Non-myopic Strategies in Prediction Markets

Stanko Dimitrov
Department of Industrial and Operations Engineering
University of Michigan, 1205 Beal Avenue, Ann Arbor, MI 48109-2117, USA
sdimitro@umich.edu

Rahul Sami
School of Information
University of Michigan, 1075 Beal Avenue, Ann Arbor, MI 48109-2112, USA
rsami@umich.edu

May 18, 2007

Abstract

One attractive feature of market scoring rules [Hanson '03] 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. Such non-myopic strategies are difficult to analyze, and the empirical results are inconclusive. In this paper, we describe a new approach to analyzing non-myopic strategies and the existence of myopic equilibria. We 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 is not a sequential equilibrium strategy for the logarithmic market scoring rule. We then illustrate how introducing a form of discounting into the prediction market can speed up information revelation by non-myopic traders.

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 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 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 [Han03] is an innovative tradable security; it is based on the idea of a *proper scoring rule* [Bri50]. Pennock's Dynamic Parimutuel Market [Pen04] 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 at least myopically optimal [Han03]. A similar (although slightly weaker) characterization of myopically optimal strategies in dynamic parimutuel markets is reported in [NS07]. However, much of the concern about manipulation in prediction markets is based on non-myopic strategies; as yet, very little is known theoretically about the existence and characterization of manipulative non-myopic strategies in these markets.

In this paper, we describe a new approach we are undertaking to analyzing these non-myopic strategies. We motivate our ideas with numerical exploration of the projection game [NS07], which is a simple analytical model that is strategically equivalent to the market scoring rule with a spherical scoring rule. Our analytical results are derived directly using the logarithmic market scoring rule.

Related Work There have been several field and experimental studies of manipulation in prediction markets. Strumpf and Rhode [SR07] conducted experiments on manipulating prices in the Iowa Electronic Market. Hanson et al [HOP06] 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.

We are not aware of prior theoretical analysis of manipulation in prediction markets specifically. However, 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 [AG92] 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 [AG92] 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 rule. Our model requires only risk-neutral traders in a prediction-market context. We refer readers to the paper by Chakraborty and Yilmaz [CY04] for references to other research on manipulation in financial markets.

Feigenbaum et al [FFPS05] 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.

Based on the intuition provided by the projection game, Nikolova and Sami [NS07] present an instance in which myopic strategies are not optimal for both players, 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.

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. In section 2.1, we analyze this model to derive conditions under which the myopically optimal truthful strategy is optimal even in a non-myopic sense. We explain how this analysis can be recursively extended to characterize finite sequences of moves that are in equilibrium and result in ultimate revelation of true information. In section 3, we present preliminary numerical results that suggest that three cases can arise in a market with discounting: (1) The myopic strategies form an equilibrium and information is rapidly incorporated into the market; (2) The myopic strategies do not form an equilibrium, but there is an equilibrium with bluffing in initial rounds followed by myopic behavior; and, (3) There is no equilibrium profile in which information is incorporated in a finite number of rounds.

Motivated by the numerical results, we present our main results in section 4: We formally prove that, for an undiscounted log-scoring rule market with two players in a Bayesian setting, there is *generically* no myopic equilibrium; in fact, there is generically no equilibrium in which all information is incorporated in a finite amount of time. Our technique, building on results from classic information theory, could be very useful for other aspects of information market analysis.

In section 5 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 describe some future steps.

2 A Simple Model

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 $p_A = p(E/x_1 = 0, x_2 = 0)$ (or $p(E/00)$ for short), $p_B = p(E/11)$, $p_C = p(E/01)$ and $p_D = p(E/10)$. We study the behavior of the market for different values of the parameters p_A, p_B, p_C, p_D and q.

We assume that the trade in security F is conducted using a market scoring rule [Han03]. 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. The market maker seeds the market with a value p^0 that is irrelevant for our analysis. Here, we primarily consider the *spherical scoring rule* function for s. We consider a sequence of alternating moves in which P1 moves first, P2 moves next, P1 potentially moves again, and so on.

Following [NS07], trade in this market (with a spherical scoring rule) can be visualized as play in a restricted *projection game*, described as follows: The game is played on a 2-dimensional board, with two axes corresponding to the two possible outcomes E and $\overline{E}$. The game position at any point of time is a point on the positive quadrant of the unit circle. Players take turns in moving from the current position to another point on the circle. A position $(\cos\theta, \sin\theta)$ corresponds to a probability assessment p such that $\frac{p}{1-p} = \tan\theta$. The profit of each move $(\cos\theta_1, \sin\theta_1) \to (\cos\theta_2, \sin\theta_2)$, conditional on a "true" probability $p*$, is determined as follows: It is the projection of the directed segment from the starting to the final point onto the line through the origin and the point corresponding to the true probability $p*$. Thus, moves towards the true probability have a positive profit, and moves away from the true probability have a negative profit.

We stress that we are using the projection game merely as a way to visualize the expected profit of different moves in a market scoring rule, in order to compare different strategies. The projection game cannot be directly implemented (because the true probabilities are unknown), but the results of [NS07] show that strategies in this game are equivalent to strategies in a market with the spherical scoring rule, and perhaps a slightly different prior distribution. Thus, although we describe myopic and bluffing strategies in the context of the projection game, there is a $1-1$ correspondence with strategies in an MSR.[1]

In the projection game representation of the prediction market, the possible signal realizations lead to the following points on the circle:

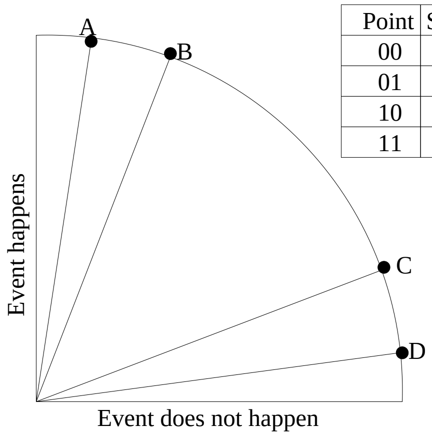

Point	Signals
00	A
01	C
10	D
11	B

Figure 1: Point realizations, a generalization of [NS07]

One can also think of each of these points as a corresponding angle the line through the origin and the point make with the x-axis. For example point C corresponds with angle θ_c, point A corresponds to angle θ_a, etc. Ideally, we would like the prediction market to converge to the point corresponding to the true signals; this would give the best possible estimate of the probability of E.

2.1 Myopic behavior

We now analyze the price dynamics if each trader followed her myopically optimal strategy. There are two additional points X and Y 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 point is B, and probability $1-q$ to the possibility that the optimal point is D. In the balance, her belief about the likelihood of E would be in between that implied by B and that implied by D. It is easy to show that (in the projection game) her optimal myopic strategy is to move to a point X defined by $\theta_x = \arctan(q\sin\theta_d + (1-q)\sin\theta_b)/(q\cos\theta_d + (1-q)\cos\theta_b)$. Likewise, if P1 saw $x_1 = 0$, she would move to a point Y defined in terms of θ_a and θ_c.

Now, P2 cannot directly see x_1, but he 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 $X \neq Y$; 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, he can calculate the best possible estimate of the conditional probability of E. Due to the myopic strategyproofness of the projection game (and market scoring rules),

[1] As a sanity check, we also calculated the numerical results in section 3 in an MSR with a logarithmic scoring rule. The results are qualitatively similar; if anything, they show a stronger tendency to manipulation.

P2 would move to A,B,C,orD respectively. 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 X or Y 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 a 1. She could move to Y in the first round, instead of her myopically optimal strategy of moving to X. Now, if P2 is expecting myopic behavior, she would incorrectly infer that $x_1 = 0$, and correspondingly move to the wrong point: A instead of D, or C instead of B. P1 can see where P2 moved to, and make a subsequent correcting move: $A \to D$ or $C \to B$ respectively. P1's incentive to bluff is thus determined by the profitability of this bluffing strategy relative to the honest (myopic) strategy. Note that 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 r. Note that P2 can analyze P1's profit in different scenarios, and thus, can infer that P1 would not necessarily be truthful. Now, we can try to equilibria in which the bluffing probability r is known to P2, who takes it into account and reacts accordingly. It must be that $0 < r < 1$, 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: He sees x_2, and assigns some probability r 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 r, x_2 and the position (X or Y) that P1 left the market at. Next, we can repeat the analysis from P2's point of view, and determine if the myopic response is optimal for P2, or if he 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 labels C and D are interchanged.

With many instances[2] of point locations A,B,C, and D, we can show that for all q such that $0 < q < 1$, the myopic strategy is not optimal; further, this holds even when C and D are interchanged. The reasoning above

[2] In fact, we conjecture that this is true for all non-degenerate locations of the points A,B, C, D; however, we have not been able to prove this yet.

leads to a startling consequence that *no finite sequence of moves by the two players leading to the optimal values (where the ith move involves bluffing with some probability* $r_i, 0 \le r_i \le 1$) *can be an equilibrium!* This follows by a backward induction argument: If there is an equilibrium strategy that leads to the exact optimal value, the last move in that strategy should reflect both x_1 and x_2. Suppose, without loss of generality, that the last (*n*th) move was by P1; it follows that P1 must know x_2 with certainty at the time of her last move, and thus, the second-last move must reveal x_2 with certainty. By assumption, if P2 had any uncertainty about x_1 when he made this move, it would be profitable for him to bluff with some probability $0 < r_{n-1} < 1$ and correct the market after P1 made the *n*th move. It follows that, for this strategy to be an equilibrium, P2 must know x_1 with certainty before the $(n-1)$th move. Repeating this argument, we reach a contradiction, and therefore, no such equilibrium strategy exists.

3 Numerical Exploration

In this section we describe our approach to numerically analyzing the profitability of addressing non-myopic strategies. The results in this section motivate our theoretical results in section 4, and give insight into the required discounting factors.

We analyze the incentive to bluff in two ways: (1) The relative profit from bluffing is influenced by the relative position of the points A, B, C, D. We conduct a numerical analysis of the profit for specific positions of four points, and all permutations of the four labels A,B,C,D over these positions, with different values of q. We use this to qualitatively infer how the propensity to bluff changes in different situations. Our results show many situations in which the myopic strategy, or any finite strategy profile terminating in the optimal position, cannot be in equilibrium. In fact, they suggest that this is *always* true. We then consider a very natural solution to this problem, which is to introduce an explicit form of time-discounting into the market rules. The numerical results show that even a small discount factor could stabilize the myopic equilibrium.[3]

(2) We can analytically study the problem for a special case of Boolean functions: The case in which p_A, p_B, p_C, p_D are each either 0 or 1, so that the event under consideration is a deterministic function of the bits x_1, x_2. Such functions were studied by Feigenbaum et al [FFPS05]; in their model, which did not include the possibility of manipulation, they found that the market would converge rapidly for certain Boolean functions, such as the AND function. Here, we show that bluffing is profitable even in this simple case of the AND function, and hence, the market convergence to the true price is not guaranteed to be rapid if agents are strategic in a non-myopic way.

[3]It could be argued that there is always an implicit discount factor influencing traders, due to impatience or uncertainty over whether the game will be terminated abruptly, as is standard in bargaining models. However, this is less natural in information market settings: players do not obtain their profits until the market cashes out, so there is little justification for impatience, and automation of trading might lead to long sequences of precisely timed moves in a short amount of time. For this reason, we prefer to introduce an explicit discount factor. Controlling the magnitude of the discounting could also control the speed of convergence of the market.

3.1 Expressing bluffing profits

As noted earlier it may not be always advantageous for the market players to act myopically. In particular if the expected score or profit from bluffing in the market is greater than the expected score of acting myopically then a rational player will bluff. Equivalently if the difference in bluffing and acting myopically is positive then a rational player will bluff.

Therefore, during the first move in the market the first player has to decide if she will bluff or will not bluff. Without loss of generality, we assume that player P1 sees $x_1 = 1$; thus, the myopically optimal position to move to is X, whereas a move to Y corresponds to a bluffing move.

P1 has to decide whether to move to X or Y in the first round. If P1 moves to Y in the first round, thus bluffing, she will play a corrective move in the third round after P2 has moved. The corrective move be either from A to D (if P2 revealed that $x_2 = 0$) or from C to B (if P2 revealed that $x_2 = 1$). We want to analyze the *difference* between the profits P1 earns from the truthful and bluffing strategies.

We analyze the first-round and third-round profit differences separately, and then combine them. Due to the path-independence property of market scoring rules, the expected profit difference between the truthful and bluffing first-round moves is equal to the expected profit of a move from X to Y. (This will be negative, as X is the optimal point given P1's beliefs in the first round). In the bluffing strategy, P1 makes an additional corrective move, whereas in the myopic strategy she does not. The additional corrective move always yields a positive profit to P1, as it ends at the optimal point given the true information. Thus, we can compare the relative profit of myopic and bluffing strategies by studying when the expected combined profit of the move from X to Y and the corrective move is positive.

3.2 Numerical calculation of profitability of bluffing

Here we discuss some of the numerical analysis applied to non-myopic strategies in information markets. Following our discussion that myopic strategies do not always lead to a market equilibrium, we first present one solution, namely a discount factor, to this issue we then present a numerical analysis for a set of games.

First we try to characterize under what conditions player P1 will bluff. She will bluff if she loses less from the initial misleading step than she gains, in expectation, from the final corrective step. We can look at the negative of the ratio, m, of the expected cost of misleading to the expected payoff of the corrective step to see if player P1 has an incentive to bluff during her first move. Given the definition of m, if $m < 1$ player P1 has at incentive to bluff during her first move, if $m = 1$ player P1 is indifferent to bluffing, and if $m > 1$ player P1 has no incentive to bluff initially. Given the position of the points A, B, C, D, and the probability q that $x_2 = 0$, we can calculate the ratio m.

3.2.1 Discounted Markets and Discount Factors

The ratio m has another, insightful interpretation. We consider modifying the rules of trading in the market to introduce *discounting* over time, as follows: Define a factor $\delta, 0 \leq \delta \leq 1$. When cashing out the market, the profit or loss made by player P1 in her ith move is scaled by δ^{i-1}. In each individual move, a player's profit from a move is proportional to her profit from the same move in an unscaled market, and therefore, the myopic properties of the prediction market are not altered. However, the bluffing strategies described above are altered, because the profit from the corrective move is scaled by δ, while the bluff move is not scaled. It is easy to observe that in any situation with ratio $m < 1$, m is exactly the maximum value of δ for which the myopic strategy is at least as good as the bluffing strategy for P1, assuming that P2 will play myopically. Thus, we plot the value of m for different situations, as q is varied.

To make this notion more concrete, assume that player 1 observes zero as her signal, we are assuming that player 1 is equally likely to observe either signal. For the notation below θ_i, $i = \{a,b,c,d\}$ corresponds to the point associated with the signals in Figure 1. Therefore in expectation her net profit for the bluffing move is:

$$\begin{aligned} E\{NP[\text{bluff move}]\} &= [q((\cos\theta_x - \cos\theta_y)\cos\theta_c + (\sin\theta_x - \sin\theta_y)\sin\theta_c) + \\ &\quad (1-q)((\cos\theta_x - \cos\theta_y)\cos\theta_a + (\sin\theta_x - \sin\theta_y)\sin\theta_a)] \end{aligned}$$

Similarly player 1 has an expected net profit from the corrective move of:

$$\begin{aligned} E\{NP[\text{corrective move}]\} &= \delta[q(1-\cos(\theta_b - \theta_c)) + \\ &\quad (1-q)(1-\cos(\theta_d - \theta_a))] \end{aligned}$$

Let θ_x, θ_y be the myopically optimal positions.

As noted above player 1 will behave myopically if:

$$E\{NP[\text{corrective move}]\} + E\{NP[\text{bluff move}]\} \leq 0$$

Let $\Delta :=$ be the expected loss due to bluffing.

Define the ratio $m = \left(-\frac{\Delta}{1-q\cos(\theta_b-\theta_c)-(1-q)\cos(\theta_d-\theta_a)}\right)$. If m is non-negative for all realizations of $\theta_a, \theta_b, \theta_c, \theta_d$, and q then there is always a δ such that player 1 will behave myopically in the first move.

Claim: $m \geq 0 \quad \forall q, \theta_i$.

Proof:

1. $1 - q\cos(\theta_b - \theta_c) - (1-q)\cos(\theta_d - \theta_a) \geq 0$ Note that $\cos(\theta) \leq 1$. As $q \in [0,1]$ it means that $q\cos(\theta_b - \theta_c) - (1-q)\cos(\theta_d - \theta_a) \leq 1$ giving the conclusion.

2. $\Delta < 0$ this follows from the fact that both of these moves are deviations away from the true p-line. Therefore the dot products must be both negative.

The above imply that $m \geq 0$. In the case that $1 - q\cos(\theta_b - \theta_c) - (1-q)\cos(\theta_d - \theta_a) = 0$ we have a case of an unbounded m assuming that we are not the degenerate case of the first player's move not revealing any information, i.e. $\theta_x = \theta_y$.

As we have just show that $m \geq 0$ for all $\theta_a, \theta_b, \theta_c, \theta_d$, and q we know that for any realization of the problem we can find a discount factor such that the expected net profit of bluffing is less than the expected net profit of acting myopically.

3.2.2 How much to Discount

Below we discuss a three scenarios, a nominal projection game scenario and the AND game. The discussion will be motivated through an numerical plots of the different scenarios.

Each of the plots below will have a title with a permutation of the letters ABCD. Assume that $a > b > c > d$, the the title corresponds to the assignment of the angles θ_a through θ_d to the values a through d. For example the title of ABCD implies that $\theta_a = a, \theta_b = b, \theta_c = c, \theta_d = d$ and ACDB implies that $\theta_a = a, \theta_b = d, \theta_c = b, \theta_d = c$. Without loss of generality, we can assume that A is the point with maximum θ_a; if it isn't, this can always be engineered by relabeling the signal values for x_1 and/or x_2. Thus, we look at configurations with A fixed.

We now consider a nominal instance of the projection game. In particular let $a = 75°, b = 60°, c = 30°$, and $d = 15°$. Looking at the graphs in Figure 2 we note that $m < 1$ for all values of $q \neq 0$ or 1. What is more surprising is that in some scenarios, in particular ACDB, with a discount factor of approximately 0.99 player 1 will behave myopically regardless of what probability player 2 observes her signal. Moreover in some of the other scenarios, ABDC for example, a discount factor of 0.2 make player 1 behave myopically for a large range of q values, $[0, 0.15]$ and $[0.85, 1]$.

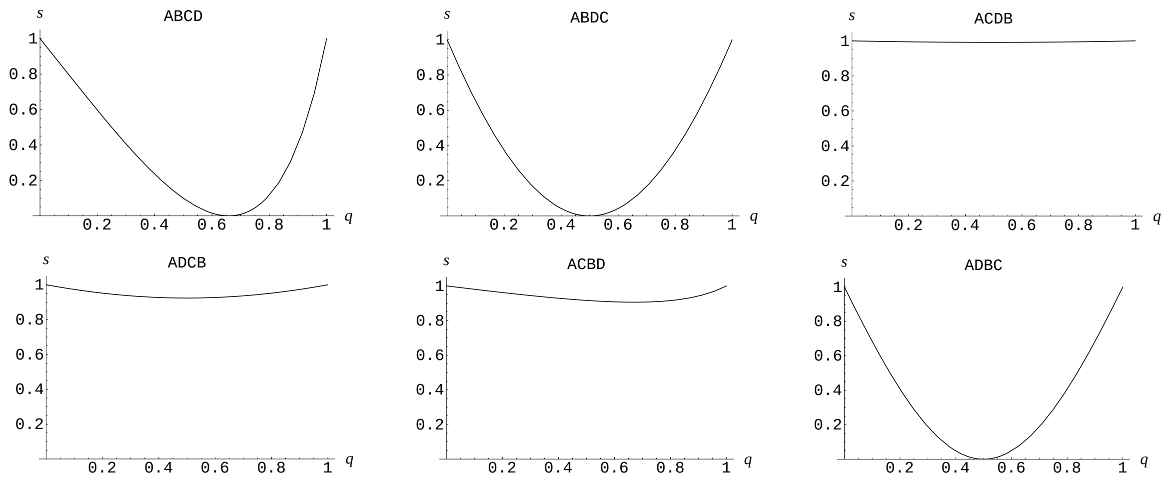

Figure 2: m vs q for the Nominal projection game with $a = 75°, b = 60°, c = 30°$, and $d = 15°$

Immediately, we can see that there are cases in which, in the absence of discounting, there is no finite equilibrium strategy in which the market converges to the correct position. For example, consider the market defined by configuration ABCD, with player 1 playing first. We can see that $m < 1$ whenever q is between 0 and 1, so player 1 would bluff with some probability $0 < r < 1$. Next, from player 2's point of view, the market would be in configuration ABDC. Again, we see that $m < 1$ whenever P2 has some suspicion that P1 bluffed in the previous round. Thus, P2 would also bluff with some probability. Following the reasoning in section 2, no finite revealing equilibrium strategy profile is possible. This result appears to hold generically, for all positions of the four points and all values of q between 0 and 1.

With even slight discounting, we can identify situations in which the myopic equilibrium is possible: for example, the ACDB configuration with a discount factor of 0.99. Further, with some discount factors δ, we can construct examples in which there is an equilibrium in which P1 bluffs with some probability r, but P2 plays myopically even knowing this, and thus the market converges in 3 rounds (after P1's corrective move). For example, consider the market in configuration ADBC, with $q = 0.4$ and $\delta = 0.8$. Inspection of the graphs shows that P1 still benefits from some bluffing in the first round. However, from P2's point of view, the market configuration will be modeled by the ACBD configuration, and thus, bluffing will not be profitable with $\delta = 0.8$.

AND function : Next, we consider a special case, the AND game, defined as follows: The security payoff is the logical AND of the two players' signals. We show that a discount factor of one is adequate only when q is either zero or one. Below in Figure 3 we see the plots of m vs. q for the AND game. To model the AND game we set $a = d = c = 0°$ and $b = 90°$ where the roles of a, b, c, d are described above. As the plots in Figure 3 describe, a player in the AND game will always bluff as long as $q \neq 0$ or 1.

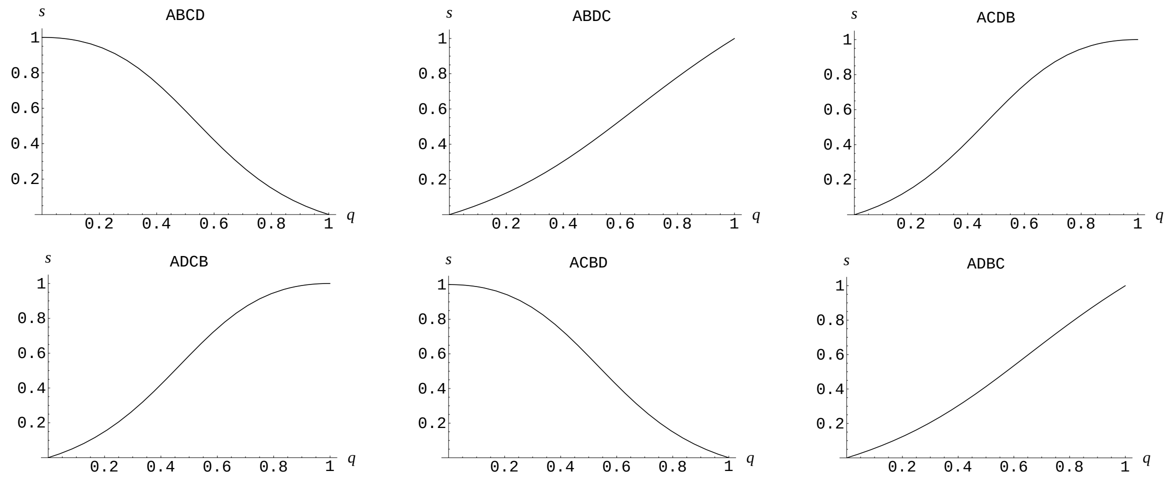

Figure 3: m vs q for the AND game

4 Analysis of the logarithmic MSR

Building on the intuition developed in section 3, we now consider an analytical proof to show that bluffing is always profitable. 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 (including the spherical scoring rule, and thus, the segment game.)

4.1 Generic Bluffing

Consider an order of play with $P1$ playing before $P2$. As player 1 is playing first we define the following:

$$\begin{array}{rcll} p_x & = & qp_b + (1-q)p_d & \textit{P1's expectation if she sees } x_1 = 1 \\ p_y & = & qp_c + (1-q)p_a & \textit{P1's expectation if she sees } x_1 = 0 \end{array}$$

The probabilities p_x and p_y determine the optimal myopic moves for player 1.

Lemma 1. *Consider any equilibrium strategy profile. If player 1 has a deterministic strategy of playing p_u when she sees $x_1 = 1$ and $p_v \neq p_u$ when she sees $x_1 = 0$, then $p_u = p_x$ and $p_v = p_y$.*

Proof. Assume that $p_u \neq p_x$. Whenever player 1 plays p_u player 2 will deduce that player 1 has observed a 1; then, player 2 will capture all the remaining surplus. Player 1 thus gets at most the profit she derives from her first move. Therefore, by definition of p_x, player 1 has a positive deviation, in expected payoff, from p_u to p_x. A similar argument holds for p_v and p_y. □

Assume that the market starts at an arbitrary point p_s as player 1 is moving first, when she moves first we can think of her always moving to p_x and if she decided to bluff she will then move to p_y. Therefore when comparing the two strategies the initial move p_s to p_y cancels out therefore we can assume that the market starts at p_x and therefore the score of the initial move is zero.

We now express the expected profits in terms of information-theoretic entropy. While the connection between proper scoring rules and entropy is well known, it has, to the best of our knowledge, not been used in the prediction markets literature. To this end, we introduce the following notation:

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

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

$D(p_i || p_j)$ The **relative entropy** of two probability mass functions $p(x)$ and $q(x)$ is defined in [CT91] 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 $qD(p_b||p_c) + (1-q)D(p_d||p_a) - D(p_x||p_y)$.*

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 p_x. Thus, the expected deviation move score for player 1 is:

$$\begin{array}{rcl} S(p_x, p_y) & = & p_x \log \frac{p_y}{p_x} + (1-p_x) \log \frac{1-p_y}{1-p_x} \\ & = & -D(p_x||p_y) \end{array}$$

As player 2 has a probability of q of seeing a one, player 1 will have to have a corrective step from p_c to p_b with probability q. Similarly with probability $1-q$ player 1 will make her corrective step from p_a to p_d. Therefore the expected score from the corrective step is $qS(p_c, p_b) + (1-q)S(p_a, p_d)$.

$$\begin{aligned} S(p_c, p_b) &= p_b \log \tfrac{p_b}{p_c} + (1-p_b) \log \tfrac{1-p_b}{1-p_c} \\ &= D(p_b || p_c) \end{aligned}$$

$$\begin{aligned} S(p_a, p_d) &= p_d \log \tfrac{p_d}{p_a} + (1-p_d) \log \tfrac{1-p_d}{1-p_a} \\ &= D(p_d || p_a) \end{aligned}$$

□

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 following conditions hold at the start of trading:*

C1*: Player 1 ascribes some belief $q \neq 0, 1$ to player 2 having seen a 1*

C2*: Player 2 ascribes some belief $q' \neq 0, 1$ to player 1 having seen a 0*

C3*: The optimal probabilities p_a, p_b, p_c, p_d are in general position, i.e., no two are equal.*

Further, restrict attention to equilibria in which P1's first round move is not deterministic strategy that does not depend on x_1. (In other words, we ignore equilibria in which P1 always moves to some value of p, regardless of her information. Such strategies merely pass the buck to P2.) Then, there is no finite sequential equilibrium strategy profile in which P1 always moves to some p_u in the first round when she sees $x_1 = 1$, and P1 always moves to $p_v \neq p_u$ in the first round when she sees $x_1 = 0$.

Proof. Let (σ_1, σ_2) be a sequential 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 p_x when $x_1 = 1$ and p_y 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_a, p_b, p_c, p_d 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:

$$qD(p_b||p_c) + (1-q)D(p_d||p_a) - D(p_x||p_y)$$

Now, from a well-known convexity property of the relative entropy [CT91, pp.32], we have:

$$\begin{aligned} qD(p_b||p_c) + (1-q)D(p_d||p_a) &\geq D(qp_b + (1-q)p_d || qp_c + (1-q)p_a) && (1) \\ \Rightarrow qD(p_b||p_c) + (1-q)D(p_d||p_a) &\geq D(p_x||p_y) && (2) \\ \Rightarrow qD(p_b||p_c) + (1-q)D(p_d||p_a) - D(p_x||p_y) &\geq 0 && (3) \end{aligned}$$

Inequality 2 follows from the definition of p_x and p_y. Thus, inequality 3 implies that bluffing is always at least as profitable as behaving myopically by lemma 2. Moreover, inequality 1 is strict when $q \neq 0, 1$ and $p_d \neq p_a, p_b$. 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 position condition we assumed is sufficient but not necessary. All that appears to be necessary is that the optimal value is not independent of either player's signal. This explains why there are no myopic equilibria even in the case of the AND function, for which the points are not in general position.

4.2 Nonexistence of finite equilibrium

Theorem 1 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 sequential 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. For this result, we make the simplifying assumption that all mixed strategies are discrete, *i.e.*, a player can only mix strategies over finitely many points in any one move. This is a very mild assumption; it should be possible to relax it with careful analysis.

Theorem 4. *Under conditions [C1]-[C3] of theorem 3, there is no sequential equilibrium strategy profile in which the market is certain to be in the optimal state after n rounds, for any finite n.*

Proof. Suppose the conditions [C1]-[C3] of theorem 3 are met. Let (σ_1, σ_2) be any sequential equilibrium. We now argue that, with nonzero probability over $P1$'s mixed strategy, the conditions [C1]-[C3] of the theorem will be met after $P1$'s first move, with the roles of $P1$ and $P2$ interchanged. First, observe that the values of the optimal probabilities p_a, p_b, p_c, p_d never change during the course of play; switching roles from $P1$ first to $P2$ first corresponds to simply swapping the labels on these points. Thus, they remain in general position throughout the execution of the strategy. (The points p_x and p_y may change with the players' beliefs, but this does not matter.) It is therefore sufficient to show that the other two properties are preserved (with nonzero probability).

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 p_u such that P1 moved to p_u with nonzero probability t when $x_1 = 0$ *and* with nonzero probability t' when $x_1 = 1$. Suppose there was no such p_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 p_u is played with finite probability. Further, conditioning on $P1$ moving to p_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 p_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 equilibrium cannot converge with certainty after n rounds. □

4.3 Possible extensions to the basic result

In this section, we briefly discuss our approach to generalizing the core results.

Our results are obtained for the logarithmic market scoring rule. We believe that the results will immediately be valid for many (if not all) proper scoring rules. The key element of the proof is the convexity of the relative entropy. Corresponding to any proper scoring rule, there is a different measure of relative entropy (see the article by Grunwald and Dawid [GD04]); if it also satisfies the convexity properties, our results would still be valid. We believe that the convexity arises naturally from the properties of proper scoring rules, but we have not yet verified this.

Another natural extension is to permit each players signal to have more than two possible values. Again, we conjecture that the core results will still be valid; a similar convexity argument seems attainable. However, the "general position" condition we rely on becomes increasingly restrictive as the number of signals grows, and so it will be useful to identify the weakest condition that be used.

A generalization to more than two players also seems feasible. Suppose a market has n players playing in turn. Then, from the first player's point of view, all the remaining players can be thought of as a single player with 2^{n-1} possible signals. Thus, the incentive to bluff should still be present. Again, it will be important to find a useful genericity condition.

5 Discussion and Future Work

This paper reports preliminary results on analyzing non-myopic strategies in a two-player prediction market setting. We find, surprisingly, that the myopic strategies are almost never optimal in a non-myopic sense. Of course, 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. We investigated a simple modification, which includes a form of discounting, to ameliorate this potential problem.

The need for discounting shows a connection to classical 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 players; 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.

This is work in progress, and there are several directions for future work. Analyzing situations with more than two players, and more than two signals, would be very useful. Finally, it would be interesting to derive a connection between the discount factor and the speed of convergence of the prediction market prices to their true value. We speculate that a reduction to concepts from information theory would be useful for this result as well, perhaps to bound the rate at which the entropy is reduced over time.

References

[AG92] Franklin Allen and Douglas Gale. Stock-price manipulation. *Review of Financial Studies*, 5(3):503–529, 1992.

[Bri50] G. Brier. Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78:1–3, 1950.

[CT91] Thomas M. Cover and Joy A. Thomas. *Elements of information theory*. Wiley-Interscience, New York, NY, USA, 1991.

[CY04] Archishman Chakraborty and Bilge Yilmaz. Informed manipulation. *Journal of Economic Theory*, 114:132–152, 2004.

[FFPS05] Joan Feigenbaum, Lance Fortnow, David M. Pennock, and Rahul Sami. Computation in a distributed information market. *Theoretical Computer Science*, 343:114–132, 2005. (A preliminary version appeared in the 2003 ACM Conference on Electronic Commerce).

[GD04] Peter Grunwald and A. Dawid. Game theory, maximum entropy, minimum discrepancy and robust bayesian decision theory. *Annals of Statistics*, 32:1367–1433, 2004.

[Han03] Robin Hanson. Combinatorial information market design. *Information Systems Frontiers*, 5(1):107–119, 2003.

[HOP06] 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.

[NS07] Evdokia Nikolova and Rahul Sami. A strategic model for information markets. In *Proceedings of the Seventh Annual ACM Conference on Electronic Commerce (EC '07) (to appear)*, 2007.

[Pen04] David Pennock. A dynamic parimutuel market for information aggregation. In *Proceedings of the Fourth Annual ACM Conference on Electronic Commerce (EC '04)*, June 2004.

[SR07] Koleman 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*.