A mechanism is sybilproof if an agent is no worse off reporting once in this mechanism than consecutively reporting more than once. Equivalently, an agent reporting in a mechanism with her private belief will not increase her expected value by reporting more than once with the same belief. For ease of analysis we consider a slightly weaker condition of reporting twice being no better than reporting once; our results hold for any number of consecutive trades. Mathematically, for any n, $i \in \{1, \ldots, n\}$ and μ_i, as well as any reports $r^{(1)}$ and $r^{(2)}$,

(5.2)

$$E_{\omega \sim p_i} V_i(\rho(p_i, \mu_i, n, \omega)) \geq E_{\omega \sim p_i} V_i(\rho(r^{(1)}, \mu_i, n+1, \omega) + \rho(r^{(2)}, \mu_{i+1}, n+1, \omega)),$$

where $\mu_{i+1} = (\mu_i, r^{(1)})$.

Note that this definition embodies a limited, myopic form of a sybil attack. In particular, it states that consecutive trading under different identities is not profitable, but does not rule out attacks involving non-consecutive trades. As discussed in the previous chapter, dishonest trading with multiple non-consecutive trades may be profitable even with current market mechanisms and risk neutral traders; further, the profitability of such attacks depends on the specific information distribution pattern as well as the mechanism chosen. In order to make a clean comparison to current market forms, we focus on simpler attacks involving consecutive trades only.

P3: Bounded Subsidy: There exists an upper bound β on the expected value of the subsidy the market maker needs to invest into the market. That is, for any number of trades in the market, n, and any collection of reports $r_i \in [0,1],\ i \in \{1,\ldots,n\}$, made,

$$E_\omega\left(\sum_{i=1}^{n}\rho(r_i,\mu_i,n,\omega)\right) < \beta.$$

Here we use E_ω to denote that the expectation is taken with respect to the true probability of the event taking place.

To summarize, we define the class of market-like mechanisms to be all mechanisms that *guarantee liquidity*, are *anonymous, myopically strategy proof and sybilproof,* and have *bounded market subsidy.*

Before we introduce our results, we must introduce the concepts of information structure, report informativeness, and expected normalized reward.

Information Structure: Recall that we define an *information structure* to consist of a set of possible signal realizations for each trader, and the posterior probability of the event given a subset of signal realizations (equivalently, the joint probability of signal realizations and the event outcome). An information structure can be represented as a tree of possible trading histories that would arise if each agent honestly reported her belief.

Informativeness: For a given information structure, we define *informativeness* of an agent k as the expected relative entropy reduction of the posteriors after including agent k's signal into the history. The relative entropy, or Kullback-Leibler divergence, between p and q is defined as (see [26, Chapter 2]):

$$D(p||q) \triangleq p \log \frac{p}{q} + (1-p) \log \frac{1-p}{1-q}.$$

Let $p(\mu_k)$ denote the posterior probability of the event after the first k signals are revealed. The informativeness of k is then defined as the expectation, over all future histories (all future traders' signal realizations) in the information structure, of $D(\omega||p(\mu_{k-1})) - D(\omega||p(\mu_k))$, where the outcome ω is treated as a distribution with all mass on the eventual outcome. This definition of informativeness measures how much each agent's report contributes to reducing the uncertainty about the event's occurrence.

Expected Normalized Reward: The informativeness and reward of an agent's report depend on the history of previous trades. Therefore, in order to compare the reward an agent receives from a report, we define the *expected normalized reward*. Note that we assumed agents are expected value maximizers; therefore, to compare the reward between agents, we normalize the expected reward they receive by the informativeness of that report. For example, if an agent with informativeness h and posterior belief p_i reports r_i, she expects to receive a reward of $E_{\omega \sim p_i} \rho(r_i, \mu_i, n, \omega)$; then her expected normalized reward is $\frac{E_{\omega \sim p_i} \rho(r_i, \mu_i, n, \omega)}{h}$.

5.3 Proposed Mechanism

In this section we review the work presented by Allen [5] and then present one mechanism that satisfies the properties outlined in Section 5.2.

5.3.1 Allen's Result

We first explore the core of Allen's result involving a single trader who is an expected value maximizer with an increasing value function $V(\cdot)$. We assume that the trader has unknown risk preference, *i.e.*, while function $V(\cdot)$ is known to be monotone, its exact form is not known to the market-maker. The trader is asked to predict the probability of an outcome ω.

Suppose the trader gives her assessment, $\hat{p}$, on the outcome of the event, while her true belief is p. If the event occurs, she receives a reward of 1 with probability $q(\hat{p}) = 1 - (1 - \hat{p})^2$, and if the event does not occur, she receives a reward of 1 with probability $\hat{q}(\hat{p}) = 1 - \hat{p}^2$.

As the trader is maximizing her expected value, she will maximize over $\hat{p}$

$$E_{\omega \sim p} V(\hat{p}) = p[q(\hat{p})V(1) + (1 - q(\hat{p}))V(0)] + (1 - p)[\hat{q}(\hat{p})V(1) + (1 - \hat{q}(\hat{p}))V(0)],$$

and thus set $\hat{p} = p$, as long as $V(1) > V(0)$. Therefore, in an Allen sweepstake (the $q(\cdot)$ and $\hat{q}(\cdot)$ functions are referred to as sweepstake functions), a trader will reveal her true belief to maximize her expected reward value. Note that this report also maximizes the probability of winning the sweepstake.

Allen's result follows from the fact that the expected value is linear in p (the probability of winning), and regardless of the form of the value function, as long as it is monotonically increasing, any sweepstakes $q(\cdot)$, $\hat{q}(\cdot)$ such that

$$p \frac{dq(\hat{p})}{d\hat{p}} + (1 - p) \frac{d\hat{q}(\hat{p})}{d\hat{p}} \bigg|_{\hat{p} = p} = 0$$

is satisfied will result in an expected value maximizing player revealing her true beliefs. The condition above means that Allen's sweepstake functions can be generalized to $q(\hat{p}) = a - b(1-\hat{p})^2$ and $\hat{q}(\hat{p}) = c - b\hat{p}^2$.

5.3.2 Proposed Sweepstakes

In this section we propose one possible sweepstakes method that is a generalization of Allen's result. Consider the following serial sweepstakes:

1. Each agent, in order, observes the previous agents' reports, and plays an individual sweepstake as defined by Allen with generalized sweepstake functions described in Section 5.3.1.

2. The outcome of the event is observed.

3. If there are n reports in the mechanism, then each player reporting $\hat{p}$ wins a reward of 1 with probability $q(\hat{p}) = \frac{1}{4^n}(1-(1-\hat{p})^2)$ if the event occurred and $\hat{q}(\hat{p}) = \frac{1}{4^n}(1-\hat{p}^2)$ if the event did not occur.

As the discussion below demonstrates, this method satisfies the desired properties outlined in Section 5.2.

Guaranteed liquidity: As each agent, by making a report, participates in an Allen sweepstake, she is guaranteed a nonnegative payoff, thus guaranteeing liquidity.

Anonymity: As each agent receives a reward independent of who she is, anonymity holds.

Bounded market subsidy: If the event occurs, and there are n agents in the mechanism, then at worst every agent would have reported 1, and each would receive a reward of 1 with probability $\frac{1}{4^n}$, meaning that in expectation the total reward given out is $n \cdot \frac{1}{4^n}$. (The case if the event does not occur is similar.) In either case, the highest expected reward given out would be $n \cdot \frac{1}{4^n} < \frac{1}{4}$ for $n \geq 1$.

Myopically strategy proof: This follows from the fact that each (myopic) agent is playing an independent Allen sweepstake with generalized sweepstake functions.

Sybilproof: In this setting, we will use the "full" definition of sybilproofness, *i.e.*, an agent is no worse off reporting once than consecutively reporting more than once. With a slight abuse of notation, we want to show that

$$EV(\rho(\text{reporting once})) \geq EV(\rho(\text{reporting } k \text{ times})),\ k \geq 1. \tag{5.3}$$

Consider an agent with true belief p. Let us define

$$\bar{P}(r) = p(1 - (1-r)^2) + (1-p)(1-r^2)$$

($\bar{P}(r)$ can be interpreted as this agent's estimate, based on her beliefs, of the probability of winning the sweepstake if she were the only trader in the mechanism and reported value r). Then, when there are n traders in the mechanism and this agent reports only once (in which case she would report her true belief, $r = p$), her expected value of payoff is

$$EV(\rho(\text{reporting once})) = \frac{1}{4^n}\bar{P}(p)\,V(1).$$

Consider now the case when, by creating sybils, this agent makes $k > 1$ (possibly different) consecutive reports, and thus increasing the number of trades in the mechanism to $n + k - 1$. Due to monotonicity of $V(\cdot)$, it can be argued that to maximize her expected value, the agent should still have each of the sybils report her true belief (since, by reporting $r = p$, each sybil plays in a sweepstake the agent believes it has the highest chance of winning). Taking into account

$V(0) = 0$, the expectation of the value received by the agent from sweepstakes of k sybils is

$$EV(\rho(\text{reporting } k \text{ times})) = \sum_{i=1}^{k} \binom{k}{i} \left(\frac{1}{4^{n+k-1}}\bar{P}(p)\right)^{i} \left(1 - \frac{1}{4^{n+k-1}}\bar{P}(p)\right)^{k-i} V(i)$$
$$\leq \sum_{i=1}^{k} \binom{k}{i} \left(\frac{1}{4^{n+k-1}}\bar{P}(p)\right)^{i} \left(1 - \frac{1}{4^{n+k-1}}\bar{P}(p)\right)^{k-i} iV(1)$$
$$= k\frac{1}{4^{n+k-1}}\bar{P}(p)\,V(1)$$
$$\leq \frac{1}{4^{n}}\bar{P}(p)\,V(1).$$

The first inequality holds because $V(\cdot)$ is concave and $V(0) = 0$, implying that $V(i) \leq iV(1)$ for $i \geq 1$.

We conclude that for $k \geq 1$, $EV(\rho(\text{reporting once})) \geq EV(\rho(\text{reporting } k \text{ times}))$.

The mechanism above possesses all of the desirable properties outlined in Section 5.2; however, the mechanism distributes rewards that are exponentially decreasing with the number of agents. Moreover, we know that if every agent makes an equally informative report, then the expected normalized reward also decreases exponentially with the number of reports. In the next section we will show that expected normalized reward must decrease exponentially, in the worst case over all information structures, whenever agents with unknown risk aversion are allowed to participate in a mechanism with the properties of Section 5.2.

5.4 Impossibility Result

In this section, we show that if agents with arbitrary risk averse preferences are allowed to participate in a mechanism of this class, then the normalized expected reward of the agents must, on at least one family of information structures, decrease exponentially with the number of agents.

An overview of our proof is as follows: We first show that if agents of arbitrary risk averse preferences participate in the mechanism, and the mechanism satisfies the individual rationality condition included in the definition of myopic strategyproofness, then the amount of reward given for any report must be non-negative. We then consider a situation with n traders who each trade once, honestly revealing their posterior probabilities. From this situation we identify a sybil attack with the same n agents, *i.e.*, the agents having the same posterior probabilities, making a total of $n+1$ reports, meaning that two of these reports are made by the same agent. Using the sybilproofness property, if follows that the expected reward of the agent making the two reports cannot be greater than her expected reward of making only one report. We then construct a different information structure with $n+1$ agents making the same $n+1$ reports made in the sybil attack setting, however the reports are truthful. We will then show that the expected rewards in the setting with $n+1$ agents are bounded by a constant multiple of the expected rewards in the sybil attack setting. Because the expected reward of the agent reporting twice in the "sybil setting" is bounded by her expected reward in the "honest setting" with n agents, it means that the expected reward of one of the agents in the setting with $n+1$ agents is bounded by the expected reward of one of the agents in the setting with n agents. We use this construction inductively, to show that the expected rewards must be exponentially decreasing in n, even if all traders are trading honestly.

5.4.1 Sybil Attack Payoff Bound

Before we show our main result, we first show that the reward in any mechanism satisfying the properties in Section 5.2 must be non-negative. We also show that the expected reward an agent receives for making the same report under different beliefs is bounded when the beliefs are bounded away from 0 and 1.

Lemma V.1. *If a mechanism satisfies the properties in Section 5.2 and allows all rational agents with arbitrary risk averse types to trade, then each agent receives non-negative reward for any report.*

Proof. Consider a mechanism that has the potential to distribute negative rewards. Now, consider an agent that receives a value of x for any payoff less than ϵ for some $\epsilon > 0$, and a value of ϵ for any payoff greater than or equal to ϵ, *i.e.*, $V(x) = \begin{cases} x & x < \epsilon \\ \epsilon & x \geq \epsilon \end{cases}$. Consider the reward distribution that is given to this agent if she trades honestly. By the myopic strategyproofness property, this agent's expected value from the rewards given to her must be non-negative. Given her value function, if any reward were negative, a small enough ϵ could be found such that the expected reward is negative.

□

For ease of analysis, we will use the following approximation of relative entropy:

$$\forall p, q \in [0.3, 0.7] \qquad 2(p-q)^2 \leq D(p||q) \leq 5(p-q)^2. \tag{5.4}$$

The left-hand inequality was proven by Okamoto [60]. The right-hand inequality follows, for $p, q \in [0.3, 0.7]$, from the inequalities of $\log(1+x) \leq x$ and $\log(1-x) \leq -x$. We make this approximation because we can now approximate the informativeness of a report (move) as the square of the difference between the posteriors after every report (move). Moreover, a constant factor will not impact the exponential decrease in the expected normalized reward of the report we are establishing.

Note that a mechanism cannot "tell" if it is faced with n agents trading honestly, or a sybil attack with $n-1$ agents with one trading twice consecutively. The only difference between these two situations is in the beliefs of the agents; in the sybil attack, the market prices after a sequence of moves may not match the posterior probability of the last trader. As all posterior probabilities are in a bounded range, we show that this has limited effect:

Lemma V.2. *For ρ_1, ρ_0 non-negative (with at least one of them positive), $p, q \in (0,1)$, and $\bar{p} = (1-p)$, and $\bar{q} = 1-q$:*

$$\min\left(\frac{p}{q}, \frac{\bar{p}}{\bar{q}}\right) \leq \frac{p\rho_1 + \bar{p}\rho_0}{q\rho_1 + \bar{q}\rho_0} \leq \max\left(\frac{p}{q}, \frac{\bar{p}}{\bar{q}}\right).$$

The proof follows from the observation that $\frac{\alpha p + \rho_o}{\alpha q + \rho_0}$ is monotone decreasing if $p < q$ and monotone increasing if $p > q$ in $\alpha > 0$, with an asymptote at 1.

Proof. Without loss of generality, we can assume $\rho_1 \geq \rho_0$.

Case 1: $\rho_0 = \rho_1 > 0$:

$$\frac{p\rho_1 + \bar{p}\rho_0}{q\rho_1 + \bar{q}\rho_0} = 1,$$

which satisfies the inequalities trivially.

Case 2: $\rho_1 > \rho_0$, *i.e.*, $\rho_1 = \rho_0 + \alpha$, where $\alpha > 0$:

$$\frac{p\rho_1 + \bar{p}\rho_0}{q\rho_1 + \bar{q}\rho_0} = \frac{p(\rho_0 + \alpha) + \bar{p}\rho_0}{q(\rho_0 + \alpha) + \bar{q}\rho_0} = \frac{\alpha p + \rho_0}{\alpha q + \rho_0}.$$

If $p \geq q$:

$$\min\left(\frac{p}{q}, \frac{\bar{p}}{\bar{q}}\right) \leq 1 \leq \frac{\alpha p + \rho_0}{\alpha q + \rho_0} \leq \frac{\alpha p}{\alpha q} \leq \max\left(\frac{p}{q}, \frac{\bar{p}}{\bar{q}}\right)$$

The first inequality holds by the definition of $p, q, \bar{p}, \bar{q}$. The remaining inequalities hold because $p \geq q$ and α is positive.

If $p < q$:

$$\min\left(\frac{p}{q}, \frac{\bar{p}}{\bar{q}}\right) \leq \frac{\alpha p}{\alpha q} \leq \frac{\alpha p + \rho_0}{\alpha q + \rho_0} \leq 1 \leq \max\left(\frac{p}{q}, \frac{\bar{p}}{\bar{q}}\right)$$

This first three inequalities hold because $p < q$ and α is positive. The fourth inequality holds by the definition of $p, q, \bar{p}, \bar{q}$.

□

Lemma V.2 states that the ratio between expectations over the same values with respect to two different probability mass functions is bounded by the ratio of the elements in the probability mass functions. We use this idea to bound the expected reward of the same report made in a mechanism under two different beliefs.

Lemma V.3. *An agent making a report r_i after history μ_i under belief $p(\mu_i, s_i)$, receives an expected reward that is a constant multiple of the expected reward she receives when making report r_i after history μ_i under belief $q(\mu_i, s_i)$, where $p(\mu_i, s_i), q(\mu_i, s_i) \in (p^l, p^u)$ with $0 < p^l < p^u < 1$. Mathematically,*

$$\begin{aligned}\min\left(\frac{p^l}{p^u}, \frac{\bar{p}^l}{\bar{p}^u}\right) &\leq \frac{p(\mu_i,s_i)\rho(r_i,\mu_i,n,\omega=1)+\bar{p}(\mu_i,s_i)\rho(r_i,\mu_i,n,\omega=0)}{q(\mu_i,s_i)\rho(r_i,\mu_i,n,\omega=1)+\bar{q}(\mu_i,s_i)\rho(r_i,\mu_i,n,\omega=0)} \\ &\leq \max\left(\frac{p^l}{p^u}, \frac{\bar{p}^l}{\bar{p}^u}\right).\end{aligned}$$

The proof follows by applying lemma V.2 in the setting of non-negative rewards, and using the fact the posteriors are bounded to (p^l, p^u) to get a bound for all posteriors. In our setting we use $p^l = 0.497$ and $p^u = 0.503$.

Proof. By Lemma V.1 we know that the reward to any agent is non-negative. This means that the expected reward for a risk neutral agent must also be non-negative, and may be written as: $p(\mu_i, s_i)\rho(r_i, \mu_i, n, \omega = 1) + (1 - p(\mu_i, s_i))\rho(r_i, \mu_i, n, \omega = 0)$ in one setting, and $q(\mu_i, s_i)\rho(r_i, \mu_i, n, \omega = 1) + (1 - q(\mu_i, s_i))\rho(r_i, \mu_i, n, \omega = 0)$ in the other.

The rewards are non-negative and $p(\mu_i, s_i), q(\mu_i, s_i) \in (p^l, p^u), (p^l, p^u) \subset (0, 1)$. We also assume that at least one of the rewards is strictly positive. From Lemma V.2 it follows that:

$$\begin{aligned}\min\left(\frac{p(\mu_i,s_i)}{q(\mu_i,s_i)}, \frac{\bar{p}(\mu_i,s_i)}{\bar{q}(\mu_i,s_i)}\right) &\leq \frac{p(\mu_i,s_i)\rho(r_i,\mu_i,n,\omega=1)+\bar{p}(\mu_i,s_i)\rho(r_i,\mu_i,n,\omega=0)}{q(\mu_i,s_i)\rho(r_i,\mu_i,n,\omega=1)+\bar{q}(\mu_i,s_i)\rho(r_i,\mu_i,n,\omega=0)} \\ &\leq \max\left(\frac{p(\mu_i,s_i)}{q(\mu_i,s_i)}, \frac{\bar{p}(\mu_i,s_i)}{\bar{q}(\mu_i,s_i)}\right)\end{aligned}$$

As both posteriors, $p(\mu_i, s_i)$ and $q(\mu_i, s_i)$, are bounded in some interval $(p^l, p^u) \subset (0,1)$, the result may be rewritten as:

$$\begin{aligned}\min(\tfrac{p^l}{p^u}, \tfrac{\bar{p}^l}{\bar{p}^u}) &\leq \tfrac{p(\mu_i,s_i)\rho(r_i,\mu_i,n,\omega=1)+\bar{p}(\mu_i,s_i)\rho(r_i,\mu_i,n,\omega=0)}{q(\mu_i,s_i)\rho(r_i,\mu_i,n,\omega=1)+\bar{q}(\mu_i,s_i)\rho(r_i,\mu_i,n,\omega=0)} \\ &\leq \max(\tfrac{p^l}{p^u}, \tfrac{\bar{p}^l}{\bar{p}^u})\end{aligned}$$

□

5.4.2 Inductive Construction of Information Structures

To reiterate, we define an *information structure* to be the set of all histories and all honest reports with no sybils for all combinations of all the agents' signal realizations. This can be represented as a tree where each level represents an agent, each node represents the belief, or market prediction, of an agent before her report is made, and each branch represents the signal realizations that lead to the current market prediction. (Note that, except for the first agent, each agent's belief is a posterior distribution, calculated based on observing the (honest) reports of the other agents and inferring their signals from their reports.) Such a representation can be seen in Figure 5.4. In Figure 5.4, there are 3 agents, each starting with a common prior on ω of 0.5. If the first agent sees a "right" signal, she makes an honest report of $0.5 + \frac{c}{2}$, and all agents calculate their posterior. If the second agent sees a "left" signal, and had a prior of $0.5 + \frac{c}{2}$, she reports $0.5 + \frac{c}{4}$, and so on.

We will now describe our inductive procedure to construct a family of information structures $I^{(n)}$. Fix a mechanism in the class described in Section 5.2. The construction of $I^{(n)}$ depends on the mechanism, but enables us to show that, for any given mechanism, there is a family of information structures on which the mechanism must exhibit exponentially decreasing normalized rewards. The inductive construction requires one additional parameter c, which bounds the agents' posteriors. In this chapter we use $c = 0.003$.

We assume that each agent has a common prior of 0.5 on the outcome of the event. Further, all agents' posteriors are bounded between $0.5 - c$ and $0.5 + c$. In the constructed structures, the first $j + 1$ agents, with probability $\frac{1}{2}$, observe one of two signals (*e.g.*, *right, left*). The remaining agents observe their signals with probabilities that make their beliefs consistent with their reports. How the value of j is selected will be explained later in this section, and we will note that the actual reports of agents after $j + 1$ do not influence the results.

The first structure, $I^{(1)}$, has only one agent. If the agent receives a *right* signal then her posterior is $0.5+c$, and if she receives a *left* signal then her posterior is $0.5-c$. (In Figure 5.1, we illustrate the information structure $I^{(1)}$; we say that the informativeness of the report in $I^{(1)}$ is $(0.5+c-0.5)^2 = c^2$). We then create $I^{(2)}$ by splitting the report in $I^{(1)}$ into two symmetric reports as illustrated in Figure 5.2. When we say a report is split, we mean that an additional level is added to the binary tree representing the information structure. Further, the history up to the split report remains the same and history after the split changes in order to make all reports consistent with the joint probability distribution underlying the resulting information structure. We say that a report is symmetric when for all signal realizations the absolute value of the differences in the prior, before the report is made, and the posterior, after the report is made, of that report are equal. For instance, in Figure 5.2, the first report is symmetric because $|0.5-(0.5-\frac{c}{2})| = |0.5-(0.5+\frac{c}{2})|$, and the second report is symmetric because $|(0.5-\frac{c}{2})-(0.5-c)| = |(0.5-\frac{c}{2})-0.5| = |(0.5+\frac{c}{2})-0.5| = |(0.5+\frac{c}{2})-(0.5+c)|$. Similarly, all of the reports in Figure 5.4 are symmetric, however the last report in Figure 5.3 may or may not be symmetric, depending on the values of α_1–α_6.

Starting with $I^{(3)}$, we construct $I^{(k+1)}$ from $I^{(k)}$ by comparing the expected reward of the two symmetric reports just added in $I^{(k)}$. Note that we will always be

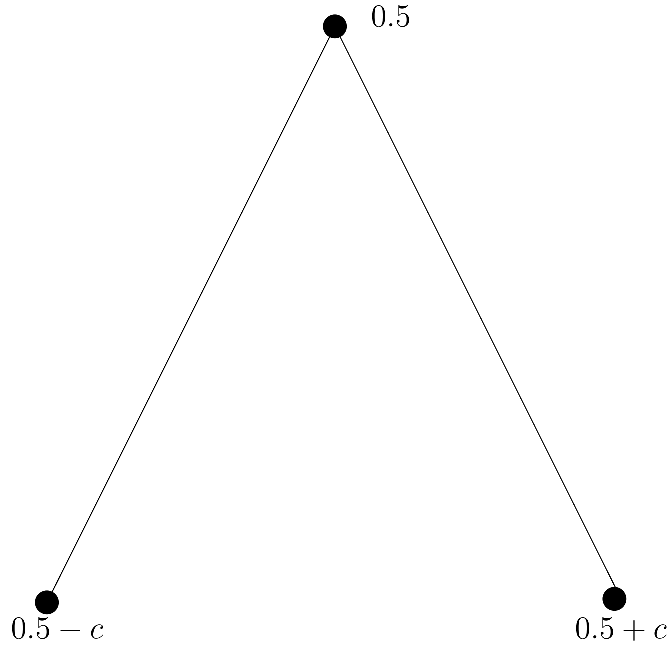

Figure 5.1: $n = 1$

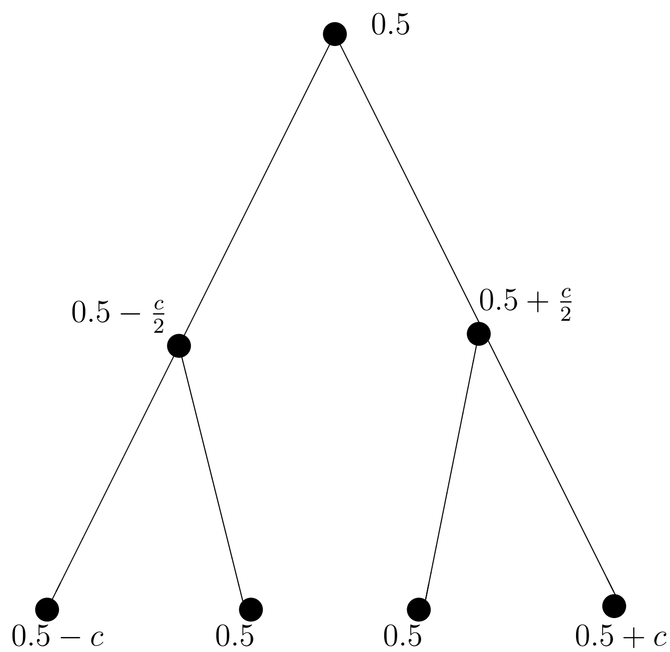

Figure 5.2: $n = 2$

comparing two symmetric reports, because by construction we are always splitting a report into two symmetric reports. Say, report j and report $j+1$ are these reports. For example, in $I^{(2)}$, $j = 1$. If the expected reward of agent j is 4 times greater than the expected reward of agent $j+1$, then $I^{(k+1)}$ is created by splitting report j into two symmetric reports (this case is exhibited when creating $I^{(3)}$ shown in Figure 5.3 from $I^{(2)}$). Assuming report j shifted the posterior by $\pm\beta$ (meaning for one signal realization the difference in posteriors will be β and in the other signal realization the difference in posteriors will be $-\beta$), then these new reports will shift the posterior by $\pm 0.9\beta$. Note that the report following the split in $I^{(k+1)}$ will become asymmetric in order to make the beliefs consistent with the reports (the exact value of this report does not matter because rewards are independent of future reports). For $I^{(k+1)}$ constructed in this manner, we set $j_{k+1} = j$, where j_{k+1} is the first report we consider for splitting in the next step of the inductive construction process.

Otherwise, *i.e.*, if the expected reward of agent j in structure $I^{(k)}$ is not 4 times greater than the expected reward of agent $j+1$, then $I^{(k+1)}$ is created by splitting report $j+1$ into two symmetric reports (this case is exhibited when creating $I^{(3)}$ shown in Figure 5.4 from $I^{(2)}$). Assuming report $j+1$ shifted the posterior by $\pm\beta$, then these new reports will shift the posterior by $\pm 0.5\beta$. For $I^{(k+1)}$ constructed in this manner, we set $j_{k+1} = j+1$; again, j_{k+1} is the first report we consider for splitting next.

By the construction of $I^{(k+1)}$ we keep the history before the split report the same as it was in $I^{(k)}$. Moreover, if we let j_k^s be the report split in $I^{(k)}$, then we know that $j_{k+1} = j_k^s$, where j_{k+1} is the first report we consider in $I^{(k+1)}$. We note that so long as the posterior before a report is between the posteriors after the report under two different signal realizations, we can always construct a feasible information structure (*i.e.*, an information structure that is consistent with Bayesian updating on some joint distribution on s_i and ω). By the construction of $I^{(k+1)}$ from $I^{(k)}$ we note that the posteriors before a report are always bounded between the posteriors after, meaning the information structures we construct are feasible.

For example, in Figure 5.2 we show $I^{(2)}$ when there are two agents in the mechanism, $n = 2$ (the informativeness for the first report is $((0.5 + \frac{c}{2}) - 0.5)^2 = \frac{c^2}{4}$, and the second is $((0.5 + c) - (0.5 + \frac{c}{2}))^2 = \frac{c^2}{4}$) ; We show $I^{(3)}$ in Figure 5.3 if report $j = 1$ was split, and in Figure 5.4 we show $I^{(3)}$ when report $j = 2$ was split.

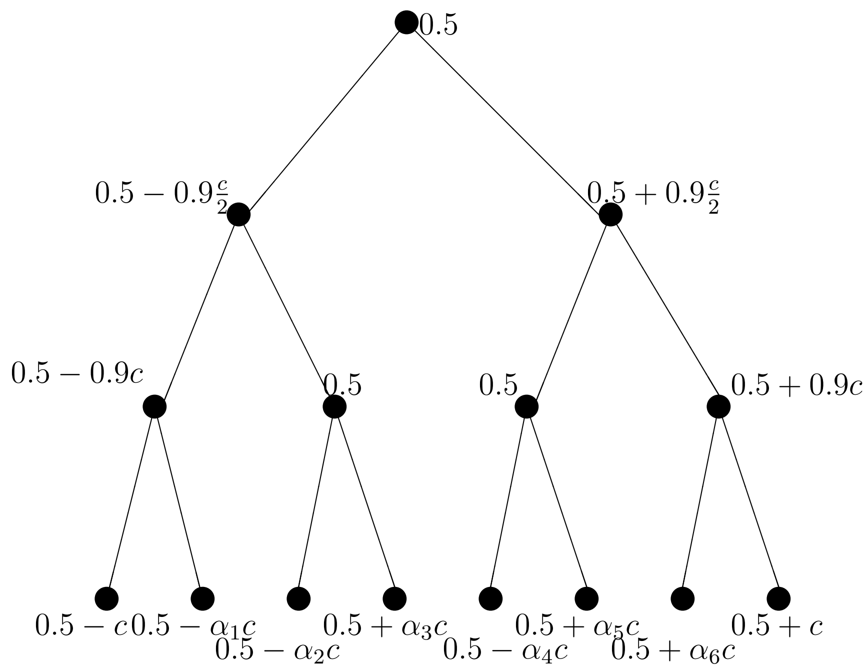

Figure 5.3: $n = 3$ case 1

5.4.3 Exponential Bound

The exponential bound that we show follows from the fact that we allow agents of arbitrary risk averse preferences to participate in a mechanism satisfying the properties in Section 5.2.

Theorem V.4. *If an anonymous, guaranteed liquidity mechanism satisfies properties* **P1–P3***, then, there is a family of information structures $I^{(n)}$, each parameterized by a number n of agents, such that, if all agents perform perfect Bayesian updating according to the structure $I^{(n)}$ and report their posteriors honestly, the minimum*

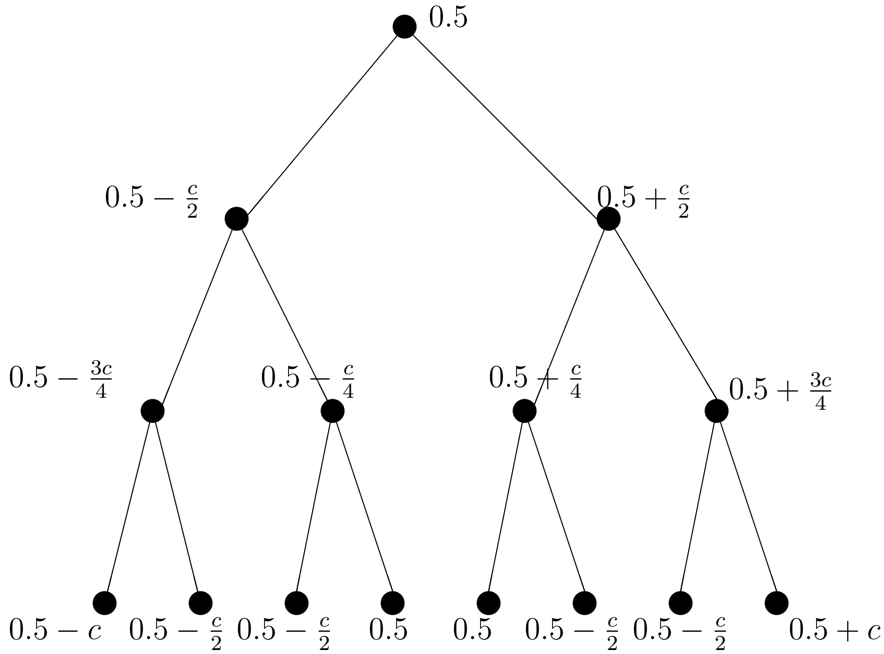

Figure 5.4: $n = 3$ case 2

normalized expected reward of an agent must decrease exponentially with n*.*

The proof follows inductively by building a family of information structures, where one of the two last added reports is split into two new reports, to simulate a sybil attack. The report split is determined by looking at the expected rewards of the two new added reports. In one case the new reports have 0.81 times the informativeness of the original report, and in the other the new reports have 0.25 times the informativeness. The sybilproofness property, specialized for a risk neutral agent, implies that the expected reward from the original report must be greater than or equal to the expected reward of the two new reports. Using the fact that the rewards are non-negative and Lemma V.3, we show that one of the new reports, made in the honest setting, has an expected normalized reward that is no more than $\frac{c+0.5}{0.5-c} \cdot \frac{80}{81}$ the expected normalized reward of the report that is split. We note that for $c = 0.003$, $\frac{c+0.5}{0.5-c} \cdot \frac{80}{81} < 1$. Iteratively applying this observation, we can show that there exists at least one trade that has an expected normalized reward exponentially smaller than the expected normalized reward of only participating in the mechanism in which $n = 1$.

Proof. In the following we will be applying the sybilproofness property at every step, while constructing a sequence of information structures $I^{(n)}$, $n = 1, 2, \ldots$, where n is the number of agents trading. During a sybil attack, a risk neutral agent would be comparing the expected payoff she may receive for trading once to the expected payoff she may receive when trading twice with the same information. In order for the sybilproof property to hold, the expected payoff over her beliefs in information structure $I^{(n)}$, from reporting once must be at least as large as the expected payoff from reporting twice. We write this as:

(5.5)

$$E_{\omega \sim P\{I^{(n)}\}} \rho(r_n, \mu_n, n, \omega) \geq E_{\omega \sim P\{I^{(n)}\}} [\rho(r'_n, \mu_n, n+1, \omega) + \rho(r'_{n+1}, \mu_{n+1}, n+1, \omega)].$$

Above we use $E_{\omega \sim P\{I^{(n)}\}}[\cdot]$ to note that the expectation of the reward is with respect to the posteriors in $I^{(n)}$. In order to avoid adding unnecessary notation we use the following convention to discern if we are considering a case of a sybil attack or a case of a honest trade:

$$\begin{aligned} \text{honest trade:} \quad & E_{\omega \sim P\{I^{(N)}\}}[\rho(r_j^N, \mu_j, N, \omega) + \rho(r_{j+1}^N, \mu_{j+1}, N, \omega)] \\ \text{sybil attack:} \quad & E_{\omega \sim P\{I^{(N)}\}}[\rho(r_{j+1}^{N+1''}, \mu''_{j+1}, N+1, \omega) + \rho(r_{j+2}^{N+1''}, \mu''_{j+1}, N+1, \omega)]. \end{aligned}$$

The key difference here is that, in the sybil attack setting, the number of traders is N, having beliefs defined in $I^{(N)}$, is not the same as the number of reports, $N + 1$, the third value in the reward function. However, in the honest trade setting they are both equal to N.

Below we present a proof by induction:

Inductive Hypothesis: There exists an information structure $I^{(N)}$ with a pair of reports, $j, j+1$ that are symmetric and each have informativeness t_N, and

(5.6)

$$E_{\omega\sim P\{I^{(N)}\}}[\rho(r_j^N,\mu_j,N,\omega)+\rho(r_{j+1}^N,\mu_{j+1},N,\omega)] \leq \left(\frac{0.5+c}{0.5-c}\right)^{N-1}\left(\frac{80}{81}\right)^{N-2} t_N 4l.$$

Here we set $l = \frac{E_{\omega\sim P\{I^{(1)}\}}\rho(r_1^1,\mu_1,1,\omega)}{c^2}$ to be the expected normalized reward of an agent reporting once in information structure $I^{(1)}$. In $I^{(1)}$ an agent starts with a prior of 0.5, with equal probability observes a value *(right, left)* of her private binary signal, and reports $0.5+c$ if she observes one value (*right*) or $0.5-c$ if she observes the other (*left*).

If we show (5.6) holds for all N, then we would have shown the result of the theorem, as we would have constructed a family of information structures where there exists at least one report that has a normalized expected reward exponentially smaller than l. In particular, as all rewards are non-negative, it follows that all expected rewards are also non-negative. This means:

$$\begin{aligned}
& E_{\omega\sim P\{I^{(N)}\}}\rho(r_j^N,\mu_j,N,\omega) && \leq E_{\omega\sim P\{I^{(N)}\}}[\rho(r_j^N,\mu_j,N,\omega)+\rho(r_{j+1}^N,\mu_{j+1},N,\omega)] \\
\implies\; & E_{\omega\sim P\{I^{(N)}\}}\rho(r_j^N,\mu_j,N,\omega) && \leq \left(\frac{0.5+c}{0.5-c}\right)^{N-1}\left(\frac{80}{81}\right)^{N-2} t_N 4l \\
\implies\; & \frac{E_{\omega\sim P\{I^{(N)}\}}\rho(r_j^N,\mu_j,N,\omega)}{t_N} && \leq \left(\frac{0.5+c}{0.5-c}\right)^{N-1}\left(\frac{80}{81}\right)^{N-2} 4l.
\end{aligned}$$

Note that the last inequality gives us the desired exponential reduction for $c \leq 0.003$.

Base case: For $N=2$, we start with $I^{(1)}$ where only one agent is making a single report with an initial prior of 0.5 and either reports $0.5+c$ with probability $\frac{1}{2}$ or $0.5-c$ with probability $\frac{1}{2}$. This agent now is considering making two reports with the same prior she had in $I^{(1)}$. The two reports considered have equal informativeness and are created by splitting the report in $I^{(1)}$ into two symmetric reports as depicted in Figure 5.2. In order for the agent not to make these two reports via two sybils, the sybilproofness property must hold:

$$E_{\omega\sim P\{I^{(1)}\}}\rho(r_1,\mu_1,1,\omega) \geq E_{\omega\sim P\{I^{(1)}\}}[\rho(r'_1,\mu_1,2,\omega)+\rho(r'_2,\mu_2,2,\omega)]. \tag{5.7}$$

Using Lemma V.3, we can compare $E_{\omega\sim P\{I^{(1)}\}}[\rho(r'_1,\mu_1,2,\omega)+\rho(r'_2,\mu_2,2,\omega)]$ to $E_{\omega\sim P\{I^{(2)}\}}[\rho(r'_1,\mu_1,2,\omega)+\rho(r'_2,\mu_2,2,\omega)]$ as two agents make the same reports under different priors, noting that by construction both reports are contained in $[0.5-c, 0.5+c]$. From Lemma V.3 we know that the ratio of the expected rewards is bounded:

$$\frac{0.5-c}{0.5+c}E_{\omega\sim P\{I^{(2)}\}}[\rho(r'_1,\mu_1,2,\omega)+\rho(r'_2,\mu_2,2,\omega)] \leq E_{\omega\sim P\{I^{(1)}\}}[\rho(r'_1,\mu_1,2,\omega)+\rho(r'_2,\mu_2,2,\omega)]. \tag{5.8}$$

By the definition of $I^{(2)}$, the informativeness of either of the two reports in $I^{(2)}$ is $\left(\frac{1}{2}\right)^2 c^2 = \frac{1}{4}c^2$, and the informativeness of the only report is $I^{(1)}$ is c^2. Note that by the definition of t_n, we can write $t_2 = \frac{1}{4}c^2$. We now examine the normalized expected reward for the reports in $I^{(2)}$:

$$\frac{E_{\omega\sim P\{I^{(2)}\}}[\rho(r'_1,\mu_1,2,\omega)+\rho(r'_2,\mu_2,2,\omega)]}{\frac{1}{4}c^2}.$$

Similarly, we calculate the normalized expected reward of the report in $I^{(1)}$ as:

$$\frac{E_{\omega\sim P\{I^{(1)}\}}\rho(r_1,\mu_1,1,\omega)}{c^2}.$$

Combining (5.7) and (5.8) we have:

$$\begin{aligned}
& \tfrac{0.5-c}{0.5+c}E_{\omega\sim P\{I^{(2)}\}}[\rho(r'_1,\mu_1,2,\omega)+\rho(r'_2,\mu_2,2,\omega)] && \leq E_{\omega\sim P\{I^{(1)}\}}\rho(r_1,\mu_1,1,\omega) \\
\Longrightarrow\quad & \tfrac{0.5-c}{0.5+c}\tfrac{E_{\omega\sim P\{I^{(2)}\}}[\rho(r'_1,\mu_1,2,\omega)+\rho(r'_2,\mu_2,2,\omega)]}{c^2} && \leq \frac{E_{\omega\sim P\{I^{(1)}\}}\rho(r_1,\mu_1,1,\omega)}{c^2} \\
\Longrightarrow\quad & \tfrac{0.5-c}{0.5+c}\tfrac{E_{\omega\sim P\{I^{(2)}\}}[\rho(r'_1,\mu_1,2,\omega)+\rho(r'_2,\mu_2,2,\omega)]}{c^2} && \leq l.
\end{aligned}$$

The last inequality follows from the definition of l. We can now multiply both sides by 4 and have the inequality:

$$\begin{aligned} \frac{0.5-c}{0.5+c}\frac{E_{\omega\sim P\{I^{(2)}\}}[\rho(r_1',\mu_1,2,\omega)+\rho(r_2',\mu_2,2,\omega)]}{\frac{1}{4}c^2} &\leq 4l \\ \implies \frac{0.5-c}{0.5+c}\frac{E_{\omega\sim P\{I^{(2)}\}}[\rho(r_1',\mu_1,2,\omega)+\rho(r_2',\mu_2,2,\omega)]}{t_2} &\leq 4l. \end{aligned}$$

The last inequality holds by the definition of t_2. We now multiply both sides by t_2 and $\frac{0.5+c}{0.5-c}$ to get:

$$\begin{aligned} E_{\omega\sim P\{I^{(2)}\}}[\rho(r_1',\mu_1,2,\omega)+\rho(r_2',\mu_2,2,\omega)] &\leq \frac{0.5+c}{0.5-c}t_2 4l \\ (5.9) \quad \implies \quad E_{\omega\sim P\{I^{(2)}\}}[\rho(r_1',\mu_1,2,\omega)+\rho(r_2',\mu_2,2,\omega)] &\leq \frac{0.5+c}{0.5-c}\left(\frac{80}{81}\right)^0 t_2 4l \end{aligned}$$

We have just shown that (5.6) hold for $N = 2$, and will now look at the inductive step.

Inductive Step: Assuming that the inductive hypothesis holds for N, will show that it holds for $N + 1$ by constructing an appropriate information structure $I^{(N+1)}$ that is derived from $I^{(N)}$. Recall that j and $j + 1$ are the two reports just created in $I^{(N)}$. We will construct $I^{(N+1)}$ by setting $I^{(N+1)} = I^{(N+1')}$ if the expected reward of report j is greater than 4 times the expected reward of report $j + 1$, and we set $I^{(N+1)} = I^{(N+1'')}$ otherwise. $I^{(N+1')}$ is defined by splitting report j in $I^{(N)}$ into two symmetric reports, each with a difference in the posteriors before and after the report of $0.9\cdot\sqrt{t_N}$, and the reports after split will become asymmetric so that the beliefs can remain consistent (note that the future reports do not matter as long as consistency is preserved and the posteriors are bounded between $0.5 - c$ and $0.5 + c$). $I^{(N+1'')}$ is defined by splitting report $j + 1$ in $I^{(N)}$ into two symmetric reports, each with a difference in the posteriors before and after the report of $\frac{1}{2}\cdot\sqrt{t_N}$, with report j in $I^{(N+1'')}$ the same as report j in $I^{(N)}$. We will now show that (5.6) holds for $I^{(N+1)}$.

Consider the case $I^{(N+1)} = I^{(N+1')}$. We know that by the sybilproofness property the following must hold:

$$E_{\omega\sim P\{I^{(N)}\}}[\rho(r_j^{N+1'},\mu'_j,N+1,\omega)+\rho(r_{j+1}^{N+1'},\mu'_{j+1},N+1,\omega)] \leq$$
$$E_{\omega\sim P\{I^{(N)}\}}\rho(r_j^N,\mu_j,N,\omega)$$
$$\implies \quad E_{\omega\sim P\{I^{(N)}\}}[\rho(r_j^{N+1'},\mu'_j,N+1,\omega)+\rho(r_{j+1}^{N+1'},\mu'_{j+1},N+1,\omega)] \leq$$
$$\tfrac{4}{5}E_{\omega\sim P\{I^{(N)}\}}[\rho(r_j^N,\mu_j,N,\omega)+\rho(r_{j+1}^N,\mu_{j+1},N,\omega)].$$

The last inequality holds by the conditions on the expected rewards of reports j and $j+1$ in $I^{(N)}$ that lead to $I^{(N+1')}$. We recall by Lemma V.3 that:

$$\tfrac{c-0.5}{c+0.5}E_{\omega\sim P\{I^{(N+1')}\}}[\rho(r_j^{N+1'},\mu'_j,N+1,\omega)+\rho(r_{j+1}^{N+1'},\mu'_{j+1},N+1,\omega)] \leq$$
$$E_{\omega\sim P\{I^{(N)}\}}[\rho(r_j^{N+1'},\mu'_j,N+1,\omega)+\rho(r_{j+1}^{N+1'},\mu'_{j+1},N+1,\omega)]$$

meaning:

$$E_{\omega\sim P\{I^{(N+1')}\}}[\rho(r_j^{N+1'},\mu'_j,N+1,\omega)+\rho(r_{j+1}^{N+1'},\mu'_{j+1},N+1,\omega)] \quad \leq$$
$$\tfrac{c+0.5}{c-0.5}\cdot\tfrac{4}{5}E_{\omega\sim P\{I^{(N)}\}}[\rho(r_j^N,\mu_j,N,\omega)+\rho(r_{j+1}^N,\mu_{j+1},N,\omega)]$$
$$\implies \quad E_{\omega\sim P\{I^{(N+1')}\}}[\rho(r_j^{N+1'},\mu'_j,N+1,\omega)+\rho(r_{j+1}^{N+1'},\mu'_{j+1},N+1,\omega)] \quad \leq$$
$$\left(\tfrac{0.5+c}{0.5-c}\right)^N\cdot\tfrac{4}{5}\left(\tfrac{80}{81}\right)^{N-2}t_N 4l.$$

The last inequality holds by the inductive hypothesis (5.6). We recall that the informativeness of each the two new reports in $I^{(N+1')}$ is $0.81t_N$, *i.e.*, $t_{N+1}=0.81t_N$ in $I^{(N+1')}$ thus,

$$E_{\omega\sim P\{I^{(N+1')}\}}[\rho(r_j^{N+1'},\mu'_j,N+1,\omega)+\rho(r_{j+1}^{N+1'},\mu'_{j+1},N+1,\omega)] \quad \leq$$
$$\left(\tfrac{0.5+c}{0.5-c}\right)^N\left(\tfrac{80}{81}\right)^{N-1}t_{N+1}4l.$$

This last inequality holds by observing $\frac{80}{81}\cdot 0.81=\frac{4}{5}$.

Consider now the case $I^{(N+1)}=I^{(N+1'')}$. By the sybilproofness property the following must hold:

$$E_{\omega\sim P\{I^{(N)}\}}[\rho(r_{j+1}^{N+1''},\mu''_{j+1},N+1,\omega)+\rho(r_{j+2}^{N+1''},\mu''_{j+1},N+1,\omega)] \leq E_{\omega\sim P\{I^{(N)}\}}\rho(r_j^N,\mu_j,N,\omega)$$

$$\implies E_{\omega\sim P\{I^{(N)}\}}[\rho(r_{j+1}^{N+1''},\mu''_{j+1},N+1,\omega)+\rho(r_{j+2}^{N+1''},\mu''_{j+1},N+1,\omega)] \leq \tfrac{1}{5}E_{\omega\sim P\{I^{(N)}\}}[\rho(r_j^N,\mu_j,N,\omega)+\rho(r_{j+1}^N,\mu_{j+1},N,\omega)].$$

The last inequality holds by the conditions on the expected rewards of reports j and $j+1$ in $I^{(N)}$ that lead to $I^{(N+1'')}$. Recall that by Lemma V.3,

$$\tfrac{c-0.5}{c+0.5}E_{\omega\sim P\{I^{(N+1'')}\}}[\rho(r_{j+1}^{N+1''},\mu''_{j+1},N+1,\omega)+\rho(r_{j+2}^{N+1''},\mu''_{j+1},N+1,\omega)] \leq E_{\omega\sim P\{I^{(N)}\}}[\rho(r_{j+1}^{N+1''},\mu''_{j+1},N+1,\omega)+\rho(r_{j+2}^{N+1''},\mu''_{j+1},N+1,\omega)]$$

meaning:

$$E_{\omega\sim P\{I^{(N+1'')}\}}[\rho(r_{j+1}^{N+1''},\mu''_{j+1},N+1,\omega)+\rho(r_{j+2}^{N+1''},\mu''_{j+1},N+1,\omega)] \leq \tfrac{c+0.5}{c-0.5}\cdot\tfrac{1}{5}E_{\omega\sim P\{I^{(N)}\}}[\rho(r_j^N,\mu_j,N,\omega)+\rho(r_{j+1}^N,\mu_{j+1},N,\omega)]$$

$$\implies E_{\omega\sim P\{I^{(N+1'')}\}}[\rho(r_{j+1}^{N+1''},\mu''_{j+1},N+1,\omega)+\rho(r_{j+2}^{N+1''},\mu''_{j+1},N+1,\omega)] \leq \left(\tfrac{c+0.5}{c-0.5}\right)^N\cdot\tfrac{1}{5}\cdot\left(\tfrac{80}{81}\right)^{N-2}t_N 4l.$$

The last inequality holds by the inductive hypothesis (5.6). We note that the informativeness of each the two new reports in $I^{(N+1'')}$ is $0.25t_N$, *i.e.*, $t_{N+1}=0.25t_N$ in $I^{(N+1'')}$. This allows us to write:

$$\implies E_{\omega\sim P\{I^{(N+1'')}\}}[\rho(r_{j+1}^{N+1''},\mu''_{j+1},N+1,\omega)+\rho(r_{j+2}^{N+1''},\mu''_{j+1},N+1,\omega)] \leq \left(\tfrac{c+0.5}{c-0.5}\right)^N\cdot\tfrac{1}{5}\cdot\left(\tfrac{80}{81}\right)^{N-1}t_{N+1}4l.$$

The last inequality holds by observing $\frac{1}{5}\leq\frac{80}{81}\cdot 0.25$.

Note that in either case the inductive hypothesis holds, meaning that the exponential decrease is preserved.

□

5.5 Discussion and Future Work

In this chapter we present one mechanism that satisfies the desirable properties of subsidized prediction markets that allows agents with unknown risk averse value functions to participate. However, this mechanism requires that the normalized expected rewards exponentially decrease with the number of agents. We then show that as long as the risk aversion structure of the agents is not known, for any mechanism in the class described in Section 5.2 that allows agents with arbitrary risk averse value functions to participate, the normalized expected reward decreases exponentially with the number of agents. This result may help in developing new forms of prediction markets, that are guaranteed to converge to the truthful equilibrium even in the presence of risk averse agents.

The presented result relies on the fact that agents can be highly risk averse, leading to the observation that only non-negative rewards are given for any report. This seems like a very restrictive setting, and thus considering a setting where the amount of risk aversion is bounded, and that does not have the exponential decrease in the normalized expected reward, is of interest. Moreover, the mechanism that we presented satisfied all of the desired properties due to the fact that the expected value of a reward is linear in the probabilities. In the future it would be interesting if a mechanism satisfying the properties in Section 5.2 could be found where the agents could be characterized by a more general decision model, such as prospect theory or uncertainty aversion, that is not necessarily linear in the probabilities.

CHAPTER VI

Conclusion

In this dissertation we address issues in information management problems using an algorithmic and economic perspective. In particular, we focus on the two topics of information procurement and delivery by way of prediction markets and network routing, respectively. To summarize, our contributions are:

- **Devise robust routing policies that take into account demand fluctuations and congestion control.** We formulate mathematical models used to generate routing policies that take into account active congestion control and are robust to demand variation. Combining random early detection (RED), multi-protocol label switching (MPLS), robust optimization, and non-linear programming we show that routing policies that take into account natural demand variation perform better than current routing policies. Though the resulting model is $\mathcal{NP}$-hard, we are able to show that solutions returned by modern non-linear solvers outperform routing policies currently in use when congestion is high.

- **Show bluffing can exist in prediction markets with non-myopic traders. Devise prediction market mechanisms robust to non-myopic traders.** In prediction markets using market scoring rules we characterize the behavior of strategic traders with complementary information. We show that traders have incentive to bluff and do not fully reveal their information. Further, we propose the δ-discounted market scoring rule to make the market robust to strategic traders.

- **Devise a prediction market mechanism that accurately aggregates the information of risk averse traders. We show that any prediction market that aggregates the beliefs of risk averse traders exponentially decreases the reward as the number of participants increases.** Current prediction markets assume risk neutral traders. However, risk aversion is exhibited by players in both real money and play money settings. We describe the desirable properties that a prediction market satisfies. We propose a sweepstakes mechanism that satisfies all of these properties even in the presence of risk averse traders. Our mechanism exponentially decreases rewards to agents as the number of agents increases. We show that this characteristic is necessary in any prediction market that accurately aggregates the beliefs of risk averse traders

Future Work

In what follows, we present some future research directions in information procurement and delivery:

- **Design quick approximations to the $\mathcal{NP}$-hard routing model accounting for congestion control.** In Chapter II, we show the usefulness of the designed routing policies as they outperform existing routing policies. However, the model to generate the designed routing policies is $\mathcal{NP}$-hard and requires large amounts of computation with current non-linear programming solvers.

Can we design fast approximation algorithms that achieve near-optimal solutions and scale well with the size of the network?

- **Design good routing heuristics, based on local knowledge instead of global optimization.** The models and results in Chapter II assume global knowledge of the entire network and its demand in designing the routing policy. In an evolving network, this information is not available. Therefore, finding a good routing policy using only local information is of interest.

- **Identify the information setting of traders in a prediction market using their trading histories.** In our analysis of non-myopic strategies in prediction markets (Chapter IV), we note that Chen *et al.* [22] show a setting where the myopic strategy is an equilibrium strategy. These two results are extended by Chen *et al.* [20] to show non-myopic strategies are equilibrium strategies when traders have complementary information and myopic strategies are equilibrium strategies when traders have substitutable information. In the future, identifying the relationship between traders' information by their trading history is of interest.

- **Analysis of prediction markets on events with uncertain ending times.** In Chapter IV we introduce the discounted market scoring rule to alleviate the issue of bluffing in prediction markets. Discounting may be characterized as traders being uncertain if they will come back to the market to trade again, therefore, they are discounting their future expected payoffs. This uncertainty naturally leads to the consideration of prediction markets on events with uncertain ending times. Consider a project manager trying to decide which of two unlikely long term projects to fund. If the success of each of these projects was traded on a current prediction market, their price would be very close to zero, making them difficult to differentiate. However, if these two projects were represented as a claim with an uncertain ending time, "Will the first project finish before the second project?", the projects could be differentiated giving the manager additional information. Such contingencies lead to new strategic questions for traders. For example, if a trader knows that the first project will be a success but is uncertain about the second project, they may not know when to reveal their information. Characterizing trader behavior in this setting is of interest for future prediction market design.

- **Analyze the necessity of exponential reward decrease in prediction markets with traders with bounded risk aversion.** In Chapter V we discuss the impact of risk averse traders in prediction markets. We use the fact traders of arbitrary risk averse preferences may participate in a market to show that the reward distributed in any mechanism must be non-negative. With the non-negative reward requirement, we show that for any mechanism to satisfy 5 desirable properties the mechanism must exponentially reduce the distributed reward as the number of traders increases. This leaves the open question of can this exponential reduction be eliminated if the risk aversion of the participating traders is bounded?

BIBLIOGRAPHY

BIBLIOGRAPHY

[1] Abilene network demand data. http://www.cs.utexas.edu/~yzhang/research/AbileneTM/.

[2] Abilene Network. http://abilene.internet2.edu.

[3] Shipra Agrawal, Erick Delage, Mark Peters, Zizhuo Wang, and Yinyu Ye. A unified framework for dynamic pari-mutuel information market design. In *EC '09: Proceedings of the tenth ACM Conference on Electronic Commerce*, pages 255–264, 2009.

[4] Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, February 1993.

[5] Franklin Allen. Discovering personal probabilities when utility functions are unknown. *Management Science*, 33(4):542–544, 1987.

[6] Franklin Allen and Douglas Gale. Stock-price manipulation. *Review of Financial Studies*, 5(3):503–529, 1992.

[7] Marco Angrisani, Antonio Guarino, Steffen Huck, and Nathan Larson. No-trade in the laboratory. *Social Science Research Network Working Paper Series*, October 2008.

[8] David Applegate and Edith Cohen. Making intra-domain routing robust to changing and uncertain traffic demands: understanding fundamental tradeoffs. In *SIGCOMM '03: Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 313–324, New York, NY, USA, 2003. ACM.

[9] Boris S. Axelrod, Benjamin J. Kulick, Charles R. Plott, and Kevin A. Roust. Design improved parimutuel-type information aggregation mechanisms: Inaccuracies and the long-shot bias as disequilibrium phenomena. Social Science Working Paper 1268, 2007.

[10] Yossi Azar, Edith Cohen, Amos Fiat, Haim Kaplan, and Harald Racke. Optimal oblivious routing in polynomial time. In *STOC '03: Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 383–388, New York, NY, USA, 2003. ACM Press.

[11] Walid Ben-Ameur and Hervé Kerivin. Routing of uncertain traffic demands. *Optimization and Engineering*, 6(3):283–313, 2005.

[12] Dimitris Bertsimas and Melvyn Sim. Robust discrete optimization and network flows. *Mathematical Programming*, 98(1):49–71, September 2003.

[13] Dimitris Bertsimas and Aurélie Thiele. A robust optimization approach to supply chain management. In *Integer Programming and Combinatorial Optimization*, pages 145–156. 2004.

[14] Dimitris Bertsimas and Aurélie Thiele. A robust optimization approach to inventory theory. *Operations Research*, 54(1):150–168, January 2006.

[15] Tilman Borgers, Angel Hernando-Veciana, and Daniel Krahmer. When are signals complements or substitutes? July 2007.

[16] Glenn W. Brier. Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78(1):1–3, January 1950.

[17] Archishman Chakraborty and Bilge Yilmaz. Informed manipulation. *Journal of Economic Theory*, 114:132–152, 2004.

[18] Chandra Chekuri. Routing and network design with robustness to changing or uncertain traffic demands. *SIGACT News*, 38(3):106–129, 2007.

[19] Kay-Yut Chen, Leslie R. Fine, and Bernardo A. Huberman. Forecasting uncertain events with small groups. In *Proceedings of the Third ACM Conference on Electronic Commerce (EC'01)*, pages 58–64, Tampa, Florida, 2001.

[20] Yiling Chen, Stanko Dimitrov, Rahul Sami, Daniel Reeves, David M. Pennock, Robin Hanson, Lance Fortnow, and Rica Gonen. Gaming prediction markets: Equilibrium strategies with a market maker. *Algorithmica*, May 2009.

[21] Yiling Chen and David M. Pennock. A utility framework for bounded-loss market makers. *Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence (UAI 2007), Vancouver, BC Canada*, pages 49–56, 2007.

[22] Yiling Chen, Daniel Reeves, David M. Pennock, Robin Hanson, Lance Fortnow, and Rica Gonen. Bluffing and strategic reticence in prediction markets. In *WINE'07, the third Workshop on Internet and Network Economics*, pages 70–81, San Diego, CA, December 2007. Springer LNCS.

[23] Yiling Chen, Daniel M. Reeves, David M. Pennock, Robin D. Hanson, and Rica Gonen. Bluffing and strategic reticence in prediction markets. In *Second Workshop on Prediction Markets EC 2007*, 2007.

[24] Steven Cherry. Bet On It! *Spectrum, IEEE*, 44(9):48–53, 2007.

[25] Millie Chu, Yuriy Zinchenko, Shane G. Henderson, and Michael B. Sharpe. Robust optimization for intensity modulated radiation therapy treatment planning under uncertainty. *Physics in Medicine and Biology*, 50(23):5463–5477, December 2005.

[26] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley-Interscience, August 1991.

[27] Mark E. Crovella and Azer Bestavros. Self-similarity in world wide web traffic: evidence and possible causes. *IEEE/ACM Transactions on Networking*, 5(6):835–846, December 1997.

[28] Stanko Dimitrov and Rahul Sami. Non-myopic strategies in prediction markets. In *Second Workshop on Prediction Markets EC 2007*, 2007.

[29] Stanko Dimitrov and Rahul Sami. Non-myopic strategies in prediction markets. In *EC '08: Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 200–209, 2008.

[30] Stanko Dimitrov, Rahul Sami, and Marina Epelman. Subsidized prediction markets for risk averse traders. In *WINE'09, the fifth Workshop on Internet and Network Economics*, Rome, Italy, December 2009. Springer LNCS.

[31] Phred Dvorak. Best Buy taps 'Prediction Market'; imaginary stocks let workers forecast whether retailer's plans will meet goals. *The Wall Street Journal*, pages B1+, September 2008.

[32] Anders Eriksson. Optimization in target movement simulations. *Computer Methods in Applied Mechanics and Engineering*, 197(49-50):4207–4215, September 2008.

[33] Joan Feigenbaum, Lance Fortnow, David M. Pennock, and Rahul Sami. Computation in a distributed information market. *Theoretical Computer Science*, 343:114–132, 2005.

[34] Sally Floyd and Van Jacobson. Random early detection gateways for congestion avoidance. *IEEE/ACM Transactions on Networking*, 1(4):397–413, 1993.

[35] Robert Fourer, David M. Gay, and Brian W. Kernighan. *AMPL: A Modeling Language for Mathematical Programming*. Duxbury Press, November 2002.

[36] Laurent E. Ghaoui and Hervé Lebret. Robust solutions to least-squares problems with uncertain data. *SIAM Journal on Matrix Analysis and Applications*, 18(4):1035–1064, 1997.

[37] Laurent E. Ghaoui, Francois Oustry, and Hervé Lebret. Robust solutions to uncertain semidefinite programs. *SIAM Journal on Optimization*, 9(1):33–52, 1998.

[38] Philip E. Gill, Walter Murray, and Michael A. Saunders. *User's Guide for SNOPT Version 7: Software for Large-Scale Nonlinear Programming*, June 2008.

[39] Donald Goldfarb and Garud Iyengar. Robust portfolio selection problems. *Mathematics of Operations Research*, 28(1):1–38, February 2003.

[40] Takao Hagishita and Makoto Ohsaki. Topology mining for optimization of framed structures. *Journal of Advanced Mechanical Design, Systems, and Manufacturing*, 2(3):417–428, 2008.

[41] Nils H. Hakansson, J. Gregory Kunkel, and James A. Ohlson. Sufficient and necessary conditions for information to have social value in pure exchange. *The Journal of Finance*, 37(5):1169–1181, 1982.

[42] Robin Hanson. Combinatorial information market design. *Information Systems Frontiers*, 5(1):107–119, January 2003.

[43] Robin Hanson. Logarithmic market scoring rules for modular combinatorial information aggregation. *The Journal of Prediction Markets*, pages 3–15, February 2007.

[44] Robin Hanson, Ryan Oprea, and Dave Porter. Information aggregation and manipulation in an experimental market. *Journal of Economic Behavior and Organization*, 60(4):449–459, 2006.

[45] The Hollywood Stock Exchange.

[46] The Iowa Electronic Markets.

[47] InTrade. http://www.intrade.com.

[48] Joseph B. Kadane and Robert L. Winkler. Separating probability elicitation from utilities. *Journal of the American Statistical Association*, 83(402):357–363, 1988.

[49] David M. Kreps. A note on "fulfilled expectations" equilibria. *Journal of Economic Theory*, 14(1):32–43, February 1977.

[50] Nicolas S. Lambert, John Langford, Jennifer Wortman, Yiling Chen, Daniel Reeves, Yoav Shoham, and David M. Pennock. Self-financed wagering mechanisms for forecasting. In *EC '08: Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 170–179, 2008.

[51] Will E. Leland, Murad S. Taqqu, Walter Willinger, and Daniel V. Wilson. On the self-similar nature of ethernet traffic (extended version). *IEEE/ACM Transactions on Networking*, 2(1):1–15, February 1994.

[52] Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, USA, June 1995.

[53] Declan McCullagh. Tech lessons learned from the wisdom of crowds. http://news.cnet.com/2100-1014_3-6143896.html, December 2006.

[54] Paul Milgrom and Nancy Stokey. Information, trade and common knowledge. *Journal of Economic Theory*, 26(1):17–27, February 1982.

[55] Multiprotocol label switching architecture.

[56] Robert F. Nau. Coherent decision analysis with inseparable probabilities and utilities. *Journal of Risk and Uncertainty*, 10(1):71–91, January 1995.

[57] Robert F. Nau. The aggregation of imprecise probabilities. *Journal of Statistical Planning and Inference*, 105(1):265–282, June 2002.

[58] Evdokia Nikolova and Rahul Sami. A strategic model for information markets. In *EC '07: Proceedings of the 8th ACM Conference on Electronic Commerce*, pages 316–325, 2007.

[59] Theo Offerman, Joep Sonnemans, Gijs van de Kuilen, and Peter P. Wakker. A truth-serum for non-Bayesians: Correcting proper scoring rules for risk attitudes. Working paper: http://www1.fee.uva.nl/creed/pdffiles/propscr30jun08.pdf, June 2008.

[60] Masashi Okamoto. Some inequalities relating to the partial sum of binomial probabilities. *Annals of the Institute of Statistical Mathematics*, 10(1):29–35, March 1959.

[61] OSPF version 2. http://www.ietf.org/rfc/rfc2328.txt.

[62] David M. Pennock. *Aggregating Probabilistic Beliefs: Market Mechanisms and Graphical Representations*. PhD thesis, The University of Michigan, 1999.

[63] David M. Pennock. A dynamic pari-mutuel market for hedging, wagering, and information aggregation. In *EC '04: Proceedings of the 5th ACM Conference on Electronic Commerce*, pages 170–179, 2004.

[64] Matthew Roughan, Mikkel Thorup, and Yin Zhang. Traffic engineering with estimated traffic matrices. In *IMC '03: Proceedings of the 3rd ACM SIGCOMM conference on Internet measurement*, pages 248–258, New York, NY, USA, 2003. ACM.

[65] Mark Rubinstein. Securities market efficiency in an Arrow-Debreu economy. *The American Economic Review*, 65(5):812–824, 1975.

[66] Esther-Mirjam Sent. The tricks of the (no-)trade (theorem). *History of Political Economy*, 38:305–321, January 2006.

[67] Yosef Sheffi. *Urban Transportation Networks: Equilibrium Analysis with Mathematical Programming Methods*. Prentice-Hall, Englewood Cliffs, NJ, 1985.

[68] Rui Z. Shen and Nick Mckeown. Designing a predictable internet backbone network. In *HotNets-III Proceedings*, 2004.

[69] Maiko Shigeno. Maximum network flows with concave gains. *Mathematical Programming*, 107(3):439–459, July 2006.

[70] Shekhar Srivastava, Gaurav Agrawal, Deep Medhi, Michal Pioro, and Michał Pióro. Determining feasible link weight systems under various objectives for OSPF networks. In *IEEE eTransactions on Network and Service Management*, volume 2, 2004.

[71] Koleman S. Strumpf and Paul W. Rhode. Manipulating political stock markets: A field experiment and a century of observational data, 2007.

[72] A. Ben Tal and A. Nemirovski. Robust truss topology design via semidefinite programming. *SIAM Journal on Optimization*, 7(4):991–1016, 1997.

[73] A. Ben Tal and A. Nemirovski. Robust convex optimization. *Mathematics of Operations Research*, 23(4):769–805, 1998.

[74] Jean Tirole. On the possibility of speculation under rational expectations. *Econometrica*, 50(5):1163–1181, 1982.

[75] H. R. Varian. Differences of opinion in financial markets. In *Financial Risk: Theory, Evidence and Implications: Proceedings of the 11th Annual Economic Policy Conference of the Federal Reserve Bank of St. Louis*, pages 3–37, 1989.

[76] Robert L. Winkler. The quantification of judgment: Some methodological suggestions. *Journal of the American Statistical Association*, 62(320):1105–1120, 1967.

[77] Robert L. Winkler. Scoring rules and the evaluation of probability assessors. *Journal of the American Statistical Association*, 64(327):1073–1078, 1969.

[78] Robert L. Winkler and Allan H. Murphy. Nonlinear utility and the probability score. *Journal of Applied Meteorology*, 9:143–148, 1970.

[79] Justin Wolfers and Eric Zitzewitz. Prediction markets. *The Journal of Economic Perspectives*, 18(2):107–126, 2004.