

Eliciting and Aggregating Truthful and Noisy Information

A dissertation presented

by

Xi Gao

to

The School of Engineering and Applied Sciences

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Computer Science

Harvard University

Cambridge, Massachusetts

November 2014

©2014 - Xi Gao

All rights reserved.

Eliciting and Aggregating Truthful and Noisy Information

Abstract

In the modern world, making informed decisions requires obtaining and aggregating relevant information about events of interest. For many political, business, and entertainment events, the information of interest only exists as opinions, beliefs, and judgments of dispersed individuals, and we can only get a complete picture by putting the separate pieces of information together. Thus, an important first step towards decision making is motivating the individuals to reveal their private information and coalescing the separate pieces of information together.

In this dissertation, I study three information elicitation and aggregation methods, *prediction markets*, *peer prediction mechanisms*, and *adaptive polling*, using both theoretical and applied approaches. These methods mainly differ by their assumptions on the participants' behavior, namely whether the participants possess noisy or perfect information and whether they strategically decide on what information to reveal. The first two methods, prediction markets and peer prediction mechanisms, assume that the participants are strategic and have perfect information. Their primary goal is to use carefully designed monetary rewards to incentivize the participants to truthfully reveal their private information. As a result, my studies of these methods focus on understanding to what extent are these methods incentive compatible in theory and in practice. The last method, adaptive polling, assumes that the participants are not strategic and have noisy information. In this case, our goal is to accurately and efficiently estimate the latent ground truth given the noisy information, and we aim to evaluate whether this goal can be achieved by using this method experimentally.

I make four main contributions in this dissertation. First, I theoretically analyze how the participants' knowledge of one another's private information affects their strategic behavior when trading in a prediction market with a finite number of participants. Each participant may trade multiple times in the market, and hence may have an incentive to withhold or misreport his information in order to mislead other participants and capitalize on their mistakes. When the participants' private information is unconditionally independent, we show that the participants reveal their information as late as possible at any equilibrium, which is arguably the worse outcome for the purpose of information aggregation. We also provide insights on the equilibria of such prediction markets when the participants' private information is both conditionally and unconditionally dependent given the outcome of the event.

Second, I theoretically analyze the participants' strategic behavior in a prediction market when a participant has outside incentives to manipulate the market probability. The presence of such outside incentives would seem to damage the information aggregation in the market. Surprisingly, when the existence of such incentives is certain and common knowledge, we show that there exist separating equilibria where all the participants' private information is revealed and fully aggregated into the market probability. Although there also exist pooling equilibria with information loss, we prove that certain separating equilibria are more desirable than many pooling equilibria because the separating equilibria satisfy domination based belief refinements, maximize the social welfare of the setting, or maximize either participant's total expected payoff. When the existence of the outside incentives is uncertain, trust cannot be established and the separating equilibria no longer exist.

Third, I experimentally investigate participants' behavior towards the peer prediction mechanisms, which were proposed to elicit information without observable ground truth. While peer prediction mechanisms promise to elicit truthful information by rewarding participants with carefully constructed payments, they also admit uninformative equilibria where coordinating participants provide no useful information. We conduct the first controlled

online experiment of the Jurca and Faltings [2009] peer prediction mechanism, engaging the participants in a multiplayer, real-time and repeated game. Using a hidden Markov model to capture players’ strategies from their actions, our results show that participants successfully coordinate on uninformative equilibria and the truthful equilibrium is not focal, even when some uninformative equilibria do not exist or result in lower payoffs. In contrast, most players are consistently truthful in the absence of peer prediction, suggesting that these mechanisms may be harmful when truthful reporting has similar cost to strategic behavior.

Finally, I design and experimentally evaluate an adaptive polling method for aggregating small pieces of imprecise information together to produce an accurate estimate of a latent ground truth. In designing this method, we make two main contributions: (1) Our method aggregates the participants’ noisy information by using a theoretical model to account for the noise in the participants’ contributed information. (2) Our method uses an active learning inspired approach to adaptively choose the query for each participant. We apply this method to the problem of ranking a set of alternatives, each of which is characterized by a latent strength parameter. At each step, adaptive polling collects the result of a pairwise comparison, estimates the strength parameters from the pairwise comparison data, and adaptively chooses the next pairwise comparison question to maximize expected information gain. Our MTurk experiment shows that our adaptive polling method can effectively incorporate noisy information and improve the estimate accuracy over time. Compared to a baseline method, which chooses a random pairwise comparison question at each step, our adaptive method can generate more accurate estimates with less cost.

Contents

Abstract	iii
Acknowledgments	xiv
1 Introduction	1
1.1 High Level Introduction to Three Methods	4
1.1.1 Prediction Markets	4
1.1.2 Peer Prediction Mechanisms	5
1.1.3 Adaptive polling	6
1.2 Connecting the Three Methods	7
1.3 My Contributions	10
1.3.1 Theoretical Studies of Prediction Markets	10
1.3.2 Experimental Evaluation of Peer Prediction Mechanisms	11
1.3.3 Designing and Evaluating Adaptive Polling	12
1.4 Dissertation Organization	12
2 Preliminaries	13
2.1 Proper Scoring Rules	13
2.2 Prediction Markets	14
2.3 Market Scoring Rules	16
3 Finite-Stage Prediction Markets	18
3.1 Our Results	19
3.2 Related Work	20
3.3 The Market Game	22
3.3.1 The Finite-Stage Market Game	22
3.3.2 Information Structure	23
3.3.3 Solution Concept and Players' Strategies	26
3.4 The 3-Stage Market Game with Any Information Structure	27
3.4.1 Describing PBE of the 3-Stage Market Game	28
3.4.2 Systematically Identify Candidate PBE Strategies	30

3.4.3	The Consistency Condition	33
3.5	PBE of the Finite-Stage I Game	34
3.5.1	Delaying PBE of 3-stage I Game	35
3.5.2	A Family of PBE for the Finite-Stage I Game	37
3.5.3	Discussion	41
3.6	The 3-Stage D Game	42
3.6.1	An Expression for Alice's Ex-Interim Expected Payoff	43
3.6.2	Three Candidate PBE Strategies for Alice	44
3.6.3	A Sufficient Condition for the Truthful PBE	45
3.7	Conclusion and Future Work	46
4	Prediction Markets with Outside Incentives	48
4.1	Our Results	50
4.2	Related Work	52
4.3	The Market Game	56
4.3.1	The 2-Stage Market Game	56
4.3.2	Solution Concept and Players' Strategies	58
4.4	Known Outside Incentive	59
4.4.1	Truthful vs. Separating PBE	60
4.4.2	A Deeper Look into Alice's Strategy Space	63
4.4.3	A Necessary and Sufficient Condition for Pure Strategy Separating PBE	65
4.4.4	Pure Strategy Separating PBE	71
4.4.5	Pooling PBE	72
4.5	Identifying Desirable PBE	75
4.5.1	Domination-based Belief Refinement	76
4.5.2	Social Welfare	79
4.5.3	Alice's Total Expected Payoff	80
4.5.4	Bob's Expected Payoff	83
4.6	Extensions	83
4.6.1	Other Market Scoring Rules	84
4.6.2	Uncertain Outside Incentive	85
4.7	Connection to Spence's Job Market Signaling Game	86
4.8	Conclusion and Future Directions	90
5	An Experimental Evaluation of a Peer Prediction Mechanism	91
5.1	Related Work	93
5.2	The Jurca and Faltings Mechanism	96
5.3	Experiment Design and Setup	98

5.4	Treatments	103
5.5	Results	106
5.5.1	Summary of Data	106
5.5.2	Learning a Hidden Markov Model	107
5.5.3	Classifying Convergence to Pure Strategy Equilibria	115
5.5.4	Non Peer-Prediction Treatment	117
5.6	Experimental Challenges	117
5.7	Conclusion and Future Work	120
6	Adaptive Polling for Information Aggregation	122
6.1	Related Work	123
6.2	Our Adaptive Method	126
6.2.1	Noisy Information Model	127
6.2.2	Maximum Likelihood Estimation	128
6.2.3	Adaptive Approach	130
6.3	Experiment Design	131
6.4	Results	133
6.5	Conclusion and Future Directions	138
7	Conclusion and Future Directions	140
	Bibliography	145
	Appendix A Appendix to Chapter 3	153
A.1	Omitted Proofs	153
A.1.1	Proof of Theorem 2	153
A.1.2	Proof of Theorem 3	157
A.1.3	Proof of Lemma 2	161
A.1.4	Proof of Lemma 3	163
A.1.5	Proof of Theorem 4	164
A.1.6	Proof of Theorem 5	166
A.1.7	Proof of Theorem 6	170
A.2	Omitted Derivations	172
A.2.1	Derivation for the expression of $u_{a_i}(r)$	172
	Appendix B Appendix to Chapter 4	174
B.1	Proof of Proposition 1	174
B.2	Example 2	175
B.3	Proof of Proposition 3	176

B.4	Example 3	177
B.5	Proof of Theorem 9	179
B.6	Proof of Theorem 10	182
B.7	Proof of Proposition 5	182
B.8	Proof of Proposition 7	184
B.9	Proof of Theorem 16	185
B.10	Proof of Theorem 17	186
Appendix C Appendix to Chapter 5		190
C.1	Estimated HMMs	190
Appendix D Appendix to Chapter 6		192

List of Tables

1.1	Comparison and contrast between the three methods	7
3.1	An example prior distribution. Each cell gives the value of $P(\omega, a_i, b_j)$ for the realized outcome ω , Alice's signal a_i and Bob's signal b_j	46
4.1	Comparison between our setting and Spence's job market signaling game . .	88
5.1	Payment rule examples. In (a), each cell gives a player's payoff if he reports r and his reference report is r_f . In (b) and (c), each cell gives a player's payoff if he reports r and m out of the n_f reference reports are s_1).	97
5.2	Treatments	103
5.3	Payment rules of treatments 1 and 2. The cell at (r, r_f) gives a player's payment if the player reports r and his reference report is r_f	104
5.4	Payment rules of treatments 3 and 4. The cell at (r, f_f) gives a player's payment if the player reports r and n_f of his reference reports are MM	105
5.5	Comparison of actual payoff with expected payoff at truthful equilibrium . .	107
5.6	Each tuple gives the estimated strategy $(\mu_j(MM, MM), \mu_j(GB, MM))$. All numbers are rounded to 2 decimal places.	110
5.7	Classification of convergence to pure strategy equilibria using the simple method. Each cell gives the number of games converging to the particular equilibrium in the specified treatment. The symbol “-” means that the equilibrium does not exist for the payment rule tested in the specified treatment.	117
B.1	An example prior distribution. Each cell gives the value of $P(\Omega, S_A, S_B)$ for the corresponding realizations of Ω , S_A , and S_B	175
B.2	An example prior distribution with $\epsilon \in (0, 0.25)$. Each cell gives the value of $P(\Omega, S_A, S_B)$ for the corresponding realizations of Ω , S_A , and S_B	177
C.1	Treatment 1 Estimated HMM ($K = 4$). $P(MM MM)$ is the probability of reporting MM given a MM signal. $P(MM GB)$ is the probability of reporting MM given a GB signal. P_i is the initial probability of state i . The cell at row i and column j gives the transition probability from state i to state j	190

C.2	Treatment 2 Estimated HMM ($K = 4$). $P(\text{MM} \mid \text{MM})$ is the probability of reporting MM given a MM signal. $P(\text{MM} \mid \text{GB})$ is the probability of reporting MM given a GB signal. P_i is the initial probability of state i . The cell at row i and column j gives the transition probability from state i to state j	190
C.3	Treatment 3 Estimated HMM ($K = 4$). $P(\text{MM} \mid \text{MM})$ is the probability of reporting MM given a MM signal. $P(\text{MM} \mid \text{GB})$ is the probability of reporting MM given a GB signal. P_i is the initial probability of state i . The cell at row i and column j gives the transition probability from state i to state j	191
C.4	Treatment 4 Estimated HMM ($K = 4$). $P(\text{MM} \mid \text{MM})$ is the probability of reporting MM given a MM signal. $P(\text{MM} \mid \text{GB})$ is the probability of reporting MM given a GB signal. P_i is the initial probability of state i . The cell at row i and column j gives the transition probability from state i to state j	191
C.5	Non-Peer Prediction Treatment Estimated HMM ($K = 4$). $P(\text{MM} \mid \text{MM})$ is the probability of reporting MM given a MM signal. $P(\text{MM} \mid \text{GB})$ is the probability of reporting MM given a GB signal. P_i is the initial probability of state i . The cell at row i and column j gives the transition probability from state i to state j	191

List of Figures

3.1	A PBE of a Finite-Stage I Game with 3 players	40
4.1	An illustration of Y_H and Y_T by partitioning Alice’s strategy space. The blue regions contain Alice’s reports that are dominated by truthful reports. The white regions contain Alice’s reports that are not dominated by truthful reports. Y_H and Y_T are the upper bound values for the white regions.	64
5.1	The Game Interface	102
5.2	Percentage of players with the specified signal and report	107
5.3	The graphical model for each player i implied by the HMM.	109
5.4	Treatment 1 Results.	112
5.5	Treatment 2 Results.	113
5.6	Treatment 3 Results.	114
5.7	Treatment 4 Results.	115
5.8	Non Peer-Prediction Treatment Results. Each row shows how a single player’s strategy evolves over multiple rounds.	118
6.1	Two example pictures. The left picture has 342 dots, and the right one has 447 dots.	131
6.2	Frequency of the left picture being selected in the 1200 pairwise comparisons of all 12 trials. The x-axis represents the difference in number of dots between the left and right pictures (left – right). The observations are grouped into 7 buckets according to the difference in dots. Each bar represents the empirical frequency for the corresponding bucket. The curve is $\Phi(0.017x)$	134
6.3	The dynamics of the estimated strength parameters for an adaptive polling trial. The x-axis is the number of iterations. The y-axis is the value of the estimated strength parameters. The rightmost part of the figure labels the value of the “gold standard” strength parameter for each picture.	135
6.4	The entropy of the estimated distribution $\mathcal{N}(\hat{s}, \hat{\Sigma})$	136
6.5	The log score — the logarithm of the pdf of $\mathcal{N}(\hat{s}, \hat{\Sigma})$ evaluated at the “gold standard” strength parameters	136

6.6	The fraction of the pairwise comparison questions that are correctly answered	137
D.1	The dynamics of the estimated strength parameters for adaptive polling trials	192
D.2	The dynamics of the estimated strength parameters for random polling trials	193

Acknowledgments

First and foremost, I would like to express my immense gratitude to my advisor Yiling Chen. I came to Harvard knowing almost nothing about research. Throughout the past six years, Yiling has given me endless support and guidance, and she has patiently witnessed my transformation into a researcher. Yiling cared about her students deeply, and I can feel her pouring her heart out all the time to teach us all that she knows about being a researcher and to guide us towards success. She taught me how to write a captivating introduction to a research paper, how to deliver an engaging presentation, and how to be unafraid of asking questions during a talk. She was there at every step of the way to support me and to cheer me on my success, and I am extremely grateful to have her as my PhD advisor.

I am grateful to many people who helped me through or witnessed the many defining moments of my PhD years. As an undergraduate, Kevin Leyton-Brown gave me my first opportunity to experience what research was all about. Avi Pfeffer gave me the opportunity to come to Harvard and to realize my PhD dreams. As a new graduate student, I became fascinated with prediction markets while sitting in my first graduate CS class taught by Yiling, and prediction markets later become a central topic in my dissertation. The help and encouragement from my qualifying committee members including Michael Mitzenmacher and David Parkes motivated me to improve my presentation and public speaking skills. David Parkes attended almost all of my seminar or practice talks and always raised insightful questions and comments. During the third year of my PhD, I had the pleasure of working with Rick Goldstein, an extremely talented Harvard undergraduate, and Ian Kash, a postdoc who patiently helped me through my paper writing struggles. With Rick, Ian and Yiling, we wrote the first research paper that I am proud of. Yoram Bachrach gave me my first opportunity to be a research intern at Microsoft Research. This internship gave me a tremendous push towards becoming a critical thinker and an independent researcher. From my officemate and good friend Andrew Mao, I learned how to conduct Mechanical Turk experiments. My collaboration with Andrew will later turn into (and hopefully continue to be)

a long term partnership as we design, implement and conduct elaborate online experiments with simultaneous participation of many human subjects. I enjoyed working with Jie Zhang and Yiling towards the goal of fully characterizing the participants' equilibrium behavior in market scoring rule prediction markets. Last summer, I was fortunate to be an intern at Microsoft Research NYC where I worked on exciting projects with David Pennock, Jenn Wortman Vaughan and Miro Dudik.

Many people at Harvard has made the past six years a memorable experience for me. Anna Huang, I will always remember meeting you for the first time during Zhenming's office hour, going to concerts with you, learning French with you, taking swing dancing classes with you, chatting with you all the time, (the list goes on and on...). It was so much fun to learn swing dancing with you. Thanks to your constant encouragement, I've finally achieved a reasonable skill level for at least one dance form (among so many that I've tried before). Andrew Mao, we hardly agreed on anything, but we wrote our trick or treat paper (that I am extremely proud of), we shared the crazily nice condo on the 53rd floor in Manhattan, and we beat the 2-person Hanabi numerous times! Jens Witkowski, thank you for being a good friend and for visiting Harvard and never leaving. It is always so easy to talk with you and you seem to always understand me on many levels. Mike Ruberry, you brought the idea of board games to a new level. Thank you for buying so many wonderful games and for always patiently explaining the games to us in such an entertaining way. Pao Siangliulue, I will miss your constant punching and your enthusiasm about aikido. Thank you for never having an opinion so I can always choose what to eat when we go to Porter square. My fellow EconCSers at Harvard, Shaili, Malvika, John, Mike, Andrew, Victor, Bo, Ming, Greg, Jens, Thibaut, it was great to know all of you and to support one another through our PhD journeys. My fellow CS Theory students, Varun, Colin, Zhenming, Thomas, Mark, Scott, Anudhyan, you are the smartest, most creative and talented people that I've ever met. Thank you for making our lives full of fun and laughter! Varun, thanks for giving me the memory that once upon a time I was able to speak French to real people.

Last but not least, I want to thank my family members, my parents and my husband Jingnan. Dear Jingnan, thank you for bringing so much sunshine, adventure, and laughter into my life, for always believing in me, supporting me, and pushing me to become a better person! Our best days are yet to come. My dear parents, you were one of the original motivations for me to pursue graduate school. Thank you for believing that I can do anything and for always loving and supporting me no matter what!

Chapter 1

Introduction

To make informed decisions in the modern world, it is crucial to obtain and aggregate relevant information about some uncertain events of interest. For instance, depending on the forecasted release dates of their products, companies such as Google and HP may want to adjust their research and development strategies. Businesses such as Yelp and Angie's list care about soliciting accurate and honest reviews of restaurants, service companies and health care professionals. Individuals speculate about and bet on the rankings of sports teams in competitions and the rankings of horses in horse races. In such domains, the information of interest only exists as opinions, beliefs, and judgements of dispersed individuals. Each individual only possesses a small piece of the whole information puzzle, and we can only get a complete picture by putting the separate pieces of information together. Therefore, motivating the individuals to reveal their private information and coalescing the separate pieces of information together is an important first step towards decision making.

Information Elicitation: There are three main challenges associated with the task of eliciting useful information from individuals. First, for individuals who have useful information, it is costly for them to reveal it, yet this action only benefits others. Thus, it is imperative to reward the participants for revealing their information. Many websites offer explicit rewards for the participants' contributions. These rewards may or may not be mon-

etary. Iowa Electronic Markets, a famous example of prediction markets, provide monetary rewards for participants who have accurate information for forecasting future events. On Amazon’s Mechanical Turk, requesters offer monetary rewards to the workers for completing many different kinds of tasks, such as labeling images, providing opinions, transcribing audios. On platforms such as StackOverflow, Yahoo Answers, Reddit, and Yelp, participants can earn non-monetary rewards such as badges and ratings for their contributions. These rewards make it rational for participants to contribute their opinions and information rather than not.

Moreover, although explicit rewards make participating and contributing rational, they do not necessarily incentivize the individuals to contribute their information truthfully, especially when the individuals strategically decide on what to contribute. In many scenarios, instead of truthfully revealing his information, a participant may be tempted to withhold information or reveal false information. For example, a prediction market participant may reveal false information inside the market in order to gain external rewards. When participating on Yelp or Amazon, a participant may want to be nice and not leave negative reviews, or he may be paid to leave a positive or negative review for a certain company. On crowdsourcing markets such as Amazon’s Mechanical Turk, a participant is typically offered a constant monetary reward for completing a task. In this case, the participant may choose to make a contribution requiring the least effort rather than contributing his true opinion or information. Thus, the second challenge is to motivate the participants to truthfully reveal his information rather than strategically withhold or misreport their information.

Finally, for incentivizing truthful contributions, the designer assumes that the participants already have the desired information. In reality, however, participants need to invest costly effort to obtain the desired information, and the quality of the participants’ information depends on the amount of effort invested. Thus, the third challenge is to motivate the participants to invest costly effort in order to acquire accurate and high quality information.

In this dissertation, I study several information elicitation methods, which offer explicit

rewards for the participants’ contributions. I focus on studying how the reward schemes of these methods address the first two challenges of incentivizing the participants to participate and to contribute truthful information. I do not study whether they can address the third challenge of motivating the participants to invest costly effort to acquire information. More specifically, I characterize the participants’ behavior towards these methods in theory and in practice and analyze the effect of these behavior on the quality of the elicited information.

Information Aggregation: Once the participants’ reports are collected, we may want to aggregate these reports together to produce a single quantity for decision making, e.g. a probability estimate of an uncertain event. Given a set of reports, some common methods for aggregating the reports include statistical methods, voting, and probabilistic inference. The aggregation and elicitation processes can be separate, but they may also be combined into a single mechanism. For example, in a prediction market, each participant is motivated to infer information from other participants’ reports and incorporate these inferred information into his own report, effectively performing the aggregation for the mechanism.

For a given set of reports, different aggregation methods may produce estimates of different accuracies. Choosing the best aggregation method critically depends on how the participants’ information relates to the ground truth we are estimating. With little knowledge of how participants formed their reports, simple methods such as averaging or majority voting can be suitable for a wide range of settings. However, if we can better capture the participants’ private information using theoretical models, then we may be able to develop more sophisticated aggregation methods to produce more accurate estimates. In this dissertation, I propose an adaptive polling method, which uses a particular statistical model to capture the noise in the participants’ reports, and demonstrate that this method can effectively aggregate the participants’ reports to produce an accurate estimate of the latent ground truth.

1.1 High Level Introduction to Three Methods

In this dissertation, I study three methods: *prediction markets*, *peer prediction mechanisms*, and *adaptive polling*. In this section, I give a high level introduction to these methods and describe the research questions that I address for these methods.

1.1.1 Prediction Markets

An important task for decision making is to generate accurate forecasts of uncertain events. Some examples of such events are weather, elections, product sales, and Oscar winner. Prediction markets are designed to elicit probabilistic estimates from participants and aggregate these probabilities together to forecast uncertain events, whose realized outcomes will be observed in the future. Participants of a prediction market can reveal their private information through trading contracts, and the market rewards the participants by evaluating their probabilistic estimates against the observable ground truth. At any time, the current market price/probability can be interpreted as a consensus probabilistic forecast for the future event given all the information that has been revealed to the market so far.

Substantial empirical research has demonstrated that prediction markets produce remarkably accurate forecasts in practice [Berg et al., 2001, Camerer, 1998, Chen and Plott, 2002, Debnath et al., 2003, Forsythe et al., 1992, 1999, Wolfers and Zitzewitz, 2004]. Yet, existing theory on prediction markets fails to explain its empirical success, especially how and why information gets aggregated in the market. Ideally, prediction markets should be designed to be incentive compatible such that every participant reveals his private information truthfully at his first opportunity to trade in the market. Unfortunately, prediction market mechanisms are not incentive compatible in general. Thus, whether the market can elicit and aggregate accurate information depends on how the strategic participants behave in response to the monetary rewards provided by the mechanism. Specifically, when a participant has multiple opportunities to trade in the market, or when he receives payoffs from

outside of the market contingent on his trade in the market, the participant may have incentives to withhold or misreport his private information when trading the market in order to mislead other participants and maximize his total expected payoff. If such manipulations occur, it would be questionable whether all the participants' information will be revealed and incorporated into the market probability.

In this dissertation, I focus on studying market scoring rules proposed by Hanson [2007a], which is one of the most popular automated market maker mechanisms for prediction markets. Market scoring rules achieves a weaker property than incentive compatibility: if all participants are rational and self-interested economic agents, then any myopic participant is incentivized to truthfully reveal his private information when trading in a market scoring rule prediction market. A myopic participant is not forward looking and decides on what to report for the current trade based on his expected reward from the current trade only.

My main goal in studying market scoring rules is to understand the behavior of non-myopic participants and the effects of these participants' behavior on the accuracy of the collected and aggregated information. Specifically I tackle the following research questions:

- In theory, how do non-myopic participants behave at game-theoretic equilibria of market scoring rules?
- How does the equilibrium behavior of non-myopic participants affect the accuracy of the elicited and aggregated probabilistic estimates for uncertain events?

1.1.2 Peer Prediction Mechanisms

For many events of interest, the outcome of the events are neither observable nor verifiable, so we cannot evaluate the participants' reports against an observable ground truth. For example, the outcome of the event may be subjective (e.g. the quality of a book), not publicly observed (e.g. the breakdown frequency of a product, which is only known by the manufacture), or not verifiable (e.g. the extinction of the human race). Peer prediction mechanisms

are designed to elicit probabilistic estimates from participants about such events without observable ground truth. Peer prediction mechanisms use carefully designed monetary rewards to induce the participants to reveal their private information truthfully. Assuming that the participants are rational and self-interested economic agents, peer prediction mechanisms have the following desirable theoretical property: a participant maximizes his expected reward by reporting truthfully if he believes that all other participants are also truthful. In other words, peer prediction mechanisms induce truth telling in equilibrium.

Despite the desirable property of peer prediction mechanisms, there is little theoretical guarantee that the participants will adhere to truth telling in practice. First of all, many peer prediction mechanisms have other uninformative equilibria where no useful information is revealed, and the theory has no prediction on which equilibrium will be played in practice. Moreover, there is a lack of understanding of the practical performance of peer prediction mechanisms since most of them have not been evaluated in a practical setting.

My main goal in studying peer prediction mechanisms is to understand the behavior of the participants towards peer prediction mechanisms in practice. Specifically I tackle the following research questions:

- In practice, how do the participants behave towards peer prediction mechanisms?
- How do the participants' behavior towards peer prediction mechanisms affect the accuracy of the elicited probabilistic estimates for uncertain events?

1.1.3 Adaptive polling

Adaptive polling aims to aggregate small pieces of imprecise information together in order to produce an accurate estimate of a latent ground truth. In designing this method, we make very different assumptions about the participants compared to the assumptions of prediction market and peer prediction mechanisms. First, we assume the participants to be non-strategic. Thus, eliciting truthful information is not our concern because the participants

will truthfully reveal their information for a constant monetary reward. Moreover, we assume that each participant provides a small and partial piece of noisy information of the latent ground truth. One challenge is to understand and account for the noise in the participants' information when aggregating their contributions. In addition, we assume to have the ability to actively query the participants for specific pieces of information. So we aim to find a way to adaptively query information from the participants to improve the accuracy of the estimate as quickly as possible. In summary, we explore the following research questions:

- How to use a theoretical model to best capture the participants' noisy observations for a particular latent ground truth?
- How to efficiently aggregate small pieces of noisy information contributed by the participants in order to produce an accurate estimate of the ground truth?
- How to adaptively query information from the participants to improve the accuracy of the estimate produced as quickly as possible?

1.2 Connecting the Three Methods

In this section, I compare and contrast the three methods in terms of which part of the problem they are targeting, whether they require observable ground truth, and what assumptions they make about the participants' behavior. By comparing and contrasting them from several perspectives, I would like to relate the three methods and provide a bigger picture for this dissertation. The following discussion is summarized in Table 1.1.

Table 1.1: Comparison and contrast between the three methods

Prediction Markets	Peer Prediction	Adaptive Polling
Elicitation and Aggregation	Elicitation	Aggregation
Observable Ground Truth	Latent Ground Truth	
Strategic Participants		Non-Strategic Participants
Perfect Information		Noisy Information

Targeted Problem First, all three methods were proposed to solve the same high level problem, but each of them was designed to target different parts of the whole problem.

Prediction markets were designed to tackle the whole information elicitation and aggregation problem. A prediction market provides carefully designed monetary rewards to incentivize the participants to truthfully reveal their private information. Moreover, the design of the market encourages each participant to improve upon the existing market estimate by inferring information from the historical forecasts and incorporating the inferred information into their own estimate. In essence, the participants are performing the aggregation for the mechanism.

In contrast, peer prediction mechanisms and adaptive polling focus on solving part of the whole problem. Peer prediction mechanisms aim to elicit truthful information from the participants, but they do not specify how to use the elicited information. Although adaptive polling has distinct elicitation and aggregation processes, it assumes away the challenge of eliciting truthful information by considering the participants to be non-strategic. Instead, adaptive polling focuses on adaptively query information from the participants and efficiently aggregate the elicited information to produce an accurate estimate of the latent ground truth.

Type of Ground Truth These three methods also differ by the type of event they are able to estimate or forecast. Prediction markets are designed to forecast events with observable ground truth, since their rewards are determined by evaluating the participants’ probabilistic estimates against the realized outcome of the event. Because an observable ground truth is available, prediction markets can achieve relatively stronger theoretical guarantee – they are incentive compatible for myopic participants.

In contrast, peer prediction mechanisms and adaptive polling are more powerful methods since they can estimate or predict events without observable ground truth. Due to the lack of observable ground truth, peer prediction mechanisms can only achieve a weak theoretical guarantee — they can induce truth telling in equilibrium only. Adaptive polling uses a theoretical model to accurately capture the noise in the participants’ information. Therefore,

it can make good use of the reported noisy information to accurately estimate the latent ground truth.

Assumptions of Participants’ Behavior One of the most important distinctions among these methods is their assumptions of the participants’ behavior. These assumptions have critical influences on the focus of these methods. For different assumptions, the designers face different challenges when attempting to solve the overall problem, and thus the resulting methods are derived with particular desirable properties which target specific aspects of the whole problem. In particular, the assumptions of these three methods capture two distinct aspects of the participants’ behavior: whether the participants make perfect or imperfect observations of the desired information and whether they are strategic or not.

Prediction markets and peer prediction mechanisms assume that the participants have “perfect information” and they are “strategic”. The participants make perfect observations of the desired information and their information are not noisy. Moreover, the participants are rational and self-interested economic agents and that they choose their actions in order to maximize their expected rewards from the mechanism. Given these assumptions, the designer aims to design the mechanism to be incentive compatible such that the participants are best off truthfully revealing their information. However, achieving complete incentive compatibility is quite challenging. For the mechanisms studied in this dissertation, there are several settings where a participant may be able to improve his expected rewards from inside or outside of the mechanism by withholding or misreporting his information.

In contrast, adaptive polling assumes that the participants possess “noisy information” but they are “not strategic”. The participants make noisy observations of the desired information, but they are always willing to truthfully reveal their information regardless of whether their expected rewards can be improved by behaving otherwise. Given these assumptions, this method focuses on developing the most effective way to aggregate the noisy information contributed by the participant to accurately estimate the latent ground truth.

The assumptions of all three methods are simplifying the participants’ behavior in one

way or another. In practice, it is reasonable to assume that the participants have “noisy information” and they are “strategic”. However, prediction markets and peer prediction mechanisms ignore the “noisy information” aspect whereas adaptive polling assumes away the “strategic” aspect. In the conclusion, I will discuss the implication of our results on these assumptions as well as some new models of participants’ behavior in recent work.

1.3 My Contributions

In this section, I outline the four main contributions of my dissertation.

1.3.1 Theoretical Studies of Prediction Markets

We theoretically analyze the participants’ behavior at game theoretic equilibria of market scoring rule prediction markets in following two settings.

In the first setting, there are a finite number of participants in the market and each participant can only trade for a finite number of times in the market. Since a participant can trade multiple times in the market, he may have incentives to misreport or withhold his information in order to mislead the other market participants and capitalize on their mistakes later on. We characterize equilibria of such prediction markets depending on how the participants’ private information relates to the event being forecasted. When the participants’ information is unconditionally independent, there exists a unique family of equilibria, where, qualitatively speaking, every participant reveals his private information as late as possible. This is arguably the worst outcome for the purpose of information aggregation. We also provide insights for the equilibria of the markets when the participants’ private information is both conditionally and unconditionally dependent on the realized outcome of the event.

In the second setting, there are two participants in the market and each trades in the market once. The final probabilistic forecast generated by the market is used to make a decision, and the first market participant receives an additional payoff from outside of

the market contingent on the decision. If the first market participant’s outside payoff is attractive, he may have incentives to misreport or withhold his information when trading inside the market in order to improve his additional payoff from outside of the market. For this setting, it would seem that information loss within the market is inevitable because of the first participant’s incentive to manipulate the market probability. However, we show that there exists a “separating” equilibria where all the participants’ private information is revealed and incorporated into the final market probability, which is the goal of running the market. We characterize sufficient and necessary conditions for such separating equilibria to exist. We also characterize “pooling” equilibria where information loss occurs in the market and show that the separating equilibria are more desirable than many other equilibria because the separating equilibria satisfy domination based belief refinements, maximize the social welfare of the setting, or maximize either participant’s total expected payoff.

1.3.2 Experimental Evaluation of Peer Prediction Mechanisms

We experimentally evaluate the performance of the Jurca and Faltings [2009] (JF) peer prediction mechanism through a controlled online experiment. Our experiment allows the participants to learn and adapt to the mechanism through a multiplayer, real-time and repeated game. Using our experimental data, we analyze the participants’ behavior in terms of convergence to game-theoretic equilibria of the JF mechanism.

In our setting, we observe that participants clearly favor the uninformative equilibria over the truthful equilibrium when paid by the JF mechanism. Eliminating some uninformative equilibria or making them less desirable successfully deterred the participants from choosing them, but did not motivate the participants to be truthful. In contrast, the majority of the participants are consistently truthful in the absence of any peer prediction mechanism. Methodology wise, our work demonstrates the promise of evaluating game theoretic mechanisms through online behavioral experiments, and using probabilistic models for analyzing experimental data to explain the participants’ behavior.

1.3.3 Designing and Evaluating Adaptive Polling

We design and experimentally evaluate an adaptive polling method for aggregating small pieces of imprecise information to produce an accurate estimate of a latent ground truth. In designing this method, our main contribution is to combine an aggregation process, which produces accurate estimates by accounting for the noise in the participants' information, with an elicitation process, which adaptively queries information from the participants in order to quickly improve the estimate accuracy.

We apply adaptive polling to the problem of ranking a set of alternatives, each of which is characterized by a latent strength parameter. Our goal is to produce accurate estimates of the strength parameters in order to correctly rank the alternatives. At each step, adaptive polling collects the result of a pairwise comparison, estimates the strength parameters from the pairwise comparison data, and adaptively chooses the next pairwise comparison question to myopically maximize expected information gain.

We evaluate our method through an experiment on Amazon Mechanical Turk. Our experimental results show that the adaptive method can effectively incorporate noisy information and improve the estimate accuracy over time. Also, adaptive polling is superior to a naive method of presenting a random pair of alternatives for each participant.

1.4 Dissertation Organization

Chapter 2 introduce preliminary technical concepts to prepare for the subsequent sections. Chapter 3 presents a theoretical analysis of participants' equilibrium behavior in prediction markets with a finite number of stages. Chapter 4 describes a theoretical analysis of prediction markets where a market participant has incentives to manipulate the market probability from outside of the market. Chapter 5 is an experimental evaluation of a peer prediction mechanism. Chapter 6 proposes and experimentally evaluates an adaptive polling method for eliciting and aggregating information without economic incentives.

Chapter 2

Preliminaries

2.1 Proper Scoring Rules

The simplest information elicitation mechanism is a scoring rule. A scoring rule provides a carefully designed payment to incentivize a single expert to truthfully report his probabilistic estimates of an event. The amount of the payment depends on the expert's reported probabilistic estimates and the realized outcome of the event [Good, 1952, Winkler, 1969, Savage, 1971, Gneiting and Raftery, 2007].

Formally, consider an event Ω with a set ω of m mutually exclusive and exhaustive outcomes. Let $\mathbf{r} = \{r_1, r_2, \dots, r_m\}$ be the probabilistic estimates reported by the expert. For a given set of probabilistic estimates $\mathbf{r}$, the scoring rule $s(\omega, \mathbf{r})$ assigns a score $s_i(\mathbf{r})$ if the i -th outcome in ω is realized. In fact, any bounded, convex and differentiable function of $\mathbf{r}$ corresponds to a proper scoring rule [Savage, 1971]. For example, a popular proper scoring rule is the logarithmic scoring rule

$$s_i(\mathbf{r}) = b \ln(r_i) \tag{2.1}$$

where b is a positive parameter.

A (strictly) proper scoring rule is incentive compatible for a risk neutral expert — a

risk neutral expert maximizes his expected score by truthfully reporting his probabilistic estimates. For example, consider the logarithmic proper scoring rule with $m = 2$. Assume that the expert believes the probabilities of the two outcomes to be p and $1 - p$, and he reports the probabilities q and $1 - q$. The expert maximizes his expected score by reporting $q = p$, as shown below:

$$\frac{\partial}{\partial q} \{pb \ln(q) + (1 - p)b \ln(1 - q)\} = 0 \quad (2.2)$$

$$\Rightarrow b \left\{ \frac{p}{q} - \frac{1 - p}{1 - q} \right\} = 0 \quad (2.3)$$

$$\Rightarrow b\{p(1 - q) - q(1 - p)\} = 0 \quad (2.4)$$

$$\Rightarrow b\{p - pq + pq - q\} = 0 \quad (2.5)$$

$$\Rightarrow q = p \quad (2.6)$$

In this dissertation, I study prediction market and peer prediction mechanisms, which are both derived using proper scoring rules.

2.2 Prediction Markets

Prediction markets are powerful mechanisms created for the purpose of forecasting future events. Similar to a financial market, a prediction market offers contracts whose payoffs are tied to outcomes of a future event. For example, to forecast whether the flu activity in Massachusetts will be widespread by May 1 this year, the contract could pay \$1 per share if the CDC characterizes the flu activity to be widespread in Massachusetts on May 1 and \$0 otherwise. Participants can express their private information about flu activity in a credible way by trading shares of this contract, and they will be rewarded based on their trades and the realized outcome of the event. If a risk neutral participant believes that the flu activity will be widespread with probability p , then he can make profits in expectation by buying the contract if its current price is lower than p and selling the contract if its current price

is higher than p . Thus, the market price of the contract may be interpreted as a consensus forecast for this event.

Prediction markets have been deployed and shown to be extremely successful in practice. Some notable examples of deployed prediction markets include Iowa Electronic Markets (<http://tippie.uiowa.edu/iem/>), Inkling Markets (<http://inklingmarkets.com/>), Hollywood Stock Exchange (<http://www.hsx.com/>) and Betfair (<http://www.betfair.com/GBR/en/>). Substantial empirical research has demonstrated that prediction markets produce remarkably accurate forecasts in practice and often outperform alternative forecasting methods in a wide range of settings [Berg et al., 2001, Wolfers and Zitzewitz, 2004, Forsythe et al., 1992, 1999, Debnath et al., 2003, Chen and Plott, 2002].

A prediction market has traditionally been run like a financial market, by setting up a double auction where traders place orders to buy or sell shares of contracts and making an auctioneer match the buy and sell orders without incurring any risk. These double auctions work well for financial markets where the number of participants is large, making it easy to match buy and sell orders. However, double auctions may suffer from the thin market problem, which prevents the participants from revealing their information through trading. Moreover, double auctions are zero-sum games for the participants, and the no-trade theorem shows that rational risk-neutral traders should not participate in such a zero-sum game. These problems motivate the use of automated market maker mechanisms.

For an automated market maker mechanism, each participant interacts with the market maker for buying or selling shares of the contracts. The market maker is always willing to accept any buy or sell orders at the price quoted by the market maker. Such an automated market maker essentially subsidizes the market for the purpose of eliciting useful information. Hanson’s market scoring rules (MSR) is the de facto automated market mechanism [Hanson, 2007a]. We formally introduce market scoring rules in the next section. MSR has the desirable property that the maximum amount of money that the market maker may lose is bounded in the worst case.

2.3 Market Scoring Rules

Market scoring rules (MSR), proposed by Hanson [2003, 2007b], are derived from proper scoring rules. Consider a proper scoring rule $s(\omega, r)$ for a binary event. The corresponding MSR market is a sequential shared version of the proper scoring rule. Suppose that we would like to forecast a binary event Ω with its realized outcome denoted $\omega \in \{0, 1\}$. The MSR market starts with an initial market probability r^0 for $\omega = 1$. For a binary event, the probability of outcome $\omega = 0$ is implicitly $1 - r^0$. Participants trade in the market in sequence and each participant can change the current probability estimate to a new value of his choice. The market closes at a predefined time. After that, the realized outcome ω is observed and participants receive their payoffs. When a participant changes the market probability for $\omega = 1$ from r^{t-1} to r^t , he is paid the scoring rule difference, $s(\omega, r^t) - s(\omega, r^{t-1})$, depending on the realized outcome ω . A participant may trade in the market multiple times. If T_i denotes the set of stages where participant i trades, then participant i 's total payoff is the sum of the payoff for each of his reports, $\sum_{t \in T_i} (s(\omega, r^t) - s(\omega, r^{t-1}))$. The logarithmic market scoring rule (LMSR), derived from the logarithmic proper scoring rule, is one of the most popular market scoring rule mechanisms.

A market scoring rule has the nice incentive property that a risk-neutral, myopic participant can maximize his expected payoff by truthfully reporting his probability estimate, because he cannot influence the market probabilities before his report. A participant is myopic if he is not forward looking and trades in each stage as if it is his only chance to trade in the market. If a participant can trade multiple times in the market and wants to maximize his *total* payoff, then he may want to misreport his estimate or withhold his information in order to mislead other participants and capitalize on their mistakes later on. Alternatively, a participant may want to manipulate the market probability in order to increase an payoff he receives from outside of the market.

We describe MSR as a mechanism where participants sequentially revise the probability estimates of event outcomes. However, 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 the market maker updates the contract prices based on the trades. In practice, a MSR market offers one contract for each outcome and the contract pays off \$1 if the corresponding outcome materializes. The prices of all contracts represent a probability distribution over the outcome space. In this work, we analyze MSR as a mechanism for changing probability estimates since abstracting away the contracts makes subsequent analyses more tractable. We refer interested readers to [Chen and Pennock, 2007] and [Abernethy et al., 2013] for more information on the equivalence of the two models.

Chapter 3

Finite-Stage Prediction Markets

A prediction market provides economic incentives for the participants to reveal their private information about some uncertain event of interest. A market participants can express his probability estimates for outcomes of the event by trading shares of financial contracts, and he will be rewarded if his probability estimate is more accurate than the previous market estimate. By observing trading activities of other participants, a rational participant can infer some information from their activities and combine such information with his private information when trading in the market. Prediction markets rely on the economic incentives provided by the mechanism and the belief updating of participants to achieve the primary goal of eliciting and aggregating information about uncertain events of interest.

To this end, arguably we desire that participants reveal their private information *truthfully* and *immediately* in prediction markets. However, how well the information elicitation and aggregation goal is achieved depends on the strategic behavior of the self-interested market participants, and the behavior of the market participants is influenced by their private information and their knowledge of others' private information. We formally refer to the relationship between the participants' private information as the *information structure* of the participants.

In this work, we model a prediction market as an extensive-form Bayesian game where

each participant has a piece of private information, there is a joint distribution of the participants’ private information and the event outcome, and this joint distribution is common knowledge to all participants. This joint distribution, which is the information structure of the market game, captures what participants know about one another’s private information. The goal of this work is to understand how and how quickly information is aggregated in the market by characterizing game-theoretic equilibria of the market game for different information structures.

We study Hanson’s logarithmic market scoring rule (LMSR) [Hanson, 2007b], which is the de facto automated market maker mechanism for prediction markets. Because participants interact with the market maker, which is the mechanism per se and behaves deterministically, we only need to model the participants side of the market. This makes the generally challenging equilibrium analysis for extensive-form Bayesian games tractable for some information structures in our setting.

Prior work [Chen et al., 2007, 2010b] has shown that when participants’ information is conditionally independent given the true outcome of the event, there exists a unique family of equilibria where every participant races to truthfully reveal all their information as soon as possible. This is arguably the most desirable outcome for the market’s goal. This work considers the market games with the two remaining classes of information structures:

- The I game: the participants’ private information is unconditionally independent, and
- The D game: the participants’ private information is both conditionally and unconditionally dependent.

3.1 Our Results

We characterize the unique family of equilibria of the I game with a finite number of participants and a finite number of stages. At any equilibrium in this family, if player i ’s last stage of participation in the market is after player j ’s, player i only reveals his information after

player j 's last stage of participation and on or before his own last stage of participation. Qualitatively speaking, at any equilibrium of this game, participants race to delay revealing their information, which is probably the least desirable outcome for the market's goal. These equilibria are in stark contrast to the equilibria of the market game when the participants' information is conditionally independent given the realized outcome of the event.

We also provide insights on equilibria of the D game. It is in general challenging to characterize equilibria of the D game because the information structure does not have any clear mathematical property that we can leverage. We provide a systematic method for identifying possible equilibrium strategies in the market game with any information structure. With this method, we identify all possible PBE strategies for the players in a restricted 3-stage D game with 2 participants Alice and Bob, the sequence of participation Alice Bob and Alice, and 2 realized signals for Alice. Moreover, we show that there exist instances of the D game that admit truthful equilibria. In particular, we give a sufficient condition for a truthful equilibrium to exist in the 3-stage D game and give a prior distribution satisfying this condition.

3.2 Related Work

Our work closely follows the prior work by Chen et al. [2007], Dimitrov and Sami [2007], and Chen et al. [2010b]. We model a prediction market as an extensive-form Bayesian game as in these prior work. Chen et al. [2010b] considered both a finite-stage, finite-player and an infinite-stage, finite-player market game. They showed that when players' information is conditionally independent given the true state of the world, for both the finite- and infinite-stage games, there is a unique type of Perfect Bayesian Equilibria (PBE), where players reveal their information truthfully and as soon as they can. When players' information is (unconditionally) independent, they proved that the truthful play is not an equilibrium for both the finite- and infinite-stage games. An earlier work by Nikolova and Sami [2007] also

presented an instance in which the truthful strategy is not optimal in an extensive-form game based on this market. However, when players have independent information, the existence of a PBE was left as an open question. In this work, we characterize all PBE of the finite-stage game with independent information and explore a special case of the setting when players' information is neither conditionally nor unconditionally independent.

Instead of characterizing equilibria, Ostrovsky [2012] studied whether information is fully aggregated *in the limit* at a PBE of an infinite-stage, finite-player market game with risk-neutral players. He characterized a separability condition under which the market price of a security converges to its expected value conditioned on all information with probability 1 at any PBE. If the security and the partition structure satisfy the separability condition, then the following situation will never occur: every player believes the security to be of one value in a particular state whereas the actual value of the security in this state is a different value. In particular, the separability condition is always satisfied by complete markets, which is the setting studied in this chapter. Thus Ostrovsky's setting does not place any restriction on the information structure of the market. Iyer et al. [2010a] extended the setting to risk-averse players and characterized the condition for full information aggregation in the limit at any PBE. However, whether a PBE exists in such market games remains an open question.

The 3-stage version of our prediction market model resembles the ones studied by Dimitrov and Sami [2010a] and our work in Chapter 4: they both study 2-player games and the first player has another chance of participation after the second player's turn in the game. However, both Dimitrov and Sami [2010a] and our work in Chapter 4 consider that the first player has utility for some event outside of the current market and the price in the current market influences the outcome of this event. In this work, players only derive utilities from their trades in the market.

Jian and Sami [2010] studied market scoring rule prediction markets in a laboratory setting. In their experiment, participants may have conditionally or unconditionally independent information and the trading sequence may or may not be structured (a trading

sequence is structured if it is pre-specified and is common knowledge to all participants). They confirmed previous theoretical predictions of the strategic behavior by Chen et al. [2010b] when the trading sequence is structured. This study suggests that the behavior of participants in a prediction market critically depends on whether they reason about the other participants’ private information.

There are experimental and empirical studies on price manipulation in prediction markets using double auction mechanisms. The results are mixed, some giving evidence for the success of price manipulation [Hansen et al., 2004b] and others showing the robustness of prediction markets to price manipulation [Camerer, 1998, Hanson et al., 2007, Rhode and Strumpf, 2004, 2007]. In the literature on financial markets, participants have been shown to manipulate market prices [Allen and Gale, 1992, Chakraborty and Yilmaz, 2004, Kumar and Seppi, 1992].

3.3 The Market Game

We model a logarithmic market scoring rule (LMSR) [Hanson, 2007b] prediction market as a Bayesian extensive-form game. Our setting is similar to that of these prior work [Chen et al., 2010b, 2007, Dimitrov and Sami, 2007, Ostrovsky, 2012]. We study the LMSR market for forecasting a binary event where $\Omega = \{0, 1\}$ denotes the outcome space and $\omega \in \Omega$ denotes the realized outcome of this event. Many real-world prediction markets focus on such binary events, for example “whether the UK economy will go into recession in 2013”, “whether the movie Lincoln will win the Academy Award for Best Picture”, and “whether a Democrat will win the US Presidential election in 2016”.

3.3.1 The Finite-Stage Market Game

Our LMSR market game has n stages and $m \leq n$ players. The players participates in one or more stages of the market game, following a pre-defined sequence, which is common

knowledge¹.

Each player i has private information about the event given by a private signal $s_i \in S_i$ with signal space S_i and $|S_i| = n_i$. Each signal is only observed by the intended player. The prior distribution of the event outcomes and the players' private signals, denoted by $\mathcal{P} : \Omega \times S_1 \times \cdots \times S_m \rightarrow [0, 1]$, is common knowledge. Before the market starts, nature draws the realized event outcome and the private signals of the players according to $\mathcal{P}$. The players receive their private signals before the market opens, and the realized event outcome is revealed after the market closes.

The players are risk-neutral Bayesian agents. The belief of the player participating in stage t can depend on the reported estimates in the first $t - 1$ stages as well as on his own private signal.

The 3-Stage Market Game. The simplest version of the market game that admits non-trivial strategic play is a 2-player 3-stage game. The two players are Alice and Bob, and the sequence of participation is Alice, Bob, and then Alice. Alice and Bob received private signals $s_A \in S_A, |S_A| = n_A$ and $s_B \in S_B, |S_B| = n_B$ respectively. The analysis of this 3-stage market game will serve as building blocks for our analysis of the finite-stage market game.

3.3.2 Information Structure

The prior distribution $\mathcal{P}$ is a critical component of each instance of the market game. It encodes the relationship between the players' private signals and the event outcome, and it enables players with private signals to reason about other players' signals and the realized event outcome. We refer to $\mathcal{P}$ as the “*information structure*” of the market game. The primary goal of this work is to characterize the strategic play in a market game with a given information structure.

¹It is an interesting future direction to consider a model where players endogenously choose when to participate. However, our equilibrium results for the D game with a pre-defined participation order imply that players will delay revealing their information as much as possible in the D game even with endogenously chosen participation order. We discuss these implications in section 3.5 after our equilibrium results.

Three Classes of Information Structures

There are three classes of information structures: conditionally independent (*CI game*), unconditionally independent (*I game*), and neither conditionally independent nor unconditionally independent (*D game*). These three classes are mutually exclusive and exhaustive. The first two types impose natural independence assumptions on the prior distribution $\mathcal{P}$, and they were first separately studied by Chen et al. [2007] and Dimitrov and Sami [2007], and later in their joint work [Chen et al., 2010b].

In a CI game, players' signals are independent conditioned on the realized event outcome. Prior work [Chen et al., 2007, 2010b] showed that there is a unique type of perfect Bayesian equilibria (PBE) for the CI game where players honestly report their estimates as early as possible. Thus, in this work, we focus on analyzing the I and D games.

For I games, players' signals are unconditionally independent from one another, but they are not independent of and may stochastically influence the event outcome. Formally, the prior distribution $\mathcal{P}$ for an I game must satisfy: $P(s_i)P(s_j) = P(s_i, s_j), \forall s_i \in S_i, s_j \in S_j$ for any two players i and j . Dimitrov and Sami [2007] and Chen et al. [2010b] showed that the I game does not have a truthful PBE where every player honestly reports his estimate as early as he can, but they left the existence of PBE as an open question.

The I information structure can be motivated by several examples. For stylized ones, consider a setting where each player independently observes a coin flip. The event to be predicted is some aggregate information about all of the independent coin flips, for example, whether more than 1/3 of the coin flips are heads. In this example, the players' signals are independent because the coin flips are independent events. For an abstract example, each player's private information can be thought of as a single piece of a jigsaw puzzle, and the event being forecasted is related to the completed picture. For a more realistic example, consider a flu prediction scenario. Several doctors live in different regions of the country. Each doctor gets information about the flu by treating his own patients living in his region. So the doctors' information about the flu is arguably independent because the patients'

health conditions in different regions are independent.

Even though the CI and I information structures capture events in some natural settings, they impose strong independence assumptions on the relationship between the players' private signals. Ideally, we would like to understand the players' strategic behavior in the market game without restricting to a particular information structure. For this reason, we study the D information structure consisting of signals that are neither conditionally independent nor unconditionally independent. In other words, the signals in a D game are both conditionally dependent and unconditionally dependent. Formally, a prior distribution $\mathcal{P}$ in a D game satisfies: $\exists s_i \in S_i, s_j \in S_j$, s.t. $P(s_i)P(s_j) \neq P(s_i, s_j)$ for two players i and j and $\exists s_{i'} \in S_i, s_{j'} \in S_j, \omega \in \Omega$, s.t. $P(s_{i'}, s_{j'}|\omega) \neq P(s_{i'}|\omega)P(s_{j'}|\omega)$ for two players i' and j' . It would be interesting to explore whether the D information structure could be further divided up into smaller classes with intuitive properties.

The Distinguishability Condition

To avoid degenerate cases in our analysis, we assume that the prior distribution $\mathcal{P}$ satisfies the following *distinguishability* condition, consisting of two parts.

Definition 1. *The prior distribution $\mathcal{P}$ satisfies the distinguishability condition if for all i it satisfies inequality (3.1)*

$$P(1|\mathbf{s}_{-i}, s_i) \neq P(1|\mathbf{s}_{-i}, s'_i), \forall \mathbf{s}_{-i} \in \mathbf{S}_{-i}, \forall s_i, s'_i \in S_i \cup \{\phi\}, s_i \neq s'_i \quad (3.1)$$

where $s_i = \phi$ means player i 's private signal is not observed, and $\mathbf{S}_{-i} = \{S_1 \cup \{\phi\}\} \times \cdots \times \{S_{i-1} \cup \{\phi\}\} \times \{S_{i+1} \cup \{\phi\}\} \times \cdots \times \{S_m \cup \{\phi\}\}$, and inequality (3.2)

$$\sum_{s_i \in S_i} p_{s_i} P(1|s_i, \mathbf{s}) \neq \sum_{s_i \in S_i} p_{s_i} P(1|s_i, \mathbf{s}') \quad (3.2)$$

where $\mathbf{s} \neq \mathbf{s}'$ are any two different vectors of realized signals of any subset of players excluding i , and the vector $(p_{s_i})_{s_i \in S_i}$ is any probability distribution over S_i .

Inequality (3.1) generalizes the general informativeness condition by Chen et al. [2010b].

The inequality is satisfied if different signal realizations of player i always lead to different posterior probabilities of $\omega = 1$, for any vector of realized signals for any subset of the other players (including unobserved signals). In other words, a player's signal always contains some information. Inequality (3.2) is similar to the distinguishability assumption used by Dimitrov and Sami [2010a]. It requires that for any two realizations of signals of a subset of players, they lead to different estimates for outcome $\omega = 1$ given any belief about player i 's signal. This condition allows other players to infer the signals of the subset of players whenever they reveal their information truthfully.

While the distinguishability condition may be a nontrivial technical restriction, it allows us to focus on interesting strategic decisions in the game play without encountering degenerated cases.

3.3.3 Solution Concept and Players' Strategies

We use the *perfect Bayesian equilibrium* (PBE), which is informally a subgame perfect refinement of the Bayesian Nash equilibrium, as our solution concept. A PBE requires specifying each player's strategy given a realized signal at each stage of the game as well as the player's belief about the signals of players participating in all of the previous stages. The strategies and the beliefs of the players form a PBE of the market game if and only if, for each player, his strategy at every stage is optimal given the beliefs, and the beliefs are derived from the strategies using Bayes' rule whenever possible.

By properties of the logarithmic scoring rule, at a player's last chance to participate in the market, the player has the strictly dominant strategy of truthfully revealing his private information. So at any PBE, all private information is fully incorporated into the market estimate at the end of the market game. Thus, the focus of our analysis is on how quickly information gets incorporated into the market estimate throughout the game. In the following paragraphs, we distinguish between truthful and non-truthful strategies for a player in terms of when the player's private information is first revealed in the market game.

We use the term *truthful strategy* (also called truthful betting) to refer to the strategy where at a player's first chance to participate in the market, the player changes the market estimate to his posterior probability of outcome $\omega = 1$ given his signal and his belief about other players' signals. The truthful strategy fully reveals a player's private information as early as possible.

In contrast to the truthful strategy, a player may choose to misreport his information and manipulate the market estimate. For instance, a player can play a mixed strategy and reveal a noisy version of his signal to the subsequent players in the game. Alternatively, a player may try to withhold his private information from the other players by not changing the market estimate at all. Such non-truthful strategies hurt information aggregation in the market by causing the market estimate to contain inaccurate information at least temporarily.

3.4 The 3-Stage Market Game with Any Information Structure

Before diving into the PBE analysis of the finite-stage market game, we describe some preliminary analysis of the 3-stage market game with any information structure. In section 3.4.1, we justify that, in order to describe a PBE of the 3-stage market game, it suffices to describe Alice's strategy in the first stage and Bob's belief in the second stage. This allows us to greatly simplify our exposition in later analyses. Next, we prove a theorem in section 3.4.2, which allows us to systematically identify candidate PBE strategies for the players. This theorem gives us a useful method to make educated guesses about the possible PBE strategies in order to tackle the PBE existence question and to construct a PBE if one exists for the 3-stage market game with a given prior distribution. Finally, in section 3.4.3, we describe a consistency condition, which must be satisfied by a player's strategy in any PBE of the 3-stage game.

3.4.1 Describing PBE of the 3-Stage Market Game

We present a preliminary analysis of the 3-stage market game and introduce some notations for our later analyses.

In the 3-stage game, Alice and Bob observe their realized signals s_A and s_B respectively at the beginning of the market. In the first stage, Alice changes the market estimate for outcome $\omega = 1$ from the initial market estimate r^0 to r_A of her choice. In the second stage, Bob observes Alice's first-stage report r_A and changes the market estimate to r_B . In the third stage, upon observing Bob's second-stage report r_B , Alice changes the market estimate from r_B to r^f , and then the market closes.

Alice's first-stage strategy is a mapping $\sigma : S_A \rightarrow \Delta([0, 1])$ where $\Delta([0, 1])$ is the set of probability distributions over $[0, 1]$. For clarity of analysis and presentation, we assume that the support of Alice's first-stage strategy is finite. The results in this work however hold even if the support of Alice's first-stage strategy is infinite. Let $\sigma_{s_A}(r_A)$ denote the probability that Alice reports r_A in the first stage after observing the signal s_A according to the strategy σ .

In the second stage, when Bob observes Alice's first-stage report r_A , he forms a belief about Alice's signals. Bob's belief specifies the likelihood that Alice received signal s_A when Alice reported r_A and Bob received signal s_B . Let $\mu_{r_A, s_B}(s_A)$ denote the probability that Bob's belief assigns for Alice's s_A signal when Alice reported r_A and Bob received signal s_B . $\mu_{r_A, s_B}(s_A)$ is defined for any $r_A \in [0, 1]$. At any PBE, we need to describe Bob's belief both on and off the equilibrium path. When r_A is in the support of Alice's first-stage PBE strategy, the game is on the equilibrium path and $\mu_{r_A, s_B}(s_A)$ is derived from Alice's strategy using Bayes' rule according to the PBE definition. However, when r_A is not in the support of Alice's equilibrium strategy, that is, the game is off the equilibrium path, $\mu_{r_A, s_B}(s_A)$ is still important for a PBE because the belief needs to ensure that Alice does not find it profitable to deviate from her PBE strategy. Off the equilibrium path, there are often more than one set of Bob's beliefs that can satisfy this requirement.

Bob only participates once, in the second stage of this game. By properties of strictly proper scoring rules, Bob has a strictly dominant strategy to report his posterior probability estimate of the event truthfully, given his belief. Thus, at any PBE, Bob must be using a pure strategy, which is fully determined by his belief, his signal, and Alice's first-stage report. Let $x_{s_B}(r_A)$ denote Bob's optimal report given his signal s_B and Alice's first-stage report r_A . At any PBE, Bob's optimal report $x_{s_B}(r_A)$ can be determined from his belief as follows:

$$x_{s_B}(r_A) = \sum_{s_A} \mu_{r_A, s_B}(s_A) P(1|s_A, s_B), \quad \forall s_A \in S_A, s_B \in S_B, r_A \in [0, 1].$$

In the third stage, Alice observes Bob's report and may change the market estimate again. At any PBE, knowing Bob's PBE strategy, Alice's belief on the equilibrium path can be derived from Bob's strategy using Bayes' rule. This is Alice's last stage of participation. Thus, by properties of strictly proper scoring rules, Alice has a strictly dominant strategy to report her probability estimate truthfully. Similar to Bob's strategy, Alice's third-stage strategy must be a pure strategy and it is fully determined by her belief, her signal, and Bob's report. We note that Alice's belief off the equilibrium path in the third stage is not important, because Bob has a dominant strategy in the second stage and will not deviate from it no matter what belief Alice has.

The above analysis shows that, to describe a PBE of the 3-stage market game, it suffices to specify Alice's strategy in the first stage and Bob's belief in the second stage. The rest of the strategic play is completely determined given them.

Moreover, for clarity in our analysis, we specify Bob's strategy rather than Bob's belief at a PBE. We can easily derive a belief of Bob such that Bob's strategy is optimal given it, shown as follows. First, Bob's strategy is valid if and only if $x_{s_B}(r_A) \in [\min_{s_A} P(1|s_A, s_B), \max_{s_A} P(1|s_A, s_B)]$ for any s_B , because for any possible belief for Bob, his posterior probability should always fall into this interval. When r_A is in the support of Alice's PBE strategy, Bob's belief is derived from Alice's PBE strategy using Bayes' rule.

When r_A is not in the support of Alice's PBE strategy, the PBE definition requires that Bob's belief be derived from a possible strategy for Alice using Bayes' rule. For such an r_A and for any s_B , we know that $\min_{s_A} P(1|s_A, s_B) \leq x_{s_B}(r_A) \leq \max_{s_A} P(1|s_A, s_B)$ holds and one of the two inequalities must be strict due to the distinguishability assumption. For a given s_B , let $s'_A = \arg \min_{s_A} P(1|s_A, s_B)$ and $s''_A = \arg \max_{s_A} P(1|s_A, s_B)$. Then consider a possible strategy satisfying $\sigma_{s'_A}(r_A) = p$, $\sigma_{s''_A}(r_A) = 1 - p$ and $\sigma_{s_A}(r_A) = 0$ for any other s_A , where $p = \frac{P(s''_A|s_B)(x - P(1|s'_A, s_B))}{P(s'_A|s_B)(P(1|s'_A, s_B) - x) + P(s''_A|s_B)(x - P(1|s''_A, s_B))}$. This strategy for Alice is valid, and thus we can derive Bob's off the equilibrium path belief for r_A from this strategy using Bayes' rule.

3.4.2 Systematically Identify Candidate PBE Strategies

To tackle the PBE existence problem and construct a PBE if one exists, it is essential that we make an educated guess of the players' possible PBE strategies. Theorem 1 below allows us to pinpoint a possible PBE strategy for Alice in the 3-stage game with any information structure, by comparing Alice's ex-ante expected total payoff (of both the first and the third stages) when using different first-stage strategies assuming that Bob knows and conditions on Alice's strategy.

For Theorem 1 below, for any of Alice's strategy σ_1 , let $\pi_A(\sigma_1, \sigma_1)$ be Alice's ex-ante expected payoff when Alice uses strategy σ_1 in the first stage, Bob knows Alice's first-stage strategy σ_1 and conditions his belief on this strategy. This means that, for any r in the support of Alice's first-stage strategy σ_1 , Bob's belief is derived from strategy σ_1 by using Bayes' rule. For any other r , there is no restriction on Bob's belief as long as it is valid.

In the proof of Theorem 1, we make an important distinction between a player's ex-ante and ex-interim expected payoff. A player's ex-ante expected payoff is his expected payoff without observing his signal, whereas his ex-interim expected payoff is his expected payoff given his signal.

Theorem 1. *For the 3-stage market game, if two different first-stage strategies σ_1 and σ_2*

for Alice satisfy inequality (3.3), then strategy σ_2 cannot be part of any PBE of this game.

$$\pi_A(\sigma_1, \sigma_1) > \pi_A(\sigma_2, \sigma_2) \quad (3.3)$$

Proof. We prove this by contradiction. Suppose that two different first-stage strategies σ_1 and σ_2 for Alice satisfy inequality (3.3), and Alice's first-stage strategy σ_2 is part of a PBE of the 3-stage market game. Let μ_B denote Bob's belief at this PBE. μ_B specifies a distribution over Alice's signals for every possible first-stage report $r \in [0, 1]$ and any of Bob's signals s_B . Alice's ex-ante expected payoff at this PBE is $\pi_A(\sigma_2, \sigma_2)$. This proof holds for any valid belief for Bob at this PBE.

Suppose that Alice deviates from this PBE to play the strategy σ_1 in the first stage and Bob has the same belief μ_B as before. Let $\pi_A(\sigma_1, \sigma_2)$ denote Alice's total ex-ante expected payoff in the game at this deviation. The expression $\pi_A(\sigma_1, \sigma_2)$ is well defined since Alice knows Bob's belief and strategy at the original PBE. Similarly, let $\pi_B(\sigma_1, \sigma_2)$ denote Bob's ex-ante expected payoff in the second stage at this deviation.

At any PBE of this game, in the third stage, Alice can always infer Bob's signal given Bob's report by the distinguishability condition. So Alice always changes the market estimate to $P(1|s_A, s_B)$ in the third stage given Alice's signal s_A and Bob's signal s_B . Thus, the total expected payoff that Alice and Bob can get at any PBE of the 3-stage market game is

$$\pi_{AB} = \sum_{s_A, s_B} \left\{ P(1, s_A, s_B) \log \frac{P(1|s_A, s_B)}{r^0} + P(0, s_A, s_B) \log \frac{P(0|s_A, s_B)}{1 - r^0} \right\}$$

which is fixed given the initial probability r^0 and the prior distribution $\mathcal{P}$. Note that the above result holds not only at a PBE but whenever Bob reveals all of his information and Alice knowing his strategy maximizes her expected payoff. Therefore, by definition of π_{AB} , we must have

$$\pi_{AB} = \pi_A(\sigma_1, \sigma_2) + \pi_B(\sigma_1, \sigma_2), \forall \sigma_1, \sigma_2 \quad (3.4)$$

Inequality (3.3) is satisfied by assumption, so we have

$$\begin{aligned}\pi_A(\sigma_1, \sigma_1) &> \pi_A(\sigma_2, \sigma_2) \\ \Rightarrow \pi_{AB} - \pi_A(\sigma_1, \sigma_1) &< \pi_{AB} - \pi_A(\sigma_2, \sigma_2)\end{aligned}\tag{3.5}$$

$$\Rightarrow \pi_B(\sigma_1, \sigma_1) < \pi_B(\sigma_2, \sigma_2)\tag{3.6}$$

where equation (3.5) is due to equation (3.4).

For a fixed first-stage strategy of Alice and for any belief of Bob, Bob's ex-ante expected payoff is maximized when his belief is derived from Alice's first-stage strategy using Bayes' rule. This can be proven as follows. When Bob's belief is derived from Alice's first-stage strategy by using Bayes' rule, then in the second stage, Bob changes the market estimate to $x_{s_B}(r_A)$ when Alice reports r_A in the first stage and Bob receives the s_B signal. Recall that by definition, $x_{s_B}(r_A) = P(1|r_A, s_B) = \sum_{s_A} P(s_A|r_A, s_B)P(1|s_A, s_B)$. In this case, Bob's expected payoff in the second stage is

$$\sum_{s_B, r_A} P(s_B, r_A) \left\{ x_{s_B}(r_A) \log \frac{x_{s_B}(r_A)}{r_A} + (1 - x_{s_B}(r_A)) \log \frac{1 - x_{s_B}(r_A)}{1 - r_A} \right\}.\tag{3.7}$$

When Bob has another belief, let $\hat{x}$ denote Bob's optimal report with this belief. Then Bob's expected payoff in the second stage is

$$\sum_{s_B, r_A} P(s_B, r_A) \left\{ x_{s_B}(r_A) \log \frac{\hat{x}}{r_A} + (1 - x_{s_B}(r_A)) \log \frac{1 - \hat{x}}{1 - r_A} \right\}.\tag{3.8}$$

The difference in Bob's ex-ante expected payoff for the two different beliefs for Bob is (3.7) - (3.8):

$$\sum_{s_B, r_A} P(s_B, r_A) \left\{ x_{s_B}(r_A) \log \frac{x_{s_B}(r_A)}{\hat{x}} + (1 - x_{s_B}(r_A)) \log \frac{1 - x_{s_B}(r_A)}{1 - \hat{x}} \right\}$$

which is nonnegative by properties of the relative entropy.

Therefore, for any two first-stage strategies σ_1 and σ_2 for Alice, we have shown that

$$\pi_B(\sigma_1, \sigma_2) \leq \pi_B(\sigma_1, \sigma_1)\tag{3.9}$$

Combining inequalities (3.6) and (3.9), we have

$$\begin{aligned}
& \pi_B(\sigma_1, \sigma_2) < \pi_B(\sigma_2, \sigma_2) \\
& \Rightarrow \pi_{AB} - \pi_A(\sigma_1, \sigma_2) < \pi_{AB} - \pi_A(\sigma_2, \sigma_2) \\
& \Rightarrow \pi_A(\sigma_1, \sigma_2) > \pi_A(\sigma_2, \sigma_2).
\end{aligned} \tag{3.10}$$

According to inequality (3.10), if Alice uses the first-stage strategy σ_2 at a PBE, then she can improve her ex-ante expected payoff by deviating to using the strategy σ_1 . Then there must exist at least one realized signal for Alice, say s_A , such that Alice's ex-interim expected payoff after receiving the s_A signal is higher when she deviates to the strategy σ_1 than when she follows the strategy σ_2 . (Otherwise, if Alice's ex-interim expected payoff for every realized signal is lower when she deviates to using the strategy σ_1 than when she follows the strategy σ_2 , then her ex-ante expected payoff must also be lower when she deviates to using the strategy σ_1 than when she follows the strategy σ_2 , contradicting inequality (3.10).) As a result, when Alice receives the s_A signal, she can improve her ex-interim expected payoff by deviating to using the strategy σ_1 and this contradicts with our assumption that Alice's first-stage strategy σ_2 is part of a PBE of the 3-stage market game. $\square$

According to Theorem 1, to find Alice's possible PBE strategies for the 3-stage market game, it suffices to compare Alice's ex-ante expected payoffs for all possible first-stage strategies assuming Bob knows Alice's strategy, and only the strategies maximizing Alice's ex-ante expected payoff can possibly be Alice's PBE strategy. This gives us a systematic way to identify possible PBE strategies without worrying about constructing Bob's off-equilibrium path beliefs.

3.4.3 The Consistency Condition

Our analyses of the 3-stage game frequently make use of a consistency condition described in Theorem 2 by Chen et al. [2010b]. For completeness, we re-state this condition as a lemma

below. The consistency condition requires that, at a PBE of the 3-stage game, for any r_A in the support of Alice's first-stage strategy σ , the posterior probability of outcome $\omega = 1$ given σ and r_A should be equal to r_A . Intuitively, this requires that, Alice's first-stage strategy must not leave free payoff for Bob to claim in the second stage. If Alice's first-stage strategy does not satisfy the consistency condition, then Bob can get positive expected payoff simply by changing the market estimate to a value satisfying the consistency condition, and Bob can claim this positive expected payoff without having any private information about the event being predicted. This is contrary to Alice's goal of minimizing Bob's expected payoff since the 3-stage market game is a constant-sum game in expectation at any PBE.

Lemma 1 (Consistency Condition for 3-Stage Market Game). *At a PBE of the 3-stage market game, if σ is Alice's first-stage strategy and r is in the support of strategy σ (i.e. $\exists s_A \in S_A, \sigma_{s_A}(r) > 0$), then σ must satisfy the following consistency condition:*

$$P(1|\sigma, r) = r$$

3.5 PBE of the Finite-Stage I Game

We characterize all PBE of the finite-stage I game in this section. Our analysis begins with the 3-stage I game. Alice participates twice in the game, so she may have incentives to manipulate the market estimate in the first stage. We first identify a unique candidate PBE strategy for Alice by showing that if a PBE exists for the 3-stage I game, then Alice's first-stage strategy must be changing the market estimate to the prior probability of the event. This is equivalent to Alice delaying her participation until the third stage if the market starts with the prior probability of the event. We refer to this strategy as Alice's *delaying* strategy for the 3-stage I game. Alice's *delaying* strategy reveals absolutely no information to Bob about her signal. Next, we explicitly construct a PBE of the 3-stage I game in which Alice uses the *delaying* strategy in the first stage. These two results together imply that, the *delaying* PBE is unique for this game, in the sense that Alice must use the delaying strategy

in every PBE of this game, even though Bob’s belief can be different off the equilibrium path.

Given the *delaying* PBE of the 3-stage I game, we construct a family of PBE for the finite-stage I game using backward induction. Suppose that the players in the finite-stage I game are ordered by their last stages of participation. Then at every PBE of the finite-stage I game, each player i withholds his private information until after player $i - 1$ finishes participating in the game, and then player i may truthfully reveal his private information in any of the subsequent stages in which he participates. In particular, there exists a particular PBE in this family where each player does not reveal any private information until his last stage of participation, and this is arguably the worst PBE of this game for the goal of information aggregation.

3.5.1 Delaying PBE of 3-stage I Game

We argue below that the delaying strategy is the only candidate PBE strategy for Alice in the 3-stage I game. Theorem 2 essentially proves that the delaying PBE of the 3-stage I game is unique with respect to Alice’s strategy, if a PBE exists for this game. Part of the proof of Theorem 2 uses the argument in the proof of Theorem 2 in Chen et al. [2010b].

Theorem 2. *If the 3-stage I game has a PBE, then Alice’s strategy at the PBE must be the delaying strategy, i.e. changing the market estimate to the prior probability of the event in the first stage.*

The proof of Theorem 2 is included in Appendix A.1.1.

Sketch. We first argue that if a PBE exists for the 3-stage I game, then Alice’s first-stage strategy at this PBE must be a deterministic strategy. We show this by contradiction by assuming that there are at least two points in the support of Alice’s first-stage PBE strategy. Then we construct another first-stage strategy achieving a better expected payoff for Alice, which means that the original strategy cannot be a PBE strategy by Theorem 1. By the

consistency condition, if Alice's first-stage strategy is deterministic, it must be the strategy of changing the market estimate to the prior probability of the event. $\square$

While the delaying strategy is the only possible PBE strategy for the 3-stage I game, we still don't know whether a PBE exists. In order for a PBE to exist, there must exist a belief of Bob to ensure that Alice does not find it profitable to deviate from the delaying strategy to any other strategy. Identifying such a belief for Bob can be challenging because essentially we need to specify what Bob will do upon observing every possible report of Alice in $[0, 1]$. In Theorem 3, we give an explicit construction of a PBE of the 3-stage I game in which Alice uses the *delaying* strategy in the first stage. At this PBE, Alice's first-stage strategy reveals no information to Bob about her private signal, and Bob's belief makes this delaying strategy the optimal choice for Alice.

Theorem 3. *There exists a PBE of the 3-stage I game where Alice's first-stage strategy is*

$$\sigma_{s_A}(P(1)) = 1, \quad \forall s_A \in S_A$$

and Bob's second-stage strategy is

$$x_{s_B}(r_A) = \begin{cases} f_{s_B}(\alpha_{s_B}^{\min}), & r_A \in [0, \alpha_{s_B}^{\min}) \\ f_{s_B}(r_A), & r_A \in [\alpha_{s_B}^{\min}, \alpha_{s_B}^{\max}] \\ f_{s_B}(\alpha_{s_B}^{\max}), & r_A \in (\alpha_{s_B}^{\max}, 1] \end{cases}, \quad \forall s_B \in S_B$$

where

$$\begin{aligned} f_{s_B}(r_A) &= \frac{P(1|s_B)P(0)r_A}{P(1)P(0|s_B) + (P(1|s_B) - P(1))r_A} \\ \beta_{s_B}^{\min} &= \min_{s_A} P(1|s_A, s_B), \beta_{s_B}^{\max} = \max_{s_A} P(1|s_A, s_B) \\ \alpha_{s_B}^{\min} &= f_{s_B}^{-1}(\beta_{s_B}^{\min}), \alpha_{s_B}^{\max} = f_{s_B}^{-1}(\beta_{s_B}^{\max}) \end{aligned}$$

The proof of Theorem 3 is included in Appendix A.1.2.

Sketch. We describe the first part of the proof below showing that Bob's strategy is a valid

PBE strategy.

First, Bob's belief on the equilibrium path is derived from Alice's first-stage strategy using Bayes' rule since $x_{s_B}(P(1)) = P(1|s_B)$. Moreover, for Bob's strategy to be a valid PBE strategy, it must satisfy $x_{s_B}(r_A) \in [\min_{s_A} P(1|s_A, s_B), \max_{s_A} P(1|s_A, s_B)]$, $\forall s_B, r_A \in [0, 1]$. To show this, note that by definition, $\beta_{s_B}^{\min} < \beta_{s_B}^{\max}$, $\alpha_{s_B}^{\min} < \alpha_{s_B}^{\max}$, and $f_{s_B}(r_A)$ is monotonically increasing in $r_A \in [0, 1]$ since

$$\frac{df_{s_B}(r_A)}{dr_A} = \frac{P(1)(1 - P(1))P(1|s_B)(1 - P(1|s_B))}{\{P(1)P(0|s_B) + (P(1|s_B) - P(1))r_A\}^2} > 0$$

Hence the domain of $x_{s_B}(r_A)$ is well-defined. In addition, we have

$$\beta_{s_B}^{\min} = f_{s_B}(\alpha_{s_B}^{\min}) \leq x_{s_B}(r_A) \leq f_{s_B}(\alpha_{s_B}^{\max}) = \beta_{s_B}^{\max}, \forall r_A \in [0, 1].$$

Thus, Bob's strategy is valid. The rest of the proof then proves that Alice's delaying strategy is a best response to Bob's strategy. $\square$

Based on Theorems 2 and 3 above, we have established both the existence and the uniqueness (with respect to Alice's first-stage strategy) of the PBE for the 3-stage I game.

3.5.2 A Family of PBE for the Finite-Stage I Game

We are ready to characterize the PBE of the finite-stage I game. By using backward induction and the delaying PBE of the 3-stage I game, we characterize a family of PBE of the finite-stage I game in Theorem 4. At any PBE in this family, players delay revealing their private information as much as possible.

We first generalize the consistency condition for the 3-stage game to the finite-stage game in Lemma 2. This consistency condition dictates that, for any stage k , the posterior probability of $\omega = 1$ given the participants' strategies and reports in the first k stages must be equal to the report of the participant in stage k at any PBE of this game.

Lemma 2 (Consistency Condition for Finite-Stage Market Game). *At a PBE of the finite-*

stage I game, suppose that σ^k and r^k are the strategy and the report for the participant of stage k respectively, then for every k , the participants' strategies and reports must satisfy equation (3.11).

$$P(1|r^1, \dots, r^k, \sigma^1, \dots, \sigma^k) = r^k \quad (3.11)$$

The proof of Lemma 2 is included in Appendix A.1.3.

In Lemma 3 below, we analyze the *tail* of the finite-stage I game starting from the second-to-last stage of participation for the last player to the last stage of the game. The theorem shows that, in terms of strategic play, this portion of the finite-stage I game essentially reduces to a 3-stage I game. Thus, at any PBE, the last player chooses to not participate in the game in his second-to-last stage of participation. This key argument will be used repeatedly in the proof of the PBE of the finite-stage I game.

For Lemma 3 and Theorem 4, let the m players of the finite-stage I game be ordered by their last stages of participation. That is, for any $1 \leq i \leq m$, let t_i denote player i 's last stage of participation, such that $t_i < t_j$ for any $1 \leq i < j \leq m$. Without loss of generality, we assume that player m has more than one stages of participations.

Lemma 3. *Let stage k be the second to last stage of participation for player m ($k < t_m$). At any PBE of the finite-stage I game, player m does not change the market estimate in stage k .*

The proof of Lemma 3 is included in Appendix A.1.4.

Finally, in Theorem 4, we prove the existence of a family of PBE of the finite-stage I game.

Theorem 4. *At any PBE of the finite-stage I game, the players use the following strategies:*

- *From stage 1 to stage $t_1 - 1$, player 1 uses any strategy that satisfies the consistency condition. In stage t_1 , player 1 truthfully reveals his signal.*

- For any $2 \leq i \leq m - 1$, from stage 1 to stage $t_{i-1} - 1$, player i does not participate in the game. From stage $t_{i-1} + 1$ to stage $t_i - 1$, player i uses any strategy that satisfies the consistency condition. In stage t_i , player i truthfully reveals his signal.
- From stage 1 to stage $t_m - 1$, player m does not participate in the game. In stage t_m , player m truthfully reveals his signal.

The proof of Theorem 4 is included in Appendix A.1.5.

Sketch. We describe the argument for player m and $m - 1$ here.

By properties of LMSR, player m truthfully reveals his signal in stage t_m , which is the last stage of the game. If stage t^* denotes the second to last stage of participation for player m , then the game from stage t^* to t_m can be reduced to a 3-stage I game (where player m is Alice and other players participating between t^* and t_m are a composite Bob). By Lemma 3, player m does not participate in stage t^* . Now remove this stage and let t^* be the *new* second to last stage of participation for player m , and the game from stage t^* to t_m again reduces to a 3-stage I game. Applying Lemma 3 again, we know that player m does not participate in stage t^* either. Inferring recursively, player m does not participate in any stage from 1 to $t_m - 1$.

For player $m - 1$, he truthfully reveals his signal in stage t_{m-1} by properties of LMSR. From stage $t_{m-2} + 1$ to $t_{m-1} - 1$, player $m - 1$ is the only participant because players 1 to $m - 2$ already finished participating and player m does not participate by our earlier argument. Thus, player $m - 1$ uses any strategy satisfying the consistency condition from stage $t_{m-2} + 1$ to $t_{m-1} - 1$. We combine the stages from $t_{m-2} + 1$ to $t_{m-1} - 1$ (denoted t^{**}) as the new last stage for player $m - 1$. Let t^* be the new second to last stage of participation for player $m - 1$, and note that $t^* < t_{m-2}$. Again, the game from stage t^* to t^{**} reduces to a 3-stage I game (where player $m - 1$ is Alice). By Lemma 3, player $m - 1$ does not participate in stage t^* . Inferring recursively, player $m - 1$ does not participate in any stage from 1 to $t_{m-2} - 1$. □

To understand Theorem 4, consider dividing the finite-stage I game into m segments with player i being the *owner* of the segment from stage $t_{i-1} + 1$ to stage t_i . At any PBE, each player does not participate in any stage before his segment, uses a strategy satisfying the consistency condition within his segment, and truthfully reveals his private signal at the last stage of his segment.

Figure 3.1 illustrates a particular PBE of a finite-stage I game. The letters A , B , and C denote the three players and their sequence of participation. A black letter means that the player truthfully reveals his signal in that stage. If the letter is gray, then the player uses a strategy satisfying the consistency condition. Note that the strategy of not changing the market estimate satisfies the consistency condition. A white letter means that the player is scheduled to participate but does not change the market estimate in that stage. The thick vertical bars mark the boundaries of the players' segments in the game.



Figure 3.1: A PBE of a Finite-Stage I Game with 3 players

The multiple PBE of the finite-stage I game differ by how early each player chooses to truthfully reveal his signal within his segment of the game. For the purpose of information aggregation, the best case is when every player chooses to truthfully reveal his signal in the first stage of his own segment. However, there exists a PBE where every player waits until the last stage of his segment to truthfully reveal his information, and this is arguably the worst PBE for the goal of information aggregation.

Although our model assumes a pre-specified participation order, our results still provide useful insights for the I game if the players endogenously choose when to participate in the game. Consider the I game with n stages and $m < n$ players where each player endogenously chooses in which stage to participate in the game. Our results for the I game suggest that, at any PBE all players will choose to delay their participation and no information is reveal in the first $n - m$ stages. The exact characterization of PBE would critically depend on how multiple

trades submitted in the same stage are executed. This dependency is generally undesirable. Our assumption of pre-specified participation order circumvents this dependency and we believe our results still provide useful insights for players' behavior in this setting.

3.5.3 Discussion

The delaying PBE for the I game is arguably the worst outcome for the purpose of information aggregation since each player's private information may be incorporated into the market at his last chance to trade in the market. For our market game, closing the market early may solve this problem by ensuring that each player's information is incorporated early in the market, as long as each player gets to trade at least once in the market. However, in practice, closing the market early may not solve this problem. First, for our model, we make the simplifying assumption that all players receive their private information before the market opens and they do not receive new information later on. In practice, players do receive new information over time, and closing the market early may prevent these new information from being incorporated into the market. Moreover, another simplifying assumption in our model is that player's participation order is exogenously determined. In practice, players endogenously determine when to participate in the market. Thus, if we reduce the time the market is open, then some players may not be able to participate in the market, and thus their information does not get incorporated into the market forecast.

When comparing the PBE of the finite-stage I game with the truthful PBE of the finite-stage CI game [Chen et al., 2010b], it is interesting to note how two different information structures can induce equilibrium behavior at the opposite ends of the spectrum: The players in the CI game race to reveal their private information as early as possible, whereas the players in the I game delay as much as possible to reveal their private information.

This difference is spiritually consistent with the concepts of complementarity and substitution of private signals defined by Chen et al. [2010b]. Consider the ex ante expected payoffs of players. In the I games, players' private signals can be intuitively considered as

complements. When the current market prediction is the prior probability, the sum of players' expected payoffs when each player reports a posterior probability conditioned only on his own private signal is strictly less than the total expected payoff that can be earned by reporting a posterior probability conditioned on all of the available private signals in any I game. This means that, every player in the I game prefers to wait for other players to make their reports first since observing more reports and thus inferring more signals improve the player's expected payoff. In contrast, in the CI games, players' private signals are substitutes. For any current market prediction, the sum of players' expected payoffs when each player reports a posterior probability conditioned only on his own private signal is strictly greater than the total expected payoff that can be earned by reporting a posterior probability conditioned on all of the available private signals. Thus, players prefer to race to capitalize on their private information early in the game.

3.6 The 3-Stage D Game

The CI and I games admit two families of PBE that seem to lie at the two extremes of the spectrum: players race to reveal information early in the CI game, but race to withhold information in the I game. It is interesting to ask whether some instances of the D game may give rise to one of these two types of equilibria too. Yet, it is challenging to perform equilibrium analysis for the D game, because the dependency among the players' signals does not provide precise mathematical conditions that we can leverage.

Our goal in this section is moderate. We would like to explore a restricted 3-stage D game and obtain insights on what the players' PBE strategies may look like for this game if a PBE exists. We do not prove the existence of a PBE for this class. Nevertheless, we provide a sufficient condition for the prior distribution, which guarantees the existence of a truthful PBE for the D game. We also provide an example distribution that satisfies this condition.

In this section, we consider the 3-stage D game where Alice's private signal has only 2 realizations. For this special case, we use a_0 and a_1 to denote Alice's two possible signals.

3.6.1 An Expression for Alice's Ex-Interim Expected Payoff

To characterize PBE of the 3-stage D game, one major challenge is to construct Bob's off-equilibrium-path belief, that is, Bob's belief for any Alice's report that is not in the support of Alice's first-stage PBE strategy. One easy way to construct Bob's belief is to assume that any Alice's report r is always in the support of Alice's first-stage PBE strategy. In other words, we can construct Bob's belief for any Alice's report r as if his belief is always on the equilibrium path. Given this assumption, as long as the consistency condition is satisfied, Bob's belief for any r is uniquely determined.

In what follows, we derive an expression for Alice's ex-interim expected payoff for a given signal a_i and a particular first-stage report r (denoted $u_{a_i}(r)$, $i = 0, 1$) at any PBE of the 3-stage market game. When deriving $u_{a_i}(r)$, we assume that Alice's first-stage payoff satisfies the consistency condition, Alice and Bob know each other's strategies and beliefs, and mostly importantly Bob's belief for any Alice's report r is derived as if the belief is on the equilibrium path for any given r . That is, for any Alice's report r , we assume that the report r is always in the support of Alice's first-stage strategy, and we construct Bob's belief for r using Bayes' rule accordingly. The expression of $u_{a_i}(r)$ is given below. The complete derivation is included in Appendix A.2.1.

$$\begin{aligned}
u_{a_i}(r) = & P(1|a_i) \log \frac{r}{P(1)} + P(0|a_i) \log \frac{1-r}{1-P(1)} \\
& + \sum_{s_B} \left\{ P(1, s_B|a_i) \log \frac{P(1|a_i, s_B)}{x_{s_B}(r)} + P(0, s_B|a_i) \log \frac{P(0|a_i, s_B)}{1-x_{s_B}(r)} \right\} \quad (3.12)
\end{aligned}$$

where $x_{s_B}(r)$ is

$$x_{s_B}(r) = \frac{P(1, s_B|a_0)(P(1|a_1) - r) + P(1, s_B|a_1)(r - P(1|a_0))}{P(s_B|a_0)(P(1|a_1) - r) + P(s_B|a_1)(r - P(1|a_0))}$$

This expression is useful for our following discussion in several ways. First, given $u_{a_i}(r)$, we can easily calculate Alice's ex-ante expected payoff for using a particular strategy and use it to identify Alice's candidate PBE strategies by using Theorem 1. Second, to construct a PBE of the market game using a particular Alice's first-stage strategy, we must check whether $u_{a_i}(r)$ satisfies the requirements of a PBE for any report r in the support of Alice's first-stage strategy. For instance, if r_1 and r_2 are both in the support of Alice's first-stage strategy for Alice's signal a_i , then the PBE requirements specify that we must have $u_{a_i}(r_1) = u_{a_i}(r_2)$. Finally, to construct a PBE of the 3-stage D game, we must specify Bob's off equilibrium path belief. One easy way to construct Bob's belief is to treat every Alice's report r as if it is on the equilibrium path. Given this assumption and the consistency condition, Bob's belief for every possible Alice's report r is uniquely determined. Given that Bob's belief is constructed in this way, we can use $u_{a_i}(r)$ to check whether a given Alice's first-stage strategy and Bob's belief can form a PBE of the 3-stage D game.

3.6.2 Three Candidate PBE Strategies for Alice

We identify three candidate PBE strategies for Alice in the 3-stage D game. These three strategies are the truthful strategy, the delaying strategy, and a mixed strategy in which Alice makes a deterministic report r for one realized signal and she mixes between reporting r and reporting her true posterior probability estimate for the other realized signal. The proof of Theorem 5 is included in Appendix A.1.6.

Theorem 5. *If there exists a PBE of the 3-stage D game, then Alice must play one of the following three strategies at the PBE ²:*

- *the truthful strategy:* $\sigma_{a_i}(P(1|a_i)) = 1, i = 0, 1$

²Technically, Alice's PBE strategy could be of the form $\sigma_{a_i}(P(1|a_i)) = 1 - p, \sigma_{a_i}(r) = p, \sigma_{a_{1-i}}(P(1|a_{1-i})) = 1 - q, \sigma_{a_{1-i}}(r) = q$, for some $p, q \in [0, 1], r \in [\min_{a_i} P(1|a_i), \max_{a_i} P(1|a_i)]$. However, if there exists a PBE of a 3-stage D game where Alice plays this mixed strategy, then there also exists a truthful PBE for this game. So we include this strategy as a special case when the 3-stage D game has a truthful PBE.

- the delaying strategy: $\sigma_{a_i}(P(1)) = 1, i = 0, 1$

- the mixed strategy:

$$\sigma_{a_i}(P(1|a_i)) = 1 - p, \sigma_{a_i}(r) = p, \sigma_{a_{1-i}}(r) = 1, \quad (3.13)$$

where $p = \frac{P(a_{1-i})(r - P(1|a_{1-i}))}{P(a_i)(P(1|a_i) - r)}$ and $u_{a_i}(P(1|a_i)) = u_{a_i}(r)$ is satisfied

for some $r \in (\min_{a_i \in S_A} P(1|a_i), P(1)) \cup (P(1), \max_{a_i \in S_A} P(1|a_i))$, $i = 0, 1$.

3.6.3 A Sufficient Condition for the Truthful PBE

In this section, we show in Theorem 6 that a monotonicity condition is sufficient for the existence of a truthful PBE of the 3-stage D game. This monotonicity condition requires that, for a fixed $i = 0, 1$, Alice's ex-interim expected payoff $u_{a_i}(r)$ is monotonically decreasing as the value of r changes from $P(1|a_i)$ to $P(1|a_{1-i})$. The proof of Theorem 6 is included in Appendix A.1.7.

Intuitively, to construct a truthful equilibrium, we need to construct Bob's belief such that Alice's ex-interim expected payoff given Bob's belief is maximized when she reports truthfully (i.e. $r = P(1|a_i), \forall i = 0, 1$). We construct Bob's belief for every Alice's report r by assuming that r is always in the support of Alice's first-stage strategy. Then, a truthful equilibrium exists if and only if $u_{a_i}(r)$ is maximized when $r = P(1|a_i), \forall i = 0, 1$ for any $r \in [\max(P(1|a_i), P(1|a_{1-i})), \min(P(1|a_i), P(1|a_{1-i}))]$. The monotonicity condition described in Theorem 6 is simply a stronger condition, which ensures that Alice's ex-interim expected payoff is maximized at her truthful report given either of her signals.

Theorem 6. *If for any $i = 0, 1$, $u_{a_i}(r)$ is monotonically decreasing as the value of r changes from $P(1|a_i)$ to $P(1|a_{1-i})$, then there exists a PBE of the 3-stage D game where Alice's first-stage strategy is*

$$\sigma_{a_i}(P(1|a_i)) = 1, \forall i = 0, 1$$

and Bob's second-stage strategy is

$$x_{s_B}(r) = \frac{P(1, s_B|a_0)(P(1|a_1) - r) + P(1, s_B|a_1)(r - P(1|a_0))}{P(s_B|a_0)(P(1|a_1) - r) + P(s_B|a_1)(r - P(1|a_0))}, \forall s_B \in S_B \quad (3.14)$$

Next, we give an example of a D information structure satisfying the monotonicity condition above. The example was found through exhaustive search of the space of the D information structures with a reasonable discretization factor. Nevertheless, it was relatively easy to identify the example because it was easy to check whether $u_{a_i}(r)$ satisfies the monotonicity condition.

Example 1. Consider an instance of the 3-stage D game where the prior distribution $\mathcal{P}$ is given by Table 3.1. For this example, Bob only has two possible realized signals b_0 and b_1 . This prior distribution satisfies the monotonicity condition specified in Theorem 6. As r increases from $P(1|a_0)$ to $P(1|a_1)$, $u_{a_0}(r)$ decreases and $u_{a_1}(r)$ increases.

	$\omega = 1$			$\omega = 0$	
	a_0	a_1		a_0	a_1
b_0	0.15	0.2	b_0	0.2	0.05
b_1	0.05	0.05	b_1	0.25	0.05

Table 3.1: An example prior distribution. Each cell gives the value of $P(\omega, a_i, b_j)$ for the realized outcome ω , Alice's signal a_i and Bob's signal b_j .

3.7 Conclusion and Future Work

We analyze how the the participants' knowledge of one another's private information, also called the information structure, affects their strategic behavior when trading in a prediction market. We model the logarithmic market scoring rule prediction market as an extensive-form Bayesian game, and characterize perfect Bayesian equilibria of this market game for different information structures. When the participants' private information is unconditionally independent (I game), we show that there exists a unique family of PBE for the market game with a finite number of players and a finite number of stages. At any PBE in this

family, assuming that the players are ordered by their last stages of participation, each player does not participate in the game before the previous player's last stage of participation. In particular, there exists a PBE where every player waits until their last stage of participation to truthfully reveal their information, and this is arguably the worst outcome for information aggregation. An immediate future direction is to determine whether a PBE exists for the I game with a finite number of players but an infinite number of stages.

We also study a restricted version of the market game with 2 players and 3 stages where the players' private information is neither conditionally independent nor unconditionally independent (D game). Our result narrows down the possible PBE strategies to three simple strategies if a PBE exists. We conjecture that, there exists a PBE of this restricted D game where the first participant plays one of these three strategies. For future work, we are interested in proving the existence of the PBE of the D game for any information structure, characterizing sufficient and necessary conditions for each type of PBE to exist, and exploring whether the PBE results extend to the game with a finite or an infinite number of stages.

Regarding assumptions of our market model, an interesting future direction is to analyze the PBE of a different model of prediction markets where the participants endogenously choose when to trade in the market, instead of following a pre-specified participation sequence. This model better captures how participants trades in practice markets and may provide more insights on how information is aggregated in prediction markets in practice.

Chapter 4

Prediction Markets with Outside Incentives

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 [Hanson, 2007a] 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 [Berg et al., 2001, Wolfers and Zitzewitz, 2004, Forsythe et al., 1992, 1999, Debnath et al., 2003, Chen and Plott, 2002].

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 [Lacomb et al., 2007]. 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 work, we investigate information aggregation in predic-

tion 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.

4.1 Our Results

We study a Bayesian model of a logarithmic market scoring rule (LMSR) [Hanson, 2007a] 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 [Spence, 1973]. This mapping provides nice intuitions for the existence of this subset of separating PBE.

4.2 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 [2010b] 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 [2010b] 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. [2010a] and our work in Chapter 3 show that the equilibria and information revelation in such settings depend on the structure of the participants' private information. Ostrovsky [2011] and Iyer et al. [2010b] focus on studying information aggregation at any PBE of a prediction market instead of directly characterizing the equilibria. Ostrovsky [2011] 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. [2010b] extend the setting of Ostrovsky [2011] 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. [2009] 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 [2010] 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 [2011] and Chen et al. [2011] 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 [Othman and Sandholm, 2010, Chen and Kash, 2011, Chen et al., 2011] assume that experts do not have an inherent interest in the decision and they derive utility only from the scoring rule payment. Boutilier [2012] 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 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 work, 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 work 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. [2010a], us in Chapter 3, and Ostrovsky [2011].

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. [2004a] and by Rhode and Strumpf [2004] 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. [2004a] 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. [2007] 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. [2007] 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 [1998] 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 [2010] 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. [2010a] 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 [Spence, 1973]. 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 4.7.

4.3 The Market Game

4.3.1 The 2-Stage Market Game

In this work, we study the MSR market with two rational, risk-neutral participants Alice and Bob. 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 $\mathcal{P}$ denote a joint prior probability distribution over Ω , S_A and S_B . We assume $\mathcal{P}$ is common knowledge and omit it in our notation for brevity.

¹Our 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 endpoints of all such intervals. The number of possible rankings is exponential in n , making the analysis challenging.

We define $f_{s_A, \emptyset} = P(\omega = 1 | S_A = s_A)$ and $f_{\emptyset, s_B} = P(\omega = 1 | S_B = s_B)$ to represent the posterior probability for $\omega = 1$ given Alice's and Bob's private signal respectively. Similarly, $f_{s_A, s_B} = P(\omega = 1 | S_A = s_A, S_B = s_B)$ represents the posterior probability for $\omega = 1$ given both signals. We assume that Alice's H signal indicates a strictly higher probability for $\omega = 1$ than Alice's T signal, for any realized signal s_B for Bob, i.e. $f_{H, s_B} > f_{T, s_B}, \forall s_B \in \{H, T\}$. In addition, we assume that without knowing Bob's signal, Alice's signal alone also predicts a strictly higher probability for $\omega = 1$ with the H signal than with the T signal and Alice's signal alone can not predict ω 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 $\omega = 1$ as the event that the flu is widespread and $\omega = 0$ 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 market game has two stages. The sequence of participation is Alice and then Bob. Alice changes the market probability from r^0 to 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.

In addition to Alice and Bob's payoffs from within the market, Alice also 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 4.6.

4.3.2 Solution Concept and Players' Strategies

Our solution concept is the *perfect Bayesian equilibrium* (PBE) [Fudenberg and Tirole, 1991], 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.

Alice's first-stage strategy is a mapping $\sigma : \{H, T\} \rightarrow \Delta([0, 1])$, where $\Delta([0, 1])$ is the set of probability distributions over $[0, 1]$. 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 changes 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

²This 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 [Cho and Kreps, 1987] for an example.

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.

In our PBE analysis, we use the notions of *separating* and *pooling* PBE, similar to the solution concepts used by Spence [Spence, 1973]. 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 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 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.

4.4 Known Outside Incentive

In our basic model, it is certain and common knowledge that Alice has the outside payoff. If Alice reports truthfully in the market and Bob believes that she is being truthful, then Alice's outside payoff when she receives the H signal is larger than her outside payoff when she receives the T signal. 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. In an equilibrium of the market, Alice balances these two conflicting forces. 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 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.

4.4.1 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

$$\sigma_H(f_{H,\emptyset}) = 1, \sigma_T(f_{T,\emptyset}) = 1, \quad (4.1)$$

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

$$\sigma_H(p_1) = 1, \sigma_T(p_2) = 1. \quad (4.2)$$

for any $p_1, p_2 \in [0, 1]$ and $p_1 \neq p_2$.

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:

$$L(f_{s_A,\emptyset}, r_A) = f_{s_A,\emptyset} \log \frac{f_{s_A,\emptyset}}{r_A} + (1 - f_{s_A,\emptyset}) \log \frac{1 - f_{s_A,\emptyset}}{1 - r_A}, \quad (4.3)$$

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 B.1. Lemma 4 below gives a sufficient condition on the prior distribution and outside payoff function for the nonexistence of the truthful PBE.

Lemma 4. *For any prior distribution $\mathcal{P}$ and outside payoff function $Q(\cdot)$, if inequality (4.4) is satisfied, Alice's truthful strategy given by (4.1) is not part of any PBE of this game.*

$$L(f_{T,\emptyset}, f_{H,\emptyset}) < E_{S_B}[Q(f_{H,S_B}) - Q(f_{T,S_B}) \mid S_A = T] \quad (4.4)$$

Proof. We prove by contradiction. Suppose that inequality (4.4) 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,

$$\mu_{s_B, f_{H,\emptyset}}(H) = 1, \mu_{s_B, f_{T,\emptyset}}(T) = 1. \quad (4.5)$$

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 $E_{S_B}[Q(f_{H,S_B}) - Q(f_{T,S_B}) \mid S_A = T]$ (RHS of inequality (4.4)) and her expected loss in market scoring rule payoff is $L(f_{T,\emptyset}, f_{H,\emptyset})$ (LHS of inequality (4.4)). Because of (4.4), 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 (4.4) 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 (4.4).

In Appendix B.2, we present and discuss Example 2, which shows a prior distribution and an outside payoff function for which inequality (4.4) 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 4.4.3, we will present another pair of prior distribution and outside payoff function in Example 3 where a truthful

PBE also fails to exist.

4.4.2 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 $\omega = 1$. Her equilibrium strategy critically depends on the relative attractiveness of her expected market scoring rule payoff and her expected outside payoff, which depends on the prior distribution and the outside payoff function. In this section, we partition Alice's strategy space using some key values in order to facilitate our equilibrium analysis.

First, to illustrate the intuition, we define the key values in Alice's strategy space by partitioning Alice's strategy space in the following way, illustrated in Figure 4.1. For each signal s_A , the blue regions contain values for Alice's reports that are dominated by her truthful reports, regardless of Bob's strategy. The white regions contain values for Alice's reports that are not dominated by her truthful reports, for some strategy for Bob. This partition shows that, given a particular realized signal, it only makes sense for Alice to consider reporting values in the white regions. For Alice's signal s_A , we define Y_{s_A} and Y_{-s_A} to be the upper and lower bound values for the white regions respectively. These values are well defined and uniquely determined given the prior joint distribution and the outside payoff function.

Next, we formally define Y_{s_A} and Y_{-s_A} . Given a prior distribution $\mathcal{P}$ 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 (4.6) and Y_{-s_A} to be the unique value in $[0, f_{s_A, \emptyset}]$ satisfying equation (4.7):

$$L(f_{s_A, \emptyset}, Y_{s_A}) = E_{S_B}[Q(f_{H, S_B}) - Q(f_{T, S_B}) \mid s_A], \quad (4.6)$$

$$L(f_{s_A, \emptyset}, Y_{-s_A}) = E_{S_B}[Q(f_{H, S_B}) - Q(f_{T, S_B}) \mid s_A]. \quad (4.7)$$

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.




Figure 4.1: An illustration of Y_H and Y_T by partitioning Alice's strategy space. The blue regions contain Alice's reports that are dominated by truthful reports. The white regions contain Alice's reports that are not dominated by truthful reports. Y_H and Y_T are the upper bound values for the white regions.

Note that the RHS of equations (4.6) and (4.7) 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} always exist and are well defined — given any pair of prior distribution and outside payoff function, there always exist $Y_{s_A} \in [f_{s_A,\emptyset}, 1)$ and $Y_{-s_A} \in (0, f_{s_A,\emptyset}]$ such that equations (4.6) and (4.7) 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 in order to convince Bob that she has the H signal, after receiving the s_A signal respectively. The RHS of equations (4.6) and (4.7) 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. So Y_{s_A} and Y_{-s_A} are reports for Alice such that her loss in market scoring rule payoff is exactly equal to her maximum and minimum expected gain in outside payoff respectively. 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 truthfully, 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:*

$$Y_H \geq f_{H,\emptyset} \geq Y_{-H}, \quad (4.8)$$

$$Y_T \geq f_{T,\emptyset} \geq Y_{-T}, \quad (4.9)$$

$$Y_H \geq Y_{-T}. \quad (4.10)$$

Proof. (4.8) and (4.9) 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 B.3.

4.4.3 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.

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:

$$SE_1 : \begin{cases} \sigma_H(\max(Y_T, f_{H,\emptyset})) = 1, \sigma_T(f_{T,\emptyset}) = 1 \\ \text{When } Y_{-T} < Y_T, \mu_{s_B, r_A}(H) = \begin{cases} 1, & \text{if } r_A \in [Y_T, 1] \\ 0, & \text{if } r_A \in (Y_{-T}, Y_T) \end{cases} \\ \text{When } Y_{-T} = Y_T, \mu_{s_B, r_A}(H) = \begin{cases} 1, & \text{if } r_A \in [0, Y_{-T}] \\ 1, & \text{if } r_A \in (Y_T, 1] \\ 0, & \text{if } r_A = Y_T = Y_{-T} \\ 1, & \text{if } r_A \in [0, Y_{-T}) \end{cases} \end{cases} \quad (4.11)$$

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 7. *If $Y_H \geq Y_T$, SE_1 described in (4.11) 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. $\square$

Necessary Condition

In Theorem 7, 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 5 states that, at any separating PBE, after receiving the T signal, Alice must report $f_{T,\emptyset}$ with probability 1. Lemma 6 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 8 that $Y_H \geq Y_T$ is a necessary condition for a separating PBE to exist.

Lemma 5. *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 5 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 4.6 when extending our results to other MSR markets.

Lemma 6. *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 5, 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 8. *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 6, 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:

$$L(f_{H, \emptyset}, Y_{-H}) = L(f_{H, \emptyset}, Y_H) < L(f_{H, \emptyset}, Y_T) \leq L(f_{H, \emptyset}, Y_{-T}) \Rightarrow Y_{-H} > Y_{-T} \quad (4.12)$$

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

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 $\mathcal{P}$ 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 B.4, we describe Example 3 illustrating a scenario where $Y_H < Y_T$ is satisfied.

Proposition 4. *For any outside payoff function Q , 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 truthful 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 $Y_H > Y_T$. Note that this argument is independent of the outside payoff function because this argument compares Alice's strategy for both signals and the outside payoff function has identical effects on both signals. We formalize this argument below.

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

$$L(f_{H, \emptyset}, Y_H) = E_{S_B}[Q(f_{H, S_B}) - Q(f_{T, S_B})] = L(f_{T, \emptyset}, Y_T). \quad (4.13)$$

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

$$L(f_{T, \emptyset}, Y_H) > L(f_{H, \emptyset}, Y_H). \quad (4.14)$$

Using (4.13) and (4.14), we can derive that

$$L(f_{T, \emptyset}, Y_H) > L(f_{H, \emptyset}, Y_H) = L(f_{T, \emptyset}, Y_T). \quad (4.15)$$

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. [2010a] and us in Chapter 3 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.

4.4.4 Pure Strategy Separating PBE

In section 4.4.3, 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 strategy ³.

By Lemma 5, at any separating PBE, Alice's strategy must be of the following form:

$$\sigma_H^S(r_A) = 1, \sigma_T^S(f_{T,\emptyset}) = 1. \quad (4.16)$$

for some $r_A \in [0, 1]$. In Lemma 7, we further narrow down the possible values of r_A in Alice's strategy at any separating PBE.

Lemma 7. *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 6, Alice does not report any $r_A \in (Y_{-T}, Y_T)$ after receiving the H signal. $\square$

Lemma 7 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. Theorems 9 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

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

either $Y_{-H} > Y_{-T}$ or $Y_{-H} \leq Y_{-T}$. If $Y_{-H} \leq Y_{-T}$, we show in Theorem 10 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 B.5 and B.6 respectively.

Theorem 9. *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 10. *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 $\sigma_H^S(r_A) = 1, \sigma_T^S(f_{T,\emptyset}) = 1$.*

4.4.5 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:

$$PE_1(k) : \begin{cases} \sigma_H^P(k) = 1, \sigma_T^P(k) = \gamma(k), \sigma_T^P(f_{T,\emptyset}) = 1 - \gamma(k) \\ \mu_{s_B, r_A}^P(H) = \begin{cases} g(\gamma(k), s_B), & \text{if } r_A = k \\ 0, & \text{if } r_A \in [0, k) \cup (k, 1] \end{cases} \end{cases} \quad (4.17)$$

where

$$g(\gamma(k), s_B) = \frac{P(S_A = H|s_B)}{P(S_A = H|s_B) + P(S_A = T|s_B)\gamma(k)}, \quad (4.18)$$

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

$$L(f_{T,\emptyset}, k) \leq E_{S_B}[Q(g(\gamma(k), S_B)f_{H,S_B} + (1 - g(\gamma(k), S_B))f_{T,S_B}) - Q(f_{T,S_B}) \mid S_A = T] \quad (4.19)$$

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 (4.19), then $\gamma(k) = 1$. Otherwise, $\gamma(k)$ is set to a value such that the LHS and RHS of equation (4.19) are equal.

It is easy to see that $\gamma(k)$ is well defined for every $k \in (\max(Y_{-H}, Y_{-T}), \min(Y_H, Y_T)) \setminus \{f_{T,\emptyset}\}$. The RHS of inequality (4.19) is strictly monotonically decreasing in $\gamma(k)$. When $\gamma(k) = 0$, the RHS equals $L(f_{T,\emptyset}, Y_T)$ and $L(f_{T,\emptyset}, Y_{-T})$. Because $Y_{-T} < k < Y_T$, we know that $\gamma(k) > 0$.

By (4.17), 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 11, we show that $PE_1(k)$ is a pooling PBE if the following inequality is satisfied:

$$L(f_{H,\emptyset}, k) \leq E_{S_B}[Q(g(\gamma(k), S_B)f_{H,S_B} + (1 - g(\gamma(k), S_B))f_{T,S_B}) - Q(f_{T,S_B}) \mid S_A = H]. \quad (4.20)$$

Inequality (4.20) ensures that when Alice receives the H signal, she is better off reporting k rather than reporting $f_{H,\emptyset}$ given Bob's belief in $PE_1(k)$. When $k = f_{H,\emptyset}$, inequality (4.20) is automatically satisfied because the LHS of inequality (4.20) is 0 and the RHS of inequality (4.20) is positive. However, for other values of k , whether inequality (4.20) is satisfied depends on the prior distribution and the outside payoff function. This means that, if

$f_{H,\emptyset} < Y_T$, which ensures the interval $(\max(Y_{-H}, Y_{-T}), \min(Y_H, Y_T))$ is nonempty and contains $f_{H,\emptyset}$, then there always exists a pooling PBE of our game where Alice reports $f_{H,\emptyset}$ with probability 1 after receiving the H signal.

Theorem 11. *If $(\max(Y_{-H}, Y_{-T}), \min(Y_H, Y_T))$ is nonempty, for any $k \in (\max(Y_{-H}, Y_{-T}), \min(Y_H, Y_T)) \setminus \{f_{T,\emptyset}\}$, $PE_1(k)$ is a pooling PBE of our game if inequality (4.20) is satisfied.*

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

When Alice receives the H signal and $k \neq f_{H,\emptyset}$, for $r_A \in [0, 1] \setminus \{k\}$, it is optimal for Alice to report $f_{H,\emptyset}$ since her outside payoff is constant and her market scoring rule payoff is maximized. By inequality (4.20), Alice weakly prefers reporting k than reporting $f_{H,\emptyset}$. 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 $k = f_{H,\emptyset}$, 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 $f_{T,\emptyset}$. So the support of Alice's equilibrium strategy after receiving the T signal includes at most $f_{T,\emptyset}$ and k . By the definition of $\gamma(k)$, either Alice is indifferent between reporting $f_{T,\emptyset}$ 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 $f_{T,\emptyset}$ 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:

$$BE_1(k) : \begin{cases} \sigma_H^B(k) = 1, \sigma_T^B(k) = 1 \\ \mu_{s_B, r_A}^B(H) = \begin{cases} P(S_A = H|s_B), & \text{if } r_A = k \\ 0, & \text{if } r_A \in [0, k) \cup (k, 1] \end{cases} \end{cases} \quad (4.21)$$

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.

4.5 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 (4.11)) and SE_2 (defined in equation (4.27)) 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.

4.5.1 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 [Mas-Colell et al., 1995], 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 2. [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.*

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 2 we can put restrictions on Bob's belief at any PBE.

⁴Situations 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.

Lemma 8. *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 2, 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 B.7.

Sketch. By Theorem 9, 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 10, 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

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

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 (4.11) 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 $\sigma_T(f_{T,\emptyset}) = 1$.*

Proof. Since $f_{H,\emptyset} \in (Y_T, Y_H]$, then by Lemma 8, 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 B.8.

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 $L(f_{H,\emptyset}, r) = L(f_{H,\emptyset}, Y_T)$.*

4.5.2 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 9 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 12.

Lemma 9. *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 $P(S_A = s_A, S_B = s_B)$.

$$\begin{aligned}
u_A^G &= \sum_{v \in V} (P(H, H)\sigma_H(v) + P(T, H)\sigma_T(v)) \\
&\quad Q\left(\frac{P(H, H)\sigma_H(v)}{P(H, H)\sigma_H(v) + P(T, H)\sigma_T(v)}f_{HH} + \frac{P(T, H)\sigma_T(v)}{P(H, H)\sigma_H(v) + P(T, H)\sigma_T(v)}f_{TH}\right) \\
&\quad + (P(H, T)\sigma_H(v) + P(T, T)\sigma_T(v)) \\
&\quad Q\left(\frac{P(H, T)\sigma_H(v)}{P(H, T)\sigma_H(v) + P(T, T)\sigma_T(v)}f_{HT} + \frac{P(T, T)\sigma_T(v)}{P(H, T)\sigma_H(v) + P(T, T)\sigma_T(v)}f_{TT}\right) \\
&\leq \sum_{v \in V} (P(H, H)\sigma_H(v)Q(f_{HH}) + P(H, T)\sigma_H(v)Q(f_{HT}) \\
&\quad + P(T, H)\sigma_T(v)Q(f_{TH}) + P(T, T)\sigma_T(v)Q(f_{TT})) \\
&= P(H, H)Q(f_{HH}) + P(H, T)Q(f_{HT}) + P(T, H)Q(f_{TH}) + P(T, T)Q(f_{TT}) \\
&= u_A^S
\end{aligned} \tag{4.22}$$

where inequality (4.22) was derived by applying the convexity of $Q(\cdot)$. $\square$

Theorem 12. *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 9, any separating PBE maximizes Alice's expected outside payoff if $Q(\cdot)$ is convex. Therefore, any separating PBE maximizes the social welfare. $\square$

4.5.3 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 13 that the separating PBE SE_1 maximizes Alice's expected payoff among all separating PBE of this game. This is easy to see when $f_{H,\emptyset} \geq Y_T$ since the separating PBE SE_1 is also the truthful PBE of this game. Otherwise, if $f_{H,\emptyset} < Y_T$, Y_T is the minimum deviation from $f_{H,\emptyset}$ that Alice can report in order to convince Bob that she has the H signal.

Theorem 13. *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 (4.11).*

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

By Lemma 5, 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.

When $f_{H,\emptyset} \geq Y_T$, according to Theorem 7, Alice reports $f_{H,\emptyset}$ after receiving the H signal and this is the maximum expected payoff she could get after receiving the H signal.

When $f_{H,\emptyset} < Y_T$, 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

$$r_A < Y_{-T} \Rightarrow L(f_{H,\emptyset}, Y_T) \leq L(f_{H,\emptyset}, Y_{-T}) < L(f_{H,\emptyset}, r_A). \quad (4.23)$$

where the inequality $L(f_{H,\emptyset}, Y_T) \leq L(f_{H,\emptyset}, Y_{-T})$ is due to Proposition 3. Therefore, when $f_{H,\emptyset} \leq Y_T$, 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 13 suggests that SE_1 is likely a desirable PBE of our game. In Theorems 14

and 15, we compare Alice's expected payoff in SE_1 with her expected payoff in the pooling PBE of this game. Again, when $f_{H,\emptyset} \geq Y_T$, 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 $f_{H,\emptyset} < Y_T$, the relationship is less clear. In Theorem 15, we show that, if $k \in (\max(Y_{-H}, Y_{-T}), Y_T) \setminus \{f_{T,\emptyset}\}$ satisfies inequality (4.24), then Alice's expected payoff in SE_1 is greater than her expected payoff in the pooling PBE $PE_1(k)$.

Theorem 14. *If $Q(\cdot)$ is convex, $Y_H \geq Y_T$ and $f_{H,\emptyset} \geq Y_T$, Alice's expected payoff is maximized in the pure strategy separating PBE SE_1 among all PBE of our game.*

Proof. By Lemma 9, any separating PBE maximizes Alice's expected outside payoff if $Q(\cdot)$ is convex. When $f_{H,\emptyset} \geq Y_T$, SE_1 is the truthful PBE and strictly maximizes Alice's expected market scoring rule payoff. $\square$

Theorem 15. *If $Q(\cdot)$ is convex, $Y_H \geq Y_T$, and $f_{H,\emptyset} < Y_T$, Alice's expected payoff in the pure strategy separating PBE SE_1 is greater than her expected payoff in $PE_1(k)$ for any $k \in (\max(Y_{-H}, Y_{-T}), Y_T) \setminus \{f_{T,\emptyset}\}$ if k satisfies inequality (4.24) below.*

$$P(s_A = H)L(f_{H,\emptyset}, Y_T) \leq P(s_A = H)L(f_{H,\emptyset}, k) + P(s_A = T)\gamma(k)L(f_{T,\emptyset}, k). \quad (4.24)$$

Proof. By Lemma 9, 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 7 is less by $P(s_A = H)L(f_{H,\emptyset}, Y_T)$, and Alice's payoff in $PE_1(k)$ is less by $P(s_A = H)L(f_{H,\emptyset}, k) + P(s_A = T)\gamma(k)L(f_{T,\emptyset}, k)$. 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 $P(s_A = H)L(f_{H,\emptyset}, Y_T) \leq P(s_A = H)L(f_{H,\emptyset}, k) + P(s_A = T)\gamma(k)L(f_{T,\emptyset}, k)$, which is stated in inequality (4.24). $\square$

Note that inequality (4.24) is automatically satisfied for any $k \leq r$ where r is the unique

value in $[0, f_{H,\emptyset}]$ satisfying $L(f_{H,\emptyset}, r) = L(f_{H,\emptyset}, Y_T)$ since

$$P(s_A = H)L(f_{H,\emptyset}, Y_T) = P(s_A = H)L(f_{H,\emptyset}, r) \leq P(s_A = H)L(f_{H,\emptyset}, k) \quad (4.25)$$

$$< P(s_A = H)L(f_{H,\emptyset}, k) + P(s_A = T)\gamma(k)L(f_{T,\emptyset}, k). \quad (4.26)$$

However, for $k \geq r$, whether k satisfies inequality (4.24) depends on the prior distribution and the outside payoff function.

4.5.4 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 (4.27), which is the same as $SE_2(Y_H)$ defined in equation (B.24) in Appendix B.5. We state SE_2 below for convenience. The proof of Theorem 16 is included in Appendix B.9.

Theorem 16. *Among all pure strategy separating PBE of our game, Bob's expected payoff is maximized in the following pure strategy separating PBE SE_2 .*

$$SE_2 : \begin{cases} \sigma_H(Y_H) = 1, \sigma_T(f_{T,\emptyset}) = 1 \\ \mu_{s_B, r_A}(H) = \begin{cases} 1, & \text{if } r_A \in [Y_H, 1] \\ 0, & \text{if } r_A \in [0, Y_H) \end{cases} \end{cases} \quad (4.27)$$

4.6 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.

4.6.1 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 4.3.1, let $s(\omega, r)$ denote a strictly proper scoring rule of a binary random variable Ω where ω is the realization of Ω and r is the reported probability of $\omega = 1$. Then the loss function $L_s(f_{s_A, \emptyset}, r_A)$ for the strictly proper scoring rule $s(\omega, r)$ can be defined as follows:

$$L_s(f_{s_A, \emptyset}, r_A) = f_{s_A, \emptyset} \{s(1, f_{s_A, \emptyset}) - s(1, r_A)\} + (1 - f_{s_A, \emptyset}) \{s(0, f_{s_A, \emptyset}) - s(0, r_A)\} \quad (4.28)$$

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

$$L_s(f_{T, \emptyset}, 1) \geq E_{S_B}[Q(f_{H, S_B}) - Q(f_{T, S_B}) \mid S_A = T] \quad (4.29)$$

$$L_s(f_{H, \emptyset}, \max(f_{H, \emptyset}, Y_T)) \leq E_{S_B}[Q(f_{H, S_B}) - Q(f_{T, S_B}) \mid S_A = H] \quad (4.30)$$

If inequality (4.29) is satisfied, we know that Y_T is well defined. Then, if inequality (4.30) 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 (4.29)

is violated, then the quantity Y_T is not well defined, so inequality (4.30) is not a well defined statement as well.

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

$$L_s(f_{T,\emptyset}, 0) \geq E_{S_B}[Q(f_{H,S_B}) - Q(f_{T,S_B}) \mid S_A = T] \quad (4.31)$$

$$L_s(f_{H,\emptyset}, Y_{-T}) \leq E_{S_B}[Q(f_{H,S_B}) - Q(f_{T,S_B}) \mid S_A = H] \quad (4.32)$$

Again, inequality (4.31) ensures that Y_{-T} is well-defined. If inequality (4.32) 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 17 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 17. *A separating PBE of our game exists if and only if at least one of the pair of inequalities (4.29) and (4.30) and the pair of inequalities (4.31) and (4.32) is satisfied.*

We include the complete proof in Appendix B.10.

4.6.2 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 18. *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. $\square$

4.7 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 [Spence, 1973]. 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 [Mas-Colell et al., 1995]. 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 $u(\omega, e \mid \theta)$, is equal to his wage less the cost of getting education, i.e. $u(\omega, e \mid \theta) = \omega - c(e, \theta)$.

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 $[\tilde{e}, e_1]$ 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 $\tilde{e}$ 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 4.1.

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

Table 4.1: 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
$\tilde{e}$, 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

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 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 $[\tilde{e}, e_1]$, 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 4.1. 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.

4.8 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 high level 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. Another interesting future direction is to consider outside payoff functions with other structures, such as threshold functions or non-monotone functions.

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.

Chapter 5

An Experimental Evaluation of a Peer Prediction Mechanism

Businesses and organizations often face the challenge of gathering accurate and informative feedback or opinion from multiple individuals. Notably, community-based websites such as Yelp, IMDb, Amazon, TripAdvisor, and Angie’s List that review products and services are largely dependent on voluntary feedback contributed by their users. The proliferation of online labor markets has also created an opportunity for outsourcing simple tasks, such as classifying images and identifying offensive content on the web, to a readily available online workforce.

In all of these settings, there are some significant barriers in getting participants to honestly reveal their information. First, it is often costly to formulate and share an honest opinion—for example, when evaluating qualities of books or restaurants—whereas uninformative contributions require little or no effort. More importantly, it may be difficult or impossible to verify individual contributions against an observable ground truth, because either the information is inherently subjective or it is too costly to be verified.

These difficulties led to the development of *peer prediction* mechanisms, pioneered by Miller et al. [2005]. To incentivize truthful reports, these mechanisms leverage the stochastic cor-

relation of participants’ information and reward each participant by comparing the participant’s report with those of his peers. Existing theory on peer prediction focuses on designing such monetary rewards to induce a *truthful Bayesian Nash equilibrium* among all participants—it is in a participant’s best interest to report his information truthfully if he believes all other participants will also be truthful.

Despite the existence of the truthful equilibrium, peer prediction theory provides little assurance that participants will adhere to it in practice. This is because peer prediction mechanisms also induce *uninformative* equilibria where participants can coordinate to make a set of reports that are independent of their information. Moreover, such uninformative equilibria are unavoidable in general [Jurca and Faltings, 2009, Waggoner and Chen, 2013]. Although Miller et al. [2005] argue that the truthful equilibrium is likely focal due to limited communication or ethical preferences of participants, there is generally little theoretical or empirical support for this conjecture.

In this work, we aim to understand how participants will behave towards the peer prediction mechanisms in the presence of multiple equilibria. Specifically, we address the following question:

Will the participants play one of the multiple equilibria of the peer prediction mechanisms? If they do, which equilibrium will they choose and why?

To tackle this question, we tested participants’ behavior towards the Jurca and Faltings [2009] (JF) mechanism by engaging them in a multiplayer, real-time and repeated game through a controlled online experiment on Amazon’s Mechanical Turk (MTurk) [Mason and Suri, 2012]. We allow the participants to repeatedly interact with the JF mechanism, revealing their behavioral dynamics over time. Our work is the first empirical evaluation of participants’ behavior towards a peer prediction mechanism in terms of convergence to game-theoretic equilibria.

In our particular experimental setting, we show that the truthful equilibrium is not focal and that participants clearly favor the uninformative equilibria when paid by peer prediction.

In contrast, a majority of the participants are consistently truthful when peer prediction is not used. Specifically, for payment rules where the uninformative equilibria have higher payoffs, a majority of the players coordinated on an uninformative equilibrium. Moreover, for payment rules where the symmetric uninformative equilibria do not exist or have less payoff than the truthful equilibrium, we successfully deterred the players from choosing the uninformative equilibria but still did not motivate truthful reports. Hence, our results suggest that adopting peer prediction mechanisms may be harmful in scenarios where the cost of being truthful is similar to that of acting strategically.

In answering our question about equilibrium play, we needed to infer the players’ intended strategies, which are not directly observed in an experiment. This is a common challenge in experimental studies of game-theoretic mechanisms. Researchers typically try to detect equilibrium play by comparing the frequencies of players’ actions with equilibrium or other strategies. In this work, we use the novel approach of a hidden Markov model (HMM) to capture players’ strategies as latent variables. The HMM is ideal for game-theoretic modeling because it allows us to infer strategies and equilibrium play from observed actions, while allowing the strategies to evolve over time. Our analytical results show that there is great potential in using probabilistic latent variable models to describe game-theoretic and other experimental data.

5.1 Related Work

There have been considerable theoretical developments on peer prediction mechanisms. The earlier mechanisms rely on a common prior assumption — the information structure, i.e. the prior joint distribution of the participant’s signals and the event outcome, must be common knowledge among the participants and known by the designer [Miller et al., 2005, Jurca and Faltings, 2009, Zhang and Chen, 2014]. Miller et al. [2005] proposed the first peer prediction method, which rewards each participant for whether his report is predictive of another

participant’s report using a proper scoring rule. While the Miller et al. [2005] mechanism induces the truthful equilibrium, it also has uninformative equilibria where all participants make the same report and reveal no information. Jurca and Faltings [2009] generalized and improved the Miller et al. [2005] mechanism to eliminate or hamper the undesirable symmetric uninformative equilibria by making each participant’s payment depend on multiple other reports instead of one other report.

The common prior assumption is quite stringent, particularly since the mechanism needs to know the common prior in order to determine the payments. Several subsequent mechanisms relax this assumption — they either do not require this assumption at all [Witkowski and Parkes, 2012b], or do not require the mechanism to know the common prior [Prelec, 2004, Witkowski and Parkes, 2012a, Radanovic and Faltings, 2013]. In particular, the robust Bayesian truth serum (RBTS) [Witkowski and Parkes, 2012a] improved upon the Bayesian truth serum (BTS) [Prelec, 2004] by being incentive compatible for small populations, and Radanovic and Faltings [2013] devised a mechanism similar to RBTS [Witkowski and Parkes, 2012a] for non-binary signals.

Regarding the elicibility of private information without observable ground truth, Waggoner and Chen [2013] proved the impossibility result that any peer prediction mechanism has at least one uninformative equilibrium where the participants’ reports are independent of their private information. Intuitively, an uninformative equilibrium can be reached if the participants play the same game as if they do not possess any private information. Moreover, Radanovic and Faltings [2013] gave several impossibility results regarding the design of peer prediction mechanisms with different assumptions on the common prior distribution, focusing on the case when the participants have conditionally independent and identically distributed signals. Zhang and Chen [2014] proved that the information structures satisfying stochastic relevance is the maximal set of information structures that are truthfully elicitable. Moreover, they generalized the Miller et al. [2005] mechanism for the maximal truthfully elicitable set of information structures when the designer knows the information

structure, and they proposed a sequential mechanism for a slightly smaller set of information structures when the designer does not know the information structure.

In recent work, Witkowski et al. [2013] and Dasgupta and Ghosh [2013] assume that a participant can invest costly effort, which improves the quality of his report, and they proposed mechanisms to incentivize both truthfulness and high efforts. They consider binary information elicitation tasks and assume that an agent could invest costly effort to achieve a certain quality or a probability of identifying the ground truth. By using negative payments to penalize disagreement, Witkowski et al. [2013] designed an output-agreement mechanism such that participants with quality above a threshold choose to participate and invest effort while those below do not participate. Dasgupta and Ghosh [2013] devised a mechanism where exerting maximum effort and truthful reporting is a Nash equilibrium with maximum payoffs to all participants. Their key contribution is a technique for penalizing low-effort agreement leading to the uninformative equilibria by using the presence of multiple tasks.

Very few experimental work has adopted peer prediction mechanisms. Among them, only the work by John et al. [2012] aim to evaluate the incentives of a peer prediction mechanism, specifically the Bayesian truth serum [Prelec, 2004]. This motivated us to experimentally evaluate the incentives of other peer prediction mechanisms. John et al. [2012] asked psychologists to report their engagements in questionable research practices in an anonymous survey study, using BTS to incentivize truth telling. They showed that the self admission rate was higher in the BTS group than in the control group.

Several experiments adopted but did not evaluate peer prediction mechanisms [Prelec and Seung, 2006, Shaw et al., 2011, Gao et al., 2012]. Prelec and Seung [2006] demonstrated that BTS can be used to infer the ground truth even if most participants' subjective judgements are wrong. Shaw et al. [2011] used the description of BTS as the contextual manipulation for one financial incentive tested in their online experiment, but did not pay the workers according to the mechanism. Gao et al. [2012] used the Witkowski and Parkes [2012b] mechanism to score judges on their evaluations of the quality of short tourism ads.

5.2 The Jurca and Faltings Mechanism

We describe the JF mechanism [Jurca and Faltings, 2009], which include the first peer prediction mechanism [Miller et al., 2005] as a special case.

We are interested in an item for which some people have opinions or subjective information, given by their private signals, and we aim to design monetary payments so that participants will truthfully reveal their signals about the item. The item has a type $\omega \in \Omega$; we focus on the binary case with $|\Omega| = 2$. Among $N \geq 3$ participants providing their opinions, participant i 's private opinion is a binary signal $s \in S, S = \{s_1, s_2\}$ that only he observes.

The JF mechanism assumes that the relationship between the participants' signals and the type of the item is common knowledge. This assumption allows inference about the likelihood of one participant's signal given another participant's signal. The participants' signals are assumed to be conditionally independent given the type of the item. Specifically, the type of the item is drawn by nature according to a probability distribution $P(\omega), \omega \in \Omega$ where $\sum_{\omega \in \Omega} P(\omega) = 1$, and each participant's signal s is independently drawn according to the conditional probability distribution $P(s | \omega), s \in S, \omega \in \Omega$. Moreover, the mechanism assumes that the prior $P(\omega)$ and the conditional distribution $P(s | \omega)$ are common knowledge for all participants and the mechanism.

With their realized signals, each participant makes a report $r \in S$ to the mechanism. The mechanism pays each participant according to the participant's report r and n_f reports of his peers ($1 \leq n_f \leq N - 1$), denoted the participant's *reference reports*. For example, participant i 's reference report could be one other random participant's report ($n_f = 1$) or all other participants' reports ($n_f = N - 1$). When $n_f = 1$, we use r_f to denote the single reference report.

The payment rule of the JF mechanism is designed to induce a truthful Bayesian Nash equilibrium among rational and risk-neutral participants: a participant maximizes his expected payment by truthfully reporting his private signal if he believes that all other par-

ticipants are reporting their signals truthfully. A payment rule that supports the truthful equilibrium must satisfy the following linear constraints:

$$\sum_{0 \leq m \leq n_f} P(m \mid s_i) (u(s_i, m) - u(s_{3-i}, m)) \geq 0, \forall i = 1, 2$$

where $u(r, m)$ ($r \in S$, $m \in [1, n_f]$) denotes a participant's payment if he reports r and m out of his n_f reference reports are s_1 and $P(m \mid s)$ denotes the probability that given a signal of $s \in S$, m out of n_f signals are s_1 . We use this linear program with additional constraints to derive payment rules in our experiment.

Although the JF mechanism supports the truthful equilibrium by design, it always induces other *uninformative* pure strategy equilibria. Furthermore, the uninformative equilibria may yield higher payoffs than the truthful equilibrium, making it questionable whether participants will choose the truthful equilibrium in practice. For example, for the payment rule in Table 5.1a, for every $s \in S$, it is an equilibrium for all participants to report s regardless of their signals. These equilibria yield higher payoff (1.20 or 1.50) than the truthful equilibrium (0.91). In general, when a participant's payment depends only on one reference report, Jurca and Faltings [2009] proved that at least one of the symmetric pure strategy coordinating equilibria must yield higher payoff than the truthful equilibrium.

	$r = s_1$	$r = s_2$
$r_f = s_1$	1.50	0.30
$r_f = s_2$	0.10	1.20

(a) Example 1.

	$r = s_1$	$r = s_2$
$m = 0$	0.90	0.80
$m = 1$	0.10	1.50
$m = 2$	1.50	0.10
$m = 3$	0.80	0.90

(b) Example 2.

	$r = s_1$	$r = s_2$
$m = 0$	0.10	0.15
$m = 1$	0.10	0.90
$m = 2$	1.50	0.15
$m = 3$	0.15	0.10

(c) Example 3.

Table 5.1: Payment rule examples. In (a), each cell gives a player's payoff if he reports r and his reference report is r_f . In (b) and (c), each cell gives a player's payoff if he reports r and m out of the n_f reference reports are s_1).

Fundamentally, the existence of such uninformative equilibria is unavoidable as participants' reports are only compared with one another and not with the ground truth. Although we cannot make truthful reporting the unique pure strategy equilibrium, it is possible to mod-

ify the payment rule to remove the symmetric pure strategy coordinating equilibria entirely or make them less desirable than the truthful equilibrium. Either can be achieved if we have at least 4 participants and a participant’s payment depends on all other participants’ reports [Jurca and Faltings, 2009].

Table 5.1b shows a payment rule without any symmetric coordinating equilibrium. If all participants make the same report, then any participant can improve his payoff from 0.80 to 0.90 by reporting the other signal. However, this payment rule has asymmetric coordinating equilibria where 3 participants always report s_1 and the remaining participant always reports s_2 (or vice versa). These asymmetric coordinating equilibria also reveal no information but may be harder to reach if the participants cannot communicate with one another.

Alternatively, we could make the symmetric coordinating equilibria less desirable than the truthful equilibrium (Table 5.1c). This is possible because a participant’s payment is maximized when his report agrees with the majority but not all of his reference reports. At either symmetric coordinating equilibrium, every participant gets 0.15, which is much less than the expected payoff of 0.50 at the truthful equilibrium.

5.3 Experiment Design and Setup

Among all peer prediction mechanisms, we chose to test the binary version of the JF mechanism for the following reasons. Compared to the Miller et al. [2005] mechanism, Jurca and Faltings [2009] generalized and improved the Miller et al. [2005] mechanism to eliminate or hamper the undesirable symmetric uninformative equilibria. One of our main goals is to test whether eliminating or hampering the uninformative equilibria will promote truthful behavior. Moreover, compared to the later peer prediction mechanisms, it is relatively easier for the participants to understand and to reason about the payment rule of the Jurca and Faltings [2009] mechanism. First, the Jurca and Faltings [2009] mechanism only requires the participant to choose a value in a small discrete report space. Most later peer prediction

mechanisms require participants to report probabilities, so they need to choose a value in a continuous real-valued report space. It is psychologically easier for the participants to reason about discrete rather than continuous values. It is known that that people have systematic bias when reasoning about probabilities. Also, the Jurca and Faltings [2009] mechanism does not use any complicated mathematical formula whereas the later peer prediction mechanisms use complicated mathematical formulae such as proper scoring rules. As a result, it is much easier to make the payment rule of the Jurca and Faltings [2009] mechanism accessible to the lay participant compared to the later peer prediction mechanisms.

We evaluated the JF mechanism in a repeated setting, although it was defined and analyzed as a one-shot mechanism. In practice, it is natural to expect participants to learn and adapt to the mechanism by providing their opinions for different products and services. Hence, a repeated setting best captures the intended uses of peer prediction and the participants' learning dynamics. Such studies of one-shot games in a repeated setting are typical in experimental economics [Shachat et al., 2012].

We conduct our experiment on MTurk. Peer prediction mechanisms are naturally suited for online experiments, as their intended application is for web-based or impersonal settings. Online participants enjoy greater anonymity and are not subject to the social norms typical of real-world interaction. Additionally, online labor markets, such as MTurk, provide immediate access to a large and diverse subject pool, which is not readily available through alternative means. We recruited over 3000 unique participants in our experiment. Recent studies have shown that online experiments are not only viable, but can be advantageous for large-scale studies of behavior [Horton et al., 2011, Mason and Suri, 2012, Rand, 2012].

Trick or treat story In our experiment, we described the mechanism using a simple and fun story about trick or treating on Halloween night:

A group of kids are trick or treating on Halloween night. There are two types of houses, A and B, giving out two types of candies, M&M's and gummy bears, in different proportions. The kids randomly choose a house to go trick or treating. The chosen house may be one of the two types with equal chance but the kids don't

know which type of house was chosen. Each kid secretly receives one randomly selected candy from the chosen house. A clown shows up and asks each kid to tell him the type of candy received, promising a payment in return. Each kid may claim to have either type of candy to the clown. The clown collects reports from all the kids, and then rewards each kid based on the kid’s claim and the other kids’ claims according to a payment rule.

This story can be easily mapped to the setting of the JF mechanism. Each player is a kid, the house is the type of the item, the candies are the signals and reports, and the clown is the mechanism.

This story serves several purposes. First, it makes the mechanism accessible to participants who may be unfamiliar with economic theory. Second, it satisfies the mechanism’s common knowledge assumptions, which are required for their theoretical properties to hold. The story also highlights the conditional independence of the signals given the type of the item by emphasizing that the proportions of candies at each house remains the same regardless of the candies given out.

The story also explains that misreporting one’s private information is perfectly acceptable. Although this is a common concept in game theory, it can be unintuitive for MTurk workers since they often associate misbehavior with rejection or punishment. To counter this bias, we stated that each player’s candy is obtained in secret and is not observed by anybody else. We used the clown as a neutral character to represent the mechanism. The action of “making a claim” is a neutral phrasing in lieu of words such as “lying” or “cheating” with negative connotations. We also emphasized that each player can claim to have either type of candy, and that the clown cannot verify whether the player’s claim matches the player’s candy.

Setup Our experiment engaged the participants in a real-time, multiplayer, and repeated game through MTurk. A fixed number of players are matched together and play a game repeatedly for a fixed number of rounds. We use TurkServer [Mao et al., 2012], a framework and API for real-time interaction between experiment participants, in our experiment.

We offer a \$1.00 base payment for completing the task. Each player also receives a bonus

payment equal to his average reward in the game (ranging from \$0.10 to \$1.50). We chose a large bonus relative to the base payment to motivate workers to focus on their performance in the game.

We controlled our subject pool in several ways. Each worker may participate only once in the experiment, so that no worker has prior experience with the game. We also restricted our tasks to US workers, for two reasons. Having US workers minimizes the likelihood of connection issues since we require synchronous connections to a US-based server. Second, controlling for geography avoids unexpected behavior if people from other regions have different behavioral norms or a language barrier in understanding the instructions. Additionally, we used common qualifications ($> 95\%$ HITs approved, > 100 HITs completed) to ensure that workers are familiar with the MTurk system.

Task Users progress through the task in several sections. The first page describes some general information and requires consent, followed by an 11-page tutorial consisting mainly of pictures. The tutorial describes the trick or treat story and the game interface. After the tutorial, the participant must pass a short quiz testing their understanding of the task. Each participant has 3 attempts to pass the quiz with a score of at least 80%. If they fail all 3 attempts, they are permanently blocked from our experiment. After passing the quiz, participants wait in a virtual lobby for enough players to start a new game. When enough players arrive, a 'READY' button appears for each player, starting a new game when enough players press this button; this ensures that all players are paying attention when the experiment starts. We explain the game interface in the next section. After the game, the participants are asked to describe the strategies and reasoning they used in a short exit survey.

Interface The game interface (Figure 5.1) describes some general information such as the number of rounds and the number of other players at the top, the steps of the current round on the left, and the history of game play on the right. Players cannot communicate with one another during the task, other than viewing reports from previous rounds.

You are playing **20 rounds** of the game with **2 other player(s)**. Your bonus (= average reward) so far: **\$0.63**.




Figure 5.1: The Game Interface

When designing the table showing the history of play, we carefully considered how to show the optimal amount of information. With an abundance of information, a participant may become distracted or confused, and pay less attention to the task. However, with insufficient information, a participant may not be able to learn or improve his strategy by observing other participants' actions. After initial trials of our experiment, we ultimately chose to display the other participants' claims as an aggregate summary, which is a concise yet still informative representation.

To control for position biases on the game interface, we randomize the row and column order of the payment rule once for each participant and show this randomized table throughout the task. We also randomize the order of questions on the quiz and the order of the radio buttons for choosing claims.

Dealing with disconnections Due to the synchronicity of the experiment, disruptions from connection issues are possible. To ensure that a game progresses smoothly when such

	T1	T2	T3	T4	T5
Number of players per game	3	3	4	4	1
Number of rounds	20	20	20	30	20
Number of games without expelled players	103	104	103	103	411

Table 5.2: Treatments

issues occur, we expel a participant from the game if he is disconnected for at least 1 minute (a reasonable threshold since a typical game lasts 5 minutes). An expelled player cannot reconnect to the game, and the server will choose truthful reports on behalf of the expelled player. This ensures that other players experience the game as normal. Our analysis excludes 37 games with expelled players (4% of the games).

5.4 Treatments

We designed and conducted five treatments in sequence (Table 5.2). Because we had no *a priori* prediction of players' empirical behavior, we used the results of the earlier treatments to design the later treatments. We allowed each participant to participate only once in any treatment, and our quiz showed that participants have similar comprehension of the task in all treatments.

For each game, we recruited a small number of players and allowed them to play for a large number of rounds. Having a small number of players minimizes players' waiting times and potential connection issues, as well as making it easier for each player to reason about other players' actions. With a large number of rounds, we hoped to give players sufficient time to explore and improve their strategies.

For all payment rules, we chose the maximum payment of 1.50 to make it more attractive than the 1.00 base payment, and the minimum payment of 0.10 to prevent extreme behavior resulting from attempting to avoid a payment of zero.

Let $\Omega = \{A, B\}$ denote the two types of houses and let $S = \{\text{MM}, \text{GB}\}$ be the common signal and report space. We use the MM (or GB) equilibrium to refer to the symmetric

coordinating equilibrium where all players always report MM (or GB), and let $x\text{MM}y\text{GB}$ denote an asymmetric coordinating equilibrium for a $(x + y)$ -player game where x players always report MM and y players always report GB.

The prior A fixed prior is used for all treatments, as shown in equation (5.1).

$$P(A) = 0.5, \quad P(\text{MM} \mid A) = 0.2, \quad P(\text{MM} \mid B) = 0.7 \quad (5.1)$$

This prior has the nice property that if one player receives a given signal, other players are more likely to have received the same signal. This reasoning may motivate players to be truthful if they believe other players are also truthful.

Treatment 1 For treatments 1 and 2, we used payment rules where each player has one reference report, chosen randomly from all other reports. For treatment 1 (Table 5.3a), all four values in the payment rule are distinct and neither report strictly dominates the other. At the truthful equilibrium, every player obtains 0.91 in expectation. Moreover, both the MM and GB equilibria yield higher payments than the truthful equilibrium. Every player gets the maximum payment of 1.50 at the MM equilibrium, making it the highest-paying choice among all equilibria.

	$r = \text{MM}$	$r = \text{GB}$
$r_f = \text{MM}$	1.50	0.30
$r_f = \text{GB}$	0.10	1.20

(a) Treatment 1 payment rule.

	$r = \text{MM}$	$r = \text{GB}$
$r_f = \text{MM}$	1.50	0.10
$r_f = \text{GB}$	0.10	1.50

(b) Treatment 2 payment rule.

Table 5.3: Payment rules of treatments 1 and 2. The cell at (r, r_f) gives a player's payment if the player reports r and his reference report is r_f .

Treatment 2 For treatment 1, it's easy to identify the MM equilibrium as having the highest payment. Thus, in treatment 2, we modified the payment rule of treatment 1 to be symmetric (Table 5.3b), such that the MM and GB equilibria have the same payments. We hypothesized that such a payment rule may deter the players from reaching either the MM or GB equilibrium, especially without direct communication.

Treatment 3 For payment rules where each player has only one reference report, we are inevitably limited in incentivizing truthful reporting because one of the MM and GB equilibria will always yield higher payment than the truthful equilibrium [Jurca and Faltings, 2009]. To overcome this limitation we tested 4-player payment rules in treatments 3 and 4 where each player’s payment depends on all other players’ reports. We aim to either eliminate the MM and GB equilibria altogether or make these equilibria yield worse payment than the truthful equilibrium.

The payment rule for treatment 3 (Table 5.4a) does not have the MM and GB equilibria. However, it does support the 3MM1GB and 1MM3GB equilibria and an equilibrium where every player always reports the signal he did not receive. Compared to the truthful equilibrium, the 3MM1GB or 1MM3GB equilibria seem more attractive. They give 3 players the maximum payment of 1.50 and the remaining player 0.90, which is comparable to the expected payment of 0.91 at the truthful equilibrium.

	$r = \text{MM}$	$r = \text{GB}$
$n_f = 0$	0.90	0.80
$n_f = 1$	0.10	1.50
$n_f = 2$	1.50	0.10
$n_f = 3$	0.80	0.90

(a) Treatment 3 payment rule.

	$r = \text{MM}$	$r = \text{GB}$
$n_f = 0$	0.10	0.15
$n_f = 1$	0.10	0.90
$n_f = 2$	1.50	0.15
$n_f = 3$	0.15	0.10

(b) Treatment 4 payment rule.

Table 5.4: Payment rules of treatments 3 and 4. The cell at (r, f_f) gives a player’s payment if the player reports r and n_f of his reference reports are MM .

Treatment 4 For treatment 4’s payment rule, the MM and GB equilibria pay very little, as shown in Table 5.4b. At the either MM or GB equilibrium, a player obtains a small payment of 0.15, which is close to the minimum payment of 0.10. In contrast, every player receives 0.50 in expectation at the truthful equilibrium.

We chose this payment rule for two reasons. First, it is impossible to eliminate all of the MM, GB, 3MM1GB and 1MM3GB equilibria for any such 4-player payment rule. Moreover, by supporting the MM and GB equilibria instead of the 3MM1GB and 1MM3GB equilibria,

it is possible to make the payments from the coordinating equilibria much less than those of the truthful equilibrium.

Non Peer-Prediction Treatment For comparison, we would like to understand how players behave when they are not paid by any peer prediction mechanism. In this treatment, each player is paid 0.90 in every round, which is comparable to a player’s expected payment at the truthful equilibrium in treatments 1 and 3. Also, each player plays the game alone, without observing other players’ reports. We believe that this setting is closest to how such a constant payment would be used in practice.

5.5 Results

We collected results on a fairly large scale, recruiting 3533 unique subjects over 65 days for both the pilot and the actual experiment. In the pilot experiment, 705 workers passed the quiz and 542 of them played a total of 181 games. In the actual experiment, 2031 workers passed the quiz and 1988 of them played a total of 861 games.

We received generally positive feedback about the design of our task, suggesting that peer prediction mechanisms can be made accessible to lay participants. Many participants remarked that the game was easy to understand, quick, smooth and enjoyable. 79% of all workers who attempted the quiz eventually passed it, suggesting that the quiz was of appropriate difficulty and the participants had adequate understanding of the mechanism.

5.5.1 Summary of Data

Figure 5.2 shows a summary of our experimental data: the percentage of players receiving a particular signal and making a specified report for each round in each treatment. In treatments 1 and 4, the percentage of players with GB signals and MM reports increases, whereas in treatment 2, the percentage of players with MM signals and GB reports increases. The total percentage of misreporting is smallest in treatment 5.




Figure 5.2: Percentage of players with the specified signal and report

	T1	T2	T3	T4
Actual payoff	1.13	1.05	0.87	0.57
Expected payoff at truthful equilibrium	0.91	0.98	0.90	0.50

Table 5.5: Comparison of actual payoff with expected payoff at truthful equilibrium

Table 5.5 compares the players' average payoffs in the game with their expected payoff at the truthful equilibrium in treatments 1-4. Compared to a player's expected payoff at the truthful equilibrium, the players' average payoff is higher in treatments 1, 2, and 4, and lower in treatment 3.

5.5.2 Learning a Hidden Markov Model

Our main goal is to answer the following question:

Will the players reach one of the multiple equilibria of the game, and if so which equilibrium will the players choose and why?

While equilibrium concepts are defined based on players' strategies, we only observe the players' signals and reports in our experiment. In fact, characterizing unobserved strategies is a common challenge for experimentally testing game-theoretic predictions. A common approach is to test whether a player's actions are consistent with an equilibrium strategy, but this heuristic assumes that a player's actions are drawn from a stationary distribution [Selten and Chmura, 2008]. Alternatively, one might directly elicit mixed strategies from players in the form of a probability distribution over actions [Noussair and Willinger, 2011], but this invasive elicitation method may significantly influence players' behavior.

We use the hidden Markov model [Rabiner, 1989] to accurately capture the players' strategies in repeated games. Widely used in speech recognition, natural language processing, and computational biology, the HMM allows us to infer strategies from actions and analyze how they evolve over time without fixing the strategies *a priori*. Recent work has adopted similar probabilistic models to detect latent behavior in repeated games [Shachat et al., 2012, Ansari et al., 2012].

We model the players' behavior as follows. There are K latent states in the HMM. The j -th state corresponds to the mixed strategy

$$(\boldsymbol{\mu}_j(\text{MM}, \text{MM}), \boldsymbol{\mu}_j(\text{GB}, \text{MM}))$$

where $\boldsymbol{\mu}_j(s, r) = \text{P}(r \mid s)$ is the probability of reporting r when receiving signal s for strategy $j \in \{1, \dots, K\}$. Each player $i \in \{1, \dots, N\}$ chooses his starting strategy from the distribution

$$\boldsymbol{\pi} = (\pi_1, \dots, \pi_K)$$

where π_j is the probability that a player adopts strategy j in the first round. Players change their strategies according a stochastic matrix $\mathbf{A}$, with

$$\mathbf{A}(j, j') = \text{Pr}(\psi_{t+1}^i = j' \mid \psi_t^i = j), \quad \forall j, j' \in \{1, \dots, K\}$$




Figure 5.3: The graphical model for each player i implied by the HMM.

where ψ_t^i is player i 's strategy in round t and $\mathbf{A}(j, j')$ is the probability that a player adopts strategy j in round t and switches to strategy j' in round $t + 1$.

This HMM makes several assumptions about players' behavior. A player only chooses strategies among the K states of the HMM. Moreover, a player's distribution over strategies in round $t + 1$ is Markovian and depends only on her strategy in round t . Thus, each player changes his strategy stochastically according to a fixed transition distribution, and this may not capture the scenario when players change their strategies in response to their peers' actions. However, by considering each player independently, the HMM describes the population as a whole instead of capturing the intentions of individual players. It is thus a natural first step in studying the evolution of unobserved strategies.

Our experimental observations are pairs (s_t^i, r_t^i) , corresponding to player i 's signal s_t^i and report r_t^i in round t . The HMM defines the following probability distribution over players' reports based on their signals and model parameters:

$$\Pr(\mathbf{r} \mid \mathbf{s}, \boldsymbol{\pi}, \boldsymbol{\mu}, \mathbf{A}) = \prod_{i=1}^N \left(\prod_{t=2}^T \boldsymbol{\mu}_{\psi_t^i}(s_t^i, r_t^i) \mathbf{A}(\psi_{t-1}^i, \psi_t^i) \right) \boldsymbol{\mu}_{\psi_1^i}(s_1^i, r_1^i) \cdot \boldsymbol{\pi}(\psi_1^i) \quad (5.2)$$

As shown in Figure 5.3, this model differs from the canonical HMM as both the observed signal s_t^i and the hidden state ψ_t^i influence the observed action r_t^i , but can be estimated using the same methods. We maximize equation (5.2) over $\boldsymbol{\pi}, \boldsymbol{\mu}$, and $\mathbf{A}$ using the Baum-

Welch Expectation-Maximization (EM) algorithm, obtaining maximum likelihood values of the parameters. These parameters reveal the set of strategies, the initial and the long-run distribution over strategies.

Robustness The likelihood function for the HMM is not log-concave and so the local optimum found by the EM algorithm depends on the initial parameters. To get closer to the global optimum, we ran the EM algorithm with 100,000 restarts with random initial parameters and chose the parameters with the highest log likelihood. Many of the restarts produced equivalent “best” solutions. The equivalent solutions contain the same hidden states and transition probabilities in different orders. Therefore, we feel confident that we found solutions equivalent to the global optimum.

Model selection The number of states (strategies) K for the HMM may significantly impact our equilibrium convergence analysis. As K increases, there is a diminishing return on the increase in the log likelihood. A common criterion for model selection is the Bayesian information criterion (BIC) [Schwarz, 1978]. For all treatments, we chose $K = 4$ in order to maximize the BIC.

	GB	MM	Truthful	Mixed	Mixed 2
Treatment 1	(0.13, 0.09)	(1.00, 0.99)	(0.99, 0.04)	(0.82, 0.45)	
Treatment 2	(0.01, 0.00)	(0.99, 0.99)	(0.96, 0.01)	(0.60, 0.32)	
Treatment 3	(0.02, 0.03)	(0.87, 0.97)	(0.97, 0.05)	(0.54, 0.42)	
Treatment 4		(0.96, 0.97)	(0.96, 0.06)	(0.73, 0.61)	(0.34, 0.37)
Constant payment	(0.02, 0.00)		(0.98, 0.02)	(0.16, 0.96)	(0.68, 0.34)

Table 5.6: Each tuple gives the estimated strategy $(\mu_j(\text{MM}, \text{MM}), \mu_j(\text{GB}, \text{MM}))$. All numbers are rounded to 2 decimal places.

Estimated HMM parameters The HMMs estimated for all treatments are described in detail in Appendix C.1. The interpretations of the 4 states of each HMM are shown in Table 5.6. The majority of the estimated strategies are close enough to one of the pure strategies (MM, GB or truthful) and are therefore interpreted correspondingly. Notably, these strategies emerged as a result of estimating our model on the data without *a priori* restrictions. Every state in each estimated HMM has a large self-transition probability

(≥ 0.9). This suggests that players rarely switch between strategies, and they play each strategy for 10 rounds on average before switching to a different strategy.

When estimating each HMM, the number of strategies K affects how the strategies capture the players' noisy behavior. When K increases, the pure strategies in the HMM become less noisy and closer to their theoretical definitions. For instance, as K increases from 3 to 4, the truthful strategy in treatment 1 changes from (0.91, 0.22) to (0.99, 0.04). When K is small, the strategies are less pure because they must incorporate players' exploratory behavior. As K increases, the pure strategies become less noisy because the noisy behavior can be captured using additional strategies.

Equilibrium convergence using HMM analysis The set of strategies in each HMM only describes the population in aggregate. To better understand the actions of each player, we use the Viterbi [1967] algorithm to estimate the most likely sequence of strategies used by each player over the multiple rounds of the game. Formally, for given parameters $\boldsymbol{\mu}, \boldsymbol{\pi}$, and $\mathbf{A}$, the Viterbi algorithm computes

$$\boldsymbol{\psi}^{i*} = (\psi_1^{i*}, \dots, \psi_T^{i*}) = \underset{\boldsymbol{\psi}^i}{\operatorname{argmax}} \left(\prod_{t=2}^T \boldsymbol{\mu}_{\psi_t^i}(s_t^i, r_t^i) \mathbf{A}(\psi_{t-1}^i, \psi_t^i) \right) \boldsymbol{\mu}_{\psi_1^i}(s_1^i, r_1^i) \cdot \boldsymbol{\pi}(\psi_1^i) \quad (5.3)$$


Having a sequence of most likely strategies for each player allows us to characterize equilibrium convergence for each game—here, we are referring to convergence to game-theoretic equilibria and not to the stationary distribution of the Markov chain for the HMM. In the following results, we use the values $\psi^{i*}, i \in \{1, \dots, N\}$, to describe strategies and equilibria in games.

Treatment 1 Results Treatment 1's results are shown in Figure 5.4. Strikingly, many players adopted more profitable strategies and converged to the uninformative equilibria during the game. In particular, they started with the truthful or mixed strategy but switched to and stayed with the MM or GB strategy until the end (Figure 5.4a). The MM and GB strategies are that are close to absorbing under the HMM interpretation.

Moreover, players particularly favored the MM equilibrium, which yields the highest

payoff. 1/3 of the players used the MM strategy throughout the game and another 1/3 of the players switched to playing the MM strategy during the game (Figure 5.4a). The percentage of games playing the MM equilibrium increased dramatically from 8% to around 46%, whereas the total percentage of games playing the GB or truthful equilibria remained less than 10% (Figure 5.4b).

Unsurprisingly, few players are truthful due to the high payoffs of the MM and GB equilibria. Around 15% of the players are truthful throughout the game, but only 5% of the games converged to the truthful equilibrium by the end.



(a) Each row shows how a single player's strategy evolves over multiple rounds.

(b) Fraction of games matching an equilibrium strategy profile in each round.

Figure 5.4: Treatment 1 Results.

Treatment 2 Results In treatment 2, we aimed to deter the players from choosing either the MM or GB equilibrium by giving them the same payoff (Figure 5.5). This was unsuccessful as 36% of the games converged to the GB equilibrium (Figure 5.5b).

Compared to treatment 1, this treatment is better at promoting truthful behavior in the short term, but not in the long run. More players in this treatment adopted the truthful strategy at the beginning—32% of the players in this treatment compared to 16% in treatment 1 (Figures 5.4a and 5.5a). We hypothesize that this happened because it is harder to coordinate on the GB equilibrium than to coordinate on the MM equilibrium in treatment 1.

However, by the end of the game, less than 4% of the games in either treatment converged to the truthful equilibrium (Figures 5.4b and 5.5b).

Interestingly, players clearly favored the GB equilibrium over the MM equilibrium although they yield the same payoffs. One reason for this seems to be a property of the prior: the probability of the GB signal (55%) is higher than that of the MM signal. If players start by being truthful, this property alone could naturally lead them to the GB equilibrium. Indeed, we observe that nearly all players starting with the truthful strategy who switched, changed to the GB rather than the MM strategy (Figure 5.5a). Moreover, players' exit survey answers revealed that they deliberately coordinated on the GB equilibrium once they recognized that other players were more likely to receive the GB candy. These suggest that players have a basic understanding of the prior and can use it to determine their strategies.





Figure 5.5: Treatment 2 Results.

Treatment 3 Results In this treatment, we aimed to promote truthful behavior by eliminating the symmetric coordinating equilibria, although there still exist asymmetric coordinating equilibria (Figure 5.6). As expected, it was empirically much more difficult for players to find and reach the asymmetric equilibria versus the symmetric equilibria. Less than 5% of the games converged to an asymmetric equilibrium (Figure 5.6b) whereas more than 1/3

of the games converged to a symmetric equilibrium in the first two treatments (Figures 5.4b and 5.5b).

However, the increased difficulty of finding a coordinating equilibrium did not promote truthful reporting. Not a single game converged to the truthful equilibrium (Figure 5.6b) despite 20% of players being truthful over the entire game (Figure 5.6a).

The mixed strategy $(0.54, 0.42)$ used by more than half of the players may capture the players' random exploration because they are unable to decide on a strategy for this complex payment rule. Alternatively, this strategy is very close to the random strategy $(0.50, 0.50)$, which is part of a symmetric mixed strategy equilibrium, and 11% of games appear to reach this mixed strategy equilibrium (Figure 5.6b).



(a) Each row shows how a single player's strategy evolves over multiple rounds.

(b) Fraction of games matching an equilibrium strategy profile in each round.

Figure 5.6: Treatment 3 Results.

Treatment 4 Results In this treatment, we made the MM and GB equilibria to have very low payoffs. Similar to treatment 3, this effectively deterred the players from choosing them (Figure 5.7). Less than 2% of the games reached the MM equilibrium and no game reached the GB equilibrium (Figure 5.7b).

However, this did not promote truthful behavior. No game reached the truthful equilibrium (Figure 5.7b) although 16% of the players were truthful during the entire game

(Figure 5.7a).

Players strongly favored the MM strategy over the GB strategy, possibly because of the higher payments in the MM column (Figure 5.4b). A large percentage (26%) of the players adopted the MM strategy by the end of the game (Figure 5.7a). In contrast, they seemed to not consider the GB strategy at all as none of the 4 estimated strategies (Table 5.6) is close to the GB strategy (This is true even for $K = 6$).




Figure 5.7: Treatment 4 Results.

5.5.3 Classifying Convergence to Pure Strategy Equilibria

In contrast to the HMM analysis, a simpler approach to detect equilibrium convergence is to compare players' reports with an equilibrium strategy. Using this simple approach, described below, we classified the convergence of each game to pure strategy equilibria of the JF mechanism. In Table 5.7, we show that the classification results using this simple method and the HMM analysis are almost identical.

Informally, we characterize equilibrium convergence as follows. For round t , we assume that a player's report is consistent with a pure strategy if the report is observed with positive probability given the strategy. A report may be consistent with more than one pure strategy.

Then, we conclude that all players in a game converged to a particular equilibrium in round t if t is the earliest round such that all players' reports from round t to round T are consistent with this equilibrium. Finally, if the players in a game converge to more than one equilibrium, we pick the equilibrium where the convergence occurred the earliest.

Formally, consider the one-shot peer prediction game with N players and T rounds, and a pure strategy equilibrium $e \in E$ of this game. Let $\psi^i(e)$ be player i 's strategy at this equilibrium e , and let ψ_t^i be player i 's strategy in round t that we infer from this analysis. Given player i 's signal s_t^i and report r_t^i that we observed in round t , we conjecture that player i is playing strategy $\psi^i(e)$ in round t if his report is observed with positive probability given the strategy and his signal, that is

$$\mu_{\psi^i(e)}(s_t^i, r_t^i) > 0 \Rightarrow \psi_t^i = \psi^i(e), \quad \forall i, t, e.$$

Note that a player's report in each round may be consistent with multiple pure strategies. Then we say that player i converged to equilibrium e at round

$$t^i(e) = \text{minimum value in } \{1, \dots, T\} \text{ such that } \psi_t^i = \psi^i(e), \forall t^i(e) \leq t \leq T,$$

and that the game converged to equilibrium e at round

$$t(e) = \max_{1 \leq i \leq N} t^i(e).$$

Players in a game may converge to different equilibria at different rounds. Hence, we find the equilibrium for which the convergence occurred earliest in the game. The game converged to the equilibrium $e^* = \operatorname{argmin}_{e \in E} t(e)$ at round $t^* = \min_{e \in E} t(e)$. If convergence occurred sufficiently early ($t^* \leq (T - 5)$), we classify that the game converged to the equilibrium e^* at round t^* . Otherwise, the game remains unclassified.

	Treatment 1		Treatment 2		Treatment 3		Treatment 4	
	HMM	Simple	HMM	Simple	HMM	Simple	HMM	Simple
Truthful	4	5	4	7	0	0	0	0
GB	5	4	37	34	-	-	0	0
MM	47	47	7	7	-	-	2	1
3MM1GB	-	-	-	-	4	4	-	-
1MM3GB	-	-	-	-	1	1	-	-
Unclassified	47	47	56	56	98	98	101	102

Table 5.7: Classification of convergence to pure strategy equilibria using the simple method. Each cell gives the number of games converging to the particular equilibrium in the specified treatment. The symbol “-” means that the equilibrium does not exist for the payment rule tested in the specified treatment.

5.5.4 Non Peer-Prediction Treatment

Remarkably, in the absence of peer prediction methods, this treatment was much more successful in incentivizing truthful reports than all other treatments. 2/3 of the players reported truthfully during the entire game (Figure 5.8), despite having no explicit incentive to do so. We conjecture that the reason for these observations is that truthful reporting is the easiest choice for players who are offered a constant payment, since the cost of exploring and adopting alternative strategies may be greater than that of being truthful in the design of our trick or treat game. When peer prediction is used in the other treatments, it may prompt and even motivate the players to explore non-truthful strategies, leading to the strategic play observed. This implies that peer prediction mechanisms may be better suited for scenarios where behaving truthfully is much more costly than acting strategically.

5.6 Experimental Challenges

In designing an online behavioral experiment to test an economic mechanism, we overcame some unique challenges. We believe that awareness of these issues is important for future experimental work.




Figure 5.8: Non Peer-Prediction Treatment Results. Each row shows how a single player’s strategy evolves over multiple rounds.

Coordinating simultaneous participation of multiple workers Our experiment requires multiple unique workers to participate in each game simultaneously. This is unusual for MTurk, where workers typically participate independently, without interaction, and complete many similar tasks (HITs) at a time.

To coordinate the simultaneous participation of several workers, we built a virtual lobby, where workers wait for enough others to start the game [Mason and Suri, 2012]. However, simple posting such tasks on MTurk is not sufficient to ensure a smooth experience for the workers and to collect high quality data. In our pilot experiment, workers experienced a long wait in the lobby because they accepted our tasks at very different times. This resulted in games with workers not paying attention or accidentally disconnecting from the server, causing further frustration for other workers.

We solved this problem by mimicking the recruitment process for lab experiments. In a separate recruitment task, we described our experiment and paid \$0.10 for workers to consent to being contacted. We invite the recruits via email to participate in our experiment at specified times. Each time, we posted the tasks during a 30-minute window to encourage the timely arrival of the participants and to minimize their waiting time. This recruitment process worked extremely well: we completed the experiment significantly more quickly and

collected higher quality data, compared to our pilot experiment. By establishing prior history with a worker, we can compensate a worker for his time using a bonus payment if he could not finish the experiment due to technical issues. Many of these features are now part of the TurkServer [Mao et al., 2012] platform, and we encourage its use for deploying similar experiments.

Ensuring attention and comprehension Using unique participants, we face the challenge of ensuring workers’ attention and comprehension with only a single interaction. Lab experiments can easily ensure participants’ attention with location and time constraints. In contrast, getting online workers’ full attention can be a luxury due to the distractions in their environment. We addressed this issue in several ways. First, we went through several design iterations to ensure that our task was clear, enjoyable and smooth. We also promised a generous bonus payment contingent on workers’ performance in the task, motivating them to pay attention. Finally, we used a built-in tutorial and quiz to ensure that workers have adequate understanding of the task.

The MTurk meta-environment A platform such as MTurk is a small part of a much bigger online environment. Workers constantly communicate through various forums ¹, review requesters on sites such as TurkOpticon², and use third party services like TurkAlert to monitor requester activities. These systems can significantly influence the results of online experiments [Chandler et al., 2013].

While many such communities have policies to protect the integrity of research data, workers may still unwittingly disclose details of our task that we do not want revealed, or they may speculate and reach incorrect conclusions about the purpose of our experiment. They may also share their confusion and frustration with other workers without notifying the requester. The unique worker requirement also makes the task less attractive than other

¹Popular sites include <http://www.turkernation.com>, <http://www.reddit.com/r/mturk>, <http://www.mturkforum.com>, <http://www.cloudmebaby.com>, and <http://www.mturkgrind.com>.

²<http://turkopticon.ucsd.edu>

tasks that allow repeated participation. Furthermore, a requester’s reputation on review sites such as TurkOpticon can seriously influence his ability to recruit participants in the future.

As a result, we carefully designed our task and conducted our experiments to provide workers with a positive experience. We promptly responded to email communications from workers at all times, especially during the time window when tasks were posted. We made sure our payments compared fairly to market wages, and paid workers promptly upon task completion. We also extensively monitored workers’ activities on the various online communities, advertising our presence to prevent intentional discussion of our task, and responding quickly to workers’ questions. This also allowed us to recruit workers from the broader online community. Despite only able to participate once, many workers left positive TurkOpticon reviews for us and said that they would be happy to work on our tasks again.

5.7 Conclusion and Future Work

In our experimental setting, we did not observe truthful behavior when participants are rewarded using the Jurca and Faltings [2009] mechanism. Players easily converged to the uninformative equilibria, and hampering these equilibria did not induce truthful behavior. In contrast, players are generally truthful in the absence of economic incentives.

This observation, however, may be due to several features of our setting. First, in our setting, the costs of truthful reporting and behaving strategically. This property makes it more likely for the participants to be truthful in the absence of economic incentives. From this perspective, we are interested in the direction of evaluating peer prediction mechanisms in a setting when truthful reporting is much more costly than behaving otherwise. Moreover, in our experiment, we assign players to fixed groups and let them play the game repeatedly with the same peers. This essentially tests the peer prediction as a repeated game. While there is an equilibrium of the repeated game where every player in every round adopts his

strategy at an equilibrium of the stage game (i.e. the one-shot peer prediction game), the repeated game may have other equilibria where players' behavior in each round is different from their behavior at any equilibrium of the stage game. For future work, to alleviate the repeated game effects, we would like to test peer prediction by recruiting a large number of players and randomly rematching the players for each round of the game. Finally, we are also excited about exploring other social, psychological and economic techniques for motivating truthful behavior. For instance, if we inform the participants of the existence of the truthful equilibrium, the participants' behavior may be very different than what we've observed.

Our work shows the promise of evaluating game theoretic mechanisms through online behavioral experiments, especially for mechanisms inherently designed for an online or crowd-sourced setting. Online infrastructures allow for conducting experiments at a much greater scale, as we show in recruiting over 3000 participants for our experiment. In addition, our results also motivate the general use of probabilistic models for analyzing game-theoretic and other experimental data, demonstrating the potential of greatly improved explanatory power of participants' behavior over existing techniques.

Chapter 6

Adaptive Polling for Information Aggregation

Decision making often relies on collecting small pieces of relevant information from many individuals and aggregating such information into a consensus to forecast some event of interest. Such information elicitation and aggregation is especially challenging when the outcome space of the event is large, due to the inherent difficulties in reasoning over and propagating information through the large outcome space in a consistent and efficient manner.

In recent years, online labor markets, such as Amazon Mechanical Turk (MTurk), have become a burgeoning platform for human computation [Law and von Ahn, 2011]. MTurk provides easy access to an ever-growing workforce that is readily available to solve complex problems such as image labeling, translation, and speech-to-text transcriptions. One salient feature of MTurk is that the tasks typically offer small monetary rewards (e.g. 10 cents) and involve simple, one-shot interactions. This leads to a natural problem solving approach where a complex problem is decomposed into many simple, manageable subtasks, such that each worker can make a small, relatively independent contribution towards the overall solution. The algorithm then takes care of integrating the solutions to the subtasks into a coherent final solution to the entire problem.

In this work, we examine whether we can leverage online labor markets’ easy access to participants to effectively solve the information elicitation and aggregation problem for an event with an exponentially large outcome space. Our proposed algorithm, through simple, one-shot interactions, adaptively collects many small pieces of potentially imprecise information from a large number of participants recruited through an online labor market, and integrates these information together into an accurate solution.

We consider a setting with n competing alternatives, each characterized by a hidden strength parameter. Our goal is to produce accurate estimates of the alternatives’ strength parameters in order to rank them. Participants have noisy information about the strengths of the alternatives. We design an adaptive algorithm which produces probabilistic estimates of the strength parameters based on collected pairwise comparison data. Moreover, our adaptive algorithm uses an active learning approach to choose each pairwise comparison question to myopically maximize the expected information gain from each participant. We then evaluate our algorithm through an MTurk experiment for a set of alternatives for which we know the underlying true ranking. Our experimental results show that the adaptive method can gradually incorporate small pieces of collected information and improve the estimates of the strength parameters over time. Compared with presenting a random pairwise comparison question at each step, adaptive questioning has the advantage of reducing the uncertainty of and increasing the accuracy of the estimates more quickly. Interestingly, this is achieved by asking more pairwise comparison questions that are less likely to be answered correctly.

6.1 Related Work

Many elaborate approaches have been developed for event forecasting. For example, prediction markets [Wolfers and Zitzewitz, 2004] allow participants to wager on the outcomes of uncertain events and make profits by improving market predictions. There have been sev-

eral attempts to design expressive prediction markets [Chen et al., 2008, Abernethy et al., 2011, Xia and Pennock, 2011, Pennock and Xia, 2011], especially for forecasting an event with a combinatorial outcome space (e.g. permutation of n alternatives). However, these combinatorial prediction markets can be computationally intractable to operate, and it is more complicated for humans to interact with the markets than participate in simpler elicitation mechanisms such as surveys. A study by Goel et al. [2010] showed that, for predicting outcomes of binary sports events, the relative advantage of using prediction markets instead of polls was very small. This suggests that methods requiring simple interactions with participants may still provide accurate results for the purpose of eliciting and aggregating information.

There is a rapidly evolving human computation literature on designing workflows for solving complex problems using crowdsourcing platforms. The simpler approaches either allow for participants to iteratively improve the solution, or to work on the same problems in parallel [Little et al., 2009, 2010]. More complex workflows attempt to break a problem down into small chunks so that the participants can make relatively independent contributions to the final solution [Kittur et al., 2011, Liem et al., 2011, Noronha et al., 2011]. Our method can be seen as a workflow that aggregates pairwise comparison results from many participants using an adaptive algorithm, and integrates these results into an accurate total ordering of the alternatives.

Our adaptive algorithm characterizes the participants’ noisy information on the strength parameters using the Thurstone-Mosteller model [Thurstone, 1927, Mosteller, 1951], which is a special case of the well known random utility model (RUM) [McFadden, 1974] in economics with Gaussian noise. The Thurstone-Mosteller model has a long history in psychology, econometrics, and statistics, and has been used in preference learning [Brochu et al., 2007, Houlby et al., 2011] and rating chess players [Elo, 1978]. In a recent work, Mao et al. [2013] use voting rules to reconstruct an underlying ranking given noisy full rankings elicited from human subjects. One of their voting rules is created using the Thurstone-Mosteller model.

Carterette et al. [2008] demonstrate from an information retrieval perspective that pairwise comparisons such as used in the Thurstone-Mosteller model are more natural and effective for human preference elicitation than absolute judgments. When the noise follows a Gumbel distribution, the RUM model becomes the Plackett-Luce model [Plackett, 1975, Luce, 2005]. For pairwise comparison, the Plackett-Luce model reduces to the well known Bradley-Terry model [Bradley and Terry, 1952]. We choose to use the Thurstone-Mosteller model because of the tractability in model estimation when using Gaussian noise.

Our algorithm takes an active learning approach to choose the pair of alternatives to query for each participant. Active learning allows us to generate more accurate estimates with fewer pairwise comparisons and less cost. Interested readers can refer to Settles [2009] for a comprehensive survey on active learning. An active learning algorithm may use one of many strategies to evaluate the informativeness of the expected data and to choose the next query. For instance, the algorithm can choose the query that would most change the current model, that would most reduce the generalization error, or that would minimize the variance [Settles, 2009]. We choose the next pair of alternatives to maximize the expected change to our estimates — maximize the expected distance between the current and updated estimated parameter distributions using the Kullback-Leibler divergence as the distance metric. Alternatively, this could be interpreted as choosing the next query to maximize expected information gain according to an information-theoretic metric. Glickman and Jensen [2005] also used this metric to optimally find pairs for tournaments using the Bradley-Terry model.

Parallel to our work is a new research area called “learning to rank” in machine learning. The goal of this area is to rank a list of items given full or partial orderings of the items. Our work is closest to the work by Azari Soufiani et al. [2013b,a]. Azari Soufiani et al. [2013b] assume that the participants’ preferences are generated by generalized random utility model and they use adaptive elicitation according to classical criteria in Bayesian experimental design. The Thurstone-Mosteller model that we use is a special case of generalized random utility model. Our active learning approach is closely related to D-optimality, which is a well

studied criterion in Bayesian experimental design. However, they elicit a full ranking from each participant whereas we only elicit a pairwise comparison. The algorithm developed by Azari Soufiani et al. [2013a] breaks the full rankings into pairwise comparisons, but they did not use active learning for elicitation. Ailon [2011] developed an algorithm to rank some items given pairwise comparisons, and the goal of the algorithm is to produce a ranking which disagrees with as a few pairwise comparisons as possible. They also use active learning to minimize the number of pairwise comparisons required. Long et al. [2010] developed an active learning framework for ranking, but they produce the ranking by eliciting a score for each item rather than pairwise comparisons.

6.2 Our Adaptive Method

We are interested in predicting the ranking of n competing alternatives, where the true ranking is determined by hidden strength parameters s_i for each alternative. If $s_i > s_j$, alternative i is ranked higher than alternative j . Participants have noisy information on the strength parameters.

Our method presents simple pairwise comparison questions to participants and elicits information only on the presented pair of alternatives. Based on the data collected, we estimate the strength parameters of all the alternatives. As it is costly to poll the participants, we adaptively choose (in each iteration) the next pair of alternatives that can provide the largest expected (myopic) improvement to the current estimation.

Let M be a $n \times n$ nonnegative matrix used to record the pairwise comparison results. $M_{i,j}$ denotes the number of times alternative i has been ranked higher than alternative j . Let $M_{i,i} = 0, \forall i$. Then, a high-level summary of our method with T iterations is presented in Algorithm 1 below.

In Section 6.2.1, we introduce the Thurstone-Mosteller model adopted for modeling the noisiness of the participants' information. We discuss the method for estimating the strength

Algorithm 1 Adaptive Information Polling and Aggregation

1. **Initialize M to a nonnegative, invertible matrix, with value 0 on the diagonal.**
 2. $t = 1$.
 - while** $t \leq T$ **do**
 3. **Estimate the strength parameters based on M .** We use the Thurstone-Mosteller model to capture the noisiness of participants' information and obtain the maximum likelihood estimates of the strength parameters. See Sections 6.2.1 and 6.2.2 for details.
 4. **Select a pair of alternatives that maximizes the expected information gain of the parameter estimation.** See Section 6.2.3 for details.
 5. **Obtain the answer to the pairwise comparison question from an participant and update the matrix M .**
 6. $t = t + 1$.
 - end while**
-

parameters of alternatives in Section 6.2.2. Together, these two parts detail how step 3 of Algorithm 1 is carried out. Finally, we explain step 4 of Algorithm 1 in Section 6.2.3.

6.2.1 Noisy Information Model

To model the noisiness of the participants' information, we adopt the Thurstone-Mosteller model or the Probit model with Gaussian noise. One may also adopt the Bradley-Terry model, also called the Logit model, by setting $P(r_i > r_j)$ to $\frac{1}{1+e^{s_j-s_i}}$, the cdf of the logistic distribution. The difference between the two models is very slight, but the Gaussian distribution of the Thurstone-Mosteller model is more tractable for the adaptive approach in our algorithm.

Let $\mathbf{s}' = (s'_1, s'_2, \dots, s'_n)$ represent the absolute strength of the n alternatives. We model a random participant's perceived absolute strength of alternative i as a random variable: $r'_i = s'_i + \epsilon'_i$, where the noise term is Gaussian, $\epsilon'_i \sim \mathcal{N}(0, \sigma^2)$ with unknown σ^2 . Thus, the probability for the participant to rank alternative i higher than alternative j is

$$P(r'_i > r'_j) = P(\epsilon'_j - \epsilon'_i < s'_i - s'_j) = \Phi\left(\frac{s'_i - s'_j}{\sqrt{2}\sigma}\right) \quad (6.1)$$

where $\Phi(\cdot)$ is the cumulative distribution function (cdf) of the standard Gaussian distribution $\mathcal{N}(0, 1)$.

We note that the σ^2 term only affects scaling. Furthermore, with a fixed number of parameters n , only their differences affect the probabilities. Without loss of generality, let

$$s_i = \frac{1}{\sqrt{2}\sigma}(s'_i - s'_k), \quad (6.2)$$

and

$$r_i = \frac{1}{\sqrt{2}\sigma}(r'_i - s'_k), \quad (6.3)$$

where k is an arbitrary reference alternative. We then have $s_k = 0$, and $r_i = s_i + \epsilon_i$, where $\epsilon_i \sim \mathcal{N}(0, 1/2)$. Effectively we only have $n - 1$ unknown parameters. The probability that a participant ranks alternative i higher than alternative j can be written as

$$P(r_i > r_j) = \Phi(s_i - s_j). \quad (6.4)$$

From now on, for simplicity, we will call $\mathbf{s}$ the strength parameters of the alternatives and $\mathbf{r}$ the perceived strength of the alternatives.

6.2.2 Maximum Likelihood Estimation

Given the pairwise comparison results M , we will obtain the maximum likelihood estimates of the strength parameters for the Thurstone-Mosteller model introduced above.

The log likelihood given M is

$$L(\mathbf{s}|M) = \sum_{i,j} M_{i,j} \log(\Phi(s_i - s_j)). \quad (6.5)$$

The maximum likelihood estimators, $\hat{\mathbf{s}}$, are the strength parameters that maximize the log likelihood, i.e. $\hat{\mathbf{s}} \in \operatorname{argmax}_{\mathbf{s}} L(\mathbf{s}|M)$.

Let $\phi(x)$ be the probability density function (pdf) of the standard Gaussian distribution: $\phi(x) = \frac{1}{\sqrt{2\pi}}e^{-\frac{x^2}{2}}$. Note that $\phi(x)$ is log-concave, that is, $\log \phi(x)$ is concave. According to Bagnoli and Bergstrom [1989], the cdf of a log-concave and differentiable pdf is also log-

concave. This means that $\log \Phi(x)$ is concave in x . Thus, the log likelihood function $L(\mathbf{s}|\mathbf{M})$ in (6.5) is a concave function of $\mathbf{s}$ and we only need to consider the first order conditions to solve the optimization problem.

The derivatives of $L(\mathbf{s}|\mathbf{M})$ are

$$\frac{\partial L(\mathbf{M}|\mathbf{s})}{\partial s_i} = \sum_j M_{i,j} \frac{\phi(s_i - s_j)}{\Phi(s_i - s_j)} - \sum_j M_{j,i} \frac{\phi(s_j - s_i)}{\Phi(s_j - s_i)}$$

for all i . Hence, $\hat{\mathbf{s}}$ is the solution to the equation system $\frac{\partial L(\mathbf{M}|\mathbf{s})}{\partial s_i} = 0, \forall i$. This does not have a closed-form solution, but can be solved using numerical methods.

The maximum likelihood estimators $\hat{\mathbf{s}}$ asymptotically follow a multivariate Gaussian distribution. The variance and covariance of $\hat{\mathbf{s}}$ can be estimated using the Hessian matrix of the log likelihood evaluated at $\hat{\mathbf{s}}$. The Hessian matrix has elements

$$\begin{aligned} \frac{\partial^2 L}{\partial s_j \partial s_i} &= M_{i,j} \frac{\phi(s_i - s_j)}{\Phi(s_i - s_j)} \left(s_i - s_j + \frac{\phi(s_i - s_j)}{\Phi(s_i - s_j)} \right) \\ &\quad + M_{j,i} \frac{\phi(s_j - s_i)}{\Phi(s_j - s_i)} \left(s_j - s_i + \frac{\phi(s_j - s_i)}{\Phi(s_j - s_i)} \right) \end{aligned}$$

for $i \neq j$, and

$$\frac{\partial^2 L}{\partial s_i^2} = - \sum_{j:j \neq i} \frac{\partial^2 L}{\partial s_j \partial s_i}$$

for all i . Let $\mathbf{H}(\hat{\mathbf{s}})$ be the Hessian matrix at $\mathbf{s} = \hat{\mathbf{s}}$. Then, the estimated covariance matrix of $\hat{\mathbf{s}}$ is the inverse of negative $\mathbf{H}(\hat{\mathbf{s}})$, i.e.

$$\hat{\Sigma} = (-\mathbf{H}(\hat{\mathbf{s}}))^{-1}.$$

Therefore, given $\mathbf{M}$, our knowledge on $\mathbf{s}$ can be approximated by the multivariate Gaussian distribution $\mathcal{N}(\hat{\mathbf{s}}, \hat{\Sigma})$.

6.2.3 Adaptive Approach

At each iteration, the most valuable poll to present to a participant is on a pair of alternatives that can best improve our current knowledge of the strength parameters.

Let M^c be the matrix of observations, $\hat{\mathbf{s}}^c$ be the estimation of $\mathbf{s}$, and $\hat{\Sigma}^c$ be the estimated covariance matrix of $\hat{\mathbf{s}}^c$ in the current round. Because $r_i = s_i + \epsilon_i$, where $\epsilon_i \sim \mathcal{N}(0, 1/2)$ is independent Gaussian noise, the predicted perceived strength of alternatives by a random participant follows a multivariate Gaussian distribution: $\hat{\mathbf{r}}^c \sim \mathcal{N}(\hat{\mathbf{s}}^c, \hat{\Sigma}^c + \Sigma^\epsilon)$, where Σ^ϵ is the covariance matrix of the ϵ_i and has value 1/2 on the diagonal and 0 everywhere else. Hence, given a pair of alternatives i and j , the predicted probability that a random participant will rank alternative i higher than alternative j is

$$\hat{p}_{i,j}^c = P(\hat{r}_i^c > \hat{r}_j^c) = \Phi \left(\frac{\hat{s}_i^c - \hat{s}_j^c}{1 + \hat{\Sigma}^c(i, i) + \hat{\Sigma}^c(j, j) - 2\hat{\Sigma}^c(i, j)} \right)$$

where $\hat{\Sigma}^c(i, j)$ is the element of $\hat{\Sigma}^c$ at row i and column j . This means that at each iteration, for each pair of alternatives i and j , we can predict how likely a random participant will rank i higher than j and similarly will rank j higher than i .

Suppose we present the pair of alternatives i and j to a participant. If the participant ranks i higher than j , our matrix of observations will become M^{ij} , which is identical to M^c everywhere except $M^{ij}(i, j) = M^c(i, j) + 1$. We denote the approximate distribution obtained from the maximum likelihood estimation given M^{ij} as $\mathcal{N}(\hat{\mathbf{s}}^{ij}, \hat{\Sigma}^{ij})$. Intuitively, if $\mathcal{N}(\hat{\mathbf{s}}^{ij}, \hat{\Sigma}^{ij})$ is very different from our current estimation $\mathcal{N}(\hat{\mathbf{s}}^c, \hat{\Sigma}^c)$, the extra observation has a large information value. Thus, we use the Kullback-Leibler divergence, also called relative entropy, to measure the information value. The Kullback-Leibler divergence between the two multivariate normal distributions is

$$\begin{aligned} D_{\text{KL}}(\mathcal{N}(\hat{\mathbf{s}}^{ij}, \hat{\Sigma}^{ij}) \parallel \mathcal{N}(\hat{\mathbf{s}}^c, \hat{\Sigma}^c)) &= \frac{1}{2} [\text{tr} \left((\hat{\Sigma}^c)^{-1} \hat{\Sigma}^{ij} \right) \\ &+ (\hat{\mathbf{s}}^c - \hat{\mathbf{s}}^{ij})^\top (\hat{\Sigma}^c)^{-1} (\hat{\mathbf{s}}^c - \hat{\mathbf{s}}^{ij}) - \log \left(\frac{|\hat{\Sigma}^{ij}|}{|\hat{\Sigma}^c|} \right) - n], \end{aligned} \quad (6.6)$$

where n is the dimension of the random vectors, which equals the number of alternatives, and $|\hat{\Sigma}^{ij}|$ is the determinant of $\hat{\Sigma}^{ij}$. Similarly, if the participant ranks j higher than i , our matrix of observations will become M^{ji} , which is identical to M^c everywhere except $M^{ji}(j, i) = M^c(j, i) + 1$. The new approximate distribution becomes $\mathcal{N}(\hat{\mathbf{s}}^{ji}, \hat{\Sigma}^{ji})$. The Kullback-Leibler divergence $D_{\text{KL}}(\mathcal{N}(\hat{\mathbf{s}}^{ji}, \hat{\Sigma}^{ji}) \parallel \mathcal{N}(\hat{\mathbf{s}}^c, \hat{\Sigma}^c))$ can be calculated analogously to (6.6).

Putting all pieces together, for each pair of alternatives i and j , we can calculate the expected information gain of polling an participant on the pair as

$$g(i, j) = \hat{p}_{i,j}^c D_{\text{KL}}(\mathcal{N}(\hat{\mathbf{s}}^{ij}, \hat{\Sigma}^{ij}) \parallel \mathcal{N}(\hat{\mathbf{s}}^c, \hat{\Sigma}^c)) + \hat{p}_{j,i}^c D_{\text{KL}}(\mathcal{N}(\hat{\mathbf{s}}^{ji}, \hat{\Sigma}^{ji}) \parallel \mathcal{N}(\hat{\mathbf{s}}^c, \hat{\Sigma}^c)). \quad (6.7)$$

At each iteration, we pick the pair with the maximum expected information gain and present it to another participant.




Figure 6.1: Two example pictures. The left picture has 342 dots, and the right one has 447 dots.

6.3 Experiment Design

We experimentally evaluate the effectiveness of our method with participants recruited from MTurk. In our experiment, each alternative was a picture containing a relatively large

number of dots [Horton, 2010]. We generated 12 different pictures, each having 318, 335, 342, 344, 355, 381, 383, 399, 422, 447, 460, and 469 non-overlapping dots respectively. The number of dots x in each picture was independently drawn according to a distribution such that $P(x) \propto 1/x$ for $x \in [300, 500]$. Figure 6.1 presents two example pictures used in the experiment. The goal was to use our method to estimate the relative number of dots in these 12 pictures in order to correctly rank the pictures in decreasing number of dots. We chose pictures with dots as the alternatives for our experiment for several reasons:

1. We know the correct ranking and can more objectively evaluate the proposed method.
2. The number of dots in each picture is large enough that counting is not an option for participants, introducing uncertainty.
3. The differences in number of dots across pictures vary and some pairs are more difficult to compare than others; for example, pictures in some adjacent pairs differ by only 2 dots, while those in some other adjacent pair are separated by 26 dots.

We ran our experiment on MTurk. For each HIT (Human Intelligence Task in MTurk’s terminology), we presented a pair of pictures, randomly placing one on the left and the other on the right, and asked a MTurk user (Turker) to choose the picture that contained more dots. The base reward for completing a HIT was \$0.05. If the Turker correctly selected the picture with more dots, we provided another \$0.05 as a bonus. Using our adaptive method, we compute an estimate of the strength parameters which reflect the relative differences between the number of dots in the pictures, and decide which pair of pictures to present to the next Turker.

The matrix M was initialized to have value 0 on the diagonal and 0.08 everywhere else. This effectively set our initial estimate of the strength parameters to be $\mathcal{N}(\mathbf{0}, \Sigma^0)$, where Σ^0 had value 1.64 on the diagonal and value 0.82 everywhere else. This can be interpreted as our prior belief of the strength parameters without any information.

We ran 6 trials of adaptive polling. For each trial, we recruited 100 participants assuming that the budget is only enough for collecting 100 correct answers. To evaluate the advantage of our adaptive approach, we ran another 6 trials (with 100 HITs in each trial) of the random polling method, where the pair in each HIT was randomly selected. In our experiment, each HIT was completed by a Turker with a unique ID. In other words, we interacted with each participating Turker only once.

6.4 Results

The actual number of dots in each picture can be considered its absolute strength parameter s'_i , the value of which we know as the experimenter. However, in order to evaluate our method, we need to establish a “gold standard” for the strength parameters relative to the strength $s_k = 0$ of a reference alternative k , as defined in equation (6.2). To transform the absolute strength parameters into these “gold standard” strength parameters, we need a good estimate of $\frac{1}{\sqrt{2}\sigma}$ according to equation (6.2), $s_i = \frac{1}{\sqrt{2}\sigma}(s'_i - s'_k)$. We run a Probit regression [McCullagh and Nelder, 1989] on the 1200 pairwise comparison results collected from all 12 trials. Specifically, let Y be 1 if the left picture is selected and 0 if the right picture is selected. Let X be the number of dots in the left picture minus the number of dots in the right picture. Then, $P(Y = 1|X) = \Phi(X\beta)$, where $\beta = \frac{1}{\sqrt{2}\sigma}$, and we have 1200 observations for (X, Y) . The Probit regression gives us an estimate $\hat{\beta} = 0.017$. Multiplying $(s'_i - s'_k)$ by $\hat{\beta}$, we obtain the “gold standard” strength parameters -0.41, -0.12, 0, 0.03, 0.22, 0.66, 0.7, 0.97, 1.36, 1.79, 2.01, and 2.16 for the 12 pictures. The picture with 342 dots (the third lowest) is used as the reference picture and hence has a strength parameter of 0. Since we only perform a linear transformation, a picture with more dots has a larger “gold standard” strength parameter.

A fair concern with our model is whether the Thurstone-Mosteller model accurately captures the participants’ information in our setting. To evaluate this assumption, we compare




Figure 6.2: Frequency of the left picture being selected in the 1200 pairwise comparisons of all 12 trials. The x-axis represents the difference in number of dots between the left and right pictures (left – right). The observations are grouped into 7 buckets according to the difference in dots. Each bar represents the empirical frequency for the corresponding bucket. The curve is $\Phi(0.017x)$.

the empirical frequencies of the Turkers’ responses with those predicted by the Thurstone-Mosteller model. By equation (6.1), the probability for a participant to select picture i in a pairwise comparison between pictures i and j is $\Phi\left(\frac{s'_i - s'_j}{\sqrt{2}\sigma}\right)$, and we estimated $\frac{1}{\sqrt{2}\sigma} = 0.017$ using all the collected data. Thus, the Thurstone-Mosteller model predicts that the empirical frequencies of the Turkers’ responses should closely follow the distribution $\Phi(0.017(s'_i - s'_j))$. Figure 6.2 plots the empirical frequency of the left picture being selected in our experiment for seven brackets of differences in dots between the left and right pictures. The empirical frequency matches the cdf well, indicating that our setting does not significantly deviate from the Thurstone-Mosteller model. We notice that Turkers have a slight bias toward selecting the picture on the right, because when the difference in number of dots is around 0, the frequency of the left picture being selected is about 40%, in contrast to the 50% predicted by the model.




Figure 6.3: The dynamics of the estimated strength parameters for an adaptive polling trial. The x-axis is the number of iterations. The y-axis is the value of the estimated strength parameters. The rightmost part of the figure labels the value of the “gold standard” strength parameter for each picture.

Next, we look into whether our method effectively incorporates information over time. Figure 6.3 shows the dynamics of the estimated strength parameters for one of the adaptive polling trials. The figures for all adaptive and random polling trials (Figures D.1 and D.2) are presented in Appendix D.

Since the strength parameter for the picture with 342 dots is set to 0, the estimates are for the other 11 pictures. The lines are colored in grayscale such that the lightest color corresponds to the picture with the most dots and the darkest line corresponds to the picture with fewest dots. We can see that all pictures start with an estimated strength parameter of 0. As more pairwise comparisons are polled, the estimated strength parameters diverge. The overall trend is that the estimated strength parameters of pictures with more dots increase and those of pictures with less dots decrease, showing that information is aggregated into the estimates over time. The right side of Figure 6.3 labels the value of the “gold standard”

strength parameter for each picture. At the end of 100 iterations, the estimated strength parameters are close to the gold standard strength parameters. The produced ranking is generally correct, except that two adjacent pairs are flipped. A closer look reveals that these two flipped pairs have the smallest difference in dots among all adjacent pairs of the 11 pictures, with 381 and 383 dots and 344 and 355 dots respectively.




Figure 6.4: The entropy of the estimated distribution $\mathcal{N}(\hat{s}, \hat{\Sigma})$




Figure 6.5: The log score — the logarithm of the pdf of $\mathcal{N}(\hat{s}, \hat{\Sigma})$ evaluated at the “gold standard” strength parameters

Finally, we compare the performance of adaptive polling with that of random polling. In addition to our collected data, we also run 100 trials of simulation for each method using the “gold standard” strength parameters to understand what we should expect to see if our model perfectly captures the noisiness of the setting and we know the strength parameters. Figures 6.4, 6.5 and 6.6 present the results of the comparison of the performance of our adaptive polling method versus the random polling method.

Intuitively, we expect the adaptive method to reduce the entropy of the estimated distribution more quickly than the random method, since the adaptive method is optimized for quickly reducing the uncertainty of the probabilistic estimates of the strength parameters. In Figure 6.4, we show a plot of the entropy of the estimated distribution $\mathcal{N}(\hat{s}, \hat{\Sigma})$, which is




Figure 6.6: The fraction of the pairwise comparison questions that are correctly answered calculated as $\log \sqrt{(2\pi e)^n |\hat{\Sigma}|}$ where $|\hat{\Sigma}|$ is the determinant of $\hat{\Sigma}$. This figure confirms that the entropy of the estimated distribution indeed decreases faster for the adaptive polling than for the random polling. The difference in the entropy produced by the two methods is statistically significant by two-tailed t-test ($p = 0.01$).

Next, Figure 6.5 presents a comparison of the log score of the estimated distributions for the two methods. The log score is often used to measure the accuracy of a probabilistic prediction, so it is a good indicator for how well our method performs in estimating the strength parameters. Having a high log score means that our method produces accurate estimates of the strength parameters. Given an estimated distribution $\mathcal{N}(\hat{s}, \hat{\Sigma})$ and the “gold standard” strength parameters $\mathbf{s}$, the log score is the logarithm of the pdf of $\mathcal{N}(\hat{s}, \hat{\Sigma})$ evaluated at $\mathbf{s}$. Figure 6.5 shows that the log scores for both adaptive and random polling increase over time. The log scores for adaptive polling are higher but the variation is large. The difference in the log score produced by the two methods is statistically significant by two-tailed t-test ($p = 0.016$).

Interestingly, according to Figure 6.6, the fraction of pairwise comparison questions that are answered correctly is lower for adaptive polling than for the random polling, and the dif-

ference is statistically significant by two-tailed t-test ($p = 0.005$). This observation suggests that adaptive polling tends to ask relatively difficult comparison questions. The answers to these questions are more valuable for improving the estimates of the strength parameters, even though the participants are less likely to answer them correctly. Moreover, since we pay Turkers a bonus only for correct answers, this implies that the cost of adaptive polling is lower than that of random polling. For our experiment, an average of 10% in bonus payment is saved per trial by using adaptive polling instead of the random method.

6.5 Conclusion and Future Directions

In conclusion, we demonstrate that eliciting and aggregating information about the ranking of n competing alternatives can be effectively achieved by adaptively polling participants recruited from an online labor market on simple pairwise comparison questions and gradually incorporating the collected information into an overall prediction. Our adaptive polling method is robust against the unpredictable noise in the participants' information and it is effective in eliciting and aggregating information while requiring only simple interactions with the participants. With the same number of participants per trial, our adaptive polling method derives estimates with higher accuracy while requiring 10% less payments compared to the random polling method.

As discussed in section 6.1, there are several different strategies in active learning to evaluate how each query affects the informativeness of the expected data. In our algorithm, we choose each query that would most change the current estimates, and we interpret a significant change to the current estimates as informational gain. However, the active learning literature suggests many other ways to choose the next query, i.e. choosing the next query that would most reduce the uncertainty or the variance in our estimates. It would be interesting to evaluate the effects of adopting different active learning approaches on the performance of our algorithm.

The baseline method in our MTurk experiment chooses each pairwise comparison question randomly among all possible such questions. This is the most naive way to choose the questions. One can imagine choosing each question using some smarter method that is less sophisticated than active learning. For example, we can choose the pairwise comparisons for which we have least data or pairs of alternatives whose estimated strength parameters are close. It would be interesting to see whether our adaptive algorithm could still outperform these smarter baseline methods.

Although the Thurstone-Mosteller model suitably captures the noisiness of participants' information in our experiments, it has some limitations. The model implicitly assumes that participants are ex-ante equally informed and their mistakes are independent. These may not hold in some settings where some participants are better informed than others and mistakes of participants are correlated. In future work, we are interested in studying how our approach performs in such settings and developing suitable methods for them.

Even though we only evaluated our method for a setting with a known underlying ranking of the alternatives, our method can be easily adapted for settings when the underlying ranking is unknown. In this case, it is crucial to decide on a suitable termination condition for our algorithm. Since our model produces probabilistic estimates of the strength parameters, we could, for instance, choose to stop the algorithm once a desired entropy of the estimated distribution is reached. It is an interesting future direction to explore different termination conditions for applying our algorithm to such settings.

Chapter 7

Conclusion and Future Directions

The goal of this dissertation is to tackle the following research question:

How to generate an accurate estimate or prediction of an event of interest by eliciting dispersed information from multiple individuals and aggregating these information together?

I study three information elicitation and aggregation methods: prediction markets, peer prediction mechanisms, and adaptive polling, using both theoretical and applied approaches.

In Chapters 3 and 4, we theoretically characterize the equilibrium behavior of the participants when trading in market scoring rule prediction markets. Myopic participants in a prediction market are incentivized to truthfully reveal their private information. However, if a non-myopic market participant receives multiple payoffs from inside and/or outside of the market, then the participant may want to withhold or misreport his information early on in order to maximize his total payoff. We show that non-myopic participants' equilibrium behavior critically depends on the information structure (i.e. how the participants' private information relates to the event of interest) and the payoff structure of the prediction market. In Chapter 3, when the non-myopic participants have multiple opportunities to trade and receive multiple payoffs in the market, we prove that the information structure determines whether the participants reveal information early or late in the market at an equilibrium. In Chapter 4, the non-myopic participants receive payoffs from both inside and outside of

the market. In this case, our results indicate that the information and the payoff structures together determine whether all private information gets fully revealed inside the market.

In Chapters 5, we experimentally study participants' behavior towards the Jurca and Faltings [2009] peer prediction mechanism, in order to understand whether the participants will play one of the multiple equilibria characterized in theory. In theory, the Jurca and Faltings [2009] mechanism not only supports the truthful equilibrium, it also induces uninformative equilibria where participants reveal no information to the mechanism. Moreover, the mechanism has not been evaluated in practice and its theory provides little support that the participants will be truthful in practice. We conduct a controlled online experiment to evaluate the Jurca and Faltings [2009] mechanism through a multi-player, real-time and repeated game. In our experimental setting, we observed that the participants are not truthful and they successfully coordinated on the uninformative equilibria. In contrast, the participants are generally truthful in the absence of economic incentives.

One unifying theme in Chapters 3, 4 and 5 is understanding how strategic participants behave when interacting with the mechanisms. The participants are assumed to be rational and self-interested economic agents who seek to maximize their rewards from the mechanism. When assuming that the participants are strategic, our ultimate goal is to design mechanisms such that the participants will always truthfully reveal their private information as soon as the information is available. However, for both prediction markets and peer prediction mechanisms, we demonstrate that the participants may not be truthful in theoretical or experimental settings. In prediction markets, the strategic participants may withhold their information or delay revealing their information. For peer prediction mechanisms, the participants may coordinate on uninformative equilibria which reveal no information to the mechanism.

These undesirable strategic behavior of the participants inspires an immediate future direction: designing new and improved mechanisms to have stronger theoretical guarantees for truthfulness. For instance, can we design prediction markets to have better truthful-

ness properties for non-myopic participants while still maintaining incentive compatibility for myopic participants? Can we design new peer prediction mechanism such that the uninformative equilibria are eliminated or the truthful equilibrium can be shown to be focal with strong theoretical and empirical evidence? Several recent work has already proposed peer prediction mechanisms with better theoretical properties.

For prediction markets and peer prediction mechanisms, we assume that the pieces of information that the participants may potentially obtain are known in advance. We also assume that each participant knows the several pieces of information he may obtain, the participant already possesses one such piece of information, and he decides whether to truthfully reveal his information or not. This model has several assumptions that seem unreasonable in practice. It may be unreasonable to believe that the potential pieces of information are fixed and known in advance. In practice, participants need to invest costly effort to gather useful information and to synthesize the information together. In the current model, a participant may decide not report his information truthfully because making a different report can improve his expected payoff. However, when the participant needs to invest effort to gather information, he may be tempted to provide a report which requires as little effort as possible. This reason for misreporting may be more reasonable and is more likely to occur in practice. Such arguments have inspired recent work to design and develop new peer prediction mechanisms which incentivize the participants to invest in costly effort to obtain accurate information [Jens Witkowski, 2013, Dasgupta and Ghosh, 2013].

There are several future directions that aim to reconcile the gap between the theoretical understanding and the practical evaluations of these mechanisms. For example, prediction markets have been shown to produce remarkably accurate forecasts in practice. However, our theoretical results show that the participants' strategic behavior in the market may damage the information aggregation process, which seemingly contradict the empirical observations. Similarly, the literature on peer prediction mechanisms focuses only on designing these mechanisms to support the truthful equilibrium. However, our experimental results suggest that

participants may coordinate on the uninformative equilibria and reveal no information to the mechanism. To address this gap between theory and practice, one future direction is to evaluate these mechanisms in experimental or practical settings in order to understand the extent to which the participants' behavior deviate from our theoretical models. Knowing the deviation, we could either develop better theoretical models to capture the market participants' behavior in practice or seek novel practical methods to motivate the participants to be truthful.

In Chapter 6, we design an adaptive polling method for estimating the outcome of an event without observable ground truth, which is the same problem we study in Chapter 5. In contrast to prediction markets and peer prediction mechanisms, we make very different assumptions about the participants' behavior. We assume that each participant, when queried, will truthfully reveal to us one partial and noisy piece of information about the latent ground truth. Using a theoretical model to capture the noise in the participants' information, the adaptive polling method estimates the latent ground truth by aggregating these partial and noisy pieces of information together and determines the piece of information to query from the next participant in order to maximize the information gain. We apply our method to the problem of ranking n competing alternatives, each characterized by a hidden strength parameter. The method queries each participant for the result of a particular pairwise comparison. Through a MTurk experiment, we show that the adaptive polling method can effectively aggregate information over time and outperforms a naive method, which chooses a random pairwise comparison question at each step.

The model of participants' behavior used in Chapter 6 is very different from the model used in Chapters 3, 4 and 5. Both models are motivated by practical scenarios. However, each model only captures a particular aspect of the participants' behavior and the aspect captured may be more or less prominent depending on the particular setting considered. The model of strategic participants used in Chapters 3, 4 and 5 is widely used in theoretical analyses but so far seems to have little supporting evidence in practice. For instance, prediction markets

are successful empirically but there are many theoretical results on how the participants' strategic behavior damage the information aggregation process. This contrast suggests that the theoretical models of the market participants may not accurately capture the market participants' behavior in practice. For peer prediction mechanisms, we obtained evidence of participants' strategic behavior in our highly controlled experimental settings, but we have yet to verify whether these strategic behavior will be observed in more realistic scenarios. In contrast, our results in Chapter 6 provide positive support for the "noisy information" model of the participants. In particular, we show that our theoretical model accurately captures the noise in the participants' information compared to the ground truth. These positive evidence, however, may be a result of the specific setting that we chose to evaluate the adaptive polling method. In summary, both models make simplifying assumptions about the participants' behavior. Each model is a reasonable approximation of the participants' behavior only for certain situations. In fact, a better model of the participants should contain both the strategic and the noisy information aspects of their behavior, thus encompassing the two current models as special cases. The challenge then is to find a suitable way to combine the two models. For example, different aspects of the participants' behavior may be more prominent in different settings, so the combined model needs to be flexible enough to allow the weights of the two aspects be adjusted for different situations.

Bibliography

- Jacob Abernethy, Yiling Chen, and Jennifer Wortman Vaughan. An optimization-based framework for automated market-making. In *Proceedings of the 12th ACM conference on Electronic commerce*, EC '11, pages 297–306, New York, NY, USA, 2011. ACM. ISBN 978-1-4503-0261-6. doi: <http://doi.acm.org/10.1145/1993574.1993621>. URL <http://doi.acm.org/10.1145/1993574.1993621>.
- Jacob Abernethy, Yiling Chen, and Jennifer Wortman Vaughan. Efficient market making via convex optimization, and a connection to online learning. In *ACM Transactions on Economics and Computation*, New York, NY, USA, 2013. ACM.
- Nir Ailon. Active learning ranking from pairwise preferences with almost optimal query complexity. In J. Shawe-Taylor, R.S. Zemel, P. Bartlett, F.C.N. Pereira, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems 24*, pages 810–818. 2011.
- Franklin Allen and Douglas Gale. Stock-price manipulation. *The Review of Financial Studies*, 5(3):503–529, 1992.
- Asim Ansari, Ricardo Montoya, and Oded Netzer. Dynamic learning in behavioral games: A hidden markov mixture of experts approach. *Quantitative Marketing and Economics*, 10(4):475–503, 2012.
- Hossein Azari Soufiani, William Chen, David C. Parkes, and Lirong Xia. Generalized Method-of-Moments for Rank Aggregation. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NIPS 2013)*, 2013a.
- Hossein Azari Soufiani, David C. Parkes, and Lirong Xia. Preference Elicitation For General Random Utility Models. In *Proceedings of the 29th Conference on Uncertainty in Artificial Intelligence (UAI-13)*, 2013b.
- M. Bagnoli and Bergstrom. Log-concave probability and its applications. *Economic Theory*, 26:445–469, 1989.
- Joyce E. Berg, Robert Forsythe, Forrest D. Nelson, and Thomas A. Rietz. Results from a dozen years of election futures markets research. In C. A. Plott and V. Smith, editors, *Handbook of Experimental Economic Results*. 2001.

- Craig Boutilier. 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, pages 737–744, 2012. ISBN 0-9817381-2-5, 978-0-9817381-2-3.
- Ralph Allan Bradley and Milton E. Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):pp. 324–345, 1952. ISSN 00063444. URL <http://www.jstor.org/stable/2334029>.
- Eric Brochu, Nando de Freitas, and Abhijeet Ghosh. Active preference learning with discrete choice data. In *Advances in Neural Information Processing Systems*, 2007.
- Colin F. Camerer. Can asset markets be manipulated? A field experiment with race-track betting. *Journal of Political Economy*, 106(3):457–482, 1998.
- Ben Carterette, Paul N. Bennett, David Maxwell Chickering, and Susan T. Dumais. Here or there: Preference judgments for relevance, 2008.
- Archishman Chakraborty and Bilge Yilmaz. Manipulation in market order models. *Journal of Financial Markets*, 7(2):187–206, 2004.
- Jesse Chandler, Pam Mueller, and Gabriele Paolacci. Nonnaïveté among amazon mechanical turk workers: Consequences and solutions for behavioral researchers. *Behavior Research Methods*, 2013. doi: 10.3758/s13428-013-0365-7. URL <http://dx.doi.org/10.3758/s13428-013-0365-7>.
- Kay-Yut Chen and Charles R. Plott. Information aggregation mechanisms: Concept, design and implementation for a sales forecasting problem. Working paper No. 1131, California Institute of Technology, 2002.
- Yiling Chen and Ian A. Kash. Information elicitation for decision making. *AAMAS'11: Proc. of the 10th Int. Conf. on Autonomous Agents and Multiagent Systems*, 2011.
- Yiling Chen and David M. Pennock. A utility framework for bounded-loss market makers. In *UAI '07: Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence*, pages 49–56, 2007.
- Yiling Chen, Daniel M. Reeves, David M. Pennock, Robin D. Hanson, Lance Fortnow, and Rica Gonen. Bluffing and strategic reticence in prediction markets. In *WINE'07: Proceedings of the 3rd international conference on Internet and network economics*, pages 70–81, Berlin, Heidelberg, 2007. Springer-Verlag. ISBN 3-540-77104-2, 978-3-540-77104-3.
- Yiling Chen, Lance Fortnow, Nicolas Lambert, David M. Pennock, and Jennifer Wortman. Complexity of combinatorial market makers. In *EC '08: Proceedings of the 9th ACM conference on Electronic commerce*, pages 190–199, New York, NY, USA, 2008. ACM. ISBN 978-1-60558-169-9. doi: <http://doi.acm.org/10.1145/1386790.1386822>.

- Yiling Chen, Stanko Dimitrov, Rahul Sami, Daniel Reeves, David M Pennock, Robin Hanson, Lance Fortnow, and Rica Gonen. Gaming prediction markets: Equilibrium strategies with a market maker. *Algorithmica*, 58(4):930–969, 2010a.
- Yiling Chen, Stanko Dimitrov, Rahul Sami, Daniel M. Reeves, David M. Pennock, Robin D. Hanson, Lance Fortnow, and Rica Gonen. Gaming prediction markets: Equilibrium strategies with a market maker. *Algorithmica*, 58:930–969, December 2010b. ISSN 0178-4617. doi: <http://dx.doi.org/10.1007/s00453-009-9323-2>. URL <http://dx.doi.org/10.1007/s00453-009-9323-2>.
- Yiling Chen, Ian Kash, Mike Ruberry, and Victor Shnayder. Decision Markets with Good Incentives. In *Proceedings of the Seventh Workshop on Internet and Network Economics (WINE)*, 2011.
- In-Koo. Cho and David M. Kreps. Signalling games and stable equilibria. *Quarterly Journal of Economics*, 102:179–221, 1987.
- Anirban Dasgupta and Arpita Ghosh. Crowdsourced judgement elicitation with endogenous proficiency. *ACM International World Wide Web Conference (WWW)*, 2013.
- Sandip Debnath, David M. Pennock, C. Lee Giles, and Steve Lawrence. Information incorporation in online in-game sports betting markets. In *EC '03: Proceedings of the 4th ACM conference on Electronic commerce*, pages 258–259, New York, NY, USA, 2003. ACM. ISBN 1-58113-679-X. doi: <http://doi.acm.org/10.1145/779928.779987>.
- S Dimitrov and Rahul Sami. Composition of markets with conflicting incentives. In *Proceedings of the 11th ACM conference on Electronic commerce (EC '10)*, pages 53–62, New York, 2010a. ACM.
- Stanko Dimitrov and Rahul Sami. Non-myopic strategies in prediction markets. In *The 2nd Workshop on Prediction Markets*, 2007.
- Stanko Dimitrov and Rahul Sami. Composition of markets with conflicting incentives. *EC'10: Proc. of the 11th ACM Conf. on Electronic Commerce*, pages 53–62, 2010b.
- Arpad Elo. *The Rating of Chess Players: Past and Present*. Acro Publishing, New York, 1978.
- Robert Forsythe, Forrest Nelson, George R. Neumann, and Jack Wright. Anatomy of an experimental political stock market. *American Economic Review*, 82(5):1142–1161, 1992.
- Robert Forsythe, Thomas A. Rietz, and Thomas W. Ross. Wishes, expectations and actions: a survey on price formation in election stock markets. *Journal of Economic Behavior & Organization*, 39(1):83–110, May 1999. URL <http://ideas.repec.org/a/eee/jeborg/v39y1999i1p83-110.html>.
- Drew Fudenberg and Jean Tirole. *Game Theory*. MIT Press, 1991.

- Xi Alice Gao, Yoram Bachrach, Peter Key, and Thore Graepel. Quality expectation-variance tradeoffs in crowdsourcing contests. In *AAAI*, 2012.
- Mark E. Glickman and Shane T. Jensen. Adaptive paired comparison design. *Journal of Statistical Planning and Inference*, 127:279–293, 2005.
- Tilmann Gneiting and Adrian E Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.
- Sharad Goel, Daniel M. Reeves, Duncan J. Watts, and David M. Pennock. Prediction without markets. In *EC '10: Proceedings of the 11th ACM Conference on Electronic Commerce*, pages 357–366, New York, NY, USA, 2010. ACM.
- Irving John Good. Rational decisions. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 107–114, 1952.
- Jan Hansen, Carsten Schmidt, and Martin Strobel. Manipulation in political stock markets—preconditions and evidence. *Applied Economics Letters*, 11(7):459–463, 2004a.
- Jan Hansen, Carsten Schmidt, and Martin Strobel. Manipulation in political stock markets - preconditions and evidence. *Applied Economics Letters*, 11(7):459–463, 2004b.
- Robin Hanson. Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets*, 1(1):3–15, 2007a. URL <http://econpapers.repec.org/RePEc:buc:jpredm:v:1:y:2007:i:1:p:3-15>.
- Robin D. Hanson. Combinatorial information market design. *Information Systems Frontiers*, 5(1):107–119, 2003.
- Robin D. Hanson. Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets*, 1(1):1–15, 2007b.
- Robin D. Hanson, Ryan Oprea, and Dave Porter. Information aggregation and manipulation in an experimental market. *Journal of Economic Behavior and Organization*, 60(4):449–459, 2007.
- John J. Horton. The dot-guessing game: A fruit fly for human computation research. *SSRN eLibrary*, 2010. doi: 10.2139/ssrn.1600372.
- John J Horton, David G Rand, and Richard J Zeckhauser. The online laboratory: conducting experiments in a real labor market. *Experimental Economics*, 14(3):399–425, February 2011.
- Neil Houlsby, Ferenc Huszár, Zoubin Ghahramani, and Máté Lengyel. Bayesian active learning for classification and preference learning. arXiv:1112.5745, 2011.
- Krishnamur Iyer, Ramesh Johari, and Ciamac C. Moallemi. Information aggregation in smooth markets. In *EC '10: Proceedings of the 11th ACM Conference on Electronic Commerce*, pages 199–205, New York, NY, USA, 2010a. ACM.

- Krishnamurthy Iyer, Ramesh Johari, and Ciamac C Moallemi. Information aggregation in smooth markets. *EC'10: Proc. of the 11th ACM Conf. on Electronic Commerce*, pages 199–206, 2010b.
- Yoram Bachrach Peter Key David C Parkes Jens Witkowski. Dwelling on the Negative: Incentivizing Effort in Peer Prediction. *HCOMP*, pages 1–8, 2013.
- Lian Jian and Rahul Sami. Aggregation and manipulation in prediction markets: effects of trading mechanism and information distribution. *EC'10: Proc. of the 11th ACM Conf. on Electronic Commerce*, pages 207–208, 2010.
- L K John, G Loewenstein, and Drazen Prelec. Measuring the Prevalence of Questionable Research Practices With Incentives for Truth Telling. *Psychological Science*, 23(5):524–532, May 2012.
- R Jurca and B Faltings. Mechanisms for making crowds truthful. *Journal of Artificial Intelligence Research*, 34(1):209–253, 2009.
- Aniket Kittur, Boris Smus, Susheel Khamkar, and Robert E. Kraut. Crowdforge: crowd-sourcing complex work. In *Proceedings of the 24th annual ACM symposium on User interface software and technology*, UIST '11, pages 43–52, New York, NY, USA, 2011. ACM.
- Praveen Kumar and Duane J. Seppi. Futures manipulation with “cash settlement”. *Journal of Finance*, 47(4):1485–1502, 1992.
- Christina Ann Lacombe, Janet Arlie Barnett, and Qimei Pan. The imagination market. *Information Systems Frontiers*, 9(2-3):245–256, July 2007.
- Edith Law and Luis von Ahn. *Human Computation*. Morgan & Claypool Publishers, 2011.
- Beatrice Liem, Haoqi Zhang, and Yiling Chen. An iterative dual pathway structure for speech-to-text transcription. In *HCOMP '11: The 3rd Human Computation Workshop*, 2011.
- G. Little, L.B. Chilton, M. Goldman, and R.C. Miller. Exploring iterative and parallel human computation processes. In *Proceedings of the ACM SIGKDD workshop on human computation*, pages 68–76, 2010.
- Greg Little, Lydia B. Chilton, Max Goldman, and Robert C. Miller. TurkIt: tools for iterative tasks on mechanical turk. In *HCOMP '09: Proceedings of the ACM SIGKDD Workshop on Human Computation*, pages 29–30, New York, NY, USA, 2009. ACM. ISBN 978-1-60558-672-4. doi: <http://doi.acm.org/10.1145/1600150.1600159>.
- Bo Long, Olivier Chapelle, Ya Zhang, Yi Chang, Zhaohui Zheng, and Belle Tseng. Active learning for ranking through expected loss optimization. In *Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval*, SIGIR '10, pages 267–274, New York, NY, USA, 2010. ACM. ISBN 978-1-4503-0153-4. doi: 10.1145/1835449.1835495. URL <http://doi.acm.org/10.1145/1835449.1835495>.

- R Luce. Individual choice behavior: A theoretical analysis. *books.google.com*, Jan 2005. URL http://books.google.com/books?hl=en&lr=&id=ERQsKkPiKkkC&oi=fnd&pg=PP1&dq=individual+choice+behavior:+a+theoretical+analysis&ots=2gpqZzZfdg&sig=6Xu_N7IAjPkXYsmVs7jNw78PLes.
- Andrew Mao, Yiling Chen, Krzysztof Z. Gajos, David Parkes, Ariel D. Procaccia, and Haoqi Zhang. TurkServer: Enabling Synchronous and Longitudinal Online Experiments. In *Proceedings of the 4th Workshop on Human Computation (HCOMP'12)*, 2012.
- Andrew Mao, Ariel D. Procaccia, and Yiling Chen. Better Human Computation Through Principled Voting. In *Proceedings of the 27th Conference on Artificial Intelligence (AAAI'13)*, 2013.
- Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, June 1995. ISBN 0195073401. URL <http://www.worldcat.org/isbn/0195073401>.
- W. Mason and S. Suri. Conducting behavioral research on Amazon's Mechanical Turk. *Behavior Research Methods*, 44(1):1–23, 2012.
- P. McCullagh and J. A. Nelder, editors. *Generalized Linear Models*. Chapman and Hall/CRC, Boca Raton, FL, USA, 1989.
- D McFadden. *Conditional logit analysis of qualitative choice behavior*, volume 1, pages 105–142. Academic Press, 1974. URL <http://elsa.berkeley.edu/pub/reprints/mcfadden/zarembka.pdf>.
- Nolan Miller, Paul Resnick, and Richard Zeckhauser. Eliciting Informative Feedback: The Peer-Prediction Method. *MANAGEMENT SCIENCE*, 51(9):1359–1373, September 2005. doi: 10.1287/mnsc.1050.0379. URL <http://dx.doi.org/10.1287/mnsc.1050.0379>.
- Frederick Mosteller. Remarks on the method of paired comparisons: I. the least squares solution assuming equal standard deviations and equal correlations. *Psychometrika*, 16(1):3–9, March 1951. URL <http://ideas.repec.org/a/spr/psycho/v16y1951i1p3-9.html>.
- Evdokia Nikolova and Rahul Sami. A strategic model for information markets. In *EC '07: Proceedings of the 8th ACM conference on Electronic commerce*, pages 316–325, New York, NY, USA, 2007. ACM. ISBN 978-1-59593-653-0. doi: <http://doi.acm.org/10.1145/1250910.1250956>.
- Jon Noronha, Eric Hysen, Haoqi Zhang, and Krzysztof Z. Gajos. Platemate: Crowdsourcing nutrition analysis from food photographs. In *Proceedings of the 24th annual ACM symposium on User interface software and technology*, UIST '11, pages 1–12, New York, NY, USA, 2011. ACM. ISBN 978-1-4503-0716-1.
- Charles Noussair and Marc Willinger. Mixed strategies in an unprofitable game: an experiment. Technical report, 2011.

- Michael Ostrovsky. Information aggregation in dynamic markets with strategic traders. *Econometrica*, 2011.
- Michael Ostrovsky. Information aggregation in dynamic markets with strategic traders. *Econometrica*, 80(6):2595–2647, November 2012.
- Abraham Othman and Tuomas Sandholm. Decision rules and decision markets. *Proc. of the 9th Int. Conf. on Autonomous Agents and Multiagent Systems*, pages 625–632, 2010.
- David Pennock and Lirong Xia. Price updating in combinatorial prediction markets with bayesian networks. In *Proceedings of the Twenty-Seventh Conference Annual Conference on Uncertainty in Artificial Intelligence (UAI-11)*, pages 581–588, Corvallis, Oregon, 2011. AUAI Press.
- R Plackett. The analysis of permutations. *Applied Statistics*, Jan 1975. URL <http://www.jstor.org/stable/10.2307/2346567>.
- Drazen Prelec. A Bayesian truth serum for subjective data. *Science*, 306(5695):462–466, October 2004.
- Drazen Prelec and HS Seung. An algorithm that finds truth even if most people are wrong. 2006.
- Lawrence R Rabiner. A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, 1989.
- Goran Radanovic and Boi Faltings. A Robust Bayesian Truth Serum for Non-binary Signals. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI’13)*, pages 833–839, 2013.
- David G Rand. The promise of Mechanical Turk: How online labor markets can help theorists run behavioral experiments. *Journal of Theoretical Biology*, 299(C):172–179, April 2012.
- Paul W. Rhode and Koleman S. Strumpf. Historical presidential betting markets. *Journal of Economic Perspectives*, 18(2):127–142, 2004.
- Paul W. Rhode and Koleman S. Strumpf. Manipulating political stock markets: A field experiment and a century of observational data. Working Paper, 2007.
- Leonard J Savage. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association*, 66(336):783–801, 1971.
- Gideon Schwarz. Estimating the Dimension of a Model. *The Annals of Statistics*, 6(2): 461–464, 1978. ISSN 00905364. doi: 10.2307/2958889. URL <http://dx.doi.org/10.2307/2958889>.
- Reinhard Selten and Thorsten Chmura. Stationary concepts for experimental 2x2-games. *The American Economic Review*, 98(3):938–966, 2008.

- Burr Settles. Active learning literature survey. Computer Sciences Technical Report 1648, University of Wisconsin–Madison, 2009.
- Jason Shachat, J Swarthout, and Lijia Wei. A hidden Markov model for the detection of pure and mixed strategy play in games. *Available at SSRN*, 2012.
- AD Shaw, JJ Horton, and DL Chen. Designing incentives for inexpert human raters. *Proceedings of the ACM 2011 conference on computer supported cooperative work*, pages 275–284, 2011.
- Peng Shi, Vincent Conitzer, and Mingyu Guo. Prediction mechanisms that do not incentivize undesirable actions. *WINE’09: Internet and Network Economics*, pages 89–100, 2009.
- Michael Spence. Job market signalling. *Quarterly Journal of Economics*, 87(3):355–374, August 1973.
- L. L. Thurstone. A law of comparative judgement. *Psychological Review*, 34:273–286, 1927.
- Andrew Viterbi. Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *Information Theory, IEEE Transactions on*, 13(2):260–269, 1967.
- Bo Waggoner and Yiling Chen. Information elicitation sans verification. In *Proceedings of the 3rd Workshop on Social Computing and User Generated Content (SC13)*, 2013.
- Robert L Winkler. Scoring rules and the evaluation of probability assessors. *Journal of the American Statistical Association*, 64(327):1073–1078, 1969.
- Jens Witkowski and David C. Parkes. A Robust Bayesian Truth Serum for Small Populations. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI ’12)*, 2012a.
- Jens Witkowski and David C Parkes. Peer Prediction without a Common Prior. *EC 2012*, pages 1–18, April 2012b.
- Jens Witkowski, Yoram Bachrach, Peter Key, and David C. Parkes. Dwelling on the Negative: Incentivizing Effort in Peer Prediction. In *Proceedings of the 1st AAAI Conference on Human Computation and Crowdsourcing (HCOMP’13)*, 2013.
- Justin Wolfers and Eric Zitzewitz. Prediction markets. *Journal of Economic Perspective*, 18(2):107–126, 2004.
- L Xia and D.M Pennock. An efficient monte-carlo algorithm for pricing combinatorial prediction markets for tournaments. *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI), Barcelona, Catalonia, Spain*, 2011.
- Peter Zhang and Yiling Chen. Elicitability and knowledge-free elicitation with peer prediction. In *AAMAS. International Foundation for Autonomous Agents and Multiagent Systems*, May 2014.

Appendix A

Appendix to Chapter 3

A.1 Omitted Proofs

A.1.1 Proof of Theorem 2

Proof. The technique used in this proof is analogous to that of Theorem 2 in Chen et al. [2010b].

Let σ be Alice's first-stage strategy at a PBE of the 3-stage I game. By Lemma 1, σ must satisfy the consistency condition. At any PBE of the 3-stage I game, for a fixed prior distribution and a fixed initial market probability, the total of Alice's ex-ante expected payoff and Bob's ex-ante expected payoff in the game is a constant. Therefore, Alice seeks to choose a first-stage strategy in order to minimize Bob's ex-ante expected payoff. We will show that σ must dictate Alice to change the market probability to the prior probability regardless of Alice's realized signal.

We first argue that σ must be a deterministic strategy, i.e. there exists a unique $r \in [0, 1]$ such that $\sigma_{s_A}(r) = 1$ for any realized signal s_A for Alice. We prove this by contradiction. Suppose that there exists a realized signal s_A such that the support of strategy σ for signal s_A has at least 2 points, r_1 , r_2 , and perhaps a set of other points R . Then we construct another strategy σ' for Alice and show that Bob's expected payoff when Alice uses the strategy σ' is

less than his expected payoff when Alice uses the strategy σ , assuming that Bob knows and conditions on Alice's first-stage strategy.

Let $r_3 = \frac{r_1+r_2}{2}$ be the midpoint of r_1 and r_2 . Let the new strategy σ' for Alice randomize over r_1 , r_3 , and the same set of remaining points R . Under σ' , the probability that Alice receives signal s_A and reports r_1 is $\frac{P(r_1)-P(r_2)}{P(r_1)}P(s_A, r_1)$, and the probability that Alice receives signal s_A and reports r_3 is $\frac{P(r_2)}{P(r_1)}P(s_A, r_1) + P(s_A, r_2)$. Under strategy σ' , Alice mixes between reporting r_1 and r_3 with probability $P(r_1) - P(r_2)$ and $2P(r_2)$ respectively. For this strategy σ' , we can compute $P(s_A|r_3)$ as follows.

$$P(s_A|r_3) = \frac{P(s_A, r_3)}{P(r_3)} = \frac{\frac{P(r_2)}{P(r_1)}P(s_A, r_1) + P(s_A, r_2)}{2P(r_2)} = \frac{1}{2}P(s_A|r_1) + \frac{1}{2}P(s_A|r_2)$$

Note that $x_{s_B}(r_3)$ has the following relationship with $x_{s_B}(r_1)$ and $x_{s_B}(r_2)$ as shown below.

$$\begin{aligned} x_{s_B}(r_3) &= \sum_{s_A} P(s_A|r_3)P(1|s_A, s_B) = \sum_{s_A} \left(\frac{1}{2}P(s_A|r_1) + \frac{1}{2}P(s_A|r_2) \right) P(1|s_A, s_B) \\ &= \frac{x_{s_B}(r_1) + x_{s_B}(r_2)}{2} \end{aligned} \tag{A.1}$$

Let $\pi^B(\sigma)$ denote Bob's ex-ante expected payoff when Alice uses strategy σ and Bob knows and conditions on Alice using the strategy σ . We derive the expression for $\pi^B(\sigma)$

below.

$$\begin{aligned}
\pi^B(\sigma) &= P(r_1) \sum_{s_B} P(s_B|r_1) \left\{ x_{s_B}(r_1) \log \frac{x_{s_B}(r_1)}{r_1} + (1 - x_{s_B}(r_1)) \log \frac{1 - x_{s_B}(r_1)}{1 - r_1} \right\} \\
&\quad + P(r_2) \sum_{s_B} P(s_B|r_2) \left\{ x_{s_B}(r_2) \log \frac{x_{s_B}(r_2)}{r_2} + (1 - x_{s_B}(r_2)) \log \frac{1 - x_{s_B}(r_2)}{1 - r_2} \right\} \\
&\quad + \text{remaining profit over R} \\
&= P(r_1) \sum_{s_B} P(s_B) \left\{ x_{s_B}(r_1) \log \frac{x_{s_B}(r_1)}{\sum_{s_B} x_{s_B}(r_1)} + (1 - x_{s_B}(r_1)) \log \frac{1 - x_{s_B}(r_1)}{1 - \sum_{s_B} x_{s_B}(r_1)} \right\} \\
&\quad + P(r_2) \sum_{s_B} P(s_B) \left\{ x_{s_B}(r_2) \log \frac{x_{s_B}(r_2)}{\sum_{s_B} x_{s_B}(r_2)} + (1 - x_{s_B}(r_2)) \log \frac{1 - x_{s_B}(r_2)}{1 - \sum_{s_B} x_{s_B}(r_2)} \right\} \\
&\quad + \text{remaining profit over R} \\
&= (P(r_1) - P(r_2)) \sum_{s_B} P(s_B) \left\{ x_{s_B}(r_1) \log \frac{x_{s_B}(r_1)}{\sum_{s_B} x_{s_B}(r_1)} + (1 - x_{s_B}(r_1)) \log \frac{1 - x_{s_B}(r_1)}{1 - \sum_{s_B} x_{s_B}(r_1)} \right\} \\
&\quad + P(r_2) \sum_{s_B} P(s_B) \left\{ x_{s_B}(r_1) \log \frac{x_{s_B}(r_1)}{\sum_{s_B} x_{s_B}(r_1)} + (1 - x_{s_B}(r_1)) \log \frac{1 - x_{s_B}(r_1)}{1 - \sum_{s_B} x_{s_B}(r_1)} \right\} \\
&\quad + P(r_2) \sum_{s_B} P(s_B) \left\{ x_{s_B}(r_2) \log \frac{x_{s_B}(r_2)}{\sum_{s_B} x_{s_B}(r_2)} + (1 - x_{s_B}(r_2)) \log \frac{1 - x_{s_B}(r_2)}{1 - \sum_{s_B} x_{s_B}(r_2)} \right\} \\
&\quad + \text{remaining profit over R}
\end{aligned}$$

When Alice uses strategy σ' , Bob's ex-ante expected payoff $\pi^B(\sigma')$ is less than his ex-ante

expected payoff $\pi^B(\sigma)$, as shown below.

$$\begin{aligned}
\pi^B(\sigma') &= (P(r_1) - P(r_2)) \sum_{s_B} P(s_B) \left\{ x_{s_B}(r_1) \log \frac{x_{s_B}(r_1)}{\sum_{s_B} x_{s_B}(r_1)} + (1 - x_{s_B}(r_1)) \log \frac{1 - x_{s_B}(r_1)}{1 - \sum_{s_B} x_{s_B}(r_1)} \right\} \\
&\quad + 2P(r_2) \sum_{s_B} P(s_B) \left\{ x_{s_B}(r_3) \log \frac{x_{s_B}(r_3)}{\sum_{s_B} x_{s_B}(r_3)} + (1 - x_{s_B}(r_3)) \log \frac{1 - x_{s_B}(r_3)}{1 - \sum_{s_B} x_{s_B}(r_3)} \right\} \\
&\quad + \text{remaining profit over R} \\
&= (P(r_1) - P(r_2)) \sum_{s_B} P(s_B) \left\{ x_{s_B}(r_1) \log \frac{x_{s_B}(r_1)}{\sum_{s_B} x_{s_B}(r_1)} + (1 - x_{s_B}(r_1)) \log \frac{1 - x_{s_B}(r_1)}{1 - \sum_{s_B} x_{s_B}(r_1)} \right\} \\
&\quad + 2P(r_2) \sum_{s_B} P(s_B) \left\{ x_{s_B}(r_3) \log \frac{\frac{x_{s_B}(r_1) + x_{s_B}(r_2)}{2}}{\sum_{s_B} \frac{x_{s_B}(r_1) + x_{s_B}(r_2)}{2}} \right. \\
&\quad \quad \left. + \left(1 - \frac{x_{s_B}(r_1) + x_{s_B}(r_2)}{2} \right) \log \frac{1 - \frac{x_{s_B}(r_1) + x_{s_B}(r_2)}{2}}{1 - \sum_{s_B} \frac{x_{s_B}(r_1) + x_{s_B}(r_2)}{2}} \right\} \\
&\quad + \text{remaining profit over R} \tag{A.2}
\end{aligned}$$

$$\begin{aligned}
&< (P(r_1) - P(r_2)) \sum_{s_B} P(s_B) \left\{ x_{s_B}(r_1) \log \frac{x_{s_B}(r_1)}{\sum_{s_B} x_{s_B}(r_1)} + (1 - x_{s_B}(r_1)) \log \frac{1 - x_{s_B}(r_1)}{1 - \sum_{s_B} x_{s_B}(r_1)} \right\} \\
&\quad + P(r_2) \sum_{s_B} P(s_B) \left\{ x_{s_B}(r_1) \log \frac{x_{s_B}(r_1)}{\sum_{s_B} x_{s_B}(r_1)} + (1 - x_{s_B}(r_1)) \log \frac{1 - x_{s_B}(r_1)}{1 - \sum_{s_B} x_{s_B}(r_1)} \right\} \\
&\quad + P(r_2) \sum_{s_B} P(s_B) \left\{ x_{s_B}(r_2) \log \frac{x_{s_B}(r_2)}{\sum_{s_B} x_{s_B}(r_2)} + (1 - x_{s_B}(r_2)) \log \frac{1 - x_{s_B}(r_2)}{1 - \sum_{s_B} x_{s_B}(r_2)} \right\} \\
&\quad + \text{remaining profit over R} \tag{A.3} \\
&= \pi^B(\sigma)
\end{aligned}$$

where equation (A.2) follows from equation (A.1), and the inequality (A.3) follows from the strict convexity of relative entropy when the signals satisfy the informativeness condition.

Therefore, for any Alice's strategy σ where for at least one realized signal the support of the strategy has two or more points in its support (i.e. deterministic), there always exists a strategy σ' such that $\pi^B(\sigma') < \pi^B(\sigma)$. This means that, at any PBE of this game, Alice's first-stage strategy must have only one point in its support. Such a strategy for Alice does not reveal any information to Bob. If Alice's first-stage strategy is deterministic, we must have $\sigma_{s_A}(r) = 1$ for all $s_A \in S_A$ and for some $r \in [0, 1]$. Then, by the consistency condition,

the only point in the support of Alice's strategy must be $P(1)$, as shown below.

$$\begin{aligned}
P(1|r, \sigma) &= r \\
\Rightarrow \frac{\sum_{s_A} P(1|s_A) \sigma_{s_A}(r) P(s_A)}{\sum_{s_A} P(s_A) \sigma_{s_A}(r)} &= r \\
\Rightarrow \frac{\sum_{s_A} P(1|s_A) P(s_A)}{\sum_{s_A} P(s_A)} &= r \\
\Rightarrow r &= \sum_{s_A} P(1|s_A) P(s_A) = P(1)
\end{aligned}$$

Therefore, at any PBE of the 3-stage I game, Alice's strategy must be $\sigma_{s_A}(P(1)) = 1, \forall s_A$. $\square$

A.1.2 Proof of Theorem 3

Proof. According to the theorem statement, Alice's first-stage strategy is

$$\sigma_{s_A}(P(1)) = 1, \forall s_A \in S_A$$

and Bob's second-stage strategy is

$$x_{s_B}(r_A) = \begin{cases} f_{s_B}(\alpha_{s_B}^{\min}), & r_A \in [0, \alpha_{s_B}^{\min}) \\ f_{s_B}(r_A), & r_A \in [\alpha_{s_B}^{\min}, \alpha_{s_B}^{\max}], \forall s_B \in S_B \\ f_{s_B}(\alpha_{s_B}^{\max}), & r_A \in (\alpha_{s_B}^{\max}, 1] \end{cases}$$

where

$$\begin{aligned}
f_{s_B}(r_A) &= \frac{P(1|s_B)P(0)r_A}{P(1)P(0|s_B) + (P(1|s_B) - P(1))r_A} \\
\beta_{s_B}^{\min} &= \min_{s_A} P(1|s_A, s_B), \beta_{s_B}^{\max} = \max_{s_A} P(1|s_A, s_B) \\
\alpha_{s_B}^{\min} &= f_{s_B}^{-1}(\beta_{s_B}^{\min}), \alpha_{s_B}^{\max} = f_{s_B}^{-1}(\beta_{s_B}^{\max})
\end{aligned}$$

To prove that Alice's and Bob's strategies form a PBE of the 3-stage I game, we need to

show 2 things:

1. Bob's strategy is valid. That is, $\forall s_B, \forall r_A \in [0, 1], x_{s_B}(r_A) \in [\beta_{s_B}^{\min}, \beta_{s_B}^{\max}]$;
2. Alice's expected payoff $u_{s_A}(r_A)$ after receiving any signal s_A is uniquely maximized by reporting $r_A = P(1)$ given Bob's strategy.

1. We first show that Bob's strategy is valid. First, $f_{s_B}(r_A)$ is monotonically increasing in $r_A \in [0, 1]$ since

$$\frac{df_{s_B}(r_A)}{dr_A} = \frac{P(1)(1 - P(1))P(1|s_B)(1 - P(1|s_B))}{\{P(1)P(0|s_B) + (P(1|s_B) - P(1))r_A\}^2} > 0$$

Second, the domain of $x_{s_B}(r_A)$ is well-defined since

$$\beta_{s_B}^{\min} < \beta_{s_B}^{\max} \Rightarrow \alpha_{s_B}^{\min} < \alpha_{s_B}^{\max}.$$

Finally, Bob's strategy is valid since

$$\beta_{s_B}^{\min} = f_{s_B}(\alpha_{s_B}^{\min}) \leq x_{s_B}(r_A) \leq f_{s_B}(\alpha_{s_B}^{\max}) = \beta_{s_B}^{\max}, \forall r_A \in [0, 1].$$

2. We now show that Alice's expected payoff after receiving any signal s_A is uniquely maximized by reporting $r_A = P(1)$.

We divide the range $[0, 1]$ of r_A into 3 subsets and analyze the properties of $u_{s_A}(r_A)$ on these subsets.

- (a) $r_A \in [\max_{s_B}\{\alpha_{s_B}^{\min}\}, \min_{s_B}\{\alpha_{s_B}^{\max}\}]$;
- (b) $r_A \in [0, \min_{s_B}\{\alpha_{s_B}^{\min}\}) \cup (\max_{s_B}\{\alpha_{s_B}^{\max}\}, 1]$;
- (c) $r_A \in [\min_{s_B}\{\alpha_{s_B}^{\min}\}, \max_{s_B}\{\alpha_{s_B}^{\min}\}) \cup (\min_{s_B}\{\alpha_{s_B}^{\max}\}, \max_{s_B}\{\alpha_{s_B}^{\max}\}]$;

These subsets are well-defined as long as $\max_{s_B}\{\alpha_{s_B}^{\min}\} \leq \min_{s_B}\{\alpha_{s_B}^{\max}\}$, as proven

below.

$$\begin{aligned}
P(1|s_B) &= \sum_{s_A} P(1|s_A, s_B)P(s_A|s_B) \leq \sum_{s_A} \beta_{s_B}^{\max} P(s_A|s_B) = \beta_{s_B}^{\max}, \forall s_B \in S_B \\
P(1|s_B) &= \sum_{s_A} P(1|s_A, s_B)P(s_A|s_B) \geq \sum_{s_A} \beta_{s_B}^{\min} P(s_A|s_B) = \beta_{s_B}^{\min}, \forall s_B \in S_B \\
\Rightarrow \quad \beta_{s_B}^{\min} &\leq P(1|s_B) \leq \beta_{s_B}^{\max}, \forall s_B \in S_B \\
\Rightarrow \quad \alpha_{s_B}^{\min} &\leq P(1) \leq \alpha_{s_B}^{\max}, \forall s_B \in S_B \\
\Rightarrow \quad \max_{s_B} \{\alpha_{s_B}^{\min}\} &\leq P(1) \leq \min_{s_B} \{\alpha_{s_B}^{\max}\}
\end{aligned}$$

(a) For any $r_A \in [\max_{s_B} \{\alpha_{s_B}^{\min}\}, \min_{s_B} \{\alpha_{s_B}^{\max}\}]$, we show that Alice's expected payoff after receiving any signal s_A is uniquely maximized at $r_A = P(1)$. Alice's expected payoff after receiving the s_A signal, denoted by $u_{s_A}(r_A)$, is

$$\begin{aligned}
u_{s_A}(r_A) &= \sum_{s_B} \left\{ P(1, s_B|s_A) \left(\log \frac{r_A}{r^0} + \log \frac{P(1|s_A, s_B)}{x_{s_B}(r_A)} \right) \right. \\
&\quad \left. + P(0, s_B|s_A) \left(\log \frac{1-r_A}{1-r^0} + \log \frac{P(0|s_A, s_B)}{1-x_{s_B}(r_A)} \right) \right\}
\end{aligned}$$

The first derivative of $u_{s_A}(r_A)$ evaluated at $r_A = P(1)$ is zero, as shown below:

$$\frac{du_{s_A}(r_A)}{dr_A} = \frac{\sum_{s_B} x_{s_B}(r_A)P(s_B) - r_A}{r_A(1-r_A)} \Rightarrow \frac{du_{s_A}(r_A)}{dr_A} \Big|_{r_A=P(1)} = 0 \quad (\text{A.4})$$

The second derivative of $u_{s_A}(r_A)$ is negative since $x_{s_B}(r_A) \neq r_A, \forall j$ by the distinguishability condition, as shown below.

$$\frac{d^2 u_{s_A}(r_A)}{dr_A^2} = -\frac{\sum_{s_B} (x_{s_B}(r_A) - r_A)^2 P(s_B)}{r_A^2 (1-r_A)^2} < 0 \quad (\text{A.5})$$

By equations (A.4) and (A.5), for all $r_A \in [\max_{s_B} \{\alpha_{s_B}^{\min}\}, \min_{s_B} \{\alpha_{s_B}^{\max}\}]$, $u_{s_A}(r_A)$ is uniquely maximized at $r_A = P(1)$.

The above argument applies for any $r_A \in (0, 1)$ as long as $x_{s_B}(r_A) = f_{s_B}(r_A)$.

For the rest of the proof, we make use of the following inequality, derived using

equation (A.4).

$$r_A < \sum_{s_B} P(s_B) f_{s_B}(r_A), \forall r_A \in (0, P(1)) \quad (\text{A.6})$$

(b) Next, we show that $u_{s_A}(r_A)$ is monotonically increasing for all $r_A \in [0, \min_{s_B} \{\alpha_{s_B}^{\min}\})$.

We omit the symmetric argument showing that $u_{s_A}(r_A)$ is monotonically decreasing for all $r_A \in (\max_{s_B} \{\alpha_{s_B}^{\max}\}, 1]$.

We define $\gamma^{\min} = \min_{s_A} \{P(1|s_A)\}$. First, we note that $\gamma^{\min} \in (0, P(1))$ since

$$P(1) = \sum_i P(s_A) P(1|s_A) > \min_{s_A} \{P(1|s_A)\} = \gamma^{\min}$$

We first prove that $\min_{s_B} \{\alpha_{s_B}^{\min}\} \leq \gamma^{\min}$. For contradiction, assume $\alpha_{s_B}^{\min} > \gamma^{\min}, \forall j$ and let $\gamma^{\min} = P(1|a_t)$ for some a_t . Then we have

$$\begin{aligned} \alpha_{s_B}^{\min} > \gamma^{\min}, \forall j &\Rightarrow \beta_{s_B}^{\min} > f_{s_B}(\gamma^{\min}), \forall j \\ \Rightarrow \gamma^{\min} = P(1|a_t) &= \sum_{s_B} P(s_B) P(1|a_t, s_B) \geq \sum_{s_B} P(s_B) \beta_{s_B}^{\min} > \sum_{s_B} P(s_B) f_{s_B}(\gamma^{\min}) \end{aligned}$$

which contradicts inequality (A.6) for $r_A = \gamma^{\min}$.

For $r_A < \min_{s_B} \{\alpha_{s_B}^{\min}\}$, $x_{s_B}(r_A)$ is a constant for all j . Thus, $u_{s_A}(r_A)$ is monotonically increasing for $r_A \in [0, \min_{s_B} \{\alpha_{s_B}^{\min}\})$. since the first derivative of $u_{s_A}(r_A)$ with respect to r_A is positive, as shown below.

$$\frac{du_{s_A}(r_A)}{dr_A} = \frac{P(1|s_A) - r_A}{r_A(1 - r_A)} \geq \frac{\gamma^{\min} - r_A}{r_A(1 - r_A)} \geq \frac{\min_{s_B} \{\alpha_{s_B}^{\min}\} - r_A}{r_A(1 - r_A)} > 0$$

(c) Finally, we show that $u_{s_A}(r_A)$ is monotonically increasing for all $r_A \in [\min_{s_B} \alpha_{s_B}^{\min}, \max_{s_B} \alpha_{s_B}^{\min})$.

We omit the symmetric argument showing that $u_{s_A}(r_A)$ is monotonically decreasing for all $r_A \in (\min_{s_B} \{\alpha_{s_B}^{\max}\}, \max_{s_B} \{\alpha_{s_B}^{\max}\}]$. For the following argument, let $s_{B,j}$ be the j -th realized signal for Bob. Without loss of generality, we assume that $\alpha_{s_{B,j}}^{\min}$ follows the increasing order, i.e. $\alpha_{s_{B,1}}^{\min} \leq \dots \leq \alpha_{s_{B,n_B}}^{\min}$. For $k \in \{0, \dots, n_B - 2\}$,

if $r_A \in [\alpha_{s_{B,k}}^{\min}, \alpha_{s_{B,k+1}}^{\min})$, then we have

$$\begin{aligned} x_{s_{B,j}}(r_A) &= f_{s_{B,j}}(r_A), \forall j = 0, \dots, k \\ x_{s_{B,j}}(r_A) &= f_{s_{B,j}}(\alpha_{s_{B,j}}^{\min}) = \beta_{s_{B,j}}^{\min}, \forall j = k+1, \dots, n_B - 1 \end{aligned}$$

We show below that the first derivative of $u_{s_A}(r_A)$ is positive.

$$\begin{aligned} \frac{du_{s_A}(r_A)}{dr_A} &= \frac{P(1|s_A) - r_A + \sum_{j=1}^k \{P(s_{B,j})f_{s_{B,j}}(r_A) - P(1, s_{B,j}|s_A)\}}{r_A(1 - r_A)} \\ &= \frac{\sum_{j=k+1}^n P(1, s_{B,j}|s_A) + \sum_{j=1}^k P(s_{B,j})f_{s_{B,j}}(r_A) - r_A}{r_A(1 - r_A)} \\ &= \frac{\sum_{j=k+1}^n \{P(1, s_{B,j}|s_A) - P(s_{B,j})f_{s_{B,j}}(r_A)\}}{r_A(1 - r_A)} + \frac{\sum_{j=1}^k P(s_{B,j})f_{s_{B,j}}(r_A) - r_A}{r_A(1 - r_A)} \\ &> \frac{\sum_{j=k+1}^n \{P(1, s_{B,j}|s_A) - P(s_{B,j})\beta_{s_{B,j}}^{\min}\}}{r_A(1 - r_A)} \tag{A.7} \\ &\geq \frac{\sum_{j=k+1}^n \{P(1, s_{B,j}|s_A) - P(s_{B,j})P(1|s_A, s_{B,j})\}}{r_A(1 - r_A)} \\ &= 0 \end{aligned}$$

where inequality (A.7) was derived by inequality (A.6). Hence, $u_{s_A}(r_A)$ is monotonically increasing for all $r_A \in [0, \min_{s_{B,j}} \{\alpha_{s_{B,j}}^{\min}\})$.

In conclusion, Bob's strategy is valid and Alice's expected payoff $u_{s_A}(r)$ is uniquely maximized at $r_A = P(1)$. Therefore, the specified strategies for Alice and Bob form a PBE of the 3-stage I game. $\square$

A.1.3 Proof of Lemma 2

Proof. At a PBE of the finite-stage I game, suppose that r^k and σ^k are the report and the strategy of the player in stage k , and suppose that for a particular k , the consistency condition given in equation (A.8) below is violated.

$$P(1|r^1, \dots, r^k, \sigma^1, \dots, \sigma^k) = r^k \tag{A.8}$$

Then we construct a perturbed strategy $\hat{\sigma}^k$ satisfying the consistency condition, and we show that the player's expected payoff by using the perturbed strategy $\hat{\sigma}^k$ is greater than her expected payoff by using the original strategy σ^k .

To construct the perturbed strategy $\hat{\sigma}^k$, we start by setting $\hat{\sigma}^k = \sigma^k$. Let $x \in [0, 1]$ be a point in the support of strategy σ^k such that the consistency condition fails for x , i.e. $P(\omega = 1 | r^1, \dots, r^{k-1}, x, \sigma^1, \dots, \sigma^k) \neq x$. Let $\hat{x} = P(\omega = 1 | r^1, \dots, r^{k-1}, x, \sigma^1, \dots, \sigma^k)$. Then, whenever the strategy σ^k dictates that the player change the market probability to x , let the strategy $\hat{\sigma}^k$ dictate that the player change the market probability to $\hat{x}$. We repeat this perturbation for each x in the support of strategy σ^k such that $x \neq \hat{x}$. By using this perturbation, the strategy $\hat{\sigma}^k$ satisfies the consistency condition.

Next, we show that the player's expected payoff by using the perturbed strategy $\hat{\sigma}^k$ is greater than her expected payoff by using her original strategy σ^k . Let x_k and $\hat{x}_k$ be the random variables that correspond to the values that the player of stage k the market probability to, and let x and $\hat{x}$ be their realizations. Note that any x has a corresponding value of $\hat{x}$, so we may write expressions like $\sum_x \hat{x}$ in which $\hat{x}$ is implicitly indexed by x .

The difference between the player's expected payoff by using strategy $\hat{\sigma}^k$ and σ^k is

$$\begin{aligned}
& \sum_{z, r^1, \dots, r^{k-1}, x} P(z, r^1, \dots, r^{k-1}, x) (\log P(z | r^1, \dots, r^{k-1}, \hat{x}) - \log x) \\
&= \sum_{r^1, \dots, r^{k-1}, x} P(r^1, \dots, r^{k-1}, x) \sum_z P(z | r^1, \dots, r^{k-1}, x) (\log \hat{x} - \log x) \\
&= \sum_{r^1, \dots, r^{k-1}, x} P(r^1, \dots, r^{k-1}, x) \sum_z \hat{x} (\log \hat{x} - \log x) \\
&= \sum_{r^1, \dots, r^{k-1}, x} P(r^1, \dots, r^{k-1}, x) D(\mathbf{p}_{(\hat{x}_k)} || \mathbf{p}_{(x_k)})
\end{aligned}$$

where $\mathbf{p}_{(\hat{x}_k)}$ and $\mathbf{p}_{(x_k)}$ are the probability distributions of $\hat{x}_k$ and x_k respectively. $D(\mathbf{p}_{(\hat{x}_k)} || \mathbf{p}_{(x_k)})$ is relative entropy, which is nonnegative and strictly positive when the two distributions are not the same. Since σ^k does not satisfy the consistency condition, there is at least one x such that $P(x) > 0$ and $\hat{x} = x$. Thus we have $D(\mathbf{p}_{(\hat{x}_k)} || \mathbf{p}_{(x_k)}) > 0$, and this contradicts

our assumption that σ^k is an PBE strategy for the player of stage k of the finite-stage I game. □

A.1.4 Proof of Lemma 3

Proof. Recall that m is the last player of the game, and t_m is the last stage of the game. Also, stage k is the second to last stage of participation for player m ($k < t_m$). Consider the part of the finite-stage I game starting from stage k to stage t_m . There must exist at least one player $j < m$ whose last stage of participation is between stage k and stage t_m . We combine the players participating after stage k and before stage t_m as one composite player, and also combine their signals to be one composite signal. Because all the signals are independent, the signal of this composite player is also independent of the signal of player m . Therefore, we can treat the part of the finite-stage I game from stage k to stage t_m as a 3-stage I game where player m is Alice and the composite player is Bob. By the distinguishability condition, at every PBE, information is fully aggregated at the same of the finite-stage I game. Thus, at any PBE of this 3-stage I game, the total expected payoff of players is constant given the market estimate at the beginning of stage k and the prior distribution. Thus, player m seeks to minimize the total expected payoff of the composite player. By Theorems 2 and 3, there exists a PBE of this 3-stage I game. At any PBE of this game, in stage k , player m changes the market probabilities to the prior probability of the event at the beginning of stage k . Since player m is a Bayesian agent, he can condition his belief of the probability of the event on the strategies and the reports of all participants in the previous stages. Thus, at the beginning of stage k , player m believes the prior probabilities of the event to be

$$P(1|r^1, \dots, r^{k-1}, \sigma^1, \dots, \sigma^{k-1}).$$

where $r^1, \dots, r^{k-1}$ and $\sigma^1, \dots, \sigma^{k-1}$ are the reports and the strategies of the participants in the first $k-1$ stages. By Lemma 2, at any PBE, the strategies and reports of all participants

must satisfy the consistency condition. Thus, we must have

$$P(1|r^1, \dots, r^{k-1}, \sigma^1, \dots, \sigma^{k-1}) = r^{k-1}.$$

Therefore, player m 's report r^{k-1} in stage k is equal to the market estimate immediately before stage k . This means that player m does not change the market estimate in stage k of the game at any PBE. $\square$

A.1.5 Proof of Theorem 4

Proof. First, we exclude degenerate cases by assuming that, if a player participates in any number of consecutive stages in this game, then these stages are combined into one stage for the player. This does not affect the players' strategic behaviors in this game because the player's total payoff in these consecutive stages only depends on the market estimate at the beginning of the first stage in this sequence, the market estimate at the end of the last stage in this sequence, and the realized outcome of the event.

By Lemma 2, at any PBE of this game, the strategy of each participant must satisfy the consistency condition.

PBE Strategy of Player m

We first consider player m , who is also the last participant of the game. Stage t_m must be the last stage of the game. By properties of LMSR, player m truthfully reveals his realized signal in stage t_m .

Let t^* denote the second to last stage of participation for player m . Consider the game starting from stage t^* to stage t_m . By Lemma 3, player m does not change the market probability in stage t^* . Let t^* denote the *new* second to last stage of participation for player m . Consider the game starting from stage t^* to stage t_m . Player m does not participate in any stage in between stages t^* and t_m . By Lemma 3, player m does not change the market estimate in stage t^* . Inferring recursively, we can show that, in any stage from stage 1 to

stage $t_m - 1$ in which player m is scheduled to participate, player m does not change the market estimate in any of these stages.

In summary, from stage 1 to stage $t_m - 1$, player m does not participate in the game. In stage t_m , player m truthfully reveals his private signal.

PBE Strategy of Player i , $2 \leq i \leq m - 1$

Consider player $m - 1$. By properties of the LMSR, player $m - 1$ truthfully reveals his signal in stage t_{m-1} .

From stage $t_{m-2} + 1$ to $t_{m-1} - 1$, by previous argument, player m does not participate in any of these stages. Also, by the way in which players are ordered, any player i where $i < m - 1$ already finished their participation in the game by the end of stage t_{m-2} . Thus, player $m - 1$ is the only participant from stage $t_{m-2} + 1$ to stage $t_{m-1} - 1$ in this game. Thus, for these stages, if player $m - 1$ is scheduled to participate, he may use any strategy as long as the strategy satisfies the consistency condition.

Next, consider stage 1 to stage t_{m-2} . Since player $m - 1$ is the only participant from stage $t_{m-2} + 1$ to stage $t_{m-1} - 1$, we can combine these stages as stage t^{**} and call it the *new* last stage of participation for player $m - 1$. Let stage t^* be the *new* second to last stage of participation for player $m - 1$. Note that we must have $k < t_{m-2}$. Consider the game from stage t^* to stage t^{**} . By Lemma 3, player $m - 1$ does not change the market estimate in stage t^* . Inferring recursively, we can show that, for any stage before t_{m-2} in which player $m - 1$ is scheduled to participate, player $m - 1$ does not participate in any stage in the game.

Using the same argument, we can summarize the strategy of player i , for any $2 \leq i \leq m - 1$, as follows: From stage 1 to stage $t_{i-1} - 1$, player i does not participate in the game. From stage $t_{i-1} + 1$ to $t_i - 1$, player i uses any strategy that satisfies the consistency condition. In stage t_i , player i truthfully reveals his private information.

PBE Strategy of Player 1

By properties of LMSR, player 1 truthfully reveals his signal in stage t_1 . By our arguments

above, from stage 1 to the stage $t_1 - 1$, none of the other players participates in any stage of the game. Thus, player 1 is the only participant from stage 1 to stage $t_1 - 1$ and he may use any strategy that satisfies the consistency condition. $\square$

A.1.6 Proof of Theorem 5

Proof. This proof has 3 main steps.

1. First, we study the function $u_{a_i}(r)$ for Alice's ex-interim expected payoff at any PBE of the 3-stage market game.

$$u_{a_i}(r) = P(1|a_i) \log \frac{r}{P(1)} + P(0|a_i) \log \frac{1-r}{1-P(1)} \\ + \sum_j \left\{ P(1, s_B|a_i) \log \frac{P(1|a_i, s_B)}{x_{s_B}(r)} + P(0, s_B|a_i) \log \frac{P(0|a_i, s_B)}{1-x_{s_B}(r)} \right\}$$

We prove that $u_{a_i}(r)$ has the following property: For any $r \in [\min_i P(1|a_i), \max_i P(1|a_i)]$, $\frac{u'_{a_i}(r)}{(P(1|a_i)-r)}$ is independent of the value of a_i .

2. Next, by using the above property of $u_{a_i}(r)$, we show that there does not exist a PBE of the 3-stage D game where Alice's strategy satisfies

$$\exists r_1, r_2 \in [\min_i P(1|a_i), \max_i P(1|a_i)], r_1 \neq r_2 \text{ s.t. } \sigma_{a_i}(r_1) > 0, \sigma_{a_i}(r_2) > 0, \forall i = 0, 1$$

Intuitively, this means that the support of Alice's PBE strategy at the 3-stage D game cannot overlap at 2 or more points.

3. Finally, we show that if there exists a PBE of the 3-stage D game, Alice must play one of the three specified strategies described at the PBE.

Step 1: The first derivative of $u_{a_i}(r)$ is

$$u'_{a_i}(r) = \frac{P(1|a_i) - r}{r(1-r)} - \sum_j \left\{ \frac{(P(1, s_B|a_i) - P(s_B|a_i)x_{s_B}(r))x'_{s_B}(r)}{x_{s_B}(r)(1-x_{s_B}(r))} \right\}$$

We would like to compare the expression of $u'_{a_i}(r)$ for $i = 0$ and $i = 1$.

First, we show the expressions of $x_{s_B}(r)$, $1-x_{s_B}(r)$, $x'_{s_B}(r)$ and $P(1, s_B|a_i) - P(s_B|a_i)x_{s_B}(r)$ as follows:

$$\begin{aligned}
x_{s_B}(r) &= \frac{[P(1, a_0|s_B)P(a_1) - P(1, a_1|s_B)P(a_0)]r + [P(1, a_1|s_B)P(1, a_0) - P(1, a_0|s_B)P(1, a_1)]}{[P(a_0|s_B)P(a_1) - P(a_1|s_B)P(a_0)]r + [P(a_1|s_B)P(1, a_0) - P(a_0|s_B)P(1, a_1)]} \\
1 - x_{s_B}(r) &= \frac{[P(0, a_0|s_B)P(a_1) - P(0, a_1|s_B)P(a_0)]r + [P(0, a_1|s_B)P(1, a_0) - P(0, a_0|s_B)P(1, a_1)]}{[P(a_0|s_B)P(a_1) - P(a_1|s_B)P(a_0)]r + [P(a_1|s_B)P(1, a_0) - P(a_0|s_B)P(1, a_1)]} \\
x'_{s_B}(r) &= \frac{(P(1|a_0, s_B) - P(1|a_1, s_B))P(a_0|s_B)P(a_1|s_B)P(a_0)P(a_1)(P(1|a_0) - P(1|a_1))}{[P(a_0|s_B)P(a_1) - P(a_1|s_B)P(a_0)]r + [P(a_1|s_B)P(1, a_0) - P(a_0|s_B)P(1, a_1)]^2} \\
P(1, s_B|a_i) - P(s_B|a_i)x_{s_B}(r) &= (P(1|a_i) - r) \frac{P(s_B)P(a_0|s_B)P(a_1|s_B)(P(1|a_0, s_B) - P(1|a_1, s_B))}{[P(a_0|s_B)P(a_1) - P(a_1|s_B)P(a_0)]r + [P(a_1|s_B)P(1, a_0) - P(a_0|s_B)P(1, a_1)]}
\end{aligned}$$

Note that these expressions have common components in their denominators. To simplify the expression of $u'_{a_i}(r)$, let $nu(f(x))$ and $de(f(x))$ denote the numerator and the denominator of the function $f(x)$ where $f(x)$ is $x_{s_B}(r)$, $1-x_{s_B}(r)$, $x'_{s_B}(r)$ or $P(1, s_B|a_i) - P(s_B|a_i)x_{s_B}(r)$. Notice that:

$$de(x_{s_B}(r)) = de(1 - x_{s_B}(r)) = de(P(1, s_B|a_i) - P(s_B|a_i)x_{s_B}(r)), \{de(x_{s_B}(r))\}^2 = de(x'_{s_B}(r)).$$

Then the expression of $u'_{a_i}(r)$ can be re-written as:

$$\begin{aligned}
u'_{a_i}(r) &= (P(1|a_i) - r) \frac{1}{r(1-r)} - \sum_j \left\{ \frac{(P(1, s_B|a_i) - P(s_B|a_i)x_{s_B}(r)) \frac{nu(x'_{s_B}(r))}{de(x'_{s_B}(r))}}{\frac{nu(x_{s_B}(r))}{de(x_{s_B}(r))} \frac{nu(1-x_{s_B}(r))}{de(1-x_{s_B}(r))}} \right\} \\
&= (P(1|a_i) - r) \left(\frac{1}{r(1-r)} - \sum_j \left\{ \frac{(P(1, s_B|a_i) - P(s_B|a_i)x_{s_B}(r)) nu(x'_{s_B}(r))}{(P(1|a_i) - r) nu(x_{s_B}(r)) nu(1-x_{s_B}(r))} \right\} \right)
\end{aligned}$$

Note that the expressions of $\frac{(P(1, s_B|a_i) - P(s_B|a_i)x_{s_B}(r))}{(P(1|a_i) - r)}$, $nu(x'_{s_B}(r))$, $nu(x_{s_B}(r))$, and $nu(1 - x_{s_B}(r))$ do not depend on the value of a_i . So $\frac{u'_{a_i}(r)}{(P(1|a_i) - r)}$ is not a function of a_i and only a function of r . Thus, $\frac{u'_{a_0}(r)}{(P(1|a_0) - r)}$ and $\frac{u'_{a_1}(r)}{(P(1|a_1) - r)}$ are the same function, and this function is independent of the value of a_i , for any $r \in [P(1|a_0), P(1|a_1)]$.

Step 2: Next, we prove the statement by contradiction. If Alice's PBE strategy in the 3-stage D game satisfies the specified condition, then by definition of a mixed strategy PBE, the following necessary condition must be satisfied:

$$u_{a_i}(r_1) = u_{a_i}(r_2), \forall i = 0, 1$$

In step 1, we showed that the expression of $u'_{a_i}(r)$ can be written as follows:

$$\begin{aligned} u'_{a_i}(r) &= (P(1|a_i) - r) f(r), \forall i = 0, 1 \\ \Rightarrow \int_{-\infty}^r u'_{a_i}(r') dr' &= P(1|a_i) \int_{-\infty}^r f(r') dr' - \int_{-\infty}^r r' f(r') dr', \forall i = 0, 1 \end{aligned} \quad (\text{A.9})$$

For convenience, we define $g(r)$ and $h(r)$ below:

$$g(r) = \int_{-\infty}^r f(r') dr', \quad h(r) = \int_{-\infty}^r r' f(r') dr'$$

From equation (A.9), we have

$$u_{a_i}(r) = P(1|a_i)g(r) - h(r) + C_i, \forall i = 0, 1$$

where C_i for $i = 0, 1$ is a constant.

By our assumption, we have

$$\begin{aligned} u_{a_i}(r_1) &= u_{a_i}(r_2), \forall i = 0, 1 \\ \Rightarrow P(1|a_i)g(r_1) - h(r_1) + C_i &= P(1|a_i)g(r_2) - h(r_2) + C_i, \forall i = 0, 1 \\ \Rightarrow P(1|a_i) &= \frac{h(r_2) - h(r_1)}{g(r_2) - g(r_1)}, \forall i = 0, 1 \\ \Rightarrow P(1|a_0) &= P(1|a_1) \end{aligned}$$

The above equation $P(1|a_0) = P(1|a_1)$ contradicts with the distinguishability condition. Therefore, the specified mixed strategy for Alice cannot be part of a PBE of the 3-stage D game.

Step 3: By the results of step 1 and 2, there are four types of strategies that can possibly be PBE strategies for Alice in the 3-stage D game. We discuss these four types of strategies separately:

1. The truthful strategy is a possible PBE strategy for Alice in the 3-stage D game, as stated in the theorem.
2. The delaying strategy is a possible PBE strategy for Alice in the 3-stage D game, as stated in the theorem.
3. The third type of strategy is the mixed strategy given by the equation below where $r \neq P(1)$.

$$\sigma_{a_i}(P(1|a_i)) = 1 - p, \sigma_{a_i}(r) = p, \sigma_{a_{1-i}}(r) = 1$$

where $p = \frac{P(a_{1-i})(r - P(1|a_{1-i}))}{P(a_i)(P(1|a_i) - r)}$ and $u_{a_i}(P(1|a_i)) = u_{a_i}(r)$ is satisfied for some $r \in (\min_i P(1|a_i), P(1)) \cup (P(1), \max_i P(1|a_i))$, $\forall i = 0, 1$.

For the above strategy to be a PBE strategy, the following necessary condition must be satisfied.

$$u_{a_i}(P(1|a_i)) = u_{a_i}(r^*),$$

where $u_m(r)$ denotes Alice's ex-ante expected payoff by using this mixed strategy in the first stage of a PBE of the 3-stage D game and $r^* = \arg \max_r u_m(r), r \in (\min_i P(1|a_i), P(1)) \cup (P(1), \max_i P(1|a_i))$.

We will show that

$$u'_m(r^*) = 0 \Rightarrow u_{a_i}(P(1|a_i)) = u_{a_i}(r^*)$$

The expression of $u_m(r)$ and its first derivative are as follows:

$$\begin{aligned}
u_m(r) &= P(a_i) \left(1 - \frac{P(a_{1-i})(r - P(1|a_{1-i}))}{P(a_i)(P(1|a_i) - r)} \right) u_{a_i}(P(1|a_i)) \\
&\quad + P(a_i) \frac{P(a_{1-i})(r - P(1|a_{1-i}))}{P(a_i)(P(1|a_i) - r)} u_{a_i}(r) + P(a_{1-i}) u_{a_{1-i}}(r) \\
\Rightarrow u'_m(r) &= P(a_{1-i})(P(1|a_i) - P(1|a_{1-i})) \frac{u_{a_i}(r) - u_{a_i}(P(1|a_i))}{(P(1|a_i) - r)^2}
\end{aligned}$$

By the distinguishability condition, we know that $P(1|a_i) - P(1|a_{1-i}) \neq 0$. Therefore, we have

$$u'_m(r) = 0 \Rightarrow u_{a_i}(r) - u_{a_i}(P(1|a_i)) = 0$$

4. The final type of strategy is the mixed strategy defined below:

$$\begin{aligned}
&\exists! r \in (P(1|a_0), P(1|a_1)), p \in (0, 1), q \in (0, 1), \\
&\text{s.t. } \sigma_{a_0}(P(1|a_0)) = 1 - p, \sigma_{a_0}(r) = p, \sigma_{a_1}(r) = q, \sigma_{a_1}(P(1|a_1)) = 1 - q \quad (\text{A.10})
\end{aligned}$$

For this mixed strategy, we observe that, if Alice uses this strategy in a PBE of the 3-stage D game, then there must also exist a PBE of where Alice uses the truthful strategy in the first stage. So we include this mixed strategy as a special case when the truthful PBE exists for this game.

□

A.1.7 Proof of Theorem 6

Proof. To show that Alice's strategy and Bob's strategy form a PBE of the 3-stage D game, we need to prove 3 things below.

1. First, we show that Bob's belief on the equilibrium path is derived from Alice's strategy using Bayes' rule.

If Alice reports $P(1|a_i)$ in the first stage, then Bob's belief should assign probability

1 to Alice's signal a_i . Thus, Bob strategy must be to change the market probability to $P(1|a_i, s_B)$ in the second stage if he receives s_B signal. By definition of $x_{s_B}(r)$ in equation (3.14), we can easily check that

$$x_{s_B}(P(1|a_i)) = P(1|a_i, s_B)$$

This means that Bob's belief satisfies this requirement.

2. Next, We show that Bob's belief is valid, i.e. $x_{s_B}(r) \in [\min_{a_i} P(1|a_i, s_B), \max_{a_i} P(1|a_i, s_B)], \forall s_B$.

First notice that $x_{s_B}(r)$ is monotonic in r since the sign of $x_{s_B}(r)$ does not depend on the value of r .

$$x'_{s_B}(r) = \frac{(P(1|a_0, s_B) - P(1|a_1, s_B))P(a_0|s_B)P(a_1|s_B)P(a_0)P(a_1)(P(1|a_0) - P(1|a_1))}{[P(a_0|s_B)P(a_1) - P(a_1|s_B)P(a_0)]r + [P(a_1|s_B)P(1, a_0) - P(a_0|s_B)P(1, a_1)]^2}$$

Thus, $x_{s_B}(r)$ achieves its maximum and minimum at $r = P(1|a_i)$. So we just need to check the value of $x_{s_B}(P(1|a_i)) \in [\min_{a_i} P(1|a_i, s_B), \max_{a_i} P(1|a_i, s_B)], \forall i = 0, 1$. From the argument above, we have $x_{s_B}(P(1|a_i)) = P(1|a_i, s_B)$ and it's within the specified range. Thus, Bob's belief is valid.

3. Finally, we prove that given Bob's strategy in the second stage, Alice maximizes her total expected payoff by reporting $P(1|a_i)$ when she receives the a_i signal. When Alice receives the signal a_i and reports r , her total expected payoff is given by $u_{a_i}(r)$. By our assumption, $u_{a_i}(r)$ is monotonically decreasing as r changes from $P(1|a_i)$ to $P(1|a_{1-i})$. Thus, when Alice receives the a_i signal, her total expected payoff is uniquely maximized by reporting $P(1|a_i)$.

□

A.2 Omitted Derivations

A.2.1 Derivation for the expression of $u_{a_i}(r)$

Let σ be Alice's first-stage strategy in any PBE of the 3-stage D game and let r be any report in the support of σ . Since σ and r satisfy the consistency condition, we must have

$$\begin{aligned} P(1|r, \sigma) &= r \\ \Rightarrow \frac{P(1|a_0)\sigma_{a_0}(r)P(a_0) + P(1|a_1)\sigma_{a_1}(r)P(a_1)}{P(a_0)\sigma_{a_0}(r) + P(a_1)\sigma_{a_1}(r)} &= r \\ \Rightarrow \sigma_{a_0}(r)P(a_0)(P(1|a_0) - r) &= \sigma_{a_1}(r)P(a_1)(r - P(1|a_1)) \end{aligned} \quad (\text{A.11})$$

By the consistency condition, it's easy to see that $r \in [\min_{a_i}\{P(1|a_i)\}, \max_{a_i}\{P(1|a_i)\}]$.

By equation (A.11), we have

$$\begin{aligned} &\sigma_{a_0}(r)P(a_0)(P(1|a_0) - r) + \sigma_{a_0}(r)P(a_1)(r - P(1|a_1)) \\ &= \sigma_{a_0}(r)P(a_1)(r - P(1|a_1)) + \sigma_{a_1}(r)P(a_1)(r - P(1|a_1)) \\ \Rightarrow \sigma_{a_0}(r)(P(a_0)(P(1|a_0) - r) &+ P(a_1)(r - P(1|a_1))) \\ &= (\sigma_{a_0}(r) + \sigma_{a_1}(r))P(a_1)(r - P(1|a_1)) \\ \Rightarrow \frac{\sigma_{a_0}(r)}{\sigma_{a_0}(r) + \sigma_{a_1}(r)} &= \frac{P(a_1)(r - P(1|a_1))}{P(a_0)(P(1|a_0) - r) + P(a_1)(r - P(1|a_1))} \end{aligned} \quad (\text{A.12})$$

At any PBE, Bob's belief on the equilibrium path is derived from Alice's strategy by using the Bayes' rule. Since Alice only has 2 realized signals, it suffices to specify $\mu_{r,s_B}(a_0)$ since $\mu_{r,s_B}(a_1) = 1 - \mu_{r,s_B}(a_0)$. Bob's belief can be derived as follows:

$$\mu_{r,s_B}(a_0) = \frac{P(a_0, r|s_B)}{P(r|s_B)} = \frac{P(a_0|s_B)\sigma_{a_0}(r)}{P(a_0|s_B)\sigma_{a_0}(r) + P(a_1|s_B)\sigma_{a_1}(r)} \quad (\text{A.13})$$

Taking equation (A.12) and plugging into equation (A.13), we have

$$\mu_{r,s_B}(a_0) = \frac{P(a_0|s_B)P(a_1)(r - P(1|a_1))}{(P(a_0|s_B)P(a_1) - P(a_1|s_B)P(a_0))r + (P(a_1|s_B)P(1, a_0) - P(a_0|s_B)P(1, a_1))} \quad (\text{A.14})$$

At any PBE, Bob's strategy $x_{s_B}(r)$ is fully determined given Bob's belief, Bob's signal, and Alice's report, as follows:

$$x_{s_B}(r) = P(1|r, s_B) = \mu_{r, s_B}(a_0)P(1|a_0, s_B) + (1 - \mu_{r, s_B}(a_0))P(1|a_1, s_B) \quad (\text{A.15})$$

Plugging the expression of Bob's belief (A.14) into the definition of Bob's strategy (A.15), we have

$$x_{s_B}(r) = \frac{P(1, s_B|a_0)(P(1|a_1) - r) + P(1, s_B|a_1)(r - P(1|a_0))}{P(s_B|a_0)(P(1|a_1) - r) + P(s_B|a_1)(r - P(1|a_0))}$$

Finally, we can write down the expression of $u_{a_i}(r)$ as follows.

$$\begin{aligned} u_{a_i}(r) = & P(1|a_i) \log \frac{r}{P(1)} + P(0|a_i) \log \frac{1-r}{1-P(1)} \\ & + \sum_{s_B} \left\{ P(1, s_B|a_i) \log \frac{P(1|a_i, s_B)}{x_{s_B}(r)} + P(0, s_B|a_i) \log \frac{P(0|a_i, s_B)}{1-x_{s_B}(r)} \right\} \end{aligned}$$

Appendix B

Appendix to Chapter 4

B.1 Proof of Proposition 1

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

$$\frac{\partial}{\partial r_A} L(f_{a_i, \emptyset}, r_A) = \frac{r_A - f_{a_i, \emptyset}}{r_A(1 - r_A)}. \quad (\text{B.1})$$

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

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

$$\frac{\partial}{\partial f_{a_i, \emptyset}} L(f_{a_i, \emptyset}, r_A) = \log \left(\frac{f_{a_i, \emptyset}}{1 - f_{a_i, \emptyset}} \frac{1 - r_A}{r_A} \right). \quad (\text{B.2})$$

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

B.2 Example 2

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

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

Table B.1: An example prior distribution. Each cell gives the value of $P(\Omega, S_A, S_B)$ for the corresponding realizations of Ω , 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

$$L(f_{T,\emptyset}, f_{H,\emptyset}) = f_{T,\emptyset} \log \frac{f_{T,\emptyset}}{f_{H,\emptyset}} + (1 - f_{T,\emptyset}) \log \frac{1 - f_{T,\emptyset}}{1 - f_{H,\emptyset}} = \left(0.54 \log \frac{0.54}{0.64} + 0.46 \log \frac{0.46}{0.36} \right) \approx 0.021. \quad (\text{B.3})$$

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

$$\begin{aligned} & E_{S_B}[Q(f_{H,S_B}) - Q(f_{T,S_B}) \mid S_A = T] \\ &= P(S_B = H \mid S_A = T)(f_{H,H} - f_{T,H}) + P(S_B = T \mid S_A = T)(f_{H,T} - f_{T,T}) \\ &= 0.6(1 - 0.9) + 0.4(0.1 - 0) \end{aligned} \quad (\text{B.4})$$

$$= 0.1. \quad (\text{B.5})$$

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 (4.4) 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:

$$L(f_{T,\emptyset}, Y_T) = E_{S_B}[Q(f_{H,S_B}) - Q(f_{T,S_B}) \mid S_A = T] \quad (\text{B.6})$$

$$\Rightarrow 0.54 \log \frac{0.54}{Y_T} + 0.46 \log \frac{0.46}{1 - Y_T} = 0.1 \quad (\text{B.7})$$

$$\Rightarrow Y_T \approx 0.747 \quad (\text{B.8})$$

Similarly, we solve for Y_H below:

$$L(f_{H,\emptyset}, Y_H) = E_{S_B}[Q(f_{H,S_B}) - Q(f_{T,S_B}) \mid S_A = H] \quad (\text{B.9})$$

$$\Rightarrow 0.64 \log \frac{0.64}{Y_H} + 0.36 \log \frac{0.36}{1 - Y_H} = 0.1 \quad (\text{B.10})$$

$$\Rightarrow Y_H \approx 0.827 \quad (\text{B.11})$$

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.

B.3 Proof of Proposition 3

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

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

$$L(f_{T,\emptyset}, Y_T) = L(f_{T,\emptyset}, Y_{-T}). \quad (\text{B.12})$$

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

$$L(f_{T,\emptyset}, Y_{-T}) < L(f_{H,\emptyset}, Y_{-T}). \quad (\text{B.13})$$

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

$$L(f_{H,\emptyset}, Y_T) < L(f_{T,\emptyset}, Y_T). \quad (\text{B.14})$$

Hence, we must have $L(f_{H,\emptyset}, Y_T) < L(f_{H,\emptyset}, Y_{-T})$ due to equation (B.12) and inequalities (B.13) and (B.14), as

$$L(f_{H,\emptyset}, Y_T) < L(f_{T,\emptyset}, Y_T) = L(f_{T,\emptyset}, Y_{-T}) < L(f_{H,\emptyset}, Y_{-T}). \quad (\text{B.15})$$

□

B.4 Example 3

Example 3. Consider the outside payoff function and the prior distribution in Table B.2.

We show below that there exists sufficiently small ϵ such that $Y_H < Y_T$.

	$\Omega = 1$			$\Omega = 0$	
	$S_A = H$	$S_A = T$		$S_A = H$	$S_A = T$
$S_B = H$	ϵ	ϵ	$S_B = H$	0	$0.5 - 2\epsilon$
$S_B = T$	ϵ	0	$S_B = T$	$0.5 - 2\epsilon$	ϵ

Table B.2: An example prior distribution with $\epsilon \in (0, 0.25)$. Each cell gives the value of $P(\Omega, S_A, S_B)$ for the corresponding realizations of Ω , S_A , and S_B .

It is easy to compute $f_{H,\emptyset} = 4\epsilon$, $f_{T,\emptyset} = 2\epsilon$, $f_{H,H} = 1$, $f_{H,T} = \frac{\epsilon}{0.5-\epsilon}$, $f_{T,H} = \frac{\epsilon}{0.5-\epsilon}$, and

$f_{T,T} = 0$. With this, we can calculate

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

As ϵ approaches 0, we have

$$\lim_{\epsilon \rightarrow 0} L(f_{H,\emptyset}, Y_H) = 0. \quad (\text{B.16})$$

Because $\lim_{\epsilon \rightarrow 0} f_{H,\emptyset} = \lim_{\epsilon \rightarrow 0} 4\epsilon = 0$, by definition of $L(f_{H,\emptyset}, Y_H)$, (B.16) implies that

$$\lim_{\epsilon \rightarrow 0} Y_H = 0. \quad (\text{B.17})$$

Similarly, we have

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

As ϵ approaches 0, we have

$$\lim_{\epsilon \rightarrow 0} L(f_{T,\emptyset}, Y_T) = 1. \quad (\text{B.18})$$

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

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

Given (B.18), we have

$$\lim_{\epsilon \rightarrow 0} Y_T = 1 - 1/e. \quad (\text{B.19})$$

Combining (B.17) and (B.19), 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

$$\lim_{\epsilon \rightarrow 0} P(S_B = H | S_A = T) = \lim_{\epsilon \rightarrow 0} (1 - 2\epsilon) = 1, \quad (\text{B.20})$$

$$\lim_{\epsilon \rightarrow 0} P(S_B = T | S_A = T) = \lim_{\epsilon \rightarrow 0} 2\epsilon = 0. \quad (\text{B.21})$$

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

$$\lim_{\epsilon \rightarrow 0} (f_{HH} - f_{TH}) = \lim_{\epsilon \rightarrow 0} (1 - \frac{\epsilon}{0.5 - \epsilon}) = 1, \quad (\text{B.22})$$

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

$$\lim_{\epsilon \rightarrow 0} (f_{HT} - f_{TT}) = \lim_{\epsilon \rightarrow 0} (\frac{\epsilon}{0.5 - \epsilon} - 0) = 0. \quad (\text{B.23})$$

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 2, 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.

B.5 Proof of Theorem 9

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 inter-

vals $[\max(Y_{-H}, Y_T), f_{H,\emptyset})$ 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:

$$SE_2(k) : \begin{cases} \sigma_H^S(k) = 1, \sigma_T^S(f_{T,\emptyset}) = 1 \\ \mu_{s_B, r_A}^S(H) = \begin{cases} 1, & \text{if } r_A \in [k, 1] \\ 0, & \text{if } r_A \in [0, k) \end{cases} \end{cases} \quad (\text{B.24})$$

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:

$$SE_3(k) : \begin{cases} \sigma_H^S(k) = 1, \sigma_T^S(f_{T,\emptyset}) = 1 \\ \mu_{s_B, r_A}^S(H) = \begin{cases} 1, & \text{if } r_A = k \\ 0, & \text{if } r_A \in [0, k) \cup (k, 1] \end{cases} \end{cases} . \quad (\text{B.25})$$

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

$$L(f_{H,\emptyset}, k) \leq L(f_{H,\emptyset}, Y_{-H}) = E_{S_B}[Q(f_{H,S_B}) - Q(f_{T,S_B}) \mid a_i] \quad (\text{B.26})$$

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. $\square$

B.6 Proof of Theorem 10

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:

$$SE_4(k) : \begin{cases} \sigma_H^S(k) = 1, \sigma_T^S(f_{T,\emptyset}) = 1 \\ \mu_{s_B, r_A}^S(H) = \begin{cases} 0, & \text{if } r_A \in (k, 1] \\ 1, & \text{if } r_A \in [0, k] \end{cases} \end{cases} \quad (\text{B.27})$$

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

$$L(f_{H,\emptyset}, k) \leq L(f_{H,\emptyset}, Y_{-H}) = E_{S_B}[Q(f_{H,S_B}) - Q(f_{T,S_B}) \mid a_i] \quad (\text{B.28})$$

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. $\square$

B.7 Proof of Proposition 5

Proof. The existence of a separating PBE requires $Y_H \geq Y_T$ by Theorem 8. By Lemma 5, we have $\sigma_T(f_{T,\emptyset}) = 1$ at any separating PBE.

By Theorem 9, 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 8, 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 8, 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 10, 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 8, 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 $r_A = Y_T + \epsilon$ for a

small $\epsilon > 0$ such that $L(f_{H,\emptyset}, Y_{-T}) > L(f_{H,\emptyset}, r_A)$. Such an ϵ must exist because as $\epsilon \rightarrow 0$, $L(f_{H,\emptyset}, r_A) \rightarrow L(f_{H,\emptyset}, Y_T)$. 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 $r_A = Y_T + \epsilon$ 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 (4.11) is such a belief. $\square$

B.8 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 $\epsilon > 0$ such that if Alice receives the H signal, then reporting any $r_A \leq r$ is strictly worse than reporting $Y_T + \epsilon$.

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 $0 < \epsilon < r - r_A$, then we must have $L(f_{H,\emptyset}, r_A) > L(f_{H,\emptyset}, r - \epsilon) = L(f_{H,\emptyset}, Y_T + \epsilon)$. 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 $Y_T + \epsilon$ is weakly better than her expected outside payoff by reporting r_A . Therefore, for any $\epsilon \in (0, r - r_A)$, Alice is strictly worse off reporting any $r_A < r$ than reporting $Y_T + \epsilon$.

(2) $r_A = r$: For any small $\epsilon > 0$, we have that $L(f_{H,\emptyset}, r) < L(f_{H,\emptyset}, Y_T + \epsilon)$. However as $\epsilon \rightarrow 0$, $L(f_{H,\emptyset}, Y_T + \epsilon) - L(f_{H,\emptyset}, r) \rightarrow 0$. 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 $\mu_{s_B, Y_T + \epsilon}(T) = 0$, for any $0 < \epsilon \leq Y_H - Y_T$. Regardless of ϵ , Alice's expected outside payoff by reporting $Y_T + \epsilon$ 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 $\epsilon > 0$ such that Alice's total expected payoff by reporting r is strictly less than her total expected payoff by reporting $Y_T + \epsilon$. $\square$

B.9 Proof of Theorem 16

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 (B.24).

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 5, 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 10 in Appendix B.6. We know that $Y_{-H} \leq Y_T$ implies $Y_H \geq Y_T$ by the proof of Theorem 8. 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. $\square$

B.10 Proof of Theorem 17

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 (4.29) and (4.30) are satisfied, we can show that SE_5 is a separating PBE of our game.

$$SE_5 : \begin{cases} \sigma_H^S(\max(Y_T, f_{H,\emptyset})) = 1, \sigma_T^S(f_{T,\emptyset}) = 1 \\ \mu_{s_B, r_A}^S(H) = \begin{cases} 1, & \text{if } r_A \in [Y_T, 1] \\ 0, & \text{if } r_A \in [0, Y_T) \end{cases} \end{cases} \quad (\text{B.29})$$

First, we show that Alice's strategy is optimal given Bob's belief. Since inequality (4.29) 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 (4.30). 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 (4.31) and (4.32) are satisfied, then we can show that SE_6 is a

separating PBE of our game.

$$SE_6 : \begin{cases} \sigma_H^S(Y_{-T}) = 1, \sigma_T^S(f_{T,\emptyset}) = 1 \\ \mu_{s_B, r_A}^S(H) = \begin{cases} 0, & \text{if } r_A \in (Y_{-T}, 1] \\ 1, & \text{if } r_A \in [0, Y_{-T}] \end{cases} \end{cases} \quad (\text{B.30})$$

First, we show that Alice's strategy is optimal given Bob's belief. Since inequality (4.31) 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 $f_{T,\emptyset}$. 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 $f_{T,\emptyset}$ 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 (4.29) and (4.30) is violated. Then, we can show that Alice does not report any value $r_A \in [f_{T,\emptyset}, 1]$ after receiving the H signal at any separating PBE. We divide the argument for this into 2 cases.

(1) If inequality (4.29) is violated, we know that Y_T is not well defined. We show by contradiction that Alice does not report any value in $[f_{T,\emptyset}, 1]$ after receiving the H signal. Suppose that at a separating PBE, Alice reports $r_A \in [f_{T,\emptyset}, 1]$ 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 5, in any separating PBE, Bob's belief must be $\mu_{s_B, f_{T, \emptyset}}(H) = 0$ and Alice must report $f_{T, \emptyset}$ after receiving the T signal. Since inequality (4.29) is violated, then we have that

$$L_s(f_{T, \emptyset}, r_A) \leq L_s(f_{T, \emptyset}, 1) < E_{S_B}[Q(f_{H, S_B}) - Q(f_{T, S_B}) \mid S_A = T], \quad (\text{B.31})$$

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

(2) Otherwise, if inequality (4.29) is satisfied but inequality (4.30) is violated, then we know that Y_T is well defined. If $f_{H, \emptyset} \geq Y_T$, then inequality (4.30) is automatically satisfied, so we must have that $f_{H, \emptyset} < Y_T$ and $L_s(f_{H, \emptyset}, Y_T) > E_{S_B}[Q(f_{H, S_B}) - Q(f_{T, S_B}) \mid S_A = H]$. Then Alice does not report any $r_A \in [Y_T, 1]$ after receiving the H signal because doing so is dominated by reporting $f_{H, \emptyset}$. Next, we can show by contradiction that Alice does not report any $r_A \in [f_{T, \emptyset}, Y_T)$ after receiving the H signal at any separating PBE. Suppose that at any separating PBE, Alice reports $r_A \in [f_{T, \emptyset}, Y_T)$ 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 5, 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.

Hence, if at least one of the inequalities (4.29) and (4.30) is violated, then at any separating PBE, Alice does not report any $r_A \in [f_{T, \emptyset}, 1]$ after receiving the H signal.

Similarly, we can show that, if at least one of the inequalities (4.31) and (4.32) is violated, Alice does not report any value $r_A \in [0, f_{T, \emptyset}]$ after receiving the H signal at any separating PBE. We again consider 2 cases:

(1) If inequality (4.31) is violated, we know that Y_{-T} is not well defined. Then we can show that Alice does not report any value in $[0, f_{T, \emptyset}]$ after receiving the H signal. We prove

by contradiction. Suppose that at a separating PBE, Alice reports $r_A \in [0, f_{T,\emptyset}]$ 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 5, in any separating PBE, Bob's belief must be $\mu_{s_B, f_{T,\emptyset}}(H) = 0$ and Alice must report $f_{T,\emptyset}$ after receiving the T signal. Since inequality (4.31) is violated, we have that

$$L_s(f_{T,\emptyset}, r_A) \leq L_s(f_{T,\emptyset}, \emptyset) < E_{S_B}[Q(f_{H,S_B}) - Q(f_{T,S_B}) \mid S_A = T], \quad (\text{B.32})$$

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

(2) Otherwise, if inequality (4.31) is satisfied but inequality (4.32) is violated, then we know that Y_{-T} is well defined. Also, we must have that $L_s(f_{H,\emptyset}, Y_{-T}) > E_{S_B}[Q(f_{H,S_B}) - Q(f_{T,S_B}) \mid S_A = H]$. Then Alice does not report any $r_A \in [0, Y_{-T}]$ after receiving the H signal because doing so is dominated by reporting $f_{H,\emptyset}$. Next, We can show by contradiction that at any separating PBE, Alice does not report any $r_A \in (Y_{-T}, f_{T,\emptyset}]$ after receiving the H signal. Suppose that at any separating PBE, Alice reports $r_A \in (Y_{-T}, f_{T,\emptyset}]$ 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 5, 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.

Hence, if at least one of the inequalities (4.31) and (4.32) is violated, in any separating PBE, Alice does not report any $r_A \in [0, f_{T,\emptyset}]$ 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. $\square$

Appendix C

Appendix to Chapter 5

C.1 Estimated HMMs

i	$P(\text{MM} \mid \text{MM})$	$P(\text{MM} \mid \text{GB})$	P_i
0	0.82	0.45	0.50
1	1.00	0.99	0.27
2	0.99	0.04	0.17
3	0.13	0.09	0.05

(a) States and initial probabilities.

	$j = 0$	$j = 1$	$j = 2$	$j = 3$
$i = 0$	0.89	0.08	0.02	0.01
$i = 1$	0.01	0.99	0.00	0.00
$i = 2$	0.01	0.02	0.96	0.01
$i = 3$	0.02	0.00	0.00	0.97

(b) Transition probabilities.

Table C.1: Treatment 1 Estimated HMM ($K = 4$). $P(\text{MM} \mid \text{MM})$ is the probability of reporting **MM** given a **MM** signal. $P(\text{MM} \mid \text{GB})$ is the probability of reporting **MM** given a **GB** signal. P_i is the initial probability of state i . The cell at row i and column j gives the transition probability from state i to state j .

i	$P(\text{MM} \mid \text{MM})$	$P(\text{MM} \mid \text{GB})$	P_i
0	0.96	0.01	0.33
1	0.99	0.99	0.04
2	0.01	0.00	0.14
3	0.61	0.31	0.49

(a) States and initial probabilities.

	$j = 0$	$j = 1$	$j = 2$	$j = 3$
$i = 0$	0.95	0.00	0.03	0.02
$i = 1$	0.01	0.99	0.00	0.00
$i = 2$	0.00	0.00	1.00	0.00
$i = 3$	0.02	0.02	0.05	0.91

(b) Transition probabilities.

Table C.2: Treatment 2 Estimated HMM ($K = 4$). $P(\text{MM} \mid \text{MM})$ is the probability of reporting **MM** given a **MM** signal. $P(\text{MM} \mid \text{GB})$ is the probability of reporting **MM** given a **GB** signal. P_i is the initial probability of state i . The cell at row i and column j gives the transition probability from state i to state j .

i	$P(\text{MM} \mid \text{MM})$	$P(\text{MM} \mid \text{GB})$	P_i
0	0.97	0.06	0.23
1	0.87	0.97	0.07
2	0.02	0.03	0.06
3	0.54	0.42	0.64

(a) States and initial probabilities.

	$j = 0$	$j = 1$	$j = 2$	$j = 3$
$i = 0$	0.98	0.00	0.01	0.01
$i = 1$	0.00	0.97	0.01	0.02
$i = 2$	0.01	0.01	0.99	0.00
$i = 3$	0.00	0.01	0.01	0.97

(b) Transition probabilities.

Table C.3: Treatment 3 Estimated HMM ($K = 4$). $P(\text{MM} \mid \text{MM})$ is the probability of reporting **MM** given a **MM** signal. $P(\text{MM} \mid \text{GB})$ is the probability of reporting **MM** given a **GB** signal. P_i is the initial probability of state i . The cell at row i and column j gives the transition probability from state i to state j .

i	$P(\text{MM} \mid \text{MM})$	$P(\text{MM} \mid \text{GB})$	P_i
0	0.96	0.06	0.22
1	0.73	0.61	0.48
2	0.96	0.97	0.06
3	0.34	0.37	0.25

(a) States and initial probabilities.

	$j = 0$	$j = 1$	$j = 2$	$j = 3$
$i = 0$	0.97	0.02	0.00	0.01
$i = 1$	0.01	0.97	0.03	0.00
$i = 2$	0.00	0.01	0.98	0.00
$i = 3$	0.01	0.02	0.01	0.97

(b) Transition probabilities.

Table C.4: Treatment 4 Estimated HMM ($K = 4$). $P(\text{MM} \mid \text{MM})$ is the probability of reporting **MM** given a **MM** signal. $P(\text{MM} \mid \text{GB})$ is the probability of reporting **MM** given a **GB** signal. P_i is the initial probability of state i . The cell at row i and column j gives the transition probability from state i to state j .

i	$P(\text{MM} \mid \text{MM})$	$P(\text{MM} \mid \text{GB})$	P_i
0	0.98	0.02	0.65
1	0.02	0.00	0.02
2	0.16	0.96	0.01
3	0.68	0.34	0.33

(a) States and initial probabilities.

	$j = 0$	$j = 1$	$j = 2$	$j = 3$
$i = 0$	1.00	0.00	0.00	0.00
$i = 1$	0.01	0.99	0.01	0.00
$i = 2$	0.02	0.00	0.92	0.06
$i = 3$	0.01	0.00	0.01	0.98

(b) Transition probabilities.

Table C.5: Non-Peer Prediction Treatment Estimated HMM ($K = 4$). $P(\text{MM} \mid \text{MM})$ is the probability of reporting **MM** given a **MM** signal. $P(\text{MM} \mid \text{GB})$ is the probability of reporting **MM** given a **GB** signal. P_i is the initial probability of state i . The cell at row i and column j gives the transition probability from state i to state j .

Appendix D

Appendix to Chapter 6







Figure D.1: The dynamics of the estimated strength parameters for adaptive polling trials




(a) Random polling trial 1




(b) Random polling trial 2




(c) Random polling trial 3



(d) Random polling trial 4



(e) Random polling trial 5



(f) Random polling trial 6

Figure D.2: The dynamics of the estimated strength parameters for random polling trials