Non-myopic Strategies in Prediction Markets*

Stanko Dimitrov
University of Michigan
Department of Industrial and Operations Engineering
sdimitro@umich.edu

Rahul Sami
University of Michigan
School of Information
rsami@umich.edu

ABSTRACT

One attractive feature of market scoring rules [Hanson, *Information Systems Frontiers*, 2003] is that they are *myopically strategyproof*: It is optimal for a trader to report her true belief about the likelihood of an event provided that we ignore the impact of her report on the profit she might garner from future trades. This does not rule out the possibility that traders may profit by first misleading other traders through dishonest trades and then correcting the errors made by other traders. In this paper, we describe a new approach to analyzing non-myopic strategies and the existence of myopic equilibria. We first use a simple model with two partially informed traders in a single information market to gain insight into the conditions under which different equilibrium behavior emerges. We prove that, under generic conditions, the myopically optimal strategy profile is *not* a weak Perfect Bayesian Equilibrium (PBE) strategy for the logarithmic market scoring rule. We show that our results extend to multiple traders and signals. We propose a simple discounted market scoring rule that reduces the opportunity for bluffing strategies. We show that in any weak PBE, myopic or otherwise, the market price converges to the optimal price, and the rate of convergence can be bounded in terms of the discounting parameter.

Categories and Subject Descriptors

J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*

General Terms

Economics, Theory

*Stanko Dimitrov was supported by NSF grants CCF-0728768 and IGERT-0114368. Rahul Sami was partially supported by NSF grant CCF-0728768. Some results in this paper were presented in preliminary form at the Second Prediction Markets Workshop [9].

EC '08, Jun 8-12, 2008, Chicago, IL, USA.

Keywords

prediction markets, information markets, strategic analysis, market scoring rule

1. INTRODUCTION

It has long been observed that, because market prices are influenced by all the trades taking place, they reflect the combined information of all the traders. Prediction markets (*e.g.*, the Iowa Electronic Market [14]; the Hollywood Stock Exchange [13]; and numerous other sites) are markets designed and deployed specifically to aggregate information about future events; in a prediction market, traders trade in securities whose ultimate value is contingent on the outcome of future events. An informed trader can use her private information to recognize inaccuracies in the current trading price and execute profitable trades. These trades in turn influence the trading price; further, the prices provide signals to other traders about the private information. Other traders learn from these signals and adjust their beliefs about the true value of the security. Ideally, this will lead to a situation in which all traders reach a consensus belief that reflects all available information.

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 [11] is an innovative tradable security; it is based on the idea of a *proper scoring rule* [4]. Pennock's Dynamic Parimutuel Market [17] is another new market design that is based on the traditional parimutuel 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 [11]. A similar (although slightly weaker) characterization of myopically optimal strategies in dynamic parimutuel markets is reported by Nikolova and Sami [16]. 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 exploit 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 paper, 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. In 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 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 certain 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 [18] conducted experiments on manipulating prices in the Iowa Electronic Market. Hanson *et al.* [12] 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 a 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 [1] describe a model in which a manipulative trader can make a deceptive trade in an 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 [1] 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 [5] for references to other research on manipulation in financial markets.

Feigenbaum *et al.* [10] also studied 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 [16] 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.* [2] also proposed a form of discounting in an experimental parimutuel market, and showed that it promoted early trades. Unlike the parimutuel 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.* [7, 6]. Chen *et al.* study also study non-myopic strategies in prediction markets; their initial results [7] were reported at the same time as the preliminary version of our results [9]. 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. Astonishingly, this difference in models leads to opposite results: they show that, under their model, following the myopic strategy is an equilibrium strategy. Further, Chen *et al.* [6] 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.* [3] 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 paper

The rest of this paper is structured as follows: In section 2 we describe the 2-player model we use to highlight deception threats, and we introduce some equilibrium concepts. In section 3, we formally analyze the simple 2-player model and show that there exists no finite equilibrium in this setting. In section 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 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 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.

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 E, by trading in a security F based on E. Two players, P1 and P2, are each endowed with some private information about E. We assume the simplest possible case, in which P1 and P2 each have a single bit of information (x_1, x_2 respectively) relevant to E. (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 E.

The prior probability distribution can then be completely specified by specifying the prior probabilities of the signals and the conditional probability of E given each combination of signals. We assume that the two signals are independent. Further, we assume for simplicity that x_1 is 0 or 1 with equal probability. The probability that x_2 is 1 is given by a parameter $0 < q < 1$. Thus, the model can be fully specified by specifying four probabilities $p_{00} = P\{E|x_1 = 0, x_2 = 0\}$ (or $P\{E|00\}$ for short), $p_{11} = P\{E|11\}$, $p_{01} = P\{E|01\}$ and $p_{10} = P\{E|10\}$. We study the behavior of the market for different values of the parameters $p_{00}, p_{11}, p_{01}, p_{10}$ and q. Note that p_{ij} may be thought of as the probability of E given signals i and j. To summarize:

We assume that the trade in security F is conducted using a market scoring rule [11]. Players make a sequence of market moves; in each move, the player reports a probability p^i. At the end, when the event E is revealed, the move earns a player a net score $s(E, p^i) - s(E, p^{i-1})$, where s is some proper scoring rule. In this paper, we assume the *logarithmic scoring rule*. The market maker seeds the market with a value p^0 that 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.

Table 1: Probability Realizations

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

In a market using the logarithmic scoring rule, the score of any one move is a constant multiple of $\log p^i - \log p^{i-1}$ if E occurs. Without loss of generality we assume that the constant multiple is 1 for our analysis of the market scoring rule. In section 5, we propose a scheme in which the constant multiplier changes over time.

2.1 Myopic behavior

We now analyze the price dynamics if each trader followed her myopically optimal strategy. There are two additional probabilities r_0 and r_1 that arise in the analysis of the myopic behavior because of P1's uncertainty of P2's signal. Suppose P1 saw $x_1 = 1$. She would then condition her prior on this information, resulting in a posterior in which she ascribes probability q to the possibility that the optimal probability is p_{11}, and probability $1 - q$ to the possibility that the optimal probability is p_{10}. In the balance, her belief about the likelihood of E would be in between that implied by p_{11} and that implied by p_{10}. Therefore, her optimal myopic strategy if she observed $x_1 = 1$ would be to report probability $r_1 = qp_{11} + (1 - q)p_{10}$, or simply the expected optimal probability conditioned on her seeing 1 as her signal. Likewise, if P1 saw $x_1 = 0$, she would move to a point r_0 defined in terms of p_{01} and p_{00}.

After P1's move P2 cannot directly see x_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 $r_0 \neq r_1$; this allows us to focus on strategic threats instead of difficulties in extracting signals from the price. Then, P2 can infer the value of x_1; combining this with the value of x_2 that P2 observed, she can calculate the best possible estimate of the conditional probability of E. 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. 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.

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 players' 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 r_1 or r_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 $x_1 = 1$. She could move to r_0 in the first round, instead of her myopically optimal strategy of moving to r_1. Now, if P2 is expecting myopic behavior, she would incorrectly infer that $x_1 = 0$, 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 x_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 s. 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 s is known to P2, who takes it into account and reacts accordingly. It must be that $0 < s < 1$, because otherwise P2 would know x_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 x_2, and assigns some probability s to $x_1 = 1$. The myopic optimal response for P2 *taking into account the probability that P1 is bluffing* can be determined: it is a function of s, x_2 and the position (r_1 or r_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 x_1 and x_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 m player n 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.

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.* [15].

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 paper, 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.

3. ANALYSIS OF THE SIMPLE MODEL

Building on the intuition developed in section 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.

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 1. *An instance of the prediction market with m players satisfies the* general informativeness condition *if there is no vector of signals for any $m-1$ players that makes the m^{th} player's signals reveal no distinguishing information about the optimal probability. Formally, for $m = 2$, 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 $m > 2$, using the notation of section 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 $m-1$ 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 q of seeing a one. As player

1 is playing first, her honest belief, conditioned on her signal, would be given by:

$$\begin{array}{lll} r_1 & = & qp_{11} + (1-q)p_{10} \quad \text{expectation if P1 sees } x_1 = 1 \\ r_0 & = & qp_{01} + (1-q)p_{00} \quad \text{expectation if P1 sees } x_1 = 0 \end{array}$$

The probabilities r_1 and r_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 1. *Consider any equilibrium strategy profile. If player 1 has a deterministic strategy of playing r_u when she receives a* 1 *signal and $r_v \neq r_u$ when she receives a* 0*, then $r_u = r_1$ and $r_v = r_0$.*

PROOF. Assume that $r_u \neq r_1$. Whenever player 1 plays r_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, r_1 would yield a higher profit to player 1 in the first round. Therefore, player 1 has a profitable deviation in expected payoff from r_u to r_1. A similar argument holds for r_v and r_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 r_1 followed by a move from r_1 to r_0. Therefore, when comparing the two strategies the initial move p_s to r_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 r_1 to r_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 [8] as:

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

LEMMA 2. *Assume that Player* 2 *expects player* 1 *to play honestly, and reacts accordingly. If Player 1 observes $x_1 =$ 1 and bluffs, her expected increase in score (over following the myopic strategy) is $qD(p_{11}||p_{01}) + (1-q)D(p_{10}||p_{00}) - D(r_1||r_0)$.*

PROOF. We analyze the change in Player 1's score due to her two moves (deviation and subsequent correction) separately. Given $P1$'s information ($x_1 = 1$), the expected probability of the event happening is r_1. Thus, the expected deviation move score for player 1 is:

$$\begin{array}{lll} S(r_1, r_0) & = & r_1 \ln \frac{r_0}{r_1} + (1-r_1) \ln \frac{1-r_0}{1-r_1} \\ & = & -D(r_1||r_0) \end{array}$$

As player 2 has a probability of q of seeing a one, player 1 will have to have a corrective step from p_{01} to p_{11} with probability q. Similarly with probability $1-q$ player 1 will make her corrective step from p_{00} to p_{10}. Therefore the expected score from the corrective step is $qS(p_{01}, p_{11}) + (1-q)S(p_{00}, p_{10})$.

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

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

□

THEOREM 3. *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 r_u in the first round when she sees $x_1 = 1$, and $P1$ always moves to $r_v \neq r_u$ in the first round when she sees $x_1 = 0$.*

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 1, $P1$ must move to r_1 when $x_1 = 1$ and r_0 when $x_1 = 0$. 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 x_2 and the inferred value of x_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 $x_1 = 1$, Lemma 2 shows that the expected additional score increase if $P1$ bluffed is given by:

$$qD(p_{11}||p_{01}) + (1-q)D(p_{10}||p_{00}) - D(r_1||r_0)$$

Now, from a well-known convexity property of the relative entropy [8, pp.30], we have:

$$\begin{array}{lll} & qD(p_{11}||p_{01}) + (1-q)D(p_{10}||p_{00}) & \geq \\ & \quad D(qp_{11} + (1-q)p_{10}||qp_{01} + (1-q)p_{00}) & (1) \\ \Rightarrow & qD(p_{11}||p_{01}) + (1-q)D(p_{10}||p_{00}) & \geq \\ & \quad D(r_1||r_0) & (2) \\ \Rightarrow & qD(p_{11}||p_{01}) + (1-q)D(p_{10}||p_{00}) - D(r_1||r_0) & \geq \\ & \quad 0 & (3) \end{array}$$

Inequality (2) follows from the definition of r_1 and r_0. Thus, inequality (3) implies that bluffing is always at least as profitable as behaving myopically by lemma 2. Moreover, inequality (3) is strict when $q \neq 0, 1$ and $p_{10} \neq p_{00}, p_{11}$; this follows directly from the log sum inequality [8]. 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.

3.2 Nonexistence of finite equilibrium

Theorem 3 and Lemma 1 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 n of rounds.

THEOREM 4. *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 n rounds, for any finite n.*

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 3, there are two cases that could arise in equilibrium.
Case (i): $P1$ plays some strategy p_u with certainty regardless of the value of x_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 x_1. Now, we claim that there is a position r_u such that P1 moved to r_u with nonzero probability t when $x_1 = 0$ *and* with nonzero probability t' when $x_1 = 1$. Suppose there was no such r_u. Then, the support of $P1$'s first-round strategy for different values of x_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 2 shows that the myopic strategy would be as good as this.

Observe that r_u is played with strictly positive probability. Further, conditioning on $P1$ moving to r_u in the first round, $P2$ would assign some probability $\hat{q} \neq 0, 1$ to $x_1 = 1$. $P1$'s beliefs about x_2 haven't changed at all after the first round. Thus, the conditions of theorem 3 still hold after the first round, conditional on r_u being played.

Repeating this argument for each of the first $n-1$ 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 n rounds. $\square$

4. GENERALIZING THE RESULTS

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

- $\mathcal{M}$ The set of all players.
- $\mathbf{i} \in \{0, n-1\}^{m-1}$ is a vector of the signals for all players *other than* player 1.
- $(j, \mathbf{i}) \in \{0, n-1\}^m$ is a vector of the signals for all players; j denotes player i's signal.
- q_j^k prior probability that player k sees signal j
- $q_\mathbf{i}$ $\prod_{k \in \mathcal{M}/\{1\}} q_{\mathbf{i}_k}^k$ is the probability of players 2 through m 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:

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

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 5. *Player 1 has an expected score increase of* $\sum_\mathbf{i} q_\mathbf{i} D(p_{(j,\mathbf{i})} || p_{(\bar{j},\mathbf{i})}) - D(r_j || r_{\bar{j}})$ *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 r_j to $r_{\bar{j}}$:

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

As each player, k, has a probability of q_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 $q_\mathbf{i}$. Summing over all possible $\mathbf{i}$ we have the expected profit of the corrective step being $\sum_\mathbf{i} q_\mathbf{i} S(p_{(\bar{j},\mathbf{i})}, p_{(j,\mathbf{i})})$

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 $\sum_\mathbf{i} q_\mathbf{i} D(p_{(j,\mathbf{i})} || p_{(\bar{j},\mathbf{i})})$ $\square$

THEOREM 6. *The myopic 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:

$$\begin{aligned} &\sum_\mathbf{i} q_\mathbf{i} D(p_{(j,\mathbf{i})} || p_{(\bar{j},\mathbf{i})}) - D(r_j || r_{\bar{j}}) > 0 \\ \iff &\sum_\mathbf{i} q_\mathbf{i} D(p_{(j,\mathbf{i})} || p_{(\bar{j},\mathbf{i})}) > D(r_j || r_{\bar{j}}) \\ \iff &\sum_\mathbf{i} q_\mathbf{i} D(p_{(j,\mathbf{i})} || p_{(\bar{j},\mathbf{i})}) > D(\sum_\mathbf{i} q_\mathbf{i} p_{(j,\mathbf{i})} || \sum_\mathbf{i} q_\mathbf{i} p_{(\bar{j},\mathbf{i})}) \end{aligned} \quad (4)$$

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

$$D(\sum_i \lambda_i p_i || \sum_i \lambda_i q_i) \leq \sum_i \lambda_i D(p_i || q_i) \quad (5)$$

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 [8, 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) holds strictly.

$\square$

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 6, 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 q_j^1. 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 6 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.

4.2 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. We believe that many, if not all, other scoring rules have an associated convex divergence function. The results above holds for all scoring rules with strictly convex divergence functions. 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.

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.

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 $\delta^{i-1}(\log(1-q) - \log(1-p))$.

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.

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 $p^i(\pi)$ after the last trade in π, and the associated ex-ante probability $P\{\pi\}$ of reaching each such signal node. Now, we define

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

where $p^*(\pi)$ 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 7. *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.

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

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

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.

$$\tilde{M}^i = \begin{cases} \delta^k(\mathcal{D}^{2k} - \mathcal{D}^{2k+1}) & \forall i = 2k \\ \delta^k(\mathcal{D}^{2k+1} - \mathcal{D}^{2(k+1)}) & \forall i = 2k+1 \end{cases} \quad (6)$$

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.

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 q_1 and q_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* (q_1, q_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 (q_1, q_2) determines an instance of the game. Let $\pi^1_M(q_1, q_2)$ 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 $\pi^2_M(q_1, q_2)$ 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 $\mathcal{D}^0(q_1, q_2)$ denote the initial profit potential of this instance: the expected *total* profit in moving from p_s to the optimal point. In other words, $\mathcal{D}^0(q_1, q_2)$ represents the sum of both players' expected profits if they both followed the myopic strategy. Now, define the complementarity coefficient $C(q_1, q_2) = \frac{\pi^1_M(q_1,q_2)+\pi^2_M(q_1,q_2)}{\mathcal{D}^0(q_1,q_2)}$. 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 $\mathcal{D}^0(q_1, q_2)$. If this is greater than the sum of what they could each have earned playing first, *i.e.*, $\pi^1_M(q_1, q_2) + \pi^2_M(q_1, q_2)$, then this indicates that their individual bits of information have increasing marginal value, *i.e.* they are complements. The *lower* $C(q_1, q_2)$ is, the *greater* the complementarity of information. Finally, let the *complementarity bound* C of the configuration be the minimum, over all values of (q_1, q_2), of $C(q_1, q_2)$.

Under our independence assumption $C(q_1, q_2)$ 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 (q_1, q_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 8[1] we note that $\pi^2_M|\sigma_1 \geq \pi^2_M$.

THEOREM 8. $\pi^2_M|\sigma_1 \geq \pi^2_M$

(where, for clarity, we have suppressed the dependence on (q_1, q_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 8, 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 $\delta^k\mathcal{D}^{2k}$:

$$\begin{aligned} \mathcal{D}^{2k} &\leq \mathcal{D}^0 - [\pi^1_M + \pi^2_M - \delta^k\mathcal{D}^{2k}] \quad (7) \\ \iff (1-\delta^k)\mathcal{D}^{2k} &\leq \mathcal{D}^0 - [\pi^1_M + \pi^2_M] \\ &= \mathcal{D}^0\left(1 - \frac{\pi^1_M + \pi^2_M}{\mathcal{D}^0}\right) \\ \iff \mathcal{D}^{2k} &\leq \frac{\mathcal{D}^0\,(1 - C(q_1, q_2))}{1-\delta^k} \quad (8) \end{aligned}$$

[1]See Appendix for proof

Note that for any configuration, by definition, $C(q_1, q_2) \geq C$. Thus, we can rewrite inequality (8) as:

$$\mathcal{D}^{2k} \leq \frac{\mathcal{D}^0(1-C)}{1-\delta^k} \tag{9}$$

From inequality (9) 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 (9) 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:

$$\begin{aligned}\mathcal{D}^n &\leq \mathcal{D}^0\left(\tfrac{1-C}{1-\delta^k}\right)^m \\ &= \mathcal{D}^0\left(\tfrac{1-C}{(1-\delta^{k/2})(1+\delta^{k/2})}\right)^m \\ &= \mathcal{D}^0(1+\delta^{k/2})^{-m} = \mathcal{D}^0(1+C)^{-\frac{n}{2k}} \\ &= \mathcal{D}^0(1+C)^{-\frac{n\log\delta}{4\log C}} = \mathcal{D}^0\delta^{n\frac{-\log(1+C)}{4\log C}}\end{aligned} \tag{10}$$

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 (10) 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.

6. DISCUSSION AND FUTURE WORK

In this paper, we analyze non-myopic strategies in a two-player prediction market setting. We find, surprisingly, 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.* [6] 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.

Acknowledgment

We are grateful to Yiling Chen, P.J. Healy, and David Pennock for helpful discussions. We would like to thank John Ledyard for pointing us to earlier work on discounting in parimutuel markets, and the anonymous reviewers for many helpful suggestions that improved this paper.

7. REFERENCES

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

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

[3] T. Borgers, A. Hernando-Veciana, and D. Krahmer. When are signals complements or substitutes? Working paper. Available from `http://www-personal.umich.edu/~tborgers/BHK.pdf`, July 2007.

[4] G. Brier. Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78:1–3, 1950.

[5] A. Chakraborty and B. Yilmaz. Informed manipulation. *Journal of Economic Theory*, 114:132–152, 2004.

[6] Y. Chen, D. M. Reeves, D. M. Pennock, R. D. Hanson, and L. Fortnow. Bluffing and strategic reticence in prediction markets. In *Workshop on Internet and Network Economics*, 2007.

[7] Y. Chen, D. M. Reeves, D. M. Pennock, R. D. Hanson, and R. Gonen. Bluffing and strategic reticence in prediction markets. In *Second Workshop on Prediction Markets (held at EC '07)*, 2007.

[8] T. M. Cover and J. A. Thomas. *Elements of information theory*. Wiley-Interscience, New York, NY, USA, 1991.

[9] S. Dimitrov and R. Sami. Non-myopic strategies in prediction markets. In *Second Workshop on Prediction Markets (held at EC '07)*, 2007.

[10] J. Feigenbaum, L. Fortnow, D. M. Pennock, and R. Sami. Computation in a distributed information market. *Theoretical Computer Science*, 343:114–132, 2005.

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

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

[13] The Hollywood Stock Exchange. `http://www.hsx.com/`.

[14] The Iowa Electronic Markets. `http://www.biz.uiowa.edu/iem/`.

[15] A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, June 1995.

[16] E. Nikolova and R. Sami. A strategic model for information markets. In *Proceedings of the Seventh Annual ACM Conference on Electronic Commerce (EC '07)*, 2007.

[17] D. Pennock. A dynamic parimutuel market for information aggregation. In *Proceedings of the Fourth Annual ACM Conference on Electronic Commerce (EC '04)*, June 2004.

[18] K. Strumpf and P. Rhode. Manipulating political stock markets: A field experiment and a century of observational data, 2007. Working paper. Available from `http://www.unc.edu/~cigar/papers/ManipHIT_Jan2007.pdf`.

APPENDIX

Proof of theorem 8

THEOREM 8. $\pi_M^2|\sigma_1 \geq \pi_M^2$

PROOF. Consider any strategy σ for player 1 such that setting $\sigma_1 = \sigma$ minimizes $\pi_M^2|\sigma_1$. 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 $q_1 = 0.5$, *i.e.*, that player 1 has equal chance of seeing either a 1 or a 0. Any other value of q_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 $\alpha = P\{x_1 = 1|A\} = \frac{u_A}{p_A}$ and $\beta = P\{x_1 = 1|B\} = \frac{u_B}{p_B}$. With these definitions we can define the myopic move of player 2 given that market is at A and $x_2 = 1$ as $r_1^\alpha = \alpha p_{11} + (1-\alpha)p_{01}$; likewise, if $x_2 = 0$ as $r_0^\alpha = \alpha p_{10} + (1-\alpha)p_{00}$. Similarly r_0^β and r_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 $\gamma = P\{x_1 = 1|C\} = \frac{\frac{p_B}{p_A}u_A + u_B}{2p_B} = \frac{1}{2}\frac{u_A}{p_A} + \frac{1}{2}\frac{u_B}{p_b} = \frac{\alpha+\beta}{2}$. We can now define the myopic move of player 2 given the market is at C and $x_2 = 1$ as $r_1^\gamma = \gamma p_{11} + (1-\gamma)p_{01}$. Likewise if $x_2 = 0$ and the market is at C the myopic move of player 2 is defined as $r_0^\gamma = \gamma p_{10} + (1-\gamma)p_{00}$.

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

$$\begin{aligned}
\pi_M^2|\sigma \;=\; & 0.5(p_A - p_B)[qD(r_1^\alpha||qr_1^\alpha + (1-q)r_0^\alpha) \\
& +(1-q)D(r_0^\alpha||qr_1^\alpha + (1-q)r_0^\alpha)] \\
& +0.5p_B[qD(r_1^\alpha||qr_1^\alpha + (1-q)r_0^\alpha) \\
& +(1-q)D(r_0^\alpha||qr_1^\alpha + (1-q)r_0^\alpha)] \\
& +0.5p_B[qD(r_1^\beta||qr_1^\beta + (1-q)r_0^\beta) \\
& +(1-q)D(r_0^\beta||qr_1^\beta + (1-q)r_0^\beta)] \\
& + \text{remaining profit over } R
\end{aligned}$$

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

$$\begin{aligned}
\pi_M^2|\sigma' \;=\; & 0.5(p_A - p_B)[qD(r_1^\alpha||qr_1^\alpha + (1-q)r_0^\alpha) \\
& +(1-q)D(r_0^\alpha||qr_1^\alpha + (1-q)r_0^\alpha)] \\
& +0.5\cdot 2p_B[qD(r_1^\gamma||qr_1^\gamma + (1-q)r_0^\gamma) \\
& +(1-q)D(r_0^\gamma||qr_1^\gamma + (1-q)r_0^\gamma)] \\
& + \text{remaining profit over } R
\end{aligned}$$

From the definitions of the myopic moves given the market states, note that $r_1^\gamma = \frac{r_1^\alpha + r_1^\beta}{2}$ and $r_0^\gamma = \frac{r_0^\alpha + r_0^\beta}{2}$. This means that $\pi_M^2|\sigma'$ can be bounded as :

$$\begin{aligned}
\pi_M^2|\sigma' \;=\; & 0.5(p_A - p_B)[qD(r_1^\alpha||qr_1^\alpha + (1-q)r_0^\alpha) \\
& +(1-q)D(r_0^\alpha||qr_1^\alpha + (1-q)r_0^\alpha)] \\
& +0.5\cdot 2p_B[qD(r_1^\gamma||qr_1^\gamma + (1-q)r_0^\gamma) \\
& +(1-q)D(r_0^\gamma||qr_1^\gamma + (1-q)r_0^\gamma)] \\
& + \text{remaining profit over } R \\
=\; & 0.5(p_A - p_B)[qD(r_1^\alpha||qr_1^\alpha + (1-q)r_0^\alpha) \\
& +(1-q)D(r_0^\alpha||qr_1^\alpha + (1-q)r_0^\alpha)] \\
& +0.5\cdot 2p_B[qD(\tfrac{r_1^\alpha + r_1^\beta}{2}||q\tfrac{r_1^\alpha + r_1^\beta}{2} + (1-q)\tfrac{r_0^\alpha + r_0^\beta}{2}) \\
& +(1-q)D(\tfrac{r_0^\alpha + r_0^\beta}{2}||q\tfrac{r_1^\alpha + r_1^\beta}{2} + (1-q)\tfrac{r_0^\alpha + r_0^\beta}{2})] \\
& + \text{remaining profit over } R \\
<\; & 0.5(p_A - p_B)[qD(r_1^\alpha||qr_1^\alpha + (1-q)r_0^\alpha) \\
& +(1-q)D(r_0^\alpha||qr_1^\alpha + (1-q)r_0^\alpha)] \\
& +0.5p_B[qD(r_1^\alpha||qr_1^\alpha + (1-q)r_0^\alpha) \\
& +(1-q)D(r_0^\alpha||qr_1^\alpha + (1-q)r_0^\alpha)] \\
& +0.5p_B[qD(r_1^\beta||qr_1^\beta + (1-q)r_0^\beta) \\
& +(1-q)D(r_0^\beta||qr_1^\beta + (1-q)r_0^\beta)] \\
& + \text{remaining profit over } R \\
=\; & \pi_M^2|\sigma
\end{aligned}$$

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 $\pi_M^2|\sigma$, 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 $\pi_M^2|\sigma \geq \pi_M^2$. $\square$