

Washington University in St. Louis

Washington University Open Scholarship

McKelvey School of Engineering Theses & Dissertations

McKelvey School of Engineering

Spring 5-15-2017

On the Aggregation of Subjective Inputs from Multiple Sources

Mithun Chakraborty

Washington University in St. Louis

Follow this and additional works at: https://openscholarship.wustl.edu/eng_etds



Part of the [Computer Sciences Commons](#)

Recommended Citation

Chakraborty, Mithun, "On the Aggregation of Subjective Inputs from Multiple Sources" (2017). *McKelvey School of Engineering Theses & Dissertations*. 237.

https://openscholarship.wustl.edu/eng_etds/237

This Dissertation is brought to you for free and open access by the McKelvey School of Engineering at Washington University Open Scholarship. It has been accepted for inclusion in McKelvey School of Engineering Theses & Dissertations by an authorized administrator of Washington University Open Scholarship. For more information, please contact digital@wumail.wustl.edu.

WASHINGTON UNIVERSITY IN ST. LOUIS

School of Engineering and Applied Science
Department of Computer Science and Engineering

Dissertation Examination Committee:

Sanmay Das, Chair

Yiling Chen

Roman Garnett

Roch Guerin

Brendan Juba

On the Aggregation of Subjective Inputs from Multiple Sources

by

Mithun Chakraborty

A dissertation presented to
The Graduate School
of Washington University in
partial fulfillment of the
requirements for the degree
of Doctor of Philosophy

May 2017

Saint Louis, Missouri

copyright by
Mithun Chakraborty
2017

Contents

List of Tables	v
List of Figures	vi
Acknowledgments	ix
Abstract	xii
1 Introduction	1
1.1 Overview	1
1.2 Contributions	3
1.3 Organization of the dissertation	6
2 Semantics of aggregation with market scoring rules	9
2.1 Introduction	9
2.2 Related work	11
2.3 Model and definitions	12
2.3.1 Opinion Pool (OP)	13
2.3.2 Market Scoring Rule (MSR)	15
2.4 MSR behavior with risk-averse myopic agents	17
2.4.1 LMSR and constant absolute risk aversion (CARA) utility: LogOP	21
2.4.2 LMSR and an atypical utility with decreasing absolute risk aversion: LinOP	24
2.4.3 LMSR and logarithmic utility	26
2.5 Discussion and future work	31
3 Outcome manipulation in incentivized collective forecasting mechanisms	33
3.1 Introduction	33
3.1.1 A model for manipulation	35
3.1.2 Contributions	36
3.2 Related work	38
3.3 Model and definitions	41
3.4 Equilibrium analysis of the two-player game	45
3.4.1 Equilibrium when Bob's non-participation is certain ($\pi = 1$)	53
3.4.2 Equilibrium when Bob's participation is certain ($\pi = 0$)	56

3.4.3	Equilibrium when Bob's participation is uncertain ($0 < \pi < 1$)	59
3.5	Implications of the equilibrium result	65
3.5.1	Dependence of crossover probability on Alice's uncertainty	65
3.5.2	Private signal revelation	66
3.5.3	HPP profit sharing	68
3.5.4	Informativeness of market prices about final outcome	69
3.6	A specific signal structure	70
3.7	Extensions to more than two agents	71
3.8	Incentivizing truthful trading and voting	75
3.9	Discussion	79
4	Aggregating censored signals from non-strategic noisy agents	81
4.1	Introduction	81
4.2	Related work	83
4.3	The learning problem	84
4.4	The intuition for our solution approach	86
4.4.1	Non-Parametric Histograms	88
4.5	The algorithm	89
4.6	Analysis	93
4.6.1	Bayesian inference with uncensored signals	93
4.6.2	Approximate Bayesian inference with censored signals	96
4.7	Experimental results	98
4.8	Discussion	99
5	Market making in practice: CDA with LMSR	101
5.1	Introduction and related work	101
5.2	Market Model	104
5.3	Evaluation	109
5.3.1	Results	112
5.4	Discussion	114
6	Conclusion	116
Appendix A	Proofs of theorems in Chapter 2	117
A.1	A general well-behaved MSR as an Opinion Pool for a general risk-averse utility	117
A.1.1	LMSR as LogOP for CARA utility agents	123
A.1.2	LMSR as LinOP for an atypical utility with decreasing absolute risk aversion	126
A.2	LMSR with logarithmic utility agents	130
Appendix B	Proofs of theorems in Chapter 3	131
B.1	Proof of Proposition 2 from Section 3.4	131

B.2	Equilibrium of the game in Section 3.4.3	133
B.2.1	Analysis for symmetric well-behaved MSRs with $0 < p^L < \frac{1}{4}$, $\frac{3}{4} < p^H < 1$, e.g. LMSR	141
B.2.2	Analysis for symmetric well-behaved MSRs with $p^L = \frac{1}{4}$, $p^H = \frac{3}{4}$, e.g. QMSR	149
B.2.3	Analysis for symmetric well-behaved MSRs with $\frac{1}{4} < p^L < \frac{1}{2}$, $\frac{1}{2} < p^H < \frac{3}{4}$, e.g. SMSR	152
Appendix C	Proofs of results in Chapter 4	162
Vita	180

List of Tables

3.1	Structural properties of the three representative market scoring rules considered in this paper.	49
3.2	Alice's best vote and Bob's best report-vote pair if he trades, given starting price p_0 and Alice's report p_A . Recall that $p^L \in (0, \frac{1}{2})$ and $p^H \in (\frac{1}{2}, 1)$ for any symmetric well-behaved MSR, hence the table covers all possible combinations of the values of p^L , p^H , and p_0 . Theorem 4 applies to the $p_0 = \frac{1}{2}$ shown in the middle of the table.	51
3.3	Alice's PBE price-report for a symmetric well-behaved MSR with $p^L \in (0, \frac{1}{4})$, $p_0 = \frac{1}{2}$	63
3.4	Alice's PBE price-report for a symmetric well-behaved MSR with $p^L = \frac{1}{4}$, $p_0 = \frac{1}{2}$	64
3.5	Alice's PBE price-report for a symmetric well-behaved MSR with $p^L \in (\frac{1}{4}, \frac{1}{2})$, $p_0 = \frac{1}{2}$	64
3.6	The crossover probability π_c as a function of q_0 over the sub-intervals into which the p^L splits the entire possible range $(0, 1)$ of q_0 values, for symmetric well-behaved MSRs with p^L smaller than, equal to, and larger than $\frac{1}{4}$	65
3.7	Alice and Bob's expected profits for being STRATEGIC vs. TRUTHFUL, with 10 intervening traders for the signal structure described in Section 3.7. Results are averaged over 10^5 simulations. The first entry in each cell corresponds to Alice, the second to Bob. The table does not list all strategies available to the players, and does not depict the equilibrium; it shows that when <i>both</i> players employ Theorem 4 strategies, each achieves at least as much expected profit as they would for any strategy-pair in which at least one is TRUTHFUL. We observed similar results for other parameter combinations.	75
B.1	133

List of Figures

2.1	Comparison of a logarithmic utility function $u_{\log}(c; B) = \ln(c + B)$, $c \geq -B$, where $B = 1$ is the (positive) budget, with various instances of the atypical decreasing absolute risk aversion utility function (A.13) $u_{\text{atyp}}(c; B, b) = \ln(\exp((c+B)/b) - 1)$ with the same budget $B = 1$ but different scaling factors $b = 0.1, 1, 10$	27
2.2	The error measure increases with increasing difference between p_{true} and $p_0 = 0.5$ for any fixed \tilde{w}_{\max} , and also with an increase in \tilde{w}_{\max} for any given p_{true} -value; nevertheless, the error appears to be small even for higher values of \tilde{w}_{\max} (less than 0.03 for $0.1 \leq p_{\text{true}} \leq 0.9$, $\tilde{w}_{\max} \leq 0.75$). Error bars are not shown since standard errors are consistently two orders of magnitude smaller than corresponding sample means.	30
2.3	Price trajectories for two sample simulations with $\tilde{w}_{\max} = 0.2$ and $\tilde{w}_{\max} = 2$ are displayed in panels (a) and (b) respectively, $p_{\text{true}} = 0.7$ for both. “Lower bound” (dashed blue curve) and “Upper bound” (dashed red curve) for each trading episode i correspond to price bounds p_i^{\min} and p_i^{\max} respectively. . .	31
2.4	Variation of the approximation error measure in our simulations with respect to the pseudo-sample size (confidence) parameter of the distribution of agent beliefs.	32
3.1	Alice’s payoff as a function of her Stage 1 action for three representative MSRs with Bob’s non-participation probability $\pi = 0$ and starting market price $p_0 = \frac{1}{2}$; the general characteristics are similar for other symmetric well-behaved MSRs: We have jump discontinuities at the lower and upper thresholds, and a value of zero at $p_A = \frac{1}{2}$. For LMSR, the figure is truncated, and both the outer segments actually continue towards $-\infty$ symmetrically, away from $\frac{1}{2}$	59
3.2	Alice’s price-reports for $q_0 > \frac{1}{2}$ and STRATEGIC Bob’s responses for HPP (left) and LPP domains. The actual magnitude of p_A in the LPP domain depends on values of π, q_0 . Results are symmetric for $q_0 < \frac{1}{2}$	61
3.3	Dependence of crossover probability on Alice’s posterior belief about Bob for the three MSRs; e.g. for QMSR, if $q_0 = 0.25$. then $\pi_c \approx 0.9537$, so we have an LPP equilibrium with $p_A = (1 + \pi(1 - q_0))/2$ for $\pi > 0.9537$, and an HPP equilibrium with $p_A = p^H = 0.75$ for $\pi < 0.9537$	66

3.4	Crossover from HPP to LPP equilibria regions for LMSR over $0 < \pi < 1$ for the signal structure described in Section 3.6 with two different sets of parameter values: the prior probability of τ being 0 is $\rho_0 = \frac{1}{2}$ for both panels, and the error probability ρ_e is 0.4 and 0.1 for the left and right panels respectively. For the left panel, $p_A^{LPP} = \mu_{0,0} = \frac{\pi(1-q_0)}{2}$, hence p_A increases with π for $\pi > \pi_c$ whereas for the right panel, $p_A^{LPP} = \mu_{0,1} = \frac{1-\pi q_0}{2}$, hence p_A decreases with π in the LPP domain. In both panels, $p_A = p^L = 0.2$ in the HPP domain. For each curve, we have shown with a dashed blue line Alice's expected average vote taking Lemmas 5 and 6 into account, which coincides with p_A in the LPP domain.	71
3.5	Alice's expected payoff for different values of her trading action p_A given her knowledge of the extensions of Lemmas 5 and 6 for the case of three outcome-deciders. For both plots, Alice has the same signal ($s_A = 0$) and all model parameters are equal, except Bob's non-participation probability π , as noted above.	73
4.1	Evolution of $p_t(v)$ using non-parametric histograms and Bayesian updates; $p_0(v)$ is $\mathcal{N}(0, 10)$, $\sigma_\epsilon = 1$, $V \approx 9.45$. Typical evolution consists of 3 main phases. (i) Truncation ($\rho_t = \sigma_t/\sigma_\epsilon \gg 1$): All observations are in the same direction (here $x_t = +1$) and almost noiseless, so the Bayesian update results in a skewed distribution. Shown for comparison is a truncated (on one side only) and renormalized Gaussian distribution (using our heuristic), which approximates this phase better than Gaussian. (ii) Collapse : When the first observation in the opposite direction arrives (here $x_t = -1$), the distribution collapses to something more symmetric, although not quite normal. Shown for comparison is the entropy-matched normal with the same mean. (iii) Convergence : ρ_t is typically small and μ_t is close to V , i.e. the algorithm has (probabilistically) bracketed V . From then on, nearly independent observations which are close to random signs cause the distribution to rapidly converge to normal, as would be expected with truly independent observations. . . .	89
4.2	Learner's state transitions.	90
4.3	A log-log plot showing the time taken by $\mathbb{E}[\mu_t]$ (approximated by the average computed over 10^6 simulations) to converge to $V - \epsilon$ when $V = 3$, $\mu_0 = 0$, $\rho_0 = 0.5$. Evidently, the convergence time approaches $O(1/\epsilon)$ as ϵ becomes small, which gives us empirical evidence that it is near-optimal.	97
4.4	Plot of average correct-prior convergence time <i>vs</i> ρ_0 , logarithmic along the vertical axis.	98
4.5	Plot of average correct-prior convergence time <i>vs</i> ϵ , logarithmic along the vertical axis.	99
5.1	Illustration of how the market maker in the INT setting places ask and bid quotes every time the state of the books changes.	108

5.2	Experimental results, averaged over 1000 simulations each. The labels along the horizontal axis indicate the number of private Bermoulli trials with success probability p_{true} observed by each trader in the respective simulation set; this number is directly related to the precision in trader beliefs.	110
C.1	An example demonstrating that, for any $\mu^H > \mu^L$, if there is a path from μ^L that crosses $V - \varepsilon$ within a time-horizon then the identical path from μ^H also crosses $V - \varepsilon$ within that horizon.	170

Acknowledgments

I owe the largest debt of gratitude to my advisor, Sanmay Das. He was the person who introduced me to the wonderful world of interdisciplinary research in computational social science that this dissertation is based on, and continues to be a source of valuable insights into not only research but all other aspects of academia as well. He has been incredibly patient, understanding, and supportive at every step, especially when I was dealing with personal tragedy. I can honestly say that joining a Ph.D. program under Sanmay's supervision is the best decision I have made in my life thus far.

I am also immensely grateful to the other members of my dissertation committee: Roman Garnett, Brendan Juba, Roch Guerin, and Yiling Chen. The feedback I received from them on my dissertation proposal (as well as on my talks at conferences and workshops, from Yiling) has been tremendously helpful in improving this dissertation. I have also enjoyed my conversations with all of them in informal settings.

Thanks are due to my publication co-authors Malik Magdon-Ismail, Aseem Brahma, Allen Lavoie, Yonatan Naamad, and Justin Peabody. The last four have also been great office mates as well as friends off-campus. I must also include Meenal Chhabra, Zhuoshu Li, Hao Yan, Gustavo Malkomes, and Shali Jiang on the latter list: Each of them has played a significant role in making my workplace feel like a home away from home!

I must also extend my gratitude to all members of the faculty and administrative staff, as well as all the students I have befriended or interacted with (too many to mention by

name) at each of the three universities where I was a graduate student. In particular, my collaborations with Malik and Brendan have been learning experiences I will always cherish.

I am thankful for financial support from NSF IIS awards 1414452 and 1527037.

Last but not least, I would like to thank my parents, Ira and Santu Chakraborty, for making sure that I received a quality education during my formative years and also for supporting my decision to pursue higher studies abroad. Thanks are also due to my cousins, Soma Basu and Ruma Banerjee, and their families for being my emotional support system for the latter part of my Ph.D. career.

Mithun Chakraborty

Washington University in Saint Louis
May 2017

To my mother

ABSTRACT OF THE DISSERTATION

On the Aggregation of Subjective Inputs from Multiple Sources

by

Mithun Chakraborty

Doctor of Philosophy in Computer Science

Washington University in St. Louis, May 2017

Research Advisor: Professor Sanmay Das

When we have a population of individuals or artificially intelligent agents possessing diverse subjective inputs (e.g. predictions, opinions, etc.) about a common topic, how should we collect and combine them into a single judgment or estimate? This has long been a fundamental question across disciplines that concern themselves with forecasting and decision-making, and has attracted the attention of computer scientists particularly on account of the proliferation of online platforms for electronic commerce and the harnessing of collective intelligence. In this dissertation, I study this problem through the lens of computational social science in three main parts: (1) Incentives in information aggregation: In this segment, I analyze mechanisms for the elicitation and combination of private information from strategic participants, particularly crowdsourced forecasting tools called prediction markets. I show that (a) when a prediction market implemented with a widely used family of algorithms called market scoring rules (MSRs) interacts with myopic risk-averse traders, the price process behaves like an opinion pool, a classical family of belief combination rules, and (b) in an MSR-based game-theoretic model of prediction markets where participants can influence the predicted

outcome but some of them have a non-zero probability of being non-strategic, the equilibrium is one of two types, depending on this probability – either collusive and uninformative or partially revealing; (2) Aggregation with non-strategic agents: In this part, I am agnostic to incentive issues, and focus on algorithms that uncover the ground truth from a sequence of noisy versions. In particular, I present the design and analysis of an approximately Bayesian algorithm for learning a real-valued target given access only to censored Gaussian signals, that performs asymptotically almost as well as if we had uncensored signals; (3) Market making in practice: This component, although tied to the two previous themes, deals more directly with practical aspects of aggregation mechanisms. Here, I develop an adaptation of an MSR to a financial market setting called a continuous double auction, and document its experimental evaluation in a simulated market ecosystem.

Chapter 1

Introduction

1.1 Overview

The main text of the *Rigveda* (Griffith, 1896), an ancient Indian anthology of hymns that is one of the earliest extant literary creations of humankind, ends with the following prayer:

samānī va ākūtiḥ samānā hṛdayāni vaḥ|
samānamastu vo mano yathā vaḥ susahāsatī||

“One and the same be your resolve, and be your minds of one accord.
United be the thoughts of all that all may happily agree.”

— *ṛgvedasamhitā* 10.191.4,
English translation by Ralph T.H. Griffith.

However, in most spheres of human life, thoughts differ and people do not agree on many questions of general interest. Voters with different political leanings prefer different electoral candidates; judges might disagree on the ranking of participants in a competition, and jury members on whether or not the accused is guilty; economists with diverse views on the impact of current government policies often come up with significantly different estimates of the future gross domestic product of a country; even experts in science and technology might disagree on questions such as “Will a manned flight to Mars occur within the next decade?” But sometimes, for the purpose of planning, decision-making, policy formulation,

and suchlike, it is necessary to obtain a single answer to such a question – an answer that can be interpreted as the collective response of the group or population under consideration. This gives rise to the following problem:

When we have a population of individuals¹ possessing potentially different *subjective inputs* about a common topic, such as predictions on uncertain events, noisy observations of hidden truths, etc., how should we *collect* and *combine* them into a single judgment or estimate?

This problem can take on multiple incarnations, and various classes of methods (vote aggregation, pooling, collective judgment, etc.) have been developed to address them; these methods have traditionally been studied in academic disciplines such as economics and finance, sociology, political science, business, management, or operations research, and have attracted the attention of computer scientists particularly on account of the proliferation of online platforms for electronic commerce, crowdsourcing, and the harnessing of collective intelligence (Surowiecki, 2005).

Researchers in machine learning / artificial intelligence have been exploring similar issues under topics such as committee machines (Tresp, 2001), ensemble learning (Opitz and Maclin, 1999), and learning from expert advice (Cesa-Bianchi et al., 1997); but, in areas that lie at the intersection of the computational and social sciences, e.g. algorithmic economics, computational finance, computational social choice etc., this problem takes on additional nuances: One often has to consider the goals and motivations of the providers of subjective inputs in addition to their knowledge structures as well as the constraints and objectives of the entity looking to aggregate these inputs.

Some of the interesting facets of the aggregation problem seen through the lens of computational social science are listed below:

- For a group with human (or, in general, strategic) members, can we take honesty in reporting for granted? If no, how should we offer appropriate *incentives* and / or handle the potential for manipulation?

¹An *individual* can refer to one of several human beings, institutions, artificially intelligent agents, etc.

- In what *form* (e.g. probability distributions over the outcome space, point estimates, monetary bets, etc.) are individual inputs extracted? How much control do we have over this form? Are the inputs known to be constrained or corrupted (censoring, additive noise, etc.) by extraneous factors?
- What is the ultimate objective of our aggregation mechanism? This ties into our criteria for evaluating individual inputs as well as the output of the aggregator (ground truth revealed after aggregation that serves as an objective standard for assessment, peer responses etc.).
- Should individual-level elicitation and aggregation be performed as two decoupled successive stages, or is it better to interweave them somehow?

Evidently, these aspects interact with each other in complex ways, and looking for a solution that addresses all of them simultaneously is often not just impracticable but also unnecessary since not all these issues may be relevant to a single context. In this dissertation, I have adopted the approach of identifying and analyzing problem domains that enable us to focus on a subset of these issues at a time while abstracting away from others; however, my overarching goal is to develop a richer and deeper understanding of methods for subjective input aggregation, and identify and address challenges encountered in practical applications of these techniques. Below is a summary of the specific contributions I intend to make with this dissertation.

1.2 Contributions

Before summarizing the contributions of this dissertation, it is worthwhile to define terminology that will appear repeatedly in the remainder of the document.

We will assume that there exists an entity (a person or an organization) that is interested in some currently uncertain event, modeled as a random variable whose realized outcome may or may not be revealed to it in the future. We will call this entity the *principal*, and sometimes the learner or the decision maker. The principal knows that there exist *agents*, also called experts (or traders in a market context), each possessing a quantifiable subjective

input bearing on that event; it does not know the values of these inputs but believes that if it could extract and aggregate them in some principled way, the result would be the best possible guess about the outcome.

The three broad themes that I address in the subsequent chapters are as follows:

- (1) *Incentives in information aggregation.* In many situations, the agents might have no intrinsic motivation to just hand over their subjective inputs to the principal; e.g. an agent might be a meteorologist expecting compensation for reporting to the principal the numerical estimate of the global average temperature ten years into the future that is the product of her education, effort, and expertise. In such cases, it is imperative to offer agents monetary or money-like incentives², but there is a catch: If these incentives are not carefully designed, an agent acting selfishly and rationally might lie to the principal if she deems lying to be in her best interest.

One real-world approach towards providing such participation and truth-telling incentives to agents in the context of aggregating forecasts is to use *prediction markets* (Wolfers and Zitzewitz, 2004; Pennock and Sami, 2007; Arrow et al., 2007) – an umbrella term for a variety of crowdsourced forecasting tools. These online mechanisms incentivize agents by allowing them to place bets for or against outcomes of the uncertain event in question, or to buy and sell shares in a specially designed *financial instrument* whose final monetary worth is tied to the realization of the event (e.g. the instrument could be worth \$1 per share if a Democrat wins the next US presidential race, and is worthless otherwise). A publicly displayed property of the prediction market – such as the betting odds or the price of the instrument – that is updated as and when agents interact with the market (or, equivalently, report to the principal) is interpreted as the principal’s collective forecast. Well-known examples of real-world prediction markets include the Iowa Electronic Markets³, PredictIt⁴, Foresight Exchange⁵, Betfair⁶, and the Hollywood Stock Exchange⁷.

²A money-like incentive can refer to anything other than compensation in real currency that an agent values and is willing to accept in exchange for her input, e.g. raffle tickets, reputation score within a community, etc.

³<https://tippie.biz.uiowa.edu/iem/>

⁴<https://www.predictit.org/>

⁵<http://www.ideosphere.com/>

⁶<https://www.betfair.com/>

⁷<http://www.hsx.com/>

Given the plethora of empirical evidence that prediction markets are at least as effective as more traditional means of forecasting such as surveys, polls, expert opinion elicitation etc. (Wolfers and Zitzewitz, 2004; Graefe and Armstrong, 2011; Cowgill and Zitzewitz, 2015), considerable effort has been devoted to figuring out, in a rigorous formal sense, how these mechanisms function as collectors, aggregators, and disseminators of information / belief, and under what circumstances they could go wrong. These are the major research questions I address in the first (and largest) part of this dissertation. To this end, I focus on a popular family of algorithms used to implement automated prediction markets, called *market scoring rules*, abbreviated as MSRs: (a) I show that, under reasonable assumptions about the strategic nature of participating agents, the market’s output (the price) behaves as an aggregate function belonging to a well-known family called *opinion pools* that can be viewed as a generalization of weighted averaging in some respects; (b) I analyze the extent to which the predictive power of the market is retained when some participating agents do not just have *information* but also *influence* on the forecast event so that the introduction of the prediction market can generate *outcome manipulation* incentives, and suggest a remedial modification to the mechanism.

- (2) *Aggregation with non-strategic agents.* The existence of selfish-rational agents who might try to manipulate an incentivized mechanism is not the only feature of the aggregation problem that makes it hard; the hardness could arise from more fundamental issues such as how agents’ inputs are formulated and what kind of queries are feasible. Suppose that the “hidden truth” that the principal wishes to uncover is a real-valued quantity, and each agent is known to possess a noisy valuation thereof; if the principal is only allowed to ask each agent whether or not her valuation lies above or below a *threshold*, how should it aggregate the binarized signals thus procured so as to arrive at a good estimate of the real-valued target efficiently? Does the principal get to choose its threshold(s) and, if so, how should it do so? These are ideas I explore in the next part of this dissertation, abstracting away from incentive issues, and focusing on learning from differentially informed agents under the assumption that each agent reports a censored noisy version of some ground truth that we are interested in.

- (3) *Markets in practice*. This component of my dissertation, although tied to both of the above strands, deals more directly with practical issues concerning aggregation mechanisms deployed “in the wild” where many aspects of agent behavior not accounted for above come into play. While discussing theme (1), I considered a principal that designs monetary incentives / compensation schemes and an agent who can only choose what input (not necessarily consistent with her actual private information) to report – within a market context, this translates to the market administrator dictating the pricing rule, and a trader only specifying an order size (the quantity she would like to buy or sell). If a trader does not like the current prices offered but does not mind waiting to see if market conditions become favorable, there should be some way to accommodate her, otherwise she might not participate, making her information unavailable to the principal. One way of doing that in real-world financial markets is to allow a trader to specify not just an order size but also a limit on the prices acceptable to her (the highest price she is willing to pay if she is a buyer / the lowest she is willing to accept if she is a seller), and then push such an order into one of two priority queues (for buy and sell orders). How should we modify the market price-setting rule so as to make the aggregation mechanism operate with such agents, and how do those modifications affect the information aggregation characteristics of the system? Moreover, now that we are in a market setting, it also makes sense to evaluate the system not just in terms of its informative / predictive power but also *market quality* properties studied in the economics and finance literatures (e.g. trading volume, social welfare, etc.). A significant methodological difference that this theme has with the first two is that evaluation of this complex environment is primarily based on simulation studies rather than the use of analytical tools.

1.3 Organization of the dissertation

Theme (1) spans Chapters 2 and 3 whereas Chapters 4 and 5 are devoted to themes (2) and (3) respectively. A survey of the existing literature relevant to each chapter as well as its relation to the content is provided in the respective chapter. Below is a gist of each chapter: I cite published material on which the content is based, and indicate additions, if any.

Chapter 2. I begin the chapter by describing the design of the prediction market algorithm called the market scoring rule (Hanson, 2003b) that I touched upon while presenting theme (1) above. I then show that, assuming the agents participating in such a market to be *risk-averse* (Mas-Colell et al., 1995), the *instantaneous price* which can be viewed as the “output” of this mechanism is formally equivalent to an *opinion pool* (Genest and Zidek, 1986) applied to the private inputs of the agents. After establishing this result under the most general definition of risk-aversion, I deduce further restrictions on our risk-averse agent model that make the price behave as one of the more familiar members of the opinion pool family (specifically a weighted arithmetic mean and a renormalized weighted geometric mean), and show that we can even interpret the market mechanism as a Bayesian learner for these models.

This is joint work with Dr. Sanmay Das, and is based mainly on our paper (Chakraborty and Das, 2015) that was accepted for a spotlight presentation at the 29th Annual Conference on Neural Information Processing Systems (NIPS 2015). In this dissertation, I provide detailed proofs of all theorems as well as additional experiments (Chapter 2 Section 2.4.3).

Chapter 3. In this chapter, I delineate a new game-theoretic model of MSR-based prediction markets that Dr. Das and I built to address the following scenario: Market participants have power to directly affect the forecast event so that the incentives brought into the picture by the introduction of a prediction market could potentially induce some of them to distort the outcome that the market was set up to predict but not alter in any way (e.g. a referee in a game of basketball could have stakes in a prediction market that intends to forecast the winner). I prove that, if some participants have a non-zero probability of being non-strategic, the game-theoretic equilibrium of this process is one of two types, depending on this probability – either collusive and uninformative or partially revealing. Finally, I show how the compensation scheme can be modified by incorporating ideas from the literature on *peer prediction* (Miller et al., 2005) – a way of providing truth-telling incentives to agents based on what their reports tell us about the reports of their peers (other agents participating in the aggregation mechanism) when the hidden truth / uncertain event is never revealed to the principal – to counteract undesirable incentives.

This contribution is also co-authored with Dr. Das; parts of it have been published (Chakraborty and Das, 2016) at the 25th International Joint Conference on Artificial Intelligence (IJCAI 2016), and presented at the 4th Workshop on Social Computing and User Generated Content (SCUGC 2014) held in conjunction with the 15th ACM Conference on Economics and Computation (EC 2014), at the 2nd Collective Intelligence Conference (CI 2014), and also as an invited talk at the 20th Conference of the International Federation of Operational Research Societies, 2014 (IFORS 2014). I give proofs of all theorems as well as interesting corollaries and generalizations of the model (Chapter 3 Section 3.7) that have not been published before.

Chapter 4. Here, I discuss the problem of zeroing in on a real-valued target by observing only a sequence of censored (binarized) Gaussian samples where the principal (or learner) gets to choose the threshold that determines how each sample is binarized; this is followed by the design and analysis of an approximately Bayesian algorithm that achieves this result in an asymptotically near-optimal time with respect to certain problem parameters (i.e. almost as good as if we had access to the uncensored signals). I share authorship for this work, which was published at the 27th Conference on Uncertainty in Artificial Intelligence (UAI 2011), with Dr. Das and Dr. Malik Magdon-Ismail; I give the proof for the main theorem in this dissertation.

Chapter 5. I describe in this chapter an adaptation of the most commonly used variety (logarithmic) of MSRs, which takes *market orders* only (i.e. trading agents can only pick a quantity to buy or sell), to a setting with *limit orders* (i.e. trading agents can state a quantity as well as a limiting price), and an experimental set-up that we used to uncover interesting properties of the resulting market ecosystem and compare to benchmarks.

The contents of this chapter were published as a paper at the 29th AAAI Conference on Artificial Intelligence (AAAI 2015) jointly authored with Dr. Das and Justin Peabody; I have added some expository notes to underscore the significance of this chapter to my dissertation.

Chapter 6. In the final chapter, I summarize the contributions made in this dissertation; possible directions for future research relevant to each theme are already pointed out in the respective chapter(s).

Chapter 2

Semantics of aggregation with market scoring rules

2.1 Introduction

Suppose that the event on which the principal needs a forecast is modeled as a random variable and each agent's private input as a personal probability (or distribution) on that variable. One simple, principled approach towards achieving aggregation here is the *opinion pool* (OP) which directly solicits inputs from these agents and then maps this vector of inputs to a single probability (or distribution) based on certain axioms (Genest and Zidek (1986)). However, this technique abstracts away from the issue of providing proper incentives to a selfish-rational agent to reveal her private information honestly. Financial markets approach the problem differently, offering financial incentives for traders to supply their information about valuations and aggregating this information into informative prices. A *prediction market* is a relatively novel tool that builds upon this idea, offering trade in a financial security whose final monetary worth is tied to the future revelation of some currently unknown ground truth.

The branch of finance that studies the rules for exchanging financial assets and their impact on observable properties of the market is called *market microstructure* (Krishnamurti, 2009), and this term also sometimes refers to a specific set of such rules, which is the sense in which we will use it in the rest of the dissertation. Hanson (2003b) introduced a class of microstructures for automated prediction markets called *market scoring rules* (MSRs) of which the Logarithmic Market Scoring Rule (LMSR) is arguably the most widely used and

well-studied variety. An MSR effectively acts as a cost function-based *market maker* – a market maker being a type of intermediary on a trading platform always willing to take the other side of a trade with any willing buyer *or* seller – re-adjusting its quoted price after every transaction.

One of the most attractive properties of an MSR is its incentive compatibility for a myopic risk-neutral trader. But this also means that, every time an MSR trades with such an agent, the updated market price is reset to the subjective probability of that agent; the market mechanism itself does not play an active role in unifying pieces of information gleaned from the entire trading history into its current price. Ostrovsky (2012) and Iyer et al. (2014) have shown that, with differentially informed Bayesian risk-neutral and risk-averse agents respectively, trading repeatedly, “information gets aggregated” in an MSR-based market in a perfect Bayesian equilibrium. However, if agent beliefs themselves do not converge, can the price process emerging out of their interaction with an MSR still be viewed as an aggregator of information in some sense? Intuitively, even if an agent does not revise her belief based on her inference about her peers’ information from market history, her conservative attitude towards risk should compel her to trade in such a way as to move the market price not all the way to her private belief but to some function of her belief and the most recent price; thus, the evolving price should always retain some memory of all agents’ information sequentially injected into the market. Therefore, the assumption of belief-updating agents may not be indispensable for providing theoretical guarantees on how the market incorporates agent beliefs. A few attempts in this vein can be found in the literature, typically embedded in a broader context (Sethi and Vaughan, 2016; Abernethy et al., 2014), but there have been few general results; see Section 2.2 for a review.

In this chapter, we develop a new unified understanding of the information aggregation characteristics of a market with risk-averse agents mediated by an MSR, with no regard to how the agents’ beliefs are formed. In fact, we demonstrate an equivalence between such MSR-mediated markets and opinion pools. We do so by first proving that for any MSR interacting with myopic risk-averse traders, the revised instantaneous price after every trade equals the latest trader’s risk-neutral probability conditional on the preceding market state. We then show that this price update rule satisfies an axiomatic characterization of opinion pooling functions from the literature, establishing the equivalence. We identify further conditions on agent strategy under which the market price behaves (exactly or approximately) as specific

types of opinion pool. Our results are reminiscent of similar findings about competitive equilibrium prices in markets with rational, risk-averse agents (Pennock (1999); Beygelzimer et al. (2012); Millin et al. (2012) etc.), but those models require that agents learn from prices and also abstract away from any consideration of microstructure and the dynamics of actual price formation (how the agents would reach the equilibrium is left open). By contrast, our results do not presuppose any kind of generative model for agent signals, and also do not involve an equilibrium analysis – hence they can be used as tools to analyze the convergence characteristics of the market price in non-equilibrium situations with potentially fixed-belief or irrational agents.

2.2 Related work

Seminal work in establishing a formal relationship between asset prices and the private information of trading agents was done by Pennock (1999) who showed that linear and logarithmic opinion pools (see Definition 1 below) arise as special cases of the equilibrium of his intuitive model of securities markets when all agents have generalized logarithmic and negative exponential utilities respectively. More recently, an important line of research (Beygelzimer et al., 2012; Millin et al., 2012; Hu and Storkey, 2014; Storkey et al., 2015) has focused on a competitive equilibrium analysis of prediction markets under various trader models, and found an equivalence between the market’s equilibrium price and the outcome of an opinion pool with the same agents. Unlike these analyses that abstract away from the microstructure, Ostrovsky (2012) and Iyer et al. (2014) show for that certain market structures satisfying mild conditions, including MSRs, the market’s belief measure converges in probability to the ground truth, when we have repeatedly trading and learning agents with risk-neutral and risk-averse utilities respectively. Our contribution, while drawing inspiration from these sources, differs in that we delve into the characteristics of the evolution of the price rather than the properties of prices in equilibrium or upon convergence, and single out the role played by the price-setting rule itself in inducing aggregation with no regard to how agent beliefs are formulated.

While there has also been significant work on market properties for other microstructures such as continuous double auctions or mediation by sophisticated market-making algorithms

(e.g. Dave Cliff (1997); Farmer et al. (2005b); Brahma et al. (2012) and references therein) when the agents are “zero-intelligence” or derivatives thereof (and therefore definitely not Bayesian), this line of literature has not looked at market scoring rules in detail, and analytical results have been rare.

In recent years, the literature focusing on the MSR family has grown substantially. Chen and Vaughan (2010) and Frongillo et al. (2012) have uncovered isomorphisms between this type of market structure and well-known machine learning algorithms. We, on the other hand, are concerned with the similarities between price evolution in MSR-mediated markets and opinion pooling methods (see e.g. Garg et al. (2004)). Our work comes close to that of Sethi and Vaughan (2016) who show analytically that the price sequence of a cost function-based market maker with budget-limited risk-averse traders is “convergent under general conditions”, and by simulation that the limiting price of LMSR with multi-shot but myopic logarithmic utility agents is approximately a linear opinion pool of agent beliefs. Abernethy et al. (2014) show that a risk-averse exponential utility agent with an exponential family belief distribution updates the state vector of a generalization of LMSR that they propose to a convex combination of the current market state vector and the natural parameter vector of the agent’s own belief distribution (see their Theorem 5.2, Corollary 5.3) – this reduces to a logarithmic opinion pool (LogOP) for classical LMSR. The LMSR-LogOP connection was also noted by Pennock and Xia (2011) (in their Theorem 1) but with respect to an artificial probability distribution based on an agent’s observed trade that the authors defined instead of considering traders’ belief structure or strategies. We show how results of this type arise as special cases of a more general MSR-OP equivalence that we establish in Section 2.4 below.

2.3 Model and definitions

The principal is interested in a binary event $X \in \{0, 1\}$, called the *forecast event*, whose outcome will be revealed publicly at a known future date; X can represent a proposition such as “A Democrat will win the next U.S. presidential election” or “The favorite will beat the underdog by more than a pre-determined point spread in a game of football” or “The next *Avengers* movie will hit a certain box office target in its opening week.” The opinion of each of n agents on this event is quantified by her subjective point probability $\pi_i \in (0, 1)$

for X turning out to be 1, $i = 1, 2, \dots, n$. For the time being, we will be agnostic to *how* these probabilities are generated. In such a setting, if the principal had direct access to these π_i -values, the problem would be one of simply unifying them into a consensus or aggregate – this has been tackled in many fields using *opinion pools*.

2.3.1 Opinion Pool (OP)

Opinion pools have been studied for a long time, and various characterizations exist thereof (Genest and Zidek, 1986); here, we present an axiomatic characterization due to Garg et al. (2004) for an aggregate operator of the form $\hat{p} = f(p_1, p_2, \dots, p_n) \in [0, 1]$ that takes as input the vector of probabilistic reports $p_i \in [0, 1]$, $i = 1, 2, \dots, n$ submitted by n agents, also called *experts* in this context.

Definition 1. *A function $f : [0, 1]^n \rightarrow [0, 1]$ is defined as a valid opinion pool for n probabilistic reports if it satisfies the following three criteria.*

1. **Unanimity:** *If all experts agree, the aggregate also agrees with them.*
2. **Boundedness:** *The aggregate is bounded by the extremes of the inputs.*
3. **Monotonicity:** *If one expert changes her opinion in a particular direction while all other experts' opinions remain unaltered, then the aggregate changes in the same direction.*

Two popular opinion pooling methods are the Linear Opinion Pool (LinOP) and the Logarithmic Opinion Pool (LogOP) which are essentially a weighted average (or convex combination) and a renormalized weighted geometric mean of the experts' probability reports respectively. For a binary event,

$$\begin{aligned} \text{LinOP}(p_1, p_2, \dots, p_n) &= \sum_{i=1}^n \omega_i^{\text{lin}} p_i, \\ \text{LogOP}(p_1, p_2, \dots, p_n) &= \prod_{i=1}^n p_i^{\omega_i^{\text{log}}} \Big/ \left[\prod_{i=1}^n p_i^{\omega_i^{\text{log}}} + \prod_{i=1}^n (1 - p_i)^{\omega_i^{\text{log}}} \right], \end{aligned}$$

where $\omega_i^{\text{lin}}, \omega_i^{\text{log}} \geq 0 \ \forall i = 1, 2, \dots, n$, $\sum_{i=1}^n \omega_i^{\text{lin}} = 1$, $\sum_{i=1}^n \omega_i^{\text{log}} = 1$.

The following result for recursively defined pooling functions will prove useful for establishing our desired equivalence.

Lemma 1. *For a two-outcome forecasting task, if $f_2(r_1, r_2)$ and $f_{n-1}(q_1, q_2, \dots, q_{n-1})$ are valid opinion pools for two probabilistic reports r_1, r_2 and $n-1$ probabilistic reports q_1, q_2, \dots, q_{n-1} respectively, then $f(p_1, p_2, \dots, p_n) = f_2(f_{n-1}(p_1, p_2, \dots, p_{n-1}), p_n)$ is also a valid opinion pool for n reports.*

Proof. Recall from Definition 1 that a valid opinion pool $\hat{p} = \phi(p_1, p_2, \dots, p_m)$, where $p_1, p_2, \dots, p_m \in [0, 1]$ are reported expert probabilities of occurrence of binary event X , must satisfy

1. **Unanimity:** If $p_i = p \ \forall i = 1, 2, \dots, m$, then $\hat{p} = p$.
2. **Boundedness:** $\min\{p_1, p_2, \dots, p_m\} \leq \hat{p} \leq \max\{p_1, p_2, \dots, p_m\}$.
3. **Monotonicity:** \hat{p} increases monotonically as p_i increases, p_j being held constant $\forall j \neq i, i = 1, 2, \dots, m$, i.e. $\frac{\partial \phi}{\partial p_i} > 0$ everywhere $\forall i$.

By the condition of the lemma, all the above three properties are possessed by each each of f_2 and f_{n-1} , and we need to prove that f has each of these properties, too.

To prove the unanimity of f : Let $p_i = p \ \forall i = 1, 2, \dots, n$. Then,

$$\begin{aligned} f(p, p, \dots, p) &= f_2(f_{n-1}(p, p, \dots, p), p) \\ &= f_2(p, p), \quad \text{by unanimity of } f_{n-1}, \\ &= p, \quad \text{by unanimity of } f_2. \end{aligned}$$

To prove the boundedness of f : Using the upper bounds on f_2 and f_{n-1} ,

$$\begin{aligned} f(p_1, p_2, \dots, p_n) &\leq \max\{f_{n-1}(p_1, p_2, \dots, p_{n-1}), p_n\} \\ &\leq \max\{\max\{p_1, p_2, \dots, p_{n-1}\}, p_n\} \\ &= \max\{p_1, p_2, \dots, p_{n-1}, p_n\}. \end{aligned}$$

Similarly, using the lower bounds on f_2 and f_{n-1} , we can show that $f(p_1, p_2, \dots, p_n) \geq \min\{p_1, p_2, \dots, p_{n-1}, p_n\}$.

To prove the monotonicity of f : The partial derivative of f with respect to each p_i , $i = 1, 2, \dots, n - 1$ is given by

$$\frac{\partial f}{\partial p_i} = \frac{\partial}{\partial p_i} f_2(f_{n-1}(p_1, p_2, \dots, p_{n-1}), p_n) = \frac{\partial f_2(f_{n-1}, p_n)}{\partial f_{n-1}} \cdot \frac{\partial f_{n-1}(p_1, p_2, \dots, p_{n-1})}{\partial p_i} > 0$$

by the monotonicity of f_2 and f_{n-1} with respect to their respective inputs. Similarly,

$$\frac{\partial f}{\partial p_n} = \frac{\partial f_2(f_{n-1}, p_n)}{\partial p_n} > 0$$

by the monotonicity of f_2 . □

2.3.2 Market Scoring Rule (MSR)

However, these experts may not be inclined to reveal their private beliefs to the principal without the promise of any reward in return. This brings us to the issue of *information elicitation*, the first step towards information aggregation from selfish-rational agents. Elicitation at a single-expert level is traditionally accomplished using *proper scoring rules* (Brier, 1950; Good, 1952; Gneiting and Raftery, 2007).

In general, a scoring rule is a function of two variables $s(\mathbf{p}, x) \in \mathbb{R} \cup \{-\infty, \infty\}$, where \mathbf{p} is an agent's probabilistic prediction (density or mass function) about an uncertain event, x is the realized or revealed outcome of that event after the prediction has been made, and the resulting value of s is the agent's *ex post* compensation for prediction. For a binary event X , a scoring rule can just be represented by the pair $(s_1(p), s_0(p))$ which is the vector of agent compensations for $\{X = 1\}$ and $\{X = 0\}$ respectively, $p \in [0, 1]$ being the agent's *reported* probability of $\{X = 1\}$ which may or may not be equal to her *true* subjective probability, say, $\pi = \Pr(X = 1)$.

Assuming the expert to be *risk-neutral*, i.e. she will choose an action that maximizes her subjective expectation of her raw ex post compensation, we know that she will report any

probability value in the set $P_{s,\pi}^* \triangleq \arg \max_{p \in [0,1]} [\pi s_1(p) + (1 - \pi)s_0(p)]$. A scoring rule s is defined to be *proper* if it is *incentive compatible* for a risk-neutral expert; incentive compatibility is the property that it is in the expert cannot do better than reporting truthfully, which in this setting translates to the condition: $\pi \in P_{s,\pi}^* \forall \pi \in [0,1]$. The properness is *strict* if π is the sole maximizer of her expected ex post score.

In addition, a two-outcome scoring rule is *regular* if $s_j(\cdot)$ is real-valued except possibly that $s_0(1)$ or $s_1(0)$ is $-\infty$; any regular strictly proper scoring rule can be written in the following form (Gneiting and Raftery (2007)):

$$s_j(p) = G(p) + G'(p)(j - p), \quad j \in \{0, 1\}, p \in [0, 1], \quad (2.1)$$

$G : [0, 1] \rightarrow \mathbb{R}$ is a strictly convex function with $G'(\cdot)$ as a sub-gradient which is real-valued except possibly that $-G'(0)$ or $G'(1)$ is ∞ ; if $G(\cdot)$ is differentiable in $(0, 1)$, $G'(\cdot)$ is simply its derivative.

A classic example of a regular strictly proper scoring rule is the logarithmic scoring rule:

$$s_1(p) = b \ln p; \quad s_0(p) = b \ln(1 - p), \quad \text{where } b > 0 \text{ is a free parameter.} \quad (2.2)$$

In principle, if there are n risk-neutral experts, the principal could promise each of them a reward according to any strictly proper scoring rule, hence elicit their honest reports separately, and then combine these reports perhaps using an opinion pooling method; however, the principal's total payout ("loss" or cost of information acquisition) would then be $O(n)$. To circumvent this issue, Hanson (2003b) introduced an extension of a scoring rule wherein the principal initiates the process of information elicitation by making a baseline report p_0 , and then elicits publicly declared reports p_i sequentially from n agents; the ex post compensation $c_x(p_i, p_{i-1})$ received by agent i from the principal, where x is the realized outcome of event X , is the difference between the scores assigned to the reports made by herself and her predecessor:

$$c_x(p_i, p_{i-1}) \triangleq s_x(p_i) - s_x(p_{i-1}), \quad x \in \{0, 1\}. \quad (2.3)$$

If each agent acts non-collusively, risk-neutrally, and myopically (as if her current interaction with the principal is her last), then the incentive compatibility property of a strictly proper

score still holds for the sequential version. Moreover, it is easy to show that the principal’s worst-case payout (loss) is bounded regardless of agent behavior. In particular, for the two-outcome logarithmic score, the loss bound for $p_0 = 1/2$ is $b \ln 2$; b can be referred to as the principal’s *loss parameter*.

Definition 2. *We call a market scoring rule well-behaved if the underlying scoring rule is regular and strictly proper, and the associated convex function $G(\cdot)$ (as in (2.1)) is continuous and thrice-differentiable, with $0 < G''(p) < \infty$ and $|G'''(p)| < \infty$ for $0 < p < 1$.*

A sequentially shared strictly proper scoring rule of the above form can also be interpreted as a cost function-based prediction market mechanism offering trade in an Arrow-Debreu (i.e. $(0, 1)$ -valued) security written on the event X , hence the name “market scoring rule”. The cost function is a strictly convex function of the total outstanding quantity of the security that determines all execution costs; its first derivative (the cost per share of buying or the proceeds per share from selling an infinitesimal quantity of the security) is called the market’s “instantaneous price”, and can be interpreted as the market maker’s current *risk-neutral probability* (Chen and Pennock (2007)) for $\{X = 1\}$, the starting price being equal to the principal’s baseline report p_0 . Trading occurs in discrete episodes $1, 2, \dots, n$, in each of which an agent orders a quantity of the security to buy or sell given the market’s cost function and the (publicly displayed) instantaneous price. Since there is a one-to-one correspondence between agent i ’s order size and p_i , the market’s revised instantaneous price after trading with agent i , an agent’s “action” or trading decision in this setting is identical to making a probability report by selecting a $p_i \in [0, 1]$. If agent i is risk-neutral, then p_i is, by design, her subjective probability π_i . This view of MSRs is useful for operational purposes but not relevant to the theoretical results in this chapter; please refer to Hanson (2003b); Chen and Pennock (2007) for further details. However, we will return to it in Chapter 5 when we describe experiments with an algorithmic trading agent based on extensions to the logarithmic market scoring rule (LMSR).

2.4 MSR behavior with risk-averse myopic agents

We first present general results on the connection between sequential trading in an MSR-mediated market with agents having risk-averse utility (see below) and opinion pooling, and

then give a more detailed picture for representative utility functions without and with budget constraints in Sections 2.4.1, 2.4.2, and 2.4.3.

Suppose that, in addition to a belief $\pi_i = \Pr(X = 1)$, each agent i has a continuous utility function of wealth $u_i(c)$, where $c \in [c_i^{\min}, \infty]$ denotes her (ex post) wealth, i.e. her net compensation from the market mechanism after the realization of X defined in (2.3), and $c_i^{\min} \in [-\infty, 0]$ is her minimum acceptable wealth (a negative value suggests tolerance of debt); $u_i(\cdot)$ satisfies the usual criteria of *non-satiation* i.e. $u'_i(c) > 0$ except possibly that $u'_i(\infty) = 0$, and *risk aversion*, i.e. $u''_i(c) < 0$ except possibly that $u''_i(\infty) = 0$, through out its domain (Mas-Colell et al., 1995); in other words $u_i(\cdot)$ is strictly increasing and strictly concave. Additionally, we require its first two derivatives to be finite and continuous on $[c_i^{\min}, \infty]$ except that we tolerate $u'_i(c_i^{\min}) = \infty$, $u''_i(c_i^{\min}) = -\infty$. Note that, by choosing a finite lower bound c_i^{\min} on the agent's wealth, we can account for any starting wealth or budget constraint that effectively restricts the agent's action space.

Lemma 2. *If $|c_i^{\min}| < \infty$, then there exist lower and upper bounds, $p_i^{\min} \in [0, p_{i-1}]$ and $p_i^{\max} \in [p_{i-1}, 1]$ respectively, on the feasible values of the price p_i to which agent i can drive the market regardless of her belief π_i , where $p_i^{\min} = s_1^{-1}(c_i^{\min} + s_1(p_{i-1}))$ and $p_i^{\max} = s_0^{-1}(c_i^{\min} + s_0(p_{i-1}))$.*

Proof. Agent i 's ex post wealth for trading in such a way as to revise the market price from p_{i-1} to any $\tilde{p} \in [0, 1]$ is $c_x(\tilde{p}, p_{i-1})$ for outcome x but, from the constraints imposed by the utility function, this wealth cannot be smaller than c_i^{\min} for any x . Thus,

$$\begin{aligned} c_1(\tilde{p}, p_{i-1}) &\geq c_i^{\min} \\ \Rightarrow s_1(\tilde{p}) - s_1(p_{i-1}) &\geq c_i^{\min} \\ \Rightarrow s_1(\tilde{p}) &\geq c_i^{\min} + s_1(p_{i-1}) \\ \Rightarrow \tilde{p} &\geq s_1^{-1}(c_i^{\min} + s_1(p_{i-1})) = p_i^{\min}, \end{aligned}$$

since $s_1(\cdot)$ is strictly increasing (hence invertible). Also, since $c_i^{\min} \leq 0$,

$$s_1(p_i^{\min}) = c_i^{\min} + s_1(p_{i-1}) \implies s_1(p_i^{\min}) \leq s_1(p_{i-1}) \implies p_i^{\min} \leq p_{i-1}.$$

Similarly, from the inequality $c_0(\tilde{p}, p_{i-1}) \geq c_i^{\min}$ and the decreasing monotonicity of $s_0(\cdot)$, we can show that $\tilde{p} \leq s_0^{-1}(c_i^{\min} + s_0(p_{i-1})) = p_i^{\max} \geq p_{i-1}$. \square

Since the latest price p_{i-1} can be viewed as the market's current "state" from myopic agent i 's perspective, the agent's final utility depends not only on her own action p_i and the extraneously determined outcome x but also on the current market state p_{i-1} she encounters. The optimal action of myopic risk-averse agent i is then given by

$$p_i = \arg \max_{p \in [0,1]} [\pi_i u_i(c_1(p, p_{i-1})) + (1 - \pi_i) u_i(c_0(p, p_{i-1}))].$$

This leads us to the main result of this section, Theorem 1. Here, we sketch and discuss the major implications of the theorem; a detailed proof can be found in Appendix A Section A.1.

Theorem 1. *If a well-behaved market scoring rule for an Arrow-Debreu security with a starting instantaneous price $p_0 \in (0, 1)$ trades with a sequence of n myopic agents with subjective probabilities $\pi_1, \dots, \pi_n \in (0, 1)$ and risk-averse utility functions of wealth $u_1(\cdot), \dots, u_n(\cdot)$ as above, then the updated market price p_i after every trading episode $i \in \{1, 2, \dots, n\}$ is equivalent to a valid opinion pool for the market's initial baseline report p_0 and the subjective probabilities $\pi_1, \pi_2, \dots, \pi_i$ of all agents who have traded up to (and including) that episode.*

Proof sketch. For every trading episode i , by setting the first derivative of agent i 's expected utility to zero, and analyzing the resulting equation, we can arrive at the following lemmas.

Lemma 3. *Under the conditions of Theorem 1, if $p_{i-1} \in (0, 1)$, then the revised price p_i after agent i trades is the unique solution in $(0, 1)$ to the fixed-point equation:*

$$p_i = \frac{\pi_i u'_i(c_1(p_i, p_{i-1}))}{\pi_i u'_i(c_1(p_i, p_{i-1})) + (1 - \pi_i) u'_i(c_0(p_i, p_{i-1}))}. \quad (2.4)$$

Since $p_0 \in (0, 1)$, and $\pi_i \in (0, 1) \forall i$, p_i is also confined to $(0, 1) \forall i$, by induction.

Lemma 4. *The implicit function $p_i(p_{i-1}, \pi_i)$ described by (A.4) has the following properties:*

1. $p_i = \pi_i$ (or p_{i-1}) if and only if $\pi_i = p_{i-1}$.
2. $0 < \min\{p_{i-1}, \pi_i\} < p_i < \max\{p_{i-1}, \pi_i\} < 1$ whenever $\pi_i \neq p_{i-1}$, $0 < \pi_i, p_{i-1} < 1$.
3. For any given p_{i-1} (resp. π_i), p_i is a strictly increasing function of π_i (resp. p_{i-1}).

Evidently, properties 1, 2, and 3 above correspond to axioms of unanimity, boundedness, and monotonicity respectively (see Definition 1). Hence, $p_i(p_{i-1}, \pi_i)$ is a valid opinion pooling function for p_{i-1}, π_i . Finally, since (A.4) defines the opinion pool p_i recursively in terms of $p_{i-1} \forall i = 1, 2, \dots, n$, we can invoke Lemma 1 to obtain the desired result. \square

There are several points worth noting about this result.

- Since the updated market price p_i is also equivalent to agent i 's action (Section 2.3.2), the R.H.S. of (A.4) is agent i 's risk-neutral probability (Pennock (1999)) of $\{X = 1\}$, given her utility function, her action, and the current market state. Thus, Lemma 3 is a natural extension of the elicitation properties of an MSR. MSRs, by design, elicit subjective probabilities from risk-neutral agents in an incentive compatible manner; we show that, in general, they elicit risk-neutral probabilities when they interact with risk-averse agents. Lemma 3 is also consistent with the observation of Pennock (1999) that, for all belief elicitation schemes based on monetary incentives, an external observer can only assess a participant's risk-neutral probability uniquely; she cannot discern the participant's belief and utility separately.
- Observe that this pooling operation is accomplished by an MSR even without direct revelation.
- Notice the presence of the market maker's own initial baseline p_0 as a component in the final aggregate; however, for the examples we study below, the impact of p_0 diminishes with the participation of more and more informed agents, and we conjecture that this is a generic property.

In general, the exact form of this pooling function is determined by the complex interaction between the MSR and agent utility, and a closed form of p_i from (A.4) might not be attainable in many cases. However, given a particular MSR, we can venture to identify agent utility functions which give rise to well-known opinion pools. Hence, for the rest of this paper, we focus on the logarithmic market scoring rule (LMSR), one of the most popular tools for implementing real-world prediction markets.

For the logarithmic market scoring rule (LMSR),

$$c_1(p_i, p_{i-1}) = b \ln \left(\frac{p_i}{p_{i-1}} \right), \quad c_0(p_i, p_{i-1}) = b \ln \left(\frac{1 - p_i}{1 - p_{i-1}} \right)$$

so that equation (A.4) can be rewritten as

$$\frac{p_i}{1 - p_i} = \frac{\pi_i}{1 - \pi_i} \cdot \frac{u'_i \left(b \ln \left(\frac{p_i}{p_{i-1}} \right) \right)}{u'_i \left(b \ln \left(\frac{1 - p_i}{1 - p_{i-1}} \right) \right)}. \quad (2.5)$$

2.4.1 LMSR and constant absolute risk aversion (CARA) utility: LogOP

Theorem 2. *The only risk-averse utility function for which myopic agent i , having a subjective belief $\pi_i \in (0, 1)$, and trading with an LMSR market with parameter b and current instantaneous price p_{i-1} , results in the market's updated price p_i being identical to a logarithmic opinion pool between the current price and the agent's subjective belief, i.e.*

$$p_i = \pi_i^{\alpha_i} p_{i-1}^{1-\alpha_i} / \left[\pi_i^{\alpha_i} p_{i-1}^{1-\alpha_i} + (1 - \pi_i)^{\alpha_i} (1 - p_{i-1})^{1-\alpha_i} \right], \quad \alpha_i \in (0, 1), \quad (2.6)$$

is given by

$$u_i(c) = \tau_i (1 - \exp(-c/\tau_i)), \quad c \in \mathbb{R} \cup \{-\infty, \infty\}, \quad \text{constant } \tau_i \in (0, \infty), \quad (2.7)$$

the aggregation weight is $\alpha_i = \frac{\tau_i/b}{1 + \tau_i/b}$.

The proof is in Appendix A Section A.1.1. Note that (A.9) is a standard formulation of the CARA (or negative exponential) utility function with *risk tolerance* τ_i ; smaller the value of τ_i , higher is agent i 's aversion to risk. The unbounded domain of $u_i(\cdot)$ indicates a lack of budget constraints; risk aversion comes about from the fact that the range of the function is bounded above (by its risk tolerance τ_i) but not bounded below.

Moreover, the LogOP equation (A.8) can alternatively be expressed as a linear update in terms of *log-odds* (i.e. *logit functions* of probabilities), another popular means of formulating

one's belief about a binary event:

$$l(p_i) = \alpha_i l(\pi_i) + (1 - \alpha_i) l(p_{i-1}), \quad l(p) = \ln \left(\frac{p}{1-p} \right) \in [-\infty, \infty] \quad \text{for } p \in [0, 1]. \quad (2.8)$$

Aggregation weight and risk tolerance: Since α_i is an increasing function of an agent's risk tolerance relative to the market's loss parameter (the latter being, in a way, a measure of how much risk the market maker is willing to take), identity (2.8) implies that the higher an agent's risk tolerance, the larger is the contribution of her belief towards the changed market price, which agrees with intuition. Also note the interesting manner in which the market's loss parameter effectively scales down an agent's risk tolerance, enhancing the inertia factor $(1 - \alpha_i)$ of the price process.

Bayesian interpretation: The Bayesian interpretation of LogOP in general is well-known (Bordley, 1982); we restate it here in a form that is more appropriate for our prediction market setting. We can recast (A.8) as

$$p_i = \frac{p_{i-1} \left(\frac{\pi_i}{p_{i-1}} \right)^{\alpha_i}}{p_{i-1} \left(\frac{\pi_i}{p_{i-1}} \right)^{\alpha_i} + (1 - p_{i-1}) \left(\frac{1-\pi_i}{1-p_{i-1}} \right)^{\alpha_i}}.$$

This shows that, over the i^{th} trading episode $\forall i$, the LMSR-CARA agent market environment is equivalent to a Bayesian learner performing inference on the point estimate of the probability of the forecast event X , starting with the common-knowledge prior $\Pr(X = 1) = p_{i-1}$, and having direct access to π_i (which corresponds to the “observation” for the inference problem), the likelihood function associated with this observation being $\mathcal{L}(X = x | \pi_i) \propto \left| \frac{1-x-\pi_i}{1-x-p_{i-1}} \right|^{\alpha_i}$, $x \in \{0, 1\}$.

Sequence of one-shot traders: If all n agents in the system have CARA utilities with potentially different risk tolerances, and trade with LMSR myopically only once each in the order $1, \dots, n$, then the “final” market log-odds after these n trades, on unfolding the recursion in (2.8), is given by $l(p_n) = \tilde{\alpha}_0^n l(p_0) + \sum_{i=1}^n \tilde{\alpha}_i^n l(\pi_i)$. This is a LogOP where $\tilde{\alpha}_0^n = \prod_{i=1}^n (1 - \alpha_i)$ determines the inertia of the market's initial price, which diminishes as more and more traders interact with the market, and $\tilde{\alpha}_j^n$, $j \geq 1$ quantifies the degree to which an individual trader impacts the final (aggregate) market belief; $\tilde{\alpha}_j^n = \alpha_j \prod_{i=j+1}^n (1 - \alpha_i)$, $j = 1, \dots, n-1$, and $\tilde{\alpha}_n^n = \alpha_n$.

Interestingly, the weight of an agent's belief depends not only on her own risk tolerance but also on those of all agents succeeding her in the trading sequence (lower weight for a more risk tolerant successor, *ceteris paribus*), and is independent of her predecessors' utility parameters. This is sensible since, by the design of an MSR, trader i 's belief-dependent action influences the action of each of (rational) traders $i + 1, i + 2, \dots$ so that the action of each of these successors, in turn, has a role to play in determining the market impact of trader i 's belief.

In particular, if $\tau_j = \tau > 0 \ \forall j \geq 1$, then the aggregation weights satisfy the inequalities $\tilde{\alpha}_{j+1}^n / \tilde{\alpha}_j^n = 1 + \tau/b > 1 \ \forall j = 1, \dots, n-1$, i.e. LMSR assigns progressively higher weights to traders arriving later in the market's lifetime when they all exhibit identical constant risk aversion. This seems to be a reasonable aggregation principle in most scenarios wherein the amount of information in the world improves over time. Moreover, in this situation, $\tilde{\alpha}_1^n / \tilde{\alpha}_0^n = \tau/b$ which indicates that the weight of the market's baseline belief in the aggregate may be higher than those of some of the trading agents if the market maker has a comparatively high loss parameter. This strong effect of the trading sequence on the weights of agents' beliefs is a significant difference between the one-shot trader setting and the market equilibrium setting where each agent's weight is independent of the utility function parameters of her peers.

Convergence: If agents' beliefs are themselves independent samples from the same distribution \mathcal{P} over $[0, 1]$, i.e. $\pi_i \sim_{\text{i.i.d.}} \mathcal{P} \ \forall i$, then by the sum laws of expectation and variance,

$$\mathbb{E}[l(p_n)] = \tilde{\alpha}_0^n l(p_0) + (1 - \tilde{\alpha}_0^n) \mathbb{E}_{\pi \sim \mathcal{P}}[l(\pi)]; \quad \text{Var}[l(p_n)] = \text{Var}_{\pi \sim \mathcal{P}}[l(\pi)] \sum_{i=1}^n (\tilde{\alpha}_i^n)^2.$$

Hence, using an appropriate concentration inequality (Boucheron et al. (2004)) and the properties of the $\tilde{\alpha}_i^n$'s, we can show that, as n increases, the market log-odds ratio $l(p_n)$ converges to $\mathbb{E}_{\pi \sim \mathcal{P}}[l(\pi)]$ with a high probability; this convergence guarantee does not require the agents to be Bayesian.

2.4.2 LMSR and an atypical utility with decreasing absolute risk aversion: LinOP

Theorem 3. *The only risk-averse utility function for which myopic agent i , having a subjective belief $\pi_i \in (0, 1)$, and trading with an LMSR market with parameter b and current instantaneous price p_{i-1} , results in the market's updated price p_i , results in the market's updated price p_i being identical to a linear opinion pool between the current price and the agent's subjective belief, i.e.*

$$p_i = \beta_i \pi_i + (1 - \beta_i) p_{i-1}, \quad \text{for some constant } \beta_i \in (0, 1), \quad (2.9)$$

is given by

$$u_i(c) = \ln(\exp((c + B_i)/b) - 1), \quad c \geq -B_i, \quad (2.10)$$

where $B_i > 0$ represents agent i 's budget, the aggregation weight being $\beta_i = 1 - \exp(-B_i/b)$.

The proof is in Appendix A Section A.1.2. To the best of our knowledge, the above atypical utility function in (A.13) has not been described before: Its domain is bounded below, and it below, and it possesses a positive, strictly decreasing Arrow-Pratt absolute risk aversion measure (Mas-Colell et al., 1995) $A_i(c) = -u_i''(c)/u_i'(c) = \frac{1}{b(\exp((c+B_i)/b)-1)}$ for any $b, B_i > 0$.

Note that, unlike in Theorem 2, the equivalence here requires the agent utility function to depend on the market maker's loss parameter b (the scaling factor in the exponential). Since the microstructure is assumed to be common knowledge, as in traditional MSR settings, the consideration of an agent utility that takes into account the market's pricing function is not unreasonable.

Since the domain of utility function (A.13) is bounded below, we can derive π_i -independent bounds on possible values of p_i from Lemma 2: $p_i^{\min} = (1 - \beta_i)p_{i-1}$, $p_i^{\max} = \beta_i + (1 - \beta_i)p_{i-1}$. Hence, equation (A.12) becomes $p_i = \pi_i p_i^{\max} + (1 - \pi_i)p_i^{\min}$, i.e. the revised price is a linear interpolation between the agent's price bounds, her subjective probability itself acting as the interpolation factor.

Aggregation weight and budget constraint: Evidently, the aggregation weight of agent i 's belief, $\beta_i = (1 - \exp(-B_i/b))$, is an increasing function of her budget normalized with respect to the market's loss parameter; it is, in a way, a measure of her relative risk tolerance. Thus, broad characteristics analogous to the ones in Section 2.4.1 apply to these aggregation weights as well, with the log-odds ratio replaced by the actual market price.

Bayesian interpretation: Under the mild technical assumption that agent i 's belief $\pi_i \in (0, 1)$ is rational, and her budget $B_i > 0$ is such that $\beta_i \in (0, 1)$ is also rational, it is possible to obtain positive integers r_i, N_i and a positive rational number m_{i-1} such that $\pi_i = r_i/N_i$ and $\beta_i = N_i/(m_{i-1} + N_i)$. Then, we can rewrite the LinOP equation (A.12) as $p_i = \frac{r_i + p_{i-1}m_{i-1}}{m_{i-1} + N_i}$, which is equivalent to the posterior expectation of a beta-binomial Bayesian inference procedure described as follows: The forecast event X is modeled as the (future) final flip of a biased coin with an unknown probability of heads. In episode i , the principal (or aggregator) has a prior distribution $\text{BETA}(\mu_{i-1}, \nu_{i-1})$ over this probability, with $\mu_{i-1} = p_{i-1}m_{i-1}$, $\nu_{i-1} = (1 - p_{i-1})m_{i-1}$. Thus, p_{i-1} is the prior mean and m_{i-1} the corresponding ‘‘pseudo-sample size’’ parameter. Agent i is non-Bayesian, and her subjective probability π_i , accessible to the aggregator, is her maximum likelihood estimate associated with the (binomial) likelihood of observing r_i heads out of a private sample of N_i independent flips of the above coin (N_i is common knowledge). Note that m_{i-1}, N_i are measures of certainty of the aggregator and the trading agent respectively, and the latter's normalized budget $B_i/b = \ln(1 + N_i/m_{i-1})$ becomes a measure of her certainty relative to the aggregator's current state in this interpretation.

Sequence of one-shot traders and convergence: If all agents have utility (A.13) with potentially different budgets, and trade with LMSR myopically once each, then the final aggregate market price is given by $p_n = \tilde{\beta}_0^n p_0 + \sum_{i=1}^n \tilde{\beta}_i^n \pi_i$, which is a LinOP where $\tilde{\beta}_0^n = \prod_{i=1}^n (1 - \alpha_i)$, $\tilde{\beta}_j^n = \beta_j \prod_{i=j+1}^n (1 - \beta_i) \forall j = 1, \dots, n-1$, $\tilde{\beta}_n^n = \beta_n$. Again, all intuitions about $\tilde{\alpha}_j^n$ from Section 2.4.1 carry over to $\tilde{\beta}_j^n$. Moreover, if $\pi_i \sim_{\text{i.i.d.}} \mathcal{P} \forall i$, then we can proceed exactly as in Section 2.4.1 to show that, as n increases, p_n converges to $\mathbb{E}_{\pi \sim \mathcal{P}} [\pi]$ with a high probability.

Implications for logarithmic utility: Theorem 3 is somewhat surprising since it is logarithmic utility that has traditionally been found to effect a LinOP in a market equilibrium (Pennock, 1999; Beygelzimer et al., 2012; Storkey et al., 2015). Of course, our results do

not pertain to an equilibrium / convergence setting, but in light of similarities (elaborated on in Section 2.4.3) between utility function (A.13) and logarithmic utility, it is perhaps not unreasonable to ask whether the logarithmic utility-LinOP connection is still maintained approximately for LMSR price evolution under some conditions.

2.4.3 LMSR and logarithmic utility

In this section, we shall explore the idea mentioned above in Section 2.4.2 that agents with logarithmic utility induce an approximate linear opinion pool in a LMSR market under certain conditions.

Comparison of utility function (A.13) with logarithmic utility: The two utility functions under consideration are

$$\begin{aligned} u_{\text{atyp}}(c; B, b) &= \ln(\exp((c + B)/b) - 1), \quad c \geq B, \\ u_{\log}(c; w) &= \ln(c + w), \quad c \geq w \end{aligned}$$

where constants $B, w \in (0, \infty)$ are the respective budgets. First note that both are strictly increasing and strictly concave functions with domains bounded below and decreasing absolute risk aversion, since the respective Arrow-Pratt measures are $A_{\text{atyp}}(c; B, b) = \frac{1}{b(\exp((c+B)/b)-1)}$ and $A_{\log}(c; w) = \frac{1}{c+w}$. Moreover, $u_{\text{atyp}}(c)$ behaves approximately as a logarithmic utility for small values of $(c + B)/b$ and as a linear utility (corresponding to risk-neutrality) for large values thereof.

$$\begin{aligned} (c + B)/b \ll 1 &\implies u_{\text{atyp}}(c; B, b) \approx \ln(1 + (c + B)/b - 1) = \ln(c + B) - \ln b \equiv \ln(c + B); \\ (c + B)/b \gg 1 &\implies u_{\text{atyp}}(c; B, b) \approx \ln \exp((c + B)/b) = (c + B)/b \equiv c, \end{aligned}$$

using first order approximations, and applying the fact that a utility function is (strategically) equivalent to any positive affine transformation of itself.

We provide a visual contrast of the above utility functions in Figure 2.1: Note that for $b = B = 1$, the two functions are very close to each other for small (negative and close to $-B$) values of wealth c . From the graphs, it appears to be a reasonable conjecture that the




Figure 2.1: Comparison of a logarithmic utility function $u_{\log}(c; B) = \ln(c + B)$, $c \geq -B$, where $B = 1$ is the (positive) budget, with various instances of the atypical decreasing absolute risk aversion utility function (A.13) $u_{\text{atyp}}(c; B, b) = \ln(\exp((c + B)/b) - 1)$ with the same budget $B = 1$ but different scaling factors $b = 0.1, 1, 10$.

two utility functions are most similar, in the sense that the switch in the nature of (A.13) from approximately logarithmic to approximately linear occurs at a higher value of wealth, for values of b that are comparable to B .

Proposition 1. *For a myopic agent with a subjective probability $\pi_i \in (0, 1)$ and a logarithmic utility function with budget $w_i \in (0, \infty)$, i.e.*

$$u_i(c) = \ln(w_i + c), \quad c \geq -w_i,$$

the updated instantaneous price of a LMSR market with loss parameter b after interaction with the agent can be written as

$$p_i = \hat{p}_i + \Delta, \quad (2.11)$$

where \hat{p}_i is a LinOP of π_i and p_{i-1} given by

$$\hat{p}_i = (1 - \exp(-\tilde{w}_i))\pi_i + \exp(-\tilde{w}_i)p_{i-1}, \quad \tilde{w}_i = w_i/b,$$

and the error term is

$$\Delta = \pi_i(1 - p_i) \sum_{j=2}^{\infty} \frac{1}{j} \left(\frac{p_i^{\max} - p_i}{1 - p_i} \right)^j - (1 - \pi_i)p_i \sum_{j=2}^{\infty} \frac{1}{j} \left(\frac{p_i - p_i^{\min}}{p_i} \right)^j,$$

with $p_i^{\min} = p_{i-1} \exp(-\tilde{w}_i)$ and $p_i^{\max} = 1 - (1 - p_{i-1}) \exp(-\tilde{w}_i)$ being the lower and upper bounds on the price p_i imposed by the budget constraint.

The proof is in Appendix 2 Section A.2.

Approximation of actual p_i by \hat{p}_i : If, instead of the Maclaurin series in the above proof of Proposition 1, we had used the first-order approximation $-\ln(1 - x) \approx x$, which is reasonable for $|x| \ll 1$,⁸ we would have obtained $p_i \approx \hat{p}_i$. Informally, the smaller the agent's normalized budget \tilde{w}_i , the smaller the range $[p_i^{\min}, p_i^{\max}]$ of feasible values of p_i , hence the smaller the fractions $(p_i - p_i^{\min})/p_i$ and $(p_i^{\max} - p_i)/(1 - p_i)$ are, hopefully leading to a better approximation. But this might not even be necessary for achieving a small magnitude of Δ_i which is the difference of two terms of comparable orders. On eyeballing the expression for Δ_i , it appears to be roughly two orders of magnitude smaller than \hat{p}_i . Since the exact dependence of the approximation error on the value of \tilde{w}_i is hard to figure out analytically, we adopt a simulation-based approach towards exploring this relationship, described below. But, before that, we perform a quick sanity check on the approximation under consideration. From (A.19), it is evident that

$$\lim_{\pi_i \searrow 0} p_i = p_i^{\min} = \lim_{\pi_i \searrow 0} \hat{p}_i; \quad \lim_{\pi_i \nearrow 1} p_i = p_i^{\max} = \lim_{\pi_i \nearrow 1} \hat{p}_i,$$

⁸Note that the relative error of the linear approximation of the logarithmic function, i.e. $\left| \frac{x - f(x)}{f(x)} \right|$, where $f(x) = -\ln(1 - x)$, is at most 10% for $x \leq 0.193$.

indicating that the actual and approximate updated market prices coincide for extreme agent beliefs.

Experiments: We ran 5×9 sets of 1000 simulations each for getting a rough idea about the quality of the approximation $p_i \approx \hat{p}_i$. For each simulation, we generated a sequence of $n = 100$ agents defined by their time-invariant belief-budget pairs $\{(\pi_i, w_i)\}_{i=1}^n$. Since the parameter of interest is the normalized budget \tilde{w}_i , the exact value of the LMSR loss parameter b is immaterial, and we set it to 1. We sampled the \tilde{w}_i 's uniformly at random from the interval $[0, \tilde{w}_{\max}]$, $\tilde{w}_{\max} \in \{0.1, 0.2, 0.25, 0.5, 0.75\}$. The beliefs were random samples from the distribution $\text{BETA}(p_{\text{true}}, 1 - p_{\text{true}})$, $p_{\text{true}} \in \{0.1, 0.2, \dots, 0.9\}$. Thus our knowledge model was that there was a “true” underlying distribution, $\Pr(X = 1) = p_{\text{true}}$, according to which nature would decide the forecast event X in the future, and each agent had some idiosyncratic noisy version π_i of this p_{true} , the variability of the agents’ beliefs being represented by the above BETA distribution with mean $\frac{\alpha}{\alpha + \beta} = p_{\text{true}}$ and pseudo-sample size (confidence) parameter $(\alpha + \beta)$ held constant at 1 (α and β denote standard parameters of a BETA distribution). Over the n trading episodes, we computed two price trajectories starting at $p_0 = 0.5$ each, one induced by each agent maximizing her myopic expected logarithmic utility⁹, and the other by the approximate price update equation that always rejects the error term in (2.11). At the end of each simulation, we evaluated the root-mean-squared deviation between these two price trajectories, and averaged these values over all 1000 simulations in the set to obtain the “mean RMSD between true and approximate price processes” which serves as our error measure for the approximation.

We report our results in Figures 2.2 and 2.3.

Figure 2.2 gives a quantification of the approximation error for various combinations of model parameter values. On eyeballing the sample trajectories in Figure 2.3, the path of approximate prices (dashed black) seems quite close to the path of true prices (solid green), more so for the lower value of \tilde{w}_{\max} , as expected; also note the high price volatility in panel (b) corresponding to the higher agent budgets, which is understandable since the agent is

⁹For agent i , we discretized the possible range of p_i , i.e. $[p_i^{\min}, p_i^{\max}]$ in steps of 10^{-4} , computed the vector of expected logarithmic utility values for these discrete p_i values, and chose the p_i -value corresponding to the maximum entry in this vector as the updated price.




Figure 2.2: The error measure increases with increasing difference between p_{true} and $p_0 = 0.5$ for any fixed \tilde{w}_{max} , and also with an increase in \tilde{w}_{max} for any given p_{true} -value; nevertheless, the error appears to be small even for higher values of \tilde{w}_{max} (less than 0.03 for $0.1 \leq p_{\text{true}} \leq 0.9$, $\tilde{w}_{\text{max}} \leq 0.75$). Error bars are not shown since standard errors are consistently two orders of magnitude smaller than corresponding sample means.

now closer to being risk-neutral. The main takeaway message from our experiments is that the approximation seems reasonable for a wide range of values of p_{true} and \tilde{w}_{max} .

We also studied the dependence of the error measure on the parameter $(\alpha + \beta)$ which is inversely related to the variance of the traders' beliefs. We fixed $p_{\text{true}} = 0.7$ and varied $(\alpha + \beta)$ over $\{0.1, 0.25, 0.5, 1, 1.5, 2.5, 5, 7.5, 10\}$. The results are reported in Figure 2.4. For both $\tilde{w}_{\text{max}} = 0.2$ and 0.5 , we see that this error measure peaks at 1 and then drops off slowly as $(\alpha + \beta)$ increases further.

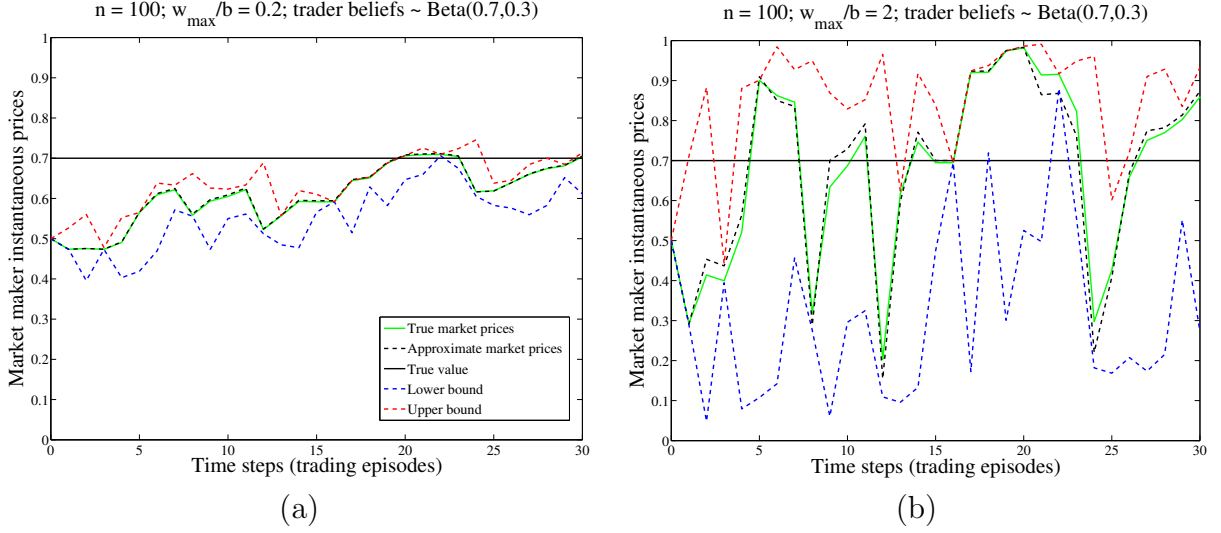


Figure 2.3: Price trajectories for two sample simulations with $\tilde{w}_{\max} = 0.2$ and $\tilde{w}_{\max} = 2$ are displayed in panels (a) and (b) respectively, $p_{\text{true}} = 0.7$ for both. “Lower bound” (dashed blue curve) and “Upper bound” (dashed red curve) for each trading episode i correspond to price bounds p_i^{\min} and p_i^{\max} respectively.

2.5 Discussion and future work

We have established the correspondence of a well-known securities market microstructure to a class of traditional belief aggregation methods and, by extension, Bayesian inference procedures in two important cases. An obvious next step is the identification of general conditions under which a MSR and agent utility combination is equivalent to a given pooling operation.

Another research direction is extending our results to a sequence of agents who trade repeatedly until “convergence”, taking into account issues such as the order in which agents trade when they return, the effects of the updated wealth after the first trade for agents with budgets, etc.

However, there is an implicit assumption common to all the results cited and presented in this chapter – the outcome of the forecast event, which serves as the ground truth for verifying all agents’ reports and hence deciding their payoffs, is extraneously determined, and is beyond all agents’ control. Prediction markets are often used in situations where this




Figure 2.4: Variation of the approximation error measure in our simulations with respect to the pseudo-sample size (confidence) parameter of the distribution of agent beliefs.

assumption is violated to a greater or lesser degree – and it is one of these situations that we will analyze in depth in Chapter 3.

Chapter 3

Outcome manipulation in incentivized collective forecasting mechanisms

3.1 Introduction

In Chapter 2, we saw how prediction markets use the idea of offering trade in financial instruments for aggregating and disseminating private information dispersed among a potentially diverse crowd. However, attention is seldom paid in the literature on prediction markets to the possibility that market participants might have some degree of control on the outcome of the forecast event, and hence the presence of a prediction market may make agents affecting the outcome act differently than they otherwise would. In fact, sometimes it is this very power to affect outcomes that gives agents the informational edge that such markets get their value from.

Consider three canonical real-world examples where prediction markets (or betting markets) have demonstrated their forecasting ability to great effect: elections / politics Berg et al. (2008), sporting events Wolfers and Zitzewitz (2004), and software product releases Cowgill and Zitzewitz (2015). In each of these cases, it is easy to see how the presence of a prediction market on the event may distort incentives. A congressional staffer or member of congress may know more about the probable outcome of a key vote than the general public, but she is also in a position to influence said outcome. A referee or player has substantial ability to influence the results of a game. A software engineer has the potential to delay (or speed up) the release of a product.

When trading agents can influence the event on which the market is written to such an extent that the outcome cannot be considered exogenous, it is natural to ask two questions:

- (1) Are market prices still *informative* of the forecast event, i.e. how much do they still tell us about the realized outcome?
- (2) Are the actions of the outcome-deciders still *truthful*, i.e. do they take the same actions that they would in the absence of the prediction market?

While it is acknowledged that prediction markets have value as aids in making business and policy decisions, they have gone through cycles of hype and bust for reasons that include regulatory concerns about manipulation, the emblematic anecdote in this vein being the failure of DARPA’s proposed policy analysis markets (Hanson, 2007b). While their actual proposed range and purpose was more complex, they became caricatured in the media as “terrorism futures” (Clifton, 2003), and the project was canceled almost as soon as information about it became publicly widespread. Stiglitz (July 31, 2003) pointed out some issues with the idea in an Op-Ed piece:

Did [Poindexter] believe there is widespread information about terrorist activity not currently being either captured or appropriately analyzed by the “experts” in the FBI and the CIA? Did he believe that the 1,000 people “selected” for the new futures program would have this information? If so, shouldn’t these people be investigated rather than rewarded?

But there are more fundamental problems with the idea. If trading is anonymous, then it could be subject to manipulation, particularly if the market has few participants, providing a false sense of security or an equally dangerous false sense of alarm. If trading is not anonymous, then anyone with information about terrorism would be, understandably, reluctant to trade on it. In that case, the market would not serve its purpose.

There are obviously prediction markets that will not work, but stock and futures markets have been used for a long time as forecasting tools, and prediction markets are similar in essence. The key is to understand when these markets may be prone to manipulation and how much to trust them. To this end, we propose and analyze a new model for studying

manipulative behavior that captures two aspects of real-world prediction markets: (1) agents directly affect the forecast event, and (2) some of the outcome-deciding agents may not participate in the prediction market (e.g. employees who have an impact on the outcome of a product launch typically would not all take part in the company’s in-house prediction market for its release date). In markets where an individual has a small effect on the outcome (like large elections), agents’ incentives for manipulation are likely to be weak. With this in mind, we mainly focus on a two-stage game-theoretic model of a market with two “players” or agents who affect the outcome and can also trade on it (Sections 3.3 and 3.4), and then discuss how our insights extend to models with more players in Section 3.7.

But, before getting into the technical details, we provide an informal overview of our model in Section 3.1.1, followed by a summary of the major contributions of this chapter in Section 3.1.2, and review of relevant literature in Section 3.2.

3.1.1 A model for manipulation

In our two-player model, the agents are called Alice and Bob; before the game commences, each of them receives a private signal about some underlying entity. In the first stage of the game, both players have the opportunity to participate (sequentially), once each, in a prediction market mediated by a (variant of a) market scoring rule (MSR) introduced in Chapter 2. Alice moves first; Bob may or may not participate in trading depending on his *type*, and, if he does, he goes second. In the second stage, the two players simultaneously (and independently) take actions which we term “votes”¹⁰ for convenience, although in general they model each participant’s role in determining the outcome. For example, for a product release date prediction market, a (binary) private signal could stand for whether an agent knows / believes she is capable of contributing her share in making sure that the launch is on time; her (binary) “vote” in this case would indicate whether she actually puts in her share of the requisite effort. The payoffs from the first-stage prediction market are determined by a simple function of the stage-two votes. If Bob has not traded, his vote is consistent with his private signal, otherwise he is strategic; Alice is always strategic.

¹⁰The nomenclature is inspired by a vote-share prediction market, e.g. Chakraborty et al. (2013).

Our model directly captures the experiments of Chakraborty et al. (2013) where prediction markets with student participants were used to forecast the fraction of “up” (vs “down”) ratings given by students to course instructors. Moreover, *Augur*, a “decentralized, open-source platform for prediction markets” slated for live release in 2017 (Peterson and Krug, 2015), is a real-world mechanism with manipulation incentives similar to those in our model: A consensus, computed from votes cast by participants called “reporters”, serves as a proxy for the payoff-deciding ground truth of a market on which these reporters can also wager.

3.1.2 Contributions

The above model yields several interesting insights. Our main result is that the equilibria of the game can be cleanly categorized into two types, depending on Bob’s probability of participation in the trading stage. Below a threshold on Bob’s participation probability, say \tilde{p} (a function of the MSR used and the signal structure), we call the equilibrium a low participation probability equilibrium (LPPE), and above \tilde{p} , we call it a high participation probability equilibrium (HPPE). In an LPPE, Alice essentially predicts Bob’s vote, and then bases her trading on the optimal combination of her own and Bob’s votes, and the prediction market price is reflective of the expected outcome. In an HPPE, on the contrary, Alice effectively expects Bob to enter and collude with her, and she chooses a market position that allows Bob and her to split the maximum extractable profit from the market mechanism, i.e. the MSR’s worst-case loss, in a (not necessarily even) ratio dependent on the MSR. The following implications of the equilibria are noteworthy:

- (a) *Informativeness of prices about outcome:* The price after Alice trades is equal to her posterior expectation of the market outcome given her signal and her trading action (hence, an efficient disseminator of information about the outcome at that point) in an LPPE, but contains only partial information about the final outcome in an HPPE. If Bob trades, the final market price is an accurate forecast of the actual outcome.
- (b) *Consistency of actions with signals:* We provide a full characterization of signal structures under which Alice’s actions are consistent with (or, at least, indicative of) her private signal. If Bob does participate in the trading stage, his actions are fully determined by Alice’s trading choice, independent of his signal.

- (c) *Effectiveness of prediction markets in the face of participants who influence outcomes substantially:* One implication of point (b) above is that, as long as some fraction of the outcome-deciders refrain from market participation and are truthful in their outcome-affecting actions (proxied in our model by Bob with his non-zero non-participation probability), the introduction of a prediction market to elicit forecasts prior to the occurrence of the outcome-deciding process is less likely to produce damaging incentives in the sense that even the potential manipulator Alice is forced to act, under certain conditions, as if she were truthful.
- (d) *Choice of market scoring rule:* The equilibrium strategies of Alice and Bob, and the resulting market properties, have a strong dependence on the the functional form of the MSR used, as shown by Tables , Sections . Thus, our analysis also provides guidelines on the the scoring rule a designer can use to implement the prediction market if she wants to achieve certain marker properties.
- (e) *Strategizing by a one-shot trader on point of entry:* The above properties are important from the perspective of a market observer and / or the market designer, but our results have implications for the agents too, particularly if they are free to choose when to interact with the market maker rather than in a pre-defined sequence. In an extreme case of our two-player model – when both Alice and Bob are deterministically strategic, and this is common knowledge – our results show that the first mover Alice may sustain a higher or lower profit (in equilibrium) than the second mover depending on the MSR used, hence it depends on the MSR whether it is better to be enter the market earlier or later.

Incentivizing truthful trading and voting: In Section 3.8, we return to the market design question, and propose a remedy for manipulation in an extension of the above two-stage game with $n \geq 2$ participants, by introducing two modifications: first, we put a budget constraint on every trader which limits the amount by which she can alter the observed market price; second, we combine the prediction market with a payment scheme based on the *peer prediction* method (Miller et al., 2005; Jurca and Faltings, 2009) for the voting mechanism, by tuning parameters, in order to obtain desirable incentive compatibility properties. We also analyze the level of subsidy needed in the combined prediction / voting mechanism.

3.2 Related work

This chapter relates most directly to three major strands of literature: (1) Incentives and manipulation in prediction markets, (2) Insider trading in financial markets, (3) Information elicitation when the “ground truth” is not revealed.

(1) The literature on incentives in prediction markets is growing ever since Hanson (2003b) introduced the concept of market scoring rules. Chen et al. (2009) and Gao et al. (2013) have studied the effect of non-myopic strategies on an LMSR market modeled as an extensive-form Bayesian game, under various private information structures of the participants who trade repeatedly. We consider agents that are still interact at most once each with the market mechanism but manipulation incentives arise from another source: their power to affect the forecast event.

Incentives for manipulation in prediction markets may arise in a number of ways. There are several contributions – both empirical / experimental (Hanson et al., 2006; Rhode and Strumpf, 2006) and theoretical (Hanson and Oprea, 2009; Boutilier, 2012; Dimitrov and Sami, 2010; Chen et al., 2011a; Huang and Shoham, 2014) – on *price manipulation* in prediction markets: tampering with the market price by belief misrepresentation, perhaps even at a monetary cost, so as to indirectly influence some decision that will be made by the principal (market organizer) or non-participating market observers based on that price, e.g. a politically motivated manipulator might make a large investment in an election prediction market to make one of the candidates *appear* stronger (Rothschild and Sethi, 2016). A related body of work pertains to *decision markets* – a collection of contingent markets set up to predict the outcomes of different decisions such that only markets contingent on decisions that are taken pay off, the rest being voided – a concept proposed by Hanson (1999) and built on by several groups (Othman and Sandholm, 2010; Chen et al., 2011b). In this chapter, we consider MSR-based prediction market settings, and not the above decision markets. Moreover, the type of manipulation we are interested in is not price manipulation but *outcome manipulation* where an agent can take an action that partially influences the outcome to be predicted, and base her trading decision on her action choice and her relevant belief about other agents.

An early formal analysis of prediction market outcome manipulation is that of Ottaviani and Sørensen (2007) on a two-outcome market model that follows the rules of the Iowa Electronic Markets (Berg and Rietz, 2006). Their results apply to a rational expectations equilibrium setting with agents having constant absolute risk aversion utility functions, and they do not take the market microstructure into account. We explicitly model the microstructure, focusing on market scoring rules – a well-studied prediction market mechanism that is widely used in practice (Jian and Sami, 2012) – under a different agent model. More recently, Shi et al. (2009) have introduced “principal-aligned scoring rules” (a refinement of proper scoring rules) to provide disincentives for extraneous manipulations by traders that can modify the probability distribution over outcomes so as to reduce the expected utility of the principal. However, this technique works only when the principal’s utility vector over outcomes is publicly known, at least in a probabilistic sense. Moreover, it provides clear incentives for manipulations that are beneficial to the principal in expectation (which might not be desirable in some cases).

(2) The second major body of literature comes from theoretical finance and market microstructure. What happens if there is an “insider” – a market participant who knows the *liquidation value* (i.e. final gross monetary worth after revelation of the outcome) of a security, and can trade on this information? Kyle (1985) proposed a seminal model of insider trading, and characterized the rate at which a monopolist insider’s information gets disseminated into market prices in the presence of noise trading and a risk-neutral market-maker. Glosten and Milgrom (1985) presented another view of how asymmetric information affects price formation, and their model has been adapted for market making in prediction markets (Das, 2008a; Brahma et al., 2012). There has also been work on extending Kyle’s model to competing traders with inside information (Holden and Subrahmanyam, 1992; Foster and Viswanathan, 1996). Ostrovsky (2012) examined information aggregation with differentially informed traders in a general information framework under both Kyle’s pricing model and market scoring rules. Again, in all of these models, the liquidation value of the market security is assumed to be exogenously determined, so there is no interaction between trading behavior and the behavior that produces the market outcome, unlike in our model.

(3) In the last part of this chapter, we use peer prediction to align the incentives of participants in the two-stage game we propose here. This idea is motivated by the literature on providing truth-telling incentives in traditional means of opinion or information gathering

such as surveys, polls, and reputation systems that can be subsumed under “information elicitation sans verification” (Waggoner and Chen, 2013). Prelec et al. have developed and experimented with the Bayesian Truth Serum or BTS (Prelec, 2004; Prelec and Seung, 2007; Weaver and Prelec, 2012). In this method, each respondent is required to answer a multiple-choice question and predict the distribution of responses from the entire population; she is then assigned a score which rewards an answer whose observed frequency is higher than the average predicted frequency and penalizes a bad prediction of population response. Under mild assumptions, this technique induces a Bayes-Nash equilibrium strategy of reporting one’s private information provided that there is a large number of participants. Several extensions of BTS have been proposed (e.g. Radanovic and Faltings (2014) and references therein).

Miller et al. (2005) proposed the peer prediction method (PP) that adapts the concept of proper scoring rules to the problem of rating. This method scores a rater with respect to the posterior belief that her report induces on another rater’s report, assuming a common prior and likelihood structure known to the mechanism. This idea has spawned an interesting line of research (e.g. Witkowski and Parkes (2012) and references therein). Witkowski and Parkes (2012) introduce two variants of PP that elicit two temporally separated reports from each agent, at least one of which is a belief report about another agent’s signal. These latter techniques can work for arbitrary subjective priors but are restricted to binary signal domains.

The situations we are interested in analyzing in this chapter are not the ones for which BTS or PP were derived. They are intended to solve the problem of getting people to vote, or give their opinion. The kinds of settings we are interested in are ones where people would already vote / give their opinion / do their work “honestly”, but the introduction of a prediction market may change their incentive to do so. Therefore, while the ideas from this literature will turn out to be useful, we cannot simply apply BTS or PP in the first place and ignore the existence of the prediction market (keeping in mind that the concept of a prediction market serves as useful shorthand in our model, but this could equally be a liquid real money futures market or the like).

It is worth mentioning here that our idea of achieving incentive compatibility by a two-stage game in Section 3.8, each stage having monetary transfers between the principal and the

participants, is similar in spirit to work on auction design (Cr  mer and McLean, 1988) that predates peer prediction. Cr  mer and McLean design a two-stage mechanism for a seller of a single item who is unaware of the potential buyers’ valuations. In the first stage, each bidder makes a payment, conditional on her peers’ announced valuations, to the seller for a “lottery” to win the item. In the second stage, bidders declare their types and the winner pays her announced valuation. The second stage by itself would not be truthful but the authors show that, under certain information structures, the combined payoff from the two stages induces dominant strategy or Bayes-Nash incentive compatibility, enabling the seller to “extract the full surplus.”

3.3 Model and definitions

Let $\tau \in T$ denote the unobservable true value of the random variable on which both the prediction market and the outcome-deciding (voting) system are predicated. At $t = 0$, the two agents, Alice and Bob (A and B in subscripts), receive private signals $s_A, s_B \in \Omega = \{0, 1\}$ respectively. The *signal structure*, comprising the prior distribution $\Pr(\tau)$ on the true value and the conditional joint distribution $\Pr(s_A, s_B|\tau)$ of the private signals given the true value, is common knowledge.

Let $q_0(\cdot)$ denote Alice’s posterior probability that Bob received the signal $s_B = 0$, given her own signal and common knowledge, i.e.

$$q_0(s) \triangleq \Pr(s_B = 0 | s_A = s) = \frac{\sum_{\tau \in T} \Pr(s_A = s, s_B = 0 | \tau) \Pr(\tau)}{\sum_{\tau \in T} \Pr(s_A = s | \tau) \Pr(\tau)} \quad \forall s \in \{0, 1\}. \quad (3.1)$$

In this chapter, we ignore the uninteresting special cases $q_0 \in \{0, 1\}$ which correspond to Alice having no uncertainty about her peer Bob’s private signal. We need no further assumptions on the signal structure for our main result (Theorem 4) since it depends only on the magnitude of q_0 regardless of how it is evaluated. We shall discuss a specific signal structure in Section 3.6. However, it is worthwhile to define here the property of *stochastic relevance* (Miller et al., 2005) which is a necessary assumption for one of our important corollaries (Corollary 1).

Definition 3. For binary random variables $s_i, s_j \in \{0, 1\}$, s_j is said to be *stochastically relevant* for s_i if and only if the posterior distribution of s_i given s_j is different for different realizations of s_j , i.e. if and only if $\Pr(s_i = 0 | s_j = 0) \neq \Pr(s_i = 0 | s_j = 1)$.

An important implication of the above definition for our model is that if s_A is stochastically relevant for s_B , there is a one-to-one correspondence between the value s of Alice’s signal and $q_0(s) = \Pr(s_B = 0 | s_A = s)$; i.e. if we could somehow learn the value of q_0 , then we would be able to extract Alice’s signal unambiguously from it.

We now describe the rules of the two-stage game comprising the market and voting mechanisms. We will call this the *trading-voting game*, and assume that its rules are known to all participants and observers.

Stage 1 (market stage): The market price at any time-step t is public, the starting price at $t = 0$ being p_0 which is the market designer’s baseline estimate of the *liquidation value*, i.e. the final gross payoff per unit of the prediction market security which, in our setting, is identical to the market outcome (see **Stage 2** below).

Because the market outcome for this problem has a different structure from that in Chapter 2, the prediction market is implemented using a slight variation of the market scoring rule (MSR) algorithm introduced in Chapter 2: As usual, we denote the underlying *strictly proper* scoring rule as $s(r, \omega)$, where ω is the true (revealed) market outcome, and r is an agent’s forecast / report on it but we now use a formulation of the rule that is used for elicitation of *personal expectations* of a continuous random variable $\omega \in [0, 1]$ and not *personal probabilities*. The principle is a natural extension to that in our Chapter 2 formulation: strict propriety implies that if an agent is promised an ex post compensation of $s(r, \omega)$, then the only way she can maximize her subjective expectation of her ex post compensation is by reporting her *expectation* of the random variable ω as her forecast r . Similar formulations were used by Ostrovsky (2012) for theoretical analyses and by Chakraborty et al. (2013) for experiments. Moreover, for a clean analysis, we shall focus on strictly proper rules satisfying some regularity and smoothness conditions (Gneiting and Raftery, 2007; Abernethy and

Frangillo, 2012):

$$s(r, \omega) = \begin{cases} f(\omega), & r = \omega \\ f(r) + f'(r)(\omega - r), & \text{otherwise} \end{cases}, \quad \omega, r \in [0, 1] \quad (3.2)$$

where $f(\cdot)$ is a continuous, finite, strictly convex function on $[0, 1]$; its first derivative $f'(\cdot)$ is continuous, monotonically increasing, and finite on $[0, 1]$ except possibly that $f'(0) = -\infty$ or $f'(1) = \infty$; its second derivative $f''(\cdot)$ is positive¹¹ on $[0, 1]$ and finite in $(0, 1)$. Additionally, we need the function to have the following symmetry:

$$f\left(\frac{1+y}{2}\right) - f\left(\frac{1-y}{2}\right) = yf'\left(\frac{1}{2}\right) \quad \forall y \in [0, 1]. \quad (3.3)$$

Henceforth, we shall refer to (market) scoring rules possessing all the above properties as *symmetric well-behaved* (market) scoring rules.¹² This covers a large family of MSRs that includes three of the most widely used and studied – (LMSR), quadratic (QMSR), and spherical (SMSR) – respectively defined as:

$$\begin{aligned} \text{LMSR:} \quad & s(r, \omega) = \omega \ln r + (1 - \omega) \ln(1 - r), \\ \text{QMSR:} \quad & s(r, \omega) = \omega^2 - (\omega - r)^2, \\ \text{SMSR:} \quad & s(r, \omega) = (r\omega + (1 - r)(1 - \omega)) / \sqrt{r^2 + (1 - r)^2}. \end{aligned}$$

At $t = 1$, Alice interacts with the market maker, and takes such a position as to change the price to p_A . At $t = 2$, Bob has an opportunity to trade but may not show up with a commonly known probability $\pi \in [0, 1]$ called Bob’s *non-participation probability*; if he does trade, he changes the price to p_B . Regardless of whether Bob trades, the market terminates after $t = 2$.

¹¹The strict positivity of $f''(\cdot)$ is sufficient but not necessary for the strict convexity of $f(\cdot)$ to hold on $[0, 1]$ (all we need is non-negativity). Nevertheless, for technical convenience, we shall restrict our presentation to convex functions with strictly positive second derivatives, and refer to such functions only when we use the expression “strict convexity”, in a slight abuse of terminology.

¹²Note that the definition of “well-behaved” used here is slightly different from that in Chapter 2 but this overloading of terms should not create a confusion since this term is not relevant to subsequent chapters, and follows the new definition throughout this chapter.

Stage 2 (voting stage): In this stage, Alice and Bob simultaneously declare their “votes” $v_A, v_B \in \Omega$. Taking part in **Stage 2** is mandatory for both agents.

We define *truthful voting* as declaring one’s private signal as one’s vote, *i.e.* $v_k = s_k$, $k \in \{A, B\}$.

We assume that, if Bob did not trade in **Stage 1**, he votes truthfully, and we call such a Bob HONEST. Any agent participating in the prediction market is Bayesian, strategic, and risk-neutral. Hence, if Bob trades, we refer to him as STRATEGIC Bob.

The liquidation value of the security, *i.e.* the market outcome, is given by the average¹³ vote $v = (v_A + v_B) / 2 \in \{0, \frac{1}{2}, 1\}$. The ex post net payoffs of Alice and Bob, which we will also sometimes refer to as their *profits*, follow from the definition of a market scoring rule as a sequentially shared proper scoring rule Hanson (2007a), and are respectively given by

$$\begin{aligned}\mathcal{R}_A(p_A, p_0, v_A, v_B) &= s(p_A, \frac{v_A + v_B}{2}) - s(p_0, \frac{v_A + v_B}{2}), \\ \mathcal{R}_B(p_B, p_A, v_A, v_B) &= s(p_B, \frac{v_A + v_B}{2}) - s(p_A, \frac{v_A + v_B}{2}).\end{aligned}\tag{3.4}$$

This completes the technical description of our model, but there are a few points worth noting here:

- Agents derive utility solely from their profits in the prediction market, and have no vested interest in any realized outcome although they have power to influence it, exactly as in the works of Ottaviani and Sørensen (2007) and Shi et al. (2009); this can be viewed as a model of agents who are inherently non-strategic in the absence of any financial incentives, as is often experimentally observed (e.g. Gao et al. (2014)).
- Bob does not strategically decide whether to take part in the prediction market; it is determined extraneously – the proclivity to trade can be viewed as one of the two independent components of Bob’s *type*, the other being his private signal s_B : The complete distribution of Bob’s type is given by the signal structure detailed above and the two-point distribution $\{\Pr(\text{HONEST Bob}) = \pi, \Pr(\text{STRATEGIC Bob}) = 1 - \pi\}$. Here,

¹³In general, we can have $v = \alpha v_A + (1 - \alpha) v_B$, $\alpha \in (0, 1)$, where α models Alice’s degree of control over the final outcome. In this chapter, we focus on the special case $\alpha = \frac{1}{2}$ as a starting point where both agents are equally powerful.

HONEST Bob models agents who remain unaffected by the introduction of the prediction market owing to social norms or some exogenous payoff.¹⁴ Among other things, we aim to study what effect, if any, the uncertainty around the market participation of some outcome-deciders has on the actions of a strategic agent when there is incentive for manipulation.

In the rest of the chapter, we will often use the terms **Stage 1** action and trading action to denote p_A , p_B , and **Stage 2** action to denote v_A , v_B . Unlike in a traditional prediction market, the true value τ of the underlying random variable is never revealed in this trading-voting game. However, if every agent were to vote truthfully (and this were common knowledge), then by the properties of MSRs, an agent’s expected net payoff would be maximized by “reporting”, in **Stage 1**, her posterior expectation of the liquidation value v . With this in mind, we will sometimes refer to p_A and p_B as the “reports” or “price-reports” of Alice and Bob respectively.

3.4 Equilibrium analysis of the two-player game

The solution concept we will use for the two-stage trading-voting game described in Section 3.3 is the *perfect Bayesian equilibrium* (PBE) which is a refinement of Nash equilibria for Bayesian games (Fudenberg and Tirole, 1991). A PBE is a specification of an *assessment*, comprising a strategy and a belief structure, for each player such that each player’s action according to her strategy and belief at any stage of the game is a “best reponse” (sequential rationality), and each player’s belief about all currently unknown aspects of the game at any stage is obtained from all players’ actions until that stage using Bayesian inference whenever possible (consistency). For this game, a strategy profile of the players Alice and Bob is a specification of the vector $((p_A, v_A), (p_B, v_B))$, and we shall denote a PBE strategy profile by $((p_A^{PBE}, v_A^{PBE}), (p_B^{PBE}, v_B^{PBE}))$.

¹⁴For an MSR-mediated prediction market, Alice’s payoff function (3.4) depends only on Bob’s **Stage 2** action v_B and not on his **Stage 1** action p_B . Hence, our results for Alice in Section 3.4 also apply to a slightly modified model where HONEST Bob can still trade strategically in **Stage 1** but is constrained to vote truthfully in **Stage 2**.

In the game under consideration, Alice makes her first move by updating the market price from p_0 to p_A which becomes publicly known; then, Bob makes his first move by either participating in the market and revising p_A to the publicly observable p_B or not participating, the latter being equivalent to setting $p_B = p_A$; after this sequential first stage, both Alice and Bob make their second moves (picking v_A and v_B respectively) simultaneously. Thus, our analysis of the game consists of the following three steps:

1. Lemma 5 delineates the inference about Alice's second move v_A that can be drawn by the rest of the world including Bob from her **Stage 1** action p_A , for any $p_0 \in (0, 1)$.
2. Lemma 6 presents, for any $p_0 \in (0, 1)$, STRATEGIC Bob's best response (p_B, v_B) to his observation of p_A and belief about v_A ; it also tells us what inference about v_B the rest of the world including Alice can draw from p_B in this case. We already know that, for HONEST Bob, $p_B = p_A$ and $v_B = s_B$.
3. Finally, Theorem 4 completes the equilibrium specification for the particular case of $p_0 = \frac{1}{2}$ by providing Alice's best choices for p_A and v_A given her signal, her knowledge of the above two lemmas, and her observation of p_B .

Lemma 5. *For the trading-voting game described in Section 3.3, if the prediction market has a starting price $p_0 \in (0, 1)$, and (p_A, v_A) denotes Alice's combined action in the two-stage game, i.e. her report-vote pair, then*

- *for any $p_A < p_0$, Alice's action $(p_A, 0)$ strictly dominates $(p_A, 1)$;*
- *for any $p_A > p_0$, her action $(p_A, 1)$ strictly dominates $(p_A, 0)$; and*
- *she is indifferent between the actions $(p_0, 0)$ and $(p_0, 1)$.*

This result holds regardless of Bob's report-vote pair (p_B, v_B) .

Informally, if Alice pulls the market price down (resp. up) from its initial value, she is “forecasting” that the final outcome will be lower (resp. higher) than the market's initial estimate and, since her payoff is higher for a prediction closer to the realized outcome (the average vote), it is in her best interest to do everything in her power to ensure a low (resp. high) average vote.

Proof. Using equations (3.4) and (3.2), we can show by simple algebra that, for any p_A , the difference between Alice's profits for voting $v_A = 1$ and $v_A = 0$ is

$$\mathcal{R}_A(p_A, p_0, 1, v_B) - \mathcal{R}_A(p_A, p_0, 0, v_B) = \frac{1}{2} (f'(p_A) - f'(p_0)), \quad \forall v_B \in \{0, 1\}.$$

From the increasing monotonicity of $f'(\cdot)$ for a symmetric well-behaved MSR, we have

$$p_A \gtrless p_0 \iff f'(p_A) \gtrless f'(p_0) \iff \mathcal{R}_A(p_A, p_0, 1, v_B) \gtrless \mathcal{R}_A(p_A, p_0, 0, v_B)$$

for $p_A \in [0, 1]$. Thus Alice's optimal vote, regardless of Bob's actions, is $v_A = 0$ if $p_A < p_0$ and $v_A = 1$ if $p_A > p_0$; she is indifferent only if $p_A = p_0$. \square

The above theorem implies that immediately after Alice has traded, one can infer that $v_A = 0$ deterministically if $p_A < p_0$, $v_A = 1$ deterministically if $p_A > p_0$. However, if $p_A = p_0$, which is equivalent to Alice not trading with the market maker, the rules of the game do not allow Bob to predict v_A deterministically: He knows that Alice does not stand to make any profit from the market regardless of the outcome, and hence, must be indifferent between voting 0 and 1. Hence, it is reasonable to assume a belief structure in which Bob's posterior belief assigns equal probabilities to $v_A = 0$ and $v_A = 1$ whenever $p_A = p_0$.¹⁵

Assumption 1. *If P^B denotes Bob's posterior belief (about Alice's vote), given his signal s_B and Alice's price report p_A , then for any signal $s \in [0, 1]$,*

$$P^B(v_A = 0 | s_B = s, p_A = p_0) = P^B(v_A = 1 | s_B = s, p_A = p_0) = \frac{1}{2}.$$

Obviously, $P^B(v_A = 0 | s_B = s, p_A < p_0) = P^B(v_A = 1 | s_B = s, p_A > p_0) = 1$, by Lemma 5.

Before we proceed further with our analysis, we will define two positive proper fractional quantities that are purely functions of the structure of the MSR used for designing the prediction market and play key roles in determining the equilibrium of the game. We can invoke the mean value theorem to show that, for a symmetric well-behaved MSR, the equation

¹⁵We shall later see in Sections 3.4.1 through 3.4.3 that, for the case $p_0 = \frac{1}{2}$ which are main focus in this paper, Alice never leaves the market price unchanged in equilibrium for any positive value of Bob's non-participation probability π . But, we still include a discussion of the case $p_A = p_0$ for the sake of completing the PBE description.

$f'(r) = (f(\frac{1}{2}) - f(0))/(\frac{1}{2} - 0)$ has a root in $(0, \frac{1}{2})$, which is unique owing to the monotonicity of $f'(\cdot)$; let this root be denoted by p^L , i.e.

$$p^L \triangleq (f')^{-1}\left(2\left(f\left(\frac{1}{2}\right) - f(0)\right)\right) \in (0, \frac{1}{2}). \quad (3.5)$$

Similarly, let p^H denote the unique root of the equation $f'(r) = (f(1) - f(\frac{1}{2}))/(\frac{1}{2})$ in $(\frac{1}{2}, 1)$, i.e.

$$p^H \triangleq (f')^{-1}\left(2\left(f(1) - f\left(\frac{1}{2}\right)\right)\right) \in (\frac{1}{2}, 1). \quad (3.6)$$

We call p^L and p^H the “lower threshold” and “upper threshold” of the MSR respectively since they mark points of discontinuity in the players’ equilibrium behavior, as we will see shortly. From the definitions, it is obvious that

$$0 < p^L < \frac{1}{2} < p^H < 1. \quad (3.7)$$

Moreover, using definitions (3.5) and (3.6) and the symmetry condition (3.3), we can obtain the following results (the detailed proof is in Appendix B Section B.1).

Proposition 2. *For a symmetric well-behaved market scoring rule, the lower and upper thresholds p^L, p^H defined in (3.5) and (3.6) satisfy the equalities*

$$f'(p^L) + f'(p^H) = 2f'(\frac{1}{2}); \quad (3.8)$$

$$p^L + p^H = 1; \quad (3.9)$$

$$f(p^H) - f(p^L) = (2p^H - 1)f'(\frac{1}{2}) = (1 - 2p^L)f'(\frac{1}{2}). \quad (3.10)$$

Table 3.1 provides the functional forms of $f(\cdot)$ and $f'(\cdot)$, and the values of the thresholds p^L, p^H for each of the three specific MSRs mentioned in Section 3.3.

Notice that, as soon as STRATEGIC Bob arrives to trade, he acquires all the information relevant to his decision making procedure that the rules of the game allow him to have (he can observe both p_0 and p_A , and draw inference about v_A in accordance with Lemma 5); Bob also knows that he and Alice are the only outcome-deciders and, even if an agent traded

	LMSR	QMSR	SMSR
$f(r)$	$\begin{cases} r \ln r + (1-r) \ln(1-r) & 0 < r < 1, \\ 0 & r \in \{0, 1\} \end{cases}$	r^2	$\sqrt{r^2 + (1-r)^2}$
$f'(r)$	$\ln\left(\frac{r}{1-r}\right)$	$2r$	$\frac{2r-1}{\sqrt{r^2 + (1-r)^2}}$
p^L	0.2	0.25	$\frac{1}{2} \left(1 - \sqrt{\frac{\sqrt{2}-1}{2}} \right) \approx 0.2725$
p^H	0.8	0.75	$\frac{1}{2} \left(1 + \sqrt{\frac{\sqrt{2}-1}{2}} \right) \approx 0.7275$

Table 3.1: Structural properties of the three representative market scoring rules considered in this paper.

after him, that agent would have no effect on his payoff. Thus, STRATEGIC Bob makes his trading and voting decisions (p_B, v_B) simultaneously.

Lemma 6. *For the trading-voting game described in Section 3.3, where the market scoring rule has lower and upper thresholds p^L, p^H , and has a starting price $p_0 \in (0, 1)$,*

- *if $p_A < p_0$, then STRATEGIC Bob's best-response vote is $v_B = 1$ (resp. $v_B = 0$) if $p_A < p^L$ (resp. $p_A > p^L$) but he is indifferent between the two possible voting choices if $p_A = p^L$, and his accompanying price-report is $p_B = \frac{1+v_B}{2}$;*
- *if $p_A > p_0$, then STRATEGIC Bob's best-response vote is $v_B = 1$ (resp. $v_B = 0$) if $p_A < p^H$ (resp. $p_A > p^H$) but he is indifferent between the two possible voting choices if $p_A = p^H$, and his accompanying price-report is $p_B = \frac{1+v_B}{2}$;*
- *if $p_A = p_0$, then STRATEGIC Bob's best-response vote is $v_B = 0$ (resp. $v_B = 1$) if $p_0 > \frac{1}{2}$ (resp. $p_0 < \frac{1}{2}$) but he is indifferent if $p_0 = \frac{1}{2}$, and his accompanying price-report is $p_B = \frac{\frac{1}{2}+v_B}{2}$.*

This result is independent of Bob's private signal s_B .

Before proving the above lemma, we present a detailed tabular representation of the results of Lemmas 5 and 6 in Table 3.2. Evidently, the quantities p^L , p^H , and p_0 split the possible range of market prices $[0, 1]$ into three or four sub-intervals depending on their relative magnitudes such that STRATEGIC Bob's best response is to “disagree with” Alice's voting

choice (revealed through p_A) in **Stage 2** if Alice's price-report p_A lies in either of the “outer” sub-intervals $[0, \min(p_0, p^L))$ or $(\max(p_0, p^H), 1]$, and to “agree with” Alice if p_A lies in the remaining one or two “inner” sub-interval(s) between $\min(p_0, p^L)$ and $\max(p_0, p^H)$; in any of these situations (except $p_A = p_0$), STRATEGIC Bob knows exactly what the market outcome v is going to be, so he can make a perfect forecast $p_B = v$. The special cases $p_A \in \{p^L, p^H, p_0\}$ represent points of transition in best-response characteristics.

Proof. If $p_A \neq p_0$, Bob knows v_A unambiguously. So, by the properties of a symmetric well-behaved MSR, Bob's payoff function $\mathcal{R}_B(p_B, p_A, v_A, v_B)$ from (3.4) is maximized uniquely at $p_B = \frac{v_A + v_B}{2}$ for any v_B , i.e. his combined two-stage action $(\frac{v_A}{2}, 0)$ dominates $(p, 0)$ for any $p \neq \frac{v_A}{2}$, and $(\frac{v_A + 1}{2}, 1)$ dominates $(p, 1)$ for any $p \neq \frac{v_A + 1}{2}$. Hence, for a known v_A , it suffices to compare $\mathcal{R}_B(\frac{v_A}{2}, p_A, v_A, 0)$ and $\mathcal{R}_B(\frac{v_A + 1}{2}, p_A, v_A, 1)$ to determine STRATEGIC Bob's optimal decision.

Case I: $p_A < p_0 \Rightarrow v_A = 0$ from Lemma 5. In this scenario, STRATEGIC Bob selects action $(p_B = 0, v_B = 0)$ if $\mathcal{R}_B(0, p_A, 0, 0)$ is larger and $(p_B = \frac{1}{2}, v_B = 1)$ if $\mathcal{R}_B(\frac{1}{2}, p_A, 0, 1)$ is larger. The difference simplifies to

$$\begin{aligned} \mathcal{R}_B(\tfrac{1}{2}, p_A, 0, 1) - \mathcal{R}_B(0, p_A, 0, 0) &= [2(f(\tfrac{1}{2}) - f(0)) - f'(p_A)] / 2 \\ &= [f'(p^L) - f'(p_A)] / 2, \quad \text{by (3.5).} \end{aligned}$$

Then, since $f'(\cdot)$ is strictly increasing,

$$p_A \gtrless p^L \iff f'(p_A) \gtrless f'(p^L) \iff \mathcal{R}_B(\tfrac{1}{2}, p_A, 0, 1) \gtrless \mathcal{R}_B(0, p_A, 0, 0).$$

Thus, if STRATEGIC Bob observes that p_A is smaller than p_0 , his best response is to report $p_B = \frac{1}{2}$ and vote $v_B = 1$ if p_A is also smaller than p^L (which is always true if $p_0 < p^L$) but to report $p_B = 0$ and vote $v_B = 0$ if p_A exceeds p^L (which is possible for a p_A smaller than p_0 only if $p^L < p_0$), and to be indifferent between these two actions if $p_A = p^L$.

p_0	p_A	Best v_A	Best (p_B, v_B) for STRATEGIC Bob
$0 < p_0 \leq p^L$ $p^L < p_0 < \frac{1}{2}$	$0 \leq p_A < p_0$	0	$(\frac{1}{2}, 1)$
	$p_A = p_0$	0 or 1	$(\frac{3}{4}, 1)$
	$p_0 < p_A < p^H$	1	$(1, 1)$
	$p_A = p^H$	1	$(\frac{1}{2}, 0)$ or $(1, 1)$
	$p^H < p_A \leq 1$	1	$(\frac{1}{2}, 0)$
	$0 \leq p_A < p^L$	0	$(\frac{1}{2}, 1)$
	$p_A = p^L$	0	$(0, 0)$ or $(\frac{1}{2}, 1)$
	$p^L < p_A < p_0$	0	$(0, 0)$
	$p_A = p_0$	0 or 1	$(\frac{3}{4}, 1)$
	$p_0 < p_A < p^H$	1	$(1, 1)$
	$p_A = p^H$	1	$(\frac{1}{2}, 0)$ or $(1, 1)$
	$p^H < p_A \leq 1$	1	$(\frac{1}{2}, 0)$
$p_0 = \frac{1}{2}$	$0 \leq p_A < p^L$	0	$(\frac{1}{2}, 1)$
	$p_A = p^L$	0	$(0, 0)$ or $(\frac{1}{2}, 1)$
	$p^L < p_A < p_0$	0	$(0, 0)$
	$p_A = p_0$	0 or 1	$(\frac{1}{4}, 0)$ or $(\frac{3}{4}, 1)$
	$p_0 < p_A < p^H$	1	$(1, 1)$
	$p_A = p^H$	1	$(\frac{1}{2}, 0)$ or $(1, 1)$
	$p^H < p_A \leq 1$	1	$(\frac{1}{2}, 0)$
$\frac{1}{2} < p_0 < p^H$ $p^H \leq p_0 < 1$	$0 \leq p_A < p^L$	0	$(\frac{1}{2}, 1)$
	$p_A = p^L$	0	$(0, 0)$ or $(\frac{1}{2}, 1)$
	$p^L < p_A < p_0$	0	$(0, 0)$
	$p_A = p_0$	0 or 1	$(\frac{1}{4}, 0)$
	$p_0 < p_A < p^H$	1	$(1, 1)$
	$p_A = p^H$	1	$(\frac{1}{2}, 0)$ or $(1, 1)$
	$p^H < p_A \leq 1$	1	$(\frac{1}{2}, 0)$
	$0 \leq p_A < p^L$	0	$(\frac{1}{2}, 1)$
	$p_A = p^L$	0	$(0, 0)$ or $(\frac{1}{2}, 1)$
	$p^L < p_A < p_0$	0	$(0, 0)$
	$p_A = p_0$	0 or 1	$(\frac{1}{4}, 0)$
	$p_0 < p_A \leq 1$	1	$(\frac{1}{2}, 0)$

Table 3.2: Alice's best vote and Bob's best report-vote pair if he trades, given starting price p_0 and Alice's report p_A . Recall that $p^L \in (0, \frac{1}{2})$ and $p^H \in (\frac{1}{2}, 1)$ for any symmetric well-behaved MSR, hence the table covers all possible combinations of the values of p^L , p^H , and p_0 . Theorem 4 applies to the $p_0 = \frac{1}{2}$ shown in the middle of the table.

Case II: $p_A > p_0 \Rightarrow v_A = 1$ from Lemma 5. Arguing as before, we now need to consider the difference

$$\begin{aligned}\mathcal{R}_B(1, p_A, 1, 1) - \mathcal{R}_B(\tfrac{1}{2}, p_A, 1, 0) &= [2(f(1) - f(\tfrac{1}{2})) - f'(p_A)] / 2 \\ &= [f'(p^H) - f'(p_A)] / 2, \quad \text{by (3.6)}.\end{aligned}$$

Again, from the increasing monotonicity of $f'(\cdot)$,

$$p_A \gtrless p^H \iff f'(p_A) \gtrless f'(p^H) \iff \mathcal{R}_B(1, p_A, 1, 1) \lesseqgtr \mathcal{R}_B(\tfrac{1}{2}, p_A, 1, 0).$$

Thus, if STRATEGIC Bob observes that p_A is larger than p_0 , his best response is to report $p_B = \frac{1}{2}$ and vote $v_B = 0$ if p_A is also larger than p^H (which is always true if $p_0 > p^H$) but to report $p_B = 1$ and vote $v_B = 1$ if p_A is smaller than p^H (which is possible for a p_A larger than p_0 only if $p^H > p_0$), and to be indifferent between these two actions if $p_A = p^H$.

Case III: $p_A = p_0$. Let us define

$$\widehat{\mathcal{R}}_B^{p_0}(p_B, v_B) \triangleq \mathbb{E}_{P^B}[\mathcal{R}_B(p_B, p_A, v_A, v_B) | p_A = p_0],$$

where P^B is Bob's posterior as defined in Assumption 1. Clearly, $\mathbb{E}_{P^B}[v_A | p_A = p_0] = \frac{1}{2}$.

On simplification using the linearity of expectation, for any v_B , we have

$$\begin{aligned}\widehat{\mathcal{R}}_B^{p_0}(p_B, v_B) &= f(p_B) + f'(p_B) \left(\frac{\frac{1}{2} + v_B}{2} - p_B \right) - f(p_0) - f'(p_0) \left(\frac{\frac{1}{2} + v_B}{2} - p_0 \right) \\ &\leq f\left(\frac{\frac{1}{2} + v_B}{2}\right) - f(p_0) - f'(p_0) \left(\frac{\frac{1}{2} + v_B}{2} - p_0 \right) \quad \forall p_B \in [0, 1], \\ &\text{equality holding only at } p_B = \frac{\frac{1}{2} + v_B}{2} \text{ due to strict convexity of } f(\cdot) \\ &= \mathcal{R}_B\left(\frac{\frac{1}{2} + v_B}{2}, p_0, \tfrac{1}{2}, v_B\right), \quad v_B \in \{0, 1\}.\end{aligned}$$

Hence, Bob's problem of choosing an action reduces to considering the difference

$$\begin{aligned}
\mathcal{R}_B\left(\frac{\frac{1}{2}+1}{2}, p_0, \frac{1}{2}, 1\right) - \mathcal{R}_B\left(\frac{\frac{1}{2}+0}{2}, p_0, \frac{1}{2}, 0\right) &= f\left(\frac{3}{4}\right) - f\left(\frac{1}{4}\right) - f'(p_0)/2 \\
&= \left[\frac{f\left(\frac{3}{4}\right) - f\left(\frac{1}{4}\right)}{\frac{1}{2}} - f'(p_0) \right] / 2 \\
&= [f'(\frac{1}{2}) - f'(p_0)] / 2 \quad \text{by symmetry condition (3.3) with } y = \frac{1}{2}.
\end{aligned}$$

Thus, by the increasing monotonicity of $f'(\cdot)$, STRATEGIC Bob should choose $(p_B = \frac{1}{4}, v_B = 0)$, or $(p_B = \frac{3}{4}, v_B = 1)$, or should remain indifferent between these two choices according to whether $p_0 (= p_A)$ is larger than, smaller than, or equal to $\frac{1}{2}$ respectively. \square

Now that we have fully characterized how the game unfolds (in a PBE) *after* Alice has taken her **Stage 1** action, the next and final step towards completing the equilibrium specification is to figure out her best-response price-report p_A . For the rest of the analysis, we will assume that the strategic selfish-rational agent Alice bases her trading decision on her belief about Bob's action, quantified by q_0 defined in (3.1) and Bob's non-participation probability $\pi \in [0, 1]$, as well as on her knowledge of Lemmas 5 and 6 above. It is germane at this point to articulate a somewhat surprising implication of Lemma 5: Although Alice can see Bob's first move p_B between making her own first and second moves, this additional piece of information has no bearing on her voting choice once she has already taken her **Stage 1** action.

To understand the impact of π on Alice's decision and hence the game outcome, let us first consider the two extreme cases $\pi = 1$ and $\pi = 0$.

3.4.1 Equilibrium when Bob's non-participation is certain ($\pi = 1$)

In this scenario, Alice knows that there is another outcome-decider Bob who will vote according to his private signal (i.e. $v_B = s_B$), and there is no way for her to influence v_B through her market action. Hence, if P^A represents Alice's subjective belief about all uncertain aspects of the game (including Bob's unobserved private signal and the future **Stage 2** outcome) given her private signal, then the "equilibrium" is fully described by specifying Alice's report-vote pair (p_A, v_A) that optimizes the expectation with respect to P^A of her

net payoff in this one-player game. For $\pi = 1$,

$$\mathbb{E}_{P^A}[v_B] = \mathbb{E}_{P^A}[s_B] = 0 \cdot q_0 + 1 \cdot (1 - q_0) = 1 - q_0, \quad (3.11)$$

where $q_0 = q_0(s)$, defined in equation (3.1), is Alice's subjective probability of Bob's signal being 0 given her own private signal. Let us define

$$\widehat{\mathcal{R}}_A(p_A, v_A) \triangleq \mathbb{E}_{P^A} [\mathcal{R}_A(p_A, p_0, v_A, v_B) | s_A, p_0].$$

Then, from equations (3.2), (3.4), and (3.11), using the linearity of expectation, we have,

$$\begin{aligned} \widehat{\mathcal{R}}_A(p_A, v_A) &= f(p_A) + f'(p_A) \left(\frac{v_A + 1 - q_0}{2} - p_A \right) - f(p_0) - f'(p_0) \left(\frac{v_A + 1 - q_0}{2} - p_0 \right) \\ &\leq \mathcal{R}_A\left(\frac{v_A + 1 - q_0}{2}, p_0, v_A, 1 - q_0\right) \quad \forall p_A \in [0, 1], \end{aligned}$$

equality holding only for $p_A = \frac{v_A + 1 - q_0}{2}$. Thus, for determining Alice's best response, it suffices to compare $\mathcal{R}_A(\frac{1 - q_0}{2}, p_0, 0, 1 - q_0)$ and $\mathcal{R}_A(1 - \frac{q_0}{2}, p_0, 1, 1 - q_0)$; the difference between them simplifies to

$$\begin{aligned} \Delta_{\mathcal{R}_A}(q_0, p_0) &\triangleq \mathcal{R}_A(1 - \frac{q_0}{2}, p_0, 1, 1 - q_0) - \mathcal{R}_A(\frac{1 - q_0}{2}, p_0, 0, 1 - q_0) \\ &= f(1 - \frac{q_0}{2}) - f(\frac{1 - q_0}{2}) - \frac{1}{2}f'(p_0) \\ &= \left[(1 - q_0)f'(\frac{1}{2}) + f(\frac{q_0}{2}) \right] - f(\frac{1 - q_0}{2}) - \frac{1}{2}f'(p_0) \\ &= f(1 - \frac{q_0}{2}) - \left[f(\frac{1 + q_0}{2}) - q_0f'(\frac{1}{2}) \right] - \frac{1}{2}f'(p_0), \end{aligned}$$

the last two equalities following from the symmetry condition (3.3), plugging in $y = 1 - q_0$ and $y = q_0$ respectively.

Obviously, Alice should take the action $(p_A = \frac{1 - q_0}{2}, v_A = 0)$ if $\Delta_{\mathcal{R}_A}(q_0, p_0) < 0$, the action $(p_A = 1 - \frac{q_0}{2}, v_A = 1)$ if $\Delta_{\mathcal{R}_A}(q_0, p_0) > 0$, and should be indifferent between these two actions otherwise. Thus, for a general value of p_0 , Alice's decision making depends in a complex manner, through $\Delta_{\mathcal{R}_A}(q_0, p_0)$, on the interaction between the properties of the convex function $f(\cdot)$ and the magnitudes of q_0 and p_0 , and so her equilibrium action cannot be expressed as a simple function of her posterior belief q_0 . However, for the special but

practically important case $p_0 = \frac{1}{2}$, we can obtain clean and insightful results as follows.

$$\Delta_{\mathcal{R}_A}(q_0, \frac{1}{2}) = (\frac{1}{2} - q_0)f'(\frac{1}{2}) + f(\frac{q_0}{2}) - f(\frac{1-q_0}{2}) = f(1 - \frac{q_0}{2}) - f(\frac{1+q_0}{2}) + (q_0 - \frac{1}{2})f'(\frac{1}{2}).$$

From the strict convexity of $f(\cdot)$, we have $f(\frac{q_0}{2}) > f(\frac{1-q_0}{2}) + (q_0 - \frac{1}{2})f'(\frac{1-q_0}{2})$ so that

$$\Delta_{\mathcal{R}_A}(q_0, \frac{1}{2}) > (\frac{1}{2} - q_0) (f'(\frac{1}{2}) - f'(\frac{1-q_0}{2})) > 0, \quad \forall q_0 \in (0, \frac{1}{2})$$

since $f'(\frac{1}{2}) > f'(\frac{1-q_0}{2}) \forall q_0 \in (0, 1)$ owing to the increasing monotonicity of $f'(\cdot)$. Likewise, we have $f(\frac{1+q_0}{2}) > f(1 - \frac{q_0}{2}) + (q_0 - \frac{1}{2})f'(1 - \frac{q_0}{2})$ so that

$$\Delta_{\mathcal{R}_A}(q_0, \frac{1}{2}) < -(q_0 - \frac{1}{2}) (f'(1 - \frac{q_0}{2}) - f'(\frac{1}{2})) < 0, \quad \forall q_0 \in (\frac{1}{2}, 1)$$

since $f'(1 - \frac{q_0}{2}) > f'(\frac{1}{2}) \forall q_0 \in (0, 1)$ owing to the increasing monotonicity of $f'(\cdot)$. Obviously, $\Delta_{\mathcal{R}_A}(\frac{1}{2}, \frac{1}{2}) = 0$.

We conclude that, for $p_0 = \frac{1}{2}$ and $\pi = 1$, Alice's equilibrium action is $(p_A = \frac{1-q_0}{2}, v_A = 0)$ if $q_0 > \frac{1}{2}$, and $(p_A = 1 - \frac{q_0}{2}, v_A = 1)$ if $q_0 < \frac{1}{2}$. She is indifferent between the report-vote pairs $(p_A = \frac{1}{4}, v_A = 0)$ and $(p_A = \frac{3}{4}, v_A = 1)$ if $q_0 = \frac{1}{2}$.

In summary, if the market starts out with a uniform prior over all possible **Stage 2** outcomes, and it is commonly known that Bob is HONEST, then in a PBE, Alice picks the *mode* of her (binary) posterior distribution over the possible values of Bob's signal as her own vote v_A^{PBE} after moving the market price to $p_A^{PBE} = \frac{v_A^{PBE} + 1 - q_0}{2}$. Although it is unsurprising that the strategic player Alice's actions are not necessarily consistent with her private signal s_A , the above equilibrium analysis has certain interesting implications for the informativeness of her price-report p_A^{PBE} that do not easily follow from intuition.

First of all, for $p_0 = \frac{1}{2}$ and any $q_0 \in (0, 1)$, note that $p_A^{PBE} \neq p_0$; hence, using Lemma 5, we can figure v_A^{PBE} out right as a simple function of p_A^{PBE} after Alice's **Stage 1** action:

$$v_A^{PBE} = \begin{cases} 0 & p_A^{PBE} < \frac{1}{2}, \\ 1 & p_A^{PBE} > \frac{1}{2}. \end{cases}$$

Next, since v_B in this scenario has no dependence on p_A^{PBE} , it be shown be from the linearity of (conditional) expectation that p_A^{PBE} satisfies the following fixed-point equation.

$$\begin{aligned}\mathbb{E}_{P^A} [v|p_A = p_A^{PBE}] &= \frac{v_A^{PBE} + \mathbb{E}_{P^A}[v_B]}{2} = \frac{v_A^{PBE} + (1 - q_0)}{2} \\ &= p_A^{PBE},\end{aligned}$$

i.e. Alice's price-report is the Bayesian estimate of the market outcome given all the information available to everyone in the world but Bob just after Alice takes her **Stage 1** action. Finally, note that

$$p_A^{PBE} = (1 - \frac{q_0}{2}) \in (\frac{3}{4}, 1) \quad \forall q_0 \in (0, \frac{1}{2}) \quad \implies \quad q_0 = 2(1 - p_A^{PBE}) \quad \forall p_A^{PBE} \in (\frac{3}{4}, 1);$$

$$p_A^{PBE} = \frac{1-q_0}{2} \in (0, \frac{1}{4}) \quad \forall q_0 \in (\frac{1}{2}, 0) \quad \implies \quad q_0 = 1 - 2p_A^{PBE} \quad \forall p_A^{PBE} \in (0, \frac{1}{4});$$

p_A^{PBE} is either $\frac{1}{4}$ or $\frac{3}{4}$ only if $q_0 = \frac{1}{2}$, and it can never lie in $(\frac{1}{4}, \frac{3}{4})$. Hence, an external observer can deduce Alice's posterior belief q_0 uniquely from p_A in an equilibrium. If we further assume the stochastic relevance of s_A for s_B (Definition 3) within a common-knowledge signal structure, then the observer can recover Alice's private signal s_A from the above value of q_0 , *regardless of her actual vote (announced signal) v_A* ! In view of these characteristics, we can call an equilibrium of this type a “partially revealing equilibrium”.

3.4.2 Equilibrium when Bob's participation is certain ($\pi = 0$)

In this scenario, Alice knows that Bob's signal has no bearing on his action and, in fact, she can fully control his actions, as indicated in Lemma 6. From Table 3.2, it follows that for p_A lying in each of the outer sub-intervals $[0, \min(p_0, p^L))$ and $(\max(p_0, p^H), 1]$, where Bob (being deterministically STRATEGIC) definitely disagrees with Alice, the average vote is $\frac{1}{2}$ so that Alice's ex post payoff is given by the function

$$\begin{aligned}\tilde{\mathcal{R}}_{p_0}(p_A) &= \mathcal{R}_A(p_A, p_0, 0, 1) = \mathcal{R}_A(p_A, p_0, 1, 0) \\ &= f(p_A) + f'(p_A) \left(\frac{1}{2} - p_A\right) - f(p_0) - f'(p_0) \left(\frac{1}{2} - p_0\right) \quad \forall p_0 \in (0, 1).\end{aligned}$$

The first derivative of the above function with respect to p_A over $[0, 1]$ is given by

$$\tilde{\mathcal{R}}'_{p_0}(p_A) = f''(p_A) \left(\frac{1}{2} - p_A\right) \gtrless 0 \iff p_A \lesseqgtr \frac{1}{2} \text{ since } f''(p_A) > 0.$$

Thus, $\tilde{\mathcal{R}}_{p_0}(\cdot)$ is a strictly convex function for $p_A \in [0, 1]$ with a unique global maximum at $p_A = \frac{1}{2}$. Since $\min(p_0, p^L) \leq p^L < \frac{1}{2} < p^H \leq \max(p_0, p^H)$, the suprema of the segments of Alice's actual overall payoff function over $[0, \min(p_0, p^L))$ and $(\max(p_0, p^H), 1]$ are at the respective "inner" extremities $\min(p_0, p^L)$ and $\max(p_0, p^H)$. However, the behavior of the payoff function over the inner interval(s) depends strongly on the relative magnitudes of p^L , p^H , and p_0 . Since this makes the general analysis technically complicated and hard to interpret qualitatively, we focus on the case $p_0 = \frac{1}{2}$, as in Section 3.4.1.

For $p_0 = \frac{1}{2}$, Table 3.2 tells us that Alice's ex post net payoff as a function of p_A over the sub-intervals $[0, p^L)$, $(p^L, \frac{1}{2})$, $(\frac{1}{2}, p^H)$, and $(p^H, 1]$ is given by the corresponding segments of $\mathcal{R}_A(p_A, \frac{1}{2}, 0, 1)$, $\mathcal{R}_A(p_A, \frac{1}{2}, 0, 0)$, $\mathcal{R}_A(p_A, \frac{1}{2}, 1, 1)$, and $\mathcal{R}_A(p_A, \frac{1}{2}, 1, 0)$; hence, the p_A that maximizes the overall payoff function, given Lemmas 5 and 6, can be obtained by analyzing these four function segments.

First note, from expressions (3.4) and (3.2), that

$$\mathcal{R}_A(\frac{1}{2}, \frac{1}{2}, 0, 1) = \mathcal{R}_A(\frac{1}{2}, \frac{1}{2}, 0, 0) = \mathcal{R}_A(\frac{1}{2}, \frac{1}{2}, 1, 1) = \mathcal{R}_A(\frac{1}{2}, \frac{1}{2}, 1, 0) = 0. \quad (3.12)$$

Moreover, by considering the first derivative of each of these functions as above, we can show that each is strictly convex over $[0, 1]$ with unique global maxima at $\frac{1}{2}$, 0, 1, and $\frac{1}{2}$ respectively. Since $0 < p^L < \frac{1}{2} < p^H < 1$, the local suprema for the sub-intervals $[0, p^L)$, $(p^L, \frac{1}{2})$, $(\frac{1}{2}, p^H)$, and $(p^H, 1]$ are at p^L , p^L , p^H , and p^H respectively; let the values of the corresponding suprema be denoted by $\mathcal{R}_{0,1}^*$, $\mathcal{R}_{0,0}^*$, $\mathcal{R}_{1,1}^*$, and $\mathcal{R}_{1,0}^*$. Note that

$$\mathcal{R}_{0,1}^* < \mathcal{R}_{0,0}^* = \mathcal{R}_{1,1}^* > \mathcal{R}_{1,0}^*,$$

which we can establish as follows.

$$\begin{aligned}
\mathcal{R}_{0,1}^* &= \mathcal{R}_A(p^L, \tfrac{1}{2}, 0, 1) \\
&< \mathcal{R}_A(\tfrac{1}{2}, \tfrac{1}{2}, 0, 1) \quad \text{since } \mathcal{R}_A(p_A, \tfrac{1}{2}, 0, 1) \text{ is uniquely maximized at } p_A = \tfrac{1}{2}, \\
&= 0 \quad \text{from (3.12),} \\
&= \mathcal{R}_A(\tfrac{1}{2}, \tfrac{1}{2}, 0, 0) \quad \text{from (3.12),} \\
&< \mathcal{R}_A(0, \tfrac{1}{2}, 0, 0) \quad \text{since } \mathcal{R}_A(p_A, \tfrac{1}{2}, 0, 0) \text{ is uniquely maximized at } p_A = 0, \\
&= \mathcal{R}_{0,0}^*;
\end{aligned}$$

by similar reasoning, $\mathcal{R}_{1,1}^* > \mathcal{R}_{1,0}^*$; and

$$\begin{aligned}
\mathcal{R}_{1,1}^* - \mathcal{R}_{0,0}^* &= \mathcal{R}_A(p^H, \tfrac{1}{2}, 1, 1) - \mathcal{R}_A(p^L, \tfrac{1}{2}, 0, 0) \\
&= [f(p^H) - f(p^L)] + p^L [f'(p^H) + f'(p^L)] - f'(\tfrac{1}{2}), \\
&\quad \text{using } p^L = 1 - p^H \text{ from (3.9),} \\
&= (1 - 2p^L)f'(\tfrac{1}{2}) + p^L \cdot 2f'(\tfrac{1}{2}) - f'(\tfrac{1}{2}), \quad \text{from (3.10) and (3.8)} \\
&= 0.
\end{aligned}$$

Figure 3.1 illustrates the variation of \mathcal{R}_A with p_A over $[0, 1]$ for LMSR, QMSR, and SMSR. From the above discussion, it is clear that this special case of the game has two PBEs with strategy profiles $((p^L, 0), (0, 0))$ and $((p^H, 1), (1, 1))$. If Alice makes the first move $p_A = p^L$, then $(p_B = 0, v_B = 0)$ is a best response by Bob according to Lemma 6, and Alice's preferable **Stage 2** action is $v_A = 0$ from Lemma 5; from Alice's perspective, if she knows that Bob will respond with $v_B = 1$ for $0 \leq p_A < p^L$ and $v_B = 0$ for $p^L \leq p_A \leq \frac{1}{2}$, then it is in her best interest to set $p_A = p^L$, and hence vote $v_A = 0$ in accordance with Lemma 5, thereby resulting in a market outcome of $v = 0$. Although $(p_B = \frac{1}{2}, v_B = 1)$ is Bob's alternative best response to Alice's $p_A = p^L$ by Lemma 6, it cannot be part of a (subgame perfect) Nash equilibrium where Alice's action is $(p^L, 0)$ because Alice would prefer $p_A = p^L + \varepsilon$ for any $\varepsilon \in (0, \frac{1}{2} - p^L)$ to $p_A = p^L$ if she knew that Bob would respond with $v_B = 1$ for $0 \leq p_A \leq p^L$ and $v_B = 0$ for $p^L < p_A \leq \frac{1}{2}$, owing to the jump discontinuity at p^L . By similar arguments, we can establish that $((p^H, 1), (1, 1))$ is another equilibrium.

Thus, Alice and Bob jointly create a fake world where they pretend to agree regardless of their signals, and thereby reap the maximum possible profit from the mechanism (the

principal organizing the market incurs maximum loss whenever Bob participates since the last trader Bob makes an accurate forecast, in accordance with Lemma 6 Hanson (2007a)). We call equilibria of this type “collusive”. Further, observe that for $\pi = 0$, Alice’s price-report gives us no information about her private signal or her belief about that of Bob; it merely indicates her vote as per Lemma 5 ($v_A = 0$ if $p_A = p^L$ and $v_A = 1$ if $p_A = p^H$).




Figure 3.1: Alice’s payoff as a function of her **Stage 1** action for three representative MSRs with Bob’s non-participation probability $\pi = 0$ and starting market price $p_0 = \frac{1}{2}$; the general characteristics are similar for other symmetric well-behaved MSRs: We have jump discontinuities at the lower and upper thresholds, and a value of zero at $p_A = \frac{1}{2}$. For LMSR, the figure is truncated, and both the outer segments actually continue towards $-\infty$ symmetrically, away from $\frac{1}{2}$.

3.4.3 Equilibrium when Bob’s participation is uncertain ($0 < \pi < 1$)

We shall now delve into the “grey area” where Alice as well as the rest of the world has some finite uncertainty about Bob’s participation. In view of the lessons learned from the relatively simpler scenarios $\pi \in \{0, 1\}$, we shall present the analysis of this more general scenario $\pi \in (0, 1)$ for a starting market price of $p_0 = \frac{1}{2}$ only¹⁶, in order to obtain clean results; in fact, we included the symmetry assumption (3.3) in our model (Section 3.3)

¹⁶The initial price $p_0 = \frac{1}{2}$ corresponds to starting the market at a uniform “prior” – a standard practice in prediction markets.

for the same reason. For a larger family of MSRs that satisfy all the criteria stated in Section 3.3 except the symmetry condition and for values of p_0 other than $\frac{1}{2}$, one can still adopt an approach similar to the one in this paper to analyze the game but the procedure will be algebraically more involved.

It is not unreasonable to conjecture that there exists some critical value of Bob's non-participation probability below which the PBE of the resulting game is of the partially revealing variety as in Section 3.4.2, and above which it is closer to the collusive equilibrium obtained in Section 3.4.1. The following theorem formalizes this intuition.

Theorem 4. *For any value of Bob's non-participation probability $\pi \in (0, 1)$ and Alice's posterior belief $q_0 \in (0, 1)$, the trading-voting game described in Section 3.3 has a perfect Bayesian equilibrium with the following attributes:*

For every q_0 , there exists a fixed value of Bob's non-participation probability, say $\pi_c(q_0)$, which we call the "crossover" probability (dependent on the MSR), on either side of which the equilibria are qualitatively different. We call the sub-interval $\pi < \pi_c$ the high participation probability (HPP) equilibrium domain, and the sub-interval $\pi > \pi_c$ the low participation probability (LPP) equilibrium domain.

• **In an HPP equilibrium:**

- In **Stage 1**, Alice moves the market price to $p_A = p^L$ if $q_0 > \frac{1}{2}$, and to $p_A = p^H$ if $q_0 < \frac{1}{2}$ where p^L, p^H are the upper and lower thresholds, independent of π, q_0 , defined in (3.5) and (3.6); STRATEGIC Bob's price-update is $p_B = 0$ if $p_A = p^L$, and $p_B = 1$ if $p_A = p^H$.
- In **Stage 2**, Alice votes $v_A = 0$ if she set $p_A = p^L$, $v_A = 1$ if $p_A = p^H$; STRATEGIC Bob votes $v_B = 0$ if he set $p_B = 0$, and $v_B = 1$ if he set $p_B = 1$.

• **In an LPP equilibrium:**

- In **Stage 1**, Alice's price-report p_A^{LPP} is equal to her posterior expectation of the market liquidation value (average vote) given the parameters π, q_0 and her report p_A^{LPP} , i.e. $p_A^{LPP} = \mathbb{E}[v|\pi, q_0, p_A = p_A^{LPP}]$. Moreover, $p_A^{LPP} < \frac{1}{2}$ if $q_0 > \frac{1}{2}$, $p_A^{LPP} > \frac{1}{2}$ if $q_0 < \frac{1}{2}$. STRATEGIC Bob's price-update is $p_B = 0$ if $p^L \leq p_A \leq \frac{1}{2}$, $p_B = 1$ if $\frac{1}{2} < p_A \leq p^H$, and $p_B = \frac{1}{2}$ otherwise.

– In **Stage 2**, Alice votes $v_A = 0$ if $p_A > \frac{1}{2}$, $v_A = 1$ if $p_A < \frac{1}{2}$; STRATEGIC Bob votes $v_B = 0$ if $p_A \in [p^L, \frac{1}{2}] \cup (p^H, 1]$, $v_B = 1$ otherwise.

Figure 3.2 illustrates the general characteristics of Alice’s and Bob’s actions in the two equilibrium domains for $q_0 > \frac{1}{2}$.




Figure 3.2: Alice’s price-reports for $q_0 > \frac{1}{2}$ and STRATEGIC Bob’s responses for HPP (left) and LPP domains. The actual magnitude of p_A in the LPP domain depends on values of π, q_0 . Results are symmetric for $q_0 < \frac{1}{2}$.

More specifically, p_A^{LPP} is one of the following:

$$\mu_{0,0} = \frac{\pi(1-q_0)}{2}, \quad \mu_{0,1} = \frac{1-\pi q_0}{2}, \quad \mu_{1,0} = \frac{1+\pi(1-q_0)}{2}, \quad \mu_{1,1} = 1 - \frac{\pi q_0}{2},$$

$$\text{where } 0 < \mu_{0,0} < \mu_{0,1} < \frac{1}{2} < \mu_{1,0} < \mu_{1,1} < 1, \quad \forall \pi, q_0 \in (0, 1). \quad (3.13)$$

Tables 3.3, 3.4, and 3.5 detail the dependence of Alice’s equilibrium trading action p_A^{PBE} on π, q_0 for any symmetric well-behaved market scoring rule with $p^L \in (0, \frac{1}{4})$, $p^L = \frac{1}{4}$, and $p^L \in (\frac{1}{4}, \frac{1}{2})$ respectively (“Domain” signifies whether the equilibrium is of the HPP or LPP type as defined above). The quantities $\pi_H^*(q_0)$ and $\pi_L^*(q_0)$ in Tables 3.3 and 3.5, and π^* in Table 3.5 are defined as follows:

- $\pi_H^*(q_0)$ is the unique root in $(0, 1)$ of the following equation in π :

$$f\left(\frac{1+\pi(1-q_0)}{2}\right) - f(p^H) - \left(1 - \frac{\pi q_0}{2} - p^H\right) f'(p^H) + \left(\frac{1-\pi}{2}\right) f'\left(\frac{1}{2}\right) = 0$$

for a given convex function $f(\cdot)$ associated with a symmetric well-behaved MSR (which has a well-defined p^H), and any given value of the parameter $q_0 \in (0, 1)$. For $0 < q_0 < \min(\frac{1}{2}, 2p^L)$, we have $\frac{1-2p^L}{1-q_0} < \pi_H^*(q_0) < 1$.

- $\pi_L^*(q_0)$ is the unique root in $(0, 1)$ of the equation

$$f\left(\frac{1-\pi q_0}{2}\right) - f(p^L) - \left(\frac{\pi(1-q_0)}{2} - p^L\right) f'(p^L) - \left(\frac{1-\pi}{2}\right) f'\left(\frac{1}{2}\right) = 0$$

for a given $f(\cdot)$ and q_0 as before. For $\max(\frac{1}{2}, 1 - 2p^L) < q_0 < 1$, we have $\frac{1-2p^L}{q_0} < \pi_L^*(q_0) < 1$.

- π^* is the unique root in $(2 - 4p^L, 1)$ of the equation

$$f\left(\frac{1}{2} - \frac{\pi}{4}\right) - f(p^L) - \left(\frac{\pi}{4} - p^L\right) f'(p^L) - \left(\frac{1-\pi}{2}\right) f'\left(\frac{1}{2}\right) = 0$$

for a given $f(\cdot)$ with $p^L \in (\frac{1}{4}, 1)$.

For the proof of existence and uniqueness of each of the above roots, refer to Appendix B Lemma 10. Table 3.6 presents the crossover probability π_c as a function of q_0 for all three mutually exclusive and exhaustive sub-classes of symmetric well-behaved MSRs, defined on the basis of the value of p^L .

Proof sketch. Here we just outline the proof of Theorem 4, the details are in Appendix B Section B.2. As in Section 3.4.1 for the case $\pi = 1$, we have to work with Alice's expected ex post profit that averages out her uncertainty with respect to two questions – what signal Bob received, and whether he will trade in the prediction market. Owing to the linearity of the profit function (3.4) in the market outcome, the expectation of the profit across outcomes is equal to the profit function evaluated at the expected outcome. The rest of the proof is similar to the analysis presented in Section 3.4.2 for the case $\pi = 0$: We need to consider the four sub-intervals $[0, p^L)$, $(p^L, \frac{1}{2}]$, $(\frac{1}{2}, p^H)$, $(p^H, 1]$, over which Alice's expected ex post payoff as a function of her price-report p_A is given by the corresponding segments of the functions $\mathcal{R}_A(p_A, \frac{1}{2}, 0, 1)$, $\mathcal{R}_A(p_A, \frac{1}{2}, 0, 0)$, $\mathcal{R}_A(p_A, \frac{1}{2}, 1, 1)$, and $\mathcal{R}_A(p_A, \frac{1}{2}, 1, 0)$ respectively, with special emphasis on the jump discontinuities at the thresholds p^L and p^H . An important difference from Section 3.4.2 is that, for $\pi \in (0, 1)$, the global maxima of these functions are no longer located at $\frac{1}{2}$, 0, 1, and $\frac{1}{2}$ respectively but at $\mu_{u,v} = \frac{u+(1-\pi)v+\pi(1-q_0)}{2}$ for $u, v \in \{0, 1\}$. So, the local suprema of some of the four segments that we are interested in may not lie at the extremities of the corresponding intervals but rather at an “interior” point, depending in a complex manner on the values of p^L , π and q_0 . For example, if $\pi < \frac{2p^L}{1-q_0}$, then the local

q_0	π	p_A^{PBE} (for $p^L < \frac{1}{4}$)	Domain
$0 < q_0 < 2p^L$	$0 < \pi < \pi_H^*(q_0)$	p^H	HPP
	$\pi = \pi_H^*(q_0)$	p^H or $\mu_{1,0}$	LPP or HPP
	$\pi_H^*(q_0) < \pi < 1$	$\mu_{1,0}$	LPP
$q_0 = 2p^L$	$0 < \pi < 1$	p^H	HPP
$2p^L < q_0 < \frac{1}{2}$	$0 < \pi < \frac{2p^L}{q_0}$	p^H	HPP
	$\pi = \frac{2p^L}{q_0}$	$p^H = \mu_{1,1}$	LPP or HPP
	$\frac{2p^L}{q_0} < \pi < 1$	$\mu_{1,1}$	LPP
$q_0 = \frac{1}{2}$	$0 < \pi < 4p^L$	p^L or p^H	HPP
	$\pi = 4p^L$	$p^L = \mu_{0,0}$ or $p^H = \mu_{1,1}$	LPP or HPP
	$4p^L < \pi < 1$	$\mu_{0,0}$ or $\mu_{1,1}$	LPP
$\frac{1}{2} < q_0 < 1 - 2p^L$	$0 < \pi < \frac{2p^L}{1-q_0}$	p^L	HPP
	$\pi = \frac{2p^L}{1-q_0}$	$p^L = \mu_{0,0}$	LPP or HPP
	$\frac{2p^L}{1-q_0} < \pi < 1$	$\mu_{0,0}$	LPP
$q_0 = 1 - 2p^L$	$0 < \pi < 1$	p^L	HPP
$1 - 2p^L < q_0 < 1$	$0 < \pi < \pi_L^*(q_0)$	p^L	HPP
	$\pi = \pi_L^*(q_0)$	p^L or $\mu_{0,1}$	LPP or HPP
	$\pi_L^*(q_0) < \pi < 1$	$\mu_{0,1}$	LPP

Table 3.3: Alice's PBE price-report for a symmetric well-behaved MSR with $p^L \in (0, \frac{1}{4})$, $p_0 = \frac{1}{2}$.

supremum over $[0, p^L)$ of the overall \mathcal{R}_A is at $\mu_{0,0} \in (0, p^L)$. Taking these issues into account, we can determine the local suprema of the four segments and compare them to establish that we have perfect Bayesian equilibria where Alice's **Stage 1** action is in accordance with Tables 3.3, 3.4, and 3.5, her **Stage 2** action is given by Lemma 5, and STRATEGIC Bob's actions are given by Lemma 6 with the restriction that $v_B = 0$ if $p_A = p^L$ and $v_B = 1$ if $p_A = p^H$ (as in Section 3.4.2). \square

Behavior at crossover probability If $\pi = \pi_c$, then Alice is indifferent between her LPP and HPP price-reports although the values of these reports are, in general, distinct. However, if the prediction market is implemented with a symmetric well-behaved MSR with $p^L < \frac{1}{4}$, and q_0 lies in $(2p^L, 1 - 2p^L)$, then Alice's LPP and HPP price-reports at the crossover probability are identical so that the two domains coincide. This is illustrated in Figure 3.4 for LMSR under a specific signal structure, the details of which are provided in Section 3.6:

q_0	π	p_A^{PBE} (for $p^L = \frac{1}{4}$)	Domain
$0 < q_0 < \frac{1}{2}$	$0 < \pi < \pi_H^*(q_0)$	p^H	HPP
	$\pi = \pi_H^*(q_0)$	p^H or $\mu_{1,0}$	LPP or HPP
$\frac{1}{2} < q_0 < 1$	$\pi_H^*(q_0) < \pi < 1$	$\mu_{1,0}$	LPP
	$0 < \pi < 1$	p^L or p^H	HPP
	$0 < \pi < \pi_L^*(q_0)$	p^L	HPP
	$\pi = \pi_L^*(q_0)$	p^L or $\mu_{0,1}$	LPP or HPP
	$\pi_L^*(q_0) < \pi < 1$	$\mu_{0,1}$	LPP

Table 3.4: Alice's PBE price-report for a symmetric well-behaved MSR with $p^L = \frac{1}{4}$, $p_0 = \frac{1}{2}$.

q_0	π	p_A^{PBE} (for $p^L > \frac{1}{4}$)	Domain
$0 < q_0 < \frac{1}{2}$	$0 < \pi < \pi_H^*(q_0)$	p^H	HPP
	$\pi = \pi_H^*(q_0)$	p^H or $\mu_{1,0}$	LPP or HPP
$\frac{1}{2} < q_0 < 1$	$\pi_H^*(q_0) < \pi < 1$	$\mu_{1,0}$	LPP
	$0 < \pi < \pi^*$	p^L or p^H	HPP
	$\pi = \pi^*$	p^L or p^H or $\mu_{0,1}$ or $\mu_{1,0}$	LPP or HPP
	$\pi^* < \pi < 1$	$\mu_{0,1}$ or $\mu_{1,0}$	LPP
$\frac{1}{2} < q_0 < 1$	$0 < \pi < \pi_L^*(q_0)$	p^L	HPP
	$\pi = \pi_L^*(q_0)$	p^L or $\mu_{0,1}$	LPP or HPP
	$\pi_L^*(q_0) < \pi < 1$	$\mu_{0,1}$	LPP

Table 3.5: Alice's PBE price-report for a symmetric well-behaved MSR with $p^L \in (\frac{1}{4}, \frac{1}{2})$, $p_0 = \frac{1}{2}$.

In the left panel, we have a signal structure where $q_0 \approx 0.52 < 1 - 2p^L = 0.6$. Hence, by Tables 3.3 and 3.6, Alice's unique PBE price-report at the crossover probability $\pi_c(q_0) = \frac{2p^L}{1-q_0} \approx 0.83$ is $p^L = \mu_{0,0} = \frac{\pi_c(1-q_0)}{2} = 0.2$. However, in the right panel, $q_0 \approx 0.82 > 0.6$, hence Alice has two alternative PBE price-reports, $p^L = 0.2$ (HPP) or $\mu_{0,1} = \frac{1-\pi_c q_0}{2} \approx 0.11$ (LPP), at the crossover probability $\pi_c(q_0) = \pi_L^*(q_0) \approx 0.96$.

Figure 3.3 depicts the crossover probability π_c as a function of q_0 for each of the three selected MSRs — logarithmic, quadratic, and spherical.

q_0	$\pi_c(q_0)$		
	$0 < p^L < \frac{1}{4}$	$p^L = \frac{1}{4}$	$\frac{1}{4} < p^L < \frac{1}{2}$
$0 < q_0 < 2p^L$	$\pi_H^*(q_0)$	$\pi_H^*(q_0)$	$\pi_H^*(q_0)$
$q_0 = 2p^L$	1	$\pi_H^*(q_0)$	$\pi_H^*(q_0)$
$2p^L < q_0 < \frac{1}{2}$	$\frac{2p^L}{q_0}$	$\pi_H^*(q_0)$	$\pi_H^*(q_0)$
$q_0 = \frac{1}{2}$	$4p^L$	1	π^*
$\frac{1}{2} < q_0 < 1 - 2p^L$	$\frac{2p^L}{1-q_0}$	$\pi_L^*(q_0)$	$\pi_L^*(q_0)$
$q_0 = 1 - 2p^L$	1	$\pi_L^*(q_0)$	$\pi_L^*(q_0)$
$1 - 2p^L < q_0 < 1$	$\pi_L^*(q_0)$	$\pi_L^*(q_0)$	$\pi_L^*(q_0)$

Table 3.6: The crossover probability π_c as a function of q_0 over the sub-intervals into which the p^L splits the entire possible range $(0, 1)$ of q_0 values, for symmetric well-behaved MSRs with p^L smaller than, equal to, and larger than $\frac{1}{4}$.

3.5 Implications of the equilibrium result

The following are some interesting corollaries to our main result, that shed light on various aspects of the operation of prediction markets in the face of outcome manipulation possibilities, including our main concerns about informativeness (Section 3.5.4) and truthfulness (Section 3.5.2).

3.5.1 Dependence of crossover probability on Alice’s uncertainty

From Table 3.6 and the representative curves in Figure 3.3, observe a peculiarity of symmetric well-behaved MSRs with $p^L < \frac{1}{4}$ such as LMSR. For $2p^L < q_0 < 1 - 2p^L$, which can be seen as a region of “high” uncertainty in Alice’s posterior about Bob’s private signal after receiving her own, the crossover probability actually decreases with Alice’s increasing uncertainty: π_c is inversely proportional to q_0 and $(1 - q_0)$ over $2p^L < q_0 < \frac{1}{2}$ and $\frac{1}{2} < q_0 < 1 - 2p^L$ respectively. This means that the partially revealing LPP domain is realized for lower values of Bob’s non-participation probability for these MSRs than for those with $p^L \geq \frac{1}{4}$ over this central region surrounding $q_0 = \frac{1}{2}$. On the flip side, outside this high uncertainty region, π_c increases faster as Alice’s certainty moves away from $q_0 \in \{0, 1\}$ for MSRs with “low” p^L -values so that there exist values of q_0 , namely $2p^L$ and $(1 - 2p^L)$ for $p^L \leq \frac{1}{4}$, for which any value of $\pi \in [0, 1]$ induces an HPP equilibrium.




Figure 3.3: Dependence of crossover probability on Alice’s posterior belief about Bob for the three MSRs; e.g. for QMSR, if $q_0 = 0.25$, then $\pi_c \approx 0.9537$, so we have an LPP equilibrium with $p_A = (1 + \pi(1 - q_0))/2$ for $\pi > 0.9537$, and an HPP equilibrium with $p_A = p^H = 0.75$ for $\pi < 0.9537$.

3.5.2 Private signal revelation

Unfortunately, STRATEGIC Bob’s report-vote pair is fully determined by Alice’s report and does not depend on s_B . There is no guarantee that Alice’s *vote* will be truthful either: In general, as Tables 3.3, 3.4, and 3.5 indicate, regardless of whether we are in the LPP or HPP domain, for $p_0 = \frac{1}{2}$,

$$\begin{aligned} q_0(s_A) < \frac{1}{2} &\implies p_A^{PBE} > \frac{1}{2} \implies v_A^{PBE} = 1; \\ q_0(s_A) > \frac{1}{2} &\implies p_A^{PBE} < \frac{1}{2} \implies v_A^{PBE} = 0, \end{aligned} \quad (3.14)$$

i.e. Alice votes the mode of her posterior distribution over Bob’s signal, as in Section 3.4.1.

However, if we invoke the additional assumption of stochastic relevance (Definition 3), then we can use p_A to uncover s_A , as in Section 3.4.1, in an LPP equilibrium. This stands in contrast to the situation where Bob’s participation is certain, the limiting case of the HPP domain, where $p_A \in \{p^L, p^H\}$ has no dependence on q_0 (Section 3.4.2).

Corollary 1. *If the signal structure is such that Alice’s signal s_A is stochastically relevant for Bob’s signal s_B , then the value of s_A can be recovered from Alice’s price-report in an LPP equilibrium $p_A^{LPP} = \mu_{u,v}(\pi, q_0)$, $u, v \in \{0, 1\}$, regardless of whether $v_A = s_A$.*

Proof. Since p_A being either p^L or p^H indicates that we are in the HPP domain, $p_A^{LPP} \notin \{p^L, p^H\}$. We also know that $p_A^{LPP} \notin \{0, \frac{1}{2}, 1\}$ from inequalities (3.13). Moreover, by observing which sub-interval, $(0, p^L)$, $(p^L, \frac{1}{2})$, $(\frac{1}{2}, p^H)$ or $(p^H, 1)$, contains Alice's LPP price-report, we can infer Alice's vote v_A and Bob's vote v_B if he participates in the market (Lemmas 5 and 6), and hence which $\mu_{u,v}$, $u, v \in \{0, 1\}$, equals p_A . Since π is common knowledge, we can solve for q_0 from the expression for $\mu_{u,v}$. For example, if $0 < p_A < p^L$, then we can be certain that $p_A = \mu_{0,1} = \frac{1-\pi q_0}{2}$ so that $q_0 = \frac{1-2p_A}{\pi}$. Under the assumption of stochastic relevance, $q_0(s_A)$ is one-to-one, so we can deduce s_A uniquely from its value. \square

In a general HPP equilibrium ($0 < \pi < \pi_c$), we can only tell whether $q_0 > \frac{1}{2}$ (if $p_A = p^L$) or $q_0 < \frac{1}{2}$ (if $p_A = p^H$), but this is insufficient for recovering s_A without further assumptions about the signal structure. The following corollary presents sufficient conditions on the signal structure for s_A to be recoverable from Alice's equilibrium action regardless of whether we are in the LPP or HPP domain.

Corollary 2. *If the signal structure is such that Alice's posterior probability q_0 of Bob obtaining signal 0, given s_A , lies strictly on different sides of $\frac{1}{2}$ for different values of $s_A \in \{0, 1\}$, i.e. either $(q_0(0) < \frac{1}{2} \text{ and } q_0(1) > \frac{1}{2})$ or $(q_0(0) > \frac{1}{2} \text{ and } q_0(1) < \frac{1}{2})$, then the value of s_A can be recovered from her PBE price-report (or, equivalently, vote).*

Proof. From (3.14), we can tell by looking at p_A (or equivalently v_A) whether $q_0 < \frac{1}{2}$ or $q_0 > \frac{1}{2}$, and can hence deduce whether $s_A = 0$ or $s_A = 1$ from the signal structure. \square

Note that the signal structure in Corollary 2 automatically implies the stochastic relevance property of Corollary 1. In light of the two corollaries above, we can conclude that whenever there is a non-zero probability of Bob not trading but voting truthfully, there are signal structures for which Alice's trading action (in equilibrium) indirectly reveals her private information!

In particular, if $q_0(0) > \frac{1}{2}$ and $q_0(1) < \frac{1}{2}$, then it follows from (3.14) that, for $p_0 = \frac{1}{2}$, $v_A = s_A$, regardless of whether the value of π puts us in the LPP or HPP domain. In other words, in our model, a sufficient condition on the binary signal structure for Alice to vote truthfully in equilibrium regardless of Bob's non-participation probability is that she believes it more likely than not for Bob to receive the same signal that she has.

3.5.3 HPP profit sharing

The HPP equilibria are a world where collusion appears with Alice as the “leader” picking the vote that both will coordinate on, and pushing the price to just the level where it makes sense for STRATEGIC Bob to push the price all the way to 0 or 1 and vote the same way as Alice. In this way, they extract the maximum profit from the market maker, and split it between the two of them in a ratio that is dependent on the functional form of the MSR. In particular, for the three major MSRs considered, Alice makes more profit than Bob in a collusive equilibrium, with the discrepancy being the least for LMSR – we omit the straightforward calculations, and present the results in the following table:

Share in total HPP profit if Bob is strategic	LMSR	QMSR	SMSR
Alice’s share	67.81%	75%	78.32%
Bob’s share	32.19%	25%	21.68%

However, it is possible to construct a symmetric well-behaved MSR with lower threshold $p^L < \frac{1}{4}$ for which Alice makes less profit than STRATEGIC Bob in an HPP equilibrium: For $f(r) = 0.99(r - \frac{1}{2})^6 + 0.01r^2$, $0 \leq r \leq 1$,¹⁷ with $p^L \approx 0.1550 < \frac{1}{4}$, Alice’s HPP profit share is approximately 46.91%!

Corollary 3. *In a trading-voting game where the prediction market is implemented by any symmetric well-behaved MSR with lower threshold $p^L \geq \frac{1}{4}$, Alice’s ex post net profit in an HPP equilibrium is greater than that of STRATEGIC Bob.*

Proof. Regardless of the value of q_0 , if we are in the HPP domain, the starting price is $p_0 = \frac{1}{2}$, and Bob trades in **Stage 1**, then Alice and Bob’s ex post net profits in an HPP equilibrium are respectively given by

$$\begin{aligned}
\mathcal{R}_A^{HPP} &= f(p^H) + (1 - p^H)f'(p^H) - f(\tfrac{1}{2}) - \tfrac{1}{2}f'(\tfrac{1}{2}) \\
&= f(p^L) - p^L f'(p^L) - f(\tfrac{1}{2}) + \tfrac{1}{2}f'(\tfrac{1}{2}); \\
\mathcal{R}_B^{HPP} &= f(1) - f(p^H) - (1 - p^H)f'(p^H) \\
&= f(0) - f(p^L) + p^L f'(p^L).
\end{aligned}$$

¹⁷The term $0.01r^2$ ensures that $f''(\frac{1}{2}) \neq 0$ so that the associated MSR satisfies all our technical desiderata.

These expressions follow readily from Alice and Bob's HPP report-vote pairs as shown in Tables 3.3, 3.4, 3.5, and 3.2, as well as definitions (3.4) and (3.2). To prove the equivalence of the two expressions for \mathcal{R}_A^{HPP} , we invoke relations (3.8), (3.9), and (3.10) from Proposition 2, and simplify. For \mathcal{R}_A^{HPP} , we additionally use the result $f(1) - f(0) = f'(\frac{1}{2})$ obtained from the symmetry condition (3.3) by plugging in $y = 1$. Hence,

$$\begin{aligned}
\mathcal{R}_A^{HPP} - \mathcal{R}_B^{HPP} &= 2 [f(p^L) - p^L f'(p^L)] - f(0) - f(\tfrac{1}{2}) + \tfrac{1}{2} f'(\tfrac{1}{2}) \\
&= 2 [f(p^L) - 2p^L (f(\tfrac{1}{2}) - f(0))] - f(0) - f(\tfrac{1}{2}) + \tfrac{1}{2} f'(\tfrac{1}{2}), \\
&\quad \text{from definition (3.5)} \\
&= 2f(p^L) + (4p^L - 1)f(0) - (1 + 4p^L)f(\tfrac{1}{2}) + \tfrac{1}{2} f'(\tfrac{1}{2}) \\
&> 2 [f(\tfrac{1}{2}) + (p^L - \tfrac{1}{2}) f'(\tfrac{1}{2})] + (4p^L - 1)f(0) - (1 + 4p^L)f(\tfrac{1}{2}) + \tfrac{1}{2} f'(\tfrac{1}{2}), \\
&\quad \text{from the strict convexity of } f(\cdot) \\
&= 4(p^L - \tfrac{1}{4}) [f(0) - f(\tfrac{1}{2}) + \tfrac{1}{2} f'(\tfrac{1}{2})], \\
&\geq 0, \quad \text{for } p^L \geq \tfrac{1}{4}
\end{aligned}$$

since $f(0) > f(\frac{1}{2}) + (0 - \frac{1}{2})f'(\frac{1}{2})$ from the strict convexity of $f(\cdot)$. Hence $\mathcal{R}_A^{HPP} > \mathcal{R}_B^{HPP}$ for any symmetric well-behaved MSR with $p^L \geq \frac{1}{4}$. \square

If Bob is HONEST, Alice's payoff is obviously a function of his private signal faithfully announced in the outcome-deciding voting stage. Corollary 3 tells us that, even if Bob is STRATEGIC and hence ends up colluding with the manipulator Alice, her profit share in a collusive equilibrium depends strongly on the MSR used – an insight that can potentially inform the choice of an MSR for market design.

3.5.4 Informativeness of market prices about final outcome

Rewriting all our results so as to focus on what the market prices p_A at $t = 1$ and p_B at $t = 2$ predict about the outcome (average vote), we obtain the following table (recall that the final price $p_B = p_A$ for HONEST Bob, and $p_B = v$ for STRATEGIC Bob):

	strategic Bob		honest Bob	
	LPP	HPP	LPP	HPP
p_A	Bayesian estimate	Predetermined	Bayesian estimate	Predetermined
p_B	Actual outcome	Actual outcome	Bayesian estimate	Predetermined
<p>“Predetermined” signifies that $p_A \in \{p^L, p^H\}$;</p> <p>“Bayesian estimate” denotes Alice’s expectation of the average vote.</p>				

3.6 A specific signal structure

Thus far, we have been non-specific about the signal structure, proving general results; we now consider a concrete example scenario to illustrate our findings: The underlying random variable takes values in the signal space itself, i.e. $T = \Omega = \{0, 1\}$, the prior probability of $\tau = 0$ being $\rho_0 \in (0, 1)$. Given τ , the agents’ signals are independently and identically distributed: for any “true” $\tau \in \{0, 1\}$, each participant gets the “correct” signal (identical to the true τ) with probability $(1 - \rho_e)$, otherwise gets the wrong signal; the error probability $\rho_e \in (0, 1) \setminus \{\frac{1}{2}\}$. Note that if and only if $\rho_e = \frac{1}{2}$, we have $q_0(0) = q_0(1) = \frac{1}{2}$ regardless of ρ_0 , hence signals are not *informative* Chen et al. (2009), i.e. the prior and posterior probabilities are equal. Then,

$$q_0(0) = \frac{(1 - \rho_e)^2 \rho_0 + \rho_e^2 (1 - \rho_0)}{(1 - \rho_e) \rho_0 + \rho_e (1 - \rho_0)}, \quad q_0(1) = \frac{(1 - \rho_e) \rho_e}{\rho_e \rho_0 + (1 - \rho_e) (1 - \rho_0)}.$$

This signal structure has multiple interesting information-revealing characteristics: First, we have $q_0(0) \neq q_0(1)$, i.e. Alice’s signal is stochastically relevant for that of Bob. Hence, Corollary 1 applies. Second, it is easy to show that, if $\rho_0 = \frac{1}{2}$ (a uniform common prior), then Alice’s vote is always truthful since, for any $\rho_e \in (0, 1)$, $s_A = 0 \Leftrightarrow q_0 > \frac{1}{2} \Leftrightarrow v_A = 0$, by (3.4).

Figure 3.4 shows Alice’s equilibrium report in a LMSR market and her expected liquidation value vs. π , for signal $s_A = 0$ and fixed ρ_0, ρ_e (hence, a fixed q_0). The HPP and LPP regions are clearly visible to the left and right of the cross-over probability, where Alice’s price-report (the dashed curve) is distinct from and coincides with her expected liquidation value (the continuous curve) respectively. The corresponding plots for the other two MSRs are qualitatively similar, hence omitted.

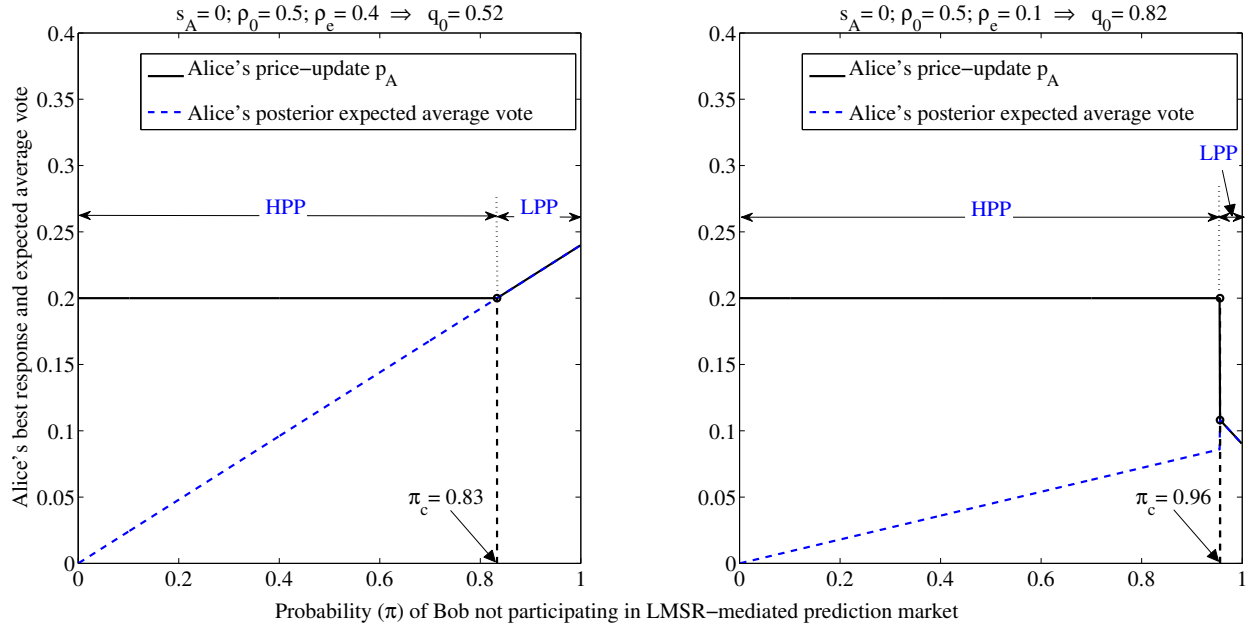


Figure 3.4: Crossover from HPP to LPP equilibria regions for LMSR over $0 < \pi < 1$ for the signal structure described in Section 3.6 with two different sets of parameter values: the prior probability of τ being 0 is $\rho_0 = \frac{1}{2}$ for both panels, and the error probability ρ_e is 0.4 and 0.1 for the left and right panels respectively. For the left panel, $p_A^{LPP} = \mu_{0,0} = \frac{\pi(1-q_0)}{2}$, hence p_A increases with π for $\pi > \pi_c$ whereas for the right panel, $p_A^{LPP} = \mu_{0,1} = \frac{1-\pi q_0}{2}$, hence p_A decreases with π in the LPP domain. In both panels, $p_A = p^L = 0.2$ in the HPP domain. For each curve, we have shown with a dashed blue line Alice's expected average vote taking Lemmas 5 and 6 into account, which coincides with p_A in the LPP domain.

3.7 Extensions to more than two agents

Our model is stylized, but the framework and methodology can be applied to more complex scenarios. Below, we sketch two specific lines of generalization.

Additional outcome-deciders who do not trade: Consider a scenario in which Alice and Bob are the only traders but jointly decide less than 100% of the outcome, say, $v = \frac{v_A + v_B + \sum_{i=1}^n v_i}{n+2}$, where $\{v_i\}_{i=1}^n$ are the votes (and also the private signals) of n non-strategic agents. We can still use the methodology of Section 3.4 to solve for the PBE, and show that it is still of two broad types – we do need further specifications for a consistent belief

structure, and the equilibria may have some additional characteristics depending on model parameters.

For example, suppose that there are three outcome-deciders Alice, Bob, and Charlie each of whom draws a signal from a structure identical to that in Section 3.6 but Charlie is deterministically HONEST; the market outcome is $v = \frac{v_A + v_B + v_C}{3}$, where C is the subscript corresponding to Charlie, and all other aspects of the model are the same as in Section 3.3. It is easy to show that the natural extension of Lemma 5 still holds, i.e. $(p_A, 0)$ dominates (resp. is dominated by) $(p_A, 1)$ for any $p_A < \frac{1}{2}$ (resp. $p_A > \frac{1}{2}$) regardless of Alice's signal and all actions taken in the game after her price update. But, now STRATEGIC Bob's best response depends not only directly on Alice's price-report p_A but also on his posterior belief about Charlie's vote / signal which is, in turn, based on his own signal s_B as well as his inference about Alice's signal s_A . Let us denote by $q_0^A(s_A)$ Alice's posterior probabilities that another agent received a signal 0, given her own signal s_A , and by $q_0^B(s_A, s_B)$ the same for Bob, given he knows not just s_B but also Alice's signal s_A .

To complete the description of a consistent belief structure for this signal structure, we assume that Bob infers $s_A = 0$ (resp. $s_A = 1$) whenever $p_A < \frac{1}{2}$ (resp. $p_A > \frac{1}{2}$). Then, the main deviation from the analysis in Section 3.4 is that Bob's thresholds p^L and p^H now become functions of $s_B \in \{0, 1\}$ so that each has two possible values, say $p_{s_B=0}^L, p_{s_B=1}^L$, and $p_{s_B=0}^H, p_{s_B=1}^H$ respectively. Hence, we can show that in any PBE, $p_A < \frac{1}{2}$ for $s_A = 0$ (i.e. $q_0^A > \frac{1}{2}$) and $p_A > \frac{1}{2}$ for $s_A = 1$ (i.e. $q_0^A < \frac{1}{2}$). Moreover, Alice's best response is, for a low enough π , to set p_A at one of the above thresholds, and, for a high enough π , at her expectation of the outcome, given her t. Consider, for example, a scenario in which the signal structure has $\rho_0 = 0.5$ and $\rho_\epsilon = 0.075$ and Alice receives signal $s_A = 0$. Calculations show that $q_0^A(s_A) = 0.8613$, and hence $p_{s_B=0}^L \approx 0.1694$, $p_{s_B=1}^L \approx 0.3257$, $p_{s_B=0}^H \approx 0.6743$, and $p_{s_B=1}^H \approx 0.8306$.

For $\pi = 0.2$, Figure 3.5(a) shows that Alice sets $p_A = p_{s_B=0}^L$ and $v_A = 0$ in equilibrium. Now, if $s_B = 0$, then STRATEGIC Bob "agrees" with Alice and votes $v_B = 0$ because p_A lies at his signal-dependent threshold $p_{s_B=0}^L$, but if $s_B = 1$, he votes $v_B = 1$ because p_A is now lower than his threshold $p_{s_B=0}^L$. i.e. in a "disagreement sub-interval". Moreover, the above result also suggests that for such a low- π equilibrium, Bob's vote is always identical to his signal, hence his signal is always revealed. For $\pi = 0.98$, Figure 3.5 (b) reveals that Alice's

expected payoff, given her knowledge of how the game unfolds after she trades, is maximized at $p_A = \frac{0 + ((1-\pi) \cdot 1 + \pi \cdot (1 - q_0^A(0)) + (1 - q_0^A(0)))}{3} \approx 0.982$, i.e. Alice's expected outcome, and this gives her PBE price-report; but now, p_A is less than both $p_{s_B=0}^L$ and $p_{s_B=1}^L$ so that, regardless of his signal, STRATEGIC Bob disagrees with Alice, and his signal cannot be inferred from his action (just as in the two-player game). In any case, STRATEGIC Bob moves the price to his posterior expectation of v , which is no longer in $\{0, 1\}$.




Figure 3.5: Alice's expected payoff for different values of her trading action p_A given her knowledge of the extensions of Lemmas 5 and 6 for the case of three outcome-deciders. For both plots, Alice has the same signal ($s_A = 0$) and all model parameters are equal, except Bob's non-participation probability π , as noted above.

Additional traders who do not affect the outcome: Agents with no control over the outcome who trade before Alice only matter in the level to which they move the price seen by Alice but, from Alice's perspective, this is equivalent to a general 'starting' price $p_0 \in (0, 1)$; if they all trade after Bob, Alice and Bob's equilibrium actions remain unchanged because, for an MSR, an agent's payoff depends on the actions of her predecessors and not on those of her successors, by design (as long as these successors are not outcome-deciders).

The game becomes more complex when there are intermediate traders between Alice and Bob. To test how the equilibrium strategies implied by Theorem 4 fare in such scenarios, we ran some simple simulations with an LMSR market maker. The market starts at $p_0 = 0.5$, and Alice moves first, but she is followed by a sequence of 10 “boundedly rational” myopic budget-limited traders, described in greater detail in the next paragraph; then Bob trades in the market (i.e. is of the STRATEGIC type) with probability $1 - \pi$. As before, the outcome is Alice and Bob’s average vote; the other traders do not vote.

Each intervening trader i obtains her private signal from the same source as Alice and Bob, and then draws her estimate w_i of the market outcome from a beta distribution: The mean of this distribution is equal to her posterior expectation of any other agent’s signal given her own realized signal, $m_i = \mathbb{E}[s_j | s_i] = \Pr(s_j = 1 | s_i)$ for any j , and the variance is $m_i(1 - m_i)/2$ or, in other words, the “pseudo-sample size” parameter of the beta distribution is 1. In addition, each trader has the same budget B , i.e. each of these agents trades in such a way that her loss (negative ex post payoff) never exceeds B . From the properties of the LMSR, it can be easily shown that such an agent i can move the market price from her observed value p_{i-1} to a maximum of $p_i^{\max} = (1 - \gamma \bar{p}_{i-1})$ and to a minimum of $p_i^{\min} = \gamma p_{i-1}$, where $\gamma = e^{-B/b}$, $\bar{p}_{i-1} = 1 - p_{i-1}$. This budget constraint captures the intuition that agents who do not influence an outcome and / or have poor knowledge thereof should be conservative in their trading decisions. The intervening trader i buys a quantity that moves the market price up to $\min(w_i, p_i^{\max})$ if $w_i > p_{i-1}$, sells a quantity that pushes the price down to $\max(w_i, p_i^{\min})$ if $w_i < p_{i-1}$, and does not trade otherwise.

As in Section 3.4, it is easy to show that Alice’s vote is revealed immediately after she trades (this is consistent with the empirical observation that insiders are often “big players”, and can be identified by other traders from their trading decisions). Thus, the problem is still easy for STRATEGIC Bob who can see his immediate predecessor’s report and infer Alice’s vote from her price update, and base his actions on this knowledge. But, how should Alice play? This is a strategically more complex game, and finding equilibrium strategies may be difficult. However, one possibility is for Alice to simply ignore the existence of the intermediate traders and use her strategy from Theorem 4 – we will call an Alice taking such an action STRATEGIC. While this is not necessarily an equilibrium (or *a priori* even a good) strategy, in our simulations, we observe that it significantly improves upon a simple alternative – declaring one’s private signal as the vote after updating the price to one’s

posterior expected average vote – which we shall designate being TRUTHFUL¹⁸. Thus, the model may well have predictive value even in this more complex setting.

Profits	STRATEGIC Bob	TRUTHFUL Bob
STRATEGIC Alice	0.303, 0.479	0.123, 0.265
TRUTHFUL Alice	0.280, 0.479	0.131, 0.265

Table 3.7: Alice and Bob’s expected profits for being STRATEGIC vs. TRUTHFUL, with 10 intervening traders for the signal structure described in Section 3.7. Results are averaged over 10^5 simulations. The first entry in each cell corresponds to Alice, the second to Bob. The table does not list all strategies available to the players, and does not depict the equilibrium; it shows that when *both* players employ Theorem 4 strategies, each achieves at least as much expected profit as they would for any strategy-pair in which at least one is TRUTHFUL. We observed similar results for other parameter combinations.

Table 3.7 shows a sample of our findings for a signal structure similar to that in Section 3.6 with $\rho_0 = 0.5$, $\rho_\epsilon = 0.4$, $\pi = 0.4$; the budget is $B = 0.5$. STRATEGIC Bob is as described in Section 3.3; the difference between TRUTHFUL Bob and HONEST Bob is that the former not only votes truthfully but also *participates* in **Stage 1**, updating the price to the realized outcome. Note that Bob’s profit in Table 3.7 appears to be independent of whether Alice is STRATEGIC or TRUTHFUL. This is due to two factors: Bob’s action (and hence his profit) depends only on whether Alice’s report p_A is above or below 0.5, indicating her vote, and for this information structure, p_A is always on the same side of 0.5 given Alice’s signal, irrespective of her strategy; his profit depends on his immediate predecessor’s report too, but with a sufficiently large number of intermediate traders, this price-report is also practically stable.

3.8 Incentivizing truthful trading and voting

Until this point, we have assumed that there is zero compensation for the agents associated with the outcome-deciding process. But, if agents derive some extra-market utility, known to the principal, from the the outcome itself or if the principal itself can offer additional payment

¹⁸For computing the posterior expected average vote, TRUTHFUL Alice just assumes that Bob will vote his private signal.

to each agent corresponding to her role in the outcome-deciding process, it is reasonable to ask whether the resulting (modified) two-stage game has an equilibrium where all $n \geq 2$ agents “vote” truthfully (and, of course, trade in a way that reflects their truthful vote).

We propose a design for a combined two-stage trading-voting mechanism (“voting” again being a metaphor for an agent’s outcome-affecting action) for the general case of n agents ($n \geq 2$) that disincentivizes manipulative behavior. The basic model is a natural extension of that described in Section 3.3 but this section is independent of the analysis in Section 3.4.

In the first stage of the game, the n agents trade in a prediction market in some pre-determined order, and then simultaneously choose their actions $v_i \in \Omega$ (not necessarily binary) in the second stage, the market liquidation value v being some function of $\{v_i\}_{i=1}^n$ (not necessarily the average). Each agent has zero non-participation probability in the market and trades exactly once. Thus, although the myopic assumption still holds, the scenario is adversarial in the sense that all outcome-deciders participate deterministically in the prediction market. The number of participants n as well as the belief structure on the underlying type τ and the agents’ private signals $\{s_i\}_{i=1}^n$ is common knowledge. Let us denote the agents’ price reports $\{p_i\}_{i=1}^n$, the starting market price being p_0 .

Here, we demonstrate our approach with a LMSR market with liquidity parameter b , the treatment for other scoring rules being similar. Although the worst-case loss in a traditional LMSR market is bounded, it is theoretically possible for an individual trader to earn an unbounded profit from it. To circumvent this issue, we place a fixed budget B on every agent (i.e. each agent trades in such a way that their loss can never exceed B).¹⁹ From the properties of LMSR, it can be easily shown that agent i can move the market price from their observed value p_{i-1} to a maximum of $p_{max} = (1 - \gamma \bar{p}_{i-1})$ and to a minimum of $p_{min} = \gamma p_{i-1}$, where $\gamma = e^{-B/b}$, $\bar{p}_{i-1} = 1 - p_{i-1}$. Thus, agent i ’s net payoff from the prediction market is

$$r_i^{PM}(p_i, p_{i-1}, v) = b \left[v \ln \left(\frac{1/p_{i-1} - 1}{1/p_i - 1} \right) - \ln \left(\frac{1 - p_{i-1}}{1 - p_i} \right) \right]. \quad (3.15)$$

¹⁹With the budget constraint, the market may lose some of its *expressiveness* (Abernethy et al., 2011), i.e. a trader may not be able to update the market price so as to coincide with their estimate of the liquidation value, but it is still *directionally expressive* in the sense that it is still rational for an agent to shift the price as close to her estimate as possible.

Simple algebra shows that for each i , the maximum and minimum possible values of the above under the budget constraint are

$$r_{i,max}^{PM} = b \ln \left(\max \left\{ \frac{1 - \gamma p_{i-1}}{1 - p_{i-1}}, \frac{1 - \gamma \bar{p}_{i-1}}{1 - \bar{p}_{i-1}} \right\} \right), \quad r_{i,min}^{PM} = -B = b \ln \gamma. \quad (3.16)$$

Let $\hat{l}_i(v_i)$ denote agent i 's posterior expectation of the liquidation value based on her decision to vote v_i (declared signal) and her inference from the common prior, her private signal s_i , and the prices p_0, p_1, \dots, p_{i-1} (assuming every other agent is Bayesian and truthful in both voting and trading), and $\hat{r}_i^{PM}(\tilde{p}, p_{i-1}, v_i)$ be her posterior expected net payoff from the market mechanism on updating the price from p_{i-1} to any reachable \tilde{p} . From the above linear dependence of r_i^{PM} on v , \hat{r}_i^{PM} is readily seen to be a simple linear function of $\hat{l}_i(v_i)$.

Our key idea is to introduce a compensation scheme for the voting mechanism such that the combined payoff from the market and the voting system when an agent is truthful in both stages (and believes that everyone else is going to be similarly truthful) exceeds the largest profit she can make by deviating from truth-telling. A promising technique for achieving this end is the *peer prediction* scheme introduced by Miller et al. (2005) proposed in the literature providing truth-telling incentives to experts in a traditional (non-market) means of information gathering like a survey or poll when the ground truth is never accessible to the principal. The following is a brief description thereof tailored to our setting.

We choose a reference participant $f(i)$ *a priori* for agent i such that i 's vote is stochastically relevant for that of $f(i)$ (the posterior is different for different realizations of the signal). At the end of voting, the transfer made to participant i by the center is a function of the posterior on participant $f(i)$'s vote $v_{f(i)}$ under the common prior, likelihood and agent i 's vote v_i , the function being a strictly proper scoring rule. Suppose we use a strictly proper scoring rule $R(\cdot)$ (not necessarily logarithmic) so that agent i 's peer-prediction score is given by $r_i^{PP}(v_{f(i)}, v_i) = \alpha_i R(g(v_{f(i)}|v_i))$, where $g(\cdot|\cdot)$ is said posterior. Thus, agent i 's expected peer-prediction score is $\hat{r}_i^{PP}(s_i, v_i) = \alpha_i \phi(s_i, v_i)$ where $\phi(s_i, v_i) = \sum_{v_{f(i)} \in \Omega} R(g(v_{f(i)}|v_i)) g(v_{f(i)}|s_i)$ and $g(\cdot|\cdot)$ is said posterior, $\alpha_i > 0, i = 1, 2, \dots, n$. Miller *et al* show that truthful voting is a strict Nash equilibrium for this mechanism. The constants α_i have no effect on the truth-telling incentives of the voting stage alone, and we show below how to tune these free parameters to ensure honest behavior. Note that while

the LMSR prediction market has a single parameter b that determines all its properties, the proposed combined mechanism has n additional parameters.

Now we can state our result: there is a way to set the parameters b and α_i (which are under the designer's control) that guarantees that the overall expected earnings $\hat{r}_i = (\hat{r}_i^{PM} + \hat{r}_i^{PP})$ for every i is uniquely maximized under the imposed constraints when their trading and voting are both consistent with their private information (in equilibrium). Let $\Delta\phi_m = \min_{s, s' \in \Omega, s' \neq s} [\phi(s, s) - \phi(s, s')]$, which is strictly positive by the incentive compatibility of the peer-prediction method and is a function of the signal structure.

Theorem 5. *Let there be n participants in the above mechanism, each with budget B . If we promise every agent i , when she arrives to trade, a peer- prediction payment with some α_i satisfying*

$$\alpha_i > \frac{b}{\Delta\phi_m} \ln \left(\max \left\{ \frac{1-\gamma p_{i-1}}{\gamma(1-p_{i-1})}, \frac{1-\gamma \bar{p}_{i-1}}{\gamma(1-\bar{p}_{i-1})} \right\} \right),$$

where all symbols have the meanings stated above, then there exists an ex interim Bayes-Nash equilibrium where each agent i announces $v_i = s_i$ in the outcome-deciding stage after having updated the market price as close as possible to her truthful expected liquidation value, i.e.

$$p_i = \begin{cases} p_{min} & \text{if } \hat{l}_i(s_i) < p_{min}, \\ \hat{l}_i(s_i) & \text{if } p_{min} \leq \hat{l}_i(s_i) \leq p_{max}, \\ p_{max} & \text{otherwise.} \end{cases}$$

Proof. Part I: For any voting choice $v_i = v'$, agent i 's expected market liquidation value assuming everyone else to be truthful is $\hat{l}_i(v')$, and since \hat{r}_i^{PP} is independent of p_i , from Equation (3.15), it follows that

$$\frac{\partial \hat{r}_i}{\partial p_i} = \frac{\partial \hat{r}_i^{PM}}{\partial p_i} = \frac{\hat{l}_i(v') - p_i}{p_i(1-p_i)} \geq 0 \iff p_i \leq \hat{l}_i(s_i).$$

Part II: For proving that when subsidies are set to yield the conditions on α_i above, it is in an agent's best interest to pick her honest vote, it suffices to show that for any possible signal values s, s' where $s' \neq s$, and any feasible prices $\tilde{p}, \tilde{p}', \tilde{p}''$, $\hat{r}_i(s_i = s, p_i = \tilde{p}', p_{i-1} = \tilde{p}, v_i = s) > \hat{r}_i(s_i = s, p_i = \tilde{p}'', p_{i-1} = \tilde{p}, v_i = s')$ which reduces to $\alpha_i[\phi(s, s) - \phi(s, s')] > \hat{r}_i^{PM}(\tilde{p}'', \tilde{p}, s') - \hat{r}_i^{PM}(\tilde{p}', \tilde{p}, s)$. The greatest lower bound on the L.H.S. is, by definition, $\alpha_i(\Delta\phi_m)$ where $\Delta\phi_m$

is a known constant and always strictly positive. An upper bound on the R.H.S. is obviously the range of all possible payoffs from the prediction market of agent i with budget B whose predecessor's price-report is p_{i-1} , i.e. $(r_{i,max}^{PM} - r_{i,min}^{PM})$, given by Equations (3.16). Thus, setting α_i to a value exceeding the (finite positive) bound specified in the theorem statement is a sufficient condition for the desired inequality to hold.

The two parts together complete the proof. \square

If $R(\cdot) \equiv \ln(\cdot)$, then the “raw” peer-prediction scores r_i^{PP} are always negative, so there is no incentive for voluntary participation. This problem can be solved, as in Miller et al. (2005), simply by subtracting from the raw score of agent i the constant $\alpha_i \min_{s,s' \in \Omega} (\ln(s|s'))$ which is a function of the prior and likelihood structures and independent of actual trader behavior. This ensures positive peer-prediction payments but also necessitates subsidization of the mechanism. For our budget-constrained LMSR market, $b \ln 2$ is a (perhaps loose) upper bound on its loss so the market subsidy is linear in b and independent of n . The amount of subsidy for the voting phase is proportional to $\sum_{i=1}^n \alpha_i$. A reasonable choice for α_i is $\alpha_i = \frac{\kappa b}{\Delta \phi_m} \ln \left(\max \left\{ \frac{1-\gamma p_{i-1}}{\gamma(1-p_{i-1})}, \frac{1-\gamma \bar{p}_{i-1}}{\gamma(1-\bar{p}_{i-1})} \right\} \right)$ where κ is a constant slightly greater than 1. It is straightforward to show that

$$\ln \left(\frac{2}{\gamma} - 1 \right) \leq \left(\frac{\Delta \phi_m}{\kappa b} \right) \alpha_i \leq \ln \left(2(1-\gamma) \left(\frac{1}{\gamma} \right)^{i+1} + \frac{1}{\gamma} \right) \forall i,$$

assuming that the starting market price is 0.5. Since $\gamma = e^{-B/b}$, it is clear that, for fixed b , α_i is $\Omega(B)$ and $O(iB)$ and, for fixed i, B , it is $\Theta(1)$. Hence the total peer prediction subsidy is linear in B , independent of b , and $\Omega(n)$ and $O(n^2)$.

3.9 Discussion

In this chapter, we have taken a significant step in exploring the crucial incentive issues that have the potential to derail the effectiveness of prediction markets for various forecasting tasks. We have introduced a new formal model for studying the incentives for and the impact of manipulation in prediction markets whose participants can affect the outcome by taking actions external to the market but there is some uncertainty about the market

participation of some outcome-deciders. We have characterized the equilibria of the induced game, discussed their properties, and outlined important extensions. Interesting avenues for future work include generalizing our results to markets with other price-setting mechanisms, richer signal structures, outcome functions other than the mean vote (such as non-linear and / or noisy functions of the agents' second-stage actions), and agents who also strategically pick the time-points at which they trade.

Chapter 4

Aggregating censored signals from non-strategic noisy agents

4.1 Introduction

In this chapter, we will abstract away from the task of providing truth-telling incentives, and concentrate on methods that can be used for combining signals obtained from agents as responses to *queries* designed by the principal, with the aim of getting as close to an unknown “ground truth” as possible; our agent model will assume a lack of strategization and consist only of the *signal structure* – the sampling distribution of agents’ signals given the ground truth – to account for the subjectivity in their inputs to the aggregator.

There exist many problem domains where the learner’s goal is to locate a certain target, given access only to a sequence of (potentially) *oracles* each of which provides a noisy binary response to the question of whether the target belongs to a sub-space (chosen by the learner) of its range of variation. Examples explored in the literature include dynamic pricing of goods and services (Harrison et al., 2012), object localization in images (Sznitman et al., 2013), and drug dosage discovery in Phase I clinical trials (Cheung and Elkind, 2010b).

Although the material in this chapter has general applicability, let us stay true to the spirit of information aggregation within a (prediction) market context presented in Chapters 2 and 3, and try to motivate the discussion by recounting the automated market maker for financial markets developed by Das and Magdon-Ismael (2009), henceforth referred to as MM. Recall that a market maker is a trading agent that places both buy and sell orders

within the same market (unlike buyers and sellers who respectively have demand and supply only), and readjusts its prices after every trade with some financial objective in mind, e.g. expected long-term profit maximization.

In the market model of Das and Magdon-Ismail (2009), the asset being traded attains a “true” (unknown) value at market inception that remains unchanged henceforth, and each trading agent acquires a noisy version of this value. At each time-step or episode, the MM publicly quotes an *ask* price and a (lower) *bid* price defined as the price at which it is willing to sell and buy one unit of the asset respectively in the episode. Exactly one agent interacts with the MM per episode and buys one unit, sells one unit, or does nothing depending on whether her valuation is higher than the ask price, lower than the bid price, or in between these quotes. Agents following such simple trading rules that are not “immediately” irrational but lack sophisticated optimization or learning components are often called *zero-intelligence traders* (Gode and Sunder, 1993), and have been used to illustrate the emergence of interesting aggregate-level properties (e.g. market efficiency) from individual properties (e.g. bounded rationality).

Since the MM sees which of the three above actions the agent took, the ask and bid prices, in addition to determining revenues, also serve as thresholds defining three (mutually exclusive and exhaustive) sub-intervals such that the MM knows which of these sub-intervals the latest trader’s valuation lies in. It can use this knowledge to adjust its quotes for the next episode with the ultimate aim of converging on the true asset value so that it stands to profit (in the long run) from the imperfectly informed traders. Das and Magdon-Ismail (2009) accomplish this task within a reinforcement learning setting by having the MM maintain a Gaussian belief state over the unknown value (which serves as the unchanging state of the world), use a moment-matching approximation to its Bayesian posterior after every agent interaction, and set bid-ask quotes as an action in its updated belief state; they show experimentally that this methodology has impressive price discovery properties. If we consider the single-threshold variant of this problem (where bid and ask price always coincide), the MM can be viewed as a learner or principal performing an aggregation of stochastic binary (thresholded) signals with its mean belief acting as the aggregate, albeit with a potential long-term profit-making aim unlike in Chapters 2 and 3. We will concern ourselves only with *what* the principal wishes to learn, and not *why* it wants to learn it.

We present an algorithm that starts with a Gaussian prior belief on a real-valued target, maintains a Gaussian belief at all times (after an initial transient phase; see below for details) by applying a moment-matching approximation to the true (complicated and non-Gaussian) posterior, and sets its threshold for querying each agent in a sequence at the mean of its current belief distribution. We show that it unconditionally converges to the target with high probability, and the asymptotic rate of convergence is near-optimal with respect to many problem parameter, optimality being defined with respect to an exact Bayesian inferential procedure that observes agents’ real-valued (unthresholded) signals.

4.2 Related work

The literature on learning with thresholded signals or binarized observations is scattered across various lines of academic research. For example, in online dynamic pricing, a seller wishes to determine the demand curve. She sets a price for a good and observes whether or not the arriving buyer chooses to purchase at that price (Harrison et al., 2010). In drug dosage discovery, the goal is typically to estimate the maximum dosage level that causes toxicity with less than some target probability (this is typically the focus of Phase I clinical trials) (Cheung and Elkind, 2010a). Threshold queries are also used in image or face localization, where classifiers are used as subroutines to determine whether or not a face or letter or character appears in the query region of some image (Sznitman and Jedynak, 2010).

Most contributions in this vein have focused on noise of a particular form: Nature generates the correct answer, but it is then sent through a noisy transmission channel (Jedynak et al., 2011). Thus, the probability of seeing the wrong signal is constant, independent of the point of measurement (the particular threshold set by the learner). Several papers have focused on proving the asymptotic optimality of policies that measure either at or around the median (Horstein, 1963; Burnashev and Zigangirov, 1974; Castro and Nowak, 2008). More recent work shows that measuring at the median is sequentially optimal for entropy reduction in the case of symmetric noise (Waeber et al., 2011). In a different vein, Karp and Kleinberg (2007) consider noisy binary search: in this problem, a finite sequence of biased coins, ordered in increasing probability of a “heads” outcome, has to be searched for the last element with a probability of heads lower than a specified target value.

The bisection problem itself can also be thought of as a version of the classic problem of stochastic root finding (Robbins and Monro, 1951), where the learner is trying to learn the root of a real-valued, decreasing function f . The model is that the learner sequentially queries at points $\theta_1, \theta_2, \dots, \theta_n$, and receives observations of $f(\theta_1), f(\theta_2), \dots, f(\theta_n)$ after addition of noise (e.g. zero-mean Gaussian noise). A natural extension to binary signals is to assume that the learner observes whether or not the corrupted signal is above or below zero. This directly corresponds to a noisy binary signal indicating whether the threshold is smaller than or larger than the root. In this case, the noise model is heavily dependent on how close the threshold is set to the target root. When the threshold is near the target, the probability of seeing the wrong signal is significantly higher and no longer bounded away from $\frac{1}{2}$.

4.3 The learning problem

We set up a formal model for studying this sequential learning problem as follows: The aim is to determine, within error tolerance ε , a fixed but unknown “target” value $V \in \mathbb{R}$, by querying a sequence of agents (or oracles) at episodes $t \in \{1, 2, \dots\}$; agent t has an idiosyncratic estimate w_t of the target due to independently and identically distributed additive (zero-mean Gaussian) noise with a fixed, known variance σ_z^2 , i.e.

$$w_t = V + z_t, \quad \text{where} \quad z_t \sim_{\text{i.i.d.}} \mathcal{N}(0, \sigma_z) \quad \forall t = 1, 2, \dots,$$

where $\mathcal{N}(\mu, \sigma)$ denotes a Gaussian distribution with location and scale parameters μ and σ respectively. If the learner had access to these raw samples w_1, w_2, \dots , the learning problem could be solved by a classic Bayesian approach: one could start with a Gaussian prior (the conjugate prior for this scenario) over the possible values of the unknown target $v \in \mathbb{R}$, i.e. $p_0(v) = \mathcal{N}(\mu_0, \sigma_0)$ for some $\mu_0 \in \mathbb{R}$ and $\sigma_0 > 0$, and hence perform a standard inference on this sequence of observations — it is well-known that the learner’s belief distribution $p_t(v)$ after every observation x_t will be Gaussian with mean μ_t and standard deviation σ_t given by the following two-dimensional updates:

$$\mu_{t+1} = \mu_t + \frac{\rho_t^2}{1 + \rho_t^2}(w_t - \mu_t), \quad \rho_{t+1}^2 = \frac{\rho_t^2}{1 + \rho_t^2}, \quad \text{where} \quad \rho_t = \sigma_t / \sigma_z, \quad \forall t = 1, 2, \dots \quad (4.1)$$

The learner returns μ_t at every t as her current estimate of V . Moreover, note that, the higher the value of the ratio ρ_t which quantifies the learner’s uncertainty relative to that of the noisily informed crowd, the larger the *step-size* $|\mu_{t+1} - \mu_t|$ for any given μ_t and w_t , i.e. as long as this Bayesian learner has a higher relative uncertainty, she takes a larger step “in the right direction” in expectation, since $\mathbb{E}_{p_t}[\mu_{t+1} - \mu_t] = \frac{\rho_t^2}{1+\rho_t^2}(V - \mu_t)$ where $\mathbb{E}_{p_t}[\cdot]$ denotes expectation with respect to the learner’s time- t belief.

The catch in our model is that the learner is constrained to ask agent t only if w_t is above or below some *threshold* θ_t of the learner’s choice, i.e. she only observes the binary signal

$$x_t = \mathbb{I}(w_t \geq \theta_t) = \text{sign}(V + z_t - \theta_t) \quad \forall t.$$

Such an observation, where we only know whether or not a data point is above a specified threshold, is said to be *censored*. With a Gaussian prior, our posterior distribution with this type of observation is not Gaussian. Our problem can thus be seen as a censored variant of stochastic root finding (Robbins and Monro, 1951) where a learner, trying to learn the root of a real-valued decreasing function $f(\cdot)$, sequentially queries at points $\theta_1, \theta_2, \dots, \theta_n$ to receive observations of $f(\theta_1), f(\theta_2), \dots, f(\theta_n)$ after addition of noise z_t ; in our scenario, the learner just sees $\mathbb{I}(f(\theta_t) + z_t \geq 0)$ where $f(\theta) = V - \theta$. Moreover, our probability of obtaining a “wrong” signal at t is

$$\Pr(w_t \geq \theta_t | V < \theta_t) \mathbb{I}(V < \theta_t) + \Pr(w_t < \theta_t | V \geq \theta_t) \mathbb{I}(V \geq \theta_t) = 1 - \Phi(\Delta_t / \sigma_z),$$

on simplification, where $\Phi(\cdot)$ is the standard normal cumulative distribution function, and $\Delta_t \triangleq |\theta_t - V|$. Evidently, the above probability increases towards $\frac{1}{2}$ as Δ_t approaches 0 or, in other words, as the threshold gets closer to the target. This is a significant deviation from earlier models for learning from noisy binary signals which assume that the correct response to every binary query at $t = 1, 2, \dots$ reaches the learner after passing through a transmission channel whose noise characteristics are independent of the query or, equivalently, the time-step (Jedynak et al., 2012).

In this setting, the learner is faced with the two-step task, at every t :

- (i) Choosing a *query threshold* θ_t given its current belief about the target (which implies that our problem can be subsumed under *active learning* (Castro and Nowak, 2008)).

- (ii) Updating its belief, and hence its estimate of the target, on receiving the response x_t .

Note that the learner is free to pick each θ_t on the basis of her prior belief, the entire history $\{\theta_1, x_1, \theta_2, x_2, \dots, \theta_{t-1}, x_{t-1}\}$, and a learning objective of her choice, e.g. minimizing her uncertainty at the end of a pre-specified time horizon. In principle, a Bayesian learner could set up a dynamic program to solve its two-part problem, where the state is an entire probability (belief) distribution and in every state, the action of computing a threshold is taken. But instead of dealing with an infinite-dimensional state space, we will aim to come up with a heuristic strategy with simple threshold-setting rules and state updates comparable to (4.1) but having provable performance guarantees. We provide in detail the underlying rationale for our algorithm in Section 4.4.

4.4 The intuition for our solution approach

The Bayesian Setting. At time $t = 0, 1, 2, \dots$, assume a (prior) distribution for V , which we denote $p_t(v)$. After observation x_t , the Bayesian update to the distribution is given by

$$p_{t+1}(v) = \Phi(x_t(v - \theta_t)/\sigma_z) p_t(v) / A_t, \quad -\infty < v < \infty,$$

where $A_t = \int_{-\infty}^{\infty} dv p_t(v) \Phi(x_t(v - \mu_t)/\sigma_z)$, and $\Phi(\cdot)$ is the standard normal cumulative distribution function (CDF). Assuming p_t is correct, which may not be the case, p_{t+1} incorporates all the new information from x_t . At time t , the best estimate of V is given by the mean of the distribution, which we define as

$$\mu_t = \mathbb{E}_{p_t}[V] = \int_{-\infty}^{\infty} dv v p_t(v).$$

If the learner had to output an estimate for V at time t , the expected cost is the variance,

$$\sigma_t^2 = \text{Var}_{p_t}[V] = \int_{-\infty}^{\infty} dv v^2 p_t(v) - \mu_t^2.$$

The Starting Prior. As with all Bayesian inference algorithms, we need to start with some prior $p_0(v)$. In our context, the noise in the signals is based on a normally distributed random

variable $z_t \sim \mathcal{N}(0, \sigma_z)$. One way to quantify the uncertainty in the learner's prior is through the learner's initial variance, which we define as $\sigma_0^2 = \rho_0^2 \sigma_z^2$ (and, in general, $\sigma_t^2 = \rho_t^2 \sigma_z^2$). Given the learner's initial variance, in accordance with the principal of maximum entropy, we adopt the least informative prior. This happens to be the normal distribution, so we assume that $p_0(v) = \mathcal{N}(0, \rho_0 \sigma_z)$ (we can always assume $\mu_0 = 0$ by translating V).

A few words about the dimensionless parameter ρ_0 , an important measure of the harshness of the learning environment, are in order. The harshest environment has $\rho_0 \rightarrow 0$, where, if the prior is correct, the learner is very sure of her belief about V , but the signals are essentially random signs, and so it is hard to make any progress in learning from the observations. This is the regime we are interested in because (i) it is the hard interesting problem; and, (ii) any inference based algorithm will eventually get more and more certain as it receives more observations, which means that $\rho_t \rightarrow 0$. Thus, if it is to succeed, any algorithm has to be able to make good progress in this harsh regime. In fact, any reasonable heuristic (and we present one) can learn when the observations are relatively noiseless; the ultimate performance of an algorithm is dependent on its behavior in this $\rho \rightarrow 0$ regime. From now on, we set the scale of the problem by choosing $\sigma_z^2 = 1$ (which is without loss of generality; the scale can always be added back through powers of σ_z using dimensional arguments).

Myopic Thresholds. Within this Bayesian setting, we perform task (i) of our problem using the the simple myopic strategy of setting $\theta_t = \mu_t$, the expectation of the learner's current belief p_t , at each t . This may or may not be sequentially optimal, but as we show in Section 4.6, it is sufficient to obtain near-optimal asymptotic performance. In the multi-threshold generalization of this problem, the selection of thresholds becomes non-trivial and is dictated by further assumptions / constraints. However, our state update procedure that we introduce in Section 4.5 can be extended to such situations and, since multiple thresholds provide strictly more information, the performance cannot be worse than that of the single-threshold algorithm which itself is near-optimal.

It is true that, to implement this myopic single-threshold strategy, one only needs to compute μ_t at every time step, and perform the Bayesian update after observing x_t . Unfortunately, this computation requires the calculation of two integrals (one to compute A_t , and one to compute the expectation) which are not analytically tractable, even for elementary starting

priors $p_0(v)$. The natural alternative is to use numerical integration. However, numerical integration leads to issues of numerical stability and efficiency. To compute p_t , one needs to store the entire history of θ_t , x_t , A_t , which is $O(t)$, and then the running time to set θ_t , if we compute the integrals numerically with N quadrature points, is $O(Nt)$. Together with the numerical instability, this rapidly becomes computationally infeasible. In addition, algorithmic issues can arise in selecting appropriate finite bounds for the integration domain. As such, it makes sense to investigate whether approximate belief updates exist that enable easier computation of thresholds and yet come with good performance guarantees. This leads us to Section 4.4.1

4.4.1 Non-Parametric Histograms

Once the choice of thresholds has been made, the main challenge is to efficiently update the prior (task (ii)). Consider the method, which we call **NonParam** henceforth, that uses a non-parametric finite-support distribution to approximate the belief state after each episode. Let $v_1 < \dots < v_N$ be N possible values for V , the prior distribution $p_0(v)$ being given by the probability-vector $(p_0(v_1), \dots, p_0(v_N))$, a truncated and renormalized discretization of our actual Gaussian prior; the Bayesian update after observing every x_t is then given by

$$p_{t+1}(v_i) = \frac{1}{A_t} \Phi(x_t(v_i - \theta_t)) p_t(v_i), \quad i = 1, 2, \dots, N, \quad A_t = \sum_{i=1}^N \Phi(x_t(v_i - \theta_t)) p_t(v_i).$$

The computation of the threshold / expectation in each episode as a finite sum $\theta_t = \mu_t = \sum_{i=1}^N v_i p_t(v_i)$ takes $O(N)$ time; to converge to within ε of V , the resolution in the finite prior should be $O(\varepsilon)$, i.e. $N = \Omega(\varepsilon^{-1})$, making this a computationally intense procedure. Another problem with this approach is that one must commit to a range for V , introducing additional assumptions, and leading to serious problems when V is outside, or in the tail of, the range. In spite of being impracticable, this method offers insights into the behavior of the Bayesian update, aiding the design of our proposed algorithm, and also serves as a (near-exact) benchmark for evaluating our accuracy. Starting from a Gaussian prior, using our myopic thresholds, and running these non-parametric histogram-based updates, we illustrate how the posterior evolves in Figure 4.1.




Figure 4.1: Evolution of $p_t(v)$ using non-parametric histograms and Bayesian updates; $p_0(v)$ is $\mathcal{N}(0, 10)$, $\sigma_\epsilon = 1$, $V \approx 9.45$. Typical evolution consists of 3 main phases. (i) **Truncation** ($\rho_t = \sigma_t/\sigma_\epsilon \gg 1$): All observations are in the same direction (here $x_t = +1$) and almost noiseless, so the Bayesian update results in a skewed distribution. Shown for comparison is a truncated (on one side only) and renormalized Gaussian distribution (using our heuristic), which approximates this phase better than Gaussian. (ii) **Collapse**: When the first observation in the opposite direction arrives (here $x_t = -1$), the distribution collapses to something more symmetric, although not quite normal. Shown for comparison is the entropy-matched normal with the same mean. (iii) **Convergence**: ρ_t is typically small and μ_t is close to V , i.e. the algorithm has (probabilistically) bracketed V . From then on, nearly independent observations which are close to random signs cause the distribution to rapidly converge to normal, as would be expected with truly independent observations.

4.5 The algorithm

We can now delineate the operation of our algorithm for belief updates (or, equivalently, learner's state transitions), illustrated in Figure 4.2.

The state of the learner at time t is completely described in terms of four parameters, (l_t, r_t, m_t, s_t) , that describe its current belief distribution, which can take on two forms: either Gaussian, or truncated Gaussian. The support of the distribution is given by (l_t, r_t) . In all cases, either the left bound $l_t = -\infty$ or the right bound $r_t = \infty$. The location and shape of the distribution are determined by m_t, s_t^2 , the mean and the variance of the underlying Gaussian. Thus the learner's belief distribution is a rescaled normal distribution

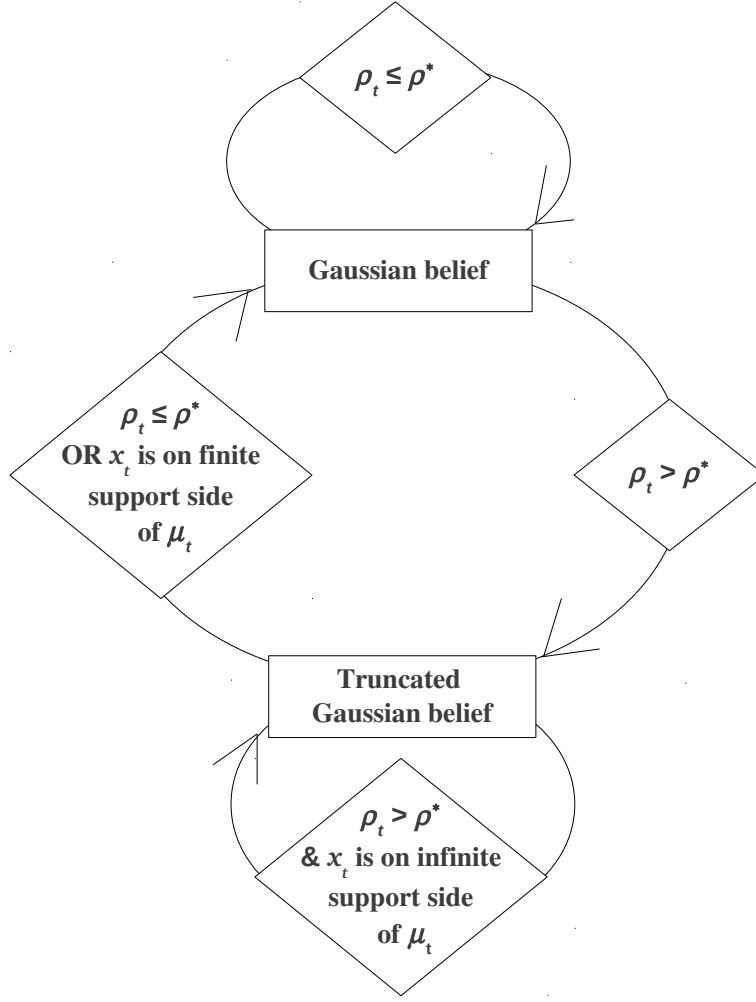


Figure 4.2: Learner's state transitions.

on the support (l_t, r_t) .

$$p_t(v) = \begin{cases} \frac{N\left(\frac{v-m_t}{s_t}\right)}{s_t \left(\Phi\left(\frac{r_t-m_t}{s_t}\right) - \Phi\left(\frac{l_t-m_t}{s_t}\right) \right)} & v \in (l_t, r_t), \\ 0 & \text{otherwise,} \end{cases}$$

where $N(\cdot)$ is the standard normal probability density function (PDF). The initial prior is normal with mean $m_0 = \mu_0 = 0$ and variance $s_0^2 = \sigma_0^2$ (equal to ρ_0^2 for $\sigma_z = 1$), which is described by the state $(-\infty, \infty, 0, \rho_0)$. These approximations to the true Bayesian evolution

allow us to compute the mean belief μ_t , identical to the threshold θ_t , by a simple known formula at every t .

Additionally, the algorithm has a *switch-over* parameter ρ^* such that, roughly speaking, a value of the relative uncertainty ρ_t , defined in (4.1), higher than ρ^* is taken to suggest that we are in a situation where the truncation heuristic (Figure 4.1 Panel (i)) is most reasonable. We use $\rho^* \geq 1$ since small ρ_t values ($\ll 1$) constitute the “challenging” regime.

A high level description of our method is given in Algorithm 1, and the general idea is as follows. We start with Gaussian belief but since initial observations are likely to be all in the same direction, we maintain a truncated normal distribution (as in Figure 4.1 Panel (i)). Upon collapse (i.e. when the learner receives the first observation in the opposite direction, as in Figure 4.1 Panel (ii)), we revert back to Gaussian, using entropy matching to set its parameters. Though the Gaussian is not very accurate at collapse, this is only a transient phase; as convergence occurs (as in Figure 4.1 Panel (iii)), the Gaussian becomes a better and better approximation, so we remain in the Gaussian world, using moment matching to update the parameters. The algebraic updates are provided below.

Algorithm 1 The Learning Algorithm

```

Initialize  $l_0 = -\infty, r_0 = \infty, m_0 = \mu_0, s_0 = \sigma_0$ .
for  $t = 0, 1, 2, \dots$  do
    Set threshold at  $\mu_t$ ;
    Receive noisy thresholded signal  $x_t$ ;
    Update  $l_t, r_t, m_t, s_t$ ;
    Compute  $\mu_{t+1}, \rho_{t+1} = \sigma_{t+1}/\sigma_z$ ;
end for
```

Approximate Gaussian Inference ($\rho_t \leq \rho^*$). As in Das and Magdon-Ismail (2008), in transitioning from Gaussian to Gaussian, we can compute the mean and variance of the true posterior, and we approximate this with the Gaussian that has the same mean and variance. So we perform *approximate moment matching inference* in this case. Das and Magdon-Ismail (2008) derive exactly such moment matching equations for two thresholds,

which we can directly specialize to the single threshold case:

$$\mu_{t+1} = \mu_t + x_t \frac{\left(\sigma_z \sqrt{2/\pi}\right) \rho_t^2}{\sqrt{1 + \rho_t^2}}; \quad (4.2)$$

$$\rho_{t+1}^2 = \rho_t^2 \left[\frac{1 + \rho_t^2 (1 - 2/\pi)}{1 + \rho_t^2} \right]. \quad (4.3)$$

Truncation ($\rho_t > \rho^*$). When ρ_t is large, we approximate the inference by truncating (as in Figure 4.1) as long as the signal is consistent with the truncation. The state updates are:

$$\begin{aligned} (l_t, \infty, m_t, s_t) &\xrightarrow{(\theta_t, x_t=+1)} (\theta_t - 2\sigma_z, \infty, m_t, s_t); \\ (-\infty, r_t, m_t, s_t) &\xrightarrow{(\theta_t, x_t=-1)} (-\infty, \theta_t + 2\sigma_z, m_t, s_t). \end{aligned}$$

Collapse. No matter what ρ_t is, if the signal is inconsistent with the truncated Gaussian, then we collapse back to Gaussian (see Figure 4.1). Unfortunately, updating to a Gaussian using moment matching would take the distribution with finite support and collapse to a distribution with infinite support and the same variance, typically producing a Gaussian that is too localized although there can be a lot of uncertainty in the learner's posterior. So a better way to capture this uncertainty is by matching the entropy. We call this approximate inference by entropy matching. To make the entropy matching analytically tractable, we first doubly truncate the distribution (as in regular truncation), compute the mean and entropy of the resulting distribution, and then collapse to the Gaussian with this mean and entropy. For finite l_t and r_t , the state updates are:

$$\begin{aligned} (l_t, \infty, m_t, s_t) &\xrightarrow{(\theta_t, x_t=-1)} (-\infty, \infty, \mu_{t+1}, \rho_{t+1}); \\ (-\infty, r_t, m_t, s_t) &\xrightarrow{(\theta_t, x_t=+1)} (-\infty, \infty, \mu_{t+1}, \rho_{t+1}). \end{aligned}$$

We abuse notation above, in that the updates are not the same in both cases. In the top case ($x_t = -1$), set $l = l_t$ and $r = \mu_t + 2\sigma_z$; in the bottom case ($x_t = +1$), set $l = \mu_t - 2\sigma_z$ and $r = r_t$. Then, a tedious but straightforward computation of the mean and entropy of the resulting rescaled doubly truncated Gaussian with support (l, r) and parameters (m_t, s_t) ,

followed by entropy matching gives:

$$\begin{aligned}\mu_{t+1} &= m_{t+1} + s_{t+1} \left[\frac{N(l') - N(r')}{\Phi(r') - \Phi(l')} \right], \\ \sigma_{t+1}^2 &= s_{t+1}^2 \left[(\Phi(r') - \Phi(l'))^2 e^{\frac{l'N(l') - r'N(r')}{\Phi(r') - \Phi(l')}} \right].\end{aligned}$$

where $l' = (l - m_t) / s_t$, $r' = (r - m_t) / s_t$.

4.6 Analysis

Note that exact Bayesian inferential procedure, described in (4.1), that we can perform for the scenario mentioned in Section 4.3 when the learner has access to raw valuations (with a Gaussian starting prior and Gaussian observations) serves as a “gold standard” for evaluating our algorithm in Section 4.5 since one is certainly getting more information from the unthresholded signals and so should be able to do better. Hence, we first analyze the raw-valuation algorithm, and then compare the asymptotic performance of our method that works with thresholded signals to it.

The proofs of all results furnished in this chapter are highly technical and relegated to Appendix C.

4.6.1 Bayesian inference with uncensored signals

Since $w_t = V + z_t$, we can rewrite state update equations (4.1) as

$$\mu_{t+1} = \frac{\mu_t + \rho_t^2(V + z_t)}{1 + \rho_t^2} \quad \text{and} \quad \rho_{t+1}^2 = \frac{\rho_t^2}{1 + \rho_t^2}.$$

Assuming that $\mu_0 = 0$, and that $\{z_t\}_{t \geq 0}$ are i.i.d. $\mathcal{N}(0, 1)$, *i.e.* $\sigma_z = 1$, we can unfold the above recursion to show that μ_t has a Gaussian distribution with

$$\mathbb{E}[\mu_t] = V - \frac{V}{1 + t\rho_0^2} \quad \text{and} \quad \text{Var}[\mu_t] = \rho_0^2 \frac{t\rho_0^2}{(1 + t\rho_0^2)^2}. \quad (4.4)$$

Theorem 6 below gives the convergence of μ_t both in expectation and with high probability. Fix an error tolerance $\varepsilon > 0$. The dependence of the expected value on t immediately implies a lower bound on the time after which the expectation of μ_t (which is our output estimate \hat{V}) is within ε of V . Further, the distribution for μ_t tells us that if we fix a small confidence parameter δ , $0 < \delta \ll 1$, and define $\zeta = -\Phi^{-1}(\delta)$, then (for $V > 0$) with probability at least δ , $\mu_t \leq \mathbb{E}[\mu_t] - \zeta \sqrt{\text{Var}[\mu_t]}$, which allows us to get a lower bound on t if we want high probability convergence.

Theorem 6. Fix $\varepsilon < \frac{|V|}{2}$, $\delta \leq \Phi(-1)$.

(i) For $t > |V|/\varepsilon\rho_0^2$, $|V| - |\mathbb{E}[\mu_t]| < \varepsilon$.

(ii) For $t > \max\left\{\frac{2|V|}{\varepsilon\rho_0^2}, \frac{4\zeta^2}{\varepsilon^2}\right\}$, where $\zeta = -\Phi^{-1}(\delta) \geq 1$,

$\Pr[\mu_t > V - \varepsilon] > 1 - \delta$ if $V > 0$; $\Pr[\mu_t < V + \varepsilon] > 1 - \delta$ if $V < 0$.

Proof. Here, we will prove the result for the case $\{V > 0\}$. Note that for $V > 0$, we must have $\frac{V}{1+t\rho_0^2} < V$ since $t, \rho_0^2 > 0$, so that $\mathbb{E}[\mu_t] > 0 \forall t \geq 1$, from result (4.4).

(i) The L.H.S. of the required inequality is

$$\begin{aligned} V - \mathbb{E}[\mu_t] &= \frac{V}{1+t\rho_0^2}, \quad \text{from result (4.4),} \\ &< \frac{V}{1+\frac{V}{\varepsilon}} \quad \text{for } t > V/\varepsilon\rho_0^2, \\ &= \frac{\varepsilon}{\frac{\varepsilon}{V} + 1} \\ &< \varepsilon \equiv \text{R.H.S.,} \end{aligned}$$

since $\frac{\varepsilon}{V} > 1$.

(ii) From the definition of the standard normal cumulative distribution function,

$$\begin{aligned}
\Pr[\mu_t < V - \varepsilon] &= \Phi\left(\frac{V - \varepsilon - \mathbb{E}[\mu_t]}{\sqrt{\text{Var}[\mu_t]}}\right), \\
\text{where } \frac{V - \varepsilon - \mathbb{E}[\mu_t]}{\sqrt{\text{Var}[\mu_t]}} &= \frac{V - \varepsilon(1 + t\rho_0^2)}{\rho_0^2\sqrt{t}}, \quad \text{from result (4.4), after simplification,} \\
&< \frac{V - \varepsilon t\rho_0^2}{\rho_0^2\sqrt{t}}, \quad \text{since } \varepsilon > 0, \\
&= \frac{V}{\rho_0^2\sqrt{t}} - \varepsilon\sqrt{t} \\
&< \frac{1}{\rho_0^2\sqrt{t}} \cdot \frac{\varepsilon t\rho_0^2}{2} - \varepsilon\sqrt{t}, \quad \text{for } t > \frac{2V}{\varepsilon\rho_0^2}, \\
&= \frac{\varepsilon\sqrt{t}}{2} - \varepsilon\sqrt{t} \\
&= -\frac{\varepsilon\sqrt{t}}{2} \\
&< -\zeta, \quad \text{for } t > \frac{4\zeta^2}{\varepsilon^2}, \\
&= \Phi^{-1}(\delta) \\
\implies \Pr[\mu_t < V - \varepsilon] &< \delta,
\end{aligned}$$

from the strict increasing monotonicity of $\Phi(\cdot)$. Now, the L.H.S. of the required inequality is

$$\Pr[\mu_t > V - \varepsilon] = 1 - \Pr[\mu_t < V - \varepsilon] > 1 - \delta \equiv \text{R.H.S.}$$

The proof for $\{V < 0\}$ is analogous. □

Part (i) of the theorem says that to get convergence in expectation, $O(V/\varepsilon\rho_0^2)$ time is needed. Note that, from the well-known *Gaussian tail inequality* (see e.g. Boucheron et al. (2004)), it is easy to show that $\zeta = -\Phi^{-1}(\delta) = O(\sqrt{\ln(1/\delta)})$; so, if one wants convergence with probability at least $1 - \delta$, then part (ii) of the theorem makes it clear that $O(V/\varepsilon\rho_0^2 + \ln \frac{1}{\delta}/\varepsilon^2)$ time is needed. These bounds will be useful in showing that our algorithm is nearly as good as exact inference with non-thresholded observations.

4.6.2 Approximate Bayesian inference with censored signals

The algorithm in Section 4.5 has two basic phases. The first is if ρ_t is large, in which case the algorithm is a heuristic that truncates the distribution until collapse into the Gaussian world, at which point the process (truncation→collapse) repeats until ρ_t gets below ρ^* . We do not go into the details of the dynamics for $\rho_t > \rho^*$ because in this relatively noiseless regime, many heuristics can “localize” the posterior quickly. The interesting regime is $\rho_t \rightarrow 0$, when the signals start to get noisy. In this regime, our algorithm will always be doing approximate Gaussian inference (since ρ_t is decreasing), updating according to Equations (4.2) and (4.3). Once in this regime, we essentially show that our algorithm is near-optimal by proving that μ_t converges quickly to V in expectation, and it also does so with high probability. For asymptotic results, we have in mind that $\rho_0 \rightarrow 0$. Proposition 3 below mirrors part (i) of Theorem 6 above, and speaks to the speed of convergence to the target in expectation for the case of censored signals; the (heuristic) proof is provided in Appendix C.

Proposition 3. *There exist absolute positive constants $C > 0$ and k , $1 \leq k < \pi\sqrt{2} \approx 4.44$ such that, if $t > C/(\rho_0^2 \varepsilon^k)$, then $|V| - |\mathbb{E}[\mu_t]| < \varepsilon$.*

Recall that the expectation is with respect to $p_0(v)$ and the i.i.d. $z_t \sim \mathcal{N}(0, 1)$, hence this result is relevant when the prior is correct. From Theorem 6, the best we could hope for, even with non-thresholded signals, for the expectation to get within ε of V is $t = O(1/\rho_0^2 \varepsilon)$. Thus, our dependence on ρ_0^2 is optimal. The theorem gives polynomial convergence in $1/\varepsilon$ but, in practice, k is almost 1, which is near-optimal asymptotic convergence, as illustrated with an example in Figure 4.3.

Our second result (the proof, again, is in Appendix C) demonstrates unconditional convergence with high probability, regardless of whether our prior $p_0(v)$ is correct. For simplicity, we assume without loss of generality that $V > 0$. Note that μ_t follows a stochastic process. We ask how long we have to wait before, with high probability, μ_t will have crossed $V - \varepsilon$. This analysis is sufficient to convey the main point of the convergence, because once you cross this barrier, the stochastic process has an attractor at V , and so will stay in this region. The tough part is getting to this region.




Figure 4.3: A log-log plot showing the time taken by $\mathbb{E}[\mu_t]$ (approximated by the average computed over 10^6 simulations) to converge to $V - \varepsilon$ when $V = 3$, $\mu_0 = 0$, $\rho_0 = 0.5$. Evidently, the convergence time approaches $O(1/\varepsilon)$ as ε becomes small, which gives us empirical evidence that it is near-optimal.

Theorem 7. Fix $0 < \delta < 1$, $0 < \varepsilon < V$, $0 < \rho_0 \leq 1$, and define $\Delta = V - \varepsilon$. There is an absolute constant $C > 0$ such that if $t > T = e^{C(\ln(1/\delta) + \Delta)/\varepsilon} / \rho_0^2$, then with probability at least $1 - \delta$, $\max_{i \leq t} \mu_i > V - \varepsilon$.

This constant T gives us an upper bound on the time at which μ_t first crosses $V - \varepsilon$. In the practical setting where the prior is ill-specified, V is very large and ε is usually specified as a percentage of V . Then, the exponent is some constant and the dependence on $1/\rho_0$ is what we are interested in. Comparing with part (ii) of Theorem 6, we see that our algorithm is asymptotically optimal with respect to $1/\rho_0$.

4.7 Experimental results

We perform experiments to evaluate the practical performance of our algorithm. Our simulations measure convergence time as the time taken by μ_t to enter the region $[V - \varepsilon, V + \varepsilon]$ for the first time. We are interested in the dependence of the convergence time on ρ_0 and ε .




Figure 4.4: Plot of average correct-prior convergence time *vs* ρ_0 , logarithmic along the vertical axis.

First, we compare the non-parametric algorithm (**NonParam**) to exact inference on non-thresholded signals, and show that noisy binary signals are almost as informative as the unthresholded signals. This is already surprising. Assuming that the prior is correct, we set the support of the non-parametric histogram to $[-6\rho_0, 6\rho_0]$ and use 1,000 histogram bins. We generate V randomly according to p_0 and any value of V outside this finite support is discarded. For our algorithm, we set the switch-over parameter $\rho^* = 2.5$. The number of steps taken by each algorithm to converge to the region $[V - \varepsilon, V + \varepsilon]$, averaged over 10^8 runs for **NonParam** and 10^9 runs for each other algorithm, is reported in Figure 4.4.

In our second set of simulations, we fix σ_0 at 0.5 and vary ε . To ensure adequate resolution for **NonParam**, we use $24\rho_0/\varepsilon$ bins (giving a resolution of $\varepsilon/2$). The number of steps to convergence is presented in Figure 4.5. The average is over 10^5 runs for **NonParam** (owing to computational burden) and 10^7 runs for each other algorithm. It is clear from the figures that not only is learning from noisy binary thresholds feasible in this Bayesian model, but can be performed almost as well as learning from non-thresholded signals, in accordance with the theory.




Figure 4.5: Plot of average correct-prior convergence time *vs* ε , logarithmic along the vertical axis.

4.8 Discussion

This contibution can be viewed as constructive proof of the claim that it is possible to learn a real-valued target from noisy valuations binarized by a threshold, asymptotically almost

as efficiently as if the actual valuations were accessible to the learner. This means that thresholding does not significantly impede our ability to learn!

These results also provide theoretical underpinnings to the Bayesian market making algorithm, introduced by Das and Magdon-Ismail (2009), and employed (with modifications to make it more practical) by Brahma et al. (2012) and Chakraborty et al. (2013).

An interesting direction for future research would be to chalk out the generalization to multiple thresholds and study how the performance converges towards the scenario where the learner observes the agents' actual (uncensored) signals as we increase the number of thresholds.

Chapter 5

Market making in practice: CDA with LMSR

5.1 Introduction and related work

In this chapter, we will review the market microstructure called the *continuous double auction* (CDA) which has traditionally been employed in many real-world financial markets, including early prediction markets, e.g. the Iowa Electronic Markets (Berg and Rietz, 2006)). This microstructure offers participating traders more flexibility in terms of the types of orders they can place, compared to the models we encountered in previous chapters where an agent / trader could only specify whether or not she wished to buy from or sell to the principal (Chapter 4) or additionally the quantity she wanted to trade (Chapters 2 and 3). We will then demonstrate how the market scoring rule concept can be modified to design an agent that operates within this CDA framework, and experimentally study the impact this MSR-based agent has on the properties of a CDA market, particularly on its power to aggregate traders' information. In doing so, we will use an extension of the zero-intelligence trader model mentioned in Chapter 4.

With this in view, let us now step outside the domain of prediction markets into the larger world of more general financial markets, such as those for stocks, bonds, and options, that provide participants with opportunities for hedging, investment, and speculation. The organizers of these markets have always had to deal with the chicken-and-egg problem of providing *liquidity*, which, roughly speaking, is a measure of ongoing trading activity or open interest in the market. To this end, many such markets employ specially designated agents, called

market makers (as we have seen in previous chapters) or, sometimes, specialists, that are responsible for providing liquidity by always being ready to transact with traders. In the last decade or so, research on algorithmic market making has become one of the interesting contact points between artificial intelligence and finance, both in the general context (Das, 2005, 2008b; Wah and Wellman, 2015), and specifically in the design of prediction markets (Hanson, 2003b; Chen and Pennock, 2007; Brahma et al., 2012; Othman et al., 2013; Abernethy et al., 2014).

Most research on algorithmic market making in both financial and prediction markets has either focused on market making as a trading strategy (Chakraborty and Kearns, 2011; Schmitz, 2011) or has modeled the market as a pure dealer market, where the market maker takes one side of every trade (Hanson, 2003b; Das, 2008b; Othman et al., 2013). The market can therefore be modeled in terms of the market maker’s quoted bid (buy) and ask (sell) prices, and traders’ decisions on whether or not to transact at these prices. However, most modern markets, ranging from big financial markets like the NYSE and NASDAQ to smaller prediction markets like the Iowa Electronic Markets, use the continuous double auction (CDA) mechanism (Forsythe et al., 1992). In CDAs, participants can place *limit orders* that specify a transaction price and are guaranteed to only execute at that price or better (although execution is, of course, no longer guaranteed). The key element of CDAs is the limit order book, which contains all active buy and sell limit orders; the highest buy and the lowest sell constitute the market bid and ask prices at any point in time.

While most practical market making algorithms (for example, those used by market makers on the NYSE and NASDAQ) are deployed in markets with limit order books, the academic literature on algorithmic market making has thus far produced almost no analysis of the impact of market making in CDA markets (with the exception of Wah and Wellman (2015)). Here we begin to tackle this problem in the context of market making in prediction markets. The logarithmic market scoring rule proposed by Hanson (2003b) is probably the most commonly deployed automated market maker in prediction markets. Hanson (2003a) also provides a scheme for integrating order books with his market making algorithm which, to the best of our knowledge, has not yet been evaluated in the literature. This scheme, as proposed, involves the market maker having special access to orders before they hit the order book, and a “parallel” implementation that looks at the incoming order, the order book, and executes portions of the trade with the market maker and portions with the existing orders

on the order book. In addition to the special system privileges this requires, it is not entirely transparent to traders, since the order books themselves never reflect the market maker’s presence (and thus give a worse impression of the state of prices and the bid-ask spread than reality).

In this paper, we propose a modification of Hanson’s scheme for integrating LMSR with CDA mechanisms that allows an LMSR-based market making agent to compute limit bid and ask prices and participate in the order books as any other trader would, while still maintaining the key desirable properties – namely improved liquidity with bounded worst-case loss. We call this the “Integrated” market maker (as opposed to the “Parallel” market maker of the original scheme). In general, analysis of the properties of market making algorithms in practice is difficult, since they affect the dynamics of the pricing mechanism itself, and therefore the standard practice of backtesting on historical data is of very limited value. However, there is evidence that simulation models with zero-intelligence (ZI) traders (Gode and Sunder, 1993) can replicate many key features of limit order book dynamics (Farmer et al., 2005a; Othman, 2008) and have practical value in assessing the properties of market making algorithms (Brahma et al., 2012). Therefore, we evaluate market properties in prediction markets populated by ZI traders; we compare the parallel and integrated implementations of LMSR with a situation where no market maker is present, and also a pure dealer market mediated by LMSR.

We are mostly interested in general market properties. In particular,

- **Information aggregation properties:** For example, how fast does the market price converge to the true underlying asset value? How far away is the price from the true value, on average?
- **Market quality properties:** For example, how liquid is the market, as measured by the bid-ask spread? How much surplus or price improvement does the market generate?

In our experiments we find that the presence of the market maker leads to generally lower bid-ask spreads and higher trader surplus (or price improvement), but, surprisingly, does not necessarily improve price discovery and market efficiency; this latter effect is more pronounced when there is higher variability in trader beliefs.

5.2 Market Model

In this section, we describe the precise market model we use and the algorithms used for trading and market making. We simulate four different market microstructures:

- (1) A Continuous Double Auction (CDA) mechanism without any market maker (pureCDA);
- (2) A CDA with the “Integrated” implementation of LMSR (INT);
- (3) A CDA with the “Parallel” implementation of LMSR (PAR);
- (4) A pure dealer framework with all trades going through a traditional LMSR market maker (pureLMSR).

Prediction market We focus on a prediction market set up to forecast whether a single extraneous uncertain event, which can be modeled as a binary random variable X , will occur at some pre-determined future date; on that date, the market terminates, and every unit (share) of the asset traded in the market is worth \$1 if the event occurs and is worthless otherwise; we call this cash equivalent of the asset its liquidation value or true value. Before that date, anyone can place orders to buy or (short-)sell any amount of the asset in the market at prices in the interval $[0, 1]$, i.e. the market institution does not impose any budget constraints on traders. We also assume that there is a fixed probability distribution with $\Pr(X = 1) = p_{\text{true}}$ from which the realization of X is drawn on the termination date so that the expected “true” value of the asset is p_{true} , but no agent in the world knows this p_{true} precisely. However, under a “rational expectations” assumption, the market price in equilibrium is expected to approach p_{true} (Pennock and Sami, 2007).

Types of orders Traders in a financial exchange can typically place buy/sell orders of two kinds: (1) *market orders* that specify only a quantity and demand immediate execution, hence accept any price offered by the other party, and (2) *limit orders* that specify both a quantity and a limit on acceptable transaction prices (called a limit price or marginal price) but are not guaranteed execution. A Continuous Double Auction (CDA) maintains two *order books*, one for buy orders (bids) and the other for sell orders (asks), which are two priority

queues for outstanding limit orders prioritized by limit price and arrival time (higher priority is assigned to a buy order with a higher bid price and a sell order with a lower ask price). Any incoming limit order is placed on the appropriate book, and the mechanism automatically checks to see if the current best (highest) bid is at least as large as the current best (lowest) ask; if yes, then the smaller of the two quantities ordered is traded at the limit price of the order that arrived earlier, the books are updated, and this is continued till the best ask exceeds the best bid. Any new market order is executed immediately, perhaps partially, against the best available outstanding order(s) or is rejected if the book on the other side is empty. In our simulations, all traders place limit orders only but some of them can become market orders effectively, e.g. if an incoming limit buy order “crosses” the books. i.e. its bid is no less than the best ask(s) on the sell order book, and its demand does not exceed the supply of said booked order(s).

Logarithmic Market Scoring Rule We now describe the cost function-based operational formulation of the LMSR market maker, that we touched upon in Chapter 2, for a single-security prediction market liquidating²⁰ in $\{0, 1\}$ (Hanson, 2003b; Chen and Pennock, 2007). Its “state” is described by a real scalar q_{mm} , interpreted as the net outstanding quantity of the security; its instantaneous price at this state, i.e. cost per share of buying/selling an infinitesimal amount from/to LMSR, is given by $p_{\text{mm}} = \frac{e^{q_{\text{mm}}/B}}{1+e^{q_{\text{mm}}/B}}$ where $B > 0$ is a parameter controlling all properties of the market maker. A trader placing a market order for buying any finite quantity Q of assets from LMSR would have to pay it a dollar amount $C(q_{\text{mm}}; Q) = B \ln \left(\frac{1+e^{(q_{\text{mm}}+Q)/B}}{1+e^{q_{\text{mm}}/B}} \right)$ and after the transaction, the market maker’s state is updated to $(q_{\text{mm}} + Q)$; for a sell order, the same formula applies by setting Q to the negative of the supplied quantity, and $-C(q_{\text{mm}}; Q) > 0$ becomes the sales proceeds. One key property of LMSR is that it’s loss is bounded (for the binary case by $B \ln 2$).

Population of traders Every agent other than the market maker is called a “background” trader (Wah and Wellman, 2015). Before every simulation, the expected true asset value p_{true} is chosen at random from a common-knowledge common prior which is a uniform distribution on $[0, 1]$. Every trader i then observes a private sequence of N_{trials} Bernoulli

²⁰The security is said to liquidate when the forecast event is realized, i.e. we know whether $X = 1$ or $X = 0$, so that the market terminates and

trials with probability of success p_{true} , and sets her idiosyncratic valuation of the asset to her Bayesian posterior expectation of the true value, $v_i = \frac{x_i+1}{N_{\text{trials}}+2}$ where x_i is the number of successes in her sample. Thus, N_{trials} is a measure of the precision of the signal that each trader receives, related to the inverse of the variance of beliefs across the population, similar to the model of Zhang et al. (2012). The implementation of a trading decision on top of the belief then follows the zero-intelligence (ZI) trader model (Gode and Sunder, 1993; Othman, 2008), with the addition of non-unit trade sizes. At each step of a simulation (a “trading episode”), a trader is picked uniformly at random and is assigned buyer or seller status with equal probability except for pureLMSR (see below). She then places her limit order, the limit price being drawn uniformly at random from $[v_i, 1]$ if she is a seller and from $[0, v_i]$ if she is a buyer, and the order quantity from a common exponential distribution with mean $\lambda = 20$ which is known to the market mechanism.

(1) pureCDA We have already fully explained the interaction between a CDA mechanism with no market making and the trading population under **Types of Orders**.

(2) PAR The parallel implementation is a single-security version of Robin Hanson’s “booked orders for market scoring rules” (Hanson, 2003a). We delineate its operation for a buy order, the treatment of sell orders being symmetric. Suppose a limit buy order for a quantity q_b at a limit price (bid) p_b arrives when the LMSR market maker’s instantaneous price is p_{mm} , and the current best bid and ask prices are $b_{\text{max}}, a_{\text{min}}$ (at market inception, both books are empty, and $p_{\text{mm}} = 0.5$). If $p_b \leq p_{\text{mm}}$ ²¹, the order cannot be immediately executed, so it is pushed on to the limit buy order book. If $p_{\text{mm}} < p_b$, and q_b is not large enough to drive p_{mm} beyond $\min\{p_b, a_{\text{min}}\}$, then the incoming order is completely executed with the market maker according to the traditional LMSR algorithm; otherwise, if $p_{\text{mm}} < p_b < a_{\text{min}}$, it is only partially executed with LMSR till p_{mm} reaches p_b , the residual order being placed on the buy order book; but if $p_b \geq a_{\text{min}}$, LMSR sells only till its instantaneous price hits a_{min} after which the incoming order executes against the best booked ask. If the top level of the book is exhausted but the incoming order is not, LMSR is invoked again, and this process recurs till either the order is finished or the new best ask exceeds the order’s bid price. The loss bound of the standard LMSR algorithm is maintained in this case.

²¹In this implementation, p_{mm} always lies between the best ask and bid prices on the books, so $p_b \leq p_{\text{mm}}$ implies that p_b does not exceed the minimum ask price either.

(3) INT In this novel “integrated” implementation that we propose, whenever the best ask and bid prices on the books change, an LMSR-based agent steps in.

1. If its instantaneous price $p_{\text{mm}} \leq b_{\text{max}}$, then LMSR generates only a limit sell order for a quantity $q_{\text{ask}} = B \ln \left(\frac{1/p_{\text{mm}} - 1}{1/a_{\text{min}} - 1} \right)$ at an ask price of $\frac{B}{q_{\text{ask}}} \ln \left(\frac{1 - p_{\text{mm}}}{1 - a_{\text{min}}} \right)$.
2. If $p_{\text{mm}} \geq a_{\text{min}}$, then it generates a buy order for $q_{\text{bid}} = B \ln \left(\frac{1/b_{\text{max}} - 1}{1/p_{\text{mm}} - 1} \right)$ at a bid of $\frac{B}{q_{\text{bid}}} \ln \left(\frac{1 - b_{\text{max}}}{1 - p_{\text{mm}}} \right)$.
3. If $b_{\text{max}} < p_{\text{mm}} < a_{\text{min}}$, both orders are generated.

Note that if fully executed immediately these orders would take the LMSR price to b_{max} and a_{min} respectively. The LMSR trader then replaces all its earlier orders with the new order(s) if this action does not immediately cross the books, otherwise it sits idle. After this step, the market is now ready to accept a new order from the background traders, or continue the execution of a partially filled outstanding order, as the case may be. Thus, this market maker can be implemented in practice as just another trader, which is a significant benefit over the PAR framework where the market maker requires some special access to incoming trades and order books. Moreover, any feasible trade with the INT market maker is executed at its actual quoted price rather than following the non-linear LMSR pricing function, which makes trading more transparent and intuitive to traders.

The original LMSR loss bound again holds. Also, we can prove that INT myopically imposes at least as high a cost on the next arriving trader as PAR, assuming that the market makers and order books are in the same state.

Proposition 4. *Suppose the LMSR market maker in both PAR and INT are in state q , and the order books are also otherwise identical. For any next arriving trade, the immediate cost incurred by the next trader is at least as high for INT as it is for PAR.*

The proof is uncomplicated, so we only sketch it here. Consider the last of the three cases for INT above, $b_{\text{max}} < p_{\text{mm}} < a_{\text{min}}$, and let Q^* be the quantity one would need to buy from LMSR to bring its price to a_{min} . Then, if the current state of the INT market maker is q , it will place a sell order of Q^* at an ask of $\frac{C(q; Q^*)}{Q^*}$. Now if a buy order for $Q < Q^*$ arrives with a




Figure 5.1: Illustration of how the market maker in the INT setting places ask and bid quotes every time the state of the books changes.

sufficiently high bid, the whole of it will execute with the market maker, and the immediate earnings of the latter will be $\frac{QC(q;Q^*)}{Q^*}$. If the PAR market maker had the same state q (hence the same p_{mm}) when the same buy order arrived, the ensuing trade would cost the trader $C(q;Q)$ which is less than INT's earnings since $\frac{C(q;Q)}{Q} < \frac{C(q;Q^*)}{Q^*}$ from the convexity of C . Similar arguments apply to the other cases.

This result suggests that INT might provide somewhat less liquidity in general than PAR, and incur less loss in doing so, but we do not expect them to be very different. However, this is a loose prediction, since the result is myopic – it says nothing about price evolution in a market; given the market maker's active role, the dynamics of the evolution of q and the order book could conceivably end up quite different. We examine this issue further in the experiments.

(4) pureLMSR In this setting, traders still place limit orders but an LMSR market maker takes one side of *every* trade. At each trading episode, a trader arrives and compares her private valuation v_i to the current market price p_{mm} . If $v_i > p_{\text{mm}}$, she decides to buy; if $v_i < p_{\text{mm}}$, she decides to sell, and leaves without placing any order otherwise. Then she picks her limit price and order size exactly as the ZI traders above. The quantity bought/sold is the minimum of the order size and the quantity needed to drive the LMSR's instantaneous price to the trader's limit price, and monetary transfers are determined by the above function $C(\cdot; \cdot)$.

Note that all components of each limit order of a trader are independent of the market state for all four settings, except for the direction of the trade (buy/sell) in pureLMSR.

5.3 Evaluation

We present an overview of the various measures we use to evaluate the properties of our market environments.

Information aggregation properties:

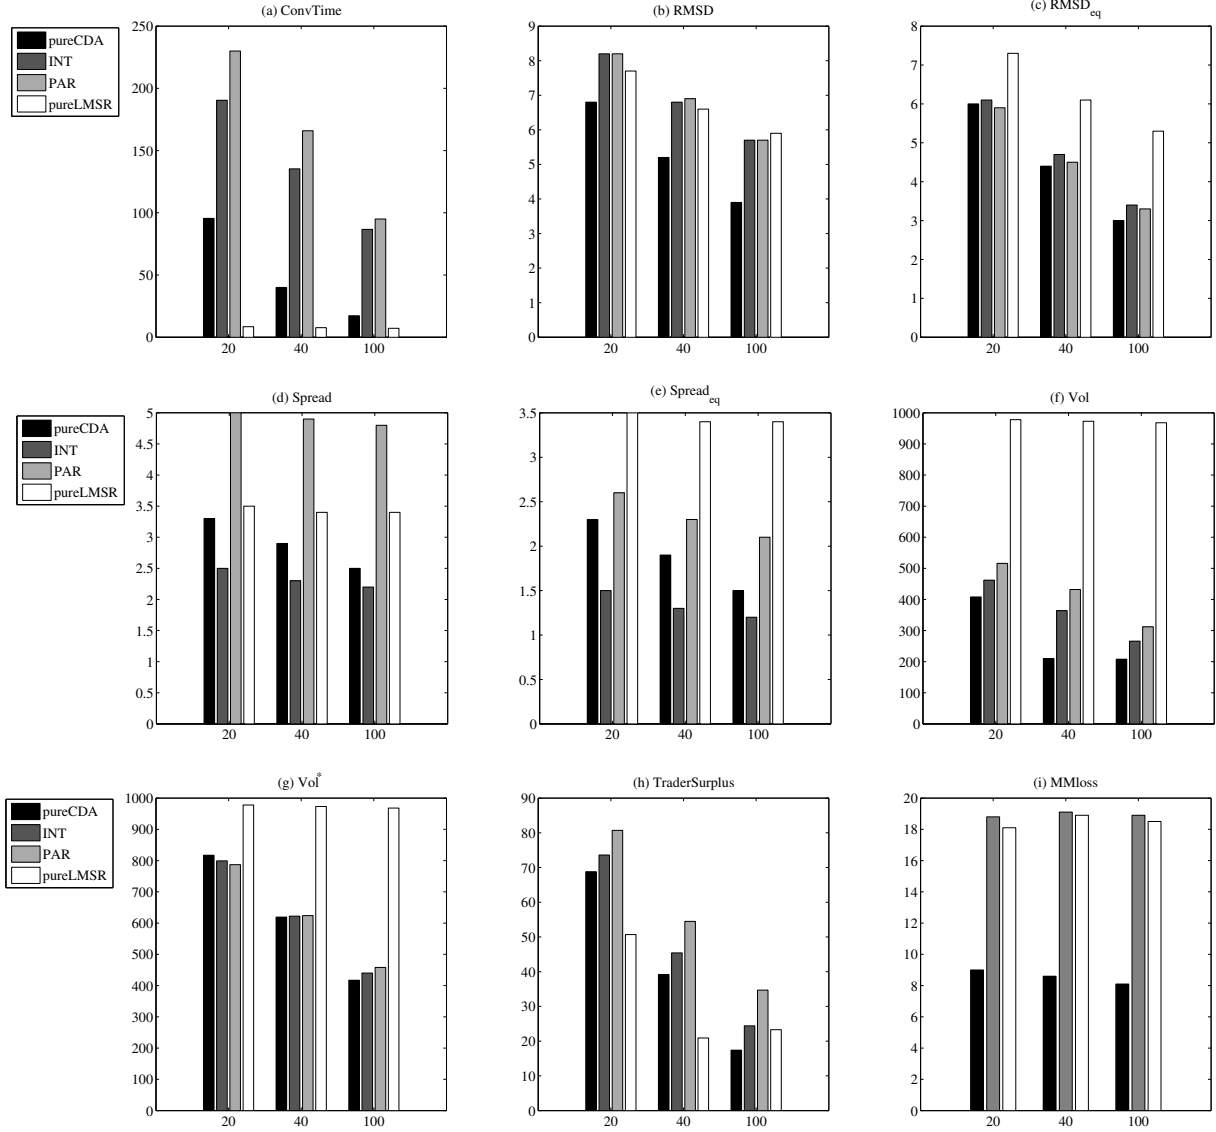


Figure 5.2: Experimental results, averaged over 1000 simulations each. The labels along the horizontal axis indicate the number of private Bernoulli trials with success probability p_{true} observed by each trader in the respective simulation set; this number is directly related to the precision in trader beliefs.

- **ConvTime** (Convergence time): This is defined as the number of trading episodes it takes for the “market price” p_M to get within a band of size ± 0.05 around the true expected asset value p_{true} for the first time; $p_M(t)$ is measured at the end of every trading episode t as the mid-point of the bid-ask spread $((b_{\text{max}}(t) + a_{\text{min}}(t))/2)$ for each of the models with CDA, and as the LMSR instantaneous price for the pure dealer case.²² Thus,

$$\text{ConvTime} = \min\{t : p_M(t) \in [p_{\text{true}} - 0.05, p_{\text{true}} + 0.05]\}.$$
A lower convergence time means that the market’s estimate (price) quickly gets close to the true expected asset value, i.e. the market is efficient.
- **RMSD** and **RMSD_{eq}**: RMSD is the root-mean-squared deviation of the market price (defined above) from p_{true} over the entire simulation (n_{trades} trading episodes). RMSD_{eq} is the root-mean-squared deviation between the same quantities but over only the “equilibrium period”, i.e. for $t \geq \text{ConvTime}$. Lower values of these measures indicate lower price volatility, another desirable property from an information aggregation perspective.

Market quality properties:

- **Spread** and **Spread_{eq}**: For each scenario with a CDA, the market bid and ask prices $b_M(t)$ and $a_M(t)$ at the end of each trading episode are the highest bid b_{max} and the lowest ask a_{min} on the books respectively (set to 0 and 1 if the corresponding book is empty). For the pure dealer setting, we assume that the market maker knows the average order size λ of the trading population, so for a current market state of q_{mm} , the effective market quotes are taken to be $a_M = \frac{C(q_{\text{mm}}; \lambda)}{\lambda}$ and $b_M = \frac{-C(q_{\text{mm}}; -\lambda)}{\lambda}$ which are the prices per share of buying and selling λ shares from and to LMSR at the current state respectively. In our notation, “Spread” denotes the bid-ask spread $(a_M(t) - b_M(t))$ averaged over all n_{trades} episodes, while “Spread_{eq}” is the average taken over the equilibrium period only, as above. The bid-ask spread is widely used as a proxy for market liquidity and smaller values are better, since they imply lower trading costs.
- (Idiosyncratic) **TraderSurplus**: If a trader with idiosyncratic valuation v places a buy order of which a quantity q goes through at an execution price p_{exec} , then the trader’s surplus is defined as $q(v - p_{\text{exec}})$ (similarly, a seller’s surplus is $q(p_{\text{exec}} - v)$). TraderSurplus

²²If the market price does not enter this band over the duration of the simulation, *ConvTime* is set to $n_{\text{trades}} = 500$; in our simulations, this is rarely observed.

denotes the sum of individual surpluses of all background traders. Also note that $(v - p_{\text{exec}})$ and $(p_{\text{exec}} - v)$ correspond loosely to the notion of price improvement, when weighted by the probability of execution at that difference. So, even in settings where the private or idiosyncratic value assumption is untenable, the surplus is still a useful measurement of how much value participants are getting from being in one particular microstructure over another. Since every order executes at a price at least as desirable as its limit price, all trader price improvements (surpluses) are positive.

- **MMloss:** This is the loss incurred by the market making mechanism, computed just like (the negative of) a trader surplus, with the private valuation replaced with the true expected asset value p_{true} . Obviously, this does not apply to pureCDA. Since the market is an ex post zero-sum game between the market maker and the trading population, this measure is also numerically equal to the true expectation of the traders' collective net payoff. This measure is particularly important when the market institution itself subsidizes the market maker.

5.3.1 Results

We ran three sets of 1000 simulations each. In each set, we used a different value of the parameter N_{trials} (20, 40, 100) controlling the precision of trader beliefs. In each simulation, we made the same random sequence of $n_{\text{trades}} = 500$ traders interact with each of our four microstructures. The LMSR parameter B is fixed at 100 for all simulations. We computed all of the above measures for each simulation, and then averaged them over all 1000 simulations. The results are presented in Figure 5.2, and the analysis follows. Note that, the values (rmsd of prices, spreads) depicted in Figures 1(c)-(f) are in cents while those in the last two figures (surplus, losses) are in dollars, for clarity.

Information aggregation: ConvTime (a) follows the pattern: pureLMSR \ll pureCDA $<$ INT $<$ PAR. However, in terms of stability (RMSD, overall (b) and in equilibrium (c)), pureCDA fares the best and the two hybrid mechanisms are very close to each other. The quick convergence and high volatility of LMSR are well-known; surprisingly, coupling it with a CDA delays convergence drastically, but it does ensure more stable prices (lower RMSD_{eq}) once the price converges. While it seems that the market maker-CDA combination might

impede the market’s learning abilities, it is likely in this case to be an artifact of the fixed beliefs held by ZI traders, who stick to their beliefs no matter what happens to the price – it’s not clear that any scoring rule style of market maker would be able to learn quickly when the signals have high variance and the traders don’t update their signals. This hypothesis is borne out by the fact that the effect diminishes as the variance in traders’ beliefs decreases.

Liquidity / Trading activity: Perhaps the biggest reason to deploy a market-maker is to reduce spreads. Figures 1 (d) and (e) show that INT serves this purpose more effectively than pureCDA. The behavior of PAR, which seems to induce very high spreads, is surprising. This behavior is because we measure the market bid and ask only after the extraneous LMSR agent has intervened and perhaps cleared some orders which would still be waiting in the books in the absence of a market maker, so the spread looks artificially large, compared with pureCDA. In addition, PAR doesn’t actually place any new orders on the books, since it waits for orders to arrive before acting, as opposed to INT, which proactively improves spreads by adding to the order book. This finding, which casts doubt on the meaningfulness of spread measurement for PAR, is problematic since many real-life traders use the spread to gauge market quality and make decisions.

To get a better idea of the market maker’s role in improving trading activity, we also computed the actual volume of trade executed. We did this in two ways: for each simulation, we maintained a ledger where each entry recorded the buyer, seller, execution price, and quantity of every market trade; after n_{trades} episodes, we added all these traded quantities together to obtain $\text{Vol} = \text{quantity absorbed by buyers and market maker (if present)} = \text{quantity supplied by sellers and market maker}$. PAR beats both pureCDA and INT with respect to this measure.

We also calculated an alternative measure of trading volume by subtracting the total residual quantity on the order books at the end of each simulation from the total quantity ordered by all traders: $\text{Vol}^* = \text{quantity absorbed by buyers (from sellers and market maker)} + \text{quantity supplied by sellers (to buyers and market maker)}$. It double-counts, perhaps appropriately, every quantity traded between background traders, and thus reflects the overall “satisfaction” of the entire background trader population in a way that the previous measure does not.²³ Strangely, for higher variability in trader beliefs, PAR gives the worst Vol^* bettered

²³Of course, for pureLMSR, $\text{Vol}^* = \text{Vol}$ since the market maker takes one side of every trade.

by INT and pureCDA, but there is a complete reversal in this behavior as the variability decreases. Based on observations of some sample trade ledgers and order book residuals, we believe that the reason is this: in any CDA with a market maker, the market maker gets the advantage of immediacy due to its continuous presence and itself undercuts some of the background traders, thereby reducing the (double-counted) quantity that changes hands between these traders. Hence pureCDA, where every trade must occur between background traders, has a higher Vol^* . But with increasing N_{trials} as trader beliefs get closer to each other, relatively more traders trade with other background traders, who now offer competitive prices themselves. This is an interesting example of how the presence of the market maker can affect the dynamics of trade in surprising ways.

Also note that regardless of the microstructure, both Vol and Vol^* decrease as the knowledge of the trading crowd gets more and more precise, which is consistent with the idea that as the noise in the beliefs of traders with a common knowledge structure reduces, trading becomes less profitable, hence less likely (an extreme case is captured by the no-trade theorems).

Welfare: Trader surplus or (weighted) price improvement decreases with increasing precision in beliefs but the presence of a market maker consistently improves the surplus as opposed to having only a CDA, PAR more so than INT. It is also noteworthy that the combination of CDA and market making performs better in this respect than each of them individually. Moreover, we consistently observe $\text{INT loss} < \text{PAR loss} \approx \text{pureLMSR loss}$, and these losses respect the known LMSR loss bound. This empirical observation supports the notion that Proposition 1 (which shows that myopic costs faced by the market maker are lower for INT than for PAR when they start from the same state in terms of q and the order books) might generalize to expected losses over *sequences* of trades from a particular starting point, an interesting direction for theoretical work on the topic (in a handful of our individual simulations, INT made slightly more loss than PAR, which shows that the sequence result cannot hold deterministically).

5.4 Discussion

In this chapter, we introduced a new LMSR-based market making algorithm that handles limit orders, and also presented one of the first contributions to the academic literature on

the experimental evaluation of algorithmic market making in a CDA setting (for another example, see Wah and Wellman (2015)). Note that we kept our agent model simple since our aim was to focus on the role of the market making entity in price discovery / information aggregation.

A natural next step in this vein would be to analyze these market settings with more sophisticated trader models, particularly (extensions of) those used by Brahma et al. (2012). To give a brief overview, these traders are of two broad types: (1) *fundamentals traders* who base their trading decisions (orders sizes and limit prices) primarily on their private information and may incorporate the observed history of market prices into their decision making, (2) *technical traders* who have no “insider” knowledge of the asset value and strategize on the basis of market price movements, serving mainly as noise traders. Among other things, it will be interesting to study how varying the proportions of these different types of trading agents affects price discovery in the market.

Chapter 6

Conclusion

To summarize this dissertation, I have studied the topic of subjective input aggregation from novel theoretical and experimental points of view, and, along the way, brought out both the similarities and differences among aggregation methods developed in diverse disciplines. In particular, I have explicated the aggregational characteristics of the price process induced by a popular prediction market making algorithm, and discussed the extent to which such mechanisms can be derailed when the potential for manipulative behavior is present. I have put forward an algorithm that can learn a fixed unknown value by only being told whether or not a Gaussian sample centered around that value is above or below a threshold in every discrete time-step, in a number of time-steps that is (asymptotically) near-optimal in certain problem parameters. I have also presented the design of a new market maker that operates in a realistic financial exchange with booked orders, and described an experimental set-up with simulated trading agents for measuring its performance. Finally, I have identified open questions that I believe still need to be addressed in the investigation of the above problem domains.

Appendix A

Proofs of theorems in Chapter 2

A.1 A general well-behaved MSR as an Opinion Pool for a general risk-averse utility

First, we shall recapitulate the mathematical properties of a *well-behaved* market scoring rule (Definition 2 in Chapter 2): The underlying (strictly proper and regular) scoring rule for such an MSR can be written as

$$s_j(p) = \begin{cases} G(p) + G'(p)(j - p), & j \in \{0, 1\}, p \in [0, 1], p \neq j, \\ G(p), & p = j \in \{0, 1\} \end{cases} \quad (\text{A.1})$$

from (2.1) in Chapter 2, where

1. $G : [0, 1] \rightarrow \mathbb{R}$ is a continuous function.
2. $G'(\cdot)$ is real-valued in $[0, 1]$ except possibly that $G'(0) = -\infty$ or $G'(1) = \infty$.
3. $G''(\cdot)$ exists and is positive in $[0, 1]$, $0 < G''(p) < \infty$ for $0 < p < 1$.
4. $G'''(\cdot)$ exists, and $|G'''(p)| < \infty$ for $0 < p < 1$.

Notice that the positivity of $G''(\cdot)$ implies the strict convexity of $G(\cdot)$ and the increasing monotonicity of $G'(\cdot)$. Property 2 ensures that $s_j(\cdot)$ is real-valued except possibly that $s_0(1) = \infty$ or $s_1(0) = \infty$. $G(p) = ps_1(p) + (1 - p)s_0(1 - p)$ is the expected score function

sometimes called the *information measure* or *generalized entropy function* associated with the scoring rule $s_x(\cdot)$ (Gneiting and Raftery (2007)).

For $x \in \{0, 1\}$, the first derivative of $s_x(p)$, $\forall p \in (0, 1)$, is

$$s'_x(p) = G''(p)(x - p) \implies s'_1(p) = G''(p)(1 - p) > 0, \quad s'_0(p) = -G''(p)p < 0,$$

since $G''(p) > 0$. Hence, $s_1(p)$ and $s_0(p)$ are strictly increasing and decreasing functions of p respectively, which is quite intuitive since the reward for predicting a higher probability for the outcome that actually materialized should be higher.

Moreover, if p_{i-1} and p_i denote respectively the instantaneous price of an MSR immediately before and after agent i interacts with it, then by the design of an MSR, the agent's ex post compensation from the market for any outcome $x \in \{0, 1\}$ is given by

$$c_x(p_i, p_{i-1}) = s_x(p_i) - s_x(p_{i-1}).$$

We can readily obtain the following properties of c_x :

$$c_1(p, p_{i-1}) - c_0(p, p_{i-1}) = G'(p) - G'(p_{i-1}); \quad (\text{A.2})$$

$$\frac{\partial}{\partial p} c_x(p, p_{i-1}) = s'_x(p) = G''(p)(x - p), \quad \forall p_{i-1} \in (0, 1), x \in \{0, 1\}. \quad (\text{A.3})$$

Hence, $c_1(p, p_{i-1})$ and $c_0(p, p_{i-1})$ are also strictly increasing and decreasing in p respectively, regardless of p_{i-1} , as expected.

Next, we shall enumerate, from Section 2.4 in Chapter 2, the criteria that an agent utility function $u_i(\cdot)$ must meet in our setting:

1. **Continuity:** $u_i(\cdot)$ is continuous over $[c_i^{\min}, \infty]$ where c_i^{\min} can attain any value in $[-\infty, 0]$.
2. **Increasing monotonicity (Non-satiation):** $u'_i(\cdot)$ is continuous and positive real-valued over $[c_i^{\min}, \infty]$ except possibly that $u'_i(c_{\min}) = \infty$ or $u'_i(\infty) = 0$.
3. **Strict concavity (Risk aversion):** $u''_i(\cdot)$ is continuous and negative real-valued over $[c_i^{\min}, \infty]$ except possibly that $u''_i(c_{\min}) = -\infty$ or $u''_i(\infty) = 0$.

We shall now provide a detailed, joint proof of Lemmas 3 and 4, for completing the proof of Theorem 1 in Section 2.4 of Chapter 2.

Restatement of Lemma 3. *If a myopic agent with subjective probability π_i and a risk-averse utility function of wealth $u_i(\cdot)$, possessing properties 1, 2, and 3 above, trades with a well-behaved market scoring rule for a single Arrow-Debreu security, and updates the market's instantaneous price from $p_{i-1} \in (0, 1)$ to p_i in the process, then p_i is the unique solution in $(0, 1)$ to the following fixed-point equation:*

$$p_i = \frac{\pi_i u'_i(c_1(p_i, p_{i-1}))}{\pi_i u'_i(c_1(p_i, p_{i-1})) + (1 - \pi_i) u'_i(c_0(p_i, p_{i-1}))}. \quad (\text{A.4})$$

Restatement of Lemma 4. *The implicit function $p_i(p_{i-1}, \pi_i)$ described by (A.4) has the following properties:*

1. $p_i = \pi_i$ if and only if $\pi_i = p_{i-1}$.
2. $0 < \min\{p_{i-1}, \pi_i\} < p_i < \max\{p_{i-1}, \pi_i\} < 1$ whenever $\pi_i \neq p_{i-1}$, $0 < \pi_i, \neq p_{i-1} < 1$.
3. For any given p_{i-1} (resp. π_i), p_i is a strictly increasing function of π_i (resp. p_{i-1}).

Proof. If agent i 's subjective probability of $\{X = 1\}$ is $\pi_i \in (0, 1)$ and her utility function is $u_i(\cdot)$, her expected myopic utility for taking a trading action that updates the market price p_{i-1} to any $p \in [0, 1]$ is given by

$$\tilde{u}(p; p_{i-1}, \pi_i) = \pi_i u_i(c_1(p, p_{i-1})) + (1 - \pi_i) u_i(c_0(p, p_{i-1})).$$

The first and second derivatives of the above with respect to p respectively simplify to

$$\begin{aligned} \tilde{u}'(p; p_{i-1}, \pi_i) &= G'''(p) f(p; p_{i-1}, \pi_i); \\ \tilde{u}''(p; p_{i-1}, \pi_i) &= G''''(p) f(p; p_{i-1}, \pi_i) + G'''(p) f'(p; p_{i-1}, \pi_i), \end{aligned}$$

where

$$\begin{aligned}
f(p; p_{i-1}, \pi_i) &= \pi_i(1-p)u'_i(c_1(p, p_{i-1})) - (1-\pi_i)pu'_i(c_0(p, p_{i-1})) \quad \text{so that} \\
f'(p; p_{i-1}, \pi_i) &= -[\pi_i u'_i(c_1(p, p_{i-1})) + (1-\pi_i)u'_i(c_0(p, p_{i-1}))] \\
&\quad + G''(p) [\pi_i u''_i(c_1(p, p_{i-1}))(1-p)^2 + (1-\pi_i)u''_i(c_0(p, p_{i-1}))p^2] \\
&< 0, \quad \forall p \in (0, 1), \text{ given any } \pi_i, p_{i-1} \in (0, 1),
\end{aligned}$$

since $G''(\cdot) > 0$, $u'_i(\cdot) > 0$, and $u''_i(\cdot) < 0$ everywhere. Hence, $f(\cdot)$ is strictly decreasing everywhere, its values at p_{i-1} and π_i being given by

$$f(p_{i-1}; p_{i-1}, \pi_i) = (\pi_i - p_{i-1})u'_i(0); \quad (\text{A.5})$$

$$f(\pi_i; p_{i-1}, \pi_i) = \pi_i(1-\pi_i) [u'_i(c_1(\pi_i, p_{i-1})) - u'_i(c_0(\pi_i, p_{i-1}))]. \quad (\text{A.6})$$

Case I $p_{i-1} < \pi_i$: From (A.2),

$$c_1(\pi_i, p_{i-1}) - c_0(\pi_i, p_{i-1}) = G'(\pi_i) - G'(p_{i-1}) > 0$$

due to the increasing monotonicity of $G'(\cdot)$. But

$$c_1(\pi_i, p_{i-1}) > c_0(\pi_i, p_{i-1}) \implies u'_i(c_1(\pi_i, p_{i-1})) < u'_i(c_0(\pi_i, p_{i-1}))$$

due to the decreasing monotonicity of $u'_i(\cdot)$. Hence, from (A.6), $f(\pi_i; p_{i-1}, \pi_i) < 0$.

Also, from (A.5), since $u'_i(0) > 0$, $f(p_{i-1}; p_{i-1}, \pi_i) > 0$.

These values, along with the decreasing monotonicity of $f(\cdot)$, imply that $f(p; p_{i-1}, \pi_i)$ has a unique zero in (p_{i-1}, π_i) .

Case II $p_{i-1} = \pi_i$: From (A.5) or (A.6),

$$f(\pi_i; p_{i-1}, \pi_i) = f(p_{i-1}; p_{i-1}, \pi_i) = 0,$$

and $\pi_i = p_{i-1}$ is the unique zero of $f(p; p_{i-1}, \pi_i)$ due to its monotonic nature.

Case III $p_{i-1} > \pi_i$: By symmetry, we can argue exactly as for **Case I** that $f(p; p_{i-1}, \pi_i)$ has a unique zero in (π_i, p_{i-1}) .

Thus for any π_i, p_{i-1} , there exists a unique solution in $(0, 1)$, say p^* , to the equation $f(p; p_{i-1}, \pi_i) = 0$. Since $|G''(p^*)|, |G'''(p^*)| < \infty$, we must have

$$\begin{aligned}\tilde{u}'(p^*; p_{i-1}, \pi_i) &= 0; \\ \tilde{u}''(p^*; p_{i-1}, \pi_i) &= G''(p^*)f'(p^*; p_{i-1}, \pi_i) < 0,\end{aligned}$$

since $G''(p^*) > 0$ and $f'(p^*; p_{i-1}, \pi_i) < 0$. In other words, rational risk-averse agent i 's price-update $p_i = \arg \max_{p \in [0, 1]} \tilde{u}(p; p_{i-1}, \pi_i)$ is given by $p_i = p^*$ so that

$$\begin{aligned}f(p_i; p_{i-1}, \pi_i) &= 0 \\ \Rightarrow \pi_i(1 - p_i)u'_i(c_1(p_i, p_{i-1})) &= (1 - \pi_i)p_i u'_i(c_0(p_i, p_{i-1})), \quad \text{from definition} \\ \Rightarrow \frac{p_i}{1 - p_i} &= \frac{\pi_i}{1 - \pi_i} \cdot \frac{u'_i(c_1(p_i, p_{i-1}))}{u'_i(c_0(p_i, p_{i-1}))} \\ \Rightarrow p_i &= \frac{\pi_i u'_i(c_1(p_i, p_{i-1}))}{\pi_i u'_i(c_1(p_i, p_{i-1})) + (1 - \pi_i)u'_i(c_0(p_i, p_{i-1}))}\end{aligned} \tag{A.7}$$

The last step facilitates the interpretation of p_i as a risk-neutral probability. However, for most subsequent proofs, we shall recall the more convenient odds ratio formulation provided in (A.7).

Moreover, it is easy to see that the findings in **Case I**, **Case II**, and **Case III** above jointly imply properties 1 and 2 in the theorem statement. To prove property 3, first note that, for $x \in \{0, 1\}$,

$$\begin{aligned}\frac{\partial}{\partial \pi_i} c_x(p_i(p_{i-1}, \pi_i), p_{i-1}) &= s'_x(p_i) \frac{\partial p_i}{\partial \pi_i} = G''(p_i)(x - p_i) \frac{\partial p_i}{\partial \pi_i}, \\ \frac{\partial}{\partial p_{i-1}} c_x(p_i(p_{i-1}, \pi_i), p_{i-1}) &= s'_x(p_i) \frac{\partial p_i}{\partial p_{i-1}} - s'_x(p_{i-1}) \\ &= G''(p_i)(x - p_i) \frac{\partial p_i}{\partial p_{i-1}} - G''(p_{i-1})(x - p_{i-1})\end{aligned}$$

Now, taking the partial derivative with respect to π_i of both sides of (A.7),

$$\begin{aligned}
\frac{1}{(1-p_i)^2} \frac{\partial p_i}{\partial \pi_i} &= \frac{1}{(1-\pi_i)^2} \frac{u'_i(c_1)}{u'_i(c_0)} + \left(\frac{\pi_i}{1-\pi_i} \right) \frac{u''_i(c_1) \frac{\partial c_1}{\partial \pi_i} u'_i(c_0) - u'_i(c_1) u''_i(c_0) \frac{\partial c_0}{\partial \pi_i}}{(u'_i(c_0))^2} \\
\Rightarrow v_1 \frac{\partial p_i}{\partial \pi_i} &= v_2 + u''_i(c_1) u'_i(c_0) G'''(p_i) (1-p_i) \frac{\partial p_i}{\partial \pi_i} + u'_i(c_1) u''_i(c_0) G'''(p_i) p_i \frac{\partial p_i}{\partial \pi_i} \\
\text{where } v_1 &= \left(\frac{1-\pi_i}{\pi_i} \right) \left(\frac{u'_i(c_0)}{1-p_i} \right)^2, \quad v_2 = \frac{u'_i(c_0) u'_i(c_1)}{\pi_i (1-\pi_i)}, \\
\Rightarrow \frac{\partial p_i}{\partial \pi_i} &= \frac{v_2}{v_1 - G''(p_i) [u''_i(c_1) u'_i(c_0) (1-p_i) + u'_i(c_1) u''_i(c_0) p_i]} \\
&> 0.
\end{aligned}$$

This is because $0 < \pi_i, p_i < 1$, $u'_i(c_1), u'_i(c_0), G'''(p_i) > 0$, and $u''_i(c_1), u''_i(c_0) < 0$ in our model so that $v_1, v_2 > 0$, hence both the numerator and denominator are positive.

Similarly, taking the partial derivative with respect to p_{i-1} of both sides of (A.7),

$$\begin{aligned}
\frac{1}{(1-p_i)^2} \frac{\partial p_i}{\partial p_{i-1}} &= \left(\frac{\pi_i}{1-\pi_i} \right) \frac{u''_i(c_1) \frac{\partial c_1}{\partial p_{i-1}} u'_i(c_0) - u'_i(c_1) u''_i(c_0) \frac{\partial c_0}{\partial p_{i-1}}}{(u'_i(c_0))^2} \\
\Rightarrow v_1 \frac{\partial p_i}{\partial p_{i-1}} &= u''_i(c_1) u'_i(c_0) \left[G'''(p_i) (1-p_i) \frac{\partial p_i}{\partial p_{i-1}} - G'''(p_{i-1}) (1-p_{i-1}) \right] \\
&\quad + u'_i(c_1) u''_i(c_0) \left[G'''(p_i) p_i \frac{\partial p_i}{\partial p_{i-1}} - G'''(p_{i-1}) p_{i-1} \right] \\
\Rightarrow \frac{\partial p_i}{\partial p_{i-1}} &= \frac{-G''(p_{i-1}) [u''_i(c_1) u'_i(c_0) (1-p_{i-1}) + u'_i(c_1) u''_i(c_0) p_{i-1}]}{v_1 - G''(p_i) [u''_i(c_1) u'_i(c_0) (1-p_i) + u'_i(c_1) u''_i(c_0) p_i]} \\
&> 0
\end{aligned}$$

for the same reasons as $\frac{\partial p_i}{\partial \pi_i}$.

Hence $p_i(p_{i-1}, \pi_i)$ is increasing in each of p_{i-1} and π_i , the other remaining constant. \square

Corollary 4. *If $\pi_i > p_{i-1}$ (resp. $\pi_i < p_{i-1}$), then $p_{i-1} < p_i < \pi_i$ (resp. $\pi_i < p_i < p_{i-1}$), i.e. a myopic risk-averse agent moves the market price in the direction of her belief but not all the way.*

This intuitive result follows from the analysis in **Case I** and **Case III** of the above proof.

Corollary 5. *The agents' beliefs as well as the market's initial price put bounds on the instantaneous price at the end of every episode:*

$$\min\{p_0, \pi_1, \pi_2, \dots, \pi_i\} \leq p_i \leq \max\{p_0, \pi_1, \pi_2, \dots, \pi_i\}, \quad \forall i = 1, 2, \dots$$

A.1.1 LMSR as LogOP for CARA utility agents

The following is the proof of Theorem 2 from Section 2.4.1 of Chapter 2.

Restatement of Theorem 2. *The only risk-averse utility function for which myopic agent i , having a subjective belief $\pi_i \in (0, 1)$, and trading with an LMSR market with parameter b and current instantaneous price p_{i-1} , results in the market's updated price p_i being identical to a logarithmic opinion pool between the current price and the agent's subjective belief, i.e.*

$$p_i = \pi_i^{\alpha_i} p_{i-1}^{1-\alpha_i} / [\pi_i^{\alpha_i} p_{i-1}^{1-\alpha_i} + (1 - \pi_i)^{\alpha_i} (1 - p_{i-1})^{1-\alpha_i}], \quad \alpha_i \in (0, 1), \quad (\text{A.8})$$

is given by

$$u_i(c) = \tau_i (1 - \exp(-c/\tau_i)), \quad c \in \mathbb{R} \cup \{-\infty, \infty\}, \quad \text{constant } \tau_i \in (0, \infty), \quad (\text{A.9})$$

the aggregation weight is $\alpha_i = \frac{\tau_i/b}{1+\tau_i/b}$.

Proof. Sufficiency: If agent i 's utility is of the form specified in the theorem, then the first and second derivatives of the utility function are respectively

$$\begin{aligned} u_i'(c) &= \exp(-c/\tau_i) > 0, \quad \text{and} \\ u_i''(c) &= -\exp(-c/\tau_i) / \tau_i < 0 \quad \forall c \in [-\infty, \infty]. \end{aligned}$$

Hence, Lemma 3 is applicable. Making appropriate substitutions in (2.5),

$$\frac{p_i}{1-p_i} = \frac{\pi_i}{1-\pi_i} \cdot \frac{\exp\left(-\frac{b}{\tau_i} \ln\left(\frac{p_i}{p_{i-1}}\right)\right)}{\exp\left(-\frac{b}{\tau_i} \ln\left(\frac{1-p_i}{1-p_{i-1}}\right)\right)} = \left(\frac{\pi_i}{1-\pi_i}\right) \left(\frac{p_i}{p_{i-1}}\right)^{-b/\tau_i} \left(\frac{1-p_i}{1-p_{i-1}}\right)^{b/\tau_i}$$

$$\text{Thus, } \left(\frac{p_i}{1-p_i} \right)^{1+b/\tau_i} = \left(\frac{\pi_i}{1-\pi_i} \right) \left(\frac{p_{i-1}}{1-p_{i-1}} \right)^{b/\tau_i}$$

Exponentiating both sides by $\frac{1}{1+b/\tau_i}$,

$$\frac{p_i}{1-p_i} = \left(\frac{\pi_i}{1-\pi_i} \right)^{\frac{1}{1+b/\tau_i}} \left(\frac{p_{i-1}}{1-p_{i-1}} \right)^{\frac{b/\tau_i}{1+b/\tau_i}} = \left(\frac{\pi_i}{1-\pi_i} \right)^{\alpha_i} \left(\frac{p_{i-1}}{1-p_{i-1}} \right)^{1-\alpha_i},$$

where $\alpha_i = \frac{1}{1+b/\tau_i} = \frac{\tau_i/b}{1+\tau_i/b}$. Simplifying, we get the required LogOP formulation in the theorem statement; alternatively, by taking the logarithm on both sides, we obtain the equivalent additive log-odds ratio formulation.

Necessity: Since we have restricted ourselves to the class of utility functions satisfying criteria 1, 2, and 3, a utility function that results in a logarithmic opinion pool on interacting with LMSR must satisfy Lemma 3 with

$$p_i = \pi_i^{\alpha_i} p_{i-1}^{1-\alpha_i} / [\pi_i^{\alpha_i} p_{i-1}^{1-\alpha_i} + (1-\pi_i)^{\alpha_i} (1-p_{i-1})^{1-\alpha_i}] \quad \text{for some constant } \alpha_i \in (0, 1),$$

or, equivalently, with

$$\frac{\pi_i}{1-\pi_i} = \left(\frac{p_i}{1-p_i} \right)^{\frac{1}{\alpha_i}} \left(\frac{1-p_{i-1}}{p_{i-1}} \right)^{\frac{1-\alpha_i}{\alpha_i}}.$$

Making the requisite substitutions in (2.5) and simplifying, we see that $u'_i(\cdot)$ must satisfy

$$\left(\frac{p_i}{p_{i-1}} \right)^{\frac{1-\alpha_i}{\alpha_i}} u'_i \left(b \ln \left(\frac{p_i}{p_{i-1}} \right) \right) = \left(\frac{1-p_i}{1-p_{i-1}} \right)^{\frac{1-\alpha_i}{\alpha_i}} u'_i \left(b \ln \left(\frac{1-p_i}{1-p_{i-1}} \right) \right) \quad \forall p_i, p_{i-1} \in (0, 1) \quad (\text{A.10})$$

since, owing to the fact that each of π_i and p_{i-1} is allowed to attain any value in $(0, 1)$, p_i defined as the LogOP above can lie anywhere in $(0, 1)$ as well.

Since $0 < \frac{p_{i-1}}{\pi_i}, \frac{1-p_{i-1}}{1-\pi_i} < \infty$, we claim that relation (A.10) is true if and only if $u'_i(\cdot)$ satisfies

$$y^{\frac{1-\alpha_i}{\alpha_i}} u'_i(b \ln(y)) = M_i, \quad \forall y \in (0, \infty), \quad \text{where constant } M_i = u'_i(0). \quad (\text{A.11})$$

The sufficiency is obvious. To establish the necessity, suppose there exists a risk-averse utility function satisfying (A.10) but not (A.11). Then, there must exist $y_1, y_2 \in (0, \infty)$, such that

$y_1 > y_2$ without loss of generality, and

$$h(y_1) \neq h(y_2), \quad \text{where} \quad h(y) = y^{\frac{1-\alpha_i}{\alpha_i}} u'_i(b \ln(y)) \quad \forall y \in (0, \infty).$$

But, if $0 < y_2 < 1 < y_1 < \infty$, we can obtain $\tilde{\pi} = y_2(y_1 - 1)/(y_1 - y_2) \in (0, 1)$ and $\tilde{p} = (y_1 - 1)/(y_1 - y_2) \in (0, 1)$ for which (A.10) is violated, giving us a contradiction. Thus, any $u_i(\cdot)$ satisfying (A.10) must also obey

$$h(y_1) = h(y_2) \quad \forall y_1, y_2 : 0 < y_2 < 1 < y_1 < \infty.$$

This also means that for any two values $y_1, y_3 \in (1, \infty)$, and any given $y_2 \in (0, 1)$, we must have $h(y_1) = h(y_2)$ as well as $h(y_3) = h(y_2)$, implying that $h(y_1) = h(y_3) \quad \forall y_1, y_3 \in (1, \infty)$. By similar reasoning, we can deduce that $h(y_2) = h(y_4) \quad \forall y_2, y_4 \in (0, 1)$. Finally, by the continuity of $h(y)$ at $y = 1$, which in turn follows from the continuity of $u'_i(c)$ at $c = 0$ in our model and the obvious continuity of $y^{\frac{1-\alpha_i}{\alpha_i}}$ at $y = 1$, we arrive at (A.11).

Now, applying the transformation $c = b \ln(y)$, we obtain the first-order ordinary differential equation

$$u'_i(c) = M_i \exp\left(-\frac{1-\alpha_i}{\alpha_i b} c\right), \quad -\infty \leq c \leq \infty$$

where the extreme values of c have been included for continuity. Solving the above, we get

$$\begin{aligned} u_i(c) &= -\frac{M_i \alpha_i b}{1-\alpha_i} \exp\left(-\frac{1-\alpha_i}{\alpha_i b} c\right) + C_i, \quad C_i \text{ being the constant of integration} \\ &= -M_i \tau_i \exp(-c/\tau_i) + C_i, \quad \text{where } \tau_i = \frac{\alpha_i b}{1-\alpha_i} \implies \alpha_i = \frac{\tau_i/b}{1 + \tau_i/b} \\ &\equiv \tau_i (1 - \exp(-c/\tau_i)) \end{aligned}$$

since a utility function is strategically equivalent to any positive-affine transformation of itself. □

A.1.2 LMSR as LinOP for an atypical utility with decreasing absolute risk aversion

Here, we present the proof of Theorem 3 from Section 2.4.2 of Chapter 2.

Restatement of Theorem 3. *If myopic rational agent i , having a subjective belief $\pi_i \in (0, 1)$ and a risk-averse utility function satisfying criteria 1, 2, and 3 in Section A.1 above, trades with a LMSR market with parameter b and current instantaneous price p_{i-1} , then the market's updated price p_i is identical to a linear opinion pool between the current price and the agent's subjective belief, i.e.*

$$p_i = \beta_i \pi_i + (1 - \beta_i) p_{i-1}, \quad \text{for some constant } \beta_i \in (0, 1) \quad (\text{A.12})$$

if and only if agent i 's utility function is of the form

$$u_i(c) = \ln(\exp((c + B_i)/b) - 1), \quad c \geq -B_i, \quad (\text{A.13})$$

where $B_i > 0$ represents agent i 's budget, with the aggregation weight being given by $\beta_i = 1 - \exp(-B_i/b)$.

Proof. If agent i 's utility is of the form specified in the theorem, then by Lemma 2, we can obtain the lower and upper bounds on the feasible values of p_i as follows:

$$\begin{aligned} s_1(p_i^{\min}) &= c_i^{\min} + s_1(p_{i-1}) \\ \Rightarrow b \ln(p_i^{\min}) &= -B_i + b \ln(p_{i-1}) \\ &= b \ln(p_{i-1} \exp(-B_i/b)) \\ &= b \ln(p_{i-1}(1 - \beta_i)), \quad \text{since } \beta_i = 1 - \exp(-B_i/b) \\ \Rightarrow p_i^{\min} &= p_{i-1}(1 - \beta_i), \quad \text{from the monotonicity of } \ln(\cdot); \quad (\text{A.14}) \\ s_0(p_i^{\min}) &= c_i^{\min} + s_0(p_{i-1}) \\ \Rightarrow b \ln(1 - p_i^{\max}) &= -B_i + b \ln(1 - p_{i-1}) = b \ln((1 - p_{i-1}) \exp(-B_i/b)) \\ \Rightarrow 1 - p_i^{\max} &= (1 - p_{i-1}) \exp(-B_i/b) = (1 - p_{i-1})(1 - \beta_i) \\ \Rightarrow p_i^{\max} &= 1 - (1 - p_{i-1})(1 - \beta_i) = \beta_i + (1 - \beta_i)p_{i-1} \quad (\text{A.15}) \end{aligned}$$

Sufficiency: For $-B_i \leq c < \infty$,

$$u'_i(c) = \frac{\exp((c + B_i)/b)}{b(\exp((c + B_i)/b) - 1)} > 0, \quad \text{and}$$

$$u''_i(c) = -\frac{\exp((c + B_i)/b)}{b^2(\exp((c + B_i)/b) - 1)^2} < 0.$$

Hence we can invoke Lemma 3. Now,

$$\begin{aligned} \exp\left(\frac{c_1(p_i, p_{i-1}) + B_i}{b}\right) &= \exp\left(\ln\left(\frac{p_i}{p_{i-1}}\right) + \frac{B_i}{b}\right) \\ &= \exp\left(\ln\left(\frac{p_i}{p_{i-1} \exp(-B_i/b)}\right)\right) \\ &= \frac{p_i}{p_{i-1}(1 - \beta_i)} \\ \Rightarrow \exp\left(\frac{c_1(p_i, p_{i-1}) + B_i}{b}\right) &= \frac{p_i}{p_i^{\min}} \quad \text{from (A.14).} \\ \text{Similarly, } \exp\left(\frac{c_0(p_i, p_{i-1}) + B_i}{b}\right) &= \frac{1 - p_i}{1 - p_i^{\max}} \quad \text{from (A.15).} \end{aligned} \tag{A.16}$$

$$\text{Hence, } \frac{u'_i(c_1(p_i, p_{i-1}))}{u'_i(c_0(p_i, p_{i-1}))} = \frac{\frac{1}{b} \cdot \frac{p_i/p_i^{\min}}{p_i/p_i^{\min} - 1}}{\frac{1}{b} \cdot \frac{(1-p_i)/(1-p_i^{\max})}{(1-p_i)/(1-p_i^{\max}) - 1}} = \frac{p_i}{1 - p_i} \cdot \frac{p_i^{\max} - p_i}{p_i - p_i^{\min}}.$$

It is precisely for obtaining the above ratio that we require the scaling factor of $1/b$, dependent on the market maker parameter, in the exponential in the utility function. Substituting in (2.5), and noting that $p_i/(1 - p_i) \neq 0$ for $0 < p_{i-1} < 1$, we get

$$\begin{aligned} 1 &= \frac{\pi_i}{1 - \pi_i} \cdot \frac{p_i^{\max} - p_i}{p_i - p_i^{\min}} \iff p_i = (1 - \pi_i)p_i^{\min} + \pi_i p_i^{\max} \\ &\iff p_i = \beta_i \pi_i + (1 - \beta_i)p_{i-1}, \end{aligned}$$

on plugging in the expressions for p_i^{\min} and p_i^{\max} from (A.14) and (A.15), and simplifying.

Necessity: Since we have restricted ourselves to the class of utility functions satisfying criteria 1, 2, and 3, a utility function that results in a linear opinion pool on interacting with

LMSR must satisfy Lemma 3 with $p_i = \beta_i \pi_i + (1 - \beta_i)p_{i-1}$ for some constant $\beta_i \in (0, 1)$. Making the requisite substitutions in (2.5) and simplifying, we see that $u'_i(\cdot)$ must satisfy

$$\frac{u'_i \left(b \ln \left(\beta_i \left(\frac{\pi_i}{p_{i-1}} \right) + 1 - \beta_i \right) \right)}{\beta_i + (1 - \beta_i) \frac{p_{i-1}}{\pi_i}} = \frac{u'_i \left(b \ln \left(\beta_i \left(\frac{1 - \pi_i}{1 - p_{i-1}} \right) + 1 - \beta_i \right) \right)}{\beta_i + (1 - \beta_i) \left(\frac{1 - p_{i-1}}{1 - \pi_i} \right)} \quad \forall p_{i-1}, \pi_i \in (0, 1). \quad (\text{A.17})$$

Since $0 < \frac{p_{i-1}}{\pi_i}, \frac{1 - p_{i-1}}{1 - \pi_i} < \infty$, we claim that relation (A.17) is true if and only if $u'_i(\cdot)$ satisfies

$$u'_i(b \ln(\beta_i y + 1 - \beta_i)) = K_i \left(\beta_i + \frac{1 - \beta_i}{y} \right), \quad \forall y \in (0, \infty), \quad (\text{A.18})$$

where constant $K_i = u'_i(0)$, and the (negative) lower bound on the domain of $u_i(\cdot)$ is given by $-B_i = b \ln(1 - \beta_i)$ with $u'_i(-B_i) = \infty$ ²⁴.

The sufficiency is obvious. To establish the necessity, suppose there exists a risk-averse utility function satisfying (A.17) but not (A.18). Then, there must exist $y_1, y_2 \in (0, \infty)$, such that $y_1 > y_2$ without loss of generality, and

$$g(y_1) \neq g(y_2), \quad \text{where} \quad g(y) = \frac{u'_i(b \ln(\beta_i y + 1 - \beta_i))}{\beta_i + \frac{1 - \beta_i}{y}} \quad \forall y \in (0, \infty).$$

But, if $0 < y_2 < 1 < y_1 < \infty$, we can obtain $\tilde{\pi} = y_2(y_1 - 1)/(y_1 - y_2) \in (0, 1)$ and $\tilde{p} = (y_1 - 1)/(y_1 - y_2) \in (0, 1)$ for which (A.17) is violated, giving us a contradiction. Thus, any $u_i(\cdot)$ satisfying (A.17) must also obey

$$g(y_1) = g(y_2) \quad \forall y_1, y_2 : 0 < y_2 < 1 < y_1 < \infty.$$

This also means that for any two values $y_1, y_3 \in (1, \infty)$, and any given $y_2 \in (0, 1)$, we must have $g(y_1) = g(y_2)$ as well as $g(y_3) = g(y_2)$, implying that $g(y_1) = g(y_3) \quad \forall y_1, y_3 \in (1, \infty)$. By similar reasoning, we can deduce that $g(y_2) = g(y_4) \quad \forall y_2, y_4 \in (0, 1)$. Finally, by the continuity of $g(y)$ at $y = 1$, which in turn follows from the continuity of $u'_i(c)$ at $c = 0$ in our model and the obvious continuity of $(\beta_i + (1 - \beta_i)/y)$ at $y = 1$, we arrive at (A.18).

²⁴This constraint is necessary since $\lim_{y \rightarrow 0^+} K_i \left(\beta_i + \frac{1 - \beta_i}{y} \right) = \infty$; also note that K_i is positive real-valued since $u'_i(c) \in (0, \infty)$ for $c \in (-B_i, \infty)$.

Now, applying the transformation $c = b \ln(\beta_i y + 1 - \beta_i)$, we obtain the first-order ordinary differential equation

$$u'_i(c) = \frac{K_i \beta_i \exp(c/b)}{\exp(c/b) - (1 - \beta_i)}, \quad b \ln(1 - \beta_i) \leq c \leq \infty$$

where the extreme values of c have been included for continuity. Solving the above, we get

$$\begin{aligned} u_i(c) &= K_i \beta_i (b \ln(\exp(c/b) - (1 - \beta_i)) + C_i), \quad C_i \text{ being the constant of integration} \\ &= K_i \beta_i (b \ln(\exp(c/b) - \exp(-B_i/b)) + C_i), \quad \text{since } -B_i = b \ln(1 - \beta_i) \\ &= K_i \beta_i b \ln(\exp((c + B_i)/b) - 1) + K_i \beta_i (C_i - B_i) \\ &\equiv \ln(\exp((c + B_i)/b) - 1) \end{aligned}$$

since a utility function is strategically equivalent to any positive-affine transformation of itself. \square

Comment on the linear price update rule: This linear price update induced in a LMSR market by a myopic agent with a static belief π_i and risk-averse utility (A.13) is indistinguishable from that due to a myopic risk-neutral non-Bayesian agent who uses a simple convex combination-based heuristic to learn from the latest market price (taking it as a proxy for the accumulated information of her partially informed peers made public till that point in time), and hence update her point estimate of $\Pr(X = 1)$ from π to $\pi'_i = \beta_i \pi_i + (1 - \beta_i) p_{i-1}$, $\beta_i \in (0, 1)$ being a measure of her confidence in her own private signal; in this non-Bayesian interpretation, too, the agent's budget $B_i = -b \ln(1 - \beta_i)$ turns out to be directly related to her confidence or certainty. This non-Bayesian learning rule can be seen as an example of “adjustment from an anchor”, a well-known heuristic in prospect theory (Tversky and Kahneman, 1974), where an agent uses her private signal π_i as an anchor and, on encountering the market as an additional information source, adjusts her belief away from her anchor by the additive term $(1 - \beta_i)(p_i - \pi_i)$ for making a trading decision.

A.2 LMSR with logarithmic utility agents

Here, we present the proof of Proposition 1 from Section 2.4.3 of Chapter 2.

Proof. Proceeding exactly as in the proof of Theorem 3 in Section 2.4.2, we can deduce the bounds $p_i^{\min} = p_{i-1} \exp(-\tilde{w}_i)$ and $p_i^{\max} = 1 - (1 - p_{i-1}) \exp(-\tilde{w}_i)$, $\tilde{w}_i = w_i/b$, on the feasible market price at the end of trading episode i , and hence rewrite

$$\hat{p}_i = (1 - \pi_i)p_i^{\min} + \pi_i p_i^{\max}.$$

For the logarithmic utility, $u'_i(c) = 1/(c + w_i) > 0$ and $u''_i(c) = -1/(c + w_i)^2 < 0$ for $-w_i \leq c < \infty$ so that we can invoke Lemma 3, and, using (2.5), show that

$$(1 - \pi_i)p_i \ln \left(\frac{p_i}{p_i^{\min}} \right) = \pi_i(1 - p_i) \ln \left(\frac{1 - p_i}{1 - p_i^{\max}} \right). \quad (\text{A.19})$$

Since $0 < \frac{p_i - p_i^{\min}}{p_i}, \frac{p_i^{\max} - p_i}{1 - p_i} < 1$, we can use the well-known Maclaurin series expansion of the logarithmic function

$$\ln(1 + x) = \sum_{j=1}^{\infty} \frac{(-1)^{j+1} x^j}{j}, \quad -1 < x \leq 1$$

to obtain the following:

$$\begin{aligned} \ln \left(\frac{p_i}{p_i^{\min}} \right) &= -\ln \left(1 - \frac{p_i - p_i^{\min}}{p_i} \right) = \frac{p_i - p_i^{\min}}{p_i} + \underline{\delta}_i; \\ \ln \left(\frac{1 - p_i}{1 - p_i^{\max}} \right) &= -\ln \left(1 - \frac{p_i^{\max} - p_i}{1 - p_i} \right) = \frac{p_i^{\max} - p_i}{1 - p_i} + \bar{\delta}_i, \end{aligned}$$

where $\underline{\delta}_i = \sum_{j=2}^{\infty} \frac{1}{j} \left(\frac{p_i - p_i^{\min}}{p_i} \right)^j$, and $\bar{\delta}_i = \sum_{j=2}^{\infty} \frac{1}{j} \left(\frac{p_i^{\max} - p_i}{1 - p_i} \right)^j$.

Substituting in Equation (A.19) and simplifying,

$$p_i = \hat{p}_i + \Delta_i,$$

where $\hat{p}_i = (1 - \pi_i)p_i^{\min} + \pi_i p_i^{\max}$, and $\Delta_i = \pi_i(1 - p_i)\bar{\delta}_i - (1 - \pi_i)p_i\underline{\delta}_i$. □

Appendix B

Proofs of theorems in Chapter 3

B.1 Proof of Proposition 2 from Section 3.4

Restatement of Lemma 2 *For a symmetric well-behaved market scoring rule, the lower and upper thresholds p^L, p^H defined in (3.5) and (3.6) satisfy the equalities*

$$f'(p^L) + f'(p^H) = 2f'(\tfrac{1}{2}); \tag{B.1}$$

$$p^L + p^H = 1; \tag{B.2}$$

$$f(p^H) - f(p^L) = (2p^H - 1)f'(\tfrac{1}{2}) = (1 - 2p^L)f'(\tfrac{1}{2}). \tag{B.3}$$

Proof. **To prove (B.1):** $f'(p^L) + f'(p^H) = 2f'(\tfrac{1}{2})$.

From the equations (3.5) and (3.6) that define p^L and p^H respectively, we get

$$\begin{aligned} f'(p^L) + f'(p^H) &= 2 \left(f(\tfrac{1}{2}) - f(0) \right) + 2 \left(f(1) - f(\tfrac{1}{2}) \right) \\ &= 2 \left(f(1) - f(0) \right) \\ &= 2f'(\tfrac{1}{2}), \end{aligned}$$

putting $y = 1$ in the symmetry condition (3.3).

To prove (B.2): $p^L + p^H = 1$.

For this result, we first establish the more general result that for any symmetric well-behaved MSR,

$$f'(x) + f'(z) = 2f'(\frac{1}{2}), \quad x, z \in (0, 1) \quad \text{if and only if } x + z = 1. \quad (\text{B.4})$$

Proof of sufficiency Setting $\frac{1+y}{2} = x$ in the symmetry condition (3.3), we obtain

$$\begin{aligned} f(x) &= f(1-x) + (2x-1)f'(\frac{1}{2}) \quad \forall x \in (\frac{1}{2}, 1) \\ \text{Thus, } f'(x) &= \lim_{\Delta x \rightarrow 0} \frac{f(x+\Delta x) - f(x)}{\Delta x} \\ &= \lim_{\Delta x \rightarrow 0} \frac{f(1-x-\Delta x) + (2x+2\Delta x-1)f'(\frac{1}{2}) - f(1-x) - (2x-1)f'(\frac{1}{2})}{\Delta x} \\ &= \lim_{\Delta x \rightarrow 0} \frac{f(1-x-\Delta x) - f(1-x) + 2\Delta x f'(\frac{1}{2})}{\Delta x} \\ &= -\lim_{\delta \rightarrow 0} \frac{f(1-x+\delta) - f(1-x)}{\delta} + \lim_{\Delta x \rightarrow 0} 2f'(\frac{1}{2}) \\ &= -f'(1-x) + 2f'(\frac{1}{2}) \end{aligned}$$

Hence, $f'(x) + f'(1-x) = 2f'(\frac{1}{2})$ for any $x \in (\frac{1}{2}, 1)$. Setting $z = 1-x$, we see that $f'(z) + f'(1-z) = 2f'(\frac{1}{2})$ holds for any $z \in (0, \frac{1}{2})$ as well; it is trivially true for $x = z = \frac{1}{2}$. ■

Proof of necessity Since $f'(\cdot)$ is monotonic on $(0, 1)$,

$$f'(x) + f'(z) \neq f'(x) + f'(1-x) \quad \text{for any } z \neq 1-x. \quad \blacksquare$$

Since, from result (3.8), we know that $f'(p^L) + f'(p^H) = 2f'(\frac{1}{2})$, we can conclude from (B.4) that $p^L + p^H = 1$.

To prove (B.3): $f(p^H) - f(p^L) = (2p^H - 1)f'(\frac{1}{2}) = (1 - 2p^L)f'(\frac{1}{2})$. If we set $p^H = \frac{1+y}{2}$, then $y = 2p^H - 1$. From (3.9), $y = 2(1 - p^L) - 1 = 1 - 2p^L$ and $p^L = \frac{1-y}{2}$. Hence, we can invoke the symmetry condition (3.3) to get the desired result. □

B.2 Equilibrium of the game in Section 3.4.3

We shall now proceed to determine Alice's choice of $p_A = p_A^{PBE}$ in a perfect Bayesian equilibrium, as stated in Theorem 4 and Tables 3.3, 3.4, and 3.5 in Section 3.4.3. Although the theorem applies to the special case $p_0 = \frac{1}{2}$, we will begin by proving results (up to and including Lemma 7) for the more general scenario $p_0 \in (p^L, p^H)$ which subsumes this special case, and then restrict ourselves to $p_0 = \frac{1}{2}$.

Propositions 5 and 6 imply that

$$p_A^{PBE} = \arg \max_p \widehat{\mathcal{R}}_A(p; p_0),$$

where $\widehat{\mathcal{R}}_A(p; p_0) = \mathbb{E}_A[\mathcal{R}_A(p_A, p_0, v_A, v_B) | p_A = p], \quad \forall p \in [0, 1],$

$\mathbb{E}_A[\cdot | p_A = p]$ denoting the expectation with respect to Alice's belief under the assumption that she will choose the best v_A given $p_A = p$ (Lemma 5), and will take into account the impact of her trading choice $p_A = p$ on STRATEGIC Bob's choices (Lemma 6).

Proposition 5. *For $p \in [0, 1] \setminus \{p^L, p^H\}$ and $p_0 \in (0, 1)$,*

$$\widehat{\mathcal{R}}_A(p; p_0) = s(p, \mathbb{E}_A[v | p_A = p] - s(p_0, \mathbb{E}_A[v | p_A = p])), \quad v = \frac{v_A + v_B}{2},$$

where $\mathbb{E}_A[v | p_A = p]$, according to our notation, is Alice's posterior expected average vote (market outcome) just before Bob has the opportunity to trade. The expressions for $\mathbb{E}_A[v | p_A = p]$ in terms of Bob's commonly known non-participation probability π and Alice's posterior probability q_0 of Bob's signal being $s_B = 0$, given her own signal s_A , for the different sub-intervals in which p may lie, are presented in Table B.1.

p	$\mathbb{E}_A[v p_A = p]$
$0 \leq p < p^L$	$\mu_{0,1} = \frac{1-\pi q_0}{2}$
$p^L < p < p_0$	$\mu_{0,0} = \frac{\pi(1-q_0)}{2}$
$p_0 < p < p^H$	$\mu_{1,1} = 1 - \frac{\pi q_0}{2}$
$p^H < p \leq 1$	$\mu_{1,0} = \frac{1+\pi(1-q_0)}{2}$

Table B.1

Proof. Using Propositions 5 and 6 and recalling that HONEST Bob votes his true signal s_B , we can easily verify the expressions for $\mathbb{E}_A[v|p_A = p]$ provided in Table B.1. Hence, using (3.2), we can show that for $0 < p^L < p_0 < p^H < 1$, and $p \in [0, 1] \setminus \{p^L, p^H\}$,

$$\begin{aligned}\widehat{\mathcal{R}}_A(p; p_0) &= \mathbb{E}_A[s(p_A, v) - s(p_0, v)|p_A = p] \\ &= \mathbb{E}_A[s(p, v) - s(p_0, v)|p_A = p] \\ &= s(p, \mathbb{E}_A[v|p_A = p]) - s(p_0, \mathbb{E}_A[v|p_A = p]),\end{aligned}$$

the key idea being the linearity of the scoring rule function $s(r, \omega)$ in ω for $r \neq \omega$.

We shall demonstrate the proof idea for the special the case $p_A = p \in [0, p^L)$ which corresponds to the first row of the Table B.1 (the treatment of the other cases being similar): In this case, Alice will certainly vote $v_A = 0$ but, from her perspective, Bob's participation is still uncertain so that $v = \frac{v_A + v_B}{2} \in \{0, \frac{1}{2}\}$. However, Alice knows that Bob will vote $v_B = 1$ definitely if is STRATEGIC (which happens with probability $(1 - \pi)$) or with probability $\Pr(s_B = 1|s_A) = (1 - q_0)$ if he is HONEST; Bob will vote $v_B = 0$ otherwise. Hence, given $p_A = p \in [0, p^L)$,

$$\begin{aligned}v &= \begin{cases} \frac{1}{2} & \text{with probability } (1 - \pi) \cdot 1 + \pi(1 - q_0) = 1 - \pi q_0, \\ 0 & \text{otherwise} \end{cases} \\ \Rightarrow \mathbb{E}_A[v|p_A = p \in [0, p^L)] &= \frac{1}{2} \cdot (1 - \pi q_0) + 0 \cdot (\pi q_0) = \frac{1 - \pi q_0}{2} = \mu_{0,1}.\end{aligned}$$

Now, for the sub-case $p \in (0, p^L)$, $p_0 \neq \frac{1}{2}$:

$$\begin{aligned}\widehat{\mathcal{R}}_A(p; p_0) &= (s(p, 0) - s(p_0, 0))\pi q_0 + (s(p, \frac{1}{2}) - s(p_0, \frac{1}{2}))(1 - \pi q_0) \\ &= (f(p) - f'(p)p - f(p_0) + f'(p_0)p_0)\pi q_0 \\ &\quad + (f(p) + f'(p)(\frac{1}{2} - p) - f(p_0) - f'(p_0)(\frac{1}{2} - p_0))(1 - \pi q_0) \quad \text{from (3.2)} \\ &= f(p) + f'(p)\left(\frac{1 - \pi q_0}{2} - p\right) - [f(p_0) + f'(p_0)\left(\frac{1 - \pi q_0}{2} - p_0\right)], \quad \text{on rearrangement} \\ &= s(p, \mu_{0,1}) - s(p_0, \mu_{0,1}) \quad (\text{Table B.1, first row}).\end{aligned}$$

But for $p \in (0, p^L)$, $p_0 = \frac{1}{2}$,

$$\begin{aligned}
\widehat{\mathcal{R}}_A(p; \tfrac{1}{2}) &= (f(p) - f'(p)p - f(\tfrac{1}{2}) + \tfrac{1}{2}f'(\tfrac{1}{2})) \pi q_0 \\
&\quad + (f(p) + f'(p)(\tfrac{1}{2} - p) - f(\tfrac{1}{2})) (1 - \pi q_0) \\
&= f(p) + f'(p) (\tfrac{1 - \pi q_0}{2} - p) - [f(\tfrac{1}{2}) + f'(\tfrac{1}{2}) (\tfrac{1 - \pi q_0}{2} - \tfrac{1}{2})] \\
&= s(p, \mu_{0,1}) - s(\tfrac{1}{2}, \mu_{0,1}).
\end{aligned}$$

Similar calculations apply to the sub-cases $p = 0, p_0 \neq \frac{1}{2}$ and $p = 0, p_0 = \frac{1}{2}$, leading to the same conclusion. \square

Note that $\widehat{\mathcal{R}}_A(p_0; p_0) = 0$, as expected. Moreover, for any $\pi, q_0 \in (0, 1)$,

$$0 < \mu_{0,0} < \mu_{0,1} < \tfrac{1}{2} < \mu_{1,0} < \mu_{1,1} < 1. \quad (\text{B.5})$$

For our subsequent analysis, it is convenient to define a family of functions

$$g_{u,v}(p; p_0) \triangleq s(p, \mu_{u,v}) - s(p_0, \mu_{u,v}), \quad u, v \in \{0, 1\}, p \in [0, 1]. \quad (\text{B.6})$$

Corollary 6. *From Proposition 5, it follows that for any $p_0 \in (p^L, p^H)$, which includes $p_0 = \frac{1}{2}$,*

$$\widehat{\mathcal{R}}_A(p; p_0) = \begin{cases} g_{0,1}(p; p_0), & 0 \leq p < p^L \\ g_{0,0}(p; p_0), & p^L < p < p_0 \\ 0, & p = p_0 \\ g_{1,1}(p; p_0), & p_0 < p < p^H \\ g_{1,0}(p; p_0), & p^H < p \leq 1 \end{cases}$$

This readily leads to the following properties of $\widehat{\mathcal{R}}_A(p; p_0)$, which turn out to be crucial in determining the global maximum of $\widehat{\mathcal{R}}_A(\cdot; p_0)$ over $[0, 1]$:

Property 1 From the properties of proper scoring rules, it is clear that the unique maximum of $g_{u,v}(\cdot; p_0) = s(p, \mu_{u,v}) - s(p_0, \mu_{u,v})$ over $[0, 1]$ for any given p_0 and (u, v) -pair is $\mu_{u,v}$; however, depending on the values of p^L (determined by the scoring

function), π , and q_0 , the value $\mu_{u,v}$ (as in Table B.1) might not lie in the sub-interval of $[0, 1]$ over which $\widehat{\mathcal{R}}_A(\cdot; p_0)$ coincides with $g_{u,v}(\cdot; p_0)$, in which case we should take into account the supremum of the relevant segment lying at one of its end-points. For example, if $\mu_{0,1} < p^L$, then $\sup_{p \in [0, p^L)} \widehat{\mathcal{R}}_A(p; p_0) = g_{0,1}(\mu_{0,1}; p_0)$ so that the maximizer of the segment of $\widehat{\mathcal{R}}_A(\cdot; p_0)$ over $[0, p^L)$ is $\mu_{0,1}$, but if $\mu_{0,1} > p^L$, then $\sup_{p \in [0, p^L)} \widehat{\mathcal{R}}_A(p; p_0) = g_{0,1}(p^L; p_0)$ which is achieved at the extremity $p = p^L$. Symmetric results hold for the other three segments.

Property 2 From definition, $g_{u,v}(p_0; p_0) = 0$ for any $u, v \in \{0, 1\}$. In particular, $g_{0,0}(p_0; p_0) = g_{1,1}(p_0; p_0) = 0$, hence $\widehat{\mathcal{R}}_A(p; p_0)$ is continuous at $p = p_0$. However, as we establish in Lemma 7 below, $\widehat{\mathcal{R}}_A(p; p_0)$ has jump discontinuities at the thresholds p^L and p^H .

Lemma 7. *For any $p_0 \in (p^L, p^H)$,*

$$g_{0,0}(p^L; p_0) > g_{0,1}(p^L; p_0), \quad \text{and} \quad g_{1,1}(p^H; p_0) > g_{1,0}(p^H; p_0),$$

regardless of $\pi, q_0 \in (0, 1)$.

Proof. From (B.6), (3.2), and Table B.1,

$$\begin{aligned} g_{0,0}(p^L; p_0) - g_{0,1}(p^L; p_0) &= \left(\frac{1-\pi}{2}\right) (f'(p_0) - f'(p^L)) \\ &> 0, \end{aligned}$$

from the increasing monotonicity of $f'(\cdot)$, since $p_0 > p^L$, and $\pi < 1$. Similarly,

$$g_{1,1}(p^H; p_0) - g_{1,0}(p^H; p_0) = \left(\frac{1-\pi}{2}\right) (f'(p^H) - f'(p_0)) > 0,$$

from the increasing monotonicity of $f'(\cdot)$, since $p^H > p_0$. □

The following lemmas hold for the particular case of $p_0 = \frac{1}{2} \in (p^L, p^H)$, and we shall invoke them repeatedly through out our equilibrium analysis.

Lemma 8. For any $\pi, q_0 \in (0, 1)$,

$$\begin{aligned} g_{1,1}(p^H; \tfrac{1}{2}) &\geq g_{0,0}(p^L; \tfrac{1}{2}), \\ \text{and } g_{1,0}(p^H; \tfrac{1}{2}) &\geq g_{0,1}(p^L; \tfrac{1}{2}) \\ \iff q_0 &\leq \tfrac{1}{2}. \end{aligned}$$

Proof. For $i, j, k, l \in \{0, 1\}$, from (B.6), (3.2), and Table B.1,

$$\begin{aligned} g_{i,j}(p^H; \tfrac{1}{2}) - g_{k,l}(p^L; \tfrac{1}{2}) &= [f(p^H) - f(p^L)] + f'(p^H)(\mu_{i,j} - p^H) - f'(p^L)(\mu_{k,l} - p^L) \\ &\quad - f'(\tfrac{1}{2})(\mu_{i,j} - \mu_{k,l}) \\ &= (p^H - p^L)f'(\tfrac{1}{2}) + f'(p^H)(\mu_{i,j} - p^H) \\ &\quad - (2f'(\tfrac{1}{2}) - f'(p^H))(\mu_{k,l} - p^L) - f'(\tfrac{1}{2})(\mu_{i,j} - \mu_{k,l}) \\ &\text{from symmetry condition (3.3) since } p^H - \tfrac{1}{2} = \tfrac{1}{2} - p^L, \\ &\text{and from the result (3.8),} \\ &= (f'(p^H) - f'(\tfrac{1}{2}))(\mu_{i,j} + \mu_{k,l} - 1), \\ &\text{on simplification, using the result } p^L + p^H = 1. \end{aligned}$$

In particular, for $\{i, j, k, l\} = \{1, 1, 0, 0\}$ and $\{i, j, k, l\} = \{1, 0, 0, 1\}$ respectively,

$$\begin{aligned} g_{1,1}(p^H; \tfrac{1}{2}) - g_{0,0}(p^L; \tfrac{1}{2}) &= g_{1,0}(p^H; \tfrac{1}{2}) - g_{0,1}(p^L; \tfrac{1}{2}) \\ &= (f'(p^H) - f'(\tfrac{1}{2}))\pi(\tfrac{1}{2} - q_0) \end{aligned}$$

Since $f'(p^L) > f'(\tfrac{1}{2})$ by the increasing monotonicity of $f'(\cdot)$ and $\pi > 0$, the above identity implies the required result. \square

Lemma 9. For any $\pi, q_0 \in (0, 1)$, and any $i, j \in \{0, 1\}$,

$$g_{1,i}(\mu_{1,i}; \tfrac{1}{2}) \geq g_{0,j}(\mu_{0,j}; \tfrac{1}{2}) \iff \mu_{1,i} + \mu_{0,j} \geq 1.$$

Proof. From (B.6) and (3.2),

$$\begin{aligned}
g_{1,i}(\mu_{1,i}; \tfrac{1}{2}) - g_{0,j}(\mu_{0,j}; \tfrac{1}{2}) &= f(\mu_{1,i}) - f(\mu_{0,j}) - f'(\tfrac{1}{2})(\mu_{1,i} - \mu_{0,j}) \\
&= [f(1 - \mu_{1,i}) + f'(\tfrac{1}{2})(2\mu_{1,i} - 1)] - f(\mu_{0,j}) \\
&\quad - f'(\tfrac{1}{2})(\mu_{1,i} - \mu_{0,j}), \\
&\quad \text{from symmetry condition 3.3 since } \mu_{1,i} > \tfrac{1}{2} \text{ by (B.5),} \\
&= [f(1 - \mu_{1,i}) - f(\mu_{0,j})] + f'(\tfrac{1}{2})(\mu_{1,i} + \mu_{0,j} - 1) \\
&> f'(\mu_{0,j})(1 - \mu_{1,i} - \mu_{0,j}) + f'(\tfrac{1}{2})(\mu_{1,i} + \mu_{0,j} - 1) \\
&\quad \text{due to the strict convexity of } f(\cdot), \\
&= (f'(\tfrac{1}{2}) - f'(\mu_{0,j}))(\mu_{1,i} + \mu_{0,j} - 1),
\end{aligned}$$

where $f'(\tfrac{1}{2}) - f'(\mu_{0,j}) > 0$ from the increasing monotonicity of $f'(\cdot)$, since $\mu_{0,j} < \tfrac{1}{2} \forall j \in \{0, 1\}$ by (B.5). Hence, the required result follows from the above inequality. \square

Lemma 10. *For any $\pi, q_0 \in (0, 1)$,*

$$g_{1,0}(\mu_{1,0}; \tfrac{1}{2}) \gtrless g_{1,1}(p^H; \tfrac{1}{2}) \iff \pi \gtrless \pi_H^*(q_0),$$

where $\pi_H^*(q_0)$ is the unique zero in $(0, 1)$ of the function

$$\mathcal{F}_H(\pi; q_0, f) \triangleq g_{1,0}(\mu_{1,0}; \tfrac{1}{2}) - g_{1,1}(p^H; \tfrac{1}{2}), \quad 0 \leq \pi \leq 1. \quad (\text{B.7})$$

Likewise,

$$g_{0,1}(\mu_{0,1}; \tfrac{1}{2}) \gtrless g_{0,0}(p^L; \tfrac{1}{2}) \iff \pi \gtrless \pi_L^*(q_0),$$

where $\pi_L^*(q_0)$ is the unique zero in $(0, 1)$ of the function

$$\mathcal{F}_L(\pi; q_0, f) \triangleq g_{0,1}(\mu_{0,1}; \tfrac{1}{2}) - g_{0,0}(p^L; \tfrac{1}{2}), \quad 0 \leq \pi \leq 1. \quad (\text{B.8})$$

The exact locations of $\pi_H^*(q_0)$ and $\pi_L^*(q_0)$ in $(0, 1)$, however, depend on the magnitude of q_0 relative to p^L, p^H , and we shall address these issues as and when necessary.

Proof. From (B.6), (3.2), and Table B.1,

$$\mathcal{F}_H(\pi; q_0, f) = f\left(\frac{1+\pi(1-q_0)}{2}\right) - f(p^H) - f'(p^H) \left(1 - p^H - \frac{\pi q_0}{2}\right) + f'\left(\frac{1}{2}\right) \left(\frac{1-\pi}{2}\right). \quad (\text{B.9})$$

Evidently, \mathcal{F}_H is continuous according to our criteria, and its first derivative of \mathcal{F}_H with respect to π is

$$\begin{aligned} \frac{\partial \mathcal{F}_H}{\partial \pi} &= f'\left(\frac{1+\pi(1-q_0)}{2}\right) \left(\frac{1-q_0}{2}\right) + \frac{q_0}{2} f'(p^H) - \frac{1}{2} f'\left(\frac{1}{2}\right) \\ &> \left(\frac{1-q_0}{2}\right) f'\left(\frac{1}{2}\right) + \frac{q_0}{2} f'(p^H) - \frac{1}{2} f'\left(\frac{1}{2}\right) \\ &\quad \text{since } \frac{1-q_0}{2} > 0 \text{ and } \frac{1+\pi(1-q_0)}{2} > \frac{1}{2} \text{ for } \pi, q_0 \in (0, 1), \\ &\quad \text{and } f'(\cdot) \text{ is strictly increasing in } (0, 1). \\ &= \frac{q_0}{2} \left(f'(p^H) - f'\left(\frac{1}{2}\right)\right) > 0 \quad \text{since } p^H > \frac{1}{2}. \end{aligned}$$

Hence, $\mathcal{F}_H(\cdot; q_0, f(\cdot))$ is a strictly increasing function for $\pi \in (0, 1)$. Moreover, for any $q \in (0, 1)$,

$$\begin{aligned} \lim_{\pi \nearrow 1} \mathcal{F}_H(\pi; q_0, f) &= f(1 - \frac{q_0}{2}) - f(p^H) - f'(p^H) \left(1 - \frac{q_0}{2} - p^H\right) \\ &> 0, \quad \text{from the strict convexity of } f(\cdot); \\ \lim_{\pi \searrow 0} \mathcal{F}_H(\pi; q_0, f) &= f\left(\frac{1}{2}\right) - f(p^H) - f'(p^H) \left(1 - p^H\right) + \frac{1}{2} f'\left(\frac{1}{2}\right) \\ &< f'\left(\frac{1}{2}\right) \left(\frac{1}{2} - p^H\right) + \frac{1}{2} f'\left(\frac{1}{2}\right) - f'(p^H) \left(1 - p^H\right) \\ &\quad \text{from the strict convexity of } f(\cdot), \\ &= - \left(1 - p^H\right) \left(f'(p^H) - f'\left(\frac{1}{2}\right)\right) \\ &< 0 \quad \text{since } \frac{1}{2} < p^H < 1, f'(\cdot) \text{ is strictly increasing.} \end{aligned}$$

From the above analysis, we conclude that $\mathcal{F}_H(\pi; q_0, f) = g_{1,0}(\mu_{1,0}; \frac{1}{2}) - g_{1,1}(p^H; \frac{1}{2})$ has a unique zero at some $\pi = \pi_H^*(q_0) \in (0, 1)$, is strictly negative for $\pi < \pi_H^*(q_0)$, and strictly positive for $\pi > \pi_H^*(q_0)$. This completes the proof.

The proof for the second part involving $\mathcal{F}_L(\pi; q_0, f)$ is analogous by symmetry, hence omitted. \square

All subsequent analysis applies to the particular case of $p_0 = \frac{1}{2}$. A perfect Bayesian equilibrium (PBE) of the two-player two-stage (trade-voting) game under consideration is a specification of a strategy profile, which in this case is the vector $(p_A^{PBE}, v_A^{PBE}, p_B^{PBE}, v_B^{PBE})$, and a consistent (Bayesian) belief system. We have already established in Lemma 6 that all relevant information about Alice's strategy that Bob needs in order to make his own decision, in case he ends up participating in the prediction market, is available directly from p_A , and there is no need for Bob to go through the process of updating his belief about Alice's signal s_A and hence reasoning about Alice's voting choice v_A . Hence, we can safely abstract away from explicitly describing Bob's belief system for our particular game. We can also abstract away from specifying how Alice updates her belief about Bob if and after Bob trades because, regardless of Bob's behavior in the market, Alice's voting choice is already fixed by the decision she makes in the first stage of the game (Lemma 5), based on her belief about Bob's actions immediately after obtaining her signal.

This still leaves us with the issue of reasoning about Alice's equilibrium price-report p_A^{PBE} . Note that, by Lemma 6, Bob is indifferent between $v_B = 0$ and $v_B = 1$ if $p_A \in \{p^L, p^H\}$ although Bob's voting choice for these values of p_A is crucial for Alice's decision making due to the jump discontinuities in $\widehat{\mathcal{R}}_A(p; p_0)$ at $p = p^L$ and $p = p^H$, as indicated by Lemma 7. But we cannot have Bob $(p_A = p^L, v_B = 1)$ or $(p_A = p^H, v_B = 0)$ as part of an equilibrium as in Section 3.4.1: If Alice knew that Bob would respond with $v_B = 1$ to $p_A \in [0, p^L]$ and $v_B = 0$ to $p_A \in (p^L, \frac{1}{2}]$ (resp. $v_B = 0$ to $p_A \in [p^H, 1]$ and $v_B = 1$ to $p_A \in (\frac{1}{2}, p^H)$), she would prefer to set p_A to a value greater than (resp. less than) but close enough to p^L (resp. p^H) so as to get a higher expected profit. Finally, observe, regardless of which $g_{u,v}(p; \frac{1}{2})$ we consider at $p = \frac{1}{2}$, Alice's payoff for $p_A = \frac{1}{2}$ is always zero. Thus, finding Alice's equilibrium price-report p_A^{PBE} reduces to the problem of figuring out the local suprema of the segments of the functions $g_{0,1}(p; \frac{1}{2})$, $g_{0,0}(p; \frac{1}{2})$, $g_{1,1}(p; \frac{1}{2})$, $g_{1,0}(p; \frac{1}{2})$, defined above, over the sub-intervals $[0, p^L)$, $[p^L, \frac{1}{2}]$, $(\frac{1}{2}, p^H]$, $(p^H, 1]$ respectively, and then comparing them to determine the global maximum of $\widehat{\mathcal{R}}_A(p; \frac{1}{2})$ over $0 \leq p \leq 1$ for different values of $\pi, q_0 \in (0, 1)$.

B.2.1 Analysis for symmetric well-behaved MSRs with $0 < p^L < \frac{1}{4}$, $\frac{3}{4} < p^H < 1$, e.g. LMSR

Case 1 $0 < q_0 < 2p^L$ In this case, $q_0 < \frac{1}{2} < 1 - 2p^L$, so that $\frac{2p^L}{1-q_0} < \frac{1-2p^L}{1-q_0} < 1 < \frac{2p^L}{q_0} < \frac{1-2p^L}{q_0}$. Moreover, since $\pi < 1$,

$$\begin{aligned}\mu_{0,1} &= \frac{1-\pi q_0}{2} > \frac{1}{2} - \pi p^L > \frac{1}{2} - p^L > p^L, \quad \text{since } p^L < \frac{1}{4}; \\ \mu_{1,1} &= 1 - \frac{\pi q_0}{2} > 1 - \pi p^L > 1 - p^L = p^H.\end{aligned}$$

Hence, regardless of π , the local suprema of $\widehat{\mathcal{R}}_A$ over the intervals $[0, p^L)$ and $(p^H, 1]$ are $g_{0,1}(p^L; \frac{1}{2})$ and $g_{1,1}(p^H; \frac{1}{2})$ at p^L and p^H respectively.

Case 1.1 $0 < q_0 < 2p^L$ and $0 < \pi \leq \frac{2p^L}{1-q_0}$. In this case,

$$\begin{aligned}\mu_{0,0} &= \frac{\pi(1-q_0)}{2} \leq p^L, \quad \text{equality holding if and only if } \pi = \frac{2p^L}{1-q_0}; \\ \mu_{1,0} &= \frac{1+\pi(1-q_0)}{2} \leq \frac{1}{2} + p^L = \frac{3}{2} - p^H < p^H, \\ &\quad \text{since } p^H = 1 - p^L > \frac{3}{4}.\end{aligned}$$

Hence, the local suprema of $\widehat{\mathcal{R}}_A$ over the intervals $[p^L, \frac{1}{2}]$ and $(p^H, 1]$ are $g_{0,0}(p^L; \frac{1}{2})$, and $g_{1,0}(p^H; \frac{1}{2})$ at p^L and p^H respectively, satisfying

$$g_{0,1}(p^L; \frac{1}{2}) < g_{0,0}(p^L; \frac{1}{2}) < g_{1,1}(p^H; \frac{1}{2}) > g_{1,0}(p^H; \frac{1}{2})$$

from Lemmas 7, 8 (since $q_0 < \frac{1}{2}$), and 7 respectively. Thus, $\widehat{\mathcal{R}}_A$ has a unique global maximum at $p = p^H$ (HPP).

Case 1.2 $0 < q_0 < 2p^L$ and $\frac{2p^L}{1-q_0} < \pi \leq \frac{1-2p^L}{1-q_0}$. In this case,

$$\begin{aligned}\mu_{0,0} &= \frac{\pi(1-q_0)}{2} > p^L; \\ \mu_{1,0} &= \frac{1+\pi(1-q_0)}{2} \leq \frac{1+(1-2p^L)}{2} = 1 - p^L = p^H, \\ &\quad \text{equality holding if and only if } \pi = \frac{1-2p^L}{1-q_0}.\end{aligned}$$

Hence, the local suprema of $\widehat{\mathcal{R}}_A$ over the intervals $[p^L, \frac{1}{2}]$ and $(p^H, 1]$ are $g_{0,0}(\mu_{0,0}; \frac{1}{2})$, and $g_{1,0}(p^H; \frac{1}{2})$ at $\mu_{0,0} > p^L$ and p^H respectively. But we already

know that $g_{1,0}(p^H; \frac{1}{2}) < g_{1,1}(p^H; \frac{1}{2})$ from Lemma 7, and that $g_{0,0}(\mu_{0,0}; \frac{1}{2}) > g_{0,0}(p^L; \frac{1}{2})$, $\mu_{0,0}$ being the global maximizer of $g_{0,0}(\cdot; \frac{1}{2})$. This implies that $g_{0,1}(p^L; \frac{1}{2}) < g_{0,0}(\mu_{0,0}; \frac{1}{2})$ since $g_{0,1}(p^L; \frac{1}{2}) < g_{0,0}(p^L; \frac{1}{2})$ from Lemma 7. We can thus conclude that the global maximum of $\widehat{\mathcal{R}}_A$ is either $g_{0,0}(\mu_{0,0}; \frac{1}{2})$ at $\mu_{0,0} > p^L$ or $g_{1,1}(p^H; \frac{1}{2})$ at p^H . Let us define

$$\begin{aligned}\mathcal{G}(\pi; q_0, f) &\triangleq g_{1,1}(p^H; \frac{1}{2}) - g_{0,0}(\mu_{0,0}; \frac{1}{2}) \\ &= f(p^H) - f(\frac{\pi(1-q_0)}{2}) + f'(p^H) (p^L - \frac{\pi q_0}{2}) \\ &\quad - f'(\frac{1}{2}) (1 - \frac{\pi}{2}), \\ &\quad \text{plugging in } \mu_{0,0} = \frac{\pi(1-q_0)}{2}, \mu_{1,1} = 1 - \frac{\pi q_0}{2}.\end{aligned}$$

$$\begin{aligned}\text{Then, } \mathcal{G}''(\pi; q_0, f) &= -f''(\frac{\pi(1-q_0)}{2}) (\frac{1-q_0}{2})^2 < 0, \\ &\quad \text{since } f''(r) > 0, 0 \leq r \leq 1;\end{aligned}$$

$$\begin{aligned}\mathcal{G}(\pi; q_0, f) \Big|_{\pi = \frac{2p^L}{1-q_0}} &= [f(p^H) - f(p^L)] + f'(p^H) \left(p^L - \left(\frac{q_0}{1-q_0} \right) p^L \right) \\ &\quad - f'(\frac{1}{2}) \left(1 - \frac{p^L}{1-q_0} \right) \\ &= [(1 - 2p^L)f'(\frac{1}{2})] + f'(p^H)p^L \left(2 - \frac{1}{1-q_0} \right) \\ &\quad - f'(\frac{1}{2}) \left(1 - \frac{p^L}{1-q_0} \right), \\ &\quad \text{from symmetry condition 3.3,} \\ &= p^L (f'(p^H) - f'(\frac{1}{2})) \left(2 - \frac{1}{1-q_0} \right) > 0,\end{aligned}$$

because $2 > \frac{1}{1-q_0}$ for $q_0 < \frac{1}{2}$, and $f'(p^H) > f'(\frac{1}{2})$ due to the increasing monotonicity of $f'(\cdot)$ (since $p^H > \frac{1}{2}$). Moreover, since $p^H = 1 - p^L$,

$$\begin{aligned}
\mathcal{G}(\pi; q_0, f) \Big|_{\pi = \frac{1-2p^L}{1-q_0}} &= f(1-p^L) - f(\tfrac{1}{2} - p^L) \\
&+ f'(p^H) \left(p^L - \left(\frac{q_0}{1-q_0} \right) \left(\tfrac{1}{2} - p^L \right) \right) \\
&- f'(\tfrac{1}{2}) \left(1 - \frac{\frac{1}{2} - p^L}{1-q_0} \right) \\
&= f(1-p^L) - [f(\tfrac{1}{2} + p^L) - 2p^L f'(\tfrac{1}{2})] \\
&+ f'(p^H) \left(\tfrac{1}{2} - \frac{\frac{1}{2} - p^L}{1-q_0} \right) - f'(\tfrac{1}{2}) \left(1 - \frac{\frac{1}{2} - p^L}{1-q_0} \right), \\
&\text{using the symmetry condition (3.3) again,} \\
&> [f(1-p^L) - f(\tfrac{1}{2} + p^L)] + f'(\tfrac{1}{2}) \left(\tfrac{1}{2} - \frac{\frac{1}{2} - p^L}{1-q_0} \right) \\
&- f'(\tfrac{1}{2}) \left(1 - 2p^L - \frac{\frac{1}{2} - p^L}{1-q_0} \right), \\
&\text{since } \tfrac{1}{2} > \frac{\frac{1}{2} - p^L}{1-q_0} \text{ for } q_0 < 2p^L, f'(p^H) > f'(\tfrac{1}{2}), \\
&> [f'(\tfrac{1}{2} + p^L) (\tfrac{1}{2} - 2p^L)] - f'(\tfrac{1}{2}) (\tfrac{1}{2} - 2p^L), \\
&\text{from the strict convexity of } f(\cdot), \\
&= 2 \left(\tfrac{1}{4} - p^L \right) (f'(\tfrac{1}{2} + p^L) - f'(\tfrac{1}{2})) \\
&> 0 \quad \text{for } p^L < \tfrac{1}{4},
\end{aligned}$$

due to the increasing monotonicity of $f'(\cdot)$.

Thus, $\mathcal{G}(\pi; q_0, f)$ is a continuous, strictly concave function over $[0, 1]$ with positive values at both $\pi = \frac{2p^L}{1-q_0}$ and $\pi = \frac{1-2p^L}{1-q_0}$, hence for $\pi \in (\frac{2p^L}{1-q_0}, \frac{1-2p^L}{1-q_0}]$,

$$\mathcal{G}(\pi; q_0, f) > 0 \implies g_{1,1}(p^H; \tfrac{1}{2}) > g_{0,0}(\mu_{0,0}; \tfrac{1}{2}).$$

Thus finally, we conclude that, for this case too, $\widehat{\mathcal{R}}_A$ has a unique global maximum at $p = p^H$ (HPP).

Case 1.3 $0 < q_0 < 2p^L$ and $\frac{1-2p^L}{1-q_0} < \pi < 1$. In this case,

$$\begin{aligned}\mu_{0,0} &= \frac{\pi(1-q_0)}{2} > \frac{1}{2} - p^L > p^L, \quad \text{since } p^L < \frac{1}{4}; \\ \mu_{1,0} &= \frac{1+\pi(1-q_0)}{2} > \frac{1+(1-2p^L)}{2} = 1 - p^L = p^H.\end{aligned}$$

Hence, the local suprema of $\widehat{\mathcal{R}}_A$ over the intervals $[p^L, \frac{1}{2}]$ and $(p^H, 1]$ are $g_{0,0}(\mu_{0,0}; \frac{1}{2})$, and $g_{1,0}(\mu_{1,0}; \frac{1}{2})$ at $\mu_{0,0} > p^L$ and $\mu_{1,0} > p^H$ respectively. But, as in *Case 1.2*, $g_{0,1}(p^L; \frac{1}{2}) < g_{0,0}(\mu_{0,0}; \frac{1}{2})$; also, since

$$\mu_{1,0} + \mu_{0,0} = \frac{1}{2} + \pi(1 - q_0) = 1 + 2(\frac{1}{4} - p^L) > 1 \text{ for } p^L < \frac{1}{4}, \pi > \frac{1-2p^L}{1-q_0},$$

Lemma 9 tells us that $g_{0,0}(\mu_{0,0}; \frac{1}{2}) < g_{1,0}(\mu_{1,0}; \frac{1}{2})$.

All these inequalities involving the local suprema of $\widehat{\mathcal{R}}_A$ over the four sub-intervals under consideration indicate that the global maximum of $\widehat{\mathcal{R}}_A$ over $(0, 1)$ is either $g_{1,1}(p^H; \frac{1}{2})$ at p^H or $g_{1,0}(\mu_{1,0}; \frac{1}{2})$ at $\mu_{1,0} > p^H$. By Lemma 10,

$$g_{1,1}(p^H; \frac{1}{2}) \gtrless g_{1,0}(\mu_{1,0}; \frac{1}{2}) \iff \pi \gtrless \pi_H^*(q_0),$$

where $\pi_H^*(q_0)$ is the unique zero in $(0, 1)$ of the continuous function $\mathcal{F}_H(\pi; q_0, f)$, defined in (B.9). We already know that $\lim_{\pi \nearrow 1} \mathcal{F}_H(\pi; q_0, f) > 0$; also, note that

$$\begin{aligned}\mathcal{F}_H(\pi; q_0, f) \Big|_{\pi = \frac{1-2p^L}{1-q_0}} &= f(1 - p^L) - f(p^H) \\ &\quad - f'(p^H) \left(p^L - \left(\frac{q_0}{1-q_0} \right) \left(\frac{1}{2} - p^L \right) \right) \\ &\quad + f'(\frac{1}{2}) \left(\frac{1}{2} - \frac{\frac{1}{2} - p^L}{1-q_0} \right) \\ &= - \left(\frac{1}{2} - \frac{\frac{1}{2} - p^L}{1-q_0} \right) (f'(p^H) - f'(\frac{1}{2})) \\ &\quad \text{since } 1 - p^L = p^H, \\ &< 0\end{aligned}$$

since $\frac{1}{2} - \frac{\frac{1}{2} - p^L}{1-q_0} > 0$ for $1 < \frac{2p^L}{q_0}$, and $f'(p^H) > f'(\frac{1}{2})$ due to the increasing monotonicity of $f'(\cdot)$. This implies that $\frac{1-2p^L}{1-q_0} < \pi_H^*(q_0) < 1$. Hence, $\widehat{\mathcal{R}}_A$ has a unique global maximum at $p = p^H$ (HPP) for $\frac{1-2p^L}{1-q_0} < \pi < \pi_H^*(q_0)$, a

unique global maximum at $p = \mu_{1,0} > p^H$ (LPP) for $\pi_H^*(q_0) < \pi < 1$, and two equivalent global maxima at $p = p^H$ and $p = \mu_{1,0}$ for $\pi = \pi_H^*(q_0)$.

In particular, in case of LMSR, for which $f(r) = r \ln r + (1-r) \ln(1-r)$ and $p^H = 4/5$, (B.9) simplifies to

$$\mathcal{F}_H(\pi; q_0, f) = \ln(25(1+x)^{1+x}(1-x)^{1-x}/64) - x \left(\frac{q_0}{1-q_0} \right) \ln 4,$$

where $x = \pi(1-q_0)$. Hence, for LMSR, $\pi_H^*(q_0) = x_L^*(q_0)/(1-q_0)$, where $x_L^*(q_0)$ is the unique root of the fixed-point equation $x = \frac{\ln(64/(25(1+x)^{1+x}(1-x)^{1-x}))}{\left(\frac{q_0}{1-q_0}\right) \ln 4}$, $x_L^*(q_0) \in (1-2p^L, 1-q_0) = (3/5, 1-q_0)$ since $\pi_H^*(q_0) \in \left(\frac{1-2p^L}{1-q_0}, 1\right)$.

Case 2 $q_0 = 2p^L$ In this case, since $\pi < 1$ and $p^L < \frac{1}{4}$,

$$\begin{aligned} \mu_{0,1} &= \frac{1-\pi q_0}{2} = \frac{1}{2} - \pi p^L > \frac{1}{2} - p^L > p^L; \\ \mu_{1,0} &= \frac{1+\pi(1-q_0)}{2} = \frac{1}{2} + \pi \left(\frac{1}{2} - p^L \right) < \frac{1}{2} + \frac{1}{2} - p^L = 1 - p^L = p^H; \\ \mu_{1,1} &= 1 - \frac{\pi q_0}{2} = 1 - \pi p^L > 1 - p^L = p^H. \end{aligned}$$

Hence, the local suprema of $\widehat{\mathcal{R}}_A$ over the intervals $[0, p^L)$, $(\frac{1}{2}, p^H]$, and $(p^H, 1]$ are $g_{0,1}(p^L; \frac{1}{2})$, $g_{1,1}(p^H; \frac{1}{2})$, and $g_{1,0}(p^H; \frac{1}{2})$ respectively, at p^L, p^H, p^H , satisfying

$$g_{1,1}(p^H; \frac{1}{2}) > g_{1,0}(p^H; \frac{1}{2}) > g_{0,1}(p^L; \frac{1}{2}),$$

from Lemmas 7 and 8 (since $q_0 < \frac{1}{2}$) respectively. This makes it impossible for the global maximum of $\widehat{\mathcal{R}}_A$ to lie in either $[0, p^L)$ or $(p^H, 1]$.

Now, for $\pi \leq \frac{2p^L}{1-2p^L}$, $\mu_{0,0} = \frac{\pi(1-2p^L)}{2} \leq p^L$ so that the local maximum of $\widehat{\mathcal{R}}_A$ over $[p^L, \frac{1}{2}]$ is $g_{0,0}(p^L; \frac{1}{2})$ at p^L . But, $g_{0,0}(p^L; \frac{1}{2}) < g_{1,1}(p^H; \frac{1}{2})$ from Lemma 8 (since $q_0 < \frac{1}{2}$). For $\pi > \frac{2p^L}{1-2p^L}$, $\mu_{0,0} > p^L$ so that the local maximum of $\widehat{\mathcal{R}}_A$ over $[p^L, \frac{1}{2}]$ is $g_{0,0}(\mu_{0,0}; \frac{1}{2})$ at $\mu_{0,0} > p^L$. In that case, we need to consider the difference $(g_{1,1}(p^H; \frac{1}{2}) - g_{0,0}(\mu_{0,0}; \frac{1}{2}))|_{q_0=2p^L} = \mathcal{G}(\pi; 2p^L, f)$ (see *Case 1.2*). We can proceed exactly as in *Case 1.2* to show that $\mathcal{G}(\pi; 2p^L, f)$ is a continuous, strictly concave function of π with $\mathcal{G}''(\pi; 2p^L, f)|_{\pi=\frac{2p^L}{1-2p^L}} > 0$ and $\mathcal{G}''(\pi; 2p^L, f)|_{\pi=1} > 0$ (note

that $\frac{1-2p^L}{1-q_0} = 1$ for $q_0 = 2p^L$; hence $\mathcal{G}(\pi; 2p^L, f) > 0 \forall \pi \left(\frac{2p^L}{1-2p^L}, 1 \right)$, implying $g_{0,0}(p^L; \frac{1}{2}) < g_{1,1}(p^H; \frac{1}{2})$, again.

Thus, for $0 < \pi < 1$, $\widehat{\mathcal{R}}_A$ has a unique global maximum at $p = p^H$ (HPP).

Case 3 $2p^L < q_0 < \frac{1}{2}$ In this case, $\frac{2p^L}{1-q_0} < \frac{2p^L}{q_0} < 1 < \frac{1-2p^L}{1-q_0} < \frac{1-2p^L}{q_0}$.

Case 3.1 $2p^L < q_0 < \frac{1}{2}$ and $0 < \pi \leq \frac{2p^L}{q_0}$. In this case,

$$\begin{aligned}\mu_{0,1} &= \frac{1-\pi q_0}{2} \geq \frac{1}{2} - p^L > p^L \quad \text{since } p^L < \frac{1}{4}; \\ \mu_{1,1} &= 1 - \frac{\pi q_0}{2} \geq 1 - p^L = p^H.\end{aligned}$$

Hence, the local suprema of $\widehat{\mathcal{R}}_A$ over the intervals $[0, p^L)$ and $(\frac{1}{2}, p^H]$ are $g_{0,1}(p^L; \frac{1}{2})$ and $g_{1,1}(p^H; \frac{1}{2})$ at p^L and p^H respectively.

Now, for $0 < \pi \leq \frac{2p^L}{1-q_0}$,

$$\begin{aligned}\mu_{0,0} &= \frac{\pi(1-q_0)}{2} \leq p^L; \\ \mu_{1,0} &= \frac{1+\pi(1-q_0)}{2} \leq \frac{1}{2} + p^L = \frac{3}{2} - p^H < p^H \quad \text{since } p^H > \frac{3}{4}.\end{aligned}$$

Hence, the local suprema of $\widehat{\mathcal{R}}_A$ over the intervals $[p^L, \frac{1}{2}]$ and $(p^H, 1]$ are $g_{0,1}(p^L; \frac{1}{2})$ and $g_{1,0}(p^H; \frac{1}{2})$ at p^L and p^H respectively, satisfying

$$g_{0,1}(p^L; \frac{1}{2}) < g_{0,0}(p^L; \frac{1}{2}) < g_{1,1}(p^H; \frac{1}{2}) > g_{1,0}(p^H; \frac{1}{2})$$

from Lemmas 7, 8 (since $q_0 < \frac{1}{2}$), and 7 respectively. Thus, $\widehat{\mathcal{R}}_A$ has a unique global maximum at $p = p^H$ (HPP).

Again, for $\frac{2p^L}{1-q_0} < \pi \leq \frac{2p^L}{q_0}$,

$$\begin{aligned}\mu_{0,0} &= \frac{\pi(1-q_0)}{2} > p^L; \\ \mu_{1,0} &= \frac{1+\pi(1-q_0)}{2} < \frac{1+(1-2p^L)}{2} = 1 - p^L = p^H \quad \text{since } \pi < \frac{1-2p^L}{1-q_0}.\end{aligned}$$

Hence, the local suprema of $\widehat{\mathcal{R}}_A$ over the intervals $[p^L, \frac{1}{2}]$ and $(p^H, 1]$ are $g_{0,0}(\mu_{0,0}; \frac{1}{2})$ and $g_{1,0}(p^H; \frac{1}{2})$ at $\mu_{0,0} > p^L$ and p^H respectively. But $g_{1,0}(p^H; \frac{1}{2}) <$

$g_{1,1}(p^H; \frac{1}{2})$ from , and since $\mu_{0,0}$ is the global maximum of $g_{0,0}(\cdot; \frac{1}{2})$, we also have $g_{0,0}(\mu_{0,0}; \frac{1}{2}) > g_{0,0}(p^L; \frac{1}{2}) > g_{0,1}(p^L; \frac{1}{2})$, the last inequality following from Lemma 7. Thus, the contention for the global maximum of $\widehat{\mathcal{R}}_A$ is between $g_{0,0}(\mu_{0,0}; \frac{1}{2})$ at $\mu_{0,0} > p^L$ and $g_{1,1}(p^H; \frac{1}{2})$ at p^H .

Again, as in *Case 1.2* and *Case 2*, we have to consider the function $\mathcal{G}(\pi; q_0, f) = g_{1,1}(p^H; \frac{1}{2}) - g_{0,0}(\mu_{0,0}; \frac{1}{2})$ for $q_0 \in (2p^L, \frac{1}{2})$. We already know that it is a continuous, strictly concave function with a positive value at $\pi = \frac{2p^L}{1-q_0}$ for $0 < q_0 < \frac{1}{2}$. Also, proceeding as in *Case 2*, we can show that

$$\begin{aligned} \mathcal{G}(\pi; q_0, f) \Big|_{\pi = \frac{2p^L}{q_0}} &> \frac{2p^L}{q_0} \left(\frac{1}{2} - q_0 \right) \left(f'(1 - \mu_{0,0}) - f'(\frac{1}{2}) \right), \quad \mu_{0,0} = \frac{p^L(1-q_0)}{q_0}, \\ &> 0, \quad \text{since } 1 - \mu_{0,0} > \frac{1}{2}, \quad q_0 < \frac{1}{2}. \end{aligned}$$

Hence, for $\frac{2p^L}{1-q_0} < \pi \leq \frac{2p^L}{q_0}$, $\mathcal{G}(\pi; q_0, f) > 0$, i.e. $g_{1,1}(p^H; \frac{1}{2}) > g_{0,0}(\mu_{0,0}; \frac{1}{2})$, so that in this case, too, $\widehat{\mathcal{R}}_A$ has a unique global maximum at $p = p^H$ (HPP).

Case 3.2 $2p^L < q_0 < \frac{1}{2}$ and $\frac{2p^L}{q_0} < \pi < 1$. In this case,

$$\begin{aligned} \mu_{0,0} &= \frac{\pi(1-q_0)}{2} > p^L \quad \text{since } \pi > \frac{2p^L}{1-q_0}; \\ \mu_{0,1} &> \mu_{0,0} > p^L \quad \text{from inequalities (B.5);} \\ \mu_{1,1} &= 1 - \frac{\pi q_0}{2} < 1 - p^L = p^H; \\ \mu_{1,0} &< \mu_{1,1} < p^H \quad \text{from inequalities (B.5).} \end{aligned}$$

Hence, the local suprema of $\widehat{\mathcal{R}}_A$ over the intervals $[0, p^L)$, $[p^L, \frac{1}{2}]$, $(\frac{1}{2}, p^H]$ and $(p^H, 1]$ are $g_{0,1}(p^L; \frac{1}{2})$, $g_{0,0}(\mu_{0,0}; \frac{1}{2})$, $g_{1,1}(\mu_{1,1}; \frac{1}{2})$, and $g_{1,0}(p^H; \frac{1}{2})$ at p^L , $\mu_{0,0} > p^L$,

$\mu_{1,1} < p^H$ and p^H respectively, satisfying

$$\begin{aligned}
g_{0,1}(p^L; \tfrac{1}{2}) &< g_{0,0}(p^L; \tfrac{1}{2}) \quad \text{from Lemma 7,} \\
&< g_{0,0}(\mu_{0,0}; \tfrac{1}{2}) \quad \text{since } \mu_{0,0} \text{ is the maximizer of } g_{0,0}(\cdot; \tfrac{1}{2}), \\
&< g_{1,1}(\mu_{1,1}; \tfrac{1}{2}) \\
&\quad \text{from Lemma 9 since } \mu_{1,1} + \mu_{0,0} = 1 + \pi(\tfrac{1}{2} - q_0) \\
&\quad \text{for } q_0 < \tfrac{1}{2}; \\
g_{1,0}(p^H; \tfrac{1}{2}) &< g_{1,1}(p^H; \tfrac{1}{2}) \quad \text{from Lemma 7,} \\
&< g_{1,1}(\mu_{1,1}; \tfrac{1}{2}) \quad \text{since } \mu_{1,1} \text{ is the maximizer of } g_{1,1}(\cdot; \tfrac{1}{2}).
\end{aligned}$$

Thus, $\widehat{\mathcal{R}}_A$ has a unique global maximum at $p = \mu_{1,1}$ (LPP).

Case 4 $q_0 = \frac{1}{2}$ Here, we just need the following two sub-cases.

Case 4.1 $q_0 = \frac{1}{2}$ and $0 < \pi \leq 4p^L$.

$$\begin{aligned}
\mu_{0,0} &= \frac{\pi}{4} \leq p^L; \\
\mu_{0,1} &= \tfrac{1}{2} - \frac{\pi}{4} \geq \tfrac{1}{2} - p^L > p^L \quad \text{since } p^L < \tfrac{1}{4}; \\
\mu_{1,0} &= \tfrac{1}{2} + \frac{\pi}{4} \leq \tfrac{1}{2} + p^L = \tfrac{3}{2} - p^H < p^H \quad \text{since } p^H > \tfrac{3}{4}; \\
\mu_{1,1} &= 1 - \frac{\pi}{4} \geq 1 - p^L = p^H.
\end{aligned}$$

Hence, the local suprema of $\widehat{\mathcal{R}}_A$ over the intervals $[0, p^L)$, $[p^L, \frac{1}{2}]$, $(\frac{1}{2}, p^H]$ and $(p^H, 1]$ are $g_{0,1}(p^L; \frac{1}{2})$, $g_{0,0}(p^L; \frac{1}{2})$, $g_{1,1}(p^H; \frac{1}{2})$, and $g_{1,0}(p^H; \frac{1}{2})$ at p^L , p^L , p^H and p^H respectively, satisfying

$$g_{0,1}(p^L; \tfrac{1}{2}) < g_{0,0}(p^L; \tfrac{1}{2}) = g_{1,1}(p^H; \tfrac{1}{2}) > g_{1,0}(p^H; \tfrac{1}{2})$$

from Lemmas 7, 8 (since $q_0 = \frac{1}{2}$), and 7 respectively. Hence, $\widehat{\mathcal{R}}_A$ has two equivalent global maxima at $p = p^L$ and $p = p^H$ (HPP).

Case 4.2 $q_0 = \frac{1}{2}$ and $4p^L < \pi < 1$.

$$\begin{aligned}\mu_{0,0} &= \frac{\pi}{4} > p^L; \\ \mu_{0,1} &> \mu_{0,0} > p^L; \\ \mu_{1,1} &= 1 - \frac{\pi}{4} < 1 - p^L = p^H; \\ \mu_{1,0} &< \mu_{1,1} < p^H.\end{aligned}$$

Hence, the local suprema of $\widehat{\mathcal{R}}_A$ over the intervals $[0, p^L)$, $[p^L, \frac{1}{2}]$, $(\frac{1}{2}, p^H]$ and $(p^H, 1]$ are $g_{0,1}(p^L; \frac{1}{2})$, $g_{0,0}(\mu_{0,0}; \frac{1}{2})$, $g_{1,1}(\mu_{1,1}; \frac{1}{2})$, and $g_{1,0}(p^H; \frac{1}{2})$ at p^L , $\mu_{0,0} > p^L$, $\mu_{1,1} < p^H$ and p^H respectively. This situation is similar to in *Case 3.2* with the only difference that $g_{0,0}(\mu_{0,0}; \frac{1}{2}) = g_{1,1}(\mu_{1,1}; \frac{1}{2})$ from Lemma 9 since $\mu_{1,1} + \mu_{0,0} = 1$ for $q_0 = \frac{1}{2}$. Hence, $\widehat{\mathcal{R}}_A$ has two equivalent global maxima at $p = \mu_{0,0}$ and $p = \mu_{1,1}$ (LPP).

Case 5 $\frac{1}{2} < q_0 < 1$ By symmetry, the analysis is similar to that for $0 < q_0 < \frac{1}{2}$ (*Cases 1, 2, and 3* combined), and is thus omitted.

B.2.2 Analysis for symmetric well-behaved MSRs with $p^L = \frac{1}{4}$, $p^H = \frac{3}{4}$, e.g. QMSR

Case 1 $0 < q_0 < \frac{1}{2}$ Note that $0 < \frac{1}{2(1-q_0)} < 1$ for these values of q_0 . We need to consider the following sub-cases:

Case 1.1 $0 < q_0 < \frac{1}{2}$ and $0 < \pi < \frac{1}{2(1-q_0)}$. In this case,

$$\begin{aligned}\mu_{0,1} &> \frac{1 - \frac{1}{2}}{2} = \frac{1}{4} = p^L \quad \text{since } \pi < 1, q_0 < \frac{1}{2}; \\ \mu_{0,0} &= \frac{1}{4} \cdot (2\pi(1 - q_0)) < \frac{1}{4} = p^L; \\ \mu_{1,1} &> 1 - \frac{1 - \frac{1}{2}}{2} = \frac{3}{4} = p^H; \\ \mu_{1,0} &= \frac{1}{4} \cdot (2 + 2\pi(1 - q_0)) < \frac{1}{4} \cdot (2 + 1) = \frac{3}{4} = p^H.\end{aligned}$$

Hence, the local suprema of $\widehat{\mathcal{R}}_A$ over the intervals $[0, p^L)$, $[p^L, \frac{1}{2}]$, $(\frac{1}{2}, p^H]$, $(p^H, 1]$ are $g_{0,1}(p^L; \frac{1}{2})$, $g_{0,0}(p^L; \frac{1}{2})$, $g_{1,1}(p^H; \frac{1}{2})$, $g_{1,0}(p^H; \frac{1}{2})$ respectively at p^L, p^L, p^H, p^H , satisfying

$$g_{0,1}(p^L; \frac{1}{2}) < g_{0,0}(p^L; \frac{1}{2}) < g_{1,1}(p^H; \frac{1}{2}) > g_{1,0}(p^H; \frac{1}{2}),$$

due to Lemmas 7, 8 (since $q_0 < \frac{1}{2}$), and 7 respectively. Hence, $\widehat{\mathcal{R}}_A$ has a unique global maximum at $p = p^H$.

Case 1.2 $0 < q_0 < \frac{1}{2}$ and $\pi = \frac{1}{2(1-q_0)}$. In this case,

$$\begin{aligned}\mu_{0,0} &= \frac{1}{2} = \frac{1}{4} = p^L; \\ \mu_{0,1} &> \mu_{0,0} = p^L \quad \text{from (B.5);} \\ \mu_{1,0} &= \frac{1+\frac{1}{2}}{2} = \frac{3}{4} = p^H; \\ \mu_{1,1} &> \mu_{1,0} = p^H \quad \text{from (B.5).}\end{aligned}$$

Hence, the local suprema of $\widehat{\mathcal{R}}_A$ over the intervals $[0, p^L)$, $[p^L, \frac{1}{2}]$, $(\frac{1}{2}, p^H]$, $(p^H, 1]$ are $g_{0,1}(p^L; \frac{1}{2})$, $g_{0,0}(p^L; \frac{1}{2})$, $g_{1,1}(p^H; \frac{1}{2})$, $g_{1,0}(p^H; \frac{1}{2})$ respectively at p^L, p^L, p^H, p^H , satisfying

$$g_{0,1}(p^L; \frac{1}{2}) < g_{0,0}(p^L; \frac{1}{2}) < g_{1,1}(p^H; \frac{1}{2}) > g_{1,0}(p^H; \frac{1}{2}),$$

due to Lemmas 7, 8 (since $q_0 < \frac{1}{2}$), and 7 respectively. Hence, $\widehat{\mathcal{R}}_A$ has a unique global maximum at $p = p^H = \mu_{1,0}$, making the two equilibrium domains indistinguishable.

Case 1.3 $0 < q_0 < \frac{1}{2}$ and $\frac{1}{2(1-q_0)} < \pi < 1$. In this case,

$$\begin{aligned}\mu_{0,0} &= \frac{1}{4} \cdot (2\pi(1 - q_0)) > \frac{1}{4} = p^L; \\ \mu_{0,1} &> \mu_{0,0} > p^L; \\ \mu_{1,0} &= \frac{1}{4} \cdot (2 + 2\pi(1 - q_0)) > \frac{1}{4} \cdot (2 + 1) = \frac{3}{4} = p^H; \\ \mu_{1,1} &> \mu_{1,0} > p^H.\end{aligned}$$

Hence, the local suprema of $\widehat{\mathcal{R}}_A$ over the intervals $[0, p^L)$, $[p^L, \frac{1}{2}]$, $(\frac{1}{2}, p^H]$, $(p^H, 1]$ are $g_{0,1}(p^L; \frac{1}{2})$, $g_{0,0}(\mu_{0,0}; \frac{1}{2})$, $g_{1,1}(p^H; \frac{1}{2})$, $g_{1,0}(\mu_{1,0}; \frac{1}{2})$ respectively at $p^L, \mu_{0,0}, p^H, \mu_{1,0}$.

Note that $\mu_{0,0} + \mu_{1,0} = \frac{1}{2} + \pi(1 - q_0) > 1$ since $\pi > \frac{1}{2(1-q_0)}$. So, from Lemma 9,

$$\begin{aligned} g_{1,0}(\mu_{1,0}; \tfrac{1}{2}) &> g_{0,0}(\mu_{0,0}; \tfrac{1}{2}) \\ &> g_{0,0}(p^L; \tfrac{1}{2}) \quad \text{since } \mu_{0,0} \text{ is the maximizer of } g_{0,0}(\cdot; \tfrac{1}{2}), \\ &> g_{0,1}(p^L; \tfrac{1}{2}), \quad \text{from Lemma 7.} \end{aligned}$$

Thus, neither $g_{0,0}(\mu_{0,0}; \frac{1}{2})$ at $\mu_{0,0}$ nor $g_{0,1}(p^L; \frac{1}{2})$ at p^L can be the global maximum – it is either $g_{1,0}(\mu_{1,0}; \frac{1}{2})$ at $\mu_{1,0}$ or $g_{1,1}(p^H; \frac{1}{2})$ at p^H . Recall, from Lemma 10, that

$$g_{1,1}(p^H; \tfrac{1}{2}) \geq g_{1,0}(\mu_{1,0}; \tfrac{1}{2}) \iff \pi \leq \pi_H^*(q_0),$$

where $\pi_H^*(q_0)$ is the unique zero in $(0, 1)$ of the continuous function $\mathcal{F}_H(\pi; q_0, f)$, defined in (B.9). We already know that $\lim_{\pi \nearrow 1} \mathcal{F}_H(\pi; q_0, f) > 0$, in general. Now, for $p^H = \frac{3}{4}$ and $0 < q_0 < \frac{1}{2}$, we obtain

$$\begin{aligned} \lim_{\pi \searrow \frac{1}{2(1-q_0)}} \mathcal{F}_H(\pi; q_0, f) &= f(\tfrac{3}{4}) - f(\tfrac{3}{4}) - f'(\tfrac{3}{4}) \left(\tfrac{1}{4} - \frac{q_0}{4(1-q_0)} \right) \\ &\quad + f'(\tfrac{1}{2}) \left(\tfrac{1}{2} - \frac{1}{4(1-q_0)} \right) \\ &< -f'(\tfrac{1}{2}) \left(\tfrac{1}{4} - \frac{q_0}{4(1-q_0)} \right) \\ &\quad + f'(\tfrac{1}{2}) \left(\tfrac{1}{2} - \frac{1}{4(1-q_0)} \right) \\ &= 0, \end{aligned}$$

the inequality following from the fact that $f'(\frac{3}{4}) > f'(\frac{1}{2})$ due to the increasing monotonicity of $f'(\cdot)$, and $\frac{1}{4} - \frac{q_0}{4(1-q_0)} > 0$ for $q_0 < \frac{1}{2}$.

Thus, $\widehat{\mathcal{R}}_A$ has a unique global maximum at p^H for $\frac{1}{2(1-q_0)} < \pi < \pi_H^*(q_0)$, a unique global maximum at $\mu_{1,0}$ for $\pi_H^*(q_0) < \pi < 1$, and two equivalent global maxima at p^H and $\mu_{1,0} > p^H$ for $\pi = \pi_H^*(q_0)$.

In particular, for QMSR, we have $f(r) = r^2$, $0 \leq r \leq 1$ in addition to $p^H = \frac{3}{4}$, so that

$$\mathcal{F}_H(\pi; q_0, f) = \frac{(1-q_0)^2}{4} \left[\pi^2 + \frac{q_0}{(1-q_0)^2} \pi - \frac{3}{4(1-q_0)^2} \right] \quad \text{on simplification,}$$

which is a quadratic polynomial in π . Using the quadratic formula and discarding the inadmissible negative root of $\mathcal{F}_H(\pi; q_0, f) = 0$, we obtain

$$\pi_H^*(q_0) = \frac{-\frac{q_0}{(1-q_0)^2} + \frac{1}{1-q_0} \sqrt{\left(\frac{q_0}{1-q_0}\right)^2 + 3}}{2} = \frac{(\sqrt{3+v^2}-v)/2}{1-q_0},$$

where $v = \frac{q_0}{1-q_0}$.

Case 2 $q_0 = \frac{1}{2}$ In this case, for $0 < \pi < 1$,

$$\begin{aligned}\mu_{0,1} &= \frac{1}{2} - \frac{\pi}{4} > \frac{1}{2} - \frac{1}{4} = \frac{1}{4} = p^L; \\ \mu_{0,0} &= \frac{\pi}{4} < \frac{1}{4} = p^L; \\ \mu_{1,1} &> 1 - \frac{1}{2} = \frac{1}{2} = p^H; \\ \mu_{1,0} &= \frac{1}{2} + \frac{\pi}{4} < \frac{1}{2} + \frac{1}{4} = \frac{3}{4} = p^H.\end{aligned}$$

Hence, the local suprema of $\widehat{\mathcal{R}}_A$ over the intervals $[0, p^L)$, $[p^L, \frac{1}{2}]$, $(\frac{1}{2}, p^H]$, $(p^H, 1]$ are $g_{0,1}(p^L; \frac{1}{2})$, $g_{0,0}(p^L; \frac{1}{2})$, $g_{1,1}(p^H; \frac{1}{2})$, $g_{1,0}(p^H; \frac{1}{2})$ respectively at p^L, p^L, p^H, p^H , satisfying

$$g_{0,1}(p^L; \frac{1}{2}) < g_{0,0}(p^L; \frac{1}{2}) = g_{1,1}(p^H; \frac{1}{2}) > g_{1,0}(p^H; \frac{1}{2}),$$

due to Lemmas 7, 8 (since $q_0 = \frac{1}{2}$), and 7 respectively. Hence, for any $\pi \in (0, 1)$, $\widehat{\mathcal{R}}_A$ has two equivalent global maxima at $p = p^L$ and $p = p^H$.

Case 3 $\frac{1}{2} < q_0 < 1$ By symmetry, the analysis is similar to that for $0 < q_0 < \frac{1}{2}$ (*Case 1*), and is thus omitted.

B.2.3 Analysis for symmetric well-behaved MSRs with $\frac{1}{4} < p^L < \frac{1}{2}$, $\frac{1}{2} < p^H < \frac{3}{4}$, e.g. SMSR

Note that, for such scoring rules, $0 < 1 - 2p^L < \frac{1}{2} < 2p^L < 1$.

Case 1 $0 < q_0 < 1 - 2p^L$ In this case, $\frac{2p^L}{q_0} > \frac{1-2p^L}{q_0} > 1$; also, $0 < \frac{1-2p^L}{1-q_0} < \frac{2p^L}{1-q_0} < 1$.

Case 1.1 $0 < q_0 < 1 - 2p^L$ and $0 < \pi \leq \frac{1-2p^L}{1-q_0}$. Then,

$$\begin{aligned}\mu_{0,0} &= \frac{\pi(1-q_0)}{2} < p^L \quad \text{since } \pi \leq \frac{1-2p^L}{1-q_0} < \frac{2p^L}{1-q_0}; \\ \mu_{0,1} &= \frac{1-\pi q_0}{2} > \frac{1-(1-2p^L)}{2} = p^L \quad \text{since } \pi < 1 < \frac{1-2p^L}{q_0}; \\ \mu_{1,0} &= \frac{1+\pi(1-q_0)}{2} \leq \frac{1+(1-2p^L)}{2} = 1 - p^L = p^H, \\ &\quad \text{equality holding only if } \pi = \frac{1-2p^L}{1-q_0}; \\ \mu_{1,1} &= 1 - \frac{\pi q_0}{2} > 1 - p^L \quad \text{since } \pi < 1 < \frac{2p^L}{q_0}, \\ &= p^H \quad \text{since } p^H = 1 - p^L.\end{aligned}$$

Hence, the local suprema of $\widehat{\mathcal{R}}_A$ over the intervals $[0, p^L)$, $[p^L, \frac{1}{2}]$, $(\frac{1}{2}, p^H]$, $(p^H, 1]$ are $g_{0,1}(p^L; \frac{1}{2})$, $g_{0,0}(p^L; \frac{1}{2})$, $g_{1,1}(p^H; \frac{1}{2})$, $g_{1,0}(p^H; \frac{1}{2})$ respectively at p^L, p^L, p^H, p^H , satisfying

$$g_{0,1}(p^L; \frac{1}{2}) < g_{0,0}(p^L; \frac{1}{2}) < g_{1,1}(p^H; \frac{1}{2}) > g_{1,0}(p^H; \frac{1}{2}),$$

due to Lemmas 7, 8 (since $q_0 < \frac{1}{2}$), and 7 respectively. Hence, $\widehat{\mathcal{R}}_A$ has a unique global maximum at $p = p^H$.

Case 1.2 $0 < q_0 < 1 - 2p^L$ and $\frac{1-2p^L}{1-q_0} < \pi \leq \frac{2p^L}{1-q_0}$. Then,

$$\begin{aligned}\mu_{0,0} &= \frac{\pi(1-q_0)}{2} \leq p^L, \quad \text{equality holding only if } \pi = \frac{2p^L}{1-q_0}; \\ \mu_{0,1} &= \frac{1-\pi q_0}{2} > \frac{1-(1-2p^L)}{2} = p^L \quad \text{since } \pi < 1 < \frac{1-2p^L}{q_0}; \\ \mu_{1,0} &= \frac{1+\pi(1-q_0)}{2} > \frac{1+(1-2p^L)}{2} = 1 - p^L = p^H; \\ \mu_{1,1} &> \mu_{1,0} = p^H \quad \text{by (B.5)}.\end{aligned}$$

Hence, the local suprema of $\widehat{\mathcal{R}}_A$ over the intervals $[0, p^L)$, $[p^L, \frac{1}{2}]$, $(\frac{1}{2}, p^H]$, $(p^H, 1]$ are $g_{0,1}(p^L; \frac{1}{2})$, $g_{0,0}(p^L; \frac{1}{2})$, $g_{1,1}(p^H; \frac{1}{2})$, $g_{1,0}(\mu_{1,0}; \frac{1}{2})$ respectively at $p^L, p^L, p^H, \mu_{1,0}$, satisfying

$$g_{0,1}(p^L; \frac{1}{2}) < g_{0,0}(p^L; \frac{1}{2}) < g_{1,1}(p^H; \frac{1}{2}),$$

due to Lemmas 7, 8 (since $q_0 < \frac{1}{2}$), and 7 respectively.

Hence, the global maximum of $\widehat{\mathcal{R}}_A$ must be either $g_{1,1}(p^H; \frac{1}{2})$ at p^H or $g_{1,0}(\mu_{1,0}; \frac{1}{2})$ at $\mu_{1,0} > p^H$. By Lemma 10 ,

$$g_{1,1}(p^H; \frac{1}{2}) \gtrless g_{1,0}(\mu_{1,0}; \frac{1}{2}) \iff \pi \lesseqgtr \pi_H^*(q_0),$$

where $\pi_H^*(q_0)$ is the unique zero in $(0, 1)$ of the continuous function $\mathcal{F}_H(\pi; q_0, f)$, defined in (B.9). Note that

$$\begin{aligned} \lim_{\pi \searrow \frac{1-2p^L}{1-q_0}} \mathcal{F}_H(\pi; q_0, f) &= f\left(\frac{1+(1-2p^L)}{2}\right) - f(p^H) \\ &\quad - f'(p^H) \left(p^L - \left(\frac{q_0}{1-q_0} \right) \left(\frac{1}{2} - p^L \right) \right) \\ &\quad + f'\left(\frac{1}{2}\right) \left(\frac{1}{2} - \frac{\frac{1}{2}-p^L}{1-q_0} \right) \\ &= -f'(p^H) \left(p^L + \left(1 - \frac{1}{1-q_0} \right) \left(\frac{1}{2} - p^L \right) \right) \\ &\quad + f'\left(\frac{1}{2}\right) \left(\frac{1}{2} - \frac{\frac{1}{2}-p^L}{1-q_0} \right), \\ &\quad \text{since } \frac{1+1-2p^L}{2} = 1 - p^L = p^H. \\ &= - \left(\frac{1}{2} - \frac{\frac{1}{2}-p^L}{1-q_0} \right) (f'(p^H) - f'(\frac{1}{2})) \\ &< 0 \end{aligned}$$

since $\frac{1}{2} - \frac{\frac{1}{2}-p^L}{1-q_0} > 0$ for $1 < \frac{2p^L}{q_0}$, and $f'(p^H) > f'(\frac{1}{2})$ due to the increasing monotonicity of $f'(\cdot)$. From this, we can conclude that $\pi_H^*(q_0) > \frac{1-2p^L}{1-q_0}$. However, we have not been able to prove (or disprove) that $\pi_H^*(q_0) < \frac{2p^L}{1-q_0}$ for a general $f(\cdot)$ satisfying $p^L > \frac{1}{4}$ and for $q_0 < 1 - 2p^L$. Our conjecture is that the relative magnitudes of $\pi_H^*(q_0)$ and $\frac{2p^L}{1-q_0}$ depend on the form of $f(\cdot)$ and the exact value of p^L , and hence, so does the value of $\pi \in \left(\frac{1-2p^L}{1-q_0}, \frac{2p^L}{1-q_0} \right]$ beyond which the global maximizer of $\widehat{\mathcal{R}}_A$ switches from p^H to $\mu_{1,0} > p^H$, if it does switch at all. But the inability to determine the conditions under which $\pi_H^*(q_0)$ lies below or above $\frac{2p^L}{1-q_0}$ does not in any way detract from our analysis, as we shall show in conjunction with the findings in *Case 1.3* below.

Case 1.3 $0 < q_0 < 1 - 2p^L$ and $\frac{2p^L}{1-q_0} < \pi < 1$. Then,

$$\begin{aligned}\mu_{0,0} &= \frac{\pi(1-q_0)}{2} > p^L \quad \text{since } \pi > \frac{2p^L}{1-q_0}; \\ \mu_{0,1} &> \mu_{0,0} > p^L \quad \text{by (B.5);} \\ \mu_{1,0} &= \frac{1+\pi(1-q_0)}{2} > \frac{1+(1-2p^L)}{2} = 1 - p^L = p^H, \\ &\quad \text{since } \pi > \frac{2p^L}{1-q_0} > \frac{1-2p^L}{1-q_0}; \\ \mu_{1,1} &> \mu_{1,0} = p^H \quad \text{by (B.5).}\end{aligned}$$

Hence, the local suprema of $\widehat{\mathcal{R}}_A$ over the intervals $[0, p^L]$, $[p^L, \frac{1}{2}]$, $(\frac{1}{2}, p^H]$, $(p^H, 1]$ are $g_{0,1}(p^L; \frac{1}{2})$, $g_{0,0}(\mu_{0,0}; \frac{1}{2})$, $g_{1,1}(p^H; \frac{1}{2})$, $g_{1,0}(\mu_{1,0}; \frac{1}{2})$ respectively at p^L , $\mu_{0,0}$, p^H , $\mu_{1,0}$. Now, for $\pi > \frac{2p^L}{1-q_0}$ where $p^L > \frac{1}{4}$, we have $\mu_{1,0} + \mu_{0,0} = \frac{1}{2} + \pi(1 - q_0) > \frac{1}{2} + 2p^L > 1$. Thus, from Lemma 9,

$$\begin{aligned}g_{1,0}(\mu_{1,0}; \tfrac{1}{2}) &> g_{0,0}(\mu_{0,0}; \tfrac{1}{2}) \\ &> g_{0,0}(p^L; \tfrac{1}{2}), \quad \text{since } \mu_{0,0} \text{ is the maximizer of } g_{0,0}(\cdot; \tfrac{1}{2}), \\ &> g_{0,1}(p^L; \tfrac{1}{2}), \quad \text{from Lemma 7.}\end{aligned}$$

Hence, the global maximum of $\widehat{\mathcal{R}}_A$ must be either $g_{1,1}(p^H; \frac{1}{2})$ at p^H or $g_{1,0}(\mu_{1,0}; \frac{1}{2})$ at $\mu_{1,0} > p^H$, putting us in a position similar to that in *Case 1.2*.

Combining the last two cases, we can conclude that $\widehat{\mathcal{R}}_A$ has a unique global maximum at p^H for $\frac{1-2p^L}{1-q_0} < \pi < \pi_H^*(q_0)$, a unique global maximum at $\mu_{1,0} > p^H$ for $\frac{1-2p^L}{1-q_0} < \pi < \pi_H^*(q_0)$, and two equivalent global maxima at p^H and $\mu_{1,0} > p^H$ for $\pi = \pi_H^*(q_0)$.

In particular, for SMSR, for which $f(r) = \sqrt{r^2 + (1-r)^2}$, $0 \leq r \leq 1$, and $p^H = \frac{1}{2} \left(1 + \sqrt{\frac{\sqrt{2}-1}{2}} \right)$, we have

$$\mathcal{F}_H(\pi; q_0, f) = \sqrt{\frac{1+\pi^2(1-q_0)^2}{2}} - \frac{1+\gamma_S-\gamma_S\pi q_0}{\sqrt{2(1+\gamma_S^2)}}, \quad \gamma_S = \sqrt{\frac{\sqrt{2}-1}{2}}.$$

By simple algebra, we can show that the equation $\mathcal{F}_H(\pi; q_0, f) = 0$ reduces to the quadratic equation

$$\pi^2 \left(1 - \left(\frac{\gamma_S^2}{1+\gamma_S^2} \right) v^2 \right) + \frac{2\gamma_S(1+\gamma_S)}{1+\gamma_S^2} \cdot \frac{v}{1-q_0} \cdot \pi - \frac{2\gamma_S}{1+\gamma_S^2} \cdot \frac{v}{(1-q_0)^2} = 0,$$

where $v = q_0/(1-q_0)$. Using the quadratic formula, discarding the inadmissible negative root, and simplifying,

$$\begin{aligned}
\pi_H^*(q_0) &= \frac{x_S^*(q_0)}{1-q_0}, \quad \text{where} \\
x_S^*(q_0) &= \frac{\sqrt{\gamma_S(1+\gamma_S^2)(2+\gamma_S v^2) - \gamma_S(1+\gamma_S)v}}{(1+\gamma_S^2) - \gamma_S^2 v^2} \\
&= \frac{\sqrt{K_1 + K_2 v^2 - K_3 v}}{1 - K_2 v^2}, \quad \text{where} \\
K_1 &= \frac{2\gamma_S}{(1+\gamma_S^2)} = 2\sqrt{2} \left(\sqrt{2} - 1 \right)^{3/2} \approx 0.7540; \\
K_2 &= \frac{\gamma_S^2}{1+\gamma_S^2} = 3 - 2\sqrt{2} \approx 0.1716; \\
K_3 &= \frac{\gamma_S(1+\gamma_S)}{1+\gamma_S^2} = 3 - 2\sqrt{2} + \sqrt{2} \left(\sqrt{2} - 1 \right)^{3/2} \approx 0.5486.
\end{aligned}$$

Case 2 $q_0 = 1 - 2p^L$ In this case, since $\pi < 1$,

$$\begin{aligned}
\mu_{0,0} &= \frac{\pi(1-q_0)}{2} = \pi p^L < p^L; \\
\mu_{0,1} &= \frac{1-\pi q_0}{2} = \frac{1-\pi(1-2p^L)}{2} > \frac{1-(1-2p^L)}{2} = p^L; \\
\mu_{1,0} &= \frac{1+\pi(1-q_0)}{2} = \frac{1}{2} + \pi p^L; \\
\mu_{1,1} &= 1 - \frac{\pi q_0}{2} = \frac{1}{2} - \pi \left(\frac{1-2p^L}{2} \right).
\end{aligned}$$

Hence, the local suprema of $\widehat{\mathcal{R}}_A$ over the intervals $[0, p^L)$ and $[p^L, \frac{1}{2}]$ are $g_{0,1}(p^L; \frac{1}{2})$ and $g_{0,0}(p^L; \frac{1}{2})$ respectively, both at p^L , satisfying $g_{0,0}(p^L; \frac{1}{2}) > g_{0,1}(p^L; \frac{1}{2})$ due to Lemma 7. But for determining the local suprema over $(\frac{1}{2}, p^H]$ and $(p^H, 1]$, we need further conditions on π .

Case 2.1 $q_0 = 1 - 2p^L$ and $0 < \pi \leq \frac{1-2p^L}{2p^L}$. Since $p^L > \frac{1}{4}$, we have

$$(1 - 2p^L)^2 = 1 - 4p^L + (2p^L)^2 < (2p^L)^2 \iff \frac{1-2p^L}{2p^L} < \frac{2p^L}{1-2p^L}.$$

Hence,

$$\begin{aligned}\mu_{1,0} &= \frac{1}{2} + \pi p^L \leq \frac{1}{2} + \frac{1-2p^L}{2} = 1 - p^L = p^H, \\ &\text{equality holding if and only if } \pi = \frac{1-2p^L}{2p^L}; \\ \mu_{1,1} &= \frac{1}{2} - \pi \left(\frac{1-2p^L}{2} \right) > 1 - p^L = p^H, \quad \text{since } \pi < \frac{2p^L}{1-2p^L}.\end{aligned}$$

Hence, the local suprema of $\widehat{\mathcal{R}}_A$ over the intervals $(\frac{1}{2}, p^H]$ and $(p^H, 1]$ are $g_{1,1}(p^H; \frac{1}{2})$ and $g_{1,0}(p^H; \frac{1}{2})$ respectively, both at p^H , satisfying $g_{1,1}(p^H; \frac{1}{2}) > g_{1,0}(p^H; \frac{1}{2})$ due to . Moreover, since $q_0 < \frac{1}{2}$, $g_{1,1}(p^H; \frac{1}{2}) > g_{0,0}(p^L; \frac{1}{2})$ from Lemma 8.

Hence, $\widehat{\mathcal{R}}_A$ has a unique global maximum at $p = p^H$.

Case 2.2 $q_0 = 1 - 2p^L$ and $\frac{1-2p^L}{2p^L} < \pi < 1$. Then,

$$\begin{aligned}\mu_{1,0} &= \frac{1}{2} + \pi p^L > \frac{1}{2} + \frac{1-2p^L}{2} = 1 - p^L = p^H; \\ \mu_{1,1} &> \mu_{1,0} > p^H.\end{aligned}$$

We already know that $g_{1,1}(p^H; \frac{1}{2}) > g_{0,0}(p^L; \frac{1}{2}) > g_{0,1}(p^L; \frac{1}{2})$, and from Lemma 10 ,

$$g_{1,1}(p^H; \frac{1}{2}) \geq g_{1,0}(\mu_{1,0}; \frac{1}{2}) \iff \pi \leq \pi_H^*(q_0),$$

where $\pi_H^*(q_0)$ is the unique zero in $(0, 1)$ of the continuous function $\mathcal{F}_H(\pi; q_0, f)$, defined in (B.9), with $\lim_{\pi \nearrow 1} \mathcal{F}_H(\pi; q_0, f) > 0$ for any $q_0 \in (0, 1)$. Now,

$$\begin{aligned}\lim_{\pi \searrow \frac{1-2p^L}{2p^L}} \mathcal{F}_H(\pi; 1 - 2p^L, f) &= f(1 - p^L) - f(p^H) - f'(p^H) \left(p^L - \frac{(1-2p^L)^2}{4p^L} \right) \\ &\quad + f'(\frac{1}{2}) \left(\frac{1}{2} - \frac{1-2p^L}{4p^L} \right) \\ &= - \left(\frac{p^L - \frac{1}{4}}{p^L} \right) (f'(p^H) - f'(\frac{1}{2})), \\ &\quad \text{since } 1 - p^L = p^H. \\ &< 0,\end{aligned}$$

since $p^L > \frac{1}{4}$ and $f'(p^H) > f'(\frac{1}{2})$ from the increasing monotonicity of $f'(\cdot)$. Thus, we can conclude that $\frac{1-2p^L}{2p^L} < \pi_H^*(q_0) < 1$.

Hence, $\widehat{\mathcal{R}}_A$ has a unique global maximum at p^H for $\frac{1-2p^L}{2p^L} < \pi < \pi_H^*(q_0)$, a unique global maximum at $\mu_{1,0}$ for $\frac{1-2p^L}{2p^L} < \pi < \pi_H^*(q_0)$, and two equivalent global maxima at p^H and $\mu_{1,0} > p^H$ for $\pi = \pi_H^*(q_0)$. For SMSR, we can proceed as in *Case 1.3* to obtain

$$\begin{aligned}\pi_H^*(q_0)\big|_{q_0=1-2p^L} &= \frac{x_S^*(q_0)\big|_{q_0=1-2p^L}}{2p^L}, \quad \text{where} \\ x_S^*(q_0)\big|_{q_0=1-2p^L} &= \frac{\sqrt{\gamma_S(1+\gamma_S^2)(2+\gamma_S v^2) - \gamma_S(1+\gamma_S)v}}{(1+\gamma_S^2) - \gamma_S^2 v^2} \bigg|_{v=2p^L/(1-2p^L)} \\ &= \gamma_S, \quad \text{since } p^L = \frac{1-\gamma_S}{2}. \\ \Rightarrow \pi_H^*(q_0)\big|_{q_0=1-2p^L} &= \frac{\gamma_S}{1-\gamma_S} \approx 0.8352.\end{aligned}$$

Case 3 $1 - 2p^L < q_0 < \frac{1}{2}$ In this case, $q_0 < 2p^L$ since $p^L > \frac{1}{4}$, hence $0 < \frac{1-2p^L}{1-q_0} < 1 < \frac{2p^L}{1-q_0}$; also, $0 < \frac{1-2p^L}{q_0} < 1 < \frac{2p^L}{q_0}$. Moreover, since $q_0 < \frac{1}{2}$, we have $\frac{1-2p^L}{1-q_0} < \frac{1-2p^L}{q_0}$.

Case 3.1 $1 - 2p^L < q_0 < \frac{1}{2}$ and $0 < \pi \leq \frac{1-2p^L}{1-q_0}$. Then,

$$\begin{aligned}\mu_{0,0} &= \frac{\pi(1-q_0)}{2} < p^L \quad \text{since } \pi < 1 < \frac{2p^L}{1-q_0}; \\ \mu_{0,1} &= \frac{1-\pi q_0}{2} > \frac{1-(1-2p^L)}{2} = p^L \quad \text{since } \pi \leq \frac{1-2p^L}{1-q_0} < \frac{1-2p^L}{q_0}; \\ \mu_{1,0} &= \frac{1+\pi(1-q_0)}{2} \leq \frac{1+(1-2p^L)}{2} = 1 - p^L = p^H, \\ &\quad \text{equality holding only if } \pi = \frac{1-2p^L}{1-q_0}; \\ \mu_{1,1} &= 1 - \frac{\pi q_0}{2} > 1 - p^L \quad \text{since } \pi < 1 < \frac{2p^L}{q_0}, \\ &= p^H \quad \text{since } p^H = 1 - p^L.\end{aligned}$$

Hence, arguing exactly as in *Case 1.1*, $\widehat{\mathcal{R}}_A$ has a unique global maximum at $p = p^H$.

Case 3.2 $1 - 2p^L < q_0 < \frac{1}{2}$ and $\frac{1-2p^L}{1-q_0} < \pi < 1$. Then,

$$\begin{aligned}\mu_{0,0} &= \frac{\pi(1-q_0)}{2} < p^L \quad \text{since } \pi < 1 < \frac{2p^L}{1-q_0}; \\ \mu_{1,0} &= \frac{1+\pi(1-q_0)}{2} > \frac{1+(1-2p^L)}{2} = 1 - p^L = p^H; \\ \mu_{1,1} &> \mu_{1,0} > p^H.\end{aligned}$$

Hence, the local suprema of $\widehat{\mathcal{R}}_A$ over the intervals $[p^L, \frac{1}{2}]$, $(\frac{1}{2}, p^H]$, $(p^H, 1]$ are $g_{0,0}(p^L; \frac{1}{2})$, $g_{1,1}(p^H; \frac{1}{2})$, $g_{1,0}(\mu_{1,0}; \frac{1}{2})$ respectively at p^L , p^H and $\mu_{1,0} > p^H$, satisfying $g_{1,1}(p^H; \frac{1}{2}) > g_{0,0}(p^L; \frac{1}{2})$ by Lemma 8 since $q_0 = \frac{1}{2}$. Note that if $\frac{1-2p^L}{1-q_0} < \pi \leq \frac{1-2p^L}{q_0}$, then $\mu_{0,1} = \frac{1-\pi q_0}{2} \geq p^L$, so that the local supremum of $\widehat{\mathcal{R}}_A$ over $[0, p^L]$ is $g_{0,1}(p^L; \frac{1}{2})$ at p^L but $g_{0,1}(p^L; \frac{1}{2}) < g_{0,0}(p^L; \frac{1}{2})$ by Lemma 7; and if $\frac{1-2p^L}{q_0} < \pi < 1$, then $\mu_{0,1} < p^L$, so that the local maximum of $\widehat{\mathcal{R}}_A$ over $[0, p^L]$ is $g_{0,1}(\mu_{0,1}; \frac{1}{2})$ at $\mu_{0,1}$ but then $g_{0,1}(\mu_{0,1}; \frac{1}{2}) < g_{1,0}(\mu_{1,0}; \frac{1}{2})$ since $\mu_{0,1} + \mu_{1,0} = 1 + \pi(\frac{1}{2} - q_0) > 1$, by Lemma 9. Thus, as in *Case 1.2*, $\widehat{\mathcal{R}}_A$ has a unique global maximum at p^H for $\frac{1-2p^L}{1-q_0} < \pi < \pi_H^*(q_0)$, a unique global maximum at $\mu_{1,0} > p^H$ for $\frac{1-2p^L}{1-q_0} < \pi < \pi_H^*(q_0)$, and two equivalent global maxima at p^H and $\mu_{1,0} > p^H$ for $\pi = \pi_H^*(q_0)$, where $\pi_H^*(q_0)$ has the same meaning as in *Case 1.2*, and the same expression for SMSR as specified after *Case 1.3*.

Case 4 $q_0 = \frac{1}{2}$ Note that $0 < 2 - 4p^L < 1$ for $\frac{1}{4} < p^L < \frac{1}{2}$, and $p^L + p^H = 1$ by Proposition 2.

Case 4.1 $q_0 = \frac{1}{2}$ and $0 < \pi \leq 2 - 4p^L = 4p^H - 2$. In this case,

$$\begin{aligned} \mu_{0,1} &= \frac{1}{2} - \frac{\pi}{4} \geq \frac{1}{2} - \frac{1}{2} + p^L = p^L, \\ &\text{equality holding if and only if } \pi = 2 - 4p^L; \\ \mu_{0,0} &= \frac{\pi}{4} < \frac{1}{4} < p^L \quad \text{for } \pi < 1; \\ \mu_{1,1} &= 1 - \frac{\pi}{4} > 1 - \frac{1}{4} = \frac{3}{4} > p^H \quad \text{for } \pi < 1; \\ \mu_{1,0} &= \frac{1}{2} + \frac{\pi}{4} \leq \frac{1}{2} + p^H - \frac{1}{2} = p^H, \\ &\text{equality holding if and only if } \pi = 4p^H - 2. \end{aligned}$$

Hence, the local suprema of $\widehat{\mathcal{R}}_A$ over the intervals $[0, p^L]$, $[p^L, \frac{1}{2}]$, $(\frac{1}{2}, p^H]$, $(p^H, 1]$ are $g_{0,1}(p^L; \frac{1}{2})$, $g_{0,0}(p^L; \frac{1}{2})$, $g_{1,1}(p^H; \frac{1}{2})$, $g_{1,0}(p^H; \frac{1}{2})$ respectively at p^L , p^L , p^H , p^H , satisfying

$$g_{0,1}(p^L; \frac{1}{2}) < g_{0,0}(p^L; \frac{1}{2}) = g_{1,1}(p^H; \frac{1}{2}) > g_{1,0}(p^H; \frac{1}{2}),$$

due to Lemmas 7, 8 (since $q_0 = \frac{1}{2}$), and 7 respectively. Hence, $\widehat{\mathcal{R}}_A$ has two equivalent global maxima at $p = p^L$ and $p = p^H$.

Case 4.2 $q_0 = \frac{1}{2}$ and $\pi = 2 - 4p^L = 4p^H - 2$.

$$\begin{aligned}\mu_{0,1} &= p^L; \\ \mu_{0,0} &< \mu_{0,1} = p^L \quad \text{by (B.5);} \\ \mu_{1,0} &= p^H; \\ \mu_{1,1} &> \mu_{1,0} = p^H \quad \text{by (B.5).}\end{aligned}$$

Hence, the local suprema of $\widehat{\mathcal{R}}_A$ over the intervals $[0, p^L]$, $[p^L, \frac{1}{2}]$, $(\frac{1}{2}, p^H]$, $(p^H, 1]$ are $g_{0,1}(p^L; \frac{1}{2})$, $g_{0,0}(p^L; \frac{1}{2})$, $g_{1,1}(p^H; \frac{1}{2})$, $g_{1,0}(p^H; \frac{1}{2})$ respectively at p^L, p^L, p^H, p^H , satisfying

$$g_{0,1}(p^L; \frac{1}{2}) < g_{0,0}(p^L; \frac{1}{2}) = g_{1,1}(p^H; \frac{1}{2}) > g_{1,0}(p^H; \frac{1}{2}),$$

due to Lemmas 7, 8 (since $q_0 = \frac{1}{2}$), and 7 respectively. Hence, $\widehat{\mathcal{R}}_A$ has two equivalent global maxima at $p = p^L = \mu_{0,1}$ and $p = p^H = \mu_{1,0}$.

Case 4.3 $q_0 = \frac{1}{2}$ and $2 - 4p^L = 4p^H - 2 < \pi < 1$. In this case,

$$\begin{aligned}\mu_{0,1} &= \frac{1}{2} - \frac{\pi}{4} < \frac{1}{2} - \frac{1}{2} + p^L = p^L \quad \text{for } \pi > 2 - 4p^L; \\ \mu_{0,0} &< \mu_{0,1} < p^L \quad \text{by (B.5);} \\ \mu_{1,0} &= \frac{1}{2} + \frac{\pi}{4} > \frac{1}{2} + p^H - \frac{1}{2} = p^H \quad \text{for } \pi < 4p^H - 2; \\ \mu_{1,1} &> \mu_{1,0} > p^H.\end{aligned}$$

Hence, the local suprema of $\widehat{\mathcal{R}}_A$ over the intervals $[0, p^L]$, $[p^L, \frac{1}{2}]$, $(\frac{1}{2}, p^H]$, $(p^H, 1]$ are $g_{0,1}(\mu_{0,1}; \frac{1}{2})$, $g_{0,0}(p^L; \frac{1}{2})$, $g_{1,1}(p^H; \frac{1}{2})$, $g_{1,0}(\mu_{1,0}; \frac{1}{2})$ respectively at $\mu_{0,1}, p^L, p^H, \mu_{1,0}$, satisfying

$$g_{0,1}(\mu_{0,1}; \frac{1}{2}) = g_{1,0}(\mu_{1,0}; \frac{1}{2}); \quad g_{0,0}(p^L; \frac{1}{2}) = g_{1,1}(p^H; \frac{1}{2}). \quad (\text{B.10})$$

due to Lemma 8 since $q_0 = \frac{1}{2}$. Again, as in *Case 1.3* and *Case 2.2*, we now need to consider the equation $\mathcal{F}_H(\pi; q_0, f)|_{q_0=\frac{1}{2}} = 0$; note that $\lim_{\pi \nearrow 1} \mathcal{F}_H(\pi; \frac{1}{2}, f) > 0$, and

$$\lim_{\pi \searrow 4p^H - 2} \mathcal{F}_H(\pi; \frac{1}{2}, f) = -2 \left(f'(p^H) - f'(\frac{1}{2}) \right) \left(\frac{3}{4} - p^H \right) < 0,$$

since $p^H < \frac{3}{4}$, and $f'(p^H) > f'(\frac{1}{2})$ due to the increasing monotonicity of $f'(\cdot)$. Hence, the root of this equation $\pi^* \in (2 - 4p^L, 1)$, so we can conclude that $\hat{\mathcal{R}}_A$ has two equivalent global maxima at p^L, p^H for $\pi < \pi^*$, two equivalent global maxima at $\mu_{0,1} < p^L, \mu_{1,0} > p^H$ for $\pi > \pi^*$, and four equivalent global maxima at $\mu_{0,1}, p^L, p^H, \mu_{1,0}$ for $\pi = \pi^*$.

For SMSR, as before,

$$\begin{aligned} \pi^*\left(\frac{1}{2}\right) &= \frac{x_S^*(q_0)\big|_{q_0=\frac{1}{2}}}{\frac{1}{2}}, \quad \text{where} \\ x_S^*(q_0)\big|_{q_0=\frac{1}{2}} &= \frac{\sqrt{\gamma_S(1+\gamma_S^2)(2+\gamma_S v^2)} - \gamma_S(1+\gamma_S)v}{(1+\gamma_S^2) - \gamma_S^2 v^2} \bigg|_{v=1} \\ \Rightarrow \quad \pi_H^*\left(\frac{1}{2}\right) &= 2 \left(\sqrt{\gamma_S(1+\gamma_S^2)(2+\gamma_S)} - \gamma_S(1+\gamma_S) \right) \\ &\approx 0.9983 \triangleq \pi_S^*, \quad \text{since } \gamma_S = \sqrt{\frac{\sqrt{2}-1}{2}}. \end{aligned}$$

Case 5 $\frac{1}{2} < q_0 < 1$ By symmetry, the analysis is similar to that for $0 < q_0 < \frac{1}{2}$ (*Cases 1, 2, and 3* combined), and is thus omitted.

Appendix C

Proofs of results in Chapter 4

We now provide the proofs of Proposition 3 and Theorem 7 from Chapter 4; please refer to the chapter for the meanings of all symbols used here.

Throughout this appendix, we assume that $\sigma_z = 1$, μ_0 , and $0 < \rho_0 \leq 1$; the latter implies that only the update equations (4.2) and (4.3) apply.

Let $\eta_t \triangleq |\mu_{t+1} - \mu_t| = \frac{\rho_t^2 \sqrt{2/\pi}}{\sqrt{1+\rho_t^2}}$, the *step-size* of the mean update. The proofs of both main results rely on the following lemma which shows that $\eta_t = \Theta(1/t)$.

Lemma 11. *For $\rho_0 \leq 1$, $\eta_t < \frac{c_1}{t}$ and*

$$\eta_t > \begin{cases} c_2 \rho_0^2 & t \leq \lfloor 1/\rho_0^2 \rfloor, \\ \frac{c_2}{t} & t \geq \lfloor 1/\rho_0^2 \rfloor + 1, \end{cases}$$

where $c_1 = \sqrt{\frac{2}{\pi}} \left(\frac{\pi + (\pi-2)\rho_0^2}{2} \right)$ and $c_2 = \frac{1}{2\sqrt{\pi}}$.

Proof. **To deduce the upper bound on η_t :** Note that, from the recursion in (4.3), it is clear that $\rho_t^2 > 0 \forall t \geq 1$ since $\rho^2 > 0$; also,

$$\rho_{t+1}^2 = \rho_t^2 \left[\frac{1 + \rho_t^2 (1 - 2/\pi)}{1 + \rho_t^2} \right] < \rho_t^2 \quad \forall t \geq 0,$$

since $0 < 1 - 2/\pi < 1$, $\rho_t^2 > 0$. Hence,

$$0 < \rho_t^2 < \rho_0^2 \leq 1, \quad \forall t \geq 1. \tag{C.1}$$

Note that $\eta_t = \frac{\rho_t^2 \sqrt{2/\pi}}{\sqrt{1+\rho_t^2}} < \rho_t^2 \sqrt{\frac{2}{\pi}} \forall t$, since $\sqrt{1+\rho_t^2} > 1$. Hence, to obtain the required upper bound c_1/t where $c_1 = \sqrt{\frac{2}{\pi}} \left(\frac{\pi+(\pi-2)\rho_0^2}{2} \right)$, it suffices to show that

$$\rho_t^2 < \left(\frac{\pi + (\pi - 2)\rho_0^2}{2} \right) \cdot \frac{1}{t}, \quad \forall t \geq 1. \quad (\text{induction hypothesis})$$

The inequality is satisfied for $t = 1$ since

$$\rho_1^2 = \frac{\rho_0^2}{1+\rho_0^2} \left(\frac{\pi + (\pi - 2)\rho_0^2}{\pi} \right) < \left(\frac{\pi + (\pi - 2)\rho_0^2}{2} \right) \cdot \frac{1}{1}, \quad (\text{base case})$$

since $\frac{\rho_0^2}{1+\rho_0^2} < 1$, $\pi > 2$. Now, assuming that the hypothesis holds for some $t \geq 1$, we have, from (4.3) again,

$$\begin{aligned} \rho_{t+1}^2 &= \rho_t^2 \left[\frac{1 + \rho_t^2 (1 - 2/\pi)}{1 + \rho_t^2 (1 - 2/\pi) + (2/\pi)\rho_t^2} \right] \\ &= \rho_t^2 / \left(1 + \frac{2/\pi}{1 + \rho_t^2 (1 - 2/\pi)} \cdot \rho_t^2 \right) \\ &= 1 / \left(\frac{1}{\rho_t^2} + \frac{2}{\pi + (\pi - 2)\rho_t^2} \right) \\ &< 1 / \left(\frac{1}{\rho_t^2} + \frac{2}{\pi + (\pi - 2)\rho_0^2} \right) \quad \text{from (C.1),} \\ &< 1 / \left(t \cdot \frac{2}{\pi + (\pi - 2)\rho_0^2} + \frac{2}{\pi + (\pi - 2)\rho_0^2} \right), \quad \text{from the induction hypothesis,} \\ &= \left(\frac{\pi + (\pi - 2)\rho_0^2}{2} \right) \cdot \frac{1}{t+1}. \quad (\text{inductive step}) \end{aligned}$$

To deduce the lower bound on η_t : First note that, by definition,

$$\eta_t > \sqrt{\frac{2}{\pi}} \cdot \frac{\rho_t^2}{\sqrt{2}} = \frac{\rho_t^2}{\sqrt{\pi}} \quad \forall t \geq 1, \quad (\text{C.2})$$

since $\rho_t^2 < \rho_0^2 \leq 1 \forall t \geq 1$. Next, we will establish that

$$\rho_t^2 > \frac{\rho_0^2}{1 + t\rho_0^2} \quad \forall t \geq 1. \quad (\text{induction hypothesis})$$

The result holds for $t = 1$ since

$$\rho_1^2 > \frac{\rho_0^2}{1 + \rho_0^2}, \quad \text{since } \frac{\pi + (\pi - 2)\rho_0^2}{\pi} > 0. \quad (\text{base case})$$

Returning to the recursion (4.3) and assuming the hypothesis to be true for some $t \geq 1$, we see that,

$$\begin{aligned} \rho_{t+1}^2 &> \frac{\rho_t^2}{1 + \rho_t^2} = \frac{1}{1/\rho_t^2 + 1} \\ &> \frac{1}{(1 + t\rho_0^2)/\rho_0^2 + 1} = \frac{\rho_0^2}{1 + t\rho_0^2 + \rho_0^2} \\ &= \frac{\rho_0^2}{1 + (t+1)\rho_0^2}. \quad (\text{inductive step}) \end{aligned}$$

Now, if $1 \leq t \leq \lfloor 1/\rho_0^2 \rfloor \leq 1/\rho_0^2 \in [1, \infty)$, then $t\rho_0^2 \leq 1$ so that $\rho_t^2 > \frac{\rho_0^2}{2}$; but, if $t \geq \lfloor 1/\rho_0^2 \rfloor + 1 > 1/\rho_0^2 > 1$, then $\rho_t^2 > \frac{1}{1/\rho_0^2 + t} > \frac{1}{t+1} = \frac{1}{2t}$. Combining these results with inequality (C.2), we obtain the desired lower bound. \square

We shall now define symbols used in proofs.

Let $\Delta\mu_t = \mu_{t+1} - \mu_t$; $\hat{p} \triangleq \Phi(\varepsilon)$ where $\varepsilon \gtrsim 0$ is a tolerance parameter, as defined in Chapter 4, and $\Phi(\cdot)$ is the standard normal cumulative distribution function; obviously, $\hat{p} > \Phi(0) = 0.5$.

Also, denote by p_t^+ the probability (from the perspective of someone who knows the true V as well as the distribution of the z_t 's) of the binary signal x_t received by the learner at time-step t is positive, given her current (public) threshold μ_t , i.e.

$$p_t^+ = \Pr[V + z_t \geq \mu_t | z_t \sim \mathcal{N}(0, 1)] = 1 - \Phi(\mu_t - V) = \Phi(V - \mu_t). \quad (\text{C.3})$$

Let $\mathbb{E}[\cdot]$ denote the expectation (from the perspective of someone who knows the true V and the learner's belief updating and threshold-setting heuristics) at time 0 with respect to the uncertainty in $\{z_i\}_{i < t}$, and $\mathbb{E}_{x_t}[\cdot]$ the same with respect to the the uncertainty in the binary signal x_t , given all the relevant available information up to the beginning of epoch t (including μ_t).

Restatement of Proposition 3. *There exist absolute positive constants $C > 0$ and k , $1 \leq k < \pi\sqrt{2} \approx 4.443$ such that, if $t > C/(\rho_0^2 \varepsilon^k)$, then $|V| - |\mathbb{E}[\mu_t]| < \varepsilon$.*

(Heuristic) proof of Proposition 3. We provide the proof for $\{V > 0\}$, that for $\{V < 0\}$ being analogous.

Evidently,

$$\mathbb{E}_{x_t}[\Delta\mu_t] = p_t^+ \eta_t + (1 - p_t^+)(-\eta_t) = (2p_t^+ - 1)\eta_t \quad \forall t,$$

where the value of η_t for every t is predetermined and known right from $t = 0$. If $\xi_t \triangleq \mathbb{E}[\mu_t]$, then

$$\xi_{t+1} = \xi_t + 2(\mathbb{E}[p_t^+ - 0.5]).$$

First note that, as long as $\mu_t < V$ (which is the scenario we are interested in) which implies that $p_t^+ > 0.5$, we have $\xi_t > 0 \quad \forall t \geq 1$ since $\xi_0 = 0$. Moreover, if we wait until $(V - \mu_t)$ is small enough (with a high probability) to use the following first-order approximation based on a Maclaurin series expansion

$$\Phi(V - \mu_t) \approx \Phi(0) + \Phi'(0)(V - \mu_t) \quad \Longleftrightarrow \quad \Phi(V - \mu_t) - 0.5 \approx \frac{V - \mu_t}{\sqrt{2\pi}},$$

then, from (C.3) and using the linearity of expectation,

$$\Delta\xi_t \approx \sqrt{\frac{2}{\pi}}(V - \xi_t)\eta_t\Delta t, \tag{C.4}$$

where $\Delta\xi_t = \xi_{t+1} - \xi_t$, $\Delta t = 1$.

Let t_α denote the number of time-steps such that

$$\max_{t \in \{0, \dots, t_\alpha\}} \mu_t \geq V - \alpha$$

with probability at least $(1 - \delta)$ where $0 < \delta \ll 1$, and α is small enough for the above linear approximation to be reasonable (but α is still larger than and independent of the tolerance parameter ε in the theorem statement). From Theorem 7 in Chapter 4, we already know²⁵

²⁵Although we have not proved this theorem yet, the proof (provided at the end of this appendix) does not assume Proposition 3; hence all our statements are consistent.

that $t_\alpha = O(e^{V/\alpha}/\rho_0^2)$ and is obviously independent of ε . Since the process has already crossed $V - \alpha$ (a value close to V) within t_α and the step-size η_t only becomes progressively smaller, we can say intuitively that μ_t is above $V - \alpha$ with a very high probability for $t > t_\alpha$. Also, for such large values of t , we can use the fact that $\eta_t = \Theta(1/t)$ to get the following approximate difference equation, from (C.4):

$$\Delta \xi_t = \frac{c\sqrt{2/\pi}}{t}(V - \xi_t)\Delta t, \quad \text{for some constant } c, c_2 \leq c \leq c_1.$$

Let t_ε be the time taken by ξ_t to “hit” (i.e. cross for the first time) $V - \varepsilon$. Since we are interested in the case $\varepsilon \rightarrow 0$, we can assume that $t_\varepsilon > t_\alpha$, and the deterministic quantity $\xi_{t_\alpha} = \mathbb{E}[\mu_{t_\alpha}]$ is less than V (otherwise, there is nothing to prove). Then, taking a continuous approximation (a differential equation) to the above difference equation and integrating between the proper limits,

$$\begin{aligned} \int_{\xi_{t_\alpha}}^{V-\varepsilon} \frac{d\xi}{V-\xi} &= c\sqrt{2/\pi} \int_{t_\alpha}^{t_\varepsilon} \frac{dt}{t} \\ \implies \ln\left(\frac{V-\xi_{t_\alpha}}{\varepsilon}\right) &= c\sqrt{2/\pi} \ln\left(\frac{t_\varepsilon}{t_\alpha}\right) \\ \implies t_\varepsilon &= t_\alpha \left(\frac{V-\xi_{t_\alpha}}{\varepsilon}\right)^k \end{aligned}$$

where $k = \frac{1}{c}\sqrt{\frac{\pi}{2}}$. Using the value of c_2 from Lemma 11, we get the upper bound on k stated in the theorem; although c_1 gives us a (loose) lower bound on k that is smaller than 1, we recall from Theorem 6 that we cannot have an asymptotic bound on the convergence time sub-linear in $1/\varepsilon$, hence we must have $k \geq 1$. \square

Restatement of Theorem 7. Fix $0 < \delta < 1$, $0 < \varepsilon < V$, $0 < \rho_0 \leq 1$, and define $\Delta = V - \varepsilon$. There is an absolute constant $C > 0$ such that if $t > T = e^{C(\ln(1/\delta) + \Delta)/\varepsilon}/\rho_0^2$, then with probability at least $1 - \delta$, $\max_{i \leq t} \mu_i > V - \varepsilon$.

Proof of Theorem 7. Given the initial belief distribution $\mathcal{N}(0, \rho_0)$, the value of ρ_t , and hence of η_t , for each $t \geq 0$ is completely determined. Thus, after t time-steps, the learner could attain any one of at most 2^t pre-defined values of μ_t each corresponding to a unique

path of the form $[(0, 0), (\mu_{(1)}, 1), \dots, (\mu_{(t)}, t)]$ in the (μ, t) -space, where $\mu_{(t)}$ denotes *one* of the possible mean beliefs that the learner could have at time t . With this insight, we define a *reinforcement learning setting* in which each such path is a state of the learner.

Define $\mathcal{S} = \{s = [(0, 0), \dots, (\mu_{(t)}, t)] ; \mu_{(t)} < V - \varepsilon\}$. Obviously, for any s in \mathcal{S} , from (C.3),

$$p_t^+ > \hat{p} = \Phi(\varepsilon) > 0.5.$$

After a binary signal is received, μ_t can only move “upward” or “downward” by the amount η_t , so that any state $s \in \mathcal{S}$ can only transition to one of two states which we denote by $[s; (\mu_{(t)} + \eta_t, t + 1)]$ and $[s; (\mu_{(t)} - \eta_t, t + 1)]$ respectively. For a given time-horizon $[0, \tau]$, let us define π as the policy which assigns to any state $s \in \mathcal{S}$ the constant probability \hat{p} of moving upward to the state $[s; (\mu_{(t)} + \eta_t, t + 1)]$, and π' the policy as that which assigns the state-contingent probability p_t^+ to the same upward transition (π' corresponds to our approximate inference algorithm) with the following exceptions: If any transtion results in a state $[(0, 0), \dots, (\mu_{(t)}, t)]$ where $\mu_t > V - \varepsilon$ or $t = \tau$, then under either policy (π or π') the process passes into a *dead state* at the next transtion and remains in that state forever. To each transtion we assign a reward 1 if the transtion results in a state with the final μ_t above $V - \varepsilon$ and 0 otherwise; all transtions to the dead state have zero reward. The state value function $\mathcal{V}_\tau^{\mathbf{p}}(s)$ for state $s \in \mathcal{S}$ under policy $\mathbf{p} \in \{\pi, \pi'\}$ is defined as the the limit, as $\gamma \rightarrow 1$, of the infinite-horizon geometrically discounted, by a factor γ , sum of expected rewards. From the definition of states and rewards, it readily follows that $\mathcal{V}_\tau^{\mathbf{p}}(s)$ is the probability that at least one of the states in the interval $[t, \tau]$ has its last μ -value exceeding $(V - \varepsilon)$, starting from state s under the policy \mathbf{p} :

$$\mathcal{V}_\tau^{\mathbf{p}}(s) = \Pr \left(\left\{ \max_{t \leq i \leq \tau} \mu_{(i)} \geq V - \varepsilon \right\} | \mathbf{p} \right).$$

The following lemma formalizes the notion that the policy π' *dominates* π .

Lemma 12. *For any $s = [(0, 0), \dots, (\mu_{(t)}, t)] \in \mathcal{S}$, where $t < \tau$, $\mathcal{V}_\tau^{\pi'}(s) \geq \mathcal{V}_\tau^\pi(s)$.*

In particular, for the initial state $\varphi = [(0, 0)]$, $\mathcal{V}_\tau^{\pi'}(\varphi) \geq \mathcal{V}_\tau^\pi(\varphi)$.

We will provide the proof of the above lemma after the main proof, but will now focus on deducing a lower bound on $\mathcal{V}_\tau^\pi(\varphi)$ for the dominated process.

$$\begin{aligned}
\mathcal{V}_\tau^\pi(\varphi) &= \Pr \left(\left\{ \max_{i \leq \tau} \mu_{(i)} \geq V - \varepsilon \right\} | \pi \right) \\
&\geq \Pr \left(\mu_{(\tau)} \geq V - \varepsilon | \pi \right) \quad \text{since } \{ \mu_{(\tau)} \geq V - \varepsilon \} \subseteq \{ \max_{i \leq \tau} \mu_{(i)} \geq V - \varepsilon \} \\
&= \Pr \left[\mu_{(\tau)} - \mathbb{E} [\mu_{(\tau)}] \geq -(2\hat{p} - 1) \sum_{t=0}^{\tau-1} \eta_t + \Delta | \pi \right], \quad \text{where } \Delta = V - \varepsilon.
\end{aligned}$$

The last equality follows from the observation that, for any $s \in \mathcal{S}$ under π , $\mathbb{E}[\Delta\mu_t] = (2\hat{p} - 1)\eta_t \forall t \geq 0$ so that $\mu_{(t)}$ is the sum of i.i.d. random variables $\{\Delta\mu_i\}_{i=0}^{t-1}$ and has expectation $(2\hat{p} - 1) \sum_{i=0}^{t-1} \eta_i$.

Using the lower bound on η_t from Lemma 11 and the inequality (which, in turn, follows from the strictly decreasing monotonicity of $1/t$)

$$\sum_{t=t_1}^{t_2} \frac{1}{t} > \int_{t_1}^{t_2+1} \frac{dt}{t} = \ln\left(\frac{t_2+1}{t_1}\right), \tag{C.5}$$

we can show that

$$(2\hat{p} - 1) \sum_{t=0}^{\tau-1} \eta_t - \Delta > 0 \quad \text{for } \tau > \tau' = \left(\left\lfloor \frac{1}{\rho_0^2} \right\rfloor + 1 \right) \exp \left(\frac{2\sqrt{\pi}\Delta}{2\hat{p} - 1} \right). \tag{C.6}$$

Hence, for $\tau > \tau'$, by Hoeffding inequality (see e.g. Boucheron et al. (2004)), we have

$$\Pr [\mu_{(\tau)} \geq V - \varepsilon | \pi] \geq 1 - \delta \quad \text{for } \delta \geq \exp \left[- \frac{2 \left((2\hat{p} - 1) \sum_{t=0}^{\tau-1} \eta_t - \Delta \right)^2}{\sum_{t=0}^{\tau-1} (2\eta_t)^2} \right].$$

Combining this result with Lemma 12, we conclude that the above inequality also holds for our algorithm, *i.e.* $\Pr [\mu_{(\tau)} \geq V - \varepsilon | \pi'] \geq 1 - \delta$.

On rearranging the inequality of interest, we find that, for any given δ , it is sufficient for τ to satisfy

$$(2\hat{p} - 1) \sum_{t=0}^{\tau-1} \eta_t - \Delta \geq \sqrt{2 \ln(1/\delta) \sum_{t=0}^{\tau-1} \eta_t^2}$$

Combining the lower bound on η_t from Lemma 11 with the inequality (C.5) as before, and using result (i) presented below, we can obtain a lower bound on the above L.H.S. that is linear in ε and logarithmic in τ ; moreover, using result (ii) stated below, we can also obtain an upper bound on the above R.H.S. that is independent of ε, τ and sub-linear in $\ln(1/\delta)$. Further, taking (C.6) into consideration (which accounts for the linear dependence on $1/\rho_0^2$ and exponential dependence on Δ), we obtain the desired asymptotic convergence time bound.

(i) From the strict concavity of $\Phi(x)$ for $x \geq 0$,

$$2\hat{p} - 1 = 2[\Phi(\varepsilon) - 0.5] = 2[\Phi(\varepsilon) - \Phi(0)] > 2N(\varepsilon)\varepsilon,$$

where $N(x) = \Phi'(x) = e^{-x^2/2}/\sqrt{2\pi}$ is the standard normal probability density function. Also, since $N(x)$ is strictly decreasing for $x \geq 0$ and it is possible to obtain a small positive constant, say λ , that exceeds all interesting values of ε , we conclude that $2\hat{p} - 1$ is bounded below by the linear (in ε) expression $2N(\lambda)\varepsilon$.

(ii) From Lemma 11, $\eta_t^2 < c_1^2/t^2$, and, for any $t_1 \in \{1, 2, \dots\}$, $t_2 \in \{t_1 + 1, t_1 + 2, \dots\}$,

$$\sum_{t=t_1}^{t_2} \frac{1}{t^2} < \sum_{t=1}^{\infty} \frac{1}{t^2} = \zeta(2) = \frac{\pi^2}{6},$$

where $\zeta(\cdot)$ denotes the Riemann zeta function (Riemann, 1859).

□

Proof of Lemma 12. For an arbitrary $s \in \mathcal{S}$, define

$$\mathcal{V}_+ = \mathcal{V}_\tau^\pi([s; (X_{(t)} + \eta_{t+1}, t + 1)]), \quad \mathcal{V}_- = \mathcal{V}_\tau^\pi([s; (X_{(t)} - \eta_{t+1}, t + 1)]).$$

By the definition of policy π ,

$$\mathcal{V}_\tau^\pi(s) = \hat{p}\mathcal{V}_+ + (1 - \hat{p})\mathcal{V}_-.$$




Figure C.1: An example demonstrating that, for any $\mu^H > \mu^L$, if there is a path from μ^L that crosses $V - \varepsilon$ within a time-horizon then the identical path from μ^H also crosses $V - \varepsilon$ within that horizon.

Given the nature of the random walk, induced by π , the probability of crossing the barrier $V - \varepsilon$ within a given time horizon starting from $\mu^H = \mu_{(t)} + \eta_{t+1}$ is at least as large as that starting from $\mu^L = \mu_{(t)} - \eta_{t+1}$ at the same epoch, i.e. $\mathcal{V}_+ \geq \mathcal{V}_-$.

To see why this is true, define a path to be a sequence of upward and downward movements (i.e. a sequence of $+1$'s and -1 's) over a fixed number of epochs and the length of the path to be the number of epochs the path spans. Under π , a path of a given length from μ^H has the same probability as the corresponding path of the same length from μ^L since the upward (resp. downward) probability at a given epoch is fixed at \hat{p} (resp. $1 - \hat{p}$).

Let $(x_t, x_{t+1}, \dots, x_T) \in \{+1, -1\}^{T-t+1}$ denote a feasible path starting from μ^L at time t and crossing $(V - \varepsilon)$ within the horizon $[t, T]$. Then the corresponding path from μ^H also crosses $(V - \varepsilon)$ within the specified horizon since

$$V - \varepsilon \leq \mu^L + \sum_{i=t}^{\tau} x_i \eta_i < \mu^H + \sum_{i=t}^{\tau} x_i \eta_i.$$

This implies that there are at least as many paths from μ^H as from μ^L that cross the barrier within a time horizon. This completes the argument since paths represent mutually exclusive ways of crossing $(V - \varepsilon)$, and hence the overall probability of crossing is the sum of the probabilities of individual paths. Figure C.1 provides an illustration.

Thus, using well-known terminology from the reinforcement learning literature, let us consider a *single deviation strategy* under which the action in an arbitrarily chosen s in \mathcal{S} is given by $\pi'(s)$ while the action in every other state $s' \in \mathcal{S}, s' \neq s$ remains $\pi(s')$. Then, the

state-action value for this state s is given by

$$\begin{aligned}
Q^\pi(s, \pi'(s)) &= p_t^+ \mathcal{V}_+ + (1 - p_t^+) \mathcal{V}_- \\
&= p_t^+ (\mathcal{V}_+ - \mathcal{V}_-) + \mathcal{V}_- \\
&> \hat{p} (\mathcal{V}_+ - \mathcal{V}_-) + \mathcal{V}_-, \quad \text{since } p_t^+ > \hat{p}, \mathcal{V}_+ - \mathcal{V}_- \geq 0, \\
&= \hat{p} \mathcal{V}_+ + (1 - \hat{p}) \mathcal{V}_- \\
&= \mathcal{V}_\tau^\pi(s).
\end{aligned}$$

Since the state s is arbitrarily chosen, we conclude that the single deviation strategy improves the value for *every* state in \mathcal{S} . Hence, by the strong form of the *policy improvement theorem* of reinforcement learning (see Sutton and Barto (1998) Sec. 4.2), if the action for every state $s \in \mathcal{S}$ is changed to $\pi'(s)$ from $\pi(s)$, then $\mathcal{V}_\tau^{\pi'}(s) \geq \mathcal{V}_\tau^\pi(s) \forall s \in \mathcal{S}$. \square

References

- J. Abernethy and R. Frongillo. A characterization of scoring rules for linear properties. In *Proceedings of the Conference on Learning Theory*, pages 27.1–27.13, 2012.
- J. Abernethy, Y. Chen, and J. Wortman Vaughan. An optimization-based framework for automated market-making. In *Proceedings of the 12th ACM Conference on Electronic Commerce*, pages 297–306, 2011.
- J. Abernethy, S. Kutty, S. Lahaie, and R. Sami. Information aggregation in exponential family markets. In *Proceedings of the 15th ACM Conference on Economics and Computation*, pages 395–412, 2014.
- K. J. Arrow, R. Forsythe, M. Gorham, R. W. Hahn, R. Hanson, D. Kahneman, J. O. Ledyard, S. Levmore, R. Litan, P. R. Milgrom, F. D. Nelson, G. R. Neumann, M. Ottaviani, C. R. Plott, T. C. Schelling, R. J. Shiller, V. L. Smith, E. Snowberg, S. Sunder, C. R. Sunstein, P. C. Tetlock, P. E. Tetlock, H. R. Varian, J. Wolfers, and E. Zitzewitz. Statement on prediction markets. *AEI-Brookings Joint Center Related Publication*, 7, 2007.
- J. Berg and T. Rietz. The Iowa Electronic Markets: Stylized facts and open issues. *Information Markets: A New Way of Making Decisions*, pages 142–169, 2006.
- J. Berg, R. Forsythe, F. Nelson, and T. Rietz. Results from a dozen years of election futures markets research. *Handbook of Experimental Economics Results*, 1:742–751, 2008.
- A. Beygelzimer, J. Langford, and D. M. Pennock. Learning performance of prediction markets with Kelly bettors. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 3*, pages 1317–1318, 2012.
- R. F. Bordley. A multiplicative formula for aggregating probability assessments. *Management Science*, 28(10):1137–1148, 1982.
- S. Boucheron, G. Lugosi, and O. Bousquet. Concentration inequalities. In *Advanced Lectures on Machine Learning*, pages 208–240. Springer, 2004.
- C. 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, pages 737–744, 2012.

- A. Brahma, M. Chakraborty, S. Das, A. Lavoie, and M. Magdon-Ismail. A Bayesian market maker. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 215–232, 2012.
- G. W. Brier. Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78(1):1–3, 1950.
- M. Burnashev and K. Zigangirov. An interval estimation problem for controlled observations. *Problemy Peredachi Informatsii*, pages 51–61, 1974. Originally in Russian.
- R. Castro and R. Nowak. Active learning and sampling. *Foundations and Applications of Sensor Management*, pages 177–200, 2008.
- N. Cesa-Bianchi, Y. Freund, D. Haussler, D. P. Helmbold, R. E. Schapire, and M. K. Warmuth. How to use expert advice. *Journal of the Association for Computing Machinery*, 44(3):427–485, 1997.
- M. Chakraborty and S. Das. Market scoring rules act as opinion pools for risk-averse agents. In *Advances in Neural Information Processing Systems*, pages 2359–2367, 2015.
- M. Chakraborty and S. Das. Trading on a rigged game: Outcome manipulation in prediction markets. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence*, pages 158–164, 2016.
- M. Chakraborty, S. Das, A. Lavoie, M. Magdon-Ismail, and Y. Naamad. Instructor rating markets. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence*, pages 159–165, 2013.
- T. Chakraborty and M. Kearns. Market making and mean reversion. In *Proceedings of the 12th ACM Conference on Electronic Commerce*, pages 307–314, 2011.
- Y. Chen and D. M. Pennock. A utility framework for bounded-loss market makers. In *Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence*, 2007.
- Y. Chen and J. W. Vaughan. A new understanding of prediction markets via no-regret learning. In *Proceedings of the 11th ACM Conference on Electronic Commerce*, pages 189–198, 2010.
- Y. Chen, S. Dimitrov, R. Sami, D. Reeves, D. Pennock, R. Hanson, L. Fortnow, and R. Gonen. Gaming prediction markets: Equilibrium strategies with a market maker. *Algorithmica*, pages 1–40, 2009.
- Y. Chen, X. Gao, R. Goldstein, and I. Kash. Market manipulation with outside incentives. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence*, 2011a.

- Y. Chen, I. Kash, M. Ruberry, and V. Shnayder. Decision markets with good incentives. In *Proceedings of the Conference on Web and Internet Economics*, pages 89–100, 2011b.
- Y. Cheung and M. Elkind. Stochastic approximation with virtual observations for dose-finding on discrete levels. *Biometrika*, 2010a.
- Y. K. Cheung and M. S. Elkind. Stochastic approximation with virtual observations for dose-finding on discrete levels. *Biometrika*, 97(1):109–121, 2010b.
- H. Clifton. Harsh criticism stymies future of pentagons futures market. *PRWeek*, 2003.
- B. Cowgill and E. Zitzewitz. Corporate prediction markets: Evidence from Google, Ford, and Firm X. *Review of Economic Studies*, 82(4):1309–1341, 2015.
- J. Crémer and R. P. McLean. Full extraction of the surplus in bayesian and dominant strategy auctions. *Econometrica: Journal of the Econometric Society*, pages 1247–1257, 1988.
- S. Das. A learning market-maker in the Glosten–Milgrom model. *Quantitative Finance*, 5(2):169–180, 2005.
- S. Das. The effects of market-making on price dynamics. In *Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems*, May 2008a.
- S. Das. The effects of market-making on price dynamics. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pages 887–894, 2008b.
- S. Das and M. Magdon-Ismail. Adapting to a market shock: Optimal sequential market-making. *Advances in Neural Information Processing Systems*, pages 361–368, 2008.
- S. Das and M. Magdon-Ismail. Adapting to a market shock: Optimal sequential market-making. In *Advances in Neural Information Processing Systems*, pages 361–368, 2009.
- J. Dave Cliff. Zero is not enough: On the lower limit of agent intelligence for continuous double auction markets. Technical report, technical Report HPL-97-141, Hewlett-Packard Laboratories Bristol. 105, 1997.
- S. Dimitrov and R. Sami. Composition of markets with conflicting incentives. In *Proceedings of the ACM Conference on Electronic Commerce*, pages 53–62, 2010.
- J. D. Farmer, P. Patelli, and I. I. Zovko. The predictive power of zero intelligence in financial markets, 2005a.
- J. D. Farmer, P. Patelli, and I. I. Zovko. The predictive power of zero intelligence in financial markets. In *Proceedings of the National Academy of Sciences of the United States of America*, volume 102, pages 2254–2259. National Academy of Sciences, 2005b.

- R. Forsythe, F. Nelson, G. R. Neumann, and J. Wright. Anatomy of an experimental political stock market. *The American Economic Review*, pages 1142–1161, 1992.
- F. Foster and S. Viswanathan. Strategic trading when agents forecast the forecasts of others. *Journal of Finance.*, 51:1437–1478, 1996.
- R. M. Frongillo, N. D. Penna, and M. D. Reid. Interpreting prediction markets: a stochastic approach. In *Advances in Neural Information Processing Systems*, pages 3266–3274, 2012.
- D. Fudenberg and J. Tirole. Perfect Bayesian equilibrium and sequential equilibrium. *Journal of Economic Theory*, 53(2):236–260, 1991.
- X. Gao, J. Zhang, and Y. Chen. What you jointly know determines how you act: Strategic interactions in prediction markets. In *Proceedings of the 14th ACM Conference on Electronic Commerce*, pages 489–506, 2013.
- X. A. Gao, A. Mao, Y. Chen, and R. P. Adams. Trick or treat: putting peer prediction to the test. In *Proceedings of the 15th ACM Conference on Economics and Computation*, pages 507–524, 2014.
- A. Garg, T. Jayram, S. Vaithyanathan, and H. Zhu. Generalized opinion pooling. In *Proceedings of the 8th International Symposium on Artificial Intelligence and Mathematics*. Citeseer, 2004.
- C. Genest and J. V. Zidek. Combining probability distributions: A critique and an annotated bibliography. *Statistical Science*, pages 114–135, 1986.
- L. Glosten and P. Milgrom. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14:71–100, 1985.
- T. Gneiting and A. E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.
- D. K. Gode and S. Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, pages 119–137, 1993.
- I. J. Good. Rational decisions. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 107–114, 1952.
- A. Graefe and J. S. Armstrong. Comparing face-to-face meetings, nominal groups, delphi and prediction markets on an estimation task. *International Journal of Forecasting*, 27(1):183–195, 2011.
- R. T. Griffith. *Hymns Of The Rigveda*. E.J. Lazarus and Co., 1896.

- R. Hanson. Decision markets. *IEEE Intelligent Systems*, 14(3):16–19, 1999.
- R. Hanson. Book orders for market scoring rules. *George Mason University*, 2003a.
- R. Hanson. Combinatorial information market design. *Information Systems Frontiers*, 5(1):107–119, 2003b.
- R. Hanson. Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets*, 1(1):3–15, 2007a.
- R. Hanson. The policy analysis market (a thwarted experiment in the use of prediction markets for public policy). *Innovations*, 2(3):73–88, 2007b.
- R. Hanson and R. Oprea. A manipulator can aid prediction market accuracy. *Economica*, 76(302):304–314, 2009.
- R. Hanson, R. Oprea, and D. Porter. Information aggregation and manipulation in an experimental market. *Journal of Economic Behavior and Organization*, 60:449–459, 2006.
- J. Harrison, N. Keskin, and A. Zeevi. Bayesian dynamic pricing policies: Learning and earning under a binary prior distribution. Working paper, Columbia and Stanford Universities, 2010.
- J. M. Harrison, N. B. Keskin, and A. Zeevi. Bayesian dynamic pricing policies: Learning and earning under a binary prior distribution. *Management Science*, 58(3):570–586, 2012.
- C. Holden and A. Subrahmanyam. Long-lived private information and imperfect competition. *Journal of Finance*, 47:247–270, 1992.
- M. Horstein. Sequential transmission using noiseless feedback. *IEEE Transactions on Information Theory*, pages 136–143, 1963.
- J. Hu and A. Storkey. Multi-period trading prediction markets with connections to machine learning. In *Proceedings of the 31st International Conference on Machine Learning*, 2014.
- E. Huang and Y. Shoham. Price manipulation in prediction markets: analysis and mitigation. In *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems*, pages 213–220, 2014.
- K. Iyer, R. Johari, and C. C. Moallemi. Information aggregation and allocative efficiency in smooth markets. *Management Science*, 60(10):2509–2524, 2014.
- B. Jedynek, P. L. Frazier, and R. Sznitman. Twenty questions with noise: Bayes optimal policies for entropy loss. Submitted, 2011.
- B. Jedynek, P. I. Frazier, and R. Sznitman. Twenty questions with noise: Bayes optimal policies for entropy loss. *Journal of Applied Probability*, 49(1):114–136, 2012.

- L. Jian and R. Sami. Aggregation and manipulation in market scoring rules: Effects of interface and information distribution. *Management Science*, 58(1):123–140, 2012.
- R. Jurca and B. Faltings. Mechanisms for making crowds truthful. *Journal of Artificial Intelligence Research*, 34(1):209, 2009.
- R. Karp and R. Kleinberg. Noisy binary search and its applications. In *Proceedings of the 18th ACM-SIAM Symposium on Discrete Algorithms*, pages 881–890, 2007.
- C. Krishnamurti. Introduction to market microstructure. In *Investment Management*, pages 13–29. Springer, 2009.
- A. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1336, 1985.
- A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic theory*, volume 1. New York: Oxford university press, 1995.
- N. Miller, P. Resnick, and R. Zeckhauser. Eliciting Informative Feedback: The Peer-Prediction Method. *Management Science*, 5(9):1359–1373, Sept. 2005.
- J. Millin, K. Geras, and A. J. Storkey. Isoelastic agents and wealth updates in machine learning markets. In *Proceedings of the 29th International Conference on Machine Learning (ICML-12)*, pages 1815–1822, 2012.
- D. Opitz and R. Maclin. Popular ensemble methods: An empirical study. *Journal of Artificial Intelligence Research*, 11:169–198, 1999.
- M. Ostrovsky. Information aggregation in dynamic markets with strategic traders. *Econometrica*, 80(6):2595–2647, 2012.
- A. Othman. Zero-intelligence agents in prediction markets. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pages 879–886, 2008.
- A. Othman and T. Sandholm. Decision rules and decision markets. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*, pages 625–632, 2010.
- A. Othman, D. M. Pennock, D. M. Reeves, and T. Sandholm. A practical liquidity-sensitive automated market maker. *ACM Transactions on Economics and Computation*, 1(3):14, 2013.
- M. Ottaviani and P. N. Sørensen. Outcome manipulation in corporate prediction markets. *Journal of the European Economic Association*, 5(2-3):554–563, 2007.

- D. Pennock and R. Sami. Computational aspects of prediction markets. In N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, editors, *Algorithmic Game Theory*. Cambridge University Press, 2007.
- D. Pennock and L. Xia. Price updating in combinatorial prediction markets with Bayesian networks. In *Proceedings of the 27th Conference on Uncertainty in Artificial Intelligence (UAI-11)*, pages 581–588, 2011.
- D. M. Pennock. *Aggregating probabilistic beliefs: Market mechanisms and graphical representations*. PhD thesis, The University of Michigan, 1999.
- J. Peterson and J. Krug. Augur: a decentralized, open-source platform for prediction markets. *arXiv preprint arXiv:1501.01042*, 2015.
- D. Prelec. A Bayesian Truth Serum for subjective data. *Science*, 306(5695):462–466, 2004.
- D. Prelec and H. Seung. An algorithm that finds truth even if most people are wrong. *Working Paper*, 2007.
- G. Radanovic and B. Faltings. Incentives for truthful information elicitation of continuous signals. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence*, 2014.
- P. Rhode and K. Strumpf. Manipulating political stock markets: A field experiment and a century of observational data. *University of Arizona, mimeo*, 2006.
- B. Riemann. Über die Anzahl der Primzahlen unter einer gegebenen Grösse. *Gesammelte Mathematische Werke und Wissenschaftlicher Nachlaß*, 2:145–155, 1859.
- H. Robbins and S. Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, pages 400–407, 1951.
- D. M. Rothschild and R. Sethi. Trading strategies and market microstructure: Evidence from a prediction market. *Journal of Prediction Markets*, 10(1), 2016.
- J. E. Schmitz. Algorithmic trading in the Iowa Electronic Markets. *Algorithmic Finance*, 1(2):157–181, 2011.
- R. Sethi and J. W. Vaughan. Belief aggregation with automated market makers. *Computational Economics*, 48(1):155–178, 2016.
- P. Shi, V. Conitzer, and M. Guo. Prediction mechanisms that do not incentivize undesirable actions. In *Proceedings of the Conference on Web and Internet Economics*, 2009.
- J. Stiglitz. Terrorism: There’s no futures in it. *Los Angeles Times*, July 31, 2003.

- A. J. Storkey, Z. Zhu, and J. Hu. Aggregation under bias: Rényi divergence aggregation and its implementation via machine learning markets. In *Machine Learning and Knowledge Discovery in Databases*, pages 560–574. Springer, 2015.
- J. Surowiecki. *The wisdom of crowds*. Anchor, 2005.
- R. S. Sutton and A. G. Barto. *Reinforcement learning: An introduction*. Cambridge: MIT Press, 1998.
- R. Sznitman and B. Jedynek. Active testing for face detection and localization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1914–1920, 2010.
- R. Sznitman, A. Lucchi, P. I. Frazier, B. Jedynek, and P. Fua. An optimal policy for target localization with application to electron microscopy. In *Proceedings of the 30th International Conference on Machine Learning (ICML)*, pages 1–9, 2013.
- V. Tresp. Committee machines. In Y. H. Hu and J.-N. Hwang, editors, *Handbook for neural network signal processing*, pages 135–151. CRC Press Cleveland, Ohio, USA, 2001.
- A. Tversky and D. Kahneman. Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131, 1974.
- R. Waeber, P. I. Frazier, and S. G. Henderson. A Bayesian Approach to Stochastic Root Finding. In *Proceedings of the Winter Simulation Conference*, pages 4033–4045. IEEE, 2011.
- B. Waggoner and Y. Chen. Information elicitation sans verification. In *Proceedings of the 3rd Workshop on Social Computing and User Generated Content (SC13)*, 2013.
- E. Wah and M. P. Wellman. Welfare effects of market making in continuous double auctions. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems*, pages 57–66, 2015.
- R. Weaver and D. Prelec. Creating truthtelling incentives with the bayesian truth serum. *Journal of Marketing Research*, pages 1–50, 2012.
- J. Witkowski and D. Parkes. Peer prediction without a common prior. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, 2012.
- J. Wolfers and E. Zitzewitz. Prediction markets. *Journal of Economic Perspectives*, 18(2): 107–126, 2004.
- H. Zhang, E. Horvitz, Y. Chen, and D. C. Parkes. Task routing for prediction tasks. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems- Volume 2*, pages 889–896, 2012.

Vita

Mithun Chakraborty

Degrees B. E. (First Class with Honours), Electronics and Telecommunication Engineering, December 2009
Ph.D., Computer Science, May 2017

Publications *Highly selective archival conferences*

Mithun Chakraborty, Sanmay Das (2016). Trading On A Rigged Game: Outcome Manipulation In Prediction Markets. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 158-164.

Mithun Chakraborty, Sanmay Das (2015). Market Scoring Rules Act As Opinion Pools For Risk-Averse Agents. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 2350-2358. **Selected for a spotlight presentation** (16.6% of accepted papers).

Mithun Chakraborty, Sanmay Das, Justin Peabody (2015). Price Evolution in a Continuous Double Auction Prediction Market With a Scoring-Rule Based Market Maker. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, pp. 835-841.

Mithun Chakraborty, Sanmay Das, Allen Lavoie, Malik Magdon-Ismael, and Yonatan Naamad (2013). *Instructor Rating Markets*. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI)*, pp. 159-165.

Aseem Brahma, Mithun Chakraborty, Sanmay Das, Allen Lavoie, and Malik Magdon-Ismael (2012). A Bayesian Market Maker. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*, pp. 215-232.

Mithun Chakraborty, Sanmay Das, and Malik Magdon-Ismael (2011). Near-Optimal Target Learning With Stochastic Binary Signals. In

Proceedings of the 27th Conference on Uncertainty in Artificial Intelligence (UAI), pp. 69-76.

Non-archival workshops and conferences

Mithun Chakraborty, Kai Yee Phoebe Chua, Sanmay Das, Brendan Juba. Coordinated Versus Decentralized Exploration In Multi-Agent Multi-Armed Bandits.

- Accepted at *8th International Workshop on Cooperative Games and Multiagent Systems* (CoopMAS) at AAMAS, 2017.
- In *2nd Workshop on Learning, Inference and Control of Multi-Agent Systems* (MALIC) at NIPS, 2016.

Mithun Chakraborty, Sanmay Das. On Manipulation in Prediction Markets When Participants Influence Outcomes Directly.

- In *4th Workshop on Social Computing and User Generated Content* (SCUGC) at EC, 2014.
- In *Proceedings of the 2nd Collective Intelligence Conference* (CI), 2014. Abstract only.

Mithun Chakraborty, Sanmay Das, Allen Lavoie, Malik Magdon-Ismail, and Yonatan Naamad. Instructor Rating Markets.

- In *1st Workshop on Social Computing and User Generated Content* (SCUGC) at EC, 2011.
- In *Proceedings of the 2nd Conference on Auctions, Market Mechanisms, and Their Applications* (AMMA), 2011. Abstract only.

Others

For a list of publications from my undergraduate projects not related to my dissertation, please visit

<http://www.cse.wustl.edu/~mithunchakraborty/Research.html>.

Research presentations

On The Aggregation Of Subjective Inputs From Multiple Sources.

- Poster at *12th Conference on Web and Internet Economics* (WINE), 2016.
- Short talk at *SIGAI Career Network and Conference* (CNC), 2016.

On Manipulation in Prediction Markets When Participants Influence Outcomes Directly.

Invited talk at *20th Conference of the International Federation of Operational Research Societies*, 2014 (IFORS'14).

Service

Journal Reviewing: Frontiers of Computer Science, Springer, 2014.

Conference Reviewing: AAAI 2014; AAAI 2015; ECAI 2014 (sub-reviewer).

PC membership: AAAI 2014; AAAI 2015; CoopMAS 2017.

Miscellaneous: Student Representative of Graduate Admissions Committee (2010-12) and Graduate Curriculum Committee (2011-12) of RPI Comp. Sci. Dept.

Teaching

WU-CIRTL Practitioner, The Teaching Center, WUSTL, 2016.

Instructor: CSE 316A Social Network Analysis, WUSTL, Fall 2015.

Guest Lecturer: CSE 516A Multi-Agent Systems, WUSTL, Springs 2015, 2016, 2017.

Teaching Assistant:

– CSE 516A Multi-Agent Systems (instructor: Sanmay Das), WUSTL, Spring 2014;

– CSCI 2500 Computer Organization (instructor: Chris Carothers), RPI, Fall 2009.

Awards

ACM SIGAI Travel Grant for SIGAI CNC, 2016.

Department Chair Award for Outstanding Teaching, Department of Computer Science and Engineering, WUSTL, 2016.

UAI Student Travel Scholarship, 2011.

M. P. Birla Sponsorship for Higher Studies, 2009.

I. B. Putatunda and S. S. Putatunda Memorial Award for **highest marks in second year (sophomore) university examination**, 2007.

Bengal Peerless Award for Academic Excellence for **highest marks in state of West Bengal in Higher Secondary (10+2) Examination**, 2005. National Merit Scholarship offered by Govt. of India, 2003.

Miscellaneous

Attended Summer School on Algorithmic Economics at Carnegie Mellon University (acceptance rate: 63/159), August 2012.

Summer Intern (Advisor: Dr. Pabitra Pal Choudhury), Applied Statistics Unit, Indian Statistical Institute, Kolkata, India, Summer 2007.

May 2017

Aggregation of Subjective Inputs, Chakraborty, Ph.D. 2017