



DIGITAL ACCESS TO SCHOLARSHIP AT HARVARD

Gaming Prediction Markets: Equilibrium Strategies with a Market Maker

The Harvard community has made this article openly available. [Please share](#) how this access benefits you. Your story matters.

Citation	Chen, Yiling, Stanko Dimitrov, Rahul Sami, Daniel M. Reeves, David M. Pennock, Robin D. Hanson, Lance Fortnow and Rica Gonen. 2009. Gaming prediction markets: equilibrium strategies with a market maker. <i>Algorithmica</i> 58(4): 930-969.
Published Version	doi:10.1007/s00453-009-9323-2
Accessed	April 22, 2018 10:39:17 AM EDT
Citable Link	http://nrs.harvard.edu/urn-3:HUL.InstRepos:5027877
Terms of Use	This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP

(Article begins on next page)

Gaming Prediction Markets: Equilibrium Strategies with a Market Maker*

Yiling Chen¹, Stanko Dimitrov², Rahul Sami³, Daniel M. Reeves⁴, David M. Pennock⁴, Robin D. Hanson⁵, Lance Fortnow⁶, Rica Gonen⁷

¹ Harvard University, School of Engineering and Applied Sciences, Cambridge, MA

² University of Michigan, Dept. of Industrial and Operations Engineering, Ann Arbor, MI

³ University of Michigan, School of Information, Ann Arbor, MI

⁴ Yahoo! Research, New York, NY

⁵ George Mason University, Department of Economics, Fairfax, VA

⁶ Northwestern University, EECS Department, Evanston, IL

⁷ Yahoo! Research, Santa Clara, CA

Received: 2 June 2008 / Accepted: 29 April 2009

Abstract We study the equilibrium behavior of informed traders interacting with *market scoring rule* (MSR) market makers. One attractive feature of MSR is that it is *myopically incentive compatible*: it is optimal for traders to report their true beliefs about the likelihood of an event outcome provided that they ignore the impact of their reports on the profit they might garner from future trades. In this paper, we analyze non-myopic strategies and examine what information structures lead to truthful betting by traders. Specifically, we analyze the behavior of risk-neutral traders with incomplete information playing in a dynamic game. We consider finite-stage and infinite-stage game models. For each model, we study the logarithmic market scoring rule (LMSR) with two different information structures: conditionally independent signals and (unconditionally) independent signals. In the finite-stage model, when signals of traders are independent conditional on the state of the world, truthful betting is a Perfect Bayesian Equilibrium (PBE). Moreover, it is the unique Weak Perfect Bayesian Equilibrium (WPBE) of the game. In contrast, when signals of traders are unconditionally independent, truthful betting is *not* a WPBE. In the infinite-stage model with unconditionally independent signals, there does not exist an equilibrium in which all information is revealed in a finite amount of time. We propose a simple discounted market scoring rule that reduces the opportunity for bluffing strategies. We show that in any WPBE for

* Preliminary versions of some of the results in this paper were presented in two conference papers, Chen *et al.* [10] and Dimitrov and Sami [13].

This is an author-created version. The original publication is available at www.springerlink.com. DOI 10.1007/s00453-009-9323-2.

the infinite-stage market with discounting, the market price converges to the fully-revealing price, and the rate of convergence can be bounded in terms of the discounting parameter. When signals are conditionally independent, truthful betting is the unique WPBE for the infinite-stage market with and without discounting.

Key words Prediction markets, game theory, bluffing, strategic betting

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. The strongest form of the *efficient markets hypothesis* [14] posits that information is incorporated into prices fully and immediately, as soon as it becomes available to anyone. A *prediction market* is a financial market specifically designed to take advantage of this property. For example, to forecast whether a product will launch on time, a company might ask employees to trade a security that pays \$1 if and only if the product launches by the planned date. Everyone from managers to developers to administrative assistants with different forms and amounts of information can bet on the outcome. The resulting price constitutes their collective probability estimate that the launch will occur on time. Empirically, such prediction markets outperform experts, group consensus, and polls across a variety of settings [16, 17, 28, 3, 4, 18, 31, 12, 8].

Yet the double-sided auction at the heart of nearly every prediction market is *not* incentive compatible. Information holders do not necessarily have incentive to fully reveal all their information right away, as soon as they obtain it. The extreme case of this is captured by the so-called *no trade theorems* [26]: When rational, risk-neutral agents with common priors interact in an unsubsidized (zero-sum) market, *the agents will not trade at all*, even if they have vastly different information and posterior beliefs. The informal reason is that any offer by one trader is a signal to a potential trading partner that results in belief revision discouraging trade.

The classic *market microstructure* model of a financial market posits two types of traders: rational traders and noise traders [24]. The existence of noise traders turns the game among rational traders into a positive-sum game, thereby resolving the no-trade paradox. However, even in this setting, the mechanism is not incentive compatible. For example, monopolist information holders will not fully reveal their information right away: instead, they will leak their information into the market gradually over time to obtain a greater profit [7].

Instead of assuming or subsidizing noise traders, a prediction market designer might choose to directly subsidize the market by employing an *automated market maker* that expects to lose money. Hanson's *market scoring rule market maker* (MSR) is one example [20, 21]. MSR requires a patron to subsidize the market but guarantees that the patron cannot lose more than a fixed amount set in advance, regardless of how many shares are exchanged or what outcome eventually occurs. The greater the subsidy, the greater the effective liquidity of the market. Since

traders face a positive-sum game, even rational risk-neutral agents have incentive to participate. In fact, even a single trader can be induced to reveal information, something impossible in a standard double auction with no market maker. Hanson proves that *myopic* risk-neutral traders have incentive to reveal all their information. However, forward-looking traders may not.

Though subsidized market makers improve incentives for information revelation, the mechanisms are still not incentive compatible. Much of the allure of prediction markets is the promise to gather information from a distributed group quickly and accurately. However, if traders have demonstrable incentives to either hide or falsify information, the accuracy of the resulting forecast may be in question. One frequent concern about incentives 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.

1.1 Our Results

In this paper, we study strategies under a logarithmic market scoring rule (LMSR) in a Bayesian extensive-form game setting with incomplete information. We show that different information structures can lead to radically different strategic properties by analyzing two natural classes of signal distributions: conditionally independent signals and unconditionally independent signals.

For conditionally independent signals, we show that truthful betting is a Perfect Bayesian Equilibrium in finite-stage and infinite-stage models. Further, we show that the truthful betting equilibrium is the unique Weak Perfect Bayesian Equilibrium in this setting. Thus, bluffing strategies are not a concern.

In contrast, when signals are unconditionally independent, we show that truthful betting is not a Weak Perfect Bayesian Equilibrium in the finite-stage model. In the infinite-stage model, we show that there is no Weak Perfect Bayesian Equilibrium that results in full information revelation in a finite number of trades. We propose and analyze a discounted version of the LMSR that mitigates the strategic hazard. In the discounted LMSR, we prove that under any WPBE strategy, the relative entropy between the market price and the optimal full-information price tends to zero at an exponential rate. The rate of convergence can be bounded in terms of the discounting parameter and a measure of complementarity in the information setting.

1.2 Related Work

Theoretical work on price manipulation in financial markets [1, 7, 23] explains the logic of manipulation and indicates that double auctions are not incentive compatible. 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 stage and then profit in later stages, even though the other traders are aware of the possibility of deception and act rationally. They use a stylized model of a multi-stage 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: Allen and Gale's model [1] needs to assume traders with different risk attitudes to get around no-trade results, which is rendered unnecessary by the inherent subsidy with a market scoring rule mechanism. 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 [7] for references to other research on manipulation in financial markets.

There are some experimental and empirical studies on price manipulation in prediction markets using double auction mechanisms. The results are mixed, some giving evidence for the success of price manipulation [19] and some showing the robustness of prediction markets to price manipulation [6, 22, 29, 30].

Feigenbaum *et al.* [15] also study prediction markets in which the information aggregation is sometimes slow, and sometimes fails altogether. In their setting, the aggregation problems arise from a completely different source: the traders are non-strategic 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 potentially non-myopic strategies. Nikolova and Sami [27] 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.

Axelrod *et al.* [2] also propose a form of discounting in an experimental parimutuel market and show that it promotes 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.

Börger *et al.* [5] study when signals are substitutes and complements in a general setting. Our equilibrium and convergence result suggests that prediction markets are one domain where this distinction is of practical importance.

This paper is a synthesis and extension of two independent sets of results, which were presented in preliminary form by Chen *et al.* [10] and Dimitrov and Sami [13].

1.3 Structure of the Paper

This article is organized as follows: In section 2, we present a little background information about market scoring rules. Section 3 details our formal model and information structures. In Section 4 we investigate how predefined sequence of play affects players' expected profits in LMSR. In Section 5 we consider, when a player can choose to play first or second, what is its equilibrium strategy in a 2-player LMSR. Results of Sections 4 and 5 serve as building blocks for our analysis in subsequent sections. Section 6 studies the equilibrium of a 2-player,

3-stage game. We generalize this result to arbitrary finite-player finite-stage games in Section 7. We analyze infinite-stage games in Section 8. Section 9 proposes a discounted LMSR and analyzes how the mechanism discourages non-truthful behavior. Finally, we conclude in Section 10.

2 Background

Consider a discrete random variable X that has n mutually exclusive and exhaustive outcomes. Subsidizing a market to predict the likelihood of each outcome, market scoring rules are known to guarantee that the market maker's loss is bounded.

2.1 Market Scoring Rules

Hanson [20,21] shows how a proper scoring rule can be converted into a market maker mechanism, called *market scoring rules* (MSR). The market maker uses a proper scoring rule, $S = \{s_1(\mathbf{r}), \dots, s_n(\mathbf{r})\}$, where $\mathbf{r} = \langle r_1, \dots, r_n \rangle$ is a reported probability estimate for the random variable X . Conceptually, every trader in the market may change the current probability estimate to a new estimate of its choice at any time as long as it agrees to pay the scoring rule payment associated with the current probability estimate and receive the scoring rule payment associated with the new estimate. If outcome i is realized, a trader that changes the probability estimate from \mathbf{r} to $\tilde{\mathbf{r}}$ pays $s_i(\mathbf{r})$ and receives $s_i(\tilde{\mathbf{r}})$.

Since a proper scoring rule is incentive compatible for risk-neutral agents, if a trader can only change the probability estimate once, this modified proper scoring rule still incentivizes the trader to reveal its true probability estimate. However, when traders can participate multiple times, they might have incentive to manipulate information and mislead other traders.

Because traders change the probability estimate in sequence, MSR can be thought of as a sequential shared version of the scoring rule. The market maker pays the last trader and receives payment from the first trader. An MSR market can be equivalently implemented as a market maker offering n securities, each corresponding to one outcome and paying \$1 if the outcome is realized [20,9]. Hence, changing the market probability of outcome i to some value r_i is the same as buying the security for outcome i until the market price of the security reaches r_i . Our analysis in this paper is facilitated by directly dealing with probabilities.

A popular MSR is the *logarithmic market scoring rule* (LMSR) where the logarithmic scoring rule

$$s_i(\mathbf{r}) = b \log(r_i) \quad (b > 0), \quad (1)$$

is used. A trader's expected profit in LMSR directly corresponds to the concept of *relative entropy* in information theory. If a trader with probability $\tilde{\mathbf{r}}$ moves the market probability from \mathbf{r} to $\tilde{\mathbf{r}}$, its expected profit (score) in LMSR is

$$S(\mathbf{r}, \tilde{\mathbf{r}}) = \sum_{i=1}^n \tilde{r}_i (s_i(\tilde{\mathbf{r}}) - s_i(\mathbf{r})) = b \sum_{i=1}^n \tilde{r}_i \log \frac{\tilde{r}_i}{r_i} = bD(\tilde{\mathbf{r}}||\mathbf{r}). \quad (2)$$

$D(\mathbf{p}||\mathbf{q})$ is the relative entropy or Kullback Leibler distance between two probability mass functions $p(x)$ and $q(x)$ and is defined in [11] as

$$D(\mathbf{p}||\mathbf{q}) = \sum_x p(x) \log \frac{p(x)}{q(x)}.$$

$D(\mathbf{p}||\mathbf{q})$ is nonnegative and equals zero only when $\mathbf{p} = \mathbf{q}$. The maximum amount a LMSR market maker can lose is $b \log n$. Since b is a scaling parameter, without loss of generality we assume that $b = 1$ in the rest of the paper.

2.2 Terminology

Truthful betting (TB) for a player in MSR is the strategy of immediately changing the market probability to the player's believed probability. In other words, it is the strategy of always buying immediately when the price is too low and selling when the price is too high. "Too low" and "too high" are determined by the player's information. The price is too low when the current expected payoff is higher than the price, and too high when current expected payoff is lower than the price. Truthful betting fully reveals a player's payoff-relevant information. *Bluffing* is the strategy of betting contrary to one's information in order to deceive future traders, with the intent of capitalizing on their resultant misinformed trading. This paper investigates scenarios where traders with incomplete information have an incentive to deviate from truthful betting.¹

3 LMSR in a Bayesian Framework

In this part, we introduce our game theoretic model of LMSR market in order to capture the strategic behavior in LMSR when players have private information.

3.1 General Settings

We consider a single event that is the subject of our predictions. $\Omega = \{Y, N\}$ is the state space of this event. The true state, $\omega \in \Omega$, is picked by nature according to a prior $\mathbf{p}^0 = \langle p_Y^0, p_N^0 \rangle = \langle \Pr(\omega = Y), \Pr(\omega = N) \rangle$. A market, aiming at predicting the true state ω , uses a LMSR market maker with initial probability estimate $\mathbf{r}^0 = \langle r_Y^0, r_N^0 \rangle$.

There are m risk neutral players in the market. Each player i gets a private signal, c_i , about the state of the world at the beginning of the market. \mathbf{C}_i is the signal space of player i with $|\mathbf{C}_i| = n_i$. The actual realization of the signal is observed only by the player receiving the signal. The joint distribution of the true state and players' signals, $\mathcal{P} : \Omega \times \mathbf{C}_1 \times \cdots \times \mathbf{C}_m \mapsto [0, 1]$, is common knowledge.

¹ With complete information, traders should reveal all information right away in MSR, because the market degenerates to a race to capitalize on the shared information first.

Players trade sequentially, in one or more stages of trading, in the LMSR market. Players are risk-neutral Bayesian agents. If a player i is designated to trade at some stage k in the sequence, it can condition its beliefs of the likelihood of an outcome on its private signals, as well as the observed prices after the first $k - 1$ trades.

Up to this point, we have not made any assumptions about the joint distribution \mathcal{P} . It turns out that the strategic analysis of the market depends critically on independence properties of players' signals. We study two models of independence – independence conditional on the true state, and unconditional independence – that are both natural in different settings. These are described below.

3.1.1 Games with Conditionally Independent (CI) Signals We start with a concrete example before formally introducing the model for conditionally independent signals. Suppose the problem for a prediction market is to predict whether a batch of product is manufactured with high quality materials or low quality materials. If the product is manufactured with high quality materials, the probability for a product to break in its first month of use is 0.01. If the product is manufactured with low quality materials, the probability for a product to break in its first month of use is 0.1. Some consumers who each bought a product have private observations of whether their products break in the first month of use. The quality of the materials will be revealed by a test in the future. This is an example where consumers have conditionally independent signals – conditional on the quality of the materials, consumers' observations are independent. Conditionally independent signals are usually appropriate for modeling situations where the true state of the world has been determined but is unknown, and signal realizations are influenced by the true state of the world.

Formally, in games with conditionally independent signals, players' signals are assumed to be independent conditional on the state of the world, *i.e.*, conditioned on the eventual outcome of the event. In other words, for any two players i and j , $\Pr(c_i, c_j | \omega) = \Pr(c_i | \omega) \Pr(c_j | \omega)$ is always satisfied by \mathcal{P} . This class can be interpreted as player i 's signal c_i is independently drawn by nature according to some conditional probability distribution $\mathbf{p}_{(c_i|Y)}$ if the true state is Y , and analogously $\mathbf{p}_{(c_i|N)}$ if the true state is N .

In order to rule out degenerated cases, we further assume that conditionally independent signals are *informative* and *distinct*. An informative signal means that after observing the signal, a player's posterior is different from its prior. Intuitively, when observing a signal does not change a player's belief, we can simply remove the signal from the player's signal space because it does not provide any information. Formally, signals are informative if and only if $\Pr(c_i = a | \omega = Y) \neq \Pr(c_i = a | \omega = N)$, $\forall 1 \leq i \leq m$ and $\forall a \in \mathbf{C}_i$. Player i has distinct signals when its posterior probability is different after observing different signals. When two signals give a player the same posterior, we can combine these two signals into one because they provide same information. Formally, signals are distinct if and only if $\Pr(\omega = Y | c_i = a) \neq \Pr(\omega = Y | c_i = a')$, $\forall a, a' \in \mathbf{C}_i$, $a \neq a'$, and $\forall i$.

Lemma 1 shows that conditional independence implies unconditional dependence.

Lemma 1 *If players have informative signals that are conditionally independent, their signals are unconditionally dependent.*

Proof Suppose the contrary, signals of players i and j , c_i and c_j , are unconditionally independent. Then, $\Pr(c_i = a, c_j = b) - \Pr(c_i = a) \Pr(c_j = b) = 0$ must be satisfied for all $a \in \mathbf{C}_i$ and $b \in \mathbf{C}_j$. By conditional independence of signals,

$$\begin{aligned}
& \Pr(c_i = a, c_j = b) - \Pr(c_i = a) \Pr(c_j = b) \\
&= \sum_z \Pr(c_i = a, c_j = b | \omega = z) \Pr(\omega = z) \\
&\quad - \sum_z \Pr(c_i = a | \omega = z) \Pr(\omega = z) \sum_z \Pr(c_j = b | \omega = z) \Pr(\omega = z) \\
&= \sum_z \Pr(c_i = a | \omega = z) \Pr(c_j = b | \omega = z) \Pr(\omega = z) \\
&\quad - \sum_z \Pr(c_i = a | \omega = z) \Pr(\omega = z) \sum_z \Pr(c_j = b | \omega = z) \Pr(\omega = z) \\
&= \Pr(Y) \Pr(N) (\Pr(c_i = a | Y) - \Pr(c_i = a | N)) (\Pr(c_j = b | Y) - \Pr(c_j = b | N)).
\end{aligned}$$

The above expression equals 0 only when $\Pr(c_i = a | Y) - \Pr(c_i = a | N) = 0$ or $\Pr(c_j = b | Y) - \Pr(c_j = b | N) = 0$ or both, which contradicts the informativeness of signals. Hence, c_i and c_j are unconditionally dependent. \square

3.1.2 Games with (Unconditionally) Independent (I) Signals In some other situations, signal realizations are not caused by the state of the world, but instead, they might stochastically influence the eventual outcome of the event. For instance, in a political election prediction market, voters' private information is their votes, which can arguably be thought as independent of each other. The election outcome – which candidate gets the majority of votes – is determined by all votes. This is an example of a game with unconditionally independent signals. Formally, for any two players i and j in games with independent signals, $\Pr(c_i, c_j) = \Pr(c_i) \Pr(c_j)$ is always satisfied by \mathcal{P} .

For games with independent signals, we primarily prove results about the *lack* of truthful equilibria. Such results require us to show that there is a strict advantage to deviate to an alternative strategy. In order to rule out degenerate cases in which the inequalities are not strict, we will often invoke the following *general informativeness condition*.

Definition 1 *An instance of the prediction market with m players and joint distribution \mathcal{P} 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: $\forall i, i' \in \mathbf{C}_1$ and $\forall j, j' \in \mathbf{C}_2$ such that $i \neq i', j \neq j'$: $\Pr(Y | c_1 = i, c_2 = j) \neq \Pr(Y | c_1 = i, c_2 = j')$ and $\Pr(Y | c_1 = i, c_2 = j) \neq \Pr(Y | c_1 = i', c_2 = j)$. For $m > 2$, we must have $\forall \mathbf{j}, \forall i' \neq i, \Pr(Y | i, \mathbf{j}) \neq$*

$\Pr(Y|i', \mathbf{j})$, where i, i' are two possible signals for any one player, and \mathbf{j} is a vector of signals for the other $m - 1$ players.

By Lemma 1, we know that unconditional independence implies conditional dependence. We note that the two information structures, conditional independence and unconditional independence, that we discuss in this paper are mutually exclusive, but not exhaustive. We do not consider the case when signals are both conditionally dependent and unconditionally dependent.

3.2 Equilibrium Concepts

The prediction market model we have described is an extensive-form game between m 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 probabilities of other players' signal realizations, contingent on reaching the node.

An assessment profile (A_1, \dots, A_m) , consisting of an assessment for each player, is a *Weak Perfect Bayesian Equilibrium* (WPBE) if and only if, 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* (PBE) or sequential equilibrium, which further require beliefs at nodes that are off the equilibrium path to be consistent, are used. In this paper, when giving a specific equilibrium, we use the refined PBE concept. When proving the nonexistence of truthful betting equilibrium and characterizing the set of all equilibria, we use the WPBE concept, because the results thus hold *a fortiori* for refinements of the WPBE concept. For a formal definition of the equilibrium concepts, we refer the reader to the book by Mas-Colell *et al.* [25].

Given the strategy components of a WPBE 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 for WPBEs. We will abuse notation slightly by simply referring to an "equilibrium strategy profile",

leaving the beliefs implicit, for WPBEs. For PBEs, we will explicitly explain players' beliefs in proofs.

4 Profits under Predefined Sequence of Play

In this section, we investigate how sequence of play affects players' expected profits in LMSR. We answer the question: when two players play myopically according to a predefined sequence of play, are the players better off on expectation by playing first or second?

In particular, we consider two players playing in sequence and define some notations of strategy and expected profits to facilitate our analysis on how the profit of the second player depends on the strategy employed by the first player. Although we are studying 2-stage games here, we use them as a building block for our results in subsequent sections, so the definitions are more general and the first player's strategy may not be truthful. In fact, our analysis apply to the first two stages of any game.

Definition 2 *Given a game and a sequence of play, a first-stage strategy σ_1 for the first player is a specification, for each possible signal received by that player, of a mixed strategy over moves.*

*Two specific first-stage strategies we will study are the **truthful** strategy (denoted σ^T) and the **null** strategy (denoted σ^N). The truthful strategy specifies that, for each signal c_i it receives, the first player changes market probability to its posterior belief after seeing the signal. The null strategy specifies that, regardless of its signal c_i , the first player leaves the market probability unchanged, the same as she found it at.*

In a 2-stage game in which Alice plays first, followed by Bob, we use notation $\pi^B(\sigma_1)$ to denote Bob's expected myopic optimal profit following a first-stage strategy σ_1 by Alice, assuming that Bob knows the strategy σ_1 and can condition on it. Likewise, in a game in which Alice moves after Bob, we use $\pi^A(\sigma_1)$ to denote Alice's expected profit.

Note that $\pi^B(\sigma^N)$ is equal to the expected myopic profit Bob would have had if he had moved first, because under the null strategy σ^N Alice does not move the market. Also, $\pi^B(\sigma^T)$ is the profit that Bob earns when Alice follows her myopically optimal strategy.

We begin with the following simple result showing that the truthful strategy is optimal for a player moving only once. This follows from the myopic optimality of LMSR, but we include a proof for completeness.

Lemma 2 *In a LMSR market, if stage t is player i 's last chance to play and μ_i is player i 's belief over actions of previous players, player i 's best response at stage t is to play truthfully by changing the market probabilities to $\mathbf{r}^t = \langle \Pr(Y|c_i, \mathbf{r}^1, \dots, \mathbf{r}^{t-1}, \mu_i), \Pr(N|c_i, \mathbf{r}^1, \dots, \mathbf{r}^{t-1}, \mu_i) \rangle$, where $\mathbf{r}^1, \dots, \mathbf{r}^{t-1}$ are the market probability vectors before player i 's action.*

Proof When a player has its last chance to play in LMSR, it is the same as the player interacting with a logarithmic scoring rule. Because the logarithmic scoring rule is strictly proper, player i 's expected utility is maximized by truthfully reporting its posterior probability estimate given the information it has. \square

Next, we provide a lemma that is very useful for obtaining subsequent theorems. Since the state of the world has only two exclusive and exhaustive outcomes, Y and N . We use the term *position* to refer the market probability for outcome Y . The market probability for outcome N is then uniquely defined because the probabilities for the two outcomes sum to 1.

Lemma 3 *Let σ_1 be a first-stage strategy for Alice that minimizes the expression $\pi^B(\sigma)$ over all possible σ . Then, σ_1 satisfies the following consistency condition: For any position x that Alice moves to under σ_1 , the probability of the outcome Y occurring conditioned on σ_1 and the fact that Alice moved to x is exactly equal to x .*

Proof Suppose that σ_1 does not satisfy this condition. We construct a perturbed first-stage strategy $\hat{\sigma}_1$, and show that $\pi^B(\hat{\sigma}_1) < \pi^B(\sigma_1)$, thus contradicting the minimality of σ_1 . Such a $\hat{\sigma}_1$ is easily constructed from any given σ_1 , as follows. We start by setting $\hat{\sigma}_1 = \sigma_1$. Consider any one position x that Alice moves to with positive probability under strategy σ_1 , and for which the consistency condition is failed. Let $q_x = \Pr(Y|x_A = x, \text{ Alice following } \sigma_1)$. Then, whenever σ_1 dictates that Alice move to x , we set $\hat{\sigma}_1$ to dictate that Alice move to $\hat{x} = q_x$ instead. We repeat this perturbation for each x such that $x \neq q_x$, thus resulting in a consistent strategy $\hat{\sigma}_1$.

The actual position that Alice moves to can be thought of as a random variable on a sample space that includes the randomly-distributed signals Alice receives as well as the randomization Alice uses to play a mixed strategy. We use x_A and \hat{x}_A to denote random variables that takes on values x and \hat{x} respectively. Similarly, the event ω is a random variable that takes on two values $\{Y, N\}$. Recall that the moves are actually probability distributions. In order to use summation notation, we define $x_Y = x$, $x_N = 1 - x$; and, likewise, $\hat{x}_Y = \hat{x}$, $\hat{x}_N = 1 - \hat{x}$. By definition of \hat{x} , $\hat{x}_z = \Pr(\omega = z|x_A = \hat{x}, \text{ Alice following } \hat{\sigma}_1)$. Note that any x has a corresponding value of \hat{x} ; thus, we may write expressions like $\sum_x \hat{x}$ in which \hat{x} is implicitly indexed by x . We use $\Pr(x, b, z)$ to represent $\Pr(x_A = x, c_B = b, \omega = z)$, and other notation is defined similarly.

The intuitive argument we use is as follows: Consider the situation in which Alice is known to be following strategy σ_1 . The total profit of two consecutive moves, in a market using the LMSR, is exactly the payoff of moving from the starting point to the end point of the second move. Hence, for any given position x that Alice leaves the market price at, we decompose Bob's response into two virtual steps. In the first virtual step, Bob moves the market to the corresponding \hat{x} . In the second step, Bob moves the market from \hat{x} to his posterior probability $\Pr(z|x, b)$. Bob's actual profit is the sum of his profit from these two steps. Note that the second step alone yields Bob at least $\pi^B(\hat{\sigma}_1)$ in expectation, because in both cases the move begins at \hat{x} , and because Bob can infer \hat{x} from x . We further argue that the first step yields a strictly positive expected profit, because for any

position x , the posterior probability of outcome Y is \hat{x}_Y . The first step thus moves the market from a less accurate value to the most accurate possible value given information x ; this always yields a positive expected profit in the LMSR.

Formalizing this argument, we compare the payoffs $\pi^B(\sigma_1)$ and $\pi^B(\hat{\sigma}_1)$ as follows.

$$\begin{aligned}
\pi^B(\sigma_1) - \pi^B(\hat{\sigma}_1) &= \sum_{x,b,z} \Pr(x, b, z) [\log \Pr(z|x, b) - \log x_z] \\
&\quad - \sum_{x,b,z} \Pr(x, b, z) [\log \Pr(z|\hat{x}, b) - \log \hat{x}_z] \\
&= \sum_{x,b,z} \Pr(x, b, z) [\log \Pr(z|x, b) - \log \Pr(z|\hat{x}, b)] \\
&\quad + \sum_{x,b,z} \Pr(x, b, z) [\log \hat{x}_z - \log x_z] \\
&= \sum_{x,b} \Pr(x, b) \sum_z \Pr(z|x, b) [\log \Pr(z|x, b) - \log \Pr(z|\hat{x}, b)] \\
&\quad + \sum_x \Pr(x) \sum_z \Pr(z|x) [\log \hat{x}_z - \log x_z] \\
&= \sum_{x,b} \Pr(x, b) D(\mathbf{p}(\omega|x_A, c_B) \| \mathbf{p}(\omega|\hat{x}_A, c_B)) \\
&\quad + \sum_x \Pr(x) \sum_z \hat{x}_z [\log \hat{x}_z - \log x_z] \\
&= \sum_{x,b} \Pr(x, b) D(\mathbf{p}(\omega|x_A, c_B) \| \mathbf{p}(\omega|\hat{x}_A, c_B)) \\
&\quad + \sum_x \Pr(x) D(\mathbf{p}(\hat{x}_A) \| \mathbf{p}(x_A))
\end{aligned}$$

where $\mathbf{p}(\omega|x_A, c_B)$ and $\mathbf{p}(\omega|\hat{x}_A, c_B)$ are the conditional distributions of ω , and $\mathbf{p}(\hat{x}_A)$ and $\mathbf{p}(x_A)$ are the probability distributions of \hat{x}_A and x_A respectively. $D(\mathbf{p}(\omega|x_A, c_B) \| \mathbf{p}(\omega|\hat{x}_A, c_B))$ and $D(\mathbf{p}(\hat{x}_A) \| \mathbf{p}(x_A))$ are relative entropy, which is known to be nonnegative and strictly positive when the two distributions are not the same. We assumed that σ_1 does not meet the consistency condition, and thus, there is at least one x such that $\Pr(x) > 0$ and $\hat{x} \neq x$. Thus, we have $D(\mathbf{p}(\hat{x}_A) \| \mathbf{p}(x_A)) > 0$. Hence, $\pi^B(\sigma_1) > \pi^B(\hat{\sigma}_1)$. This contradicts the minimality assumption of σ_1 . \square

4.1 Profits in CI Games

When Alice and Bob have conditionally independent signals and Alice plays first, we will prove that, for any strategy σ_1 that Alice chooses, if Bob is aware that Alice is following strategy σ_1 , his expected payoff from following the myopic strategy (after conditioning his beliefs on Alice's actual move) is at least as much as he could expect if Alice had chosen to play truthfully.

First, we show that in CI games, observing Alice's posterior probabilities is equally informative to Bob as observing Alice's signal directly.

Lemma 4 *When players have conditionally independent signals, if player i knows player j 's posterior probabilities $\langle \Pr(Y|c_j), \Pr(N|c_j) \rangle$, player i can infer the posterior probabilities conditional on both signals. More specifically,*

$$\Pr(\omega|c_i, c_j) = \frac{\Pr(c_i|\omega) \Pr(\omega|c_j)}{\Pr(c_i|Y) \Pr(Y|c_j) + \Pr(c_i|N) \Pr(N|c_j)},$$

where $\omega \in \{Y, N\}$.

Proof Using Bayes rule, we have

$$\begin{aligned} \Pr(\omega|c_i, c_j) &= \frac{\Pr(\omega, c_i|c_j)}{\Pr(c_i|c_j)} \\ &= \frac{\Pr(c_i|c_j, \omega) \Pr(\omega|c_j)}{\Pr(c_i|c_j, Y) \Pr(Y|c_j) + \Pr(c_i|c_j, N) \Pr(N|c_j)} \\ &= \frac{\Pr(c_i|\omega) \Pr(\omega|c_j)}{\Pr(c_i|Y) \Pr(Y|c_j) + \Pr(c_i|N) \Pr(N|c_j)}. \end{aligned}$$

The third equality comes from the CI condition. \square

We can now prove our result on Bob's myopic profit in the CI setting.

Theorem 1 *For any CI game G with Alice playing first and Bob playing second, and any σ_1 , $\pi^B(\sigma_1) \geq \pi^B(\sigma^T)$. When signals are informative and distinct, the equality holds if and only if $\sigma_1 = \sigma^T$.*

Proof We use an information-theoretic argument to prove this result. We can view Alice's signal c_A , Alice's actual move x_A (for outcome Y), Bob's signal c_B , the event ω as random variables. We then show that Bob's expected payoff can be characterized as the entropy of c_B conditioned on observing x_A , plus another term that does not depend on Alice. The more information that x_A reveals about c_B , the lower the conditional entropy $H(c_B|x_A)$, and thus, the lower Bob's expected profit. However, note that x_A is a (perhaps randomized) function of Alice's signal c_A alone. A fundamental result from information theory states that x_A can reveal no more information about c_B than c_A can, and thus, Bob's profit is minimized when x_A reveals as much information as c_A . By Lemma 4, the truthful strategy for Alice reveals all information that is in c_A , and thus, attains the minimum.

We now formalize this argument, retaining the notation above. By Lemma 3, we can assume that σ_1 is consistent, without loss of generality. Let us analyze Bob's expected payoff when Alice follows this strategy σ_1 . The unit of our analysis is a particular realization of a combination $c_A = a, x_A = x, c_B = b, \omega = z$. Bob will move to a position $y(x, \sigma_1, b) = \Pr(Y|x_A = x, \text{ Alice following } \sigma_1, c_B = b)$. For conciseness, we drop σ_1 from the notation, and write $y(x, b)$.

As before use x_z and y_z to denote the probability of $\omega = z$ inferred from positions x and y respectively. By definition of the LMSR, we have:

$$\pi^B(\sigma_1) = \sum_{x,b,z} \Pr(x, b, z) [\log y_z(x, b) - \log x_z]$$

Next, we note that $y_z(x, b) = \Pr(\omega = z|x, b)$. Further, it is immediate that $\sum_b \Pr(b|x, z) = 1$. Expanding, we have

$$\begin{aligned}
\pi^B(\sigma_1) &= \sum_{x,b} \Pr(x, b) \sum_z \Pr(z|x, b) \log y_z(x, b) \\
&\quad - \sum_x \Pr(x) \sum_z \Pr(z|x) \sum_b \Pr(b|x, z) \log x_z \\
&= \sum_{x,b} \Pr(x, b) \left[\sum_z \Pr(z|x, b) \log \Pr(z|x, b) \right] \\
&\quad - \sum_x \Pr(x) \left[\sum_z \Pr(z|x) \log \Pr(z|x) \right] \\
&= -H(\omega|x_A, c_B) + H(\omega|x_A)
\end{aligned}$$

where we have identified the terms with the standard definition of conditional entropy of random variables [11, pg. 17].

Using the relation $H(X, Y) = H(X) + H(Y|X)$, we can write

$$\begin{aligned}
\pi^B(\sigma_1) &= -H(\omega, x_A, c_B) + H(x_A, c_B) + H(\omega, x_A) - H(x_A) \\
&= [H(x_A, c_B) - H(x_A)] - [H(\omega, x_A, c_B) - H(\omega, x_A)] \\
&= H(c_B|x_A) - H(c_B|\omega, x_A) \\
&= H(c_B|x_A) - H(c_B|\omega)
\end{aligned}$$

The last transformation $H(c_B|\omega, x_A) = H(c_B|\omega)$ follows because c_B is conditionally independent of x_A conditioned on ω , and thus, knowledge of x_A does not alter its conditional distribution or its conditional entropy.

The second term is clearly independent of σ_1 . For the first term, we note that $H(c_B|c_A) \leq H(c_B|x_A)$ because x_A is a function of c_A [11, pg. 35]. More specifically,

$$\begin{aligned}
H(c_B|x_A) &= - \sum_x \Pr(x) \left[\sum_b \Pr(b|x) \log \Pr(b|x) \right] \\
&= - \sum_x \Pr(x) \left[\sum_b \left(\sum_a \Pr(a|x) \Pr(b|a) \right) \log \left(\sum_a \Pr(a|x) \Pr(b|a) \right) \right] \\
&\geq - \sum_x \Pr(x) \sum_b \sum_a [\Pr(a|x) \Pr(b|a) \log \Pr(b|a)] \\
&= - \sum_x \Pr(x) \sum_a \left[\Pr(a|x) \sum_b [\Pr(b|a) \log \Pr(b|a)] \right] \\
&= - \sum_a \Pr(a) \left[\sum_b \Pr(b|a) \log \Pr(b|a) \right] \\
&= H(c_B|c_A).
\end{aligned}$$

The inequality follows from the strict convexity of the function $x \log x$. The equality holds only when the distribution of $c_B|a$ is the same as the distribution of $c_B|a'$ for all $a, a' \in \mathbf{C}_A$, or for any x there exists some a such that $\Pr(c_A = a|x_A = x) = 1$. If $c_B|a$ and $c_B|a'$ have the same distribution for all a and a' , c_A and c_B are unconditionally independent which contradicts with the CI condition by Lemma 1. Thus, the equality holds only when σ_1 is a strategy such that for any x there exists some a that satisfies $\Pr(c_A = a|x_A = x) = 1$.

The truthful strategy σ^T satisfies the above condition. Given any signal a that Alice might see, let x_a be the position that Alice would move to if she followed the truthful strategy, and let x_A^T denote the corresponding random variable. Then $x_a = \Pr(Y|c_A = a)$. Because signals are distinct, the value of x_a uniquely corresponds to signal a . Hence, the condition is satisfied with $\Pr(c_A = a|x_A^T = x_a) = 1$.

All that remains to be shown is that there is no other strategy $\sigma_1 \neq \sigma^T$ that satisfies the condition. For some strategy σ_1 , let \mathbf{X}_A denote the value space of x_A . Suppose the contrary, for any $x \in \mathbf{X}_A$ there exists a_x such that $\Pr(c_A = a_x|x_A = x) = 1$. We first show that when $x \neq x'$ and $x, x' \in \mathbf{X}_A$, it must be that $a_x \neq a_{x'}$. According to Lemma 3, $\Pr(Y|x_A = x, \text{Alice follows } \sigma_1) = x$, which gives

$$\sum_{a'} \Pr(Y|c_A = a') \Pr(c_A = a'|x_A = x) = x \implies \Pr(Y|c_A = a_x) = x.$$

This means that $\Pr(Y|c_A = a_x) \neq \Pr(Y|c_A = a_{x'})$ when $x \neq x'$. Hence, $a_x \neq a_{x'}$ because signals are distinct. Next, we show that when $c_A = a$, it must be the case that $x_A = \Pr(Y|c_A = a)$ with probability 1 according to strategy σ_1 . If with some positive probability $x_A = x' \neq \Pr(Y|c_A = a)$, then, because $\Pr(c_A = a_{x'}|x_A = x') = 1$, $c_A = a_{x'} \neq a$, and contradiction arises. Thus, for all $a \in \mathbf{C}_A$, when $c_A = a$, Alice moves to $\Pr(Y|c_A = a)$ with probability 1 under σ_1 . Hence, strategy σ_1 is the truthful strategy σ^T . \square

The following corollary immediately follows from Theorem 1 and answers that question we proposed at the beginning of the section for CI games.

Corollary 1 *In CI games, when two players play myopically according to a pre-defined sequence of play, it is better off for a player to play first than second.*

Proof We compare Bob's expected profit in Alice-Bob and Bob-Alice cases when both players play myopically. Bob's expected profit in Alice-Bob case equals $\pi^B(\sigma^T)$, while his expected profit in Bob-Alice cases equals $\pi^B(\sigma^N)$. According to Theorem 1, $\pi^B(\sigma^N) > \pi^B(\sigma^T)$. Hence, Bob is better off when the Bob-Alice sequence is used. \square

4.2 Profits in I Games

When Alice and Bob have unconditionally independent signals and Alice plays first, we will prove that, for any strategy σ_1 that Alice chooses, if Bob is aware that Alice is following strategy σ_1 , his expected payoff from following the myopic strategy (after conditioning his beliefs on Alice's actual move) is at least as much as he could expect if Alice had chosen to play the null strategy σ^N .

Theorem 2 For any I game G with Alice playing first and Bob playing second and any first-stage strategy σ_1 , if the market starts with the prior probabilities for the two outcomes, i.e. $\mathbf{r}^0 = \langle \Pr(\omega = Y), \Pr(\omega = N) \rangle$, we have $\pi^B(\sigma_1) \geq \pi^B(\sigma^N)$. The equality holds if and only if $\sigma_1 = \sigma^N$.

Proof Assume that there is a strategy σ_1 for Alice that minimizes $\pi^B(\sigma_1)$ over all possible first-stage strategies. According to Lemma 3, we can assume that σ_1 is consistent, without loss of generality. We will show that σ_1 must involve Alice not moving the market probability at all.

Let x_A be the random variable of Alice's move for outcome Y under strategy σ_1 . We first argue that x_A is deterministic, i.e. it can only have a single value. Suppose the contrary, then σ_1 would have support over a set of points: at least two points E, F , and perhaps a set of other points R . In this case, we show that we can construct a strategy σ'_1 such that $\pi^B(\sigma'_1) < \pi^B(\sigma_1)$ by "mixing" point E and another point H . Define u_{E_i} and u_{F_i} as the probability (under σ_1) that Alice has signal $c_A = a_i$ and sets $x_A = E$ and $x_A = F$ respectively. Let p_E be the probability that Alice plays E and similarly p_F be the probability that Alice plays F . Without loss of generality let $p_F < p_E$. Define $\alpha_i = \Pr(c_A = a_i | x_A = E) = \frac{u_{E_i}}{p_E}$ and $\beta_i = \Pr(c_A = a_i | x_A = F) = \frac{u_{F_i}}{p_F}$. With these definitions we can define the myopic move of Bob given that market is at E and $c_B = b_j$ as $r_j^\alpha = \sum_i \alpha_i \Pr(Y | c_A = a_i, c_B = b_j)$. Similarly r_j^β is defined as the myopic move of Bob given the market is at F . We use \mathbf{r}_j^α and \mathbf{r}_j^β to denote the probability vectors for two outcomes $\langle r_j^\alpha, 1 - r_j^\alpha \rangle$ and $\langle r_j^\beta, 1 - r_j^\beta \rangle$ respectively.

Now, let $H = (E + F)/2$ be the midpoint of E and F , and consider a new strategy σ'_1 over points E, H , and the same set of remaining points R . Under σ'_1 , the probability that Alice has signal $c_A = a_i$ and sets $x_A = E$ is $\frac{p_E - p_F}{p_E} u_{E_i}$, and the probability that Alice has signal $c_A = a_i$ and sets $x_A = H$ equals $\frac{p_F}{p_E} u_{E_i} + u_{F_i}$. Hence, Alice mixes over E and H with probability $p_E - p_F$ and $2p_F$ respectively. As before we can define $\gamma_i = \Pr(c_A = a_i | H) = \frac{\frac{p_E - p_F}{p_E} u_{E_i} + u_{F_i}}{2p_F} = \frac{1}{2} \frac{u_{E_i}}{p_E} + \frac{1}{2} \frac{u_{F_i}}{p_F} = \frac{\alpha_i + \beta_i}{2}$. We can now define the myopic move of Bob given the market is at H and $c_B = b_j$ as $r_j^\gamma = \sum_i \gamma_i \Pr(Y | c_A = a_i, c_B = b_j)$. Similarly, we use \mathbf{r}_j^γ to represent the vector for two outcomes.

We now characterize $\pi^B(\sigma_1)$ as follows, writing p_E as $(p_E - p_F) + p_F$ to facilitate comparison with $\pi^B(\sigma'_1)$.

$$\begin{aligned} \pi^B(\sigma_1) &= (p_E - p_F) \left[\sum_j \Pr(c_B = b_j) D(\mathbf{r}_j^\alpha \| \sum_j \Pr(c_B = b_j) \mathbf{r}_j^\alpha) \right] \\ &\quad + p_F \left[\sum_j \Pr(c_B = b_j) D(\mathbf{r}_j^\alpha \| \sum_j \Pr(c_B = b_j) \mathbf{r}_j^\alpha) \right] \\ &\quad + p_F \left[\sum_j \Pr(c_B = b_j) D(\mathbf{r}_j^\beta \| \sum_j \Pr(c_B = b_j) \mathbf{r}_j^\beta) \right] \\ &\quad + \text{remaining profit over } R \end{aligned}$$

We also characterize $\pi^B(\sigma'_1)$ as:

$$\begin{aligned}\pi^B(\sigma'_1) &= (p_E - p_F) \left[\sum_j \Pr(c_B = b_j) D(\mathbf{r}_j^\alpha \parallel \sum_j \Pr(c_B = b_j) \mathbf{r}_j^\alpha) \right. \\ &\quad \left. + 2p_F \left[\sum_j \Pr(c_B = b_j) D(\mathbf{r}_j^\gamma \parallel \sum_j \Pr(c_B = b_j) \mathbf{r}_j^\gamma) \right] \right. \\ &\quad \left. + \text{remaining profit over } R \right.\end{aligned}$$

From the definitions of the myopic moves given the market states, note that $\mathbf{r}_j^\gamma = \frac{\mathbf{r}_j^\alpha + \mathbf{r}_j^\beta}{2}$. This means that $\pi^B(\sigma'_1)$ can be bounded as :

$$\begin{aligned}\pi^B(\sigma'_1) &= (p_E - p_F) \left[\sum_j \Pr(c_B = b_j) D(\mathbf{r}_j^\alpha \parallel \sum_j \Pr(c_B = b_j) \mathbf{r}_j^\alpha) \right. \\ &\quad \left. + 2p_F \left[\sum_j \Pr(c_B = b_j) D(\mathbf{r}_j^\gamma \parallel \sum_j \Pr(c_B = b_j) \mathbf{r}_j^\gamma) \right] \right. \\ &\quad \left. + \text{remaining profit over } R \right. \\ &= (p_E - p_F) \left[\sum_j \Pr(c_B = b_j) D(\mathbf{r}_j^\alpha \parallel \sum_j \Pr(c_B = b_j) \mathbf{r}_j^\alpha) \right. \\ &\quad \left. + 2p_F \left[\sum_j \Pr(c_B = b_j) D\left(\frac{\mathbf{r}_j^\alpha + \mathbf{r}_j^\beta}{2} \parallel \sum_j \Pr(c_B = b_j) \frac{\mathbf{r}_j^\alpha + \mathbf{r}_j^\beta}{2}\right) \right] \right. \\ &\quad \left. + \text{remaining profit over } R \right. \\ &< (p_E - p_F) \left[\sum_j \Pr(c_B = b_j) D(\mathbf{r}_j^\alpha \parallel \sum_j \Pr(c_B = b_j) \mathbf{r}_j^\alpha) \right] \\ &\quad + p_F \left[\sum_j \Pr(c_B = b_j) D(\mathbf{r}_j^\alpha \parallel \sum_j \Pr(c_B = b_j) \mathbf{r}_j^\alpha) \right] \\ &\quad + p_F \left[\sum_j \Pr(c_B = b_j) D(\mathbf{r}_j^\beta \parallel \sum_j \Pr(c_B = b_j) \mathbf{r}_j^\beta) \right] \\ &\quad + \text{remaining profit over } R \\ &= \pi^B(\sigma_1)\end{aligned}$$

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

Therefore, for any strategy σ_1 with two or more points in its support, there always exists a strategy σ'_1 such that $\pi^B(\sigma'_1) < \pi^B(\sigma_1)$. This means that for any strategy of Alice that minimized $\pi^B(\sigma_1)$, the strategy must have only one point in its support. Thus, the strategy does not reveal any information to Bob. Suppose that the point in the support, x , is such that $x \neq \Pr(\omega = Y)$. This again contradicts the fact that σ_1 minimizes $\pi^B(\sigma_1)$, as Bob will always make a positive payoff in expectation if he moves from x to $\Pr(\omega = Y)$. Thus, he would have a larger payoff overall if Alice left the market at x instead of $\Pr(\omega = Y)$. Therefore the strategy that minimizes the expected payoff of Bob is for Alice to report $\Pr(\omega = Y)$. However, because the market starts with $\Pr(\omega = Y)$, it is equivalent to her not trading at all in the first stage. Therefore we have shown that $\pi^B(\sigma_1) \geq \pi^B(\sigma^N)$. Only when $\sigma_1 = \sigma^N$ the equality holds. \square

The following corollary immediately follows from Theorem 2 and answers the question we proposed at the beginning of the section for I games.

Corollary 2 *In I games that start with the prior probabilities for the two outcomes, i.e. $\mathbf{r}^0 = \langle \Pr(\omega = Y), \Pr(\omega = N) \rangle$, when two players play myopically according to a predefined sequence of play, it is better off for a player to play second than first.*

Proof We compare Bob's expected profit in Alice-Bob and Bob-Alice cases when both players play myopically. Bob's expected profit in Alice-Bob case equals $\pi^B(\sigma^T)$, while his expected profit in Bob-Alice cases equals $\pi^B(\sigma^N)$. According to Theorem 2, $\pi^B(\sigma^T) \geq \pi^B(\sigma^N)$. Hence, Bob is better off when the Alice-Bob sequence is used. \square

5 Sequence Selection Game

In this section, we examine the following strategic question: Suppose that each player will get to report only once. If a player has a choice between playing first and playing second after observing its signal, which one should it choose?

To answer the question, we define a simple 2-player *sequence selection game*. Suppose that Alice and Bob are the only players in the market. Alice gets a signal c_A . Similarly, Bob gets a signal c_B . After getting their signals, Alice and Bob play the sequence selection game as follows. In the first stage, Alice chooses who—herself or Bob—plays first. The selected player then changes the market probabilities as they see fit in the second stage. In the third stage, the other player gets the chance to change the market probabilities. Then, the market closes and the true state is revealed.

In the rest of this section, we provide equilibria of the sequence selection game under CI setting and I setting respectively.

5.1 Sequence Selection Equilibrium of CI Games

Consider the sequence selection game, and assume that Alice and Bob have conditionally independent signals. The following theorem gives a PBE of the sequence selection game when the CI condition is satisfied.

Theorem 3 *When Alice and Bob have conditionally independent signals in LMSR, at a PBE of the sequence selection game*

- Alice selects herself to be the first player in the first stage;
- Alice changes the market probability to $\langle \Pr(Y|c_A), \Pr(N|c_A) \rangle$ in the second stage;
- Bob changes the market probability to $\langle \Pr(Y|c_A, c_B), \Pr(N|c_A, c_B) \rangle$ in the third stage.

Proof Since we use the PBE concept, we first describe Bob's belief μ_B for the off-equilibrium paths. We denote $\mathbf{x} = \langle x_1, \dots, x_{n_A} \rangle$ as the vector of Alice's possible posteriors for the outcome Y . That is, $x_i = \Pr(Y|c_A = a_i)$, where a_i is the i -th element of \mathbf{C}_A . Without loss of generality, assume that $x_i < x_j$ iff $i < j$. Off the equilibrium path, when Alice selects Bob to be the first player and the initial market probability is \mathbf{r}^0 , Bob's belief μ_B gives

- If $x_1 \leq r_Y^0 \leq x_{n_A}$, $\Pr(Y|\mathbf{r}^0, \mu_B) = r_Y^0$.
- If $0 \leq r_Y^0 < x_1$, $\Pr(Y|\mathbf{r}^0, \mu_B) = x_1$.
- If $x_{n_A} < r_Y^0 \leq 1$, $\Pr(Y|\mathbf{r}^0, \mu_B) = x_{n_A}$.

Bob's belief μ_B is consistent with Alice's off-equilibrium mixed strategy:

- If $c_A = a_i$ and $x_i = r_Y^0$, select Bob as the first player.
- If $c_A = a_i$ or $c_A = a_{i+1}$ and $x_i < r_Y^0 < x_{i+1}$, mix between selecting herself as the first player and selecting Bob as the first player. When $c_A = a_i$, with probability α select Bob as the first player and $(1 - \alpha)$ select herself as the first player. When $c_A = a_{i+1}$, with probability β select Bob as the first player and $(1 - \beta)$ select herself as the first player. The mixing probabilities α and β satisfy that $\Pr(Y|\text{select Bob as the first player}) = r_Y^0$.
- If $c_A = a_1$ and $0 \leq r_Y^0 < x_1$, select Bob as the first player.
- If $c_A = a_{n_A}$ and $x_{n_A} < r_Y^0 \leq 1$, select Bob as the first player.
- For all other cases, select herself as the first player.

Each player has only one chance to change the market probabilities. Hence, by Lemma 2, both of them will truthfully reveal all information that they have given their beliefs when it's their turn to play. If Alice is the first to play, she will change the market probabilities to $\langle \Pr(Y|c_A), \Pr(N|c_A) \rangle$ in the second stage. Bob, believing that prices in the second stage are Alice's posteriors, can calculate his posteriors based on both Alice's and his own signals. Bob will further change the market probabilities to $\langle \Pr(Y|c_A, c_B), \Pr(N|c_A, c_B) \rangle$ in the third stage. On the contrary, if Bob is selected as the first player, he will change the market probabilities to $\langle \Pr(Y|\mathbf{r}^0, c_B, \mu_B), \Pr(N|\mathbf{r}^0, c_B, \mu_B) \rangle$ in the second stage. Let $\hat{a}_{\mathbf{r}^0}$ denote a fictitious signal realization of Alice that satisfies $\langle \Pr(Y|\hat{a}_{\mathbf{r}^0}), \Pr(N|\hat{a}_{\mathbf{r}^0}) \rangle = \mathbf{r}^0$. $\hat{a}_{\mathbf{r}^0}$ does not necessarily belong to Alice's signal space \mathbf{C}_A . According to Lemma 4, $\langle \Pr(Y|\mathbf{r}^0, c_B, \mu_B), \Pr(N|\mathbf{r}^0, c_B, \mu_B) \rangle$ equals one of the following

1. $\langle \Pr(Y|\hat{a}_{\mathbf{r}^0}, c_B), \Pr(N|\hat{a}_{\mathbf{r}^0}, c_B) \rangle$ when $x_1 \leq r_Y^0 \leq x_{n_A}$.
2. $\langle \Pr(Y|a_{n_A}, c_B), \Pr(N|a_{n_A}, c_B) \rangle$ when $x_{n_A} < r_Y^0 \leq 1$.
3. $\langle \Pr(Y|a_1, c_B), \Pr(N|a_1, c_B) \rangle$ when $0 \leq r_Y^0 < x_1$.

To make her sequence selection in the first stage, Alice essentially compares her expected utilities conditional on her own signal in the Alice-Bob and Bob-Alice subgames.

Without loss of generality, suppose Alice has signal a_k . Let U_A^I denote Alice's expected profit conditional on her signal when the Alice-Bob subgame is picked. U_A^{II} denotes Alice's expected profit conditional on her signal when the Bob-Alice subgame is picked. Then, for case 1,

$$U_A^I = \frac{1}{\Pr(c_A = a_k)} \sum_{l,z} \Pr(a_k, b_l, z) \log \frac{\Pr(\omega = z|c_A = a_k)}{\Pr(\omega = z|\hat{a}_{\mathbf{r}^0})}, \quad \text{and} \quad (3)$$

$$U_A^{II} = \frac{1}{\Pr(c_A = a_k)} \sum_{l,z} \Pr(a_k, b_l, z) \log \frac{\Pr(\omega = z|c_A = a_k, c_B = b_l)}{\Pr(\omega = z|\hat{a}_{\mathbf{r}^0}, c_B = b_l)}, \quad (4)$$

where b_l represents the l -th signal element in \mathbf{C}_B , and $\Pr(a_k, b_l, z)$ represents the joint probability of $c_A = a_k$, $c_B = b_l$, and $\omega = z$. For conciseness, we use $\Pr(a_k)$

to represent $\Pr(c_A = a_k)$ and similarly for other terms. The difference in expected profits for Alice in the two subgames is:

$$U_A^I - U_A^{II} = \frac{1}{\Pr(a_k)} \sum_{l,z} \Pr(a_k, b_l, z) \log \frac{\Pr(z|a_k) \Pr(z|\hat{a}_{\mathbf{r}^0}, b_l)}{\Pr(z|\hat{a}_{\mathbf{r}^0}) \Pr(z|a_k, b_l)} \quad (5)$$

$$\begin{aligned} &= \frac{1}{\Pr(a_k)} \sum_{l,z} \Pr(a_k, b_l, z) \log \frac{\Pr(a_k, b_l) \Pr(\hat{a}_{\mathbf{r}^0})}{\Pr(\hat{a}_{\mathbf{r}^0}, b_l) \Pr(a_k)} \\ &= \sum_l \Pr(b_l|a_k) \log \frac{\Pr(b_l|a_k)}{\Pr(b_l|\hat{a}_{\mathbf{r}^0})} \\ &= D\left(\mathbf{p}_{(c_B|a_k)} \parallel \mathbf{p}_{(c_B|\hat{a}_{\mathbf{r}^0})}\right), \end{aligned} \quad (6)$$

where $\mathbf{p}_{(c_B|a_k)}$ is the probability distribution of Bob's signal conditional on signal a_k and $\mathbf{p}_{(c_B|\hat{a}_{\mathbf{r}^0})}$ is the probability distribution of Bob's signal conditional on the fictitious signal $\hat{a}_{\mathbf{r}^0}$. The second equality comes from Bayes rule and the conditional independence of signals. $D(\mathbf{p} \parallel \mathbf{q}) \geq 0$. The equality holds only when distributions \mathbf{p} and \mathbf{q} are the same. We thus have $U_A^I - U_A^{II} \geq 0$. When $\langle \Pr(Y|a_k), \Pr(N|a_k) \rangle \neq \mathbf{r}^0$, $\mathbf{p}_{(c_B|a_k)}$ is different from $\mathbf{p}_{(c_B|\hat{a}_{\mathbf{r}^0})}$ and hence $U_A^I - U_A^{II}$ is strictly greater than 0.

Since the market only has two exclusive outcomes, cases 2 and 3 are symmetric. We only need to prove one of them. Without loss of generality, consider case 2 where $x_{n_A} < r_Y^0 \leq 1$, Alice's expected profit conditioned on her signal for the Bob-Alice subgame becomes

$$\widetilde{U}_A^{II} = \frac{1}{\Pr(c_A = a_k)} \sum_{l,z} \Pr(a_k, b_l, z) \log \frac{\Pr(\omega = z|c_A = a_k, c_B = b_l)}{\Pr(\omega = z|c_A = a_{n_A}, c_B = b_l)}, \quad (7)$$

while U_A^I remains the same as in (3). We obtain

$$\begin{aligned} &U_A^I - \widetilde{U}_A^{II} \\ &= \frac{1}{\Pr(a_k)} \sum_{l,z} \Pr(a_k, b_l, z) \log \frac{\Pr(z|a_k)}{\Pr(z|a_{n_A})} + \sum_z \Pr(z|a_k) \log \frac{\Pr(z|a_{n_A})}{r_z^0} - \widetilde{U}_A^{II} \\ &= D\left(\mathbf{p}_{(c_B|a_k)} \parallel \mathbf{p}_{(c_B|a_{n_A})}\right) + \sum_z \Pr(z|a_k) \log \frac{\Pr(z|a_{n_A})}{r_z^0}. \end{aligned}$$

The second term in the above expression is positive because $r_Y^0 > \Pr(Y|a_{n_A}) \geq \Pr(Y|a_k)$. Hence, $U_A^I - \widetilde{U}_A^{II} > 0$. Alice does not want to deviate from selecting herself as the first player. \square

Corollary 1 in Section 4 indicates that if myopic players follow predefined sequence of play then on expectation a player is better off by playing first than second in CI games. Theorem 3 further strengthens this result and states that if a strategic player have a choice, at a PBE the player will always choose to play first in CI games no matter what signal it gets.

5.2 Sequence Selection Equilibrium of I games

In this part, we consider the alternative model in which players have independent signals. The following theorem gives a PBE of the sequence selection game when the I condition is satisfied.

Theorem 4 *If the market starts with the prior probabilities for the two outcomes, i.e. $\mathbf{r}^0 = \langle \Pr(\omega = Y), \Pr(\omega = N) \rangle$, and Alice and Bob have independent signals, at a PBE of the sequence selection game*

- Alice in the first stage selects Bob to be the first player;
- Bob changes the market probability to $\langle \Pr(Y|c_B), \Pr(N|c_B) \rangle$ in the second stage;
- Alice changes the market probability to $\langle \Pr(Y|c_A, c_B), \Pr(N|c_A, c_B) \rangle$ in the third stage.

Proof If Alice always selects Bob as the first player, this action does not reveal any information to Bob. Knowing Alice's strategy, according to Lemma 2, Bob's best responds in the second stage is to truthfully report the posterior probability conditioned on his own signal and Alice will report truthfully conditioned on both signals in the third stage. Thus, we need to prove that Alice does not want to deviate from selecting herself as the first player. Without loss of generality, assume Alice has signal a_k .

If Alice selects herself as the first player, her best response in the second stage is to report truthfully. Hence, the expected profit conditioned on her signal in this case, denoted as U_A^I , is

$$U_A^I = \sum_z \Pr(\omega = z | c_a = a_k) \log \frac{\Pr(\omega = z | c_a = a_k)}{\Pr(\omega = z)} = D(\mathbf{p}_{(\omega|a_k)} || \mathbf{p}_{(\omega)}) \quad (8)$$

If Alice selects Bob as the first player, her expected profit conditioned on her signal, denoted as U_A^{II} is

$$\begin{aligned} U_A^{II} &= \sum_{l,z} \Pr(c_b = b_l | a_k) \Pr(\omega = z | a_k, b_l) \log \frac{\Pr(\omega = z | a_k, b_l)}{\Pr(\omega = z | b_l)} \\ &= \sum_l \Pr(c_b = b_l | a_k) D(\mathbf{p}_{(\omega|a_k, b_l)} || \mathbf{p}_{(\omega|b_l)}) \end{aligned} \quad (9)$$

Because signals are independent, we have $\Pr(b_l | a_k) = \Pr(b_l)$. Hence, $\Pr(z | a_k) = \sum_l \Pr(b_l | a_k) \Pr(z | a_k, b_l) = \sum_l \Pr(b_l) \Pr(z | a_k, b_l)$. By Bayes rule, $\Pr(z) = \sum_l \Pr(b_l) \Pr(z | b_l)$. The difference in expected profits for Alice is:

$$\begin{aligned} U_A^I - U_A^{II} &= D(\mathbf{p}_{(\omega|a_k)} || \mathbf{p}_{(\omega)}) - \sum_l \Pr(b_l) D(\mathbf{p}_{(\omega|a_k, b_l)} || \mathbf{p}_{(\omega|b_l)}) \\ &= D\left(\sum_l \Pr(b_l) \mathbf{p}_{(\omega|a_k, b_l)} || \sum_l \Pr(b_l) \mathbf{p}_{(\omega|b_l)}\right) - \sum_l \Pr(b_l) D(\mathbf{p}_{(\omega|a_k, b_l)} || \mathbf{p}_{(\omega|b_l)}) \\ &\leq 0. \end{aligned}$$

The inequality follows from the convexity of relative entropy [11]. Thus, Alice does not want to deviate from selecting Bob as the first player.

The strategies and beliefs off the equilibrium path are straightforward. If Alice selects herself to be the first players. In the second stage, her best response is to report her true posterior. Bob believes that Alice reports truthfully. \square

Corollary 2 in Section 4 indicates that if myopic players follow predefined sequence of play then on expectation a player is better off by playing second than first in I games. Theorem 4 further strengthens this result and states that if a strategic player have a choice, at a PBE the player will always choose to play second in I games no matter what signal it gets.

6 The Alice-Bob-Alice Game

We now consider a 3-stage Alice-Bob-Alice game, where Alice plays in the first and third stages and Bob plays in the second stage. Alice may change the market probabilities however she wants in the first stage. Observing Alice's action, Bob may change the probabilities in the second stage. Alice can take another action in the third stage. Then, the market closes and the true state is revealed.

Let $\mathbf{r}^t = \langle r_Y^t, r_N^t \rangle$ be the market probabilities in stage t . Lemma 5 characterizes the equilibrium strategy of Alice in the third stage.

Lemma 5 *In a 3-stage Alice-Bob-Alice game in LMSR, at a WPBE Alice changes the market probabilities to $\mathbf{r}^3 = \langle r_Y^3, r_N^3 \rangle = \langle \Pr(Y|a_k, b_l), \Pr(N|a_k, b_l) \rangle$ in the third stage, when Alice has signal a_k and Bob has signal b_l .*

Proof This is proved by applying Lemma 2 to both Bob and Alice. At a WPBE, beliefs are consistent with strategies. Alice and Bob act as if they know each other's strategy. Since Bob only gets one chance to play, according to Lemma 2 Bob plays truthfully and fully reveals his information. Because the third stage is Alice's last chance to change the probabilities, Alice behaves truthfully and fully reveals her information, including her own signal and Bob's signal inferred from Bob's action in the second stage. \square

6.1 Truthful Equilibrium of CI Games

We study the PBE of the game when Alice and Bob have conditionally independent signals. Built on our equilibrium result for the sequence selection game, Theorem 5 describes a PBE of the Alice-Bob-Alice game.

Theorem 5 *When Alice and Bob have conditionally independent signals in LMSR, at a PBE of the 3-stage Alice-Bob-Alice game*

- Alice changes the market probability to $\mathbf{r}^1 = \langle \Pr(Y|c_A), \Pr(N|c_A) \rangle$ in the first stage;
- Bob in the second stage changes the market probability to $\mathbf{r}^2 = \langle \Pr(Y|c_A, c_B), \Pr(N|c_A, c_B) \rangle$;

– Alice does nothing in the third stage.

Proof Bob's belief μ_B is similar to that of the sequence selection game as described in the proof for Theorem 3. We denote $\mathbf{x} = \langle x_1, \dots, x_{n_A} \rangle$ as the vector of Alice's possible posteriors for the outcome Y . That is, $x_i = \Pr(Y|c_A = a_i)$, where a_i is the i -th element of \mathbf{C}_A . Without loss of generality, assume that $x_i < x_j$ iff $i < j$. When Alice changes market probability to \mathbf{r}^1 in the first stage, Bob's belief μ_B gives

- If $x_1 \leq r_Y^1 \leq x_{n_A}$, $\Pr(Y|\mathbf{r}^1, \mu_B) = r_Y^1$.
- If $0 \leq r_Y^1 < x_1$, $\Pr(Y|\mathbf{r}^1, \mu_B) = x_1$.
- If $x_{n_A} < r_Y^1 \leq 1$, $\Pr(Y|\mathbf{r}^1, \mu_B) = x_{n_A}$.

When observing a $r_Y^1 = y \neq x_i$ for all $1 \leq i \leq n_A$, Bob's belief μ_B is consistent with Alice's off-equilibrium mixed strategy:

- If $c_A = a_i$ or $c_A = a_{i+1}$ and $x_i < y < x_{i+1}$, mix between truthful betting and reporting $r_Y^1 = y$. When $c_A = a_i$, with probability α change market probability to $r_Y^1 = y$ and $(1 - \alpha)$ play truthful betting. When $c_A = a_{i+1}$, with probability β change market probability to $r_Y^1 = y$ and $(1 - \beta)$ play truthful betting. The mixing probabilities α and β satisfy that $\Pr(Y|r_Y^1 = y) = y$.
- If $c_A = a_1$ and $0 \leq y < x_1$, change market probability to $r_Y^1 = y$.
- If $c_A = a_{n_A}$ and $x_{n_A} < y \leq 1$, change market probability to $r_Y^1 = y$.
- For all other cases, report truthfully.

By Lemma 2, Bob does not want to deviate from truthful betting in the second stage given that Alice truthfully reports her posteriors in the first stage.

We show that Alice does not want to deviate by changing market probabilities to $\mathbf{r}^1 \neq \langle \Pr(Y|c_A), \Pr(N|c_A) \rangle$. Without loss of generality, assume that Alice has signal $c_A = a_k$. Consider the two cases:

1. *When Alice does not deviate:* Alice changes market probabilities to her true posteriors $\langle \Pr(Y|a_k), \Pr(N|a_k) \rangle$ in the first stage; Bob changes the probabilities to $\langle \Pr(Y|a_k, c_B), \Pr(N|a_k, c_B) \rangle$ in the second stage; Alice does nothing in the third stage.
2. *When Alice deviates:* Alice changes market probabilities to \mathbf{r}^1 that is different from $\langle \Pr(Y|a_k), \Pr(N|a_k) \rangle$; Bob changes market probabilities to $\langle \Pr(Y|\mathbf{r}^1, c_B, \mu_B), \Pr(N|\mathbf{r}^1, c_B, \mu_B) \rangle$ in the second stage, and Alice plays a best response according to Lemma 5 by changing market probabilities to $\langle \Pr(Y|a_k, c_B), \Pr(N|a_k, c_B) \rangle$ in the third stage.

We compare Alice's expected profits conditional on her signal in these two cases with the aid of the sequence selection game. The total payoff of two consecutive moves, in a market using the LMSR, is exactly the payoff of moving from the starting point to the end point of the second move. Hence, the expected profit that Alice gets from case 1 is the same as what she gets from the following sequence of actions: (a) Alice changes market probabilities to \mathbf{r}^1 that are different from her posteriors in the first stage; (b) A sequence selection game starts with initial market probabilities \mathbf{r}^1 ; Alice selects herself to be the first player; (c) Alice changes

market probabilities to $\langle \Pr(Y|a_k), \Pr(N|a_k) \rangle$; (d) Bob changes market probabilities to $\langle \Pr(Y|a_k, c_B), \Pr(N|a_k, c_B) \rangle$. Similarly, the expected profit that Alice gets from case 2 is the same as what she gets from the following sequence of actions: (a') Alice changes market probabilities to \mathbf{r}^1 that are different from her posteriors in the first stage; (b') A sequence selection game starts with initial market probabilities \mathbf{r}^1 ; Alice selects Bob to be the first player; (c') Bob changes market probabilities to $\langle \Pr(Y|\mathbf{r}^1, c_B, \mu_B), \Pr(N|\mathbf{r}^1, c_B, \mu_B) \rangle$; (d') Alice changes market probabilities to $\langle \Pr(Y|a_k, c_B), \Pr(N|a_k, c_B) \rangle$. Alice's expected profit from (a) is the same as that from (a'). But according to Theorem 3, Alice's expected profit from (b), (c), and (d) is greater than or equal to that from (b'), (c'), and (d'). Hence, Alice does not want to deviate from truthful betting. \square

A natural question to ask, after we know that Alice being truthful is a PBE, is whether there are other equilibria with Alice bluffing in the first stage. Theorem 6 says that there are not when players have informative and distinct signals.

Theorem 6 *When Alice and Bob have conditionally independent signals that are informative and distinct, the truthful betting PBE is the unique WPBE of the 3-stage Alice-Bob-Alice game.*

Proof According to Lemma 5, at a WPBE the market probability in the third stage always reveals information of both Alice and Bob. Thus, the total expected profit of Alice and Bob from the game is

$$\pi^{AB} = \sum_{k,l,z} \Pr(c_A = a_k, c_B = b_l, \omega = z) \log \frac{\Pr(\omega = z | c_A = a_k, c_B = b_l)}{r_z^0},$$

which is fixed given \mathbf{r}^0 and the joint probability distribution \mathcal{P} . Let σ be Alice's strategy at a WPBE. Then, Alice's expected profit by following σ is $\pi_\sigma^A = \pi^{AB} - \pi^B(\sigma)$, where $\pi^B(\sigma)$ is Bob's expected myopic optimal profit following a first round strategy σ as defined in Section 4. According to Theorem 1, $\pi^B(\sigma) \geq \pi^B(\sigma^T)$ and the equality holds only when $\sigma = \sigma^T$. Thus, for any WPBE strategy σ of Alice, at the equilibrium $\pi_\sigma^A \leq \pi^{AB} - \pi^B(\sigma^T)$, with equality holds only when $\sigma = \sigma^T$. However, any strategy σ such that $\pi_\sigma^A < \pi^{AB} - \pi^B(\sigma^T)$ can not be a WPBE strategy, because Alice can simply deviate to the truthful strategy and get higher expected profit. Hence, truthful betting is the only WPBE of the game. \square

6.2 Nonexistence of Truthful Equilibrium in I Games

We study the equilibrium of the game when Alice and Bob have independent signals and the general informativeness condition holds. In particular, we show that the truthful strategy in this setting is *not* a WPBE. Alice has an incentive to deviate from the truthful strategy on her first trade in the market.

The truthful strategy for Alice is the same as that in Section 6.1: change market probability to $\mathbf{r}^1 = \langle \Pr(Y|c_A = a_k), \Pr(N|c_A = a_k) \rangle$ in the first stage. We define

the bluffing strategy for Alice as: change market probability to $\tilde{\mathbf{r}}^1 = \langle \Pr(Y|c_A = \tilde{a}_k), \Pr(N|c_A = \tilde{a}_k) \rangle$, where $\tilde{a}_k \neq a_k$ is another signal in \mathbf{C}_A .

Since the total payoff of two consecutive moves in LMSR is exactly the payoff of moving from the starting point to the end point of the second move, we may now consider a bluffing strategy by Alice during the first trade as a move from \mathbf{r}^0 to \mathbf{r}^1 immediately followed by a move from \mathbf{r}^1 to $\tilde{\mathbf{r}}^1$. Due to the overlap in the expected score in the truthful strategy and the bluffing strategy for Alice, we eliminate the \mathbf{r}^0 to \mathbf{r}^1 move in our comparison analysis and treat the myopic move as if it had zero expected profit and analyze the bluffing strategy as a move from \mathbf{r}^1 to $\tilde{\mathbf{r}}^1$.

Lemma 6 *Assume that Bob expects Alice to play truthfully, and reacts accordingly. If Alice observes her signal, $c_A = a_k$, and bluffs, her expected change in profit over following the truthful strategy, is*

$$\Delta = \sum_l \Pr(c_B = b_l | c_A = a_k) D(\mathbf{p}_{(\omega|a_k, b_l)} || \mathbf{p}_{(\omega|\tilde{a}_k, b_l)}) - D(\mathbf{p}_{(\omega|a_k)} || \mathbf{p}_{(\omega|\tilde{a}_k)}),$$

where $\mathbf{p}_{(\omega|\cdot)}$ denotes the probability distribution of ω conditional on signal realizations.

Proof For conciseness, we use $\Pr(a_k)$ to represent $\Pr(c_A = a_k)$ and similarly for other terms. Alice's first trade is a move from \mathbf{r}^1 to $\tilde{\mathbf{r}}^1$, by the definition of LMSR her expected score from the first move is:

$$\Pr(Y|a_k) \log \frac{\Pr(Y|\tilde{a}_k)}{\Pr(Y|a_k)} + \Pr(N|a_k) \log \frac{\Pr(N|\tilde{a}_k)}{\Pr(N|a_k)} = -D(\mathbf{p}_{(\omega|a_k)} || \mathbf{p}_{(\omega|\tilde{a}_k)}).$$

As Bob has a probability of $\Pr(b_l|a_k)$ of seeing b_l as his signal, Alice will move the market from $\mathbf{r}^2 = \langle \Pr(Y|\tilde{a}_k, b_l), \Pr(N|\tilde{a}_k, b_l) \rangle$ to $\mathbf{r}^3 = \langle \Pr(Y|a_k, b_l), \Pr(N|a_k, b_l) \rangle$ with probability $\Pr(b_l|a_k)$. Therefore, the expected score of Alice's second trade in the market is:

$$\begin{aligned} & \sum_l \Pr(b_l|a_k) \left[\Pr(Y|a_k, b_l) \log \frac{\Pr(Y|a_k, b_l)}{\Pr(Y|\tilde{a}_k, b_l)} + \Pr(N|a_k, b_l) \frac{\Pr(N|a_k, b_l)}{\Pr(N|\tilde{a}_k, b_l)} \right] \\ &= \sum_l \Pr(b_l|a_k) D(\mathbf{p}_{(\omega|a_k, b_l)} || \mathbf{p}_{(\omega|\tilde{a}_k, b_l)}) \end{aligned}$$

Thus, Alice's total expected payoff conditional on her signal is the sum of the above two expressions, which gives Δ in our theorem. \square

Theorem 7 *If players have independent signals that satisfy the general informativeness condition, truthful betting is not a WPBE for the Alice-Bob-Alice game.*

Proof Lemma 6 gives the expected profit change of Alice if she were to deviate from truthful betting. Below we show that the expected profit change of deviating from the truthful strategy is greater than zero, thus Alice has an incentive to deviate from the truthful strategy. We have

$$\begin{aligned} & \sum_l \Pr(b_l|a_k) D(\mathbf{p}_{(\omega|a_k, b_l)} || \mathbf{p}_{(\omega|\tilde{a}_k, b_l)}) \\ & \geq D\left(\sum_l \Pr(b_l|a_k) \mathbf{p}_{(\omega|a_k, b_l)} || \sum_l \Pr(b_l|a_k) \mathbf{p}_{(\omega|\tilde{a}_k, b_l)}\right) \\ & = D(\mathbf{p}_{(\omega|a_k)} || \mathbf{p}_{(\omega|\tilde{a}_k)}) \end{aligned}$$

The inequality follows from the convexity of relative entropy [11]. We note that by the general informativeness condition the inequality is strict. The equality follows from the independence of the two signals, $\Pr(b_l|a_k) = \Pr(b_l)$, and the law of total probability. Therefore, $\Delta > 0$. Truthful betting is not a WPBE strategy. \square

6.2.1 Explicit Bluffing Equilibria in a Restricted Game We now introduce a simple model of independent signals and show that bluffing can be an equilibrium. In our model, Alice and Bob each see an independent coin flip and then participate in a LMSR prediction market with outcomes corresponding to whether or not both coins came up heads. Thus, $\omega \in \{\text{HH}, (\text{HT}|\text{TH}|\text{TT})\}$. We again consider an Alice-Bob-Alice game structure.

Theorem 8 *Consider the Alice-Bob-Alice LMSR coin-flipping game, where the probability of heads is p . Now restrict Alice's first stage strategies to either play truthful betting (TB) or as if her coin is heads (\hat{H}). A WPBE in this game has Alice play TB with probability $t = 1 + \frac{p}{(1-(1-p)^{-1/p})(1-p)}$, and otherwise play \hat{H} .*

Proof To show that playing TB with probability t is an equilibrium, we first compute Bob's best response to such a strategy and then show that Bob's strategy makes Alice indifferent between her pure strategies. Bob's best response is, if he has heads, to set the probability of HH to the probability that Alice has heads given that she bets as if she does (we denote Alice betting as if she had heads as \hat{H}):

$$\begin{aligned} \Pr(\text{HH} | \hat{\text{H}}\text{H}) &= \frac{\Pr(\hat{\text{H}}\text{H} | \text{HH}) \Pr(\text{HH})}{\Pr(\hat{\text{H}}\text{H} | \text{HH}) \Pr(\text{HH}) + \Pr(\hat{\text{H}}\text{H} | \text{TH}) \Pr(\text{TH})} \\ &= \frac{1 \cdot p^2}{1 \cdot p^2 + (1-t)(1-p)p} \\ &= 1 - (1-p)^{\frac{1}{p}}. \end{aligned} \tag{10}$$

If Bob has tails he sets the probability of HH to zero. Assuming such a strategy for Bob, we can compute Alice's expected profit for playing TB. This is done by computing, for each outcome in $\{\text{HH}, \text{HT}, \text{TH}, \text{TT}\}$, her profit for moving the probability from p_0 to p or 0, plus her profit for moving the probability to 0 or 1 from where Bob moves it. We have

$$\begin{aligned} &p^2 \left(\log \frac{p}{p_0} + \log \frac{1}{x} \right) + p(1-p) \log \frac{1-p}{1-p_0} + (1-p)p \log \frac{1-0}{1-p_0} \\ &+ (1-p)^2 \log \frac{1-0}{1-p_0} \end{aligned} \tag{11}$$

where x is Bob's posterior probability when he has heads and Alice appears to have heads. Similarly, Alice's expected profit for always pretending to have heads in the first stage is:

$$\begin{aligned} &p^2 \left(\log \frac{p}{p_0} + \log \frac{1}{x} \right) + p(1-p) \log \frac{1-p}{1-p_0} \\ &+ (1-p)p \left(\log \frac{1-p}{1-p_0} + \log \frac{1-0}{1-x} \right) + (1-p)^2 \log \frac{1-p}{1-p_0}. \end{aligned} \tag{12}$$

Since (11) and (12) are equal when x is set according to (10), Alice is indifferent between truthfulness and bluffing when Bob expects her to play TB with probability t . It is therefore in equilibrium for Alice to play TB with probability t , that is, Alice should, with $1 - t$ probability, pretend to have seen heads regardless of her actual information. \square

7 Finite-Player Finite-Stage Game

We extend our results for the Alice-Bob-Alice game to games with a finite number of players and finite stages in LMSR. Each player can change the market probabilities multiple times and all changes happen in sequence.

7.1 Truthful Equilibrium of CI games

We first consider the CI setting where players have conditionally independent signals. The following theorem characterizes the PBE of the game.

Theorem 9 *In the finite-player, finite-stage game with LMSR, if players have conditionally independent signals, at a PBE of the game all players report their posterior probabilities in their first round of play. If signals are informative and distinct, the truthful betting PBE is the unique WPBE of the game.*

Proof Given that every player believes that all players before it act truthfully, we prove the theorem recursively using backward induction. If it's player i 's last chance to play, it will truthfully report its posterior probabilities by Lemma 2. If it's player i 's second to last chance to play, there are other players standing in between its second to last chance to play and its last chance to play. We can combine the signals of those players standing in between as one signal and treat those players as one composite player. Because signals are conditionally independent, the signal of the composite player is conditionally independent of the signal of player i . The game becomes an Alice-Bob-Alice game for player i and at a PBE player i reports truthfully at its second to last chance to play according to Theorem 5. Inferring recursively, it is a PBE at which any player reports truthfully at its first chance to play.

If signals are informative and distinct, Theorem 6 states that the truthful betting PBE is the unique WPBE of the Alice-Bob-Alice game. This implies that truthful betting PBE is also a unique WPBE for the finite-player finite-stage game. Suppose the contrary, there is a bluffing WPBE. Select the last player who bluffs in the sequence. Without loss of generality, suppose the player is Alice and she bluffs at some stage t . Then, Alice must play truthfully at some stage $t' > t + 1$, and does not play in stages between t and t' . Combining signals of players between t and t' as one composite signal. The game from stage t to stage t' becomes an Alice-Bob-Alice game. Alice bluffing at stage t contradicts with the uniqueness of the WPBE of the Alice-Bob-Alice game. \square

7.2 No Truthful Equilibrium in I Games

For the I setting where players have independent signals, the result for finite-player finite-stage game directly follow from our result on the Alice-Bob-Alice game.

Theorem 10 *In the finite-player, finite-stage game with LMSR, if players have independent signals that satisfy the general informativeness condition, and there is at least one player who trades more than once, truthful betting is not a WPBE strategy.*

Proof Consider the last two trades of a player who plays more than once, for example Alice. If t and $t' > t + 1$ are the last two stages for Alice to trade and all other players play truthfully, we can combine signals of all players who trade between t and t' as a composite signal possessed by a composite agent. Because signals are independent, Alice's signal is independent of the composite signal. The game from stage t to stage t' can be viewed as an Alice-Bob-Alice game. By Theorem 7, we know that Alice is better off by bluffing at stage t . Hence, truthful betting is not a WPBE. \square

8 Finite-Player Infinite-Stage Games

In this section, we analyze a model in which the number of potential moves can be infinite. This is motivated by the observation that there is often no pre-defined last stage for a player. A player may decide to have one more trade at any time before the market closes.

We extend our results for finite-stage games to infinite-stage games and examine the equilibrium for CI and I settings respectively.

8.1 Truthful Betting Equilibrium in CI Games

Games with conditionally independent signals remain to have a unique equilibrium, truthful betting PBE, as is shown in the following theorem.

Theorem 11 *In a potentially infinite-stage game with conditionally independent signals, at a finite PBE all players truthfully report their posterior probabilities at their first chance of play. If signals are informative and distinct, the truthful betting equilibrium is the unique finite WPBE of the game.*

Proof If there are m players, truthful betting reveals all information after m stages. Knowing that other players are playing truthfully, a player who plays in stage $i < m$ does not want to deviate to bluffing in stage i and come to trade again in stage $m + 1$, because the game between stages i and $m + 1$ can be viewed as an Alice-Bob-Alice game. According to Theorem 5 player i would prefer to report truthfully.

If there exists a finite equilibrium that reveals all information at stage t , suppose the last time for a player to bluff is in stage i , then it must report truthfully in some

stage j and $i + 1 < j \leq t$. The game between stages i and j can be viewed as an Alice-Bob-Alice game. If signals are informative and distinct, according to Theorem 6, truthful betting equilibrium is the only WPBE of the Alice-Bob-Alice game. Contradiction arises. Truthful betting is the only finite equilibrium. \square

8.2 Nonexistence of Finite Equilibrium in I Games

Theorem 12 *In an infinite-stage game with independent signals that satisfy the general informativeness condition, there is no WPBE that reveals all information after finite number of trades.*

Proof For contradiction, consider any WPBE strategy profile σ that always results in the optimal market prediction after some number t of stages in a potentially infinite-stage game. Consider the last two players, Alice and Bob without loss of generality. They play in stages $t - 1$ and t respectively. Because the WPBE reveals all information, both Alice and Bob reports truthfully at stages $t - 1$ and t . Then, according to Theorem 7, Alice would want to deviate to bluff in stage $t - 1$ and play again in stage $t + 1$. Thus, σ can not be a WPBE strategy profile. There is not finite equilibrium. \square

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

9 Discounted LMSR and Entropy Reduction

For infinite-stage games, when players have independent signals, the market can not fully aggregate information of all players with finite trades. We propose a discounted logarithmic market scoring rule and show that the I games can converge to full information aggregation exponentially.

We first introduce the discounted LMSR mechanism in Section 9.1. Then, we characterize the infinite-stage games as continuation games in Section 9.2 and introduce the concepts of complementarity and substitution of signals in Section 9.3. We present a natural metric to quantify degree of information aggregation in Section 9.4. Facilitated with results in Sections 9.2, 9.3, and 9.4, we present our convergence analysis for discounted LMSR in Section 9.5.

9.1 Discounted LMSR

One way to address the incentives that traders with independent signals have to bluff in a market using the logarithmic market scoring rule is to reduce future payoffs using a discount parameter, perhaps resulting in an incentive to play the truthful strategy. Based on this intuition, we propose the discounted LMSR 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 just as they would in the regular market scoring rule. However, the value (positive or negative) of trade is discounted over time.² For simplicity, we assume a strict alternating sequence of trades between Alice and Bob. Suppose a trader moves the prediction for outcome Y from p to q in the i -th stage of the market, and the outcome Y 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 outcome Y did not happen, and the trader moves the prediction from p to q in the i -th stage the trader would earn a payoff of $\delta^{i-1}(\log(1 - q) - \log(1 - p))$. The regular market scoring rule corresponds to $\delta = 1$.

Clearly, the myopic strategic properties of the market scoring rule are retained in the discounted form. We note that for CI games, truthful betting is still the unique WPBE in the discounted LMSR. Intuitively, for an Alice-Bob-Alice CI game in the discounted LMSR, Alice does not want to deviate from the truthful betting strategy as in the non-discounted LMSR because for any deviating strategy her expected profit in the first stage is the same in both the non-discounted LMSR and the discounted LMSR but her expected profit in the third stage in the discounted LMSR is strictly less than that in the non-discounted LMSR, making deviation even more undesirable. As truthful betting is the unique WPBE in an Alice-Bob-Alice CI game in the discounted LMSR, the same arguments in Sections 7.1 and 8.1 can show that truthful betting is the unique WPBE for finite-stage and infinite-stage CI games with finite players in the discounted LMSR. Thus, from now on, we only focus on I games in the discounted LMSR.

We will show that the discounted LMSR can have better non-myopic strategic properties for I games. For simplicity, we assume that there are only two players, Alice and Bob. They alternately play in the market, with Alice being the first player. Even though that the discounted LMSR may admit equilibria in which players bluff with some probability, we will show that, in any WPBE strategy profile σ , the market price will converge towards the optimal probability 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 at any WPBE.

9.2 Continuation Games

It is useful to think about an infinite-stage market game G as a first-stage move followed by a subsequent game H that perhaps depends on the first-stage move. This decomposition is possible in equilibrium analysis, because the actual first-stage strategy σ_1 is common knowledge in an equilibrium profile. Then, the beliefs of the players after observing the first move x are consistent with private signals

² It is possible that traders intrinsically discount future profits; for example, they may be uncertain about how long their information will remain private. We do not, however, assume any intrinsic discounting by traders; rather, we are proposing a mechanism that explicitly discounts gains and losses in later rounds.

and a common prior \mathbf{p}' that depends on the initial prior \mathbf{p}^0 , the first-stage strategy σ_1 and the observed move x .

Definition 3 We say that H is a continuation game of G if there is a first-stage strategy σ_1 and a move x such that x is played with nonzero probability under σ , the prior \mathbf{p}' in H is consistent with \mathbf{p}^0 , σ_1 , and x , and the sequence of play in H is the same as the subsequence of play in G after the first move in G . We denote this $G \rightarrow H$.

We define the relation $\xrightarrow{*}$ as the reflexive transitive closure of \rightarrow , i.e., we write $G \xrightarrow{*} H$ and say G reduces to H if there is a sequence $G = H_0, H_1, H_2, \dots, H_k = H$ such that $H_i \rightarrow H_{i+1}$ for all i . We define $G \xrightarrow{*} G$ to always be true.

Observe that if $G \xrightarrow{*} H$, then, the full-information optimal probabilities $\Pr(Y|c_A = a, c_B = b)$ for any given pair of signal realizations (a, b) are the same in G and H . Only the prior over different signals changes after a first-round strategy, not the full-information posterior probabilities.

Lemma 7 shows that continuation games preserve the independence property of the original game.

Lemma 7 Suppose $G \xrightarrow{*} H$. Then,

- If G satisfies the conditional independence (CI) property, so does H .
- If G satisfies the independence (I) property, so does H .

Proof We need to show that if $H_i \rightarrow H_{i+1}$, H_{i+1} satisfies the same independence property as H_i .

Let σ be a WPBE strategy for game H_i . Without loss of generality, suppose that x is one of the points that the first player in H_i plays with positive probability according to σ . Let \mathcal{P}_i be the joint probability distribution for H_i . We use $\mathbf{P}_i\{\cdot\}$ to represent probability of some outcome under \mathcal{P}_i . Let $\mathbf{P}_i\{c_1 = k|x, \sigma\}$ be the probability that player 1 has signal k conditional on that it plays x . Then, knowing player 1 follows σ , player 2 updates its beliefs after seeing x , the continuation game H_{i+1} has joint probability distribution \mathcal{P}_{i+1} such that

$$\begin{aligned} \mathbf{P}_{i+1}\{c_1 = k, c_2 = l, \omega = z\} &= \mathbf{P}_i\{c_1 = k, c_2 = l, \omega = z|x, \sigma\} \\ &= \mathbf{P}_i\{c_1 = k|x, \sigma\} \mathbf{P}_i\{c_2 = l, \omega = z|c_1 = k\} \\ &= \frac{\mathbf{P}_i\{c_1 = k|x, \sigma\}}{\mathbf{P}_i\{c_1 = k|\sigma\}} \mathbf{P}_i\{c_1 = k, c_2 = l, \omega = z\} \end{aligned}$$

where k , l , and z represent signal realization of player 1, signal realization of player 2, and realized outcome respectively.

We first show that if \mathcal{P}_i of game H_i satisfies CI condition, \mathcal{P}_{i+1} of game H_{i+1} also satisfies CI condition.

$$\begin{aligned} \mathbf{P}_{i+1}\{c_1 = k, c_2 = l | \omega = z\} &= \frac{\mathbf{P}_{i+1}\{c_1 = k, c_2 = l, \omega = z\}}{\mathbf{P}_{i+1}\{\omega = z\}} \\ &= \frac{\frac{\mathbf{P}_i\{c_1=k|x,\sigma\}}{\mathbf{P}_i\{c_1=k|\sigma\}} \mathbf{P}_i\{c_1 = k, c_2 = l, \omega = z\}}{\mathbf{P}_{i+1}\{\omega = z\}} \\ &= \frac{\frac{\mathbf{P}_i\{c_1=k|x,\sigma\}}{\mathbf{P}_i\{c_1=k|\sigma\}} \mathbf{P}_i\{c_1 = k, c_2 = l | \omega = z\}}{\mathbf{P}_{i+1}\{\omega = z\} / \mathbf{P}_i\{\omega = z\}} \\ &= \frac{\frac{\mathbf{P}_i\{c_1=k|x,\sigma\}}{\mathbf{P}_i\{c_1=k|\sigma\}} \mathbf{P}_i\{c_1 = k | \omega = z\} \mathbf{P}_i\{c_2 = l | \omega = z\}}{\mathbf{P}_{i+1}\{\omega = z\} / \mathbf{P}_i\{\omega = z\}}. \end{aligned}$$

$$\begin{aligned} \mathbf{P}_{i+1}\{c_1 = k | \omega = z\} &= \frac{\frac{\mathbf{P}_i\{c_1=k|x,\sigma\}}{\mathbf{P}_i\{c_1=k|\sigma\}} \mathbf{P}_i\{c_1 = k, \omega = z\}}{\mathbf{P}_{i+1}\{\omega = z\}} \\ &= \frac{\frac{\mathbf{P}_i\{c_1=k|x,\sigma\}}{\mathbf{P}_i\{c_1=k|\sigma\}} \mathbf{P}_i\{c_1 = k | \omega = z\}}{\mathbf{P}_{i+1}\{\omega = z\} / \mathbf{P}_i\{\omega = z\}}. \end{aligned}$$

$$\begin{aligned} \mathbf{P}_{i+1}\{c_2 = l | \omega = z\} &= \frac{\sum_k \frac{\mathbf{P}_i\{c_1=k|x,\sigma\}}{\mathbf{P}_i\{c_1=k|\sigma\}} \mathbf{P}_i\{c_1 = k, c_2 = l, \omega = z\}}{\mathbf{P}_{i+1}\{\omega = z\}} \\ &= \frac{\sum_k \frac{\mathbf{P}_i\{c_1=k|x,\sigma\}}{\mathbf{P}_i\{c_1=k|\sigma\}} \mathbf{P}_i\{c_1 = k, c_2 = l | \omega = z\}}{\mathbf{P}_{i+1}\{\omega = z\} / \mathbf{P}_i\{\omega = z\}} \\ &= \frac{\sum_k \frac{\mathbf{P}_i\{c_1=k|x,\sigma\}}{\mathbf{P}_i\{c_1=k|\sigma\}} \mathbf{P}_i\{c_1 = k, c_2 = l | \omega = z\}}{\sum_k \frac{\mathbf{P}_i\{c_1=k|x,\sigma\}}{\mathbf{P}_i\{c_1=k|\sigma\}} \mathbf{P}_i\{c_1 = k, \omega = z\} / \mathbf{P}_i\{\omega = z\}} \\ &= \mathbf{P}_i\{c_2 = l | \omega = z\}. \end{aligned}$$

Thus, we also have

$$\mathbf{P}_{i+1}\{c_1 = k, c_2 = l | \omega = z\} = \mathbf{P}_{i+1}\{c_1 = k | \omega = z\} \mathbf{P}_{i+1}\{c_2 = l | \omega = z\}.$$

CI condition holds for H_{i+1} .

We then show that if \mathcal{P}_i of game H_i satisfies I condition, \mathcal{P}_{i+1} of game H_{i+1} also satisfies I condition.

$$\begin{aligned} \mathbf{P}_{i+1}\{c_1 = k, c_2 = l\} &= \sum_{\omega} \mathbf{P}_{i+1}\{c_1 = k, c_2 = l, \omega = z\} \\ &= \sum_{\omega} \frac{\mathbf{P}_i\{c_1=k|x,\sigma\}}{\mathbf{P}_i\{c_1=k|\sigma\}} \mathbf{P}_i\{c_1 = k, c_2 = l, \omega = z\} \\ &= \frac{\mathbf{P}_i\{c_1 = k | x, \sigma\}}{\mathbf{P}_i\{c_1 = k | \sigma\}} \mathbf{P}_i\{c_1 = k, c_2 = l\} \\ &= \frac{\mathbf{P}_i\{c_1 = k | x, \sigma\}}{\mathbf{P}_i\{c_1 = k | \sigma\}} \mathbf{P}_i\{c_1 = k | \sigma\} \mathbf{P}_i\{c_2 = l\} \\ &= \mathbf{P}_i\{c_1 = k | x, \sigma\} \mathbf{P}_i\{c_2 = l\} \end{aligned}$$

The second to the last equality is because of the independence condition of H_i . We also have

$$\begin{aligned}
\mathbb{P}_{i+1}\{c_1 = k\} &= \sum_l \mathbb{P}_{i+1}\{c_1 = k, c_2 = l\} \\
&= \sum_l \frac{\mathbb{P}_i\{c_1 = k|x, \sigma\}}{\mathbb{P}_i\{c_1 = k|\sigma\}} \mathbb{P}_i\{c_1 = k, c_2 = l\} \\
&= \sum_l \frac{\mathbb{P}_i\{c_1 = k|x, \sigma\}}{\mathbb{P}_i\{c_1 = k|\sigma\}} \mathbb{P}_i\{c_1 = k|\sigma\} \mathbb{P}_i\{c_2 = l\} \\
&= \mathbb{P}_i\{c_1 = k|x, \sigma\}
\end{aligned}$$

and

$$\begin{aligned}
\mathbb{P}_{i+1}\{c_2 = l\} &= \sum_k \mathbb{P}_{i+1}\{c_1 = k, c_2 = l\} \\
&= \sum_k \frac{\mathbb{P}_i\{c_1 = k|x, \sigma\}}{\mathbb{P}_i\{c_1 = k|\sigma\}} \mathbb{P}_i\{c_1 = k, c_2 = l\} \\
&= \sum_k \frac{\mathbb{P}_i\{c_1 = k|x, \sigma\}}{\mathbb{P}_i\{c_1 = k|\sigma\}} \mathbb{P}_i\{c_1 = k|\sigma\} \mathbb{P}_i\{c_2 = l\} \\
&= \mathbb{P}_i\{c_2 = l\}.
\end{aligned}$$

Hence, $\mathbb{P}_{i+1}\{c_1 = k, c_2 = l\} = \mathbb{P}_{i+1}\{c_1 = k\} \mathbb{P}_{i+1}\{c_2 = l\}$. H_{i+1} satisfies independence.

By induction, if $G \xrightarrow{*} H$, then H preserves the same independence property as G . \square

9.3 Complementarity or Substitution of Signals

In this section, we introduce the concept of complementarity for signals. Intuitively, if signals of Alice and Bob are complements, the expected profit from knowing both signals is greater than the sum of the expected profits from knowing only the individual signals. If signals are substitutes, the relation is reversed.

For a game G participated by Alice and Bob, the total profit potential in G is $\pi^{AB} = \pi^A(\sigma^N) + \pi^B(\sigma^T)$. In other words, π^{AB} is the sum of Alice's and Bob's expected profit if they both follow the truthful strategy, and Alice moves first. We note that under either the CI condition or the I condition, the order of play is not important, i.e., $\pi^{AB} = \pi^B(\sigma^N) + \pi^A(\sigma^T) = H(\omega|c_A, c_B)$.

Now we define the complementarity coefficient of a game that captures the complementarity of players' signals.

Definition 4 The complementarity coefficient $C(G)$ of a market game G is defined as

$$C(G) = \frac{\pi^A(\sigma^N) + \pi^B(\sigma^N)}{\pi^{AB}},$$

where the right-hand side profits are implicitly a function of the game G under consideration.

The complementarity bound $\hat{C}_\sigma(G)$ is the minimum value of $C(H)$ over all games H that G could reduce to by following the strategy profile σ :

$$\hat{C}_\sigma(G) = \min_{H: G \xrightarrow{\sigma} H} C(H)$$

We note that if $C(G) < 1$, signals are complements. If $C(G) > 1$, signals are substitutes.

Lemma 8 *Under the CI model, any game G with informative and distinct signals satisfies $\hat{C}_\sigma(G) > 1$ for any strategy profile σ .*

Proof From Theorem 1, it follows that $\pi^B(\sigma^N) > \pi^B(\sigma^T)$ for any CI game G with informative and distinct signals. Thus,

$$\pi^A(\sigma^N) + \pi^B(\sigma^N) > \pi^A(\sigma^N) + \pi^B(\sigma^T) = \pi^{AB} \Rightarrow C(G) > 1.$$

By lemma 7, we know that any H such that $G \xrightarrow{*} H$ itself satisfies the CI condition no matter what strategy the reduction follows. Thus, we must have $\hat{C}_\sigma(G) > 1$ for any σ . \square

Lemma 9 *Under the I model, any game G that starts with the prior probability satisfies $C(G) < 1$, and thus, $\hat{C}_\sigma(G) < 1$.*

Proof From Theorem 2, it follows that $\pi^B(\sigma^T) > \pi^B(\sigma^N)$ for any I game G starting with the prior probability. Thus,

$$\pi^A(\sigma^N) + \pi^B(\sigma^N) < \pi^A(\sigma^N) + \pi^B(\sigma^T) = \pi^{AB} \Rightarrow C(G) < 1.$$

By definition of \hat{C} , we have $\hat{C}_\sigma(G) \leq C(G)$, and thus we must have $\hat{C}_\sigma(G) < 1$. \square

9.4 Profit Potential and Entropy

We now present a natural metric, the *information deficit* \mathcal{D}^i , that quantifies the degree of aggregation in a prediction market after i trades under 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 probability (given the realization of the signals) and the actual probability after i stages.

Formally, for a strategy profile σ and a number of stages t , an *information node* ϕ consists of a realization of the signals of the two players and a sequence of t trades in the market. Let $p^i(\phi) = \Pr(Y|\text{first } i \text{ trades in } \phi, \sigma)$, and denote $\mathbf{p}^i(\phi) = \langle p^i(\phi), 1 - p^i(\phi) \rangle$. Let $p^*(\phi) = \Pr(Y|\text{realization of signals in } \phi)$, and denote $\mathbf{p}^*(\phi) = \langle p^*(\phi), 1 - p^*(\phi) \rangle$. The aggregative effect of the strategy profile σ after i trades is summarized by the collection of information nodes ϕ that can be reached, the market probability $\mathbf{p}^i(\phi)$ in stage i , and the associated ex-ante

probability $\Pr(\phi)$ of reaching each such information node. Now for $i \leq t$, we define

$$\begin{aligned} \mathcal{D}^i(\sigma) &= E_\sigma[D(\mathbf{p}^*(\phi) || \mathbf{p}^i(\phi))] \\ &= \sum_{\phi: \phi \in \sigma} \Pr(\phi) D(\mathbf{p}^*(\phi) || \mathbf{p}^i(\phi)) . \end{aligned}$$

When $t = 0$, the information nodes ϕ correspond to different realizations of the signals. Note that \mathcal{D}^0 is the same as π^{AB} defined in the previous section.

If $\mathcal{D}^i = 0$, it implies that the market will always have reached its optimal probability for the realized signals by the i -th stages. If $\mathcal{D}^i > 0$, it indicates that, with positive probability, the market has not yet reached the optimal probability. \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 information 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 i -th stage move in the *non-discounted* (standard) LMSR:

Lemma 10 *Let M^i denote the expected profit (over all signal nodes ϕ^i) of the i -th trade under profile σ . Then, $M^i = \mathcal{D}^{i-1} - \mathcal{D}^i$*

Proof

$$\begin{aligned} M^i &= \sum_{\phi: \phi \in \sigma} \Pr(\phi) \left[p^*(\phi) \log \frac{p^i(\phi)}{p^{i-1}(\phi)} + (1 - p^*(\phi)) \log \frac{1 - p^i(\phi)}{1 - p^{i-1}(\phi)} \right] \\ &= \sum_{\phi: \phi \in \sigma} \Pr(\phi) [D(\mathbf{p}^* || \mathbf{p}^{i-1}) - D(\mathbf{p}^* || \mathbf{p}^i)] \\ &= \mathcal{D}^{i-1} - \mathcal{D}^i \end{aligned}$$

The 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 i -th trade.

Given the definition of M^i , we can now define \tilde{M}^i as the expected profit of the i -th trade in the discounted LMSR. We assume discounting after every even trade, *i.e.*, after every trade by Bob.

$$\tilde{M}^i = \begin{cases} \delta^k (\mathcal{D}^{2k} - \mathcal{D}^{2k+1}) & \forall i = 2k + 1 \\ \delta^k (\mathcal{D}^{2k+1} - \mathcal{D}^{2(k+1)}) & \forall i = 2k + 2 \end{cases} \quad (13)$$

Using the definition of \tilde{M}^i we write the total profit for Alice in the δ -discounted LMSR as:

$$\text{Alice's payoff} = \sum_{i: i=2k, k \in \mathbb{Z}^+} \tilde{M}^{i-1} = \mathcal{D}^0 - \mathcal{D}^1 + \delta(\mathcal{D}^2 - \mathcal{D}^3) + \dots$$

We can similarly rewrite Bob's expected payoff as:

$$\text{Bob's payoff} = \sum_{i: i=2k, k \in \mathbb{Z}^+} \tilde{M}^i = \mathcal{D}^1 - \mathcal{D}^2 + \delta(\mathcal{D}^3 - \mathcal{D}^4) + \dots$$

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 .

9.5 Asymptotic Convergence in Discounted LMSR for I Games

In this part, we bound the value of \mathcal{D}^n for large n , in any WPBE, for I games that start with the prior probability. We will express our convergence bound in terms of complementarity coefficient, $C(G)$. Note that if $C(G) = 0$, 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 $C(G) > 0$.

Now, fix a discount parameter δ , and a particular instance of the two-person market game induced by the δ -discounted LMSR. Let σ be any WPBE of this market game. Without loss of generality we assume that Bob moves second in the market.

In the first stage, Alice will take some (perhaps mixed) action x dictated by σ . Under the equilibrium strategy profile σ , Bob will revise his beliefs consistent with the profile σ and the realized action x . According to Theorem 2, we have $\pi^B(\sigma) \geq \pi^B(\sigma^N)$.

Recall that after stage 2, the total expected payoff of both players is at most $\delta\mathcal{D}^2$. We also know that the total expected payoff of Alice in equilibrium is at least $\pi^A(\sigma^N)$, because if not, a simple deviation to the truthful strategy would be beneficial. Because $\pi^B(\sigma) \geq \pi^B(\sigma^N)$, the total expected payoff of Bob in equilibrium is also at least $\pi^B(\sigma^N)$. This means that the total payoff of the first two rounds in the market is at least $\pi^A(\sigma^N) + \pi^B(\sigma^N) - \delta\mathcal{D}^2$. Therefore we can bound \mathcal{D}^2 as:

$$\mathcal{D}^2 \leq \mathcal{D}^0 - [\pi^A(\sigma^N) + \pi^B(\sigma^N) - \delta\mathcal{D}^2].$$

Note that $\mathcal{D}^0 = \pi^{AB}$.

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 $\delta^k\mathcal{D}^{2k}$:

$$\begin{aligned} \mathcal{D}^{2k} &\leq \mathcal{D}^0 - [\pi^A(\sigma^N) + \pi^B(\sigma^N) - \delta^k\mathcal{D}^{2k}] & (14) \\ \iff (1 - \delta^k)\mathcal{D}^{2k} &\leq \mathcal{D}^0 - [\pi^A(\sigma^N) + \pi^B(\sigma^N)] \\ &= \mathcal{D}^0 \left(1 - \frac{\pi^A(\sigma^N) + \pi^B(\sigma^N)}{\mathcal{D}^0} \right) \end{aligned}$$

$$\iff \mathcal{D}^{2k} \leq \frac{\mathcal{D}^0(1 - C(G))}{1 - \delta^k} \quad (15)$$

Note that, by definition, $C(G) \geq \hat{C}_\sigma(G)$. Thus, we can rewrite inequality (15) as:

$$\mathcal{D}^{2k} \leq \frac{\mathcal{D}^0(1 - \hat{C}_\sigma(G))}{1 - \delta^k} \quad (16)$$

From inequality (16) we see that the bound on \mathcal{D}^{2k} depends only on δ , as \mathcal{D}^0 and $\hat{C}_\sigma(G)$ are both constants for any instance of the market and equilibrium profile σ . Now consider the remainder of the game after $2k$ rounds, denoted as H_{2k} , then $G \xrightarrow{*} H_{2k}$. According to Lemma 7, H_{2k} preserves the independence property of G . Hence, the equilibrium profile σ will also induce an equilibrium profile on H_{2k} . We can repeat the argument to bound \mathcal{D}^{4k} in terms of \mathcal{D}^{2k} , etc.

In this way, we rewrite inequality (16) in terms of δ and for a round $n = 2km$. We set k such that $\delta^{k/2} = \hat{C}_\sigma(G)$, i.e. $k = \frac{2 \log \hat{C}_\sigma(G)}{\log \delta}$. Using this value of k and a value $m = \frac{n}{2k}$ we note that:

$$\begin{aligned}
\mathcal{D}^n &\leq \mathcal{D}^0 \left(\frac{1 - \hat{C}_\sigma(G)}{1 - \delta^k} \right)^m \\
&= \mathcal{D}^0 \left(\frac{1 - \hat{C}_\sigma(G)}{(1 - \delta^{k/2})(1 + \delta^{k/2})} \right)^m \\
&= \mathcal{D}^0 (1 + \delta^{k/2})^{-m} \\
&= \mathcal{D}^0 (1 + \hat{C}_\sigma(G))^{-\frac{n}{2k}} \\
&= \mathcal{D}^0 (1 + \hat{C}_\sigma(G))^{-\frac{n \log \delta}{4 \log \hat{C}_\sigma(G)}} \\
&= \mathcal{D}^0 \delta^{n \frac{-\log(1 + \hat{C}_\sigma(G))}{4 \log \hat{C}_\sigma(G)}}
\end{aligned} \tag{17}$$

Note that $\frac{-\log(1 + \hat{C}_\sigma(G))}{4 \log \hat{C}_\sigma(G)}$ depends only on the complementarity bound of the game. According to Lemma 9, for any game G with independent signal, $\hat{C}_\sigma(G) < 1$. Hence, $\frac{-\log(1 + \hat{C}_\sigma(G))}{4 \log \hat{C}_\sigma(G)} > 0$. Therefore, inequality (17) shows that the relative entropy of the probabilities with respect to the optimal probabilities reduces exponentially over time.

To summarize, we have proven the following result:

Theorem 13 *Let G be a market game in the I model, with a δ -discounted logarithmic market scoring rule. For this game, let σ be any WPBE strategy profile. Then, for all $n > 1$, the information deficit \mathcal{D}^n after n rounds of trading is bounded by*

$$\mathcal{D}^n \leq \mathcal{D}^0 \delta^{n \frac{-\log(1 + \hat{C}_\sigma(G))}{4 \log \hat{C}_\sigma(G)}}$$

□

10 Conclusion

We have investigated the strategic behavior of traders in LMSR prediction markets using dynamic games with incomplete information. Specifically, we examine different scenarios where traders at equilibrium bet truthfully or bluff.

Two different information structures, conditional independence and unconditional independence of signals, are considered. Both finite-stage and infinite-stage games are investigated. We show that traders with conditionally independent signals may be worse off by bluffing in LMSR. Moreover, for both finite-stage and infinite-stage games, truthful betting is not only a PBE strategy but also a unique WPBE strategy for all traders in LMSR.

On the other hand, when the signals of traders are unconditionally independent, truthful betting is not a WPBE for finite-stage games. Traders have incentives to strategically mislead other traders with the intent of correcting the errors made by others in a later stage. We provide a bluffing equilibrium for a restricted two-player game. For infinite-stage games with unconditionally independent signals, there does not exist any WPBE that fully reveals information of all players in finite stages.

We use a simple modification to the market scoring rule, which includes a form of discounting, to ameliorate this potential problem. This allows us to prove a bound on the rate at which the error of the market, as measured by the relative entropy between full information aggregation and the actual market probability, reduces exponentially over time. The exponent depends on the complementarity coefficient of the market instance.

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 profit from the market. Moreover, players can pool their information together to make sharper predictions than either could alone and thus extract an even larger profit. They might engage in bluffing strategies to bargain over how this subsidy is divided. Explicit discounting can make this bargaining more efficient.

Acknowledgments

Stanko Dimitrov was supported by NSF grants CCF-0728768 and IGERT-0114368. Rahul Sami was partially supported by NSF grant CCF-0728768.

We are grateful to Paul J. Healy and Mohammad Mahdian for helpful discussions. We thank John Ledyard for pointing us to earlier work on discounting in parimutuel markets. We thank anonymous reviewers for their valuable comments.

References

1. F. Allen and D. Gale. Stock-price manipulation. *The Review of Financial Studies*, (5):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. *Journal of Economic Behavior and Organization*, 69:170–181, 2009.
3. J. E. Berg, R. Forsythe, F. D. Nelson, and T. A. Rietz. Results from a dozen years of election futures markets research. In C. A. Plott and V. Smith, editors, *Handbook of Experimental Economics Results*. North Holland, Amsterdam, 2008.
4. J. E. Berg and T. A. Rietz. Prediction markets as decision support systems. *Information Systems Frontier*, 5:79–93, 2003.
5. T. Börgers, A. Hernando-Veciana, and D. Krämer. When are signals complements or substitutes? Working paper, 2007.
6. C. F. Camerer. Can asset markets be manipulated? A field experiment with race-track betting. *Journal of Political Economy*, (106):457–482, 1998.

7. A. Chakraborty and B. Yilmaz. Manipulation in market order models. *Journal of Financial Markets*, 7(2):187–206, 2004.
8. K.-Y. Chen and C. R. Plott. Information aggregation mechanisms: Concept, design and implementation for a sales forecasting problem. Working paper No. 1131, California Institute of Technology, Division of the Humanities and Social Sciences, 2002.
9. Y. Chen and D. M. Pennock. A utility framework for bounded-loss market makers. In *Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence (UAI 2007)*, pages 49–56, Vancouver, BC Canada, 2007.
10. Y. Chen, D. M. Reeves, D. M. Pennock, R. D. Hanson, L. Fortnow, and R. Gonen. Bluffing and strategic reticence in prediction markets. In *Lecture Notes in Computer Science, The Third Workshop on Internet and Network Economics (WINE 2007)*, pages 70–81, 2007.
11. T. M. Cover and J. A. Thomas. *Elements of Information Theory*. John Wiley & Sons, Inc., 1991.
12. S. Debnath, D. M. Pennock, C. L. Giles, and S. Lawrence. Information incorporation in online in-game sports betting markets. In *Proceedings of the Fourth Annual ACM Conference on Electronic Commerce (EC'03)*, San Diego, CA, 2003.
13. S. Dimitrov and R. Sami. Non-myopic strategies in prediction markets. In *Proceedings of the Ninth ACM Conference on Electronic Commerce (EC'08)*, New York, July 2008. ACM.
14. E. F. Fama. Efficient capital market: A review of theory and empirical work. *Journal of Finance*, 25:383–417, 1970.
15. J. Feigenbaum, L. Fortnow, D. M. Pennock, and R. 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).
16. R. Forsythe, F. Nelson, G. R. Neumann, and J. Wright. Anatomy of an experimental political stock market. *American Economic Review*, 82(5):1142–1161, 1992.
17. R. Forsythe, T. A. Rietz, and T. W. Ross. Wishes, expectations, and actions: A survey on price formation in election stock markets. *Journal of Economic Behavior and Organization*, 39:83–110, 1999.
18. J. M. Gandar, W. H. Dare, C. R. Brown, and R. A. Zuber. Informed traders and price variations in the betting market for professional basketball games. *Journal of Finance*, LIII(1):385–401, 1999.
19. J. Hansen, C. Schmidt, and M. Strobel. Manipulation in political stock markets - pre-conditions and evidence. Technical Report, 2001.
20. R. D. Hanson. Combinatorial information market design. *Information Systems Frontiers*, 5(1):107–119, 2003.
21. R. D. Hanson. Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets*, 1(1):1–15, 2007.
22. R. D. Hanson, R. Oprea, and D. Porter. Information aggregation and manipulation in an experimental market. *Journal of Economic Behavior and Organization*, 60(4):449–459, 2007.
23. P. Kumar and D. J. Seppi. Futures manipulation with cash settlement. *Journal of Finance*, (47):1485–1502, 1992.
24. A. S. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1336, 1985.
25. A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomics Theory*. Oxford University Press, New York, NY, 1995.
26. P. Milgrom and N. L. Stokey. Information, trade and common knowledge. *Journal of Economic Theory*, 26(1):17–27, 1982.

27. 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.
28. K. Oliven and T. A. Rietz. Suckers are born, but markets are made: Individual rationality, arbitrage and market efficiency on an electronic futures market. *Management Science*, 50(3):336–351, 2004.
29. P. W. Rhode and K. S. Strumpf. Historical presidential betting markets. *Journal of Economic Perspectives*, 18(2):127–142, 2004.
30. P. W. Rhode and K. S. Strumpf. Manipulating political stock markets: A field experiment and a century of observational data. Working Paper, 2007.
31. R. H. Thaler and W. T. Ziemba. Anomalies: Parimutuel betting markets: Racetracks and lotteries. *Journal of Economic Perspectives*, 2(2):161–174, 1988.