

Market Manipulation with Outside Incentives

Yiling Chen · Xi Alice Gao · Rick Goldstein · Ian
A. Kash

Received: date / Accepted: date

Abstract Much evidence has shown that prediction markets can effectively aggregate dispersed information about uncertain future events and produce remarkably accurate forecasts. However, if the market prediction will be used for decision making, a strategic participant with a vested interest in the decision outcome may manipulate the market prediction to influence the resulting decision. The presence of such incentives outside of the market would seem to damage the market's ability to aggregate information because of the potential distrust among market participants. While this is true under some conditions, we show that, if the existence of such incentives is certain and common knowledge, in many cases, there exist separating equilibria where each participant changes the market probability to different values given different private signals and information is fully aggregated in the market. At each separating equilibrium, the participant with outside incentives makes a costly move to gain trust from other participants. While there also exist pooling equilibria where a participant changes the market probability to the same value given different private signals and information loss occurs, we give evidence suggesting that two separating equilibria are more natural and desirable than many other equilibria of this game by considering domination-based belief refinement, social welfare, and the expected payoff of either participant in the game. When the existence of outside incentives is uncertain, however, trust cannot be established

A preliminary version of this work appeared in the proceedings of the 24th Conference on Artificial Intelligence (AAAI'11) [6].

Yiling Chen
School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, USA.
E-mail: yiling@seas.harvard.edu

Xi Alice Gao
School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, USA.
E-mail: xagao@fas.harvard.edu

Rick Goldstein
Analytics Operations Engineering, Boston, MA, USA.
E-mail: rickgoldstein12@gmail.com

Ian A. Kash
Microsoft Research, Cambridge, UK.
E-mail: iankash@microsoft.com

between players if the outside incentive is sufficiently large and we lose the separability at equilibria.

Keywords Market manipulation · Equilibrium analysis · Prediction market · Information aggregation

1 Introduction

Prediction markets are powerful tools created to aggregate information from individuals about uncertain events of interest. As a betting intermediary, a prediction market allows traders to express their private information by wagering on event outcomes and rewards their contributions based on the realized outcome. The reward scheme in a prediction market is designed to offer incentives for traders to reveal their private information. For instance, Hanson's market scoring rule [18] incentivizes risk-neutral, myopic traders to truthfully reveal their probabilistic estimates by ensuring that truthful betting maximizes their expected payoffs. Substantial empirical work has shown that prediction markets produce remarkably accurate forecasts [1, 29, 13, 14, 11, 4].

In many real-world applications, the ultimate purpose to adopt prediction markets is to inform decision making. If a forecast gives early warning signs for a suboptimal outcome, companies may want to take actions to try to influence and improve the outcome. For example, if the forecasted release date of a product is later than expected, the company may want to assign more resources to the manufacturing of the product. If the box office revenue for a movie is forecasted to be less than expected, the production company may decide to increase its spending on advertising for the movie. In 2005 and 2006, GE Energy piloted what was called Imagination Markets where employees traded securities on new technology ideas and the ideas with the highest average security price during the last five days of the trading period were awarded research funding [22]. Subsequently, the GE-wide Imagination Market was launched in 2008. In these scenarios, little is understood of how the decision making process affects the incentives for the participants of the prediction market. If a market participant stands to benefit from a particular decision outcome, then he/she may have conflicting incentives from inside and outside of the market. Moreover, when the potential outside incentive is relatively more attractive than the payoff from inside the market, the participant may have strong incentives to strategically manipulate the market probability and deceive other participants.

We use flu prevention as a specific motivating example. Suppose that in anticipation of the upcoming flu season, the US Centers for Disease Control and Prevention (CDC) would like to purchase an appropriate number of flu vaccines and distribute them before the flu season strikes. To accomplish this, the CDC could run a prediction market to generate a forecast of the flu activity level for the upcoming flu season, and decide on the number of flu vaccines to purchase and distribute based on the market forecast. In this case, suppliers of flu vaccines, such as pharmaceutical companies, may have conflicting incentives inside and outside of the market. A pharmaceutical company can maximize its payoff within the market by truthfully reporting its information in the market or increase its profit from selling flu vaccines by driving up the final market probability. This outside incentive may cause the pharmaceutical company to manipulate the market probability in order to mislead the CDC about the expected flu activity level.

When participants have outside incentives to manipulate the market probability, it is questionable whether information can be fully aggregated in the prediction market, leading

to an accurate forecast. In this paper, we investigate information aggregation in prediction markets when such outside incentives exist. We characterize multiple perfect Bayesian equilibria (PBE) of our game and try to identify a desirable equilibrium among them. In particular, many of these equilibria are separating PBE, where the participant with the outside incentive makes a costly move in order to credibly reveal her private information and information is fully aggregated at the end of the market. Our results are summarized in the next section.

1.1 Our Results

We study a Bayesian model of a logarithmic market scoring rule (LMSR) [18] prediction market with two participants. Following a predefined sequence, each participant makes a single trade. The first participant has an outside incentive, which is certain and common knowledge. Specifically, the first participant receives an additional payoff from outside of the market, which is a result of a decision made based on the final market probability before the outcome of the event is realized. Due to the presence of this outside incentive, the first participant may want to mislead the other participant in order to maximize her total payoff from inside and outside of the market. Surprisingly, we show that there may exist a separating PBE, where every participant changes the market probability to different values when they receive different private information. In general, a separating equilibrium is desirable because all the private information gets incorporated into the final market probability. For our model, the existence of a separating PBE requires that the prior distribution and the outside incentive satisfy a particular condition and a separating PBE is achieved because the first participant makes a costly move in order to gain trust of the other participant.

When a separating PBE exists, we characterize all pure strategy separating PBE of our game. However, regardless of the existence of separating PBE, there also exist pooling PBE, where the first participant changes the market probability to the same value after receiving different private information. At a pooling PBE, information loss occurs because the first participant is unable to convince the other participant of her intention to be honest, even if she intends to be honest. We characterize a set of pooling equilibria of our game in which the behavior of the first participant varies from revealing most of her private information to revealing nothing.

Although it is difficult to conclude which PBE will be reached in practice, we show that, under certain conditions, two separating PBE, denoted SE_1 and SE_2 , are more desirable than many other PBE. By applying domination-based belief refinement, we show that in every separating PBE satisfying the refinement, the first participant's strategy is identical to her strategy in SE_1 . Under certain conditions, this belief refinement also excludes a subset of the pooling PBE of our game. Moreover, we establish that any separating PBE maximizes the total expected payoffs of the participants, if the outside incentive is an increasing convex function of the final market probability. In addition, we analyze the PBE from the perspective of a particular participant. The expected payoff of the first participant who has the outside incentive is maximized in the separating PBE SE_1 , among all separating PBE of our game. Under certain conditions, the first participant also gets a larger expected payoff in the separating PBE SE_1 compared to a set of pooling PBE of our game. For the second participant, his expected payoff is maximized in the separating PBE SE_2 among all separating PBE of our game. Such evidence suggests that the separating PBE SE_1 and SE_2 are more desirable than other equilibria of our game.

Finally, we examine more general settings. Our results of the basic model are extended to other market scoring rules. When the existence of the outside incentive is uncertain, we derive a negative result that there does not exist a separating PBE where information is fully aggregated. When a separating PBE exists for our game, we discuss a mapping from a subset of the separating PBE of our game to the set of separating PBE of Spence's job market signaling game [28]. This mapping provides nice intuitions for the existence of this subset of separating PBE.

1.2 Related Work

In a prediction market, participants may have incentives from inside or outside of the market to manipulate the market probability. Our work analyzes the strategic behavior of market participants due to outside incentives. In the literature, the work by Dimitrov and Sami [12] is the closest to our own. They study a model of two market scoring rule prediction markets for correlated events with two participants, Alice and Bob. Alice trades in the first market, and then trades in the second market after Bob. When considering the first market, Alice has an outside incentive because her trade in the first market can mislead Bob and she can obtain a higher profit in the second market by correcting Bob's mistake. In our model with only one market, the first participant also has an outside incentive, but the incentive is a payoff that monotonically increases with the final market probability. In addition, Dimitrov and Sami [12] focus on deriving properties of the players' equilibrium payoffs, whereas we explicitly characterize equilibria of our game and analyze the players' payoffs at these equilibria.

Even if there is no outside incentive, a participant in a prediction market may still have incentive from within the market to behave strategically. For instance, if a participant has multiple opportunities to trade in a market scoring rule prediction market, he may choose to withhold information in the earlier stages in order to make a larger profit later on, causing information loss in the process. Chen et al. [5] and Gao et al. [16] show that the equilibria and information revelation in such settings depend on the structure of the participants' private information. Ostrovsky [24] and Iyer et al. [20] focus on studying information aggregation at any PBE of a prediction market instead of directly characterizing the equilibria. Ostrovsky [24] analyzes an infinite-stage, finite-player market game with risk-neutral players. He characterized a condition under which the market price of a security converges in probability to its expected value conditioned on all information at any PBE. Iyer et al. [20] extend the setting of Ostrovsky [24] to risk-averse players and characterized the condition for full information aggregation in the limit at any PBE. In this work, to isolate the effect of outside incentives, we focus on settings where participants do not have incentives inside the market to manipulate the market probability.

Some recent studies consider incentives for participants to misreport their probability estimates in different models of information elicitation and decision making. Shi et al. [27] consider a setting in which a principal elicits information about a future event while participants can take hidden actions outside of the market to affect the event outcome. They characterize all proper scoring rules that incentivize participants to honestly report their probability estimates but do not incentivize them to take undesirable actions. Othman and Sandholm [25] pair a scoring rule with a decision rule. In their model, a decision maker needs to choose an action among a set of alternatives; he elicits from an expert the probability of a future event conditioned on each action being taken; the decision maker then deterministically selects an action based on the expert's prediction. They find that for the

max decision rule that selects the action with the highest reported conditional probability for the event, no scoring rule strictly incentivizes the expert to honestly report his conditional probabilities. Chen and Kash [8] and Chen et. al. [7] extend the model of Othman and Sandholm to settings of stochastic decision rules with a single expert and decision markets with multiple experts respectively and characterized all scoring rules that incentivize honest reporting of conditional probabilities. The above three studies [25, 8, 7] assume that experts do not have an inherent interest in the decision and they derive utility only from the scoring rule payment. Boutilier [2] however considers the setting in which an expert has an inherent utility in the decision and develop a set of compensation rules that when combined with the expert's utility induces proper scoring rules. Our work in this paper does not intend to design mechanisms to achieve good incentive properties in the presence of outside incentives. Instead, we study the impact of outside incentives on trader behavior and information aggregation in prediction markets using standard mechanisms.

In this paper, we model a participant's outside incentive as a function of the final market price. This is to capture scenarios where the participant's utility will be affected by some external decision, which will be made based on the final market price but prior to the realization of the event outcome. In some other scenarios, however, a participant may simply have preferences over event outcomes, i.e. the participant's utility is state-dependent. For example, a pharmaceutical company may make more profit when the flu activity level is widespread than when it is sporadic. In such scenarios, the participant with state-dependent utility, if risk averse, may trade in the prediction market for risk hedging and potentially affect the information aggregation in the market. We assume that all participants are risk neutral and hence this paper does not capture the risk hedging setting. If the participant with state-dependent utility is risk neutral, her payoff inside the market is independent of her utility outside of the market. The problem then reduces to market manipulation without outside incentives studied by Chen et al. [5], Gao et al. [16], and Ostrovsky [24].

There are some experimental and empirical studies on price manipulation in prediction markets due to incentives from outside of the market. The studies by Hansen et al. [17] and by Rhode and Strumpf [26] analyze historical data of political election betting markets. Both studies observe that these markets are vulnerable to price manipulations because media coverage of the market prices may influence the population's voting behavior. For instance, Hansen et al. describe an email communication in which a party encouraged its members to acquire contracts for the party in order to influence the voters' behaviors in the 1999 Berlin state elections, and it had temporary effects on the contract price. Manipulations in these studies were attempts not to derive more profit within the market but instead to influence the election outcome. These studies inspire us to theoretically study price manipulation due to outside incentives.

In a similar spirit, Hanson et al. [19] conducted a laboratory experiment to simulate an asset market in which some participants have an incentive to manipulate the prices. In their experiment, subjects receive different private information about the common value of an asset and they trade in a double auction mechanism. In their Manipulation treatment, half of the subjects receive an additional payoff based on the median transaction prices, so they (i.e. manipulators) have an incentive to raise the prices regardless of their private information. Hanson et al. observed that, although the manipulators attempted to raise the prices, they did not affect the information aggregation process and the price accuracy because the non-manipulators accepted trades at lower prices to counteract these manipulation attempts. This experiment closely resembles our setting because the incentive to manipulate is a payoff as a function of the market prices. However, there are two important differences. First, the additional payoff depends on the transaction prices throughout the entire trading period

whereas in our setting the additional payoff depends only on the final market price. Second, in Hanson’s experiment, although the existence of manipulators is common knowledge, the identities of these manipulators are not known. In our model, we assume that the manipulators’ identities are common knowledge. These differences may account for the different results in the two settings where manipulations did not have significant effect in Hanson’s experiment whereas in our model there exist pooling equilibria where manipulations can cause information loss. In particular, the separating equilibria in our setting may not be achievable in Hanson’s experiment because the anonymous manipulators cannot establish credibility with the other participants.

There are also experiments studying the effects of price manipulations on the information aggregation process in prediction markets without specifying the reasons for such manipulations. Camerer [3] tried to manipulate the price in a racetrack parimutuel betting market by placing large bets. These attempts were unsuccessful and he conjectured the reason to be that not all participants tried to make inferences from these bets. In their laboratory experiment, Jian and Sami [21] set up several market scoring rule prediction markets where participants may have complementary or substitute information and the trading sequence may or may not be structured. They found that previous theoretical predictions of strategic behavior by Chen et al.[5] are confirmed when the trading sequence is structured. Both studies suggest that whether manipulation can have a significant impact on price accuracy depends critically on the extent to which the participants know about other participants and reason about other participants’ actions. In our setting, we assume that all information is common knowledge except each participant’s private information, so manipulation can have a significant impact on price accuracy because participants can make a great amount of inference about each other and about the market price.

When separating PBE of our game exist, our game has a surprising connection to Spence’s job market signaling game [28]. In the signaling game, there are two types of workers applying for jobs. They have different productivity levels that are not observable and they can choose to acquire education, the level of which is observable. Spence show that, there exist separating PBE where the high productivity workers can use costly education as a signal to the employers in order to distinguish themselves from the low productivity workers. In our setting, we derive a similar result that at a separating PBE, one type of the first participant takes a loss by misreporting her information as a signal to the second participant in order to distinguish herself from her other type. We discuss this connection in detail in Section 6.

2 Model

2.1 Market Setup

Consider a binary random variable X . We run a prediction market to predict its realization $x \in \{0, 1\}$. Our market uses a market scoring rule (MSR) [18], which is a sequential shared version of a proper scoring rule, denoted $m(x, p)$.

A scoring rule $m(x, p)$ for a binary random variable is a mapping $m : \{0, 1\} \times [0, 1] \rightarrow (-\infty, \infty)$, where x is the realization of X and p is the reported probability of $x = 1$. The scoring rule is strictly proper if and only if the expected score of a risk-neutral participant with a particular belief q for the probability of $x = 1$ is uniquely maximized by reporting the probabilistic forecast $p = q$.

An MSR market with a scoring rule s starts with an initial market probability f_0 for $x = 1$ and sequentially interacts with each participant to collect his probability assessment. When

a participant changes the market probability for $x = 1$ from p to p' , he is paid the scoring rule difference, $m(x, p') - m(x, p)$, depending on the value of x . Given any strictly proper scoring rule, the corresponding MSR also incentivizes risk-neutral, myopic participants to truthfully reveal their probability assessments as they can not influence the market probabilities before their reports. We call a trader myopic if he is not forward looking and trades in each round as if it is his only chance to participate in the market.

Even though we describe MSR as a mechanism for updating probabilities, it is known that under mild conditions, MSR can be equivalently implemented as an automated market maker mechanism where participants trade shares of contracts with the market maker and, as a result, change market prices of the contracts [18,9]. For each outcome, there is a contract that pays off \$1 per share if the outcome materializes. The prices of all contracts represent a probability distribution over the outcome space. Hence, under mild conditions, trading contracts to change market prices is equivalent to changing market probabilities. We adopt the probability updating model of MSR in this paper to ease our analysis.

Our basic model considers the logarithmic market scoring rule (LMSR) which is derived from the logarithmic proper scoring rule

$$m(x, p) = \begin{cases} b \log(p), & \text{if } x = 1 \\ b \log(1 - p), & \text{if } x = 0 \end{cases} \quad (1)$$

where b is a positive parameter and p is a reported probability for $x = 1$. LMSR market maker subsidizes the market as it can incur a loss of $b \log 2$ if the traders predict the realized outcome with certainty. The parameter b scales the traders' payoffs and the market maker's subsidy but does not affect the incentives within the market. Without loss of generality, we assume $b = 1$ for the rest of the paper. In Section 5, we extend our results for LMSR to other market scoring rules.

Alice and Bob are two rational, risk-neutral participants in the market. They receive private signals described by the random variables S_A and S_B with realizations $s_A, s_B \in \{H, T\}$ ¹. Let π denote a joint prior probability distribution over X , S_A and S_B . We assume π is common knowledge and omit it in our notation for brevity.

We define $f_{s_A, \emptyset} = P(x = 1 | S_A = s_A)$ and $f_{\emptyset, s_B} = P(x = 1 | S_B = s_B)$ to represent the posterior probability for $x = 1$ given Alice's and Bob's private signal respectively. Similarly, $f_{s_A, s_B} = P(x = 1 | S_A = s_A, S_B = s_B)$ represents the posterior probability for $x = 1$ given both signals. We assume that Alice's H signal indicates a strictly higher probability for $x = 1$ than Alice's T signal, for any realized signal s_B for Bob, i.e. $f_{H, s_B} > f_{T, s_B}$ for any $s_B \in \{H, T\}$. In addition, we assume that without knowing Bob's signal, Alice's signal alone also predicts a strictly higher probability for $x = 1$ with the H signal than with the T signal and Alice's signal alone can not predict x with certainty, i.e. $0 < f_{T, \emptyset} < f_{H, \emptyset} < 1$.

In the context of our flu prediction example, we can interpret the realization $x = 1$ as the event that the flu is widespread and $x = 0$ as the event that it is not. Then the two private signals can be any information acquired by the participants about the flu activity, such as the person's own health condition.

In our basic model, the game has two stages. Alice and Bob receive their private signals at the beginning of the game. Then, Alice changes the market probability from f_0 to some

¹ Our results can be easily extended to a more general setting in which Bob's private signal has a finite number n of realizations where $n > 2$. However, it is non-trivial to extend our results to the setting in which Alice's private signal has any finite number n of possible realizations. The reason is that our analysis relies on finding an interval for each of Alice's signals, where the interval represents the range of reports that do not lead to a guaranteed loss for Alice when she receives this signal, and ranking all upper or lower endpoints of all such intervals. The number of possible rankings is exponential in n , making the analysis challenging.

value r_A in stage 1 and Bob, observing Alice's report r_A in stage 1, changes the market probability from r_A to r_B in stage 2. The market closes after Bob's report. The sequence of play is common knowledge.

Both Alice and Bob can profit from trading in the LMSR market. Moreover, Alice has an outside payoff $Q(r_B)$, which is a real-valued, non-decreasing function of the final market probability r_B . In the flu prediction example, this outside payoff may correspond to the pharmaceutical company's profit from selling flu vaccines. The outside payoff function $Q(\cdot)$ is common knowledge.

Even though our described setting is simple, with two participants, two realized signals for each participant, and two stages, our results of this basic model are applicable to more general settings. For instance, Bob can represent a group of participants who only participate after Alice and do not have the outside payoff. Also, our results remain the same if another group of participants come before Alice in the market as long as these participants do not have the outside payoff and they only participate in the market before Alice's stage of participation. We examine more general settings in Section 5.

2.2 Solution Concept

Our solution concept is the *perfect Bayesian equilibrium* (PBE) [15], which is a subgame-perfect refinement of Bayesian Nash equilibrium. Informally, a strategy-belief pair is a PBE if the players' strategies are optimal given their beliefs at any time in the game and the players' beliefs can be derived from other players' strategies using Bayes' rule whenever possible.

In our game, Alice's strategy is a specification of her report r_A in stage 1, given all realizations of her signal s_A . We denote her strategy as a mapping $\sigma : \{H, T\} \rightarrow \Delta([0, 1])$, where $\Delta(S)$ denotes the space of distributions over a set S . When a strategy maps to a report with probability 1 for both signals, the strategy is a *pure strategy*; otherwise, it is a *mixed strategy*. We use $\sigma_{s_A}(r_A)$ to denote the probability for Alice to report r_A after receiving the s_A signal. We further assume that the support of Alice's strategy is finite². If Alice does not have an outside payoff, her optimal equilibrium strategy facing the market scoring rule would be to report $f_{s_A, \emptyset}$ with probability 1 after receiving the s_A signal, since she only participates once. However, Alice has the outside payoff in our model. So she may find reporting other values more profitable if by doing so she can affect the final market probability in a favorable direction.

In stage 2 of our game, Bob moves the market probability from r_A to r_B . We denote Bob's belief as a mapping $\mu : \{H, T\} \times [0, 1] \rightarrow \Delta(\{H, T\})$, and we use $\mu_{s_B, r_A}(s_A)$ to denote the probability that Bob assigns to Alice having received the s_A signal given that she reported r_A and Bob's signal is s_B . Since Bob participates last and faces a strictly proper scoring rule in our game, his strategy at any equilibrium is uniquely determined by Alice's report r_A , his realized signal s_B and his belief μ ; he will report $r_B = \mu_{s_B, r_A}(H)f_{H, s_B} + \mu_{s_B, r_A}(T)f_{T, s_B}$.

Thus, to describe a PBE of our game, it suffices to specify Alice's strategy and Bob's belief because Alice is the first participant in the market and Bob has a dominant strategy which is uniquely determined by his belief. To show that Alice's strategy and Bob's belief form a PBE of our game, we only need to show that Alice's strategy is optimal given Bob's belief and Bob's belief can be derived from Alice's strategy using Bayes' rule whenever possible.

² This assumption is often used to avoid the technical difficulties that perfect Bayesian equilibrium has for games with a continuum of strategies. See the work by Cho and Kreps [10] for an example.

In our PBE analysis, we use the notions of *separating* and *pooling* PBE, similar to the solution concepts used by Spence [28]. These PBE notions mainly concern Alice's equilibrium strategy because Bob's optimal PBE strategy is always a pure strategy. In general, a PBE is *separating* if for any two types of each player, the intersection of the supports of the strategies of these two types is an empty set. For our game, Alice has two possible types, determined by her realized signal. A *separating* PBE of our game is characterized by the fact that the supports of Alice's strategies for the two signals, $\sigma(H)$ and $\sigma(T)$, do not intersect with each other. At a separating PBE, information is fully aggregated since Bob can accurately infer Alice's signal from her report and always make the optimal report. In contrast, a PBE is *pooling* if there exist at least two types of a particular player such that, the intersection of the supports of the strategies of these two types is not empty. At a *pooling* PBE of our game, the supports of Alice's strategies $\sigma(H)$ and $\sigma(T)$ have a nonempty intersection and Bob may not be able to infer Alice's signal from her report.

For our analysis on separating PBE, we focus on characterizing pure strategy separating PBE. These pure strategy equilibria have succinct representations, and they provide clear insights into the participants' strategic behavior in our game.

3 Known Outside Incentive

In our basic model, it is certain and common knowledge that Alice has the outside payoff. Due to the presence of the outside payoff, Alice may want to mislead Bob by pretending to have the signal H when she actually has the unfavorable signal T , in order to drive up the final market probability and gain a higher outside payoff. Bob recognizes this incentive, and in equilibrium should discount Alice's report accordingly. Therefore, we naturally expect information loss in equilibrium due to Alice's manipulation. However, from another perspective, Alice's welfare is also hurt by her manipulation since she incurs a loss in her outside payoff when having the favorable signal H due to Bob's discounting. In an equilibrium of the market, Alice balances these two conflicting forces.

In the following analysis, we characterize (pure strategy) separating and pooling PBE of our basic model. We emphasize on separating PBE because they achieve full information aggregation at the end of the market. By analyzing Alice's strategy space, we derive a succinct condition that is necessary and sufficient for a separating PBE to exist for our game. If this condition is satisfied, at any separating PBE of our game, Alice makes a costly statement, in the form of a loss in her market scoring rule payoff, in order to convince Bob that she is committed to fully revealing her private signal, despite the incentive to manipulate. If the condition is violated, there does not exist any separating PBE and information loss is inevitable.

3.1 Truthful vs. Separating PBE

The ideal outcome of this game is a truthful PBE where each trader changes the market probability to the posterior probability given all available information. A truthful PBE is desirable because information is immediately revealed and fully aggregated. However, we focus on separating PBE. The class of separating PBE corresponds exactly to the set of PBE achieving full information aggregation, and the truthful PBE is a special case in this class. Even when a truthful PBE does not exist, some separating PBE may still exist. We describe an example of the nonexistence of truthful PBE below.

At a truthful PBE, Alice's strategy is

$$\sigma_H(f_{H,\emptyset}) = 1, \sigma_T(f_{T,\emptyset}) = 1, \quad (2)$$

whereas at a (pure strategy) separating PBE, Alice's strategy can be of the form

$$\sigma_H(X) = 1, \sigma_T(Y) = 1. \quad (3)$$

for any $X, Y \in [0, 1]$ and $X \neq Y$.

In our market model, Alice maximizes her expected market scoring rule payoff in the first stage by reporting $f_{s_A, \emptyset}$ after receiving the s_A signal. If she reports r_A instead, then she incurs a loss in her expected payoff. We use $L(f_{s_A, \emptyset}, r_A)$ to denote Alice's expected loss in market scoring rule payoff by reporting r_A rather than $f_{s_A, \emptyset}$ after receiving the s_A signal as follows:

$$L(f_{s_A, \emptyset}, r_A) = f_{s_A, \emptyset} \log \frac{f_{s_A, \emptyset}}{r_A} + (1 - f_{s_A, \emptyset}) \log \frac{1 - f_{s_A, \emptyset}}{1 - r_A}, \quad (4)$$

which is the Kullback-Leibler divergence $D_{KL}(\mathbf{f}_{s_A} || \mathbf{r})$ where $\mathbf{f}_{s_A} = (f_{s_A, \emptyset}, 1 - f_{s_A, \emptyset})$ and $\mathbf{r} = (r_A, 1 - r_A)$. The following proposition describes some useful properties of $L(f_{s_A, \emptyset}, r_A)$ that will be used in our analysis in later sections.

Proposition 1 *For any $f_{s_A, \emptyset} \in (0, 1)$, $L(f_{s_A, \emptyset}, r_A)$ is a strictly increasing function of r_A and has range $[0, +\infty)$ in the region $r_A \in [f_{s_A, \emptyset}, 1)$; it is a strictly decreasing function of r_A and has range $[0, +\infty)$ in the region $r_A \in (0, f_{s_A, \emptyset}]$. For any $r_A \in (0, 1)$, $L(f_{s_A, \emptyset}, r_A)$ is a strictly decreasing function of $f_{s_A, \emptyset}$ for $f_{s_A, \emptyset} \in [0, r_A]$ and a strictly increasing function of $f_{s_A, \emptyset}$ for $f_{s_A, \emptyset} \in [r_A, 1]$.*

The proposition can be easily proven by analyzing the first-order derivatives of $L(f_{s_A, \emptyset}, r_A)$. For completeness, we include the proof in Appendix A. Lemma 1 below gives a sufficient condition on the prior distribution and outside payoff function for the nonexistence of the truthful PBE.

Lemma 1 *For any prior distribution π and outside payoff function $Q(\cdot)$, if inequality (5) is satisfied, Alice's truthful strategy given by (2) is not part of any PBE of this game.*

$$L(f_{T,\emptyset}, f_{H,\emptyset}) < E_{S_B}[Q(f_{H,S_B}) - Q(f_{T,S_B}) | S_A = T] \quad (5)$$

Proof We prove by contradiction. Suppose that inequality (5) is satisfied and there exists a PBE of our game in which Alice uses her truthful strategy. At this PBE, Bob's belief on the equilibrium path must be derived from Alice's strategy using Bayes' rule, that is,

$$\mu_{S_B, f_{H,\emptyset}}(H) = 1, \mu_{S_B, f_{T,\emptyset}}(T) = 1. \quad (6)$$

Given Bob's belief, Alice can compare her expected payoff of reporting $f_{H,\emptyset}$ with her expected payoff of reporting $f_{T,\emptyset}$ after receiving the T signal. If Alice chooses to report $f_{H,\emptyset}$ with probability 1 after receiving the T signal, then her expected gain in outside payoff is $E_{S_B}[Q(f_{H,S_B}) - Q(f_{T,S_B}) | S_A = T]$ (RHS of inequality (5)) and her expected loss in market scoring rule payoff is $L(f_{T,\emptyset}, f_{H,\emptyset})$ (LHS of inequality (5)). Because of (5), Alice has a positive net gain in her total expected payoff if she reports $f_{H,\emptyset}$ instead of $f_{T,\emptyset}$ after receiving the T signal. This contradicts the assumption that the truthful strategy is an equilibrium strategy. \square

Intuitively, the RHS of inequality (5) computes Alice's maximum possible gain in outside payoff when she has the T signal assuming Bob (incorrectly) believes that Alice received the H signal. Thus, if the outside payoff increases rapidly with the final market probability, Alice's maximum potential gain in outside payoff can outweigh her loss inside the market due to misreporting, which is given by the LHS of inequality (5).

In Appendix B, we present and discuss Example 1, which shows a prior distribution and an outside payoff function for which inequality (5) is satisfied and thus the truthful PBE does not exist. This is one of many examples where the truthful PBE does not exist. When we discuss the nonexistence of any separating PBE in section 3.3, we will present another pair of prior distribution and outside payoff function in Example 2 where a truthful PBE also fails to exist.

3.2 A Deeper Look into Alice's Strategy Space

Alice's strategy space is the interval $[0, 1]$ as she is asked to report a probability for $x = 1$. Her equilibrium strategy depends on the relative attractiveness of the market scoring rule payoff and outside payoff, which depend on the prior distribution and the outside payoff function. In this section, for a given pair of prior distribution and outside payoff function, we define some key values that are used to partition Alice's strategy space to facilitate our equilibrium analysis.

Given a prior distribution π and an outside payoff function Q , for $s_A \in \{H, T\}$, we define Y_{s_A} to be the unique value in $[f_{s_A, \emptyset}, 1]$ satisfying equation (7) and Y_{-s_A} to be the unique value in $[0, f_{s_A, \emptyset}]$ satisfying equation (8):

$$L(f_{s_A, \emptyset}, Y_{s_A}) = E_{S_B}[Q(f_{H, S_B}) - Q(f_{T, S_B}) \mid s_A], \quad (7)$$

$$L(f_{s_A, \emptyset}, Y_{-s_A}) = E_{S_B}[Q(f_{H, S_B}) - Q(f_{T, S_B}) \mid s_A]. \quad (8)$$

The RHS of the above two equations take expectations over all possible realizations of Bob's signal given Alice's realized signal s_A . Thus, the values of Y_{s_A} and Y_{-s_A} depend only on Alice's realized signal s_A and are independent of Bob's realized signal.

Note that the RHS of equations (7) and (8) are nonnegative because $f_{H, S_B} > f_{T, S_B}$ for all s_B and $Q(\cdot)$ is a non-decreasing function. By the properties of the loss function $L(f_{s_A, \emptyset}, r_A)$ described in Proposition 1, Y_{s_A} and Y_{-s_A} are well defined — given any pair of prior distribution and outside payoff function, there exists $Y_{s_A} \in [f_{s_A, \emptyset}, 1]$ and $Y_{-s_A} \in (0, f_{s_A, \emptyset}]$ such that equations (7) and (8) are satisfied. We note that $Y_{s_A} < 1$ and $Y_{-s_A} > 0$ because $L(f_{s_A, \emptyset}, r) \rightarrow \infty$ as $r \rightarrow 0$ or $r \rightarrow 1$.

Intuitively, Y_{s_A} and Y_{-s_A} are the maximum and minimum values that Alice might be willing to report after receiving the s_A signal respectively. The RHS of equations (7) and (8) are Alice's maximum possible expected gain in outside payoff by reporting some value r_A when she has the s_A signal. This maximum expected gain would be achieved if Bob had the belief that Alice has the H signal when she reports r_A and the T signal otherwise. Thus, for any realized signal s_A , Alice would not report any value outside of the range $[Y_{-s_A}, Y_{s_A}]$ because doing so is strictly dominated by reporting $f_{s_A, \emptyset}$, regardless of Bob's belief.

For each realized signal s_A , Alice's strategy space is partitioned into three distinct ranges, $[0, Y_{-s_A}]$, (Y_{-s_A}, Y_{s_A}) , and $[Y_{s_A}, 1]$. However, the partition of Alice's entire strategy space depends on the relative positions of Y_H , Y_{-H} , Y_T , and Y_{-T} , which in turn depend on the prior distribution and the outside payoff function. In the proposition below, we state several relationships of Y_H , Y_{-H} , Y_T , Y_{-T} , $f_{H, \emptyset}$, and $f_{T, \emptyset}$ that hold for all prior distributions and outside payoff functions.

Proposition 2 For all prior distributions and outside payoff functions, the following inequalities are satisfied:

$$Y_H \geq f_{H,\emptyset} \geq Y_{-H}, \quad (9)$$

$$Y_T \geq f_{T,\emptyset} \geq Y_{-T}, \quad (10)$$

$$Y_H \geq Y_{-T}. \quad (11)$$

Proof (9) and (10) hold by definition of Y_{s_A} and Y_{-s_A} . Because we assume $f_{H,\emptyset} > f_{T,\emptyset}$, we have $Y_H \geq f_{H,\emptyset} > f_{T,\emptyset} \geq Y_{-T}$. Thus, $Y_H \geq Y_{-T}$. \square

The relationships between Y_H and Y_T , Y_T and Y_{-H} , and Y_{-H} and Y_{-T} depend on the prior distribution and the outside payoff function. Next, we prove Proposition 3 below, which is useful for later analyses.

Proposition 3 $L(f_{H,\emptyset}, Y_T) \leq L(f_{H,\emptyset}, Y_{-T})$ and the equality holds only when $Y_T = Y_{-T}$.

This proposition is a direct consequence of Proposition 1. We include the proof in Appendix C.

3.3 A Necessary and Sufficient Condition for Pure Strategy Separating PBE

If a separating PBE exists for our game, it must be the case that when Alice receives the H signal, she can choose to report a particular value which convinces Bob that she is revealing her H signal truthfully. We show that this is possible if and only if the condition $Y_H \geq Y_T$ is satisfied. When $Y_H \geq Y_T$, if Alice receives the T signal, reporting $r_A \in [Y_T, Y_H]$ is dominated by reporting $f_{T,\emptyset}$. (Alice may be indifferent between reporting Y_T and $f_{T,\emptyset}$. Otherwise, the domination is strict.) So by reporting a high enough value $r_A \in [Y_T, Y_H]$ after receiving the H signal, Alice can credibly reveal to Bob that she has the H signal. However, when $Y_H < Y_T$, this is not possible. We show below that $Y_H \geq Y_T$ is necessary and sufficient for a separating PBE to exist for this game.

3.3.1 Sufficient Condition

To show that $Y_H \geq Y_T$ is a sufficient condition for a separating PBE to exist, we characterize a particular separating PBE, denoted SE_1 when $Y_H \geq Y_T$. At this separating PBE, Alice's strategy σ and Bob's belief μ are given below:

$$SE_1 : \begin{cases} \sigma_H(\max(Y_T, f_{H,\emptyset})) = 1, \sigma_T(f_{T,\emptyset}) = 1 \\ \text{When } Y_{-T} < Y_T, \mu_{s_B, r_A}(H) = \begin{cases} 1, & \text{if } r_A \in [Y_T, 1] \\ 0, & \text{if } r_A \in (Y_{-T}, Y_T) \end{cases} \\ \text{When } Y_{-T} = Y_T, \mu_{s_B, r_A}(H) = \begin{cases} 1, & \text{if } r_A \in [0, Y_{-T}] \\ 0, & \text{if } r_A = Y_T = Y_{-T} \\ 1, & \text{if } r_A \in [0, Y_{-T}) \end{cases} \end{cases} \quad (12)$$

The special case $Y_{-T} = Y_T$ only happens when $Y_{-T} = f_{T,\emptyset} = Y_T$, where SE_1 is a truthful betting PBE. Intuitively, when $f_{H,\emptyset} < Y_T$, Alice is willing to incur a high enough cost by reporting Y_T after receiving the H signal, to convince Bob that she has the H signal. Since

Bob can perfectly infer Alice's signal by observing her report, he would report f_{s_A, s_B} in stage 2 and information is fully aggregated. Alice lets Bob take a larger portion of the market scoring rule payoff in exchange for a larger outside payoff.

In SE_1 , Bob's belief says that if Alice makes a report that is too high to be consistent with the T signal ($r_A > Y_T$), Bob believes that she received the H signal. This is reasonable since Alice has no incentive to report a value that is greater than Y_T when she receives the T signal by the definition of Y_T . Similarly, if Alice makes a report that is too low to be consistent with the T signal ($r_A < Y_{-T}$), Bob also believes that she received the H signal. If Alice reports a value such that reporting this value after receiving the T signal is not dominated by reporting $f_{T, \emptyset}$ ($r_A \in (Y_{-T}, Y_T)$), then Bob believes that she received the T signal.

Theorem 1 *If $Y_H \geq Y_T$, SE_1 described in (12) is a separating PBE of our game.*

Proof First, we show that if $Y_H \geq Y_T$, then Alice's strategy is optimal given Bob's belief.

When Alice receives the T signal, by definition of Y_T , Alice would not report any $r_A > Y_T$, and furthermore she is indifferent between reporting Y_T and $f_{T, \emptyset}$. By definition of Y_{-T} , Alice would not report any $r_A < Y_{-T}$, and she is indifferent between reporting Y_{-T} and $f_{T, \emptyset}$. Any other report that is less than Y_T and greater than Y_{-T} is dominated by a report of $f_{T, \emptyset}$ given Bob's belief. Therefore, it is optimal for Alice to report $f_{T, \emptyset}$ after receiving the T signal.

When Alice receives the H signal and $Y_{-T} < Y_T$, given Bob's belief, she maximizes her expected outside payoff by reporting any $r_A \in [0, Y_{-T}] \cup [Y_T, 1]$. Now we consider Alice's expected market scoring rule payoff. By Proposition 3, if $f_{H, \emptyset} < Y_T$, reporting any $r_A \leq Y_{-T}$ is strictly dominated by reporting Y_T and Alice maximizes her expected market scoring rule payoff by reporting Y_T . Otherwise, if $f_{H, \emptyset} \geq Y_T$, then Alice maximizes her expected market scoring rule payoff by reporting $f_{H, \emptyset}$. When Alice receives the H signal and $Y_{-T} = Y_T$, it must be that $f_{H, \emptyset} > Y_T$. Given Bob's belief in this case, Alice maximizes her expected market scoring rule payoff by reporting $f_{H, \emptyset}$. Therefore, when Alice receives the H signal, it is optimal for her to report $\max(Y_T, f_{H, \emptyset})$.

Moreover, we can show that Bob's belief is consistent with Alice's strategy by mechanically applying Bayes' rule (argument omitted). Thus, SE_1 is a PBE of this game. \square

3.3.2 Necessary Condition

In Theorem 1, we characterized a separating PBE when $Y_H \geq Y_T$. In this part, we show that if $Y_H < Y_T$, there no longer exists a separating PBE. Intuitively, when $Y_H < Y_T$, even if Alice is willing to make a costly report of Y_H — which is the maximum value she would be willing to report after receiving the H signal — she still cannot convince Bob that she will report her T signal truthfully since her costly report is not sufficient to offset her incentive to misreport when having the T signal.

We first prove two useful lemmas. Lemma 2 states that, at any separating PBE, after receiving the T signal, Alice must report $f_{T, \emptyset}$ with probability 1. Lemma 3 says that at any separating PBE, after receiving the H signal, Alice does not report any $r_A \in (Y_{-T}, Y_T)$. Then we show in Theorem 2 that $Y_H \geq Y_T$ is a necessary condition for a separating PBE to exist.

Lemma 2 *In any separating PBE of our game, Alice must report $f_{T, \emptyset}$ with probability 1 after receiving the T signal.*

Proof Suppose that Alice reports $r_A \neq f_{T, \emptyset}$ after receiving the T signal. At any separating PBE, Bob's belief must be $\mu_{s_B, r_A}(H) = 0$, and $\mu_{s_B, f_{T, \emptyset}}(H) \geq 0$ in order to be consistent with

Alice's strategy. However, if Alice reports $f_{T,\emptyset}$ instead, she can strictly improve her market scoring rule payoff and weakly improves her outside payoff, which is a contradiction. \square

Note that Lemma 2 does not depend on the specific scoring rule that the market uses. It holds for any MSR market using a strictly proper scoring rule. In fact, we will use this lemma in Section 5 when extending our results to other MSR markets.

Lemma 3 *In any separating PBE of our game, Alice does not report any $r_A \in (Y_{-T}, Y_T)$ with positive probability after receiving the H signal.*

Proof We show this by contradiction. Suppose that at a separating PBE, Alice reports $r_A \in (Y_{-T}, Y_T)$ with positive probability after receiving the H signal. Since this PBE is separating, Bob's belief must be that $\mu_{s_B, r_A}(H) = 1$ to be consistent with Alice's strategy. By Lemma 2, in any separating PBE, Alice must report $f_{T,\emptyset}$ after receiving the T signal and Bob's belief must be $\mu_{s_B, f_{T,\emptyset}}(H) = 0$. Thus, for $r_A \in (Y_{-T}, Y_T)$, by definitions of Y_T and Y_{-T} , Alice would strictly prefer to report r_A rather than $f_{T,\emptyset}$ after receiving the T signal, which is a contradiction. \square

Theorem 2 *If $Y_H < Y_T$, there does not exist a separating PBE of our game.*

Proof We prove this by contradiction. Suppose that $Y_H < Y_T$ and there exists a separating PBE of our game. At this separating PBE, suppose that Alice reports some $r_A \in [0, 1]$ with positive probability after receiving the H signal.

By definitions of Y_H and Y_{-H} , we must have $r_A \in [Y_{-H}, Y_H]$. By Lemma 3, we know that $r_A \notin (Y_{-T}, Y_T)$. Next, we show that $Y_H < Y_T$ implies $Y_{-H} > Y_{-T}$.

By definitions of Y_H and Y_{-H} , we have $L(f_{H,\emptyset}, Y_{-H}) = L(f_{H,\emptyset}, Y_H)$. By Proposition 1 and $Y_H < Y_T$, we have $L(f_{H,\emptyset}, Y_H) < L(f_{H,\emptyset}, Y_T)$. By Proposition 3, we have $L(f_{H,\emptyset}, Y_T) \leq L(f_{H,\emptyset}, Y_{-T})$. To summarize, we have the following:

$$L(f_{H,\emptyset}, Y_{-H}) = L(f_{H,\emptyset}, Y_H) < L(f_{H,\emptyset}, Y_T) \leq L(f_{H,\emptyset}, Y_{-T}) \Rightarrow Y_{-H} > Y_{-T} \quad (13)$$

Thus, $r_A \in [Y_{-H}, Y_H]$ and $r_A \notin (Y_{-T}, Y_T)$ can not hold simultaneously. We have a contradiction. \square

3.3.3 When is $Y_H \geq Y_T$ satisfied?

Since $Y_H \geq Y_T$ is a necessary and sufficient condition for a separating PBE to exist, it is natural to ask when this condition is satisfied. The values of Y_H and Y_T , and whether $Y_H \geq Y_T$ is satisfied depend on the prior probability distribution π and the outside payoff function $Q(\cdot)$. When Alice's realized signal is $s_A \in \{H, T\}$, Y_{s_A} is the highest value that she is willing to report if by doing so she can convince Bob that she has the H signal. The higher the expected gain in outside payoff for Alice to convince Bob that she has the H signal rather than the T signal, the higher the value of Y_{s_A} . Below we first show that $Y_H > Y_T$ is satisfied when the signals of Alice and Bob are independent. In Appendix D, we describe Example 2 illustrating a scenario where $Y_H < Y_T$ is satisfied.

Proposition 4 *If the signals of Alice and Bob, S_A and S_B , are independent, i.e. $P(s_B|s_A) = P(s_B), \forall s_B, s_A$, then $Y_H > Y_T$ is satisfied.*

Proof When the signals of Alice and Bob are independent, Alice's expected maximum gain in outside payoff is the same, regardless of her realized signal. If we use the loss function as an intuitive distance measure from $f_{s_A, \emptyset}$ (the honest report) to Y_{s_A} (the maximum value that Alice is willing to report), then the amount of deviation from $f_{s_A, \emptyset}$ to Y_{s_A} is the same for the two realized signals. The monotonicity properties of the loss function and $f_{H, \emptyset} > f_{T, \emptyset}$ then imply $Y_H > Y_T$. We formalize this argument below.

By definitions of Y_H and Y_T and the independence of S_A and S_B , we have

$$L(f_{H, \emptyset}, Y_H) = E_{S_B}[Q(f_{H, S_B}) - Q(f_{T, S_B})] = L(f_{T, \emptyset}, Y_T). \quad (14)$$

By Proposition 1 and $f_{T, \emptyset} < f_{H, \emptyset} \leq Y_H$, we know

$$L(f_{T, \emptyset}, Y_H) > L(f_{H, \emptyset}, Y_H). \quad (15)$$

Using (14) and (15), we can derive that

$$L(f_{T, \emptyset}, Y_H) > L(f_{H, \emptyset}, Y_H) = L(f_{T, \emptyset}, Y_T). \quad (16)$$

Because $Y_T \geq f_{T, \emptyset}$ and $Y_H > f_{T, \emptyset}$, applying Proposition 1 again, we get $Y_H > Y_T$. \square

The information structure with independent signals has been studied by Chen et al. [5] and Gao et al. [16] in analyzing players' equilibrium behavior in LMSR without outside incentives. It is used to model scenarios where players obtain independent information but the outcome of the predicted event is stochastically determined by their aggregated information. Examples include the prediction of whether a candidate will receive majority vote and win an election, in which case players' votes can be viewed as independent signals and the outcome is determined by all votes.

3.4 Pure Strategy Separating PBE

In section 3.3, we described SE_1 , a particular pure strategy separating PBE of our game. There are in fact multiple pure strategy separating PBE of our game when $Y_H \geq Y_T$. In this section, we characterize all of them according to Alice's equilibrium strategy³.

By Lemma 2, at any separating PBE, Alice's strategy must be of the following form:

$$\sigma_H^S(r_A) = 1, \sigma_T^S(f_{T, \emptyset}) = 1. \quad (17)$$

for some $r_A \in [0, 1]$. In Lemma 4, we further narrow down the possible values of r_A in Alice's strategy at any separating PBE.

Lemma 4 *If $Y_H \geq Y_T$, at any separating PBE, Alice does not report any $r_A \in [0, Y_{-H}) \cup (Y_{-T}, Y_T) \cup (Y_H, 1]$ with positive probability after receiving the H signal.*

Proof By definitions of Y_H and Y_{-H} , Alice does not report any $r_A < Y_{-H}$ or $r_A > Y_H$ after receiving the H signal. By Lemma 3, Alice does not report any $r_A \in (Y_{-T}, Y_T)$ after receiving the H signal. \square

³ There exist other separating PBE where Alice plays the same equilibrium strategies as in our characterization but Bob has different beliefs off the equilibrium path.

Lemma 4 indicates that, at any separating PBE, it is only possible for Alice to report $r_A \in [\max(Y_{-H}, Y_T), Y_H]$ or, if $Y_{-H} \leq Y_{-T}$, $r_A \in [Y_{-H}, Y_{-T}]$ with positive probability after receiving the H signal.

The next two theorems characterize all separating PBE of our game when $Y_H \geq Y_T$ is satisfied. Theorem 3 shows that for every $r_A \in [\max(Y_{-H}, Y_T), Y_H]$ there is a separating PBE where Alice reports r_A after receiving the H signal. Given $Y_H \geq Y_T$, we may have either $Y_{-H} > Y_{-T}$ or $Y_{-H} \leq Y_{-T}$. If $Y_{-H} \leq Y_{-T}$, we show in Theorem 4 that for every $r_A \in [Y_{-H}, Y_{-T}]$, there exists a separating PBE at which Alice reports r_A after receiving the H signal. The proofs of these two theorems are provided in Appendices E and F respectively.

Theorem 3 *If $Y_H \geq Y_T$, for every $r_A \in [\max(Y_{-H}, Y_T), Y_H]$, there exists a pure strategy separating PBE of our game in which Alice's strategy is $\sigma_H^S(r_A) = 1, \sigma_T^S(f_{T,\emptyset}) = 1$.*

Theorem 4 *If $Y_H \geq Y_T$ and $Y_{-H} \leq Y_{-T}$, for every $r_A \in [Y_{-H}, Y_{-T}]$, there exists a pure strategy separating PBE in which Alice's strategy is $\sigma_H^S(r_A) = 1, \sigma_T^S(f_{T,\emptyset}) = 1$.*

3.5 Pooling PBE

Regardless of the existence of separating PBE, there may exist pooling PBE for our game in which information is not fully aggregated at the end of the market. If $f_{H,\emptyset} < Y_T$, there always exists a pooling PBE in which Alice reports $f_{H,\emptyset}$ with probability 1 after receiving the H signal. In general, if the interval $(\max(Y_{-H}, Y_{-T}), \min(Y_H, Y_T))$ is nonempty, for every $r_A \in (\max(Y_{-H}, Y_{-T}), \min(Y_H, Y_T)) \setminus \{f_{T,\emptyset}\}$, if r_A satisfies certain conditions, there exists a pooling PBE of our game in which Alice reports r_A with probability 1 after receiving the H signal. However, it is possible that no pooling PBE exists for a particular prior distribution and outside payoff function. We characterize a sufficient condition for pooling PBE to exist for our game in this section.

For any $k \in (\max(Y_{-H}, Y_{-T}), \min(Y_H, Y_T)) \setminus \{f_{T,\emptyset}\}$, consider the following pair of Alice's strategy and Bob's belief:

$$PE_1(k) : \begin{cases} \sigma_H^P(k) = 1, \sigma_T^P(k) = \gamma(k), \sigma_T^P(f_{T,\emptyset}) = 1 - \gamma(k) \\ \mu_{s_B, r_A}^P(H) = \begin{cases} g(\gamma(k), s_B), & \text{if } r_A = k \\ 0, & \text{if } r_A \in [0, k) \cup (k, 1] \end{cases} \end{cases} \quad (18)$$

where

$$g(\gamma(k), s_B) = \frac{P(S_A = H | s_B)}{P(S_A = H | s_B) + P(S_A = T | s_B)\gamma(k)}, \quad (19)$$

and $\gamma(k)$ is defined to be the maximum value within $[0, 1]$ such that the following inequality is satisfied.

$$L(f_{T,\emptyset}, k) \leq E_{s_B}[Q(g(\gamma(k), s_B)f_{H,s_B} + (1 - g(\gamma(k), s_B))f_{T,s_B}) - Q(f_{T,s_B}) | S_A = T] \quad (20)$$

Intuitively, $\gamma(k)$ represents the probability weight that Alice shifts from reporting $f_{T,\emptyset}$ to reporting k after receiving the T signal. The choice of $\gamma(k)$ ensures that Alice's expected loss in her market scoring rule payoff by misreporting is less than or equal to the expected potential gain in her outside payoff. So if $\gamma(k)$ satisfies equation (20), then $\gamma(k) = 1$. Otherwise, $\gamma(k)$ is set to a value such that the LHS and RHS of equation (20) are equal.

It is easy to see that $\gamma(k)$ is well defined for every $k \in (\max(Y_{-H}, Y_{-T}), \min(Y_H, Y_T)) \setminus \{f_{T,0}\}$. The RHS of inequality (20) is strictly monotonically decreasing in $\gamma(k)$. When $\gamma(k) = 0$, the RHS equals $L(f_{T,0}, Y_T)$ and $L(f_{T,0}, Y_{-T})$. Because $Y_{-T} < k < Y_T$, we know that $\gamma(k) > 0$.

By (18), Bob believes that Alice received the T signal if her report is not equal to k . If Alice reports k and Bob receives s_B signal, Bob believes that Alice received the H signal with probability $g(\gamma(k), s_B)$.

In Theorem 5, we show that $PE_1(k)$ is a pooling PBE if the following inequality is satisfied:

$$L(f_{H,0}, k) \leq E_{S_B}[Q(g(\gamma(k), S_B)f_{H,S_B} + (1 - g(\gamma(k), S_B))f_{T,S_B}) - Q(f_{T,S_B}) | S_A = H]. \quad (21)$$

Inequality (21) ensures that when Alice receives the H signal, she is better off reporting k rather than reporting $f_{H,0}$ given Bob's belief in $PE_1(k)$. When $k = f_{H,0}$, inequality (21) is automatically satisfied because the LHS of inequality (21) is 0 and the RHS of inequality (21) is positive. However, for other values of k , whether inequality (21) is satisfied depends on the prior distribution and the outside payoff function. This means that, if $f_{H,0} < Y_T$, which ensures the interval $(\max(Y_{-H}, Y_{-T}), \min(Y_H, Y_T))$ is nonempty and contains $f_{H,0}$, then there always exists a pooling PBE of our game where Alice reports $f_{H,0}$ with probability 1 after receiving the H signal.

Theorem 5 *If $(\max(Y_{-H}, Y_{-T}), \min(Y_H, Y_T))$ is nonempty, for any $k \in (\max(Y_{-H}, Y_{-T}), \min(Y_H, Y_T)) \setminus \{f_{T,0}\}$, $PE_1(k)$ is a pooling PBE of our game if inequality (21) is satisfied.*

Proof We'll first show that Alice's strategy is optimal given Bob's belief.

When Alice receives the H signal and $k \neq f_{H,0}$, for $r_A \in [0, 1] \setminus \{k\}$, it is optimal for Alice to report $f_{H,0}$ since her outside payoff is constant and her market scoring rule payoff is maximized. By inequality (21), Alice weakly prefers reporting k than reporting $f_{H,0}$. Enforcing the consistency with Bob's belief, we know that Alice's optimal strategy must be reporting k . When Alice receives the H signal and $k = f_{H,0}$, it is also optimal for Alice to report k because by doing so she maximizes both the expected market scoring rule payoff and the outside payoff given Bob's belief.

When Alice receives the T signal, for $r_A \in [0, 1] \setminus \{k\}$, Alice maximizes her total payoff by reporting $f_{T,0}$. So the support of Alice's equilibrium strategy after receiving the T signal includes at most $f_{T,0}$ and k . By the definition of $\gamma(k)$, either Alice is indifferent between reporting $f_{T,0}$ and k , or she may strictly prefer reporting k when $\gamma(k) = 1$. Enforcing the consistency of Bob's belief, we know that Alice's optimal strategy must be reporting k with probability $\gamma(k)$ and reporting $f_{T,0}$ with probability $1 - \gamma(k)$.

Moreover, we can show that Bob's belief is consistent with Alice's strategy by mechanically applying Bayes' rule (argument omitted). Given the above arguments, Alice's strategy and Bob's belief form a PBE of our game. \square

Babbling PBE

For Bob's belief in $PE_1(k)$, it is possible that for some k , $\gamma(k) = 1$. In this case, Alice's strategy and Bob's belief become the following:

$$BE_1(k) : \begin{cases} \sigma_H^B(k) = 1, \sigma_T^B(k) = 1 \\ \mu_{s_B, r_A}^B(H) = \begin{cases} P(s_A = H | s_B), & \text{if } r_A = k \\ 0, & \text{if } r_A \in [0, k) \cup (k, 1] \end{cases} \end{cases} \quad (22)$$

This special case of the pooling PBE is often alluded to as a *babbling* PBE. At this babbling PBE, if Alice reports k , then Bob believes that she received the H signal with the prior probability $P(s_A = H|s_B)$. Otherwise, if Alice reports any other value, then Bob believes that she received the T signal for sure. This belief forces Alice to make a completely uninformative report by always reporting k no matter what her realized signal is. This PBE is undesirable since Alice does not reveal her private information.

4 Identifying Desirable PBE

The existence of multiple equilibria is a common problem to many dynamic games of incomplete information. This is undesirable because there is no clear way to identify a single equilibrium that the players are likely to adopt and hence it is difficult to predict how the game will be played in practice. In our setting, this problem arises because we have a great deal of freedom in choosing beliefs off the equilibrium path. A common way to address this problem is to use some criteria to identify one or more equilibria to be more desirable than others. An equilibrium is more desirable than other equilibria if it satisfies reasonable belief refinements or optimizes certain desirable objectives.

In this section, we give evidence suggesting that two separating PBE SE_1 (defined in equation (12)) and SE_2 (defined in equation (28)) are more desirable than many other PBE of our game, according to several different objectives. First, in every separating PBE that satisfies the domination-based belief refinement, Alice plays the same strategy as her strategy in SE_1 . This refinement also excludes a subset of pooling PBE of our game under certain conditions. With the goal of maximizing social welfare, we show that any separating PBE maximizes the social welfare of our game among all PBE if Alice's outside payoff function $Q(\cdot)$ is convex⁴. This shows that both SE_1 and SE_2 are more desirable than pooling equilibria. Finally, we compare the multiple separating equilibria from the perspective of a particular player. In terms of maximizing Alice's total expected payoff, the PBE SE_1 is more desirable than all other separating PBE and many pooling PBE of our game. From the perspective of Bob, the PBE SE_2 maximizes Bob's total expected payoff among all separating PBE of our game.

4.1 Domination-based Belief Refinement

There has been a large literature in economics devoted to identifying focal equilibria through refinements. One simple PBE refinement, as discussed by Mas-Colell, Whinston and Green [23], arises from the idea that reasonable beliefs should not assign positive probability to a player taking an action that is strictly dominated for her. Formally, we define this refinement for our game as follows:

Definition 1 [*Domination-based belief refinement*] If possible, at any PBE satisfying domination-based belief refinement, Bob's belief should satisfy $\mu_{s_B, r_A}(\theta) = 0$ if reporting r_A for Alice's type θ is strictly dominated by reporting $r'_A \in [0, 1]$ where $r'_A \neq r_A$ for any valid belief of Bob.

⁴ Situations with a convex $Q(\cdot)$ function arise, for example, when manufactures have increasing returns to scale, which might be the case in our flu prediction example.

The qualification “if possible” covers the case that reporting r_A for all of Alice’s types is strictly dominated by reporting some other r'_A for any valid belief for Bob. In this case, if we apply the refinement to Bob’s belief, then Bob’s belief must set $\mu_{s_B, r_A}(H) = 0$ and $\mu_{s_B, r_A}(T) = 0$, which does not result in a valid belief for Bob. Therefore, in this case the refinement would not apply and Bob’s belief is unrestricted when Alice reports such a r_A . Using Definition 1 we can put restrictions on Bob’s belief at any PBE.

Lemma 5 *At any PBE satisfying the domination-based belief refinement, if $Y_H \geq Y_T$, then Bob’s belief should satisfy $\mu_{s_B, r_A}(T) = 0$ for any $r_A \in (Y_T, Y_H] \cap [Y_{-H}, Y_H]$. If $Y_{-H} \leq Y_{-T}$, then Bob’s belief should satisfy $\mu_{s_B, r_A}(T) = 0$ for any $r_A \in [Y_{-H}, Y_{-T}]$.*

Proof By definition of Y_T and Y_{-T} , reporting any $r_A > Y_T$ or $r_A < Y_{-T}$ after receiving the T signal is strictly dominated by reporting $f_{T, \emptyset}$ for Alice. By definition of Y_H and Y_{-H} , reporting any $r_A > Y_H$ or $r_A < Y_{-H}$ after receiving the H signal is strictly dominated by reporting $f_{H, \emptyset}$ for Alice.

For any $r_A \in [0, \min\{Y_{-H}, Y_{-T}\}) \cup (\max\{Y_T, Y_H\}, 1]$, Bob’s belief is unrestricted because the domination-based belief refinement does not apply. By Definition 1, it is straightforward to verify that Bob’s belief should satisfy $\mu_{s_B, r_A}(T) = 0$ for any $r_A \in (Y_T, Y_H] \cap [Y_{-H}, Y_H]$ when $Y_H \geq Y_T$, and for any $r_A \in [Y_{-H}, Y_{-T}]$ when $Y_{-H} \leq Y_{-T}$. \square

Given this refinement on Bob’s belief at the PBE, we show below that at every separating PBE of our game, Alice’s strategy must be the same as that in the separating PBE SE_1 ⁵.

Proposition 5 *At every separating PBE satisfying the domination-based belief refinement, Alice’s strategy must be $\sigma_H(\max(f_{H, \emptyset}, Y_T)) = 1$, and $\sigma_T(f_{T, \emptyset}) = 1$.*

We provide the complete proof in Appendix G.

Proof (Sketch) By Theorem 3, for every $r_A \in [\max(Y_{-H}, Y_T), Y_H]$, there exists a pure strategy separating PBE in which Alice reports r_A with probability 1 after receiving the H signal. We show that Alice would not report $r_A \in [\max(Y_{-H}, Y_T), Y_H] \setminus \max(f_{H, \emptyset}, Y_T)$ after receiving the H signal at any PBE satisfying the domination-based belief refinement.

By Theorem 4, if $Y_H \geq Y_T$ and $Y_{-H} \leq Y_{-T}$, for every $r_A \in [Y_{-H}, Y_{-T}]$, there exists a pure strategy separating PBE in which Alice reports r_A with probability 1 after receiving the H signal. First, we show that Alice would not report $r_A \in [Y_{-H}, Y_{-T}]$ at any PBE satisfying the domination-based belief refinement. Then we show that Alice also would not report Y_{-T} at any such PBE.

Finally, we show that SE_1 described in (12) satisfies the domination-based belief refinement. \square

If $f_{H, \emptyset} > Y_T$, the domination-based refinement can also exclude all pooling PBE and the unique PBE satisfying the refinement is the truthful PBE. We show below that, when $f_{H, \emptyset} > Y_T$, at every PBE of our game, Alice’s strategy is $\sigma_H(f_{H, \emptyset}) = 1$, $\sigma_T(f_{T, \emptyset}) = 1$, which is Alice’s strategy in the separating PBE SE_1 .

Proposition 6 *At every PBE of our game satisfying the domination-based refinement, if $f_{H, \emptyset} > Y_T$, then Alice’s strategy must be $\sigma_H(f_{H, \emptyset}) = 1$ and $\sigma_T(f_{T, \emptyset}) = 1$.*

⁵ Bob’s belief can be different from that in SE_1 .

Proof Since $f_{H,\emptyset} \in (Y_T, Y_H]$, then by Lemma 5, Bob's belief must set $\mu_{s_B, f_{H,\emptyset}}(T) = 0$. If Alice receives the H signal, then her market scoring rule payoff is strictly maximized by reporting $f_{H,\emptyset}$ and her outside payoff is weakly maximized by reporting $f_{H,\emptyset}$. Therefore, it is optimal for Alice to report $f_{H,\emptyset}$ after receiving the H signal.

If Alice receives the T signal, reporting $f_{H,\emptyset}$ is strictly dominated by reporting $f_{T,\emptyset}$ for any valid belief for Bob because $f_{H,\emptyset} > Y_T$. Therefore, Alice does not report $f_{H,\emptyset}$ after receiving T signal, and any PBE of the game must be a separating PBE. By Proposition 5, any separating PBE satisfying the refinement has Alice play the strategy $\sigma_H(f_{H,\emptyset}) = 1$ and $\sigma_T(f_{T,\emptyset}) = 1$. \square

If $f_{H,\emptyset} \leq Y_T$, applying the domination-based refinement does not exclude all pooling PBE of this game. In the proposition below, we show that the domination-based refinement excludes a subset of pooling PBE in which Alice reports a low enough value after receiving the H signal. The proof of the proposition is provided in Appendix H.

Proposition 7 *At every PBE of our game satisfying the domination-based refinement, if $f_{H,\emptyset} \leq Y_T$, then Alice does not report any $r_A \leq r$ after receiving the H signal where r is the unique value in $[0, f_{H,\emptyset}]$ satisfying $L(f_{H,\emptyset}, r) = L(f_{H,\emptyset}, Y_T)$.*

4.2 Social Welfare

We analyze the *social welfare* achieved in the PBE of our game. In general, *social welfare* refers to the total expected payoffs of all players in the game. In our setting, the *social welfare* of our game is defined to be the total ex-ante expected payoff of Alice and Bob excluding any payoff for the market institution. Alice's total expected payoff includes her market scoring rule payoff and her outside payoff.

Since all separating PBE fully aggregate information, they all result in the same (maximized) total ex-ante expected payoff inside the market – all that changes is how Alice and Bob split this payoff – and the same outside payoff for Alice. If the outside payoff function $Q(\cdot)$ is convex, we show in Lemma 6 that Alice's expected outside payoff is also maximized in any separating PBE of our game. Therefore, given a convex $Q(\cdot)$, social welfare is maximized at any separating PBE. We prove this claim in Theorem 6.

Lemma 6 *If $Q(\cdot)$ is convex, among all PBE of the game, Alice's expected outside payoff is maximized in any separating PBE.*

Proof Consider an arbitrary PBE of this game. Let V denote the union of the supports of Alice's strategy after receiving the H and the T signals at this PBE. Let u_A^G denote Alice's expected outside payoff at this PBE and let u_A^S denote Alice's expected outside payoff at any separating PBE. We'll prove below that $u_A^G \leq u_A^S$. We simplify our notation by using

$P(S_A, S_B)$ to denote $P(S_A = s_A, S_B = s_B)$.

$$\begin{aligned}
u_A^G &= \sum_{v \in V} (P(H, H)\sigma_H(v) + P(T, H)\sigma_T(v)) \\
&\quad Q\left(\frac{P(H, H)\sigma_H(v)}{P(H, H)\sigma_H(v) + P(T, H)\sigma_T(v)}f_{HH} + \frac{P(T, H)\sigma_T(v)}{P(H, H)\sigma_H(v) + P(T, H)\sigma_T(v)}f_{TH}\right) \\
&\quad + (P(H, T)\sigma_H(v) + P(T, T)\sigma_T(v)) \\
&\quad Q\left(\frac{P(H, T)\sigma_H(v)}{P(H, T)\sigma_H(v) + P(T, T)\sigma_T(v)}f_{HT} + \frac{P(T, T)\sigma_T(v)}{P(H, T)\sigma_H(v) + P(T, T)\sigma_T(v)}f_{TT}\right) \\
&\leq \sum_{v \in V} (P(H, H)\sigma_H(v)Q(f_{HH}) + P(H, T)\sigma_H(v)Q(f_{HT}) \\
&\quad + P(T, H)\sigma_T(v)Q(f_{TH}) + P(T, T)\sigma_T(v)Q(f_{TT})) \\
&= P(H, H)Q(f_{HH}) + P(H, T)Q(f_{HT}) + P(T, H)Q(f_{TH}) + P(T, T)Q(f_{TT}) \\
&= u_A^S
\end{aligned} \tag{23}$$

where inequality (23) was derived by applying the convexity of $Q(\cdot)$. \square

Theorem 6 *If $Q(\cdot)$ is convex, among all PBE of the game, social welfare is maximized at any separating PBE.*

Proof By definition, at any separating PBE, the total market scoring rule payoff is maximized since information is fully aggregated. By Lemma 6, any separating PBE maximizes Alice's expected outside payoff if $Q(\cdot)$ is convex. Therefore, any separating PBE maximizes the social welfare. \square

4.3 Alice's Total Expected Payoff

In this section, we compare the multiple PBE of our game in terms of Alice's total expected payoff. If Alice's total expected payoff at a particular PBE is greater than her total expected payoff in many other PBE of this game, it gives us confidence that she is likely to choose to play this particular PBE in practice.

First, we compare Alice's expected payoff in the multiple separating PBE of our game. We show in Theorem 7 that the separating PBE SE_1 maximizes Alice's expected payoff among all separating PBE of this game. This is easy to see when $f_{H,0} \geq Y_T$ since the separating PBE SE_1 is also the truthful PBE of this game. Otherwise, if $f_{H,0} < Y_T$, Y_T is the minimum deviation from $f_{H,0}$ that Alice can report in order to convince Bob that she has the H signal.

Theorem 7 *Among all pure strategy separating PBE of our game, Alice's expected payoff is maximized in the pure strategy separating PBE SE_1 as stated in (12).*

Proof In all separating PBE, Alice's expected outside payoff is the same.

By Lemma 2, in any separating PBE, Alice must report $f_{T,0}$ after receiving the T signal. Therefore, Alice's expected payoff after receiving the T signal is the same at any separating PBE.

When $f_{H,0} \geq Y_T$, according to Theorem 1, Alice reports $f_{H,0}$ after receiving the H signal and this is the maximum expected payoff she could get after receiving the H signal.

When $f_{H,0} < Y_T$, after receiving the H signal, Alice's strategy in SE_1 is to report Y_T . She is strictly worse off reporting any value greater than Y_T after receiving the H signal in any PBE. For $r_A < Y_T$, if $Y_{-H} < Y_{-T}$, it is only possible for Alice to report $r_A \in [Y_{-H}, Y_{-T})$ after receiving the H signal at any separating PBE. However, reporting $r_A \in [Y_{-H}, Y_{-T})$ makes Alice strictly worse off than reporting Y_T because

$$r_A < Y_{-T} \Rightarrow L(f_{H,0}, Y_T) \leq L(f_{H,0}, Y_{-T}) < L(f_{H,0}, r_A). \quad (24)$$

where the inequality $L(f_{H,0}, Y_T) \leq L(f_{H,0}, Y_{-T})$ is due to Proposition 3. Therefore, when $f_{H,0} \leq Y_T$, the separating PBE in which Alice reports Y_T maximizes Alice's expected payoff after receiving the H signal.

Hence, the separating PBE SE_1 maximizes Alice's expected payoff among all separating PBE of our game. \square

Theorem 7 suggests that SE_1 is likely a desirable PBE of our game. In Theorems 8 and 9, we compare Alice's expected payoff in SE_1 with her expected payoff in the pooling PBE of this game. Again, when $f_{H,0} \geq Y_T$, SE_1 is essentially the truthful PBE and therefore Alice's expected payoff is higher in SE_1 than in any pooling PBE for convex $Q(\cdot)$. When $f_{H,0} < Y_T$, the relationship is less clear. In Theorem 9, we show that, if $k \in (\max(Y_{-H}, Y_{-T}), Y_T) \setminus \{f_{T,0}\}$ satisfies inequality (25), then Alice's expected payoff in SE_1 is greater than her expected payoff in the pooling PBE $PE_1(k)$.

Theorem 8 *If $Q(\cdot)$ is convex, $Y_H \geq Y_T$ and $f_{H,0} \geq Y_T$, Alice's expected payoff is maximized in the pure strategy separating PBE SE_1 among all PBE of our game.*

Proof By Lemma 6, any separating PBE maximizes Alice's expected outside payoff if $Q(\cdot)$ is convex. When $f_{H,0} \geq Y_T$, SE_1 is the truthful PBE and strictly maximizes Alice's expected market scoring rule payoff. \square

Theorem 9 *If $Q(\cdot)$ is convex, $Y_H \geq Y_T$, and $f_{H,0} < Y_T$, Alice's expected payoff in the pure strategy separating PBE SE_1 is greater than her expected payoff in $PE_1(k)$ for any $k \in (\max(Y_{-H}, Y_{-T}), Y_T) \setminus \{f_{T,0}\}$ if k satisfies inequality (25) below.*

$$P(s_A = H)L(f_{H,0}, Y_T) \leq P(s_A = H)L(f_{H,0}, k) + P(s_A = T)\gamma(k)L(f_{T,0}, k). \quad (25)$$

Proof By Lemma 6, if $Q(\cdot)$ is convex, then any separating PBE maximizes Alice's expected outside payoff.

Fix a particular $k \in (Y_{-T}, \min\{Y_H, Y_T\})$. Compared to Alice's expected payoff when using a truthful strategy, Alice's expected payoff in SE_1 given in Theorem 1 is less by $P(s_A = H)L(f_{H,0}, Y_T)$, and Alice's payoff in $PE_1(k)$ is less by $P(s_A = H)L(f_{H,0}, k) + P(s_A = T)\gamma(k)L(f_{T,0}, k)$. Therefore, if Alice's expected payoff SE_1 is greater than or equal to Alice's expected payoff in $PE_1(k)$, then we must have $P(s_A = H)L(f_{H,0}, Y_T) \leq P(s_A = H)L(f_{H,0}, k) + P(s_A = T)\gamma(k)L(f_{T,0}, k)$, which is stated in inequality (25). \square

Note that inequality (25) is automatically satisfied for any $k \leq r$ where r is the unique value in $[0, f_{H,0}]$ satisfying $L(f_{H,0}, r) = L(f_{H,0}, Y_T)$ since

$$P(s_A = H)L(f_{H,0}, Y_T) = P(s_A = H)L(f_{H,0}, r) \leq P(s_A = H)L(f_{H,0}, k) \quad (26)$$

$$< P(s_A = H)L(f_{H,0}, k) + P(s_A = T)\gamma(k)L(f_{T,0}, k). \quad (27)$$

However, for $k \geq r$, whether k satisfies inequality (25) depends on the prior distribution and the outside payoff function.

4.4 Bob's Expected Payoff

In this section, we compare all separating PBE of our game from Bob's perspective. If Bob's expected payoff at a particular PBE is greater than his expected payoff in many other PBE of this game, then Bob is more likely to choose to play this particular PBE in practice. We show below that among all separating PBE of our game, Bob's expected payoff is maximized in the separating PBE SE_2 in equation (28), which is the same as $SE_2(Y_H)$ defined in equation (57) in Appendix E. We state SE_2 below for convenience. The proof of Theorem 10 is included in Appendix I.

Theorem 10 *Among all pure strategy separating PBE of our game, Bob's expected payoff is maximized in the following pure strategy separating PBE SE_2 .*

$$SE_2 : \begin{cases} \sigma_H(Y_H) = 1, \sigma_T(f_{T,\emptyset}) = 1 \\ \mu_{s_B, r_A}(H) = \begin{cases} 1, & \text{if } r_A \in [Y_H, 1] \\ 0, & \text{if } r_A \in [0, Y_H) \end{cases} \end{cases} \quad (28)$$

5 Extensions

We have developed our results for the basic setting, with LMSR, two players, two stages, and binary signals for each player. In this section, we extend our separating PBE results to other market scoring rules. We also consider an extension of our setting where the outside incentive is uncertain, but occurs with a fixed probability. We show that this uncertainty is detrimental to information aggregation and there does not exist any separating PBE in this setting.

5.1 Other Market Scoring Rules

For our basic model using LMSR, we characterize a necessary and sufficient condition for a separating PBE to exist. In this section, we generalize this condition for other MSR markets using strictly proper scoring rules. The main difficulty in this generalization is that, for an arbitrary market scoring rule, Y_{s_A} and Y_{-s_A} for $s_A \in \{H, T\}$ may not be well defined, whereas they are always well defined for LMSR because the loss function is not bounded from above. As a result, when generalizing the condition, we need to take into account of the cases when Y_{s_A} and/or Y_{-s_A} are not well defined.

As defined in section 2.1, let $m(x, p)$ denote a strictly proper scoring rule of a binary random variable X where x is the realization of X and p is the reported probability of $x = 1$. Then the loss function $L_m(f_{s_A, \emptyset}, r_A)$ for the strictly proper scoring rule $m(x, p)$ can be defined as follows:

$$L_m(f_{s_A, \emptyset}, r_A) = f_{s_A, \emptyset} \{m(1, f_{s_A, \emptyset}) - m(1, r_A)\} + (1 - f_{s_A, \emptyset}) \{m(0, f_{s_A, \emptyset}) - m(0, r_A)\} \quad (29)$$

For a particular market scoring rule, a sufficient condition for a separating PBE to exist can be expressed by the the following two inequalities.

$$L_m(f_{T, \emptyset}, 1) \geq E_{S_B} [Q(f_{H, S_B}) - Q(f_{T, S_B}) \mid S_A = T] \quad (30)$$

$$L_m(f_{H, \emptyset}, \max(f_{H, \emptyset}, Y_T)) \leq E_{S_B} [Q(f_{H, S_B}) - Q(f_{T, S_B}) \mid S_A = H] \quad (31)$$

If inequality (30) is satisfied, we know that Y_T is well defined. Then, if inequality (31) is also satisfied, reporting $\max(f_{H,\emptyset}, Y_T)$ for Alice is not dominated by reporting $f_{H,\emptyset}$ after receiving the H signal. So if both inequalities are satisfied, then there exists a separating PBE where Alice reports $\max(f_{H,\emptyset}, Y_T)$ after receiving the H signal. Note that if inequality (30) is violated, then the quantity Y_T is not well defined, so inequality (31) is not a well defined statement as well.

Similarly, another sufficient condition for a separating PBE is given by the following two inequalities.

$$L_m(f_{T,\emptyset}, \emptyset) \geq E_{S_B}[Q(f_{H,S_B}) - Q(f_{T,S_B}) \mid S_A = T] \quad (32)$$

$$L_m(f_{H,\emptyset}, Y_{-T}) \leq E_{S_B}[Q(f_{H,S_B}) - Q(f_{T,S_B}) \mid S_A = H] \quad (33)$$

Again, inequality (32) ensures that Y_{-T} is well-defined. If inequality (33) is satisfied, then there exists a belief for Bob such that for Alice, reporting Y_{-T} is not dominated by reporting $f_{H,\emptyset}$ after receiving the H signal. Therefore, these two inequalities ensure that there exists a separating PBE where Alice reports Y_{-T} after receiving the H signal.

We show below in Theorem 11 that satisfying at least one of these two pairs of inequalities is necessary and sufficient for a separating PBE to exist for any market scoring rule.

Theorem 11 *A separating PBE of our game exists if and only if at least one of the pair of inequalities (30) and (31) and the pair of inequalities (32) and (33) is satisfied.*

We include the complete proof in Appendix J.

5.2 Uncertain Outside Incentive

In our basic model, Alice's outside incentive is certain and common knowledge. In this section, however, we show that any uncertainty about Alice's outside incentive could be detrimental to information aggregation. Suppose that there is a fixed probability $\alpha \in (0, 1)$ for Alice to have the outside payoff. Even if the value of α is common knowledge, information loss in equilibrium is inevitable if Alice has a sufficiently large outside incentive. In particular, when Alice has an outside payoff and has received the T signal, she can report $f_{H,\emptyset}$ to pretend not to have the outside payoff and to have received the H signal. This results in these two types pooling, so the overall equilibrium is, at best, semi-separating and there is information loss.

Theorem 12 *Suppose that Alice has the outside payoff with a fixed probability $\alpha \in (0, 1)$, which is common knowledge. If $f_{H,\emptyset} < Y_T$, then there does not exist any PBE in which Alice's type with the H signal and no outside payoff separates from her type with the T signal and the outside payoff.*

Proof Proof by contradiction. Suppose that a separating PBE exists. At this separating PBE, with probability $(1 - \alpha)$, Alice reports $f_{H,\emptyset}$ after receiving the H signal and reports $f_{T,\emptyset}$ after receiving the T signal. To be consistent with Alice's strategy, Bob's belief on the equilibrium path must be $\mu_{S_B, f_{H,\emptyset}}(H) = 1$ and $\mu_{S_B, f_{T,\emptyset}}(H) = 0$. Given this belief, however, when Alice has the outside payoff, she strictly prefers to report $f_{H,\emptyset}$ after receiving the T signal since $Y_T > f_{H,\emptyset}$, which is a contradiction. \square

6 Connection to Spence's Job Market Signaling Game

In this section, we describe the connection between a subset of separating PBE of our game and a set of separating PBE of Spence's job market signaling game [28]. When a separating PBE exists for our game, $Y_H \geq Y_T$ holds and there exists a set of separating PBE where Alice reports $r_A \in [\max(Y_{-H}, Y_T), Y_H]$ after receiving H signal. If in addition $Y_{-H} \leq Y_{-T}$ also holds, then there also exists a set of separating PBE where Alice reports $r_A \in [Y_{-H}, Y_{-T}]$ after receiving the H signal. In the following analysis, we consider a set of separating PBE where Alice reports $r_A \in [\max(f_{H,\emptyset}, Y_T), Y_H]$ after receiving the H signal, which is a subset of the first set of separating PBE described above, and we map these separating PBE to the separating PBE of the signaling game.

We first describe the setting of a signaling game using the notation of Mas-Colell, Whinston, and Green [23]. In the signaling game, there are two types of workers with productivities θ_H and θ_L , with $\theta_H > \theta_L > 0$, and these productivities are not observable. Before entering the job market, each worker can get some amount of education, and the amount of education that a worker receives is observable. Getting education does not affect a worker's productivity, but the high-productivity workers in the job market may use education to distinguish them from the low-productivity workers. The cost of obtaining education level e for a type θ worker is given by the twice continuously differentiable function $c(\theta, e)$, with $c(\theta, 0) = 0$, $\frac{\partial}{\partial e} c(\theta, e) > 0$, $\frac{\partial^2}{\partial e^2} c(\theta, e) > 0$, $c(\theta_H, e) < c(\theta_L, e)$, for all $e > 0$ and $\frac{\partial}{\partial e} c(\theta_H, e) < \frac{\partial}{\partial e} c(\theta_L, e), \forall e > 0$. Both the cost and the marginal cost of education are lower for workers with productivity θ_H . Each worker can choose to work at home or work for an employer. Working at home earns the worker no wage. If the worker chooses to work for an employer, then his wage depends on the employer's belief about the worker's productivity based on the worker's education level. If a type θ worker chooses education level e and receives wage ω , then his payoff, denoted by $u(\omega, e | \theta)$, is equal to his wage less the cost of getting education, i.e. $u(\omega, e | \theta) = \omega - c(e, \theta)$.

In separating PBE of the signaling game, many education levels for the high productivity worker are possible and the low productivity worker chooses no education. In particular, any education level in some range $[\bar{e}, e_1]$ for the high productivity workers can be sustained at a PBE of this game. Intuitively, the education level of the high productivity worker cannot be below \bar{e} in a separating PBE because, if it were, the low productivity worker would find it profitable and pretend to be of high productivity by choosing the same education level. On the other hand, the education level of the high productivity worker cannot be above e_1 because, if it were, the high productivity worker would prefer to get no education instead, even if this meant that he would be considered to be of low productivity.

Consider our setting when a separating PBE exists (i.e. $Y_H \geq Y_T$), we can map elements of our game to the signaling game. We can also map separating PBE of our game where Alice reports $r_A \in [\max(f_{H,\emptyset}, Y_T), Y_H]$ to the separating PBE of the signaling game. We outline details of this mapping in Table 1.

Alice's two types H and T in our setting correspond to the two types of workers with productivities θ_H and θ_L . If Alice chooses to report a value $r_A > f_{H,\emptyset}$, she incurs a loss in the market scoring rule payoff for either type. This loss is the cost of misreporting and corresponds to the cost of getting education for either worker type in the signaling game. Moreover, the loss function and the cost function have similar properties: they are increasing and convex in education level/report and they are lower for the θ_H/H type. Also the marginal loss and cost functions are lower for the θ_H/H type. As a result of these properties, in both settings, there exists a range of possible values for the education level/report of the θ_H/H

Table 1 Comparison between our setting and Spence's job market signaling game

Spence's Job Market Signaling Game	Our Setting
θ_H , high productivity worker	H , Alice's H type
θ_L , low productivity worker	T , Alice's T type
$e > 0$, education level	$r_A > f_{H,\emptyset}$, Alice's report r_A
$c(\theta, e)$, cost of education as a function of the level e and the worker type θ	$L(f_{s_A,\emptyset}, r_A)$, loss function with respect to report r_A and type s_A
$\frac{\partial}{\partial e} c(\theta, e) > 0, \forall e > 0$, cost of education is increasing in education level	$\frac{\partial}{\partial r_A} L(f_{s_A,\emptyset}, r_A) > 0, \forall r_A > f_{H,\emptyset}$, loss is increasing in Alice's report
$\frac{\partial^2}{\partial e^2} c(\theta, e) > 0, \forall e > 0$, cost is convex in education level	$\frac{\partial^2}{\partial r_A^2} L(f_{s_A,\emptyset}, r_A) > 0, \forall r_A > f_{H,\emptyset}$, loss is convex in Alice's report
$c(\theta_H, e) < c(\theta_L, e), \forall e > 0$, cost is lower for high productivity worker	$L(f_{H,\emptyset}, r_A) < L(f_{T,\emptyset}, r_A), \forall r_A > f_{H,\emptyset}$, loss is lower for Alice's H type
$\frac{\partial}{\partial e} c(\theta_H, e) < \frac{\partial}{\partial e} c(\theta_L, e), \forall e > 0$, marginal cost is lower for high productivity worker	$\frac{\partial}{\partial r_A} L(f_{H,\emptyset}, r_A) < \frac{\partial}{\partial r_A} L(f_{T,\emptyset}, r_A), \forall r_A > f_{H,\emptyset}$, marginal loss is lower for Alice's H type
e_1 , highest education level for high productivity worker among all separating PBE	Y_H , highest report for H type among all separating PBE
\bar{e} , lowest education level for high productivity worker among all separating PBE	$\max(f_{H,\emptyset}, Y_T)$, lowest report for H type among the subset of separating PBE

type reports whereas the θ_L/T type chooses no education/does not misreport at separating PBE.

In the signaling game, the fundamental reason that education can serve as a signal is that the marginal cost of education depends on a worker's type. The marginal cost of education is lower for a high-productivity worker ($\frac{\partial}{\partial e} c(\theta_H, e) < \frac{\partial}{\partial e} c(\theta_L, e)$). As a result, a θ_H type worker may find it worthwhile to get some positive level of education $e > 0$ to raise her wage by some positive amount whereas a type θ_L worker may not be willing to get the same level of education in return for the same amount of wage increase. As a result, by getting an education in the range $[\bar{e}, e_1]$, a high-productivity worker could distinguish themselves from their low-ability counterparts. Analogously, in our setting, the fundamental reason that a separating PBE where Alice reports $r_A \in [\max(f_{H,\emptyset}, Y_T), Y_H]$ exists is that reporting a value $r_A > f_{H,\emptyset}$ has a marginally lower expected loss in market scoring rule payoff for Alice's H type than for Alice's T type. Thus, Alice's H type may be willing to report a value r_A much higher than $f_{H,\emptyset}$ in order to increase her outside payoff whereas Alice's T type may not be willing to report r_A for the same amount of increase in her outside payoff. Therefore, when $Y_H \geq Y_T$, there exists a range of reports, $[\max(f_{H,\emptyset}, Y_T), Y_H]$, such that Alice's H type can distinguish herself from Alice's T type in our game.

Note that, if $Y_{-H} \leq Y_{-T}$ holds in addition to $Y_H \geq Y_T$, it is also possible to map the set of separating PBE where Alice reports $r_A \in [Y_{-H}, Y_{-T}]$ to the separating PBE of the signaling game. The only caveat is that, instead of mapping education e directly to Alice's report r_A , we need to map education e to the distance between Alice's report r_A and $f_{H,\emptyset}$. We omit the description of the mapping because it is nearly identical to Table 1. However, many instances of our market game have separating PBE that cannot be mapped to the separating PBE of the signaling game. For example, when $Y_H > f_{H,\emptyset} > Y_T > Y_{-H}$, the set of separating PBE where Alice reports $r_A \in [Y_T, f_{H,\emptyset})$ after receiving the H signal is left unmapped. As a class of games, our market game in general has more equilibria than the signaling game.

7 Conclusion and Future Directions

We study the strategic play of prediction market participants when there exist outside incentives for the participants to manipulate the market probability. The main insight from our analysis is that conflicting incentives inside and outside of a prediction market do not necessarily damage information aggregation in equilibrium. In particular, under certain conditions, there are equilibria in which full information aggregation can be achieved. However, there are also many situations where information loss is inevitable.

Although we only consider a 2-player model, our results remain valid for a much more general setting. Our results can be easily extended to a setting in which multiple participants trade in the market after Alice and before the end of the market, as long as each participant only trades once in the market. Moreover, if there are participants trading before Alice in the market, our results can be extended to this setting if all of the private information of the participants trading before Alice are completely revealed before Alice's stage of participation.

An immediate future direction is to consider a more general setting when Alice's signal has more than 2 realizations. As suggested by our analysis, with more realized signals, Alice's equilibrium behavior could become much more complicated depending on how these realized signals influence her payoffs from inside and outside of the market.

More broadly, one important future direction is to better understand information aggregation mechanisms in the context of decision making, and design mechanisms to minimize or control potential loss in information aggregation and social welfare when there are conflicting incentives within and outside of the mechanism.

Acknowledgements

This material is based upon work supported by NSF Grant No. CCF-0953516. Xi Alice Gao is partially supported by an NSERC Postgraduate Scholarship and a Siebel Scholarship. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors alone.

References

1. Berg, J.E., Forsythe, R., Nelson, F.D., Rietz, T.A.: Results from a dozen years of election futures markets research. In: C.A. Plott, V. Smith (eds.) *Handbook of Experimental Economic Results* (2001)
2. Boutilier, C.: Eliciting forecasts from self-interested experts: scoring rules for decision makers. In: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 2, AAMAS '12*, pp. 737–744 (2012)
3. Camerer, C.F.: Can asset markets be manipulated? A field experiment with race-track betting. *Journal of Political Economy* **106**(3), 457–482 (1998)
4. Chen, K.Y., Plott, C.R.: Information aggregation mechanisms: Concept, design and implementation for a sales forecasting problem (2002). Working paper No. 1131, California Institute of Technology
5. Chen, Y., Dimitrov, S., Sami, R., Reeves, D., Pennock, D.M., Hanson, R., Fortnow, L., Gonen, R.: Gaming prediction markets: Equilibrium strategies with a market maker. *Algorithmica* **58**(4), 930–969 (2010)
6. Chen, Y., Gao, X.A., Goldstein, R., Kash, I.A.: Market Manipulation with Outside Incentives. In: *Proceedings of the 25th AAAI Conference on Artificial Intelligence (AAAI'11)* (2011)
7. Chen, Y., Kash, I., Ruberry, M., Shnayder, V.: Decision Markets with Good Incentives. In: *Proceedings of the Seventh Workshop on Internet and Network Economics (WINE)* (2011)
8. Chen, Y., Kash, I.A.: Information elicitation for decision making. *AAMAS'11: Proc. of the 10th Int. Conf. on Autonomous Agents and Multiagent Systems* (2011)
9. Chen, Y., Pennock, D.M.: A utility framework for bounded-loss market makers. *Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence* pp. 49–56 (2007)

10. Cho, I.K., Kreps, D.M.: Signalling games and stable equilibria. *Quarterly Journal of Economics* **102**, 179–221 (1987)
11. Debnath, S., Pennock, D.M., Giles, C.L., Lawrence, S.: Information incorporation in online in-game sports betting markets. In: *EC '03: Proceedings of the 4th ACM conference on Electronic commerce*, pp. 258–259. ACM, New York, NY, USA (2003). DOI <http://doi.acm.org/10.1145/779928.779987>
12. Dimitrov, S., Sami, R.: Composition of markets with conflicting incentives. *EC'10: Proc. of the 11th ACM Conf. on Electronic Commerce* pp. 53–62 (2010)
13. Forsythe, R., Nelson, F., Neumann, G.R., Wright, J.: Anatomy of an experimental political stock market. *American Economic Review* **82**(5), 1142–61 (1992). URL <http://ideas.repec.org/a/aea/aecrev/v82y1992i5p1142-61.html>
14. Forsythe, R., Rietz, T.A., Ross, T.W.: Wishes, expectations and actions: a survey on price formation in election stock markets. *Journal of Economic Behavior & Organization* **39**(1), 83–110 (1999). URL <http://ideas.repec.org/a/eee/jeborg/v39y1999i1p83-110.html>
15. Fudenberg, D., Tirole, J.: *Game Theory*. MIT Press (1991)
16. Gao, X.A., Zhang, J., Chen, Y.: What you jointly know determines how you act: strategic interactions in prediction markets. In: *Proceedings of the fourteenth ACM conference on Electronic commerce, EC '13*, pp. 489–506 (2013)
17. Hansen, J., Schmidt, C., Strobel, M.: Manipulation in political stock markets—preconditions and evidence. *Applied Economics Letters* **11**(7), 459–463 (2004)
18. Hanson, R.: Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets* **1**(1), 3–15 (2007). URL <http://econpapers.repec.org/RePEc:buc:jpredm:v:1:y:2007:i:1:p:3-15>
19. Hanson, R.D., Oprea, R., Porter, D.: Information aggregation and manipulation in an experimental market. *Journal of Economic Behavior and Organization* **60**(4), 449–459 (2007)
20. Iyer, K., Johari, R., Moallemi, C.C.: Information aggregation in smooth markets. *EC'10: Proc. of the 11th ACM Conf. on Electronic Commerce* pp. 199–206 (2010)
21. Jian, L., Sami, R.: Aggregation and manipulation in prediction markets: effects of trading mechanism and information distribution. *EC'10: Proc. of the 11th ACM Conf. on Electronic Commerce* pp. 207–208 (2010)
22. Lacombe, C.A., Barnett, J.A., Pan, Q.: The imagination market. *Information Systems Frontiers* **9**(2-3), 245–256 (2007)
23. Mas-Colell, A., Whinston, M.D., Green, J.R.: *Microeconomic Theory*. Oxford University Press (1995). URL <http://www.worldcat.org/isbn/0195073401>
24. Ostrovsky, M.: Information aggregation in dynamic markets with strategic traders. *Econometrica* (2011)
25. Othman, A., Sandholm, T.: Decision rules and decision markets. *Proc. of the 9th Int. Conf. on Autonomous Agents and Multiagent Systems* pp. 625–632 (2010)
26. Rhode, P.W., Strumpf, K.S.: Historical presidential betting markets. *Journal of Economic Perspectives* **18**(2), 127–142 (2004)
27. Shi, P., Conitzer, V., Guo, M.: Prediction mechanisms that do not incentivize undesirable actions. *WINE'09: Internet and Network Economics* pp. 89–100 (2009)
28. Spence, M.: Job market signalling. *Quarterly Journal of Economics* **87**(3), 355–374 (1973)
29. Wolfers, J., Zitzewitz, E.: Prediction markets. *Journal of Economic Perspective* **18**(2), 107–126 (2004)

A Proof of Proposition 1

Proof The partial derivative of the loss function with respect to r_A is

$$\frac{\partial}{\partial r_A} L(f_{s_A, \theta}, r_A) = \frac{r_A - f_{s_A, \theta}}{r_A(1 - r_A)}. \quad (34)$$

It is negative for $r_A < f_{s_A, \theta}$, zero for $r_A = f_{s_A, \theta}$ and positive for $r_A > f_{s_A, \theta}$. Thus, the loss function is strictly increasing for $r_A \in [f_{s_A, \theta}, 1)$ and strictly decreasing for $r_A \in (0, f_{s_A, \theta}]$. In addition, note that $L(f_{s_A, \theta}, r_A) \rightarrow \infty$ as $r_A \rightarrow 0$ or $r_A \rightarrow 1$ for any fixed $f_{s_A, \theta}$. Hence, the loss function has the range $[0, \infty)$ for both $r_A \in [f_{s_A, \theta}, 1)$ and $r_A \in (0, f_{s_A, \theta}]$.

The partial derivative of the loss function with respect to $f_{s_A, \theta}$ is

$$\frac{\partial}{\partial f_{s_A, \theta}} L(f_{s_A, \theta}, r_A) = \log \left(\frac{f_{s_A, \theta}}{1 - f_{s_A, \theta}} \frac{1 - r_A}{r_A} \right). \quad (35)$$

It equals zero when $f_{s_A, \theta} = r_A$, negative when $f_{s_A, \theta} < r_A$ and positive when $f_{s_A, \theta} > r_A$. Therefore, for a fixed $r_A \in [0, 1]$, $L(f_{s_A, \theta}, r_A)$ is strictly decreasing for $f_{s_A, \theta} \in [0, r_A]$ and strictly increasing for $f_{s_A, \theta} \in [r_A, 1]$. \square

B Example 1

Example 1 Suppose the outside payoff function is $Q(r_B) = r_B$, and the prior distribution is given by Table 2.

	$X = 1$		$X = 0$		
	$S_A = H$	$S_A = T$		$S_A = H$	$S_A = T$
$S_B = H$	0.54	0.054	$S_B = H$	0	0.006
$S_B = T$	0.036	0	$S_B = T$	0.324	0.04

Table 2 An example prior distribution. Each cell gives the value of $P(X, S_A, S_B)$ for the corresponding realizations of X , S_A , and S_B .

It is easy to compute $f_{H,\emptyset} = 0.64$, $f_{T,\emptyset} = 0.54$, $f_{H,H} = 1$, $f_{H,T} = 0.1$, $f_{T,H} = 0.9$, and $f_{T,T} = 0$.

Alice's expected loss in market scoring rule payoff when receiving the T signal but changing the market probability to $f_{H,\emptyset}$ is

$$L(f_{T,\emptyset}, f_{H,\emptyset}) = f_{T,\emptyset} \log \frac{f_{T,\emptyset}}{f_{H,\emptyset}} + (1 - f_{T,\emptyset}) \log \frac{1 - f_{T,\emptyset}}{1 - f_{H,\emptyset}} = \left(0.54 \log \frac{0.54}{0.64} + 0.46 \log \frac{0.46}{0.36} \right) \approx 0.021. \quad (36)$$

Alice's expected gain in outside payoff when receiving the T signal but convincing Bob that she has the H signal is

$$\begin{aligned} & E_{S_B}[Q(f_{H,S_B}) - Q(f_{T,S_B}) \mid S_A = T] \\ &= P(S_B = H \mid S_A = T)(f_{H,H} - f_{T,H}) + P(S_B = T \mid S_A = T)(f_{H,T} - f_{T,T}) \\ &= 0.6(1 - 0.9) + 0.4(0.1 - 0) \end{aligned} \quad (37)$$

$$= 0.1. \quad (38)$$

It is clear that $L(f_{T,\emptyset}, f_{H,\emptyset}) < E_{S_B}[Q(f_{H,S_B}) - Q(f_{T,S_B}) \mid S_A = T]$. Thus, inequality (5) is satisfied and a truthful PBE does not exist.

In addition to the above derivation, we note that even though a truthful PBE does not exist for this example, a separating PBE does exist. The intuition behind this can be shown by calculating and comparing the quantities Y_H , Y_T , and $f_{H,\emptyset}$, as illustrated below. We solve for Y_T by solving the following equation:

$$L(f_{T,\emptyset}, Y_T) = E_{S_B}[Q(f_{H,S_B}) - Q(f_{T,S_B}) \mid S_A = T] \quad (39)$$

$$\Rightarrow 0.54 \log \frac{0.54}{Y_T} + 0.46 \log \frac{0.46}{1 - Y_T} = 0.1 \quad (40)$$

$$\Rightarrow Y_T \approx 0.747 \quad (41)$$

Similarly, we solve for Y_H below:

$$L(f_{H,\emptyset}, Y_H) = E_{S_B}[Q(f_{H,S_B}) - Q(f_{T,S_B}) \mid S_A = H] \quad (42)$$

$$\Rightarrow 0.64 \log \frac{0.64}{Y_H} + 0.36 \log \frac{0.36}{1 - Y_H} = 0.1 \quad (43)$$

$$\Rightarrow Y_H \approx 0.827 \quad (44)$$

The above calculations show that we have $f_{H,\emptyset} < Y_T < Y_H$. Thus, a truthful PBE does not exist because if Alice reports $f_{H,\emptyset}$ in the first stage, then Bob will believe that there is positive probability that Alice actually received a T signal but is trying to pretend that she received a H signal, since $f_{H,\emptyset} < Y_T$. However, since $Y_H > Y_T$, a separating equilibrium exists because Alice can establish credibility with Bob by reporting any value in $[Y_T, Y_H]$ in the first stage.

Lastly, note that this example illustrates a prior distribution for which the signals of Alice and Bob are independent. In Proposition 4, we will prove that when Alice and Bob have independent signals, $Y_H > Y_T$ must be satisfied. \square

C Proof of Proposition 3

Proof If $f_{H,0} \geq Y_T \geq f_{T,0} \geq Y_{-T}$, then it is easy to see that $L(f_{H,0}, Y_T) \leq L(f_{H,0}, Y_{-T})$ and the equality holds only when $Y_T = f_{T,0} = Y_{-T}$. The remainder of the proof focuses on the case when $f_{H,0} < Y_T$.

By definitions of Y_T and Y_{-T} , we have

$$L(f_{T,0}, Y_T) = L(f_{T,0}, Y_{-T}). \quad (45)$$

By Proposition 1 and $Y_{-T} \leq f_{T,0} < f_{H,0}$, we have

$$L(f_{T,0}, Y_{-T}) < L(f_{H,0}, Y_{-T}). \quad (46)$$

By Proposition 1 and $f_{T,0} < f_{H,0} \leq Y_T$, we have

$$L(f_{H,0}, Y_T) < L(f_{T,0}, Y_T). \quad (47)$$

Hence, we must have $L(f_{H,0}, Y_T) < L(f_{H,0}, Y_{-T})$ due to equation (45) and inequalities (46) and (47), as

$$L(f_{H,0}, Y_T) < L(f_{T,0}, Y_T) = L(f_{T,0}, Y_{-T}) < L(f_{H,0}, Y_{-T}). \quad (48)$$

□

D Example 2

Example 2 Consider the outside payoff function and the prior distribution in Table 3. We show below that there exists sufficiently small ε such that $Y_H < Y_T$.

	$X = 1$		$X = 0$		
	$S_A = H$	$S_A = T$	$S_A = H$	$S_A = T$	
$S_B = H$	ε	ε	$S_B = H$	0	$0.5 - 2\varepsilon$
$S_B = T$	ε	0	$S_B = T$	$0.5 - 2\varepsilon$	ε

Table 3 An example prior distribution with $\varepsilon \in (0, 0.25)$. Each cell gives the value of $P(X, S_A, S_B)$ for the corresponding realizations of X , S_A , and S_B .

It is easy to compute $f_{H,0} = 4\varepsilon$, $f_{T,0} = 2\varepsilon$, $f_{H,H} = 1$, $f_{H,T} = \frac{\varepsilon}{0.5-\varepsilon}$, $f_{T,H} = \frac{\varepsilon}{0.5-\varepsilon}$, and $f_{T,T} = 0$. With this, we can calculate

$$\begin{aligned} L(f_{H,0}, Y_H) &= E_{S_B}[Q(f_{H,S_B}) - Q(f_{T,S_B}) | S_A = H] \\ &= P(S_B = H | S_A = H)(f_{H,H} - f_{T,H}) + P(S_B = T | S_A = H)(f_{H,T} - f_{T,T}) \\ &= (2\varepsilon) \left(1 - \frac{\varepsilon}{0.5-\varepsilon}\right) + (1-2\varepsilon) \left(\frac{\varepsilon}{0.5-\varepsilon} - 0\right). \end{aligned}$$

As ε approaches 0, we have

$$\lim_{\varepsilon \rightarrow 0} L(f_{H,0}, Y_H) = 0. \quad (49)$$

Because $\lim_{\varepsilon \rightarrow 0} f_{H,0} = \lim_{\varepsilon \rightarrow 0} 4\varepsilon = 0$, by definition of $L(f_{H,0}, Y_H)$, (49) implies that

$$\lim_{\varepsilon \rightarrow 0} Y_H = 0. \quad (50)$$

Similarly, we have

$$\begin{aligned} L(f_{T,0}, Y_T) &= E_{S_B}[Q(f_{H,S_B}) - Q(f_{T,S_B}) | S_A = T] \\ &= P(S_B = H | S_A = T)(f_{H,H} - f_{T,H}) + P(S_B = T | S_A = T)(f_{H,T} - f_{T,T}) \\ &= (1-2\varepsilon) \left(1 - \frac{\varepsilon}{0.5-\varepsilon}\right) + 2\varepsilon \left(\frac{\varepsilon}{0.5-\varepsilon} - 0\right). \end{aligned}$$

As ε approaches 0, we have

$$\lim_{\varepsilon \rightarrow 0} L(f_{T,\emptyset}, Y_T) = 1. \quad (51)$$

Because $\lim_{\varepsilon \rightarrow 0} f_{T,\emptyset} = \lim_{\varepsilon \rightarrow 0} 2\varepsilon = 0$, by definition of $L(f_{T,\emptyset}, Y_T)$,

$$\lim_{\varepsilon \rightarrow 0} L(f_{T,\emptyset}, Y_T) = -\log(1 - \lim_{\varepsilon \rightarrow 0} Y_T).$$

Given (51), we have

$$\lim_{\varepsilon \rightarrow 0} Y_T = 1 - 1/e. \quad (52)$$

Combining (50) and (52), we know that when ε is sufficiently small, $Y_H < Y_T$.

In addition to the above derivation, we describe some qualitative properties of the given prior distribution, which may be helpful in highlighting the intuitions behind the $Y_H < Y_T$ condition. For this prior distribution, Alice is willing to report a higher value after receiving the T signal due to the combined effect of two factors. First, note that when Alice has the T signal, Bob is far more likely to have the H signal than the T signal for sufficiently small ε . This is shown by

$$\lim_{\varepsilon \rightarrow 0} \mathbb{P}(S_B = H | S_A = T) = \lim_{\varepsilon \rightarrow 0} (1 - 2\varepsilon) = 1, \quad (53)$$

$$\lim_{\varepsilon \rightarrow 0} \mathbb{P}(S_B = T | S_A = T) = \lim_{\varepsilon \rightarrow 0} 2\varepsilon = 0. \quad (54)$$

Second, Alice's maximum gain in outside payoff when she has the T signal but manages to convince Bob that she has the H signal is much higher when Bob has the H signal than when he has the T signal for sufficiently small ε . When Bob has the H signal, the maximum gain for Alice is

$$\lim_{\varepsilon \rightarrow 0} (f_{HH} - f_{TH}) = \lim_{\varepsilon \rightarrow 0} \left(1 - \frac{\varepsilon}{0.5 - \varepsilon}\right) = 1, \quad (55)$$

which is greater than the maximum gain for Alice when Bob has the T signal,

$$\lim_{\varepsilon \rightarrow 0} (f_{HT} - f_{TT}) = \lim_{\varepsilon \rightarrow 0} \left(\frac{\varepsilon}{0.5 - \varepsilon} - 0\right) = 0. \quad (56)$$

Thus, when Alice has the T signal, Bob is more likely to have the H signal, resulting in a higher expected gain in outside payoff for Alice by convincing Bob that she has the H signal. This intuitively explains why Y_T is high.

In Example 1, we describe a prior distribution and outside function and show that a truthful PBE does not exist when $f_{H,\emptyset} < Y_T \leq Y_H$. Note that guaranteeing $f_{H,\emptyset} < Y_T \leq Y_H$ is not the only way for a truthful PBE to fail to exist. For instance, this example shows that, when $f_{H,\emptyset} \leq Y_H < Y_T$, a truthful PBE also fails to exist. \square

E Proof of Theorem 3

Proof If $Y_T \geq f_{H,\emptyset}$, the interval $[\max(Y_{-H}, Y_T), Y_H]$ can be written as $[\max(f_{H,\emptyset}, Y_T), Y_H]$ because $Y_{-H} \leq f_{H,\emptyset}$. If $Y_T < f_{H,\emptyset}$, the interval $[\max(Y_{-H}, Y_T), Y_H]$ can be split into two intervals $[\max(Y_{-H}, Y_T), f_{H,\emptyset}]$ and $[\max(f_{H,\emptyset}, Y_T), Y_H]$. In the following, we first consider the case $r_A \in [\max(f_{H,\emptyset}, Y_T), Y_H]$; then, for $Y_T < f_{H,\emptyset}$, we consider the case $r_A \in [\max(Y_{-H}, Y_T), f_{H,\emptyset}]$.

First, suppose that Alice reports $r_A \in [\max(f_{H,\emptyset}, Y_T), Y_H]$ after receiving the H signal. Fix a particular $k \in [\max(f_{H,\emptyset}, Y_T), Y_H]$. We prove that the following pair of Alice's strategy and Bob's belief forms a separating PBE of our game:

$$SE_2(k) : \begin{cases} \sigma_H^S(k) = 1, \sigma_T^S(f_{T,\emptyset}) = 1 \\ \mu_{s_B, r_A}^S(H) = \begin{cases} 1, & \text{if } r_A \in [k, 1] \\ 0, & \text{if } r_A \in [0, k) \end{cases} \end{cases} \quad (57)$$

We'll show that Alice's strategy is optimal given Bob's belief. If Alice receives the T signal, she does not report any $r_A > Y_T$ by definition of Y_T . She may be indifferent between reporting Y_T and $f_{T,\emptyset}$. For any $r_A < Y_T$, we have $r_A < k$ and Bob's belief sets $\mu_{s_B, r_A}^S(H) = 0$ for any $r_A < k$. So for $r_A < Y_T$, reporting $r_A = f_{T,\emptyset}$ dominates reporting any other value. Thus, it is optimal for Alice to report $f_{T,\emptyset}$ when having the T signal.

If Alice receives the H signal, according to the definitions of Y_{-H} and Y_H , she would only report values in $[Y_{-H}, Y_H]$. Given Bob's belief, Alice would only report some $r_A \in [k, Y_H]$. Because $f_{H,\emptyset} \leq k$, Alice maximizes

her expected market scoring rule payoff by reporting $r_A = k$ for any $r_A \in [k, Y_H]$. Therefore, it is optimal for Alice to report $r_A = k$ after receiving the H signal.

We can show that Bob's belief is consistent with Alice's strategy by mechanically applying Bayes' rule (argument omitted). Hence, for each $k \in [\max(Y_T, f_{H,\emptyset}), Y_H]$, $SE_2(k)$ is a separating PBE of our game.

Next, we assume $Y_T < f_{H,\emptyset}$ and consider that Alice reports $r_A \in [\max(Y_{-H}, Y_T), f_{H,\emptyset}]$ after receiving the H signal. For every $k \in [\max(Y_{-H}, Y_T), f_{H,\emptyset}]$, we prove that the following pair of Alice's strategy and Bob's belief forms a separating PBE of our game:

$$SE_3(k) : \begin{cases} \sigma_H^S(k) = 1, \sigma_T^S(f_{T,\emptyset}) = 1 \\ \mu_{s_B, r_A}^S(H) = \begin{cases} 1, & \text{if } r_A = k \\ 0, & \text{if } r_A \in [0, k) \cup (k, 1] \end{cases} \end{cases} . \quad (58)$$

We'll show that Alice's strategy is optimal given Bob's belief. If Alice receives the T signal, she does not report any $r_A > Y_T$ by definition of Y_T and is at best indifferent between reporting Y_T and reporting $f_{T,\emptyset}$. For any $r_A \in [0, Y_T)$, Bob's belief sets $\mu_{s_B, r_A}^S(H) = 0$ since $k \geq Y_T$. For any $r_A \in [0, Y_T)$, Alice maximizes her expected market scoring rule payoff by reporting $r_A = f_{T,\emptyset}$. Thus, it is optimal for Alice to report $f_{T,\emptyset}$ when having the T signal.

If Alice receives the H signal, for any $r_A \in [0, 1] \setminus \{k\}$, Alice maximizes her expected market scoring rule payoff by reporting $f_{H,\emptyset}$. By definition, we know that $Y_{-H} \leq k < f_{H,\emptyset}$. Given Bob's belief

$$L(f_{H,\emptyset}, k) \leq L(f_{H,\emptyset}, Y_{-H}) = E_{s_B}[Q(f_{H, s_B}) - Q(f_{T, s_B}) \mid s_A] \quad (59)$$

By switching from reporting $f_{H,\emptyset}$ to reporting k , Alice's expected gain in outside payoff is greater than or equal to her loss in her expected market scoring rule payoff. So she weakly prefers reporting k to reporting $f_{H,\emptyset}$. By enforcing the consistency with Bob's belief, Alice's strategy must be to report k after receiving the H signal.

We can show that Bob's belief is consistent with Alice's strategy by mechanically applying Bayes' rule (argument omitted). Hence, if $Y_T < f_{H,\emptyset}$, for each $k \in [\max(Y_{-H}, Y_T), f_{H,\emptyset}]$, $SE_3(k)$ is a separating PBE of this game. \square

F Proof of Theorem 4

Proof If $Y_{-H} \leq Y_{-T}$, for every $k \in [Y_{-H}, Y_{-T}]$, we prove the following pair of Alice's strategy and Bob's belief forms a separating PBE of our game:

$$SE_4(k) : \begin{cases} \sigma_H^S(k) = 1, \sigma_T^S(f_{T,\emptyset}) = 1 \\ \mu_{s_B, r_A}^S(H) = \begin{cases} 0, & \text{if } r_A \in (k, 1] \\ 1, & \text{if } r_A \in [0, k] \end{cases} \end{cases} . \quad (60)$$

We'll show that Alice's strategy is optimal given Bob's belief. If Alice receives the T signal, she does not report any $r_A < Y_{-T}$ by definition of Y_{-T} and is at best indifferent between reporting Y_{-T} and $f_{T,\emptyset}$. For any $r_A > Y_{-T}$, because Bob's belief sets $\mu_{s_B, r_A}^S(H) = 0$, reporting $f_{T,\emptyset}$ dominates reporting any other value in this range. Thus, it is optimal for Alice to report $f_{T,\emptyset}$ when having the T signal.

If Alice receives the H signal, for any $r_A \in (k, 1]$, Alice maximizes her expected market scoring rule payoff by reporting $r_A = f_{H,\emptyset}$. For any $r_A \in [Y_{-H}, k]$, Alice maximizes her expected market scoring rule payoff by reporting $r_A = k$. By definition of Y_{-H} , Alice is better off reporting k than reporting $f_{H,\emptyset}$ since

$$L(f_{H,\emptyset}, k) \leq L(f_{H,\emptyset}, Y_{-H}) = E_{s_B}[Q(f_{H, s_B}) - Q(f_{T, s_B}) \mid s_A] \quad (61)$$

Therefore, it is optimal for Alice to report k after receiving the H signal.

We can show that Bob's belief is consistent with Alice's strategy by mechanically applying Bayes' rule (argument omitted). Hence, if $Y_{-H} \leq Y_{-T}$, for every $k \in [Y_{-H}, Y_{-T}]$, $SE_4(k)$ is a separating PBE. \square

G Proof of Proposition 5

Proof The existence of a separating PBE requires $Y_H \geq Y_T$ by Theorem 2. By Lemma 2, we have $\sigma_T(f_{T,\emptyset}) = 1$ at any separating PBE.

By Theorem 3, for every $r_A \in [\max(Y_{-H}, Y_T), Y_H]$, there exists a pure strategy separating PBE in which Alice reports r_A with probability 1 after receiving the H signal. Now suppose that Bob's belief satisfies the domination-based refinement. Consider 2 cases.

(1) Assume that $f_{H,\emptyset} > Y_T$. Then we must have $f_{H,\emptyset} \in [\max(Y_{-H}, Y_T), Y_H]$. By Lemma 5, Bob's belief must set $\mu_{s_B, f_{H,\emptyset}}(T) = 0$. Thus, reporting $f_{H,\emptyset}$ is strictly optimal for Alice since reporting $f_{H,\emptyset}$ strictly maximizes Alice's expected market scoring rule payoff and weakly maximizes Alice's expected outside payoff. Therefore, there are no longer pure strategy separating PBE in which Alice reports $r_A \in [\max(Y_{-H}, Y_T), Y_H] \setminus \{f_{H,\emptyset}\}$ after receiving the H signal.

(2) Assume that $f_{H,\emptyset} \leq Y_T$. Then we must have $Y_{-H} < Y_T$, and the interval $[\max(Y_{-H}, Y_T), Y_H]$ can be reduced to $[Y_T, Y_H]$. By Lemma 5, Bob's belief must set $\mu_{s_B, r_A}(T) = 0$ for any $r_A \in (Y_T, Y_H]$. If Alice receives the H signal, given Bob's belief, Alice would not report any $r_A \in (Y_T, Y_H]$ because there always exists a $r'_A \in (Y_T, r_A)$ such that reporting r'_A is strictly better than reporting r_A for Alice. Therefore, there no longer exist pure strategy separating PBE in which Alice reports $r_A \in [\max(Y_{-H}, Y_T), Y_H] \setminus \{Y_T\}$ after receiving the H signal.

Hence, Alice would not report $r_A \in [\max(Y_{-H}, Y_T), Y_H] \setminus \max(f_{H,\emptyset}, Y_T)$ after receiving the H signal at any separating PBE satisfying the domination-based belief refinement.

By Theorem 4, if $Y_H \geq Y_T$ and $Y_{-H} \leq Y_{-T}$, for every $r_A \in [Y_{-H}, Y_{-T}]$, there exists a pure strategy separating PBE in which Alice reports r_A with probability 1 after receiving the H signal. By Lemma 5, Bob's belief must set $\mu_{s_B, r_A}(T) = 0$ for any $r_A \in [Y_{-H}, Y_{-T}]$. Then, if Alice receives the H signal, given Bob's belief, Alice would not report any $r_A \in [Y_{-H}, Y_{-T}]$ because there always exists a $r'_A \in (r_A, Y_{-T})$ such that reporting r'_A is strictly better than reporting r_A for Alice.

Also, Alice would not report Y_{-T} after receiving the H signal for the following reasons. We consider 2 cases. If $f_{H,\emptyset} \geq Y_T$, then Alice's market scoring rule payoff is strictly better by reporting $f_{H,\emptyset}$ than reporting Y_{-T} . Otherwise, if $f_{H,\emptyset} < Y_T$, we know that $Y_{-T} < Y_T$ and hence, by Proposition 3, $L(f_{H,\emptyset}, Y_{-T}) < L(f_{H,\emptyset}, Y_T)$. Consider $r_A = Y_T + \varepsilon$ for a small $\varepsilon > 0$ such that $L(f_{H,\emptyset}, Y_{-T}) > L(f_{H,\emptyset}, r_A)$. Such an ε must exist because as $\varepsilon \rightarrow 0$, $L(f_{H,\emptyset}, r_A) \rightarrow L(f_{H,\emptyset}, Y_T)$. Alice's market scoring rule payoff is strictly better by reporting r_A than reporting Y_{-T} . Given Bob's belief, we know that $\mu_{s_B, r_A}(T) = 0$ and $\mu_{s_B, Y_{-T}}(T) \geq 0$. So Alice's outside payoff is weakly better when reporting r_A than reporting Y_{-T} . Therefore, reporting $r_A = Y_T + \varepsilon$ strictly dominates reporting Y_{-T} .

Hence, there are no longer pure strategy separating PBE in which Alice reports $r_A \in [Y_{-H}, Y_{-T}]$ after receiving the H signal.

It remains to show that there exists a belief for Bob satisfying the refinement so that Alice's strategy $\sigma_H(\max(f_{H,\emptyset}, Y_T)) = 1$, $\sigma_T(f_{T,\emptyset}) = 1$ and Bob's belief form a PBE. It is straightforward to verify that Bob's belief in the PBE SE_1 described in (12) is such a belief. \square

H Proof of Proposition 7

Proof Let r be the unique value in $[0, f_{H,\emptyset}]$ satisfying $L(f_{H,\emptyset}, r) = L(f_{H,\emptyset}, Y_T)$. Consider a PBE satisfying the domination-based refinement. We will show that there exists an $\varepsilon > 0$ such that if Alice receives the H signal, then reporting any $r_A \leq r$ is strictly worse than reporting $Y_T + \varepsilon$.

By definition of r , we have that $L(f_{H,\emptyset}, r_A) \geq L(f_{H,\emptyset}, r) = L(f_{H,\emptyset}, Y_T), \forall r_A \leq r$. We consider 2 cases.

(1) $r_A < r$: Choose any $0 < \varepsilon < r - r_A$, then we must have $L(f_{H,\emptyset}, r_A) > L(f_{H,\emptyset}, r - \varepsilon) = L(f_{H,\emptyset}, Y_T + \varepsilon)$. Since the PBE satisfies the domination-based refinement, then Bob's belief must set $\mu_{s_B, r_A}(T) = 0, \forall r_A \in (Y_T, Y_H]$. Alice's expected outside payoff by reporting $Y_T + \varepsilon$ is weakly better than her expected outside payoff by reporting r_A . Therefore, for any $\varepsilon \in (0, r - r_A)$, Alice is strictly worse off reporting any $r_A < r$ than reporting $Y_T + \varepsilon$.

(2) $r_A = r$: For any small $\varepsilon > 0$, we have that $L(f_{H,\emptyset}, r) < L(f_{H,\emptyset}, Y_T + \varepsilon)$. However as $\varepsilon \rightarrow 0$, $L(f_{H,\emptyset}, Y_T + \varepsilon) - L(f_{H,\emptyset}, r) \rightarrow 0$. Since $r < f_{H,\emptyset} < Y_T$, if Alice reports r after receiving H signal at any PBE, then Bob's belief must set $\mu_{s_B, r}(T) > 0$. Since the PBE satisfies the domination-based refinement, then Bob's belief must set $\mu_{s_B, Y_T + \varepsilon}(T) = 0, \text{ for any } 0 < \varepsilon \leq Y_H - Y_T$. Regardless of ε , Alice's expected outside payoff by reporting $Y_T + \varepsilon$ is strictly better than her expected outside payoff by reporting r . However, as ε approaches 0, the difference between Alice's expected market scoring rule payoff for these two reports goes to 0. Hence, there

must exist $\varepsilon > 0$ such that Alice's total expected payoff by reporting r is strictly less than her total expected payoff by reporting $Y_T + \varepsilon$. \square

I Proof of Theorem 10

Proof We will show that among all pure strategy separating PBE of our game, Bob's expected payoff is maximized in $SE_2(Y_H)$, defined in equation (57).

In all separating PBE, the sum of Alice and Bob's expected payoffs inside the market is the same. Thus, the separating PBE that maximizes Bob's payoff is also the separating PBE that minimizes Alice's payoff.

By Lemma 2, in any separating PBE, Alice must report $f_{T,\emptyset}$ after receiving the T signal. Therefore, Alice's expected payoff after receiving the T signal is the same at any separating PBE.

For any separating PBE, Alice may report $r \in [Y_{-H}, Y_H]$ after receiving the H signal. In $[Y_{-H}, Y_H]$, reporting Y_H or Y_{-H} maximizes Alice's loss in her market scoring rule payoff and thus minimizes Alice's expected payoff after receiving the H signal. Reporting Y_H corresponds to the separating PBE $SE_2(Y_H)$ and reporting Y_{-H} corresponds to the separating PBE $SE_4(Y_{-H})$.

If $Y_{-H} \leq Y_{-T}$, the separating PBE $SE_4(Y_{-H})$ exists, by the proof of Theorem 4 in Appendix F. We know that $Y_{-H} \leq Y_{-T}$ implies $Y_H \geq Y_T$ by the proof of Theorem 2. Thus, when the separating PBE $SE_4(Y_{-H})$ exists, the separating PBE $SE_2(Y_H)$ also exists and Alice's total expected payoff at these two separating PBE are the same.

If $Y_{-H} > Y_{-T}$, the separating PBE $SE_4(Y_{-H})$ does not exist. However, if any separating PBE exists, then we must have $Y_H \geq Y_T$, and the separating PBE $SE_2(Y_H)$ must exist.

Hence, the separating PBE $SE_2(Y_H)$ maximizes Bob's expected payoff among all separating PBE of our game. \square

J Proof of Theorem 11

Proof Sufficient condition

First, we show that satisfying at least one of the two pairs of inequalities is a sufficient condition for a separating PBE to exist for our game.

If inequalities (30) and (31) are satisfied, we can show that SE_5 is a separating PBE of our game.

$$SE_5 : \begin{cases} \sigma_H^S(\max(Y_T, f_{H,\emptyset})) = 1, \sigma_T^S(f_{T,\emptyset}) = 1 \\ \mu_{s_B, r_A}^S(H) = \begin{cases} 1, & \text{if } r_A \in [Y_T, 1] \\ 0, & \text{if } r_A \in [0, Y_T] \end{cases} \end{cases} \quad (62)$$

First, we show that Alice's strategy is optimal given Bob's belief. Since inequality (30) is satisfied, Y_T is a well defined value in $[f_{T,\emptyset}, 1]$. If $f_{H,\emptyset} < Y_T$, then it is optimal for Alice to report Y_T after receiving the H signal because her gain in outside payoff is greater than her loss in the market scoring rule payoff by inequality (31). Otherwise, if $f_{H,\emptyset} \geq Y_T$, then it's optimal for Alice to report $f_{H,\emptyset}$ after receiving the H signal. Therefore, Alice's optimal strategy after receiving the H signal is to report $\max(f_{H,\emptyset}, Y_T)$. When Alice receives the T signal, Alice would not report any $r_A \geq Y_T$ by definition of Y_T . Any other report $r_A \in [0, Y_T]$ is dominated by a report of $f_{T,\emptyset}$ given Bob's belief. Therefore, it is optimal for Alice to report $f_{T,\emptyset}$ after receiving the T signal. Moreover, we can show that Bob's belief is consistent with Alice's strategy by mechanically applying Bayes' rule (argument omitted). Given the above arguments, SE_5 is a separating PBE of this game.

Similarly, if inequalities (32) and (33) are satisfied, then we can show that SE_6 is a separating PBE of our game.

$$SE_6 : \begin{cases} \sigma_H^S(Y_{-T}) = 1, \sigma_T^S(f_{T,\emptyset}) = 1 \\ \mu_{s_B, r_A}^S(H) = \begin{cases} 0, & \text{if } r_A \in (Y_{-T}, 1] \\ 1, & \text{if } r_A \in [0, Y_{-T}] \end{cases} \end{cases} \quad (63)$$

First, we show that Alice's strategy is optimal given Bob's belief. Since inequality (32) is satisfied, Y_{-T} is well defined. If Alice receives the H signal, reporting any $r_A \in [0, Y_{-T}]$ gives her higher outside payoff than reporting any $r_A \in (Y_{-T}, 1]$. For any $r_A \in [0, Y_{-T}]$, her outside payoff is fixed and reporting $r_A = Y_{-T}$ maximizes her market scoring rule payoff. Therefore, it is optimal for Alice to report $r_A = Y_{-T}$ after receiving the H signal. If Alice receives the T signal, she does not report any $r_A < Y_{-T}$ by definition of Y_{-T} . Given Bob's

belief, she is indifferent between reporting Y_{-T} and $f_{T,0}$. For any $r_A > Y_{-T}$, Bob's belief sets $\mu_{s_B, r_A}^S(H) = 0$, so it is optimal for Alice to report $f_{T,0}$ to maximize her market scoring rule payoff. We can show that Bob's belief is consistent with Alice's strategy by mechanically applying Bayes' rule (argument omitted). Hence, SE_6 is a separating PBE of our game.

Necessary condition

Second, we show that, if there exists a separating PBE of our game, then at least one of the two pairs of inequalities must be satisfied. We prove this by contradiction. Suppose that there exists a separating PBE of our game but at least one of the two inequalities in each of the two pairs of inequalities is violated.

Suppose that at least one of the inequalities (30) and (31) is violated. Then, we can show that Alice does not report any value $r_A \in [f_{T,0}, 1]$ after receiving the H signal at any separating PBE. We divide the argument for this into 2 cases.

(1) If inequality (30) is violated, we know that Y_T is not well defined. We show by contradiction that Alice does not report any value in $[f_{T,0}, 1]$ after receiving the H signal. Suppose that at a separating PBE, Alice reports $r_A \in [f_{T,0}, 1]$ with positive probability after receiving the H signal. Since this PBE is separating, Bob's belief must be that $\mu_{s_B, r_A}(H) = 1$ to be consistent with Alice's strategy. By Lemma 2, in any separating PBE, Bob's belief must be $\mu_{s_B, f_{T,0}}(H) = 0$ and Alice must report $f_{T,0}$ after receiving the T signal. Since inequality (30) is violated, then we have that

$$L_m(f_{T,0}, r_A) \leq L_m(f_{T,0}, 1) < E_{S_B}[Q(f_{H, S_B}) - Q(f_{T, S_B}) | S_A = T], \quad (64)$$

so Alice would strictly prefer to report r_A rather than $f_{T,0}$ after receiving the T signal, which is a contradiction.

(2) Otherwise, if inequality (30) is satisfied but inequality (31) is violated, then we know that Y_T is well defined. If $f_{H,0} \geq Y_T$, then inequality (31) is automatically satisfied, so we must have that $f_{H,0} < Y_T$ and $L_m(f_{H,0}, Y_T) > E_{S_B}[Q(f_{H, S_B}) - Q(f_{T, S_B}) | S_A = H]$. Then Alice does not report any $r_A \in [Y_T, 1]$ after receiving the H signal because doing so is dominated by reporting $f_{H,0}$. Next, we can show by contradiction that Alice does not report any $r_A \in [f_{T,0}, Y_T)$ after receiving the H signal at any separating PBE. Suppose that at any separating PBE, Alice reports $r_A \in [f_{T,0}, Y_T)$ with positive probability after receiving the H signal. Since this PBE is separating, Bob's belief must be that $\mu_{s_B, r_A}(H) = 1$ to be consistent with Alice's strategy. By Lemma 2, in any separating PBE, Alice must report $f_{T,0}$ after receiving the T signal and Bob's belief must be $\mu_{s_B, f_{T,0}}(H) = 0$. Thus, for $r_A \in (Y_{-T}, Y_T)$, by definitions of Y_T and Y_{-T} , Alice would strictly prefer to report r_A rather than $f_{T,0}$ after receiving the T signal, which is a contradiction.

Hence, if at least one of the inequalities (30) and (31) is violated, then at any separating PBE, Alice does not report any $r_A \in [f_{T,0}, 1]$ after receiving the H signal.

Similarly, we can show that, if at least one of the inequalities (32) and (33) is violated, Alice does not report any value $r_A \in [0, f_{T,0}]$ after receiving the H signal at any separating PBE. We again consider 2 cases:

(1) If inequality (32) is violated, we know that Y_{-T} is not well defined. Then we can show that Alice does not report any value in $[0, f_{T,0}]$ after receiving the H signal. We prove by contradiction. Suppose that at a separating PBE, Alice reports $r_A \in [0, f_{T,0}]$ with positive probability after receiving the H signal. Since this PBE is separating, Bob's belief must be that $\mu_{s_B, r_A}(H) = 1$ to be consistent with Alice's strategy. By Lemma 2, in any separating PBE, Bob's belief must be $\mu_{s_B, f_{T,0}}(H) = 0$ and Alice must report $f_{T,0}$ after receiving the T signal. Since inequality (32) is violated, we have that

$$L_m(f_{T,0}, r_A) \leq L_m(f_{T,0}, 0) < E_{S_B}[Q(f_{H, S_B}) - Q(f_{T, S_B}) | S_A = T], \quad (65)$$

so Alice would strictly prefer to report r_A rather than $f_{T,0}$ after receiving the T signal, which is a contradiction.

(2) Otherwise, if inequality (32) is satisfied but inequality (33) is violated, then we know that Y_{-T} is well defined. Also, we must have that $L_m(f_{H,0}, Y_{-T}) > E_{S_B}[Q(f_{H, S_B}) - Q(f_{T, S_B}) | S_A = H]$. Then Alice does not report any $r_A \in [0, Y_{-T}]$ after receiving the H signal because doing so is dominated by reporting $f_{H,0}$. Next, we can show by contradiction that at any separating PBE, Alice does not report any $r_A \in (Y_{-T}, f_{T,0}]$ after receiving the H signal. Suppose that at any separating PBE, Alice reports $r_A \in (Y_{-T}, f_{T,0}]$ with positive probability after receiving the H signal. Since this PBE is separating, Bob's belief must be that $\mu_{s_B, r_A}(H) = 1$ to be consistent with Alice's strategy. By Lemma 2, in any separating PBE, Alice must report $f_{T,0}$ after receiving the T signal and Bob's belief must be $\mu_{s_B, f_{T,0}}(H) = 0$. Thus, for $r_A \in (Y_{-T}, Y_T)$, by definitions of Y_T and Y_{-T} , Alice would strictly prefer to report r_A rather than $f_{T,0}$ after receiving the T signal, which is a contradiction.

Hence, if at least one of the inequalities (32) and (33) is violated, in any separating PBE, Alice does not report any $r_A \in [0, f_{T,0}]$ after receiving the H signal.

Therefore, if at least one of the two inequalities in the two pairs of inequalities is violated, then at any separating PBE, Alice does not report any $r_A \in [0, 1]$ after receiving the H signal. This contradicts our assumption that a separating PBE exists for our game. \square