

An Axiomatic Characterization of Wagering Mechanisms ^{*}

Nicolas S. Lambert[†] John Langford[‡] Jennifer Wortman Vaughan[§] Yiling Chen[¶]
Daniel M. Reeves^{||} Yoav Shoham^{**} David M. Pennock^{††}

March 24, 2014

Abstract

We construct a budget-balanced wagering mechanism that flexibly extracts information about event probabilities, as well as the mean, median, and other statistics from a group of individuals whose beliefs are immutable to the actions of others. We show how our mechanism, called the Brier betting mechanism, arises naturally from a modified parimutuel betting market. We prove that it is essentially the *unique* wagering mechanism that is *anonymous*, *proportional*, *sybilproof*, and *homogeneous*. While the Brier betting mechanism is designed for individuals with immutable beliefs, we find that it continues to perform well even for Bayesian individuals who learn from the actions of others.

Keywords: betting market, wagering mechanism, parimutuel betting, scoring rule, belief elicitation, prediction market

JEL classification: D47, D53, D82

^{*}Part of this work was conducted while Chen, Lambert, Langford, Pennock, Reeves, and Vaughan were working at Yahoo! Research, New York. We are grateful to Alina Beygelzimer, Francis Bloch, Drew Fudenberg, Alfred Galichon, Antoine Lallour, Michael Ostrovsky, Marco Ottaviani, Eduardo Perez, Bethany Soule, Yevgeniy Vorobeychik, Robert Wilson, the participants of the ninth ACM Conference of Electronic Commerce, the editors and the referees for helpful comments and suggestions. This material is based upon work supported by Google Research and the National Science Foundation under grants No. CCF-1101209, IIS-1054911, and CCF-0953516. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding agencies. A preliminary version circulated under the title “Self-Financed Wagering Mechanisms for Forecasting” and was awarded an ACM Conference on Electronic Commerce Best Paper Award.

[†]Corresponding author. Address: Graduate School of Business, 655 Knight Way, Stanford University, Stanford, CA 94305, USA. Email: nlambert@stanford.edu. Phone: (+1) 650-723-6283.

[‡]Microsoft Research, 4 Chateaux Circle 4B, Scarsdale, NY 10583, USA. Email: jcl@microsoft.com.

[§]Microsoft Research, 641 Avenue of the Americas, 7th Floor, New York, NY 10011, USA. Email: jenn@microsoft.com.

[¶]School of Engineering and Applied Sciences, Harvard University, 33 Oxford Street, Cambridge, MA 02138, USA. Email: yiling@seas.harvard.edu.

^{||}Beeminder Inc., 330 SE Martin Luther King Blvd, 2nd Floor, Portland, OR 97214, USA. Email: dreeves@beeminder.com.

^{**}Department of Computer Science, 353 Serra Street, Stanford University, Stanford, CA 94305, USA. Email: shoham@stanford.edu.

^{††}Microsoft Research, 641 Avenue of the Americas, 7th Floor, New York, NY 10011, USA. Email: dpennock@microsoft.com.

1 Introduction

Consider a society of individuals endowed with heterogeneous information—and thus private beliefs—about an uncertain event, for example, a future sporting event, a stock market event, or the success of a project. Eliciting and fully aggregating their information in a market is theoretically possible under certain conditions (Lucas, 1972; Grossman, 1981; Ostrovsky, 2012) and works in simple laboratory tests (Plott and Sunder, 1982), but can fail in more realistic scenarios (Plott and Sunder, 1988; Forsythe and Lundholm, 1990). Many market mechanisms, including continuous double auctions (Wolfers and Zitzewitz, 2004), parimutuel markets (Ali, 1977; Thaler and Ziemba, 1988; Hurley and McDonough, 1995; Plott et al., 2003; Ottaviani and Sørensen, 2009), and market scoring rules (Hanson, 2003, 2007), have been used for information elicitation and aggregation. In theory, the aggregation of information (even partial) in a market often relies on the assumption that participants derive posterior beliefs from a common prior, acting as Bayesian agents and rationally updating their beliefs upon observing others’ actions. In practice, however, individuals may not be sophisticated enough to reason about the actions of others and incorporate the beliefs of other agents into their own. The frequent violation of the Bayesian-rationality assumption arguably contributes to the observed markets’ failure in aggregating information in experiments and real-world scenarios.

On the other hand, some simple opinion pooling methods (French, 1985; Genest and Zidek, 1986; Clemen and Winkler, 1999) have been shown to perform well in producing an aggregated forecast in practice. For example, taking an average or a weighted average of private beliefs—in essence treating individual information sources as independent—is a robust and practical alternative, with accuracy improving in the number and diversity of individuals (Forsythe et al., 1992; Jacobs, 1995; Surowiecki, 2004; Chen et al., 2005; Dani et al., 2006; Reeves and Pennock, 2007; Page, 2007). These methods post-process individual beliefs that have already been elicited. Extracting beliefs individually offers some comparative advantages relative to market environments. For instance, the elicitor has more control on the information contained in the final aggregate. She can discard reports to eliminate outliers, and reduce the impact of forecasts by uninformed individuals. She can weight an individual’s estimate to account for the historical performance, or as a function of its position relative to the overall distribution. If the elicitor is a decision maker, she can use, in addition to the computed final aggregate, the distribution of individual estimates, which reflects the dispersion of information among the individuals.

Proper scoring rules (Brier, 1950; Savage, 1971; Winkler et al., 1996; Gneiting and Raftery, 2007) and shared scoring rules (Kilgour and Gerchak, 2004) have been designed for the purpose of eliciting individual beliefs. They are *incentive compatible* for agents who have private beliefs that are immutable to the actions of others (i.e., who do not perform Bayesian updating), meaning they reward agents maximally for reporting their true beliefs. Some scoring rules are *distance sensitive* or *effective* and reward the accuracy of the reported posterior belief, in the sense that if the final outcome is drawn at random according to some probability distribution p , then the individual who reports a distribution q receives a payment that, on expectation, decreases with the

distance between p and q . Here the distance is a metric that depends on the scoring rule that is applied; it can be, for example, the Euclidian metric. Rewarding the accuracy of heterogeneous beliefs is generally perceived as desirable (Friedman, 1983; Nau, 1985; Jose et al., 2009). First, it reduces the dispersion of reported beliefs which reduces the amount of noise in the computed aggregate. It also improves the performance of aggregation methods that discard outliers or weight each announcement according to some centrality criterion. Second, it encourages individual to seek (potentially costly) information so as to produce good forecasts (see, for example, Osband, 1989). However, if proper scoring rules can reward agents for reporting true beliefs and for being accurate, they are not *budget balanced*. That is, inducing an individual to deliver truthful forecasts requires an outside subsidy. When several individuals are asked to provide information, there is a possibility to self-finance the mechanism, by having the bad predictors sponsor the good ones. Shared scoring rules are an instance of such a budget-balanced mechanism. But these scoring rules are not *sybilproof*: Individuals can benefit from submitting several reports under different pseudonyms. Shared scoring rules also do not allow individuals to specify the maximum loss they are willing to incur.

Focusing on eliciting private information from individuals, we propose a new mechanism, the *Brier betting mechanism*, that has a number of desirable properties, including the four above.

The mechanism proceeds in two stages. First, individuals wager money on event probabilities. Once all bets are collected and the event outcome is known, the total amount wagered is redistributed among the bets. The fractions of the total allocated to bets are expressed as quadratic functions of the outcome realization and individual predictions (essentially a Brier score), in which each term is weighted by the fraction of money invested by each participant.

The Brier betting mechanism is incentive compatible for individuals with immutable beliefs, budget balanced and:

- **Proportional:** An individual’s expected payoff increases as his or her absolute accuracy increases, and decreases as the accuracy of other forecasts increases.
- **Anonymous:** The market does not discriminate among participants.
- **Sybilproof:** Market outputs do not depend on whether participants divide their wager among one or several bets.
- **Homogeneous:** Market outputs depend solely on the fraction each player invests relative to the total amount invested.

A natural question to ask is whether there are other mechanisms satisfying these properties. We show that the answer is *no*: The Brier betting mechanism is essentially the only mechanism that fulfills all these properties.

Our mechanism is also ex-ante individually rational (IR)—agents who make correct predictions have a nonnegative expected value for participating—but it is not *ex-post* IR. In a budget-balanced mechanism, if one agent can gain, another must risk losing. A bounded scoring rule with sufficient

subsidy can be ex-post IR, ensuring no participant can ever lose. But an ex-post IR mechanism rewards uninformed and even misinformed experts, encourages sybils, and isn't practical for large, open-membership groups.

The Brier betting mechanism is designed for individuals with immutable, heterogeneous prior beliefs, hence may have robust performance in some real-world scenarios where individuals are highly confident and do not change their mind upon seeing others' actions. Perhaps surprisingly, if individuals are instead rational, Bayesian agents who learn from the actions of others, our mechanism continues to perform well. Specifically, under a common prior assumption, a slight bias away from truthful reporting may arise under asymmetric information, through the correlation between the total wealth wagered and the event outcome. This bias is driven towards zero as the fraction of any individual's wealth compared to the group's converges to zero.

The paper proceeds as follows. Section 2 describes related work. Section 3 introduces the Brier betting mechanism, and shows it satisfies a number of desirable properties. Section 4 shows how the Brier betting mechanism can be constructed heuristically from the horse-race parimutuel markets. Section 5 takes an axiomatic approach. It presents the Brier betting mechanism as a member of a large class of mechanisms in which individuals can put wagers on probabilistic predictions, and shows that the Brier betting mechanism is the only member of the class that satisfies all the properties. Section 6 analyzes the performance of the Brier mechanism under a common prior setting with rational Bayesian agents. Section 7 discusses an alternative interpretation of the Brier mechanism in the context of eliciting expert advice, as well as an extension for the extraction of general statistics. Finally, Section 8 concludes.

2 Background and Related Work

The Brier betting mechanism combines features of the horse-race parimutuel market and the scoring rule elicitation method. As in the scoring rule method, market participants announce probabilistic predictions and get rewarded accordingly. As in the parimutuel market, participants specify their degree of participation, that is, their stake, in the market. Johnstone (2007) explores a similar combination but the resulting mechanism is not incentive compatible.

Parimutuel markets are commonly used to wager on sports including horse racing and jai alai. Players decide how much to wager on a categorical outcome of an uncertain variable: No probability assessment is given. After the actual outcome materializes, the pool of money is divided among the winning players who chose the correct outcome in proportion of the amount they bet. Parimutuel betting platforms have been thoroughly studied (Eisenberg and Gale, 1959; Ali, 1977; Quandt, 1986; Thaler and Ziemba, 1988; Watanabe et al., 1994; Hurley and McDonough, 1995; Takahiro, 1997; Plott et al., 2003; Ottaviani and Sørensen, 2005, 2008, 2009; Koessler et al., 2008).

Parimutuel markets enable players to confront their predictions, but they do so in a quite restricted sense. For one thing, the parimutuel game is complex. The action best played for a given subjective belief depends on the action chosen by other players: There are generally no

dominant strategies (Ali, 1977). Even in simple Bayesian game theoretic settings that distribute binary signals uniformly across players, multiple non-trivial equilibria arise (Watanabe et al., 1994; Koessler et al., 2008). Additionally, the action space can be too coarse for separating equilibria to exist. If the signal space is dense enough, different posteriors may yield the same equilibrium strategies (Ottaviani and Sørensen, 2009, 2010). This impedes both information revelation and incentives: Better information may not generate better payoffs, and some information is kept private.

Scoring rules offer an alternative technique to eliciting forecasts. Scoring rules assign scores to a probabilistic forecast as a function of the outcome realization. When the scoring rule is proper, an individual paid according to the scoring rule maximizes her expected earnings when telling the truth (Brier, 1950; Savage, 1971; Winkler et al., 1996; Gneiting and Raftery, 2007). The method elicits subjective probabilities separately for each individual, but requires an external subsidy. In contrast, in a market structure, the less informed participants typically sponsor the more informed ones. Kilgour and Gerchak (2004) extend the scoring rule method to a group of forecasters, to benefit from such a self-financing property. With these *competitive scoring rules*, each group member is evaluated according to their own score, minus the average score of the other members. Competitive scoring rules do not need an external subsidy but require that all participants be entitled to the same liability and encourage pseudonyms. In contrast, the mechanism we introduce allows each individual to specify an upper bound on the loss she is willing to incur, since an individual can never lose more than the amount she chooses to gamble.

Our mechanism also relates to prediction markets. In these markets, participants trade on event-contingent securities, for example a contract that pays off \$1 if the Red Sox win the game against the Yankees, and zero if they lose (Wolfers and Zitzewitz, 2004). Prediction markets induce aggregate forecasts through market prices. They are fundamentally dynamic: Traders may take actions at different points in time. They observe price movements and update their beliefs. This dynamic structure enables the aggregation of private beliefs, but it may also induce strategic traders to delay the revelation of private information, or to manipulate the market and propagate false beliefs (Chen et al., 2010; Lallour, 2011; Ostrovsky, 2012). As in horse-race parimutuel markets, these prediction markets combine participation and prediction together in a single dimension. Consequently, an informed trader may lack the wealth necessary to move the market, which prevents the release of private information. Symmetrically, if a trader has a belief close to the current market price, there is little or no incentive to participate. Additionally, many prediction markets are prone to the thin-market problem: With few traders, one market side may not find matching orders, hindering price discovery (Milgrom and Stokey, 1982). In contrast this paper focuses on static mechanisms that elicit individual predictions, and allows participants to decide on a prediction and participation level independently.

3 Brier Betting Mechanism

In this section we introduce a betting mechanism based on the Brier score, also known as the quadratic score or squared loss. Consider a binary outcome, given by a random variable X taking values in $\{0, 1\}$. X is interpreted as the indicator variable of a binary event E , such as the democrats winning the next presidential elections, or an economic growth greater than 3% the next year. The event is uncertain: let p be the actual probability that E occurs.¹ All expectations are with respect to the true probability p unless otherwise noted.

The Brier betting mechanism involves a group of n individuals, or *players*, indexed $1, \dots, n$, who wager money on their assessment of the probability of E . It proceeds in two stages:

- In a first stage, every player i places a bet. The bet consists of a wager $w_i \geq 0$ and a probabilistic estimate $p_i \in [0, 1]$ for event E . Bets are submitted simultaneously.
- In a second stage, the event outcome is publicly observed. Denote by x the realization of X . Each player i receives a fraction of the total amount wagered, based on the Brier score $s(q, x) = 1 - (x - q)^2$, according to the formula:

$$\pi_i((p_1, w_1), \dots, (p_n, w_n), x) = \frac{w_i}{W} + \frac{w_i}{W} \left(s(p_i, x) - \sum_j \frac{w_j}{W} s(p_j, x) \right), \quad (1)$$

where $W = \sum_j w_j$ is the total amount wagered.

The (net) payoff for player i is therefore $W\pi_i - w_i$. Note that the shares sum to one and are nonnegative.

Let us briefly examine the form Equation (1), which is discussed in more detail in Section 4. The first term, w_i/W , is the share of the collected wealth W contributed by player i . The second term determines whether player i will receive a smaller or larger share than what she initially contributed by comparing her score with the weighted average of the scores of the other players. It is easy to show that this is equivalent to

$$\begin{aligned} \pi_i((p_1, w_1), \dots, (p_n, w_n), x) &= \frac{w_i}{W} + \frac{w_i}{W} \sum_{j \neq i} \frac{w_j}{W} ((p_j - x)^2 - (p_i - x)^2) \\ &= \frac{w_i}{W} + 2 \frac{w_i}{W} \sum_{j \neq i} \frac{w_j}{W} (p_j - p_i) \left(\frac{p_i + p_j}{2} - x \right). \end{aligned}$$

This formulation offers a natural interpretation of the mechanism in terms of monetary transfer

¹It is common in the forecasting literature to assume outcomes are drawn at random from some true, *objective* probability distribution (Dawid, 1982; Foster and Vohra, 1998; Olszewski and Sandroni, 2008). One can also interpret the probability p as an immutable *subjective* probability with respect to which an individual seeks to maximize her expected payoff, as in Eisenberg and Gale (1959) and Ali (1977) who endow players with heterogeneous prior beliefs.

between players. The output of the mechanism corresponds to each player i transferring a fraction

$$\frac{w_j}{W} ((p_i - x)^2 - (p_j - x)^2) = 2\frac{w_j}{W}(p_i - p_j) \left(\frac{p_i + p_j}{2} - x \right) \quad (2)$$

of her wager to player j . If the fraction (2) is negative, the transfer is from player j to player i . Hence, for any given action profile $(p_1, w_1), \dots, (p_n, w_n)$, in expectation when $x = 1$ with probability p , player i gives away a fraction

$$\frac{w_j}{W} ((p_i - p)^2 - (p_j - p)^2) = 2\frac{w_j}{W}(p_i - p_j) \left(\frac{p_i + p_j}{2} - p \right)$$

of her wager to player j . In expectation, player i receives money from player j if and only if $|p_i - p| < |p_j - p|$. In other words, player i receives money from player j if and only if she announces a better prediction, in the sense that the absolute difference between the prediction and the truth is lower. Furthermore, her share is proportional to both the extent of the disagreement $|p_i - p_j|$ between the two players, and the quality of the average prediction of the two players, represented by the absolute difference $|(p_i + p_j)/2 - p|$.

Observing that the expected share of player i is

$$\frac{w_i}{W} + \frac{w_i}{W} \sum_{j \neq i} \frac{w_j}{W} ((p_j - p)^2 - (p_i - p)^2)$$

we can establish a list of properties that the Brier betting mechanism satisfies:

- (1) For every permutation σ over $\{1, \dots, n\}$,

$$\pi_{\sigma^{-1}(i)}((p_{\sigma(1)}, w_{\sigma(1)}), \dots, (p_{\sigma(n)}, w_{\sigma(n)}), x) = \pi_i((p_1, w_1), \dots, (p_n, w_n), x) .$$

for every x , i , and every action profile $(p_1, w_1), \dots, (p_n, w_n)$.

- (2) The expected share of player i increases as $|p_i - p|$ decreases. That is, for all i , and all $(p_1, w_1), \dots, (p_n, w_n)$, the value of

$$\mathbb{E}[\pi_i((p_1, w_1), \dots, (p_n, w_n), X)]$$

increases as $|p_i - p|$ decreases.

- (3) The expected share of player i decreases as $|p_j - p|$ decreases, for every other player j . That is, for all i , $j \neq i$, all $(p_1, w_1), \dots, (p_n, w_n)$, the value of

$$\mathbb{E}[\pi_i((p_1, w_1), \dots, (p_n, w_n), X)]$$

decreases as $|p_j - p|$ decreases.

- (4) If two or more players make identical forecasts, the mechanism output for everyone else de-

depends on the identical forecasters wagers only through the sum of the wagers of these identical forecasters, but does not depend on each wager individually.

- (5) The share to every player remains unchanged when all the players multiply their wager by the same factor.

Let us briefly discuss the meaning of these properties. Property (1) states that the mechanism is anonymous: It does not discriminate between players, in that the payoffs do not depend on the players' identities. Properties (2) and (3) state that the mechanism is proportional. According to property (2), better predictions yield better payoffs. It is worth noting that the mechanism not only rewards truth telling, it also rewards accuracy, in the sense that the closer the forecast is to the truth, the larger the expected payoff. This feature of our mechanism encourages forecasts of better quality, in the event that the player has the possibility to access more accurate information. In doing so, it reduces the spread of the distribution of the different forecasts, which increases the precision of the computed aggregate, if an aggregation method is being used as a post-processing step. With property (3), a player's payoff decreases as another player makes better predictions (such as when another player has access to more accurate information) as can be expected of any competitive mechanism. In electronic platforms, it is important to prevent players from manipulating identities. Property (4) fulfills this requirement: The mechanism is sybilproof. There is no incentive for a player to create multiple fake identities, as outputs do not depend on whether participants divide their wager among one or several bets. Finally, property (5) states that mechanism is homogeneous in that the share allocated to each player depends solely on the fraction each player invests relative to the total amount invested. In other words, the share to every player does not depend on the currency being used.

In our mechanism the wager determines a player's stake or risk exposure. Both net payoffs and expected net payoffs are monotone in the player's wager. If a player loses money (in expectation), her (expected) loss increases with the amount of her wager. If however a player earns money (in expectation), her (expected) gain also increases with the amount of her wager. While a player's loss can never exceed her wager, a player's gain never exceeds her wager either. By increasing her wager, a player can potentially earn a larger payoff but it comes at the risk of a greater loss. It is worth noting that in many instances, players win and lose far smaller amounts than their full budget.

Note that in our mechanism, the quality of a prediction is measured with respect to the absolute difference between forecasted probabilities and actual ones. It is a direct consequence of the use of the Brier score. Such measure of prediction quality is suggested by Friedman (1983) or Nau (1985), among others. It implies that the payoff is symmetric in forecast errors since forecasts $p + \epsilon$ and $p - \epsilon$ must have the same payoff. Section 7 expands the mechanism to general scoring rules and demonstrates that the above properties continue to hold with weaker analogs for properties (2) and (3), thereby allowing for different measures of prediction quality.

In Section 5, we show that, subject to a smoothing condition, the Brier betting mechanism is essentially the *only* mechanism that satisfies (1)–(5). First, we build some intuition for the

mechanism.

4 Construction of the Brier Betting Mechanism

The Brier betting mechanism can be obtained by extending the standard horse-race parimutuel betting markets. In the present context, the horse-race parimutuel market is described as follows:

- In a first stage, every player $i = 1, \dots, n$ submits a bet. The bet consists of a wager $w_i \geq 0$ and a prediction x_i of the outcome. That is, $x_i = 1$ to predict that E occurs and $x_i = 0$ to predict that E does not occur. Bets are submitted simultaneously.
- In a second stage, the realization x is publicly observed. Players whose prediction is wrong lose their entire wager. Players who predict the outcome correctly share the total amount wagered. Winners get an amount proportional to their own wager.

The (net) payoff to player i is therefore

$$s_i \frac{w_i}{\sum_j s_j w_j} \left(\sum_j w_j \right) - w_i, \quad (3)$$

where $s_i = 1$ if player i is correct (that is, if $x_i = x$) and $x_i = 0$ otherwise.

Following Eisenberg and Gale (1959), Norvig (1967) and Ali (1977), suppose that players have heterogeneous prior beliefs. That is, players hold immutable subjective beliefs on outcome probabilities, and seek to maximize their expected payoffs with respect to these beliefs. Even in this simple configuration, the parimutuel betting game lacks dominant strategies: A player's optimal action typically depends on the actions of the other players.

To illustrate this point, consider the choice of a single player i , who must decide whether to place a small bet w_i on outcome 0 or outcome 1. Let W_0 be the total amount wagered on outcome 0 by other players, and W_1 be the total amount wagered by other players on outcome 1. From Equation 3, we can see that if a player i believes that outcome 1 occurs with probability p_i , she is better off putting her wager w_i on outcome 1 if

$$p_i \frac{w_i}{W_1 + w_i} > (1 - p_i) \frac{w_i}{W_0 + w_i},$$

or

$$\frac{p_i}{1 - p_i} > \frac{W_1 + w_i}{W_0 + w_i}.$$

Therefore her optimal prediction depends on the values W_0 and W_1 , which are unknown at the time she places her bet.

To get around this problem, we modify the structure of the payoffs to ensure that a player's optimal prediction depends solely on her subjective probability. In particular, it is independent of her own wager, and is independent of the predictions and wagers of the other players. The lack of

dominant strategies owes to the term $\sum_j s_j w_j$ in Equation 3, which is the total wager of the winners. It appears as a denominator which introduces a non-linearity in the payoffs. To restore linearity, we multiply the net payoffs by $\sum_j s_j w_j / \sum_j w_j$, scaling down the amount of money transferred. The net payoff for player i becomes

$$w_i \left(s_i - \frac{\sum_j s_j w_j}{\sum_j w_j} \right). \quad (4)$$

Note that payoffs remain bounded below by $-w_i$. In this transformed redistribution rule, no matter how much she decides to wager, a player is always better off predicting the most likely outcome according to her own belief, independently of the actions of others. However the space of possible predictions remains too coarse, impeding the full revelation of private beliefs. For example, a player who believes that there is a 60% chance event E will occur announces the same prediction she would announce if she instead believed there was a 90% chance.

To overcome the limitation, in a second step we extend the space of the possible predictions. We now allow players to announce outcome probabilities. Specifically, player i 's prediction is an estimate p_i of the probability that E occurs. Players can still predict an outcome if they wish, by announcing probabilities 0 or 1, but they can now express uncertainty as part of the bet. In the payoffs given by (3) and (4), the variable s_i is essentially a measure of performance. The greater the value s_i , the greater the payoff of player i . In the original parimutuel market, s_i is restricted to values 0 and 1. This dichotomy reflects the idea that a player is either entirely correct or completely wrong. As we now allow for a continuum of predictions, it is natural to also measure performance on a continuum: The more precise the forecast, the greater the performance.

Scoring rules provide a classic measure of performance for probabilistic predictions. A *scoring rule* s assigns a score $s(q, x)$ as a function of a probabilistic forecast $q \in [0, 1]$ and the realized outcome $x = 0, 1$. Accurate forecasts match empirical probabilities. Proper scoring rules ensure that the empirical average score is maximized under this condition. A scoring rule is (*strictly*) *proper* when the expected score, under the true probability p , is (strictly) maximized when $q = p$. Proper scoring rules and their constructions have been thoroughly studied. McCarthy (1956), Shuford et al. (1966), Hendrickson and Buehler (1971), Savage (1971), Friedman (1983), and Schervish (1989) offer various characterizations of proper and strictly proper scoring rules. Gneiting and Raftery (2007) provide a summary of the results.

No matter the amount she wagers, and no matter the actions of the other players, player i maximizes the expectation of her gains (as in Equation (4)) by maximizing the expected value of s_i , according to her subjective beliefs. Therefore, if we substitute for s_i in Equation (4) any proper scoring rule taking values in $[0, 1]$, player i optimally chooses to report her true subjective probability. After substitution, the net payoff for player i becomes

$$w_i \left[s(p_i, x) - \frac{\sum_j s(p_j, x) w_j}{\sum_j w_j} \right].$$

While we could, in principle, use any proper scoring rule, the Brier score is a simple and widely used strictly proper scoring rule that satisfies several nice properties. It is defined in this context as $s(q, x) = 1 - (q - x)^2$. Unlike any other scoring rule, the Brier score has the property that, the closer the forecast is to the true probability (in absolute difference), the greater the expected score will be. Section 5 investigates this property in greater detail. Plugging in the Brier score, the net payoff for agent i is

$$w_i \left[-(p_i - x)^2 + \frac{\sum_j (p_j - x)^2 w_j}{\sum_j w_j} \right],$$

which can be rewritten as

$$\sum_{j \neq i} \frac{w_i w_j}{W} ((p_j - x)^2 - (p_i - x)^2),$$

where $W = \sum_j w_j$. Adding the wager w_i of player i and dividing by W gives the share assigned to player i .

5 Axiomatic Approach

The Brier betting mechanism is part of a larger collection of mechanisms in which agents wager money on probabilistic predictions. We refer to these as *wagering mechanisms*.

Specifically, a wagering mechanism is defined by a *redistribution rule* π . The rule specifies how the total amount wagered is divided among the players, for any number of players. The mechanism operates in two stages. In stage 1, every player i bets an amount w_i on a prediction p_i for the probability that E occurs, i.e., the probability that $X = 1$. In stage 2, the realization x of X becomes known, and the total amount wagered W is divided among the players. Player i gets a share $\pi_i((p_1, w_1), \dots, (p_n, w_n), x)$ of W . That is, player i receives a net payoff of $W\pi_i((p_1, w_1), \dots, (p_n, w_n), x) - w_i$. Shares must be nonnegative and sum to one. We require throughout that every function π_i be twice continuously differentiable in the vector of predictions (p_1, \dots, p_n) .

The main result of this paper shows that, under this requirement, the Brier betting mechanism is essentially the only wagering mechanism that satisfies the properties (1)–(5) enumerated in Section 3. We lay out the proof of this theorem in several steps. First, we show that the Brier score is the only scoring rule to satisfy the property that the expected score decreases in the absolute difference between the report and the true event probability. Next, we use this feature to provide a uniqueness result for sharing rules that divide a fixed amount of money among a set of forecasters, a result of independent interest. Finally, we build on this result to prove our main theorem.

Theorem 1. *Consider a wagering mechanism defined by the redistribution rule (π_1, \dots, π_n) in which every function π_i is twice continuously differentiable. The mechanism satisfies the list of*

desiderata (1)–(5) enumerated in Section 3 if and only if

$$\pi_i((p_1, w_1), \dots, (p_n, w_n), x) = \frac{w_i}{W} + \kappa \frac{w_i}{W} \sum_{j \neq i} \frac{w_j}{W} ((p_j - x)^2 - (p_i - x)^2)$$

for some constant $\kappa \in [0, 1]$.

The presence of the quadratic term is due to the properties (2) and (3). In Section 7 we claim that under weaker properties, the quadratic term can be replaced by any proper scoring rule. Note the introduction of the constant κ , absent from the original formulation of the payoffs in Equation (1). Properties (1)–(5) indeed hold for all $\kappa \in [0, 1]$, however it is natural to set $\kappa = 1$. For any other choice of κ players always get to keep some fixed fraction of their own wager, even in the worst circumstances. Choosing $\kappa = 1$ ensures that a player can lose arbitrarily close to her wager in some situations. Also note that property (5) is useful but not essential to the characterization. The proof we offer in the following subsections shows that without property (5), the characterization generalizes only marginally, with κ being then a function of the total amount wagered W . Such a flexibility is not of particular interest as $\kappa = 1$ remains the most natural and desirable choice.

5.1 Characterization of the Brier Score

The quadratic or Brier scoring rules, defined by $s(q, x) = -c(q - x)^2 + f(x)$ for an arbitrary function f and positive constant c , are arguably the best known strictly proper scoring rules (Brier, 1950). We begin by showing that these scoring rules are the only ones satisfying a simple monotonicity property: The expected score increases as $|p - q|$ decreases, where p is the probability assigned to event E or outcome $x = 1$. The fact that the Brier scoring rules satisfy this property was shown by Friedman (1983); the result below shows uniqueness. As most of our analysis requires some smoothness of the functions being used, we prove the result for continuously differentiable scoring rules, that is, functions $s(q, x)$ continuously differentiable in q , for every choice of x .

Proposition 1. *The Brier scoring rules are the unique scoring rules $s(q, x)$ such that $\mathbb{E}[S(q, X)]$ is strictly decreasing in the absolute difference $|q - p|$.*

Proof of Proposition 1: We begin by a characterization of all the continuously differentiable proper scoring rules. We show that $s(q, x)$ is a proper scoring rule if and only if $s'(q, x) = \gamma(q)(x - q)$ for some continuously differentiable function $\gamma(\cdot) \geq 0$. Note that this statement is a direct consequence of the more general Schervish characterization (Schervish, 1989). For this special case we include a simple proof for completeness.

To show sufficiency, observe that if $s'(q, x) = \gamma(q)(x - q)$, then

$$\mathbb{E}[s(p, X)] - \mathbb{E}[s(q, X)] = \int_q^p \gamma(t)(p - t) dt$$

which is non-negative when $\gamma(t) \geq 0$.

For necessity, remark that if a scoring rule $s(q, x)$ is proper then $q = p$ maximizes $E[s(q, X)]$, where the expectation is taken under the true probability p that $X = 1$. The first-order condition yields, for all $p \in (0, 1)$,

$$ps'(p, 1) + (1 - p)s'(p, 0) = 0 .$$

Let

$$\gamma(q, x) = \frac{s'(q, x)}{q - x}$$

whenever $q \notin \{0, 1\}$. The first-order condition can then be rewritten as

$$p(p - 1)\gamma(p, 1) + p(1 - p)\gamma(p, 0) = 0, \quad \forall p \in (0, 1) ,$$

and so $\gamma(p, 1) = \gamma(p, 0) = \gamma(p)$ for some function $\gamma(\cdot)$ on $(0, 1)$. The function $\gamma(\cdot)$ can be extended to the full interval $[0, 1]$ by continuity. The fact that $\gamma(\cdot) \geq 0$ follows from the inequality

$$E[s(p, X)] - E[s(q, X)] = \int_q^p \gamma(t)(p - t)dt \geq 0 , \quad \forall p, q \in [0, 1] .$$

We now build on this result to show that $E[s(q, X)]$ is strictly decreasing in $|q - p|$ if and only if $s'(q, x) = \gamma(x - q)$ for some constant $\gamma > 0$.

For sufficiency, observe that if $s'(q, x) = \gamma(x - q)$, then $s(q, x) = -\gamma(x - q)^2/2 + \rho(x)$ for some function ρ . Hence $E[s(q, X)] - E[s(p, X)] = \gamma(p - q)^2/2$, and the result follows.

We prove necessity starting from the equality

$$E[s(p - \delta, X)] = E[s(p + \delta, X)] , \tag{5}$$

for all p and δ such that $p + \delta$ and $p - \delta$ are within the range $[0, 1]$, equality that is implied by the monotonicity condition. We already know that

$$s(q, x) = s(0, x) + \int_0^q \gamma(t)(x - t)dt .$$

Substituting this expression of $s(q, x)$ into (5) yields

$$\int_{p-\delta}^{p+\delta} (p - t)\gamma(t)dt = 0 , \quad \forall p, \delta .$$

Differentiating with respect to δ , we get

$$(p - (p + \delta))\gamma(p + \delta) + (p - (p - \delta))\gamma(p - \delta) = -\delta\gamma(p + \delta) + \delta\gamma(p - \delta) = 0 ,$$

and so $\gamma(p + \delta) = \gamma(p - \delta)$. This equality is valid whenever $0 < p - \delta < p + \delta < 1$. By continuity of $\gamma(\cdot)$, we conclude that $\gamma(p) = \gamma$ for all p , that is, γ is constant. That γ is positive follows immediately. \square

5.2 A Uniqueness Result for Sharing Rules

We next provide a uniqueness result for mechanisms that divide a *fixed* amount of money M among a set of forecasters. These mechanisms can be interpreted as scoring rules that apply to multiple forecasters, with the constraint that scores must be nonnegative and sum to a constant M .

As in the preceding sections, X denotes the indicator variable of an event E and p the probability that E occurs. There are n forecasters. In the first stage, each forecaster i reports her probability assessment p_i . In the second stage, after observing the realization x of X , each forecaster i gets a share $\pi_i(p_1, \dots, p_n, x)$ of the total amount M . A *sharing rule* is given by the vector of shares (π_1, \dots, π_n) . Sharing rules are, by definition, nonnegative, and such that the total shares $\pi_1 + \dots + \pi_n$ sum to one. We focus the analysis on sharing rules that are twice continuously differentiable with respect to the vector of probability inputs (p_1, \dots, p_n) .

We seek sharing rules that satisfy properties analogous to (1)–(3) from Section 3. In the context of sharing rules, these properties are expressed as follows:

- (a) The sharing rule should be symmetric, i.e., for every permutation σ of the forecasters,

$$\pi_{\sigma^{-1}(i)}(p_{\sigma(1)}, \dots, p_{\sigma(n)}, x) = \pi_i(p_1, \dots, p_n, x) .$$

- (b) For all i and $j \neq i$, $\mathbf{E}[\pi_i(p_1, \dots, p_n, X)]$ should be increasing as $|p_i - p|$ decreases: A forecaster gets a larger share as her report gets closer to the truth.
- (c) For all i , $\mathbf{E}[\pi_i(p_1, \dots, p_n, X)]$ should be decreasing as $|p_j - p|$ decreases: A forecaster gets a smaller share as someone else's report gets closer to the truth.

Proposition 2. *The sharing rule (π_1, \dots, π_n) defined by*

$$\pi_i(p_1, \dots, p_n, x) = \frac{1}{n} + \frac{\alpha}{n-1} \sum_{j \neq i} ((p_j - x)^2 - (p_i - x)^2)$$

for $\alpha \in [0, 1/n]$ satisfies properties (a)–(c). Furthermore, sharing rules of this form are the only twice continuously differentiable sharing rules that satisfy them all.

Proof of Proposition 2: Sufficiency is direct, observing that

$$\mathbf{E}[\pi_i(p_1, \dots, p_{i-1}, p, p_{i+1}, \dots, p_n, X)] - \mathbf{E}[\pi_i(p_1, \dots, p_n, X)] = \alpha |p_i - p|^2 ,$$

and similarly

$$\mathbf{E}[\pi_i(p_1, \dots, p_{j-1}, p, p_{j+1}, \dots, p_n, X)] - \mathbf{E}[\pi_i(p_1, \dots, p_n, X)] = -\frac{\alpha}{n-1} |p_j - p|^2 .$$

The key to proving necessity is the following lemma:

Lemma 1. *If π_i is a twice continuously differentiable sharing rule that satisfies conditions (2) and (3) above, then there exist twice continuously differentiable functions π_i^j such that*

$$\pi_i(p_1, \dots, p_n, x) = \pi_i^i(p_i, x) + \sum_{j \neq i} \pi_i^j(p_j, x) .$$

Proof of Lemma 1: The proof makes use of the following lemma, whose proof is relegated to the Appendix.

Lemma 2. *If $f : (a, b)^n \mapsto \mathbb{R}$ is twice continuously differentiable, and if*

$$\frac{\partial^2 f(x_1, \dots, x_n)}{\partial x_i \partial x_j} = 0$$

for all x_1, \dots, x_n and all $i \neq j$, then there exist $f_i : (a, b) \mapsto \mathbb{R}$, $i = 1, \dots, n$ such that

$$f(x_1, \dots, x_n) = \sum_{i=1}^n f_i(x_i) .$$

Let \mathbb{E}_p denote the expectation over X when $X = 1$ with probability p . By property (b),

$$p = \arg \max_{p_i} \mathbb{E}_p[\pi_i(p_1, \dots, p_n, X)] , \quad (6)$$

and by property (c), for all $j \neq i$,

$$p = \arg \min_{p_i} \mathbb{E}_p[\pi_j(p_1, \dots, p_n, X)] . \quad (7)$$

The first-order condition in (6) and (7) yields, for all k and all i ,

$$\mathbb{E}_{p_i} \left[\frac{\partial \pi_k(p_1, \dots, p_n, X)}{\partial p_i} \right] = 0 .$$

As the equality remains true for all values of p_j , we can differentiate a second time and get, for all k and all $i, j, i \neq j$,

$$\mathbb{E}_{p_i} \left[\frac{\partial^2 \pi_k(p_1, \dots, p_n, X)}{\partial p_i \partial p_j} \right] = 0 .$$

Note that as π_i is twice continuously differentiable function, $\partial^2 \pi_i / \partial p_i \partial p_j = \partial^2 \pi_i / \partial p_j \partial p_i$. Hence by symmetry we also have

$$\mathbb{E}_{p_j} \left[\frac{\partial^2 \pi_k(p_1, \dots, p_n, X)}{\partial p_i \partial p_j} \right] = 0 .$$

Observe that, for any random variable Y that takes value in $\{y^L, y^H\}$, if for any p, q with $p \neq q$, the expected value of Y when $Y = y^H$ has probability p is the same as the expected value of Y when $Y = y^L$ has probability p , then it must be the case that $y^L = y^H$. Applying this to the two

preceding inequalities yields, for all $i \neq j$, all (p_1, \dots, p_n) with $p_i \neq p_j$, and all x ,

$$\frac{\partial^2 \pi_k(p_1, \dots, p_n, X)}{\partial p_i \partial p_j} = 0 .$$

By continuity, the equality remains true for $p_i = p_j$. We conclude by a direct application of Lemma 2. \square

We now return to the proof of the proposition. We apply Lemma 1, and write $\pi_i = \sum_j \pi_i^j$ with the functions π_i^j as defined by the Lemma. Using property (b) and applying Proposition 1, there exist some constants $\alpha_i > 0$ and some functions g_i^j such that, for all i ,

$$\pi_i^i(p_i, x) = -\alpha_i(p_i - x)^2 + g_i^i(x) .$$

Also, by property (c), and again in application of Proposition 1, there exist some constants $\beta_i^j > 0$ and some functions g_i^j such that, for all i, j with $i \neq j$,

$$\pi_i^j(p_i, x) = \beta_i^j(p_j - x)^2 + g_i^j(x) .$$

These two equalities imply that π_i can be re-written under the form

$$\pi_i(p_1, \dots, p_n, x) = -\alpha_i(p_i - x)^2 + \sum_{j \neq i} \beta_i^j(p_j - x)^2 + g_i(x)$$

where $g_i = \sum_j g_i^j$. Fix some forecaster i . Applying property (a), the payoff of forecaster i when the reports of the other forecasters are permuted should remain the same, which yields $\beta_i^j = \beta_i$ for some β_i and all j . When can then apply property (a) to forecaster i and some other forecaster j . When these agents are permuted, the payoffs are permuted as well, which yields $\alpha_i = \alpha$, $\beta_i = \beta$, and $g_i = g$ for all i . Hence, for all p_1, \dots, p_n , and all x ,

$$\pi_i(p_1, \dots, p_n, x) = -\alpha(p_i - x)^2 + \sum_{j \neq i} \beta(p_j - x)^2 + g(x) .$$

The last step makes use of the fact that the forecasters' shares sum to one:

$$\sum_i -\alpha(p_i - x)^2 + \sum_i \sum_{j \neq i} \beta(p_j - x)^2 + ng(x) = 1 .$$

As the equality remains true for every value of p_1 , by we can differentiate each side of the equality with respect to p_1 to get

$$-2\alpha(p_1 - x) + 2(n - 1)\beta(p_1 - x) = 0 ,$$

which gives $\beta = \alpha/(n - 1)$ and so $0 + ng(x) = 1$, hence $g(x) = 1/n$.

In summary,

$$\pi_i(p_1, \dots, p_n, x) = \frac{1}{n} + \frac{\alpha}{n-1} \sum_{j \neq i} [(p_j - x)^2 - (p_i - x)^2] .$$

Property (b) implies $\alpha \geq 0$. As shares must remain nonnegative, $\alpha \leq 1/n$. \square

To prove our main theorem, we make use of Proposition 2 in its full generality. However, when using sharing rules as an independent mechanism to reward a group of forecasters, it is natural to take $\alpha = 1$, so as to get the additional property (d):

(d) For all i , there exists a profile of reports p_1, \dots, p_n and a realization x such that the share of forecaster i is zero: $\pi_i(p_1, \dots, p_n, x) = 0$.

We have the following immediate corollary:

Corollary 1. *The sharing rule (π_1, \dots, π_n) defined by*

$$\pi_i(p_1, \dots, p_n, x) = \frac{1}{n} + \frac{1}{n(n-1)} \sum_{j \neq i} ((p_j - x)^2 - (p_i - x)^2)$$

satisfies properties (a)–(d). Furthermore, it is the only twice continuously differentiable sharing rule that satisfies them all.

5.3 Return to Wagering Mechanisms

We now return our attention to wagering mechanisms and provide the proof for the uniqueness of the Brier betting mechanism.

Proof of Theorem 1: As we argued in Section 3, sufficiency is immediate and follows from the observation that

$$\mathbb{E}[\pi_i((p_1, w_1), \dots, (p_n, w_n), X)] = \frac{w_i}{W} + \kappa \frac{w_i}{W} \sum_{j \neq i} \frac{w_j}{W} ((p_j - p)^2 - (p_i - p)^2) .$$

To show necessity, we make a parallel between three different scenarios: One in which all wagers are identical, one in which wagers may vary but remain multiple of a base wager ϵ , and, finally, the real setting of interest in which wagers are arbitrary. The rules that determine the payoffs of the first setting are derived from Proposition 2. The payoffs of the second setting are then obtained using property (4) of Section 3. Finally, we get the payoffs of the third setting by viewing this general setting as a limit case of the second setting.

Let (p_1, \dots, p_n) be a profile of n predictions.

Setting 1: In the first setting, there are N players divided into n groups with k_i members each, $\sum_j k_j = N$. Each player of the group i announces the same prediction p_i , and all players in all groups wager the same amount ϵ . The total amount of money wagered is then $W = N\epsilon$. Index

players by capital letters I, J and denote by p_I the prediction of player I . The properties (1), (2), and (3) applied to this first setting imply, by proposition 2, that player I should get a share

$$\frac{1}{N} + \frac{\alpha}{N-1} \sum_{J \neq I} [(q_J - x)^2 - (q_I - x)^2]$$

of the total wager W for some $\alpha \in [0, 1/N]$. Note that α may depend on N and ϵ . In what follows, we write the dependence explicitly as $\alpha(N, \epsilon)$.

Setting 2: Now consider a second setting with only n players. Each player i chooses prediction p_i (as above) and a wager $w_i = k_i \epsilon$ for the same values of k_i used in setting 1.

The total amount of money wagered in this setting is $W = N\epsilon$, the same as in setting 1. Furthermore, each player i in this setting announces the same prediction as every member of group i in setting 1, and wagers an amount that is the sum of the wagers of all members of this group. Therefore, by application of property (4), the share that player i receives in this setting must equal the total share that group i receives in setting 1, which is

$$\frac{k_i}{N} + \frac{\alpha(N, \epsilon)k_i}{N-1} \sum_{j \neq i} k_j [(q_j - x)^2 - (q_i - x)^2] .$$

This expression can be rewritten as

$$\frac{w_i}{W} + \beta(N, \epsilon) \frac{w_i}{W} \sum_{j \neq i} \frac{w_j}{W} [(q_j - x)^2 - (q_i - x)^2] \quad (8)$$

where we define

$$\beta(N, \epsilon) = \frac{N^2}{N-1} \alpha(N, \epsilon) .$$

By property (5), the shares do not depend on ϵ so that $\beta(N, \epsilon)$ is only a function of N , and we write $\beta(N, \epsilon) = \beta(N)$.

Setting 3: In the final step, we return to the initial problem of computing shares for players with unrestricted wagers. This third setting can be seen as a limit case of the second setting. Each player i wagers any amount $w_i > 0$ (the case $w_i = 0$ can be obtained by continuity) and predicts p_i . We compute the share of player i as follows.

Let $\epsilon^t = 1/t$ and for each player i , consider any sequence of positive integers $\{k_i^t\}_{t=1,2,\dots}$ such that $k_i^t/t \rightarrow w_i$ as $t \rightarrow +\infty$. Let $N^t = \sum_i k_i^t$. Each particular value of t corresponds to setting 2 with the corresponding values $\epsilon = \epsilon^t$ and $k_i = k_i^t$. Naturally $N^t \rightarrow +\infty$ as $t \rightarrow +\infty$. By (8), and the continuity assumption of the sharing rule, the limit of

$$\frac{w_i}{W} + \beta(N^t) \frac{w_i}{W} \sum_{j \neq i} \frac{w_j}{W} [(q_j - x)^2 - (q_i - x)^2]$$

as $t \rightarrow \infty$ exists and equals the share of player i in this third setting. Hence the limit $\lim_{N \rightarrow \infty} \beta(N)$ exists, and we conclude the proof by taking $\kappa = \lim_{N \rightarrow \infty} \beta(N)$. That $\kappa \in [0, 1]$ follows from the

necessity to have nonnegative shares and by property (2). □

6 The Case of Bayesian Individuals

The main framework of this paper assumes that individuals hold independent, immutable prior beliefs, as in the works of Eisenberg and Gale (1959) and Ali (1977) on betting markets. An individual's subjective probability is not affected by the actions of other players; everybody acts as if they are very confident about their own prediction. The assumption of immutable beliefs facilitates the analysis as we need not be concerned about the formation of the posteriors.

Our framework can be interpreted as an extreme case of information structure. There are, of course, many other relevant structures, in which individuals can learn from others' actions. In one other extreme, individuals hold a common prior about the states of the world, a setting common to many models of dynamic market mechanisms. While our mechanism is designed for one particular information structure, it is of natural interest to investigate its performance and robustness when one departs from the assumption of differing immutable prior beliefs.

The purpose of the current section is to illustrate the functioning of our mechanism in the case of players who share a common prior belief. Such an information structure is used for example in papers by Koessler et al. (2008) and Ottaviani and Sørensen (2009, 2010) in a similar context. In this configuration players are directly influenced by the actions of others. In general common prior settings, we show that players no longer announce a truthful prediction. This is true of the Brier betting mechanism, but also of any non-trivial wagering mechanism. However, in the case of the Brier betting mechanism, we are able to quantify the bias and show that predictions tend to reflect true personal assessments as the size of the betting market grows. This asymptotic property is independent of the information partitions/common prior being used.

Model. Let (Ω, \mathcal{F}, P) be a probability space, where Ω denotes the set of possible states of the world and P is a common prior. All expectations of this section are taken with respect to P . There are n risk-neutral players. Player i holds private information on the true state of the world, which leads to heterogeneous private beliefs about the likelihood of the event. As is common, to allow for full generality, we represent this private information by a sigma algebra \mathcal{F}_i . The set \mathcal{F}_i can be interpreted as the collection of all the events that player i observes privately. For example, if player i observes a real-valued signal S_i , then the set \mathcal{F}_i consists in the sigma algebra generated by the family of events $\{S_i \in [a, b]\}$, where $[a, b]$ is any interval.

Players are offered to bet in the Brier betting mechanism. Denote by E the event of interest. A player's prediction consists in a probability assessment that E occurs. Denote by X the indicator variable of E . By convention, each player may decide not to participate by wagering a zero amount. The mechanism then disregards the prediction. However if she decides to participate, the player must wager an amount within the range $[W^L, W^H]$, where W^L and W^H are mechanism parameters that specify respectively the lowest and highest amount players can wager. Imposing lower and

upper bounds on wagers is common in practice. Limiting the gains and losses makes our analysis tractable: It permits to assume risk neutrality, an hypothesis empirically validated in the context of horse-race parimutuel betting (Ali, 1977). To induce voluntary participation, and circumvent No-Trade Theorems (Milgrom and Stokey, 1982), we treat the betting game as a consumption good. Players derive a positive utility from betting, owing for example to entertainment value. If she participates, player i gets, in addition to her payoff, a utility amount equivalent to U_i dollars.

Equilibrium Analysis. In this setting a natural solution concept is that of *Bayes-Nash equilibrium*. A *pure strategy* for player i is a pair of random variables (p_i, w_i) , where $p_i : \Omega \mapsto [0, 1]$ gives a prediction—the assessment of event E 's likelihood—and $w_i : \Omega \mapsto \{0\} \cup [W^L, W^H]$ returns the wager. Both functions depend on the state and are required to be \mathcal{F}_i -measurable,² so that both the prediction and the wager of player i depend only on her private information. A *mixed strategy* is a distribution over pure strategies. For simplicity of exposition we focus on pure strategy equilibria; the analysis extends directly to equilibria in mixed-strategy. A strategy profile $\{(p_i, w_i)\}_{i=1}^n$ is an (ex-interim) Bayes-Nash equilibrium when, for every player i , and every alternative strategy $(\tilde{p}_i, \tilde{w}_i)$ of that player,

$$\begin{aligned} \mathbb{E} \left[\sum_{j \neq i} \frac{w_i w_j}{w_i + W_{-i}} ((p_j - X)^2 - (p_i - X)^2) \middle| \mathcal{F}_i \right] + U_i \mathbb{1}\{w_i > 0\} \geq \\ \mathbb{E} \left[\sum_{j \neq i} \frac{\tilde{w}_i w_j}{\tilde{w}_i + W_{-i}} ((p_j - X)^2 - (\tilde{p}_i - X)^2) \middle| \mathcal{F}_i \right] + U_i \mathbb{1}\{\tilde{w}_i > 0\} , \quad (9) \end{aligned}$$

where $W_{-i} = \sum_{j \neq i} w_j$ is the total amount other players wager.

Rearranging the terms and using $W = \sum_j w_j$ as the total wager, we get that

$$\begin{aligned} \mathbb{E} \left[\sum_{j \neq i} \frac{w_i w_j}{W} ((p_j - X)^2 - (p_i - X)^2) \middle| \mathcal{F}_i \right] = \\ \mathbb{E} \left[\sum_{j \neq i} \frac{w_i w_j}{W} (p_j - X)^2 \middle| \mathcal{F}_i \right] - \mathbb{E} \left[w_i \left(1 - \frac{w_i}{W}\right) (p_i - X)^2 \middle| \mathcal{F}_i \right] , \end{aligned}$$

so that at equilibrium p_i minimizes the last term

$$\mathbb{E} \left[w_i \left(1 - \frac{w_i}{W}\right) (p_i - X)^2 \middle| \mathcal{F}_i \right] .$$

The first order condition yields

$$\mathbb{E} \left[w_i \left(1 - \frac{w_i}{W}\right) (p_i - X) \middle| \mathcal{F}_i \right] = 0 .$$

²Recall that a function is \mathcal{G} -measurable when observing the occurrence or nonoccurrence of each event of \mathcal{G} is enough to fully determine the value of that function.

Hence, at equilibrium, the prediction of player i is

$$p_i = \frac{\mathbb{E}[(1 - w_i/W)X \mid \mathcal{F}_i]}{1 - \mathbb{E}[w_i/W \mid \mathcal{F}_i]} = \mathbb{E}[X \mid \mathcal{F}_i] - \frac{\text{Cov}(w_i/W, X \mid \mathcal{F}_i)}{1 - \mathbb{E}[w_i/W \mid \mathcal{F}_i]},$$

where $\text{Cov}(w_i/W, X \mid \mathcal{F}_i)$ denotes the covariance between w_i/W and X conditionally on information \mathcal{F}_i . It is easily verified that $0 \leq p_i \leq 1$.

Therefore, at equilibrium, each participating player i announces a prediction that is the sum of her true probability assessment $\mathbb{E}[X \mid \mathcal{F}_i]$ and a bias

$$\epsilon_i = -\frac{\text{Cov}(w_i/W, X \mid \mathcal{F}_i)}{1 - \mathbb{E}[w_i/W \mid \mathcal{F}_i]}. \quad (10)$$

The bias depends only on how the total amount wagered correlates with the outcome. If participation increases with the occurrence of the event, players tend to overestimate probabilities. Conversely if the total amount wagered decreases with the occurrence of the event, players tend to underestimate probabilities. The bias owes to the structure of the payoffs. A prediction has greater impact on a player's payoff when the total amount wagered is large. Depending on the information structure, the total amount wagered may provide information on the state of the world. Consequently the optimal prediction may depend on the total amount wagered. As the total amount wagered is uncertain, the optimal prediction is a weighted average of the predictions that are separately optimal conditionally on every possible total wager W . Since greater amounts W have stronger impact on a player's payoff, the player should put more weight on the predictions optimal for these amounts, which introduces a bias.

However biases are typically small can often be ignored. As we argue in the two examples below, they are only noticeable under some cases of asymmetric information. Besides if m is the number of participants, observing that

$$|\epsilon_i| \leq \frac{\mathbb{E}[w_i/W \mid \mathcal{F}_i]}{1 - \mathbb{E}[w_i/W \mid \mathcal{F}_i]} = O(1/m),$$

the bias becomes rapidly negligible as the number of participants grows. Thus predictions tend to reflect accurately the players' beliefs in large markets, no matter the underlying information structure. Note that the problem of biased predictions is intrinsic to all non-trivial wagering mechanisms. It is due to the budget-balance constraint, which implies that the same predictions impact payoffs differently for different given wagers of the other players.

At equilibrium, the wager for player i , w_i , maximizes

$$\mathbb{E} \left[\sum_{j \neq i} \frac{w_i w_j}{W} ((p_j - X)^2 - (p_i - X)^2) \mid \mathcal{F}_i \right] + U_i \mathbb{1}\{w_i > 0\} \text{ with } p_i = \frac{\mathbb{E}[(1 - w_i/W)X \mid \mathcal{F}_i]}{1 - \mathbb{E}[w_i/W \mid \mathcal{F}_i]}$$

Because the net payoffs from betting are always within the range $[-w_i, w_i]$, if player i participates she is always guaranteed utility at least $U_i - W^L$. In particular player i is better off participating

whenever $U_i > W^L$. If $U_i < W^L$, the player may or may not decide to participate, depending on how her private information compares to that of others. She will only participate if she finds herself sufficiently knowledgeable compared to the other players.

If equilibrium wagers are independent of the state of the world, the monotonicity of the payoffs discussed in Section 3 implies that each player should wager either 0, W^L or W^H . She wagers 0 if she is poorly informed and has low value for participation, she wagers W^L if she is poorly informed and has high value for participation, and she wagers W^H if she is well informed, no matter her value for participation. However in general equilibrium wagers need not be restricted to these three values. They can lie *a priori* anywhere in between the minimum and maximum wagers W^L and W^H . Indeed, for a given state of the world, marginal payoffs have the same sign. Hence if a player finds it best to participate, the optimal wager is either to bet the maximum amount W^H , and get a positive wealth transfer from the other players, or to bet the minimum amount W^L and subsidize the other players. In contrast, when the state of the world is uncertain, the marginal payoffs may change sign, because the marginal payoffs for the different states of the world need not vary at the same rate. Therefore conditional on a player's private information, at a particular wager w , betting one extra dollar may generate additional gains, while at a wager $w' > w$, betting one extra dollar may cause a loss. Naturally more informed players always wager more than less informed players, in the sense that, if $\mathcal{F}_i \subseteq \mathcal{F}_j$, then $w_j \geq w_i$.

An example with biased predictions. Biased predictions can only appear under asymmetric information when a player's participation depends on her private observations. For instance, when a player can receive good information in some states of the world, say \mathcal{S} , and poor information in others, to the extent that she is willing to participate only when she receives good information. A large total wager W is then an indicator that the states of the world \mathcal{S} are more likely. At the time bets are being submitted, players do not know the total wager W . However, conditional on W being large, a player should update her belief towards higher probabilities for \mathcal{S} . As we argued above, predictions conditional on large W should carry more weight than those conditional on small W . Hence a player tends to report probabilities under the assumption that \mathcal{S} is more likely than her personal assessment would suggest, causing a bias.

To illustrate our argument, consider a binary outcome X and three conditionally independent binary signals S_1, S_2, S_3 . (In this setting a state of world is simply a joint realization of all the signals and the binary outcome.) The common prior P assigns equal probability to $X = 1$ and $X = 0$. There are three players. Every player i privately observes signal S_i . Signal S_1 is completely uninformative with $P(S_1 = 1|X = 1) = P(S_1 = 1|X = 0) = 1/2$. Signal S_3 is perfectly informative with $P(S_3 = 1|X = 1) = P(S_3 = 0|X = 0) = 1$. Finally, signal S_2 is informative in some states and uninformative in some others: $P(S_2 = 1|X = 1) = 1$ but $P(S_2 = 1|X = 0) = 1/2$. Take $W^L = 1$ and $W^H = 10$ as lower and upper wager limits, respectively. Let the participation utility for player 1 be $U_1 > W^L$ (player 1 is here for entertainment purpose, in effect she subsidizes the market) and for the other two players $U_2 = U_3 = 0$ (other players only play for profit).

We first observe that in a Bayes-Nash equilibrium, the wagers of player 1 and 3 are respectively $w_1 = 1$ and $w_3 = 10$. Indeed player 1 is always better off participating because of her high entertainment value. However, since the other two players are more informed, player 1 is always losing money in equilibrium, and so bets the minimum amount allowed to minimize her losses. On the contrary player 3 is the best informed of all players. Since she always knows the true outcome she never loses money, and should wager the maximum amount to maximize her gains. In equilibrium player 2 takes an action that depends on how much she knows. If she observes signal $S_2 = 0$, then she knows that $X = 0$ and is, like player 3, perfectly informed. She should wager the maximum to get the most she can from player 1. If however $S_2 = 1$, she is poorly informed. Using the fact that, at equilibrium, $p_3 = X$, player 2's expected net incoming transfer conditional on her information $S_2 = 1$ is therefore

$$\mathbb{E} \left[\frac{w_2}{W} \sum_{j=1,3} w_j ((p_j - X)^2 - (p_i - X)^2) \middle| S_2 = 1 \right] \leq \frac{w_2 w_1}{W} - \frac{w_2 w_3}{W} \cdot \frac{2}{9} < 0 .$$

No matter how player 1 chooses her prediction, player 2 is guaranteed to lose money on average when she receives signal $S_2 = 1$, and therefore should abstain from participation. Hence, in this particular setting, the total wager at equilibrium depends on the true state of the world.

We now turn our attention to the equilibrium predictions. Player 3 being perfectly informed, she reports truthfully the probabilities 0 or 1. When player 2 participates she is perfectly informed and reports truthfully probability 0 or 1 as well. Player 1's true assessment of the probability that $X = 1$ is 50%. However in equilibrium she announces the prediction p_1 that minimizes her expected loss

$$\begin{aligned} & P(S_2 = 0) \mathbb{E} \left[\frac{w_1}{W} \sum_{j=2,3} w_j ((p_i - X)^2 - (p_j - X)^2) \middle| S_2 = 0 \right] \\ & \quad + P(S_2 = 1) \mathbb{E} \left[\frac{w_1}{W} \sum_{j=2,3} w_j ((p_i - X)^2 - (p_j - X)^2) \middle| S_2 = 1 \right] \\ & = P(S_2 = 0) \cdot \frac{20}{21} \cdot p_1^2 \\ & \quad + P(S_2 = 1) \cdot \frac{10}{11} (P(X = 0 | S_2 = 1) p_1^2 + P(X = 1 | S_2 = 1) (1 - p_1)^2) , \end{aligned}$$

which yields $p_1 = 42/85 \approx 49\% < 50\%$. Therefore in equilibrium player 1 is not truthful, even though the bias remains small enough for most practical purposes. Indeed when the total wager W is low player 1's participation is rewarded less than when W is high. High wagers indicate a signal $S_2 = 0$. Since the probability of outcome $X = 1$ conditional on $S_2 = 0$ is less than when the probability of $X = 1$ conditioned upon $S_2 = 1$, player 1 should bias her reports towards lower probabilities.

An example with truthful predictions. In contrast, when the information structure displays enough symmetry, equilibrium predictions are unbiased. Predictions are truthful even if the quality of information is heterogeneous *across* players, as long as it does not vary *within* players. Consider for example a setting with n players. As before, X is a binary outcome and P a common prior that assigns equal probabilities to both outcomes. Each player i has access to a binary signal S_i , where $P(S_i = 1|X = 1) = P(S_i = 0|X = 0) = \mu_i > 1/2$. High values μ_i indicate better information. Without loss of generality we rank the players in order of increasing knowledge by imposing $\mu_1 \leq \dots \leq \mu_n$. Signals are conditionally independent. For simplicity consider $U_i > W^L$, i.e., every player is always willing to participate, no matter how informed.

We show that there exists an equilibrium in which every player predicts truthfully according to her private information; that is, player i 's prediction strategy is $p_i = P(X = 1|S_i)$. As predictions are always unbiased in equilibria in which every player wagers independently of her signal, we seek to construct such an equilibrium.

Players with imprecise signals get positive utility from entertainment value only, but lose money on average. In effect they subsidize the market and transfer part of their wealth to the players with precise signals. Hence if such an equilibrium exists, players who are poorly informed, that is, players with low indices, should wager the minimum amount W^L . In contrast players who are well informed, those with high indices, should wager the maximum amount W^H . (Recall that when wagers do not depend on the state of the world a player's payoff is monotone in the wager.)

We examine each of these strategy profiles separately and show one of them is indeed an equilibrium. Specifically, consider n truthful strategy profiles indexed by $I = 1, \dots, n$. In each of these profiles players always predict truthfully, but wager different amounts. In profile I , every player $i \leq I$ wagers W^L while every player $i > I$ wagers W^H .

The wagers are independent of the players' signals, and by symmetry,

$$\mathbb{E}[(p_j - X)^2 - (p_i - X)^2 | S_i = 1] = \mathbb{E}[(p_j - X)^2 - (p_i - X)^2 | S_i = 0].$$

Therefore, given any such a strategy profile I , the expected net incoming transfer to player i , conditional on her information, equals

$$\begin{aligned} t_i^I &= \mathbb{E} \left[\sum_j \frac{w_i w_j}{W} ((p_j - X)^2 - (p_i - X)^2) \mid S_i \right] \\ &= w_i \sum_j \frac{w_j}{W} \mathbb{E}[(p_j - X)^2 - (p_i - X)^2 \mid S_i] \\ &= w_i \sum_j \frac{w_j}{W} \mathbb{E}[(p_j - X)^2 - (p_i - X)^2] \\ &= w_i \left(\sum_j \frac{w_j}{W} \mu_j (1 - \mu_j) \right) - w_i \mu_i (1 - \mu_i) \end{aligned}$$

where we observe that $\mathbb{E}[(p_i - X)^2] = \mu_i(1 - \mu_i)$ for any player i .

For every strategy profile I , define $f(I)$ to be the highest index i such that $t_i \leq 0$. By the above equality, t_i^I/w_i is increasing with i , and $t_1^I \leq 0$ while $t_n^I \geq 0$. In other words, a player's expected share increases with the precision of her information, and players who are less informed make a nonpositive profit while players who are more informed make a nonnegative profit. This implies the index $f(I)$ is such that, for all $i \leq f(I)$, $t_i^I \leq 0$, and for all $i > f(I)$, $t_i^I > 0$. For the strategy profile I to be an equilibrium, it suffices that all the players with an index up to $f(I)$ wager the minimum amount allowed W^L (as these players do not make a positive profit from betting), while the players with an index greater than $f(I)$ wager the maximum amount W^H (as these players make a positive profit from betting that is maximized when they wager the maximum amount). Hence if we can find some J with $f(J) = J$, it means strategy profile J is an equilibrium. To show existence of such an index value, let $g(I) = f(I) - I$. As $g(1) \geq 0$ and $g(n) \leq 0$, there must exist some I such that $g(I) \geq 0$ and $g(I+1) \leq 0$. Noting that f is increasing as t_i^I/w_i is also increasing in I , we have that either $f(I) = I$ or $f(I+1) = I+1$. Therefore at least one of the strategy profiles $I \in \{1, \dots, n\}$ is an equilibrium.

7 Discussion

7.1 Eliciting Expert Advice

Although our analysis was made in the context of betting markets, the Brier betting mechanism can also be viewed as a mechanism to elicit information from a group of experts. In this alternative context, an elicitor queries a group of n forecasters regarding probability estimates of an event of interest E . As before, X , indicator variable of the event E , denotes the outcome. Each forecaster i announces a probability estimate p_i of E . In contrast to the setting considered in Section 5.2, we allow for the forecaster to incur a loss—as we argue below, this is a desirable feature in that it discourages the uninformed forecasters to participate. In order to do so, forecaster i must provide a deposit $w_i \geq 0$ (instead of a wager). She may decide not to participate by setting her deposit to zero. The deposit can be interpreted as a worst-case loss or a maximum liability amount, in that the forecaster is guaranteed to lose no more than the amount deposited. As forecasters may have different wealths, it is natural to grant forecasters the right to specify their own liability amount—in particular, the mechanism does not prevent forecasters with low liability from participating. The larger the declared liability, the larger the potential losses, but also the larger the potential gains.

The expert elicitation procedure differs from betting markets in that we permit another party, the elicitor, to subsidize the market. A subsidy is often necessary to induce participation: In this new context a forecaster's utility reduces to her expertise fee, forecasters typically do not get the additional entertainment value of the betting markets.³ The subsidy can also incentivize forecasters to exert a potentially costly effort to retrieve high-quality information. However a complete model with costly information acquisition is outside the scope of this paper.

³Although in some cases forecasters who want to establish a reputation for themselves may wish to participate even when they incur a loss.

There are two natural ways to introduce the subsidy. First, the elicitor can give away a fixed lump sum payment M , as in Section 5.2. Each forecaster i receives the net payoff $(M + W)\pi_i - w_i$, where the share π_i is defined as before according to Equation (1), and W is the total sum deposited by the forecasters. In this variant, we only alter the budget-balanced property of the Brier betting mechanism; the net payoffs now sum to M . The properties (1)–(5) remain valid and our uniqueness result still applies. Low values of M dissuade uninformed forecasters to participate, because their deposit would finance the informed forecasters. As M grows larger, the mechanism attracts more forecasters, as even a small share can generate large positive profits if the amount to share, $M + W$, is large enough.

Alternatively, the elicitor can play a more active part in the mechanism. Acting as a forecaster, she can provide a prior probability estimate p_0 that E comes true, and can deposit $w_0 = M$ to subsidize the market. The payoff functions remain the same as in the Brier betting mechanism, accounting for this new player. As opposed to the other forecasters, the elicitor does not expect to generate a profit. The subsidy plays the same role as in the above mechanism. However in this second mechanism the elicitor may not lose all of M . The actual loss depends on how the forecasters' predictions compare with her own. Specifically, if E occurs with probability p , the elicitor gives away the expected amount

$$M \frac{w_i}{W} [(p_0 - p)^2 - (p_i - p)^2]$$

to forecaster i , where $W = M + \sum_{i \geq 1} w_i$. So forecaster i can only benefit from the subsidy when she is more accurate than the elicitor, with respect to the absolute difference between the probability estimate and the truth. Consequently if forecasters derive no personal value from participation, this alternative mechanism deters all forecasters who are less informed than the elicitor. Note that transfers between forecasters also occur, so that the actual reward of a forecaster depends not only on how well she performs against the elicitor, but also how well she performs against the other forecasters. However as forecasters can only get positive transfers from the less informed forecasters, eventually no forecaster less informed than the elicitor will want to participate. If M is low, the mechanism can even deter some forecasters who are slightly more informed than the elicitor, but significantly less informed than some other forecasters. Transfers between forecasters become negligible only when the subsidy grows large in comparison to the forecasters' aggregate liability.

7.2 Generalization of the Brier Betting Mechanism

The Brier betting mechanism is based on the Brier score which is a proper scoring rule. In principle we can, however, plug any proper scoring rule $s(q, x)$ taking values in the range $[0, 1]$ into Equation (1) to derive an alternative betting mechanism. Properties (1), (4) and (5) continue to hold, but properties (2) and (3) are specific to the Brier score. Using the fact that $s(q, x)$ is a proper scoring rule, we immediately get that the following weaker analogs are satisfied:

(2*) For a given action profile $(p_1, w_1), \dots, (p_n, w_n)$, if player i changes her prediction p_i to the true probability p , then her expected share under p increases.

(3*) For a given action profile $(p_1, w_1), \dots, (p_n, w_n)$, the expected share of player i decreases if any other player j changes her prediction from p_j to p .

If the scoring rule is strictly proper, all increases/decreases are strict. Proposition 3 of Nau (1985) shows that these two properties are respectively equivalent to seemingly stronger properties:

(2') For a given action profile $(p_1, w_1), \dots, (p_n, w_n)$, if player i changes her prediction p_i to a convex combination of p_i and p , then her expected share increases.

(3') For a given action profile $(p_1, w_1), \dots, (p_n, w_n)$, the expected share of player i decreases when another player j changes her prediction p_j to a convex combination of p_j and p .

If the scoring rule is strictly proper and the convex combinations are strict, all increases and decreases are strict. These last two properties are a natural generalization of the properties (2) and (3) without using a notion of distance between predictions and true probabilities.

The proofs of Proposition 2 and our main result Theorem 1 extend to this more general setting with minor modifications and yield the following result.

Theorem 2. *Consider a wagering mechanism defined by the redistribution rule (π_1, \dots, π_n) in which every function π_i is twice continuously differentiable. The mechanism satisfies the properties (1), (2*), (3*), (4) and (5) if and only if*

$$\pi_i((p_1, w_1), \dots, (p_n, w_n), x) = \frac{w_i}{W} + \kappa \frac{w_i}{W} \sum_{j \neq i} \frac{w_j}{W} (s(p_i, x) - s(p_j, x))$$

for some $\kappa \in [0, 1]$ and some twice continuously differentiable proper scoring rule $s(q, x)$ taking values in $[0, 1]$.

In other instances a market designer may be interested in statistics such as the mean, the median or the 95th quantile. While probability scoring rules elicit event probabilities, a library of more general scoring rules which are proper for general statistics is well-known and easily used here. The Brier betting mechanism and its set of properties can be adapted to allow for predictions of such statistics by plugging the appropriate scoring rule in Equation (1). We start by presenting the extension under its most general form, then we provide more concrete examples.

Let Ω be a set of outcomes and \mathcal{D} be a set of probability distributions over Ω . Following Lambert et al. (2008), we define a general *statistic* as a function $\Gamma : \mathcal{D} \mapsto \Theta$, where Θ is the set in which the statistic takes values, and $\Gamma(P)$ is the value the statistic takes under P . The set Θ can be arbitrary, typically \mathbb{R} or \mathbb{R}^k for single or multi-dimensional real-valued statistics. For example, if Γ denotes the mean of a random variable X , Γ can be expressed as $\Gamma(P) = \int X dP$.

A *scoring rule* for Γ is a function $s(\theta, \omega)$ that returns a real-valued score as a function of the realized outcome ω and a statistic estimate θ of the distribution over outcomes. The scoring rule s

is *proper* if reporting true statistic values maximize the expected score, that is, for all $P \in \mathcal{D}$ and all $\theta \neq \Gamma(P)$,

$$\mathbb{E}_P[s(\Gamma(P), \omega)] \geq \mathbb{E}_P[s(\theta, \omega)]$$

where \mathbb{E}_P is the expectation taken under the distribution P . The scoring rule s is *strictly proper* if the maximization is strict.⁴

Let Γ be a statistic and $s(\theta, \omega)$ be a proper scoring rule for Γ that takes values in $[0, 1]$. We extend the mechanism of Section 3 as follows:

- In stage one, every player i places a bet, which consists of a wager $w_i \geq 0$ and an estimate θ_i for statistic Γ .
- In stage two, the realized outcome ω is publicly observed. Every player i receives a fraction of the total amount wagered, based on the scoring rule s , according to the formula:

$$\pi_i((\theta_1, w_1), \dots, (\theta_n, w_n), x) = \frac{w_i}{W} + \frac{w_i}{W} \left(s(\theta_i, x) - \sum_j \frac{w_j}{W} s(\theta_j, x) \right) \quad (11)$$

where $W = \sum_j w_j$ is the total amount wagered.

By construction, the resulting mechanism is budget balanced, and properties (1), (4), and (5) continue to hold. Neither properties (2) and (3) nor their analogs (2') and (3') hold in general, but the properties (2*) and (3*) continue to hold:

- (2*) For a given action profile $(\theta_1, w_1), \dots, (\theta_n, w_n)$, if player i changes her prediction θ_i to the true value θ , then her expected share under P increases.
- (3*) For a given action profile $(\theta_1, w_1), \dots, (\theta_n, w_n)$, the expected share of player i decreases if any other player j changes her prediction from θ_j to θ .

If the scoring rule is strictly proper, all increases/decreases are strict.⁵

Consider, for example, as set of outcomes x the interval $[0, 1]$. Let \mathcal{D} be the set of probability distributions that are represented by a positive density function. If the mechanism designer is interested in the median m , she can employ the strictly proper scoring rule

$$s(m, x) = 1 - |m - x|$$

⁴Lambert et al. (2008) and Lambert (2013) provide the statistics for which strictly proper scoring rules exist, and show how to construct these scoring rules.

⁵Using the characterizations of Lambert (2013), it can be shown that our uniqueness result in Theorem 1 of the current paper continues to hold for continuous real-valued statistics, under either properties (2') and (3') or properties (2*) and (3*).

to which corresponds the redistribution rule

$$\begin{aligned}\pi_i((m_1, w_1), \dots, (m_n, w_n), x) &= \frac{w_i}{W} + \frac{w_i}{W} \sum_j \frac{w_j}{W} (s(m_i, x) - s(m_j, x)) \\ &= \frac{w_i}{W} + \frac{w_i}{W} \sum_j \frac{w_j}{W} (|m_j - x| - |m_i - x|) .\end{aligned}$$

If instead she is interested in the mean \bar{m} and variance v , she can use the strictly proper scoring rule

$$s(\bar{m}, v, x) = 2 - |\bar{m} - x|^2 - |\bar{m}^2 + v - x^2|^2 ,$$

or if she is interested in the 5% and 95% quantiles α_5 and α_{95} she can take the strictly proper scoring rule

$$s(\alpha_5, \alpha_{95}, x) = \frac{2}{3} + \frac{3}{10}(f(\alpha_{95}) - f(\alpha_5)) - \frac{1}{3}|f(\alpha_5) - x| - \frac{1}{3}|f(\alpha_{95}) - x| .$$

8 Conclusion

In a broad study conducted over two decades, Tetlock (2006) finds that the predictions of individual forecasters, even domain experts, consistently perform worse than simple baselines. Empirically, the average opinion of experts handily beats the average expert's opinion, with accuracy increasing in the number of inputs (Forsythe et al., 1992; Jacobs, 1995; Surowiecki, 2004; Reeves and Pennock, 2007; Chen et al., 2005; Dani et al., 2006; Page, 2007). Traditional scoring rules can elicit each expert's private belief, but the elicitor's cost grows with the number of participants. Market scoring rules (Hanson, 2003, 2007), or sequential shared scoring rules, where each new trader accepts a scoring rule and pays off the previous, are more cost-effective but still require a subsidy to pay off the final trader and may fail to induce information revelation in some scenarios.

We propose a class of wagering mechanisms: The Brier betting mechanism and its extensions. Agents are called to report some information about an uncertain outcome. Along with their report, they deposit an amount of money in a common pot. Upon realization of the outcome, agents receive a payment that depends on the true outcome and their own report. Because they are self-financing, wagering mechanisms like ours naturally cater to the kind of large, open groups that appear most effective in practice.

We show that the Brier betting mechanism satisfies a number of desirable properties that scoring rules and prediction markets do not, and show that it is the only mechanism that satisfies all of these properties. In a Bayesian setting in which actions of some agents may influence the assessments of some other agent, we show that the Brier betting mechanism exhibits a bias when agents are asymmetrically informed. However the bias, intrinsic to all wagering mechanisms, is typically small and vanishes to zero at a rate inversely proportional to the number of agents, and independently of the information sets of the agents.

References

- M.M. Ali. Probability and utility estimates for racetrack bettors. *Journal of Political Economy*, 85(4):803–815, 1977.
- G.W. Brier. Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78(1):1–3, 1950.
- Y. Chen, C-H. Chu, T. Mullen, and D.M. Pennock. Information markets vs. opinion pools: An empirical comparison. In *ACM Conference on Electronic Commerce*, pages 58–67, 2005.
- Y. Chen, S. Dimitrov, R. Sami, D.M. Reeves, D.M. Pennock, R.D. Hanson, L. Fortnow, and R. Gonen. Gaming prediction markets: Equilibrium strategies with a market maker. *Algorithmica*, 58(4):930–969, 2010.
- Robert T. Clemen and Robert L. Winkler. Combining probability distributions from experts in risk analysis. *Risk Analysis*, 19(2):187–203, 1999.
- V. Dani, O. Madani, D.M. Pennock, S.K. Sanghai, and B. Galebach. An empirical comparison of algorithms for aggregating expert predictions. In *Conference on Uncertainty in Artificial Intelligence*, 2006.
- A.P. Dawid. The well-calibrated Bayesian. *Journal of the American Statistical Association*, pages 605–610, 1982.
- E. Eisenberg and D. Gale. Consensus of subjective probabilities: The parimutuel method. *The Annals of Mathematical Statistics*, 30(1):165–168, 1959.
- R. Forsythe and R. Lundholm. Information aggregation in an experimental market. *Econometrica*, 58(2):309–347, 1990.
- R. Forsythe, F. Nelson, G.R. Neumann, and J. Wright. Anatomy of an experimental political stock market. *American Economic Review*, 82(5):1142–1161, 1992.
- D.P. Foster and R.V. Vohra. Asymptotic calibration. *Biometrika*, 85(2):379–390, 1998.
- Simon French. Group consensus probability distributions: A critical survey. *Bayesian Statistics*, 2: 183–202, 1985.
- D. Friedman. Effective scoring rules for probabilistic forecasts. *Management Science*, 29(4):447–454, 1983.
- Christian Genest and James V. Zidek. Combining probability distributions: A critique and an annotated bibliography. *Statistical Science*, 1(1):114–148, 1986.
- T. Gneiting and A.E. Raftery. Strictly proper scoring rules, prediction and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.

- S.J. Grossman. An introduction to the theory of rational expectations under asymmetric information. *Review of Economic Studies*, 48(4):541–559, 1981.
- R. Hanson. Combinatorial information market design. *Information Systems Frontiers*, 5(1):107–119, 2003.
- R. Hanson. Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets*, 1(1):3–15, 2007.
- A.D. Hendrickson and R.J. Buehler. Proper scores for probability forecasters. *The Annals of Mathematical Statistics*, 42(6):1916–1921, 1971.
- W. Hurley and L. McDonough. A note on the Hayek Hypothesis and the favorite-longshot bias in parimutuel betting. *American Economic Review*, 85(4):949–955, 1995.
- R.A. Jacobs. Methods for combining experts’ probability assessments. *Neural Computation*, 7(5):867–888, 1995.
- D.J. Johnstone. The parimutuel Kelly probability scoring rule. *Decision Analysis*, 4(2):66–75, 2007.
- Victor Richmond R Jose, Robert F Nau, and Robert L Winkler. Sensitivity to distance and baseline distributions in forecast evaluation. *Management Science*, 55(4):582–590, 2009.
- D.M. Kilgour and Y. Gerchak. Elicitation of probabilities using competitive scoring rules. *Decision Analysis*, 1(2):108–113, 2004.
- F. Koessler, C. Noussair, and A. Ziegelmeyer. Parimutuel betting under asymmetric information. *Journal of Mathematical Economics*, 44(7-8):733–744, 2008.
- A. Lallour. Speed of information aggregation in an automated market mechanism. *Working paper*, 2011.
- N.S. Lambert. Elicitation and evaluation of statistical forecasts. *Working paper*, 2013.
- N.S. Lambert, D.M. Pennock, and Y. Shoham. Eliciting properties of probability distributions. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 129–138, 2008.
- R.E. Lucas. Expectations and the neutrality of money. *Journal of Economic Theory*, 4(2):103–24, 1972.
- J. McCarthy. Measures of the value of information. *Proceedings of the National Academy of Sciences of the United States of America*, 42(9):654–655, 1956.
- P. Milgrom and N. Stokey. Information, trade and common knowledge. *Journal of Economic Theory*, 26:17–27, 1982.
- R.F. Nau. Should scoring rules be ‘effective’? *Management Science*, 31(5):527–535, 1985.

- T. Norvig. Consensus of subjective probabilities: A convergence theorem. *Annals of Mathematical Statistics*, 38(1):221–225, 1967.
- W. Olszewski and A. Sandroni. Manipulability of future-independent tests. *Econometrica*, 76(6):1437–1466, 2008.
- Kent Osband. Optimal forecasting incentives. *Journal of Political Economy*, 97(5):1091–1112, 1989.
- M. Ostrovsky. Information aggregation in dynamic markets with strategic traders. *Econometrica*, 80(6):2595–2648, 2012.
- M. Ottaviani and P.N. Sørensen. Parimutuel versus fixed-odds markets. *Mimeo*, 2005.
- M. Ottaviani and P.N. Sørensen. The favorite-longshot bias: An overview of the main explanations. *Handbook of Sports and Lottery markets*, pages 83–101, 2008.
- M. Ottaviani and P.N. Sørensen. Surprised by the parimutuel odds? *American Economic Review*, 99(5):2129–2134, 2009.
- M. Ottaviani and P.N. Sørensen. Noise, information, and the favorite-longshot bias in parimutuel predictions. *American Economic Journal: Microeconomics*, 2(1):58–85, 2010.
- S.E. Page. *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton University Press, Princeton, New Jersey, 2007.
- C.R. Plott and S. Sunder. Efficiency of experimental security markets with insider information: An application of rational-expectations models. *Journal of Political Economy*, 90(4):663–98, 1982.
- C.R. Plott and S. Sunder. Rational expectations and the aggregation of diverse information in laboratory security markets. *Econometrica*, 56(5):1085–1118, 1988.
- C.R. Plott, J. Wit, and W.C. Yang. Parimutuel betting markets as information aggregation devices: Experimental results. *Economic Theory*, 22(2):311–351, 2003.
- R.E. Quandt. Betting and equilibrium. *Quarterly Journal of Economics*, 101(1):201–207, 1986.
- D. Reeves and D. Pennock. *How and when to listen to the crowd (blog entry)*, 2007. URL http://www.overcomingbias.com/2007/02/how_and_when_to.html.
- L.J. Savage. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association*, 66(336):783–801, 1971.
- M.J. Schervish. A general method for comparing probability assessors. *The Annals of Statistics*, 17(4):1856–1879, 1989.
- E.H. Shuford, A. Albert, and H. Edward Massengill. Admissible probability measurement procedures. *Psychometrika*, 31(2):125–145, 1966.

- J. Surowiecki. *The Wisdom of Crowds*. Anchor Books, New York, 2004.
- W. Takahiro. A parimutuel system with two horses and a continuum of bettors. *Journal of Mathematical Economics*, 28(1):85–100, 1997.
- P.E. Tetlock. *Expert Political Judgment: How Good Is It? How Can We Know?* Princeton University Press, Princeton, New Jersey, 2006.
- R.H. Thaler and W.T. Ziemba. Anomalies: Parimutuel betting markets: Racetracks and lotteries. *Journal of Economic Perspectives*, 2(2):161–174, 1988.
- T. Watanabe, H. Nonoyama, and M. Mori. A model of a general parimutuel system: Characterizations and equilibrium selection. *International Journal of Game Theory*, 23(3):237–260, 1994.
- R.L. Winkler, J. Muñoz, J.L. Cervera, J.M. Bernardo, G. Blattenberger, J.B. Kadane, D.V. Lindley, A.H. Murphy, R.M. Oliver, and D. Ríos-Insua. Scoring rules and the evaluation of probabilities. *TEST*, 5(1):1–60, 1996.
- J. Wolfers and E. Zitzewitz. Prediction markets. *Journal of Economic Perspectives*, 18(2):107–126, 2004.

Appendix: Proof of Lemma 2

Proof of Lemma 2: We first prove by induction that for all $i \in \{1, \dots, n\}$, there exist twice continuously differentiable functions $g_i : (a, b)^{n-i+1} \mapsto \mathbb{R}$ and $h_i : (a, b)^{n-1} \mapsto \mathbb{R}$ such that

$$f(x_1, \dots, x_n) = g_i(x_1, x_{i+1}, \dots, x_n) + h_i(x_2, \dots, x_n), \quad (12)$$

and then apply this result to show that the lemma statement holds.

When $i = 1$, the claim in Equation (12) holds trivially with $g_1 = f$ and $h_1 = 0$. Now, suppose that for some $i \geq 1$, Equation (12) holds. Then since

$$\frac{\partial^2 f(x_1, \dots, x_n)}{\partial x_1 \partial x_{i+1}} = 0$$

for all values of x_1, \dots, x_n , and

$$\frac{\partial^2 h_i(x_{i+1}, \dots, x_n)}{\partial x_1 \partial x_{i+1}} = 0,$$

it must be the case that

$$\frac{\partial^2 g_i(x_1, x_{i+1}, x_{i+2}, \dots, x_n)}{\partial x_1 \partial x_{i+1}} = 0$$

for all $x_1, x_{i+1}, \dots, x_n \in (a, b)$.

Observe that, as g_i is twice continuously differentiable,

$$\frac{\partial^2 g_i(x_1, x_{i+1}, \dots, x_n)}{\partial x_1 \partial x_{i+1}} = \frac{\partial^2 g_i(x_1, x_{i+1}, \dots, x_n)}{\partial x_{i+1} \partial x_1}. \quad (13)$$

Hence, by twice integrating the right member of Equation (13), first with respect to x_1 then with respect to x_{i+1} , we find that there exist twice continuously differentiable functions $g_{i+1} : (a, b)^{n-i} \mapsto \mathbb{R}$ and $\tilde{h}_{i+1} : (a, b)^{n-i} \mapsto \mathbb{R}$ such that

$$\begin{aligned} g_i(x_1, x_{i+1}, \dots, x_n) \\ = g_{i+1}(x_1, x_{i+2}, \dots, x_n) + \tilde{h}_{i+1}(x_{i+1}, \dots, x_n). \end{aligned}$$

We can then satisfy Equation (12) at $i + 1$ by setting

$$h_{i+1}(x_2, \dots, x_n) = h_i(x_2, \dots, x_n) + \tilde{h}_{i+1}(x_{i+1}, \dots, x_n) .$$

Taking $i = n$, we find that

$$f(x_1, \dots, x_n) = g_n(x_1) + h_n(x_2, \dots, x_n) .$$

Let $f_1 = g_n$. Now, since h_n is twice continuously differentiable, we can repeat the same process for x_2 to find f_2 . Repeating the process recursively for x_3, \dots, x_{n-1} shows that there exist a set of functions f_3, \dots, f_n such that

$$f(x_1, \dots, x_n) = f_1(x_1) + \dots + f_n(x_n) .$$

□