

# STRATEGY-PROOF INCENTIVES FOR PREDICTIONS<sup>☆</sup>

Amir Ban

*Faculty of Mathematics and Computer Science, Weizmann Institute of Science, Rehovot, Israel*

---

## Abstract

Our aim is to design mechanisms that motivate all agents to reveal their predictions truthfully and promptly. For myopic agents, proper scoring rules induce truthfulness. However, as has been described in the literature, when agents take into account long-term effects of their actions, deception and reticence may appear. No simple rules exist to distinguish between the truthful and the untruthful situations, and a determination has been done in isolated cases only.

This is of relevance to prediction markets, where the market value is a common prediction, and more generally in informal public prediction forums, such as stock-market estimates by analysts. We describe three mechanisms that are strategy-proof with non-myopic considerations, and show that one of them meets all our requirements from a mechanism in almost all prediction settings. We formulate rules to distinguish truthful from untruthful settings, and use them to extensively classify prediction settings with continuous outcomes. We show how our proposed mechanism restores prompt truthfulness where incumbent mechanisms fail, and offer guidelines to implementing it in a prediction market.

---

## 1. Introduction

Mechanisms that motivate all agents to reveal their information truthfully and promptly are desirable in many situations.

Consider, for example, the estimation of company earnings by stock-market analysts, a long-standing Wall Street institution. Publicly-traded companies announce their earnings for the latest quarter or year, on dates set well in advance. Each company is typically covered by several stock-market analysts, the larger ones by dozens. These analysts issue reports containing predictions of a company's future earnings. The timing of such predictions may range from several years to days before earnings announcement, and every analyst typically updates his prediction several times in the interval. These predictions eventually become publicly available, and a consensus calculated from all predictions in force may be viewed on several popular finance websites. Not least, the analysts themselves are aware of, and are no doubt influenced by the actions and opinions of their peers.

In essence, this earnings estimation functions as a public prediction forum with an evolving consensus, that terminates when a company announces its true earnings for the forecast period, which we call the *outcome*. It acts as a sort of advisory forum for the public of investors, and this

---

<sup>☆</sup>A short version of this paper was presented in WINE'18 and published in proceedings  
Email address: [amirban@netvision.net.il](mailto:amirban@netvision.net.il) (Amir Ban)

public's interest is best served if analysts share their information and judgement truthfully and promptly.

Prediction markets are public prediction forums organized as markets. A de-facto standard for organizing prediction markets is due to Hanson (2003), using a *market scoring rule*. In such a market, probability estimates are rewarded by a *proper scoring rule* an amount  $S(\mathbf{p}, r)$ , where  $\mathbf{p}$  is a probability distribution of the outcome, and  $r$  is the outcome. A trader in Hanson's markets not only makes her probability estimate public, she changes the market price to it. She then stands to be rewarded by the *market maker* for her prediction (when the outcome becomes known), but she also commits to compensate the *previous* trader for his prediction. Her total compensation is therefore the difference  $S(\mathbf{p}, r) - S(\mathbf{p}', r)$  where  $\mathbf{p}'$  is the replaced market prediction. When the logarithmic scoring rule ( $S(\mathbf{p}, r) = \log p_r$ )<sup>1</sup> is used in a prediction market, the mechanism is called LMSR (Logarithmic Market Scoring Rule). Hanson also demonstrated how a market maker can facilitate, at a bounded cost to himself, such an LMSR market and provide liquidity by selling and buying shares of each outcome.

Proper scoring rules, are, by their definition, incentive compatible for *myopic* agents<sup>2</sup>. That is, an agent maximizes her expected score by announcing her true belief, provided longer-term effects of the prediction, if any, are ignored. The incremental, market variation of the scoring rule does not affect this incentive compatibility, because the previous agent's score does not depend on the current prediction. Furthermore, it is a straightforward generalization of Hanson (2003) to apply market scoring rules to multiple-choice, or continuous outcomes (such as in our earnings estimate forum).

When an agent is *not* myopic, and *does* take into account all consequences of her action, truthfulness will, in many cases, *not* be her optimal strategy, and incentive compatibility is lost. Such scenarios have been described in the literature, and our paper adds many further examples. As we will show, the damage to incentive compatibility caused by long-term strategic considerations is extensive, and it is *a priori* unclear whether any remedy is available.

Our aim is the design of mechanisms for rewarding predictions that are strategy-proof. We demonstrate the problem with proper scoring rules, as often leading to reticence or deception. We formulate criteria for determining which prediction settings are truthful, and apply these criteria for a complete classification of the important class of prediction settings with normal and lognormal signals. We suggest three strategy-proof mechanisms, and identify one of them, *discounting*, as having all desirable properties. We prove the applicability and effectiveness of the discounting mechanism, and offer guidelines to implementing it.

### 1.1. The Problem with Scoring Rules

A scoring rule  $S : \Delta(\mathbb{R}) \times \mathbb{R} \mapsto \mathbb{R}$  scores a prediction  $\mathbf{p}$ , representing a probability distribution of the outcome, a value  $S(\mathbf{p}, r)$  when the outcome is  $r$ . An agent whose belief of the outcome distribution is  $\mathbf{q}$  has score expectation  $S(\mathbf{p}, \mathbf{q}) := \mathbb{E}_{r \sim \mathbf{q}} S(\mathbf{p}, r)$  for prediction  $\mathbf{p}$ . A *proper* scoring rule is one for which  $S(\mathbf{q}, \mathbf{q}) \geq S(\mathbf{p}, \mathbf{q})$  for every  $\mathbf{p}, \mathbf{q} \in \Delta(\mathbb{R})$ , so that predicting one's true belief has maximal score expectation. A *strictly proper* scoring rule is one where the inequality is tight only for  $\mathbf{p} = \mathbf{q}$ . The *logarithmic* scoring rule  $S(\mathbf{p}, r) = \log p_r$ , and the *quadratic* (a.k.a. *Brier*)

---

<sup>1</sup>We use the notation  $p_x$  for the density of distribution  $\mathbf{p}$  at  $x$ .

<sup>2</sup>The incentive compatibility is also restricted to *risk-neutral* agents. Offerman et al. (2009) demonstrate how proper scoring rules may be corrected for other risk attitudes.

scoring rule  $S(\mathbf{p}, r) = 2p_r - \mathbf{p} \cdot \mathbf{p} - 1$  are examples of strictly proper scoring rules. More background on scoring rules may be found, e.g., in Gneiting and Raftery (2007).

The following generic example illustrates the problem when non-myopic considerations apply.

**Example 1.** *The public wants to predict a variable, whose outcome is  $x$ . A market scoring rule compensates predictions. Every signal of  $x$  is, i.i.d.,  $x + \epsilon$  with probability  $1/2$ , and  $x - \epsilon$  with probability  $1/2$ . The distributions of  $x$  and  $\epsilon$  are such that a single signal does not reveal  $x$ . There is an expert, who gets private signals. The public gets public signals.*

- *On Sunday, expert gets a signal.*
- *On Monday, public gets a signal.*
- *On Tuesday, expert gets another signal.*
- *On Wednesday, outcome  $x$  is revealed.*

*Question: Should expert reveal his information truthfully on Sunday?*

*Answer: NO. Whoever sees two different truthful signals is able to calculate the outcome  $x = (x + \epsilon)/2 + (x - \epsilon)/2$  exactly. For any distribution of  $\epsilon$ , and for the logarithmic and almost<sup>3</sup> every other scoring rule, the expert should not tell the truth on Sunday. This prevents the 50% probability that the market will know  $x$  on Monday, preserving a 75% probability that the expert can announce  $x$  on Tuesday.*

The canonical case, to which this example belongs, is “Alice-Bob-Alice”, where Alice speaks before and after Bob’s single speaking opportunity, both are awarded by a proper scoring rule for each prediction, and both maximize their total score. Chen et al. (2010) as well as Chen and Waggoner (2016) studied situations where several agents, each having private information, are given more than one opportunity to make a public prediction. The situations are reducible to the Alice-Bob-Alice game. The proper scoring rule assures that each will tell the truth on their last prediction, and the open question is whether Alice, when going first, will tell the truth, lie, or keep her silence. Chen et al. (2010) make the key observation that truthfulness is optimal if, in a different setup, namely, a single-prediction Alice-Bob game where Alice chooses whether to go first or second, she will always prefer going first. Building on that insight, Chen and Waggoner (2016) show that when the players’ information is what they define as “perfect informational substitutes”, they will predict truthfully and as early as allowed, when they are “perfect informational complements”, they will predict truthfully and as *late* as allowed, while when players are neither substitutes nor complements, untruthfulness can and will occur. While this characterization is helpful, few concrete cases have been settled. The most significant of those was to show that when signals are independent conditional on the outcome, and the logarithmic scoring rule (LMSR) is used, the signals are informational substitutes, meaning that in such a case, Alice will reveal all her information truthfully in her first round.

A strategy-proof mechanism of the Alice-Bob-Alice setting easily generalizes to a strategy-proof mechanism for any number of experts and prediction order, because whenever an expert (call her Alice) makes more than one prediction, one can roll together all experts making predictions between

---

<sup>3</sup>The example is true for the logarithmic scoring rule because its scores are unbounded. Every scoring rule that values exact predictions over inexact ones sufficiently will do.

Alice’s successive predictions into one expert (call him Bob), who shares their information<sup>4</sup>. This is formally proved in Theorem 15.

### 1.2. Goals of the Mechanism

We seek a mechanism with several desirable traits.

1. **TRUTHFULNESS:** The mechanism should motivate all experts to make truthful predictions, that is, to reveal their true subjective distributions of the outcome. At minimum, this means that truth-telling should be a best-response to truth-telling by all other experts, according to the player’s beliefs at the time of prediction, i.e., it is a perfect Bayesian equilibrium.<sup>5</sup>

An untruthful mechanism may still be *locally truthful* by which we mean that infinitesimal variations from the truth are suboptimal, but telling a sufficiently big lie may be advantageous.

2. **FULL DISCLOSURE:** Every expert makes a (truthful) prediction some time after getting his last signal. This is a necessary<sup>6</sup> condition for all the information on the outcome possessed in the expert’s signals to reach the public.
3. **PROMPTNESS:** Experts should reveal their signals by a truthful public prediction as soon as the prediction schedule allows, and make an updated prediction whenever receiving a new signal, again, at the earliest opportunity. Considerations by experts of when to make a prediction are contrary to the interest of the public. We shall require a *strict* preference for promptness. Indifference to timing shall not count as prompt.
4. **BOUNDED LOSS:** Hanson (2003) notes that in his market scoring rule mechanism, the market maker effectively subsidizes traders to motivate their truthfulness, and he shows that market maker’s expected loss due to that is *bounded*: Since the market maker initializes the prediction (to some prior  $\pi$ ), and needs to compensate only the last prediction (as every trader compensates the previous one), his cost (=loss) expectation is, for non-positive scoring rules (such as LMSR)

$$\mathbb{E}[Loss] = S(\mathbf{p}_{final}, \pi) - S(\pi, \pi) \leq -S(\pi, \pi) \quad (1)$$

We seek mechanisms that have bounded loss. Note that scoring rules can be arbitrarily scaled without effect on their motivation, so reducing the *size* of a bounded loss is not a goal.

### 1.3. Our Results

We propose three different incentive mechanisms, all of which are based on proper scoring rules, all of which achieve truthfulness, and the third and last also achieves promptness. They are

1. **GROUP PREDICTION:** All agents receive the final prediction’s (non-incremental) score. Since all agents have a stake in the final prediction, all will reveal their information truthfully. On the negative side, they are not motivated to be prompt about it. Another problem is

---

<sup>4</sup>Chen et al. (2010) use the same construction to generalize from Alice-Bob-Alice to a finite-players game.

<sup>5</sup>Note that the ideal of truth-telling as dominant strategy is not attainable here, because if a player is aware of another player’s distortion, the correct Bayesian response is to compensate for the distortion.

<sup>6</sup>But possibly not sufficient, e.g. when experts predict the XOR of their bit signals.

freeloaders, since agents with no information can participate and gain without contributing anything.

2. **ENFORCE SINGLE PREDICTION:** Score each of agent’s prediction with an incremental scoring rule, and award each agent the minimum score. Agents are therefore motivated to predict once only, since having made a prediction, a further prediction can only lower their reward expectation. With a proper scoring rule, this assures incentive compatibility with truthfulness. Agents are not motivated to be prompt, but instead need to find the optimal timing to make their single prediction. A major drawback is that when agents receive a time-varying signal, they will not reveal all their information.
3. **DISCOUNTING:** Discount each of agent’s incremental prediction scores by a monotonically increasing factor of the time. The idea is that if signals are not informational substitutes, they will become ones if a sufficiently steep negative time gradient is applied. When successful, this mechanism achieves the ideal result of motivating all agents to reveal their information truthfully and promptly, including when they receive time-varying signals.

We show that, under some light conditions, discounting will always work unless signals are perfectly correlated, i.e., have a deterministic relation given the outcome (as in Example 1).

Table 1 summarizes how the incumbent mechanism, and our three proposed mechanisms, measure up against each of the traits we described as desirable in the previous section.

We provide (Corollary 2) a criterion for the truthfulness of the Alice-Bob-Alice game, as well as sufficient conditions for truthfulness and for untruthfulness (Corollary 3).

We put these to use to thoroughly investigate the Alice-Bob-Alice game (and, by extension, multi-player, multi-signal games) with both the logarithmic and the quadratic (Brier) scoring rule, when player signals have a multivariate normal distribution or a multivariate lognormal distribution.

These distributions are among the best-known continuous distributions and naturally arise in many situations. They are characterized, *inter alia*, by the correlation coefficient ( $\rho \in [-1, 1]$ ) between Alice’s and Bob’s signals. When the scoring rule used is the logarithmic one, we find that when these signals are too well-correlated (whether positively or negatively), prompt truthfulness is *not* optimal. On the other hand, if the correlation is low, the game will be truthful and prompt. This includes the case  $\rho = 0$ , where, as is well-known, the signals are conditionally independent, confirming the Chen et al. (2010) result for conditionally independent signals.

However, when the quadratic scoring rule, one of the oldest and most commonly used scoring rules, is used with these multivariate distributions, it is *never* truthful for repeated predictions.

In all settings with either the logarithmic or the quadratic scoring rules, we show that our discounting mechanism restores prompt truthfulness, with the single exception of perfect correlation ( $|\rho| = 1$ ) of the players’ signals.

We show that the discounting mechanism can be effectively implemented with an automated market maker, as in Hanson’s markets, thus showing that it may be practically applied in prediction markets. In the Discussion, we offer guidelines on how to design such a mechanism.

#### 1.4. Related Literature

Scoring rules have a very long history, going back to De Finetti (1937), Brier (1950) and Good (1952). Proper scoring rules are often used for incentive-compatible belief elicitation of risk-neutral

agents (e.g. Armantier and Treich (2013)). Market scoring rules for prediction markets were introduced by Hanson (2003).

The role of Chen et al. (2010) (which is based on earlier papers by Chen et al. (2007) and Dimitrov and Sami (2008)) and Chen and Waggoner (2016) in investigating the strategy-proofness of prediction markets was already described. Gao et al. (2013) and Kong and Schoenebeck (2018) resolve some more scenarios. Conitzer (2009) embarks on a program similar to ours, citing mechanism design as a guiding principle. Accordingly, he strives to achieve the Revelation Principle, where all experts announce their private information to some organizing entity that makes the appropriate Bayesian aggregation. As we discuss in Section 2.2 below, we do not share that vision: Experts often do not know what part of their belief stems from truly private information, and even when they do, they cannot afford to go on record with a prediction which is not their best judgement. His “Group-Rewarding Information Mechanism” is similar to our Group Prediction mechanism, and its lack of fairness is pointed out. Conitzer does not propose a mechanism that achieves prompt truthfulness.

Chen et al. (2010) also suggest discounting, that “reduces the opportunity for bluffing”, in their words, but does not prevent it (Section 9.1), so their discounting mechanism does not achieve our basic requirement of truthfulness. The reason is that their formulation is different from ours, applying same discount to *before* and *after* scores. On the other hand, we discount every prediction score according to the time its prediction was made. The difference is crucial, because theirs does not result in a true market scoring rule, as defined by Hanson (2003). In consequence, our Section 2.4, on which our results rest, as well as our Section 5.2, do not apply to their formulation.

We shall occasionally rely on well-known facts of the normal and the multivariate normal distributions. The reader will find the basis for these in, e.g., Tong (2012).

The rest of this paper is organized as follows: In Section 2 we formulate the problem. In Section 3 we investigate which predictions settings are already truthful and prompt. In Section 4 we offer strategy-proof mechanisms, and show how they can solve the gaps we have found. In Section 5 we summarize and offer concluding remarks. Most proofs are to be found in the Appendix.

## 2. Problem Formulation

### 2.1. Basics

Two players, Alice and Bob, make public predictions of a real variable  $\lambda$ , whose prior distribution is  $\pi$ . The *outcome*  $x$  will be revealed after all predictions have been made. A prediction consists of revealing one’s belief of the distribution of  $\lambda$ . All agents (Alice, Bob, and the public) are Bayesian, and each prediction causes them to update their beliefs, i.e. the posterior distribution, of  $\lambda$ .<sup>7</sup>

In a truthful equilibrium, truthfulness is a best response of a player to truthfulness by all others. So we start by assuming that all agents take others’ predictions as truthful, and investigate whether the best response is to be truthful.

The posterior beliefs are distributions of the variable  $\lambda$  which are inferred from priors and likelihood functions using Bayesian techniques. In our discussion we find it more convenient and

---

<sup>7</sup>This formulation is different from the mechanism of prediction markets, but equivalent to it. In prediction markets, an agent replaces the current market prediction by his own. In our formulation, the agent merely announces a prediction, which, assuming the agent is truthful, becomes the market prediction by Bayesian inference. This is because all rational agents reach the same beliefs from the same data.

succinct to represent beliefs, without loss of generality, by a real number, rather than by a probability distribution. We use the fact that, in Bayesian analysis, when the likelihood functions belong to some family of distributions (e.g. exponential), all posterior beliefs belong to another family of distributions (Gamma distribution for the exponential family)  $Q(\mathbf{Y}) \in \Delta(\lambda)$ , called the *conjugate prior* of the first family.  $\mathbf{Y}$  is a set of real *parameters* of the inferred distribution  $Q$ . We will assume models where one, and only one of these parameters is dependent on previous predictions, while the rest  $\mathbf{Y} \setminus \{x\}$  is known from the model, the timing and the identity of the believer, but *does not depend on any previous prediction*. An example illustrates this:

**Example 2.** Assume Alice's belief of  $\lambda$  to be normally distributed  $N(\mu_A, 1/\tau_A)$  where  $\mu_A$  is the mean and  $\tau_A$  the accuracy (i.e. inverse of the variance), and Bob's is  $N(\mu_B, 1/\tau_B)$  and independent of Alice's.  $\tau_A$  and  $\tau_B$  are set by the model and are commonly known. Assume an uninformative prior. Using well-known aggregation rules for independent normal observations, if Alice announces  $\mu_A$ , Bob's belief changes to the normally distributed  $N(\mu_{AB}, 1/\tau_{AB})$ , where

$$\begin{aligned}\mu_{AB} &= \frac{\tau_A \mu_A + \tau_B \mu_B}{\tau_A + \tau_B} \\ \tau_{AB} &= \tau_A + \tau_B\end{aligned}$$

Notice that  $\tau_{AB}$  can be calculated without knowing any of the means  $\mu_A, \mu_B$ , while  $\mu_{AB}$  can be evaluated once  $\mu_A$  and  $\mu_B$  is known.

In this context we are therefore able to describe a prediction by a single real number (the mean) rather than by a probability distribution. We shall say that Alice's prior belief  $A_1$  is  $\mu_A$  and Bob's prior belief  $B_1$  is  $\mu_B$ . After Alice makes her prediction, Bob's belief changes to  $\mu_{AB}$ . After Alice and Bob both make a prediction, the public's belief is  $\mu_{AB}$ . In context, these statements are unambiguously equivalent to specifying the probability distributions in full.

The prior  $\pi$  may be uninformative, assigning equal probabilities to all possibilities<sup>8</sup>, or, if not, as also representable by a parameter. For example, if Alice and Bob participate in a prediction market, the prior parameter is the market value before Alice's first prediction.

Time is discrete,  $t = 0, 1, 2, \dots, T$ , with  $T$  the time the outcome is known.  $A_t, B_t, C_t \in \mathbb{R}$  are, respectively, Alice's, Bob's and the public's (or market's) beliefs at time  $t$ . At  $t = 1$ ,  $A_1, B_1, C_1$  are their respective prior beliefs. Any prediction takes place at  $t > 1$ .  $t = 0$  is "pre-prior" time, when players beliefs are equal to their *private* signals, so that  $A_0, B_0$  are respectively, Alice and Bob's private signals. At  $t = 1$ , each player is additionally aware of the public prior  $C_1$ , so that  $A_1$  is an inference from  $A_0$  and  $C_1$  while  $B_1$  is an inference from  $B_0$  and  $C_1$ . If the public prior is uninformative, then we have  $A_0 = A_1$  and  $B_0 = B_1$ . In other words, the players' priors equal their private signals. For completeness, we define  $C_0 = C_1$ .

To avoid degenerate exceptions, we assume the players' signals are *informative*. This means that  $A_1 \neq C_1$  and  $B_1 \neq C_1$ .

The signals have a common-knowledge joint distribution  $f(a, b; \lambda)$  conditional on the variable

$$f(a, b; \lambda) := \Pr(A_1 = a, B_1 = b | \lambda)$$

The order of predictions is Alice, Bob, then Alice again, and then the outcome  $x$  is revealed.

---

<sup>8</sup>Technically, an uninformative prior may be envisioned as the limit of a uniform or normal distribution as the variance goes to infinity.

A twice-differentiable, w.r.t.  $\lambda$ , proper scoring rule  $S(\mathbf{p}, \lambda)$  incrementally rewards each prediction made, i.e., if a player's prediction changed the public's belief from  $\mathbf{p}'$  to  $\mathbf{p}$ , the player's reward for this prediction, calculated when  $x$  is known, is  $S(\mathbf{p}, x) - S(\mathbf{p}', x)$ . Each player seeks to maximize their total reward.

As the scoring rule is proper, Bob will tell the truth on his only prediction, and Alice will tell the truth on her second and last prediction. The remaining question is whether Alice will tell the truth on her first prediction. More accurately, the question is of equilibrium: If Bob is truthful, and Bob and the public take Alice's predictions as truthful, is truth-telling Alice's best response? If the answer is affirmative when all players use all available information to update their beliefs, we have a truthful Perfect Bayesian Equilibrium.

## 2.2. Knowledge Model

As will be shown, the players behavior is affected by their *common-knowledge prior*, by which we mean a belief distribution which is explicitly known to both players and from which each inferred his or her current belief. That the common prior is commonly known is significant, because it is quite possible, and even likely, that the players share a prior but do not know what it is. For example, in predicting a poll, Alice and Bob may be basing themselves on knowledge of how their acquaintances voted, but they may not know which acquaintances they have in common. Or, they both may be basing themselves on a paper they read, but neither is aware that the other has read it. If the players do not know what their common prior is, they cannot infer anything from having one.

In prediction markets, and more generally in public prediction forums where all communication is done in public, the common-knowledge prior is known. In the context of this paper, it is the initial market value  $C_1$  (or, equivalently, distribution  $\pi$ ), when the Alice-Bob-Alice game starts.

## 2.3. Inferences from Predictions

Assume that inference functions are invertible, so that if a player's prediction is known, her signal can be computed. If Alice announces  $A_1 = a$ , the posterior outcome distribution can be calculated from her marginal distribution,  $f_A(a; \lambda) := \Pr(A_1 = a | \lambda) = \int_{-\infty}^{\infty} f(a, b'; \lambda) db'$ . Mark it  $g(a)$ .

$$g(a)_\lambda = \Pr(\lambda | A_1 = a) = \frac{\pi_\lambda f_A(a; \lambda)}{\int_{-\infty}^{\infty} \pi_{\lambda'} f_A(a; \lambda') d\lambda'} \quad (2)$$

Similarly, if Alice announces  $A_1 = a$ , and Bob privately observes  $B_1 = b$ , Bob's posterior outcome is inferred from  $f(a, b; \lambda)$ . It will become the public prediction when Bob announces it. Mark it  $h(a, b)$ .

$$h(a, b)_\lambda = \Pr(\lambda | A_1 = a, B_1 = b) = \frac{\pi_\lambda f(a, b; \lambda)}{\int_{-\infty}^{\infty} \pi_{\lambda'} f(a, b; \lambda') d\lambda'} \quad (3)$$

## 2.4. Maximizing the Reward

How does Alice maximize her total reward for both her predictions? And is this maximum achieved by telling the truth on both predictions? We will show that Alice maximizes her reward by *minimizing* Bob's reward, and therefore is truthful if truth minimizes Bob's reward.

**Theorem 1.** *Alice maximizes her expected total reward by making a first prediction that minimizes Bob's expected reward, where expectations are taken according to Alice's beliefs on her first prediction, and assuming Bob is truthful and takes Alice's prediction as truthful.*



PROOF. See Appendix.

**Corollary 2.** *If a truthful first prediction minimizes Alice’s expectation of Bob’s reward, i.e., if*

$$\mathbb{E}_{b \sim B_1 | (A_1=a)} \mathbb{E}_{\lambda \sim h(a,b)} \Pi_B(\lambda; a, \hat{a}, b) = \left\{ \mathbb{E}_{b \sim B_1 | (A_1=a)} S(h(\hat{a}, b), h(a, b)) \right\} - S(g(\hat{a}), g(a))$$

*is minimized at  $\hat{a} = a$ , then truth is Alice’s best policy.*

A corollary that we will find useful is the following.

**Corollary 3.** *Let  $c := \hat{a} - a$  be Alice’s deviation from truth, and define*

$$\Delta(c; a, b) := \left[ S(h(\hat{a}, b), h(a, b)) - S(h(a, b), h(a, b)) \right] - \left[ S(g(\hat{a}), g(a)) - S(g(a), g(a)) \right]$$

*If, for every  $c \neq 0$ , and every  $b$ ,  $\Delta(c; a, b) > 0$ , Alice will be promptly truthful.*

*Alternatively, if, for every  $c \neq 0$ , and every  $b$ ,  $\Delta(c; a, b) < 0$ , Alice will not be promptly truthful.*

PROOF. If Alice assumes  $B_1 = b$  when her belief is  $g(a)$ , then Alice’s assumed belief is  $h(a, b)$ , and so her expectation of  $\Pi_B(x; a, \hat{a}, b)$  is (see (A.2) in Appendix),  $S(h(\hat{a}, b), h(a, b)) - S(g(\hat{a}), g(a))$ . So the corollary assumes that truth ( $\hat{a} = a$ ) minimizes Alice’s expectation of Bob’s reward for every  $b$ , and therefore also Alice’s total expectation of Bob’s reward. The corollary follows by Corollary 2.  $\square$

### 3. Which Prediction Settings are Already Truthful?

We described above (Example 1) an elementary setting that is not truthful, and a handful of other settings have been settled either way in the literature. But in the landscape of prediction settings that are of interest, the coverage has been very sparse. Beyond Chen and Waggoner (2016)’s criterion of “Informational Substitutes”, which does not amount to an explicit algorithm<sup>9</sup>, and Kong and Schoenebeck (2018), who offer a relatively-efficient algorithm to settle some scenarios, we have no procedure to settle any given case, and the problem remains opaque.

With the results we derived in Section 2.4, we now have such procedures. We shall apply them to classify settings belonging to the most commonly-met continuous distributions and the most commonly used scoring rules. The distributions are:

- The multivariate normal distribution  $N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ . In contrast to general joint distributions, where the interdependence of components may take practically any form, the interdependence of joint multivariate normal components is completely determined by their covariance matrix. This means that in the Alice-Bob-Alice game, the signals interdependence is completely determined by a single real parameter: their correlation coefficient  $-1 \leq \rho \leq 1$ . Our investigation boils down to finding for what values of  $\rho$  Alice will be promptly truthful.
- The multivariate lognormal distribution  $\mathbf{Y} = \exp(\mathbf{X})$ , where  $\mathbf{X}$  is a multivariate normal distribution as described in the previous item. Stock prices, and their derivatives, are commonly modelled by a lognormal distribution. Since taking logs transform multivariate lognormal random variables to multivariate normal random variables, our classification of the multivariate normal distribution readily provides a classification of this distribution too.

---

<sup>9</sup>A submodularity property is required of the signal lattice, in a context described in their article.

The scoring rules we cover are

- The Logarithmic Scoring Rule  $S(\mathbf{p}, r) = \log p_r$ . As mentioned, this proper scoring rule is employed as Hanson's LMSR market-maker. It has strong information theory roots, and is uniquely *local*, i.e. depending only on the distribution at the outcome ( $r$ ).
- The Brier/Quadratic Scoring Rule  $S(\mathbf{p}, r) = 2p_r - \mathbf{p} \cdot \mathbf{p} - 1$ . It is the earliest scoring rule, introduced by Brier (1950), in the context of weather forecasting.

### 3.1. Multivariate Normal Distribution

Let the public prior  $\boldsymbol{\pi}$  be distributed  $N(C_0, 1/\tau_C)$ , where  $\tau_C = 1/\sigma_C^2$  is the accuracy.  $C_0$  is the public's prediction mean at the start of the game. Note that an uninformative prior is characterized by having an arbitrarily large variance  $\sigma_C^2$ , or equivalently  $\tau_C = 0$ , and in such a case, the value of  $C_0$  is inconsequential.

Similarly, Alice's and Bob's priors are distributed  $N(A_1, 1/\tau_{AC})$  and  $N(B_1, 1/\tau_{BC})$ , respectively, where  $\tau_{AC} = 1/\sigma_{AC}^2$  and  $\tau_{BC} = 1/\sigma_{BC}^2$  are the respective accuracies.  $A_1, B_1$  are the means of, respectively, Alice's and Bob's prior predictions. Alice's prediction  $A_1$  is inferred from the public prior  $C_0$  and Alice's signal  $A_0$ , and similarly for Bob.

We lose no generality in assuming that each of the player's signals is conditionally independent of the prior. Because, if Alice, e.g., is aware of the public prior  $C_0$  and her prior prediction  $A_1$ , she can consistently assume that her signal  $\boldsymbol{\pi}_A$  is distributed  $N(A_0, 1/\tau_A)$ , where

$$A_0 = \frac{\tau_{AC}A_1 - \tau_C C_0}{\tau_{AC} - \tau_C} \quad \tau_A = \tau_{AC} - \tau_C \quad (4)$$

It is easy to verify that  $\boldsymbol{\pi}_A$  and  $\boldsymbol{\pi}$  are mutually independent. Note that  $\tau_{AC} > \tau_C$  since we assume Alice's signal is informative, hence  $\tau_A > 0$ .

For an uninformative prior  $\tau_C = 0$  this degenerates to  $A_0 = A_1, B_0 = B_1$ , as can be expected.

By assumption the private signals  $\boldsymbol{\pi}_A, \boldsymbol{\pi}_B$  and the public prior  $\boldsymbol{\pi}$  belong to a jointly multivariate normal distribution, and as such, their joint distribution is completely determined by their means and by their covariance matrix. Since  $\boldsymbol{\pi}_A, \boldsymbol{\pi}_B$  are, w.l.o.g., conditionally independent of  $\boldsymbol{\pi}$ , the distributions are uncorrelated with  $\boldsymbol{\pi}$ , and the covariance matrix  $M$  for  $(\boldsymbol{\pi}_A, \boldsymbol{\pi}_B, \boldsymbol{\pi})$  is

$$M = \begin{pmatrix} \sigma_A^2 & \rho\sigma_A\sigma_B & 0 \\ \rho\sigma_A\sigma_B & \sigma_B^2 & 0 \\ 0 & 0 & \sigma_C^2 \end{pmatrix} \quad (5)$$

where  $\rho \in [-1, 1]$  is the correlation coefficient of  $\boldsymbol{\pi}_A$  and  $\boldsymbol{\pi}_B$ , defined as:

$$\rho = \frac{\text{Cov}(\boldsymbol{\pi}_A, \boldsymbol{\pi}_B)}{\sqrt{\text{Var}(\boldsymbol{\pi}_A)\text{Var}(\boldsymbol{\pi}_B)}}$$

As is well-known (e.g. Tong (2012)), components of a multivariate normal distribution are independent iff they are uncorrelated, i.e. have  $\rho = 0$ .

**Proposition 4.** Assume that the model, including covariance matrix  $M$  (5), is common knowledge. Then, knowing everyone's prior prediction  $A_1, B_1$  and  $C_0$ , the posterior, marked  $\boldsymbol{\pi}_{ABC}$ , has distribution  $N(\mu_{ABC}, 1/\tau_{ABC})$ , with

$$\mu_{ABC} = \frac{(\tau_A - \rho\sqrt{\tau_A\tau_B})A_0 + (\tau_B - \rho\sqrt{\tau_A\tau_B})B_0 + (1 - \rho^2)\tau_C C_0}{\tau_A - 2\rho\sqrt{\tau_A\tau_B} + \tau_B + (1 - \rho^2)\tau_C} \quad (6)$$

$$\tau_{ABC} = \frac{\tau_A - 2\rho\sqrt{\tau_A\tau_B} + \tau_B}{1 - \rho^2} + \tau_C \quad (7)$$

PROOF. See Appendix.

### 3.1.1. With Logarithmic Scoring

**Proposition 5.** *Let Alice's, Bob's and the public signals in an Alice-Bob-Alice game belong to a jointly multivariate normal distribution, with covariance matrix given in (5). When using the logarithmic scoring rule, the players' best-response strategy to truthfulness in others is truthfulness and promptness iff*

$$\left(1 - \rho^2\right)^2 \left(1 + \frac{\tau_C}{\tau_B}\right) \geq \left(\rho^2 + \frac{\tau_C}{\tau_A}\right) \left(\sqrt{\frac{\tau_A}{\tau_B}} - \rho\right)^2 \quad (8)$$

PROOF. See Appendix.

In particular, we note in Proposition 5 that the case  $|\rho| = 1$  is untruthful, and more generally the same is true when the correlation coefficient is either too positive or too negative. The middle ground is truthful. In particular, substituting  $\rho = 0$ , where the players' signals are conditionally independent, is found to be truthful. This generalizes the result of Chen et al. (2010) (Theorem 5) for conditionally independent signals with the logarithmic scoring rule, because it shows that conditionally independent signals induce truthfulness even when modified by a common-knowledge prior.

Also in particular, if the common prior is uninformative, or unknown, we get the following corollary by substituting  $\tau_C = 0$  in Proposition 5 and solving for  $\rho$ .

**Corollary 6.** *Let Alice's and Bob's signals in an Alice-Bob-Alice game belong to a jointly multivariate normal distribution, with covariance matrix given in (5), and an uninformative or unknown common prior ( $\tau_C = 0$ ). Then, with the logarithmic scoring rule, the players' best-response strategy to truthfulness in others is truthfulness and promptness iff*

$$\left\{ \begin{array}{l} \rho \leq \frac{\sigma_A}{\sigma_B} \\ \frac{1}{4} \left( \frac{\sigma_B}{\sigma_A} - \sqrt{\frac{\sigma_B^2}{\sigma_A^2} + 8} \right) \leq \rho \leq \frac{1}{4} \left( \frac{\sigma_B}{\sigma_A} + \sqrt{\frac{\sigma_B^2}{\sigma_A^2} + 8} \right) \end{array} \right. \quad \text{and}$$

In particular, if  $\sigma_A = \sigma_B$ , truthfulness and promptness holds whenever  $\rho \geq -\frac{1}{2}$ .

### 3.1.2. With Quadratic Scoring

On the other hand, when using the Quadratic Scoring Rule, when the signals belong to a multivariate normal distribution, truthfulness is *never* optimal. This is summarized by the following proposition.

**Proposition 7.** *Let Alice's, Bob's and the public signals in an Alice-Bob-Alice game belong to a jointly multivariate normal distribution, with covariance matrix given in (5). Then, when using the quadratic scoring rule, truthfulness is never Alice's best-response strategy to truthfulness in others. Alice is, however, locally truthful for  $0 < \rho < \frac{\sigma_A}{\sigma_B}$ .*

PROOF. See Appendix.

Remarkably, the result is unaffected by the prior ( $\sigma_C$ ).

Table 1: Mechanism Score Card

| Mechanism                 | Truthful | Full Disclosure | Prompt | Bounded Loss |
|---------------------------|----------|-----------------|--------|--------------|
| Market Scoring Rule       | ×        | ✓               | ×      | ✓            |
| Group Prediction          | ✓        | ✓               | ×      | ×            |
| Enforce Single Prediction | ✓        | ×               | ×      | ✓            |
| Discounting               | ✓        | ✓               | ✓      | ✓            |

## 4. Strategy-Proof Mechanisms

### 4.1. Group Prediction

Our first strategy-proof mechanism scores the last prediction made by a proper scoring rule, and awards the score to *each* of the participating experts.

**Proposition 8.** *Let  $x$  be the outcome, and let the last prediction made before the outcome is revealed be distribution  $\mathbf{p}$ . Then the mechanism that awards  $S(\mathbf{p}, x)$  to each participating player, where  $S$  is a proper scoring rule, is truthful. Furthermore, the mechanism elicits full disclosure.*

PROOF. See Appendix.

The mechanism does not motivate promptness, which we defined as a strict preference: Players may predict as late as possible without harming their welfare.

Another drawback is unfairness: The mechanism awards all experts the same, regardless of their contribution. Indeed, a so-called expert who has no information of his own may reap the same reward as other experts by simply repeating the current public prediction. This means that, unless the number of experts is bounded, the mechanism is not loss-bounded. Attempts to fix this would be counterproductive, compromising truthfulness. For example, rewarding nothing to “predictions” that merely repeat the public prediction, will motivate an uninformed expert to pretend knowledge by “tweaking” the current public prediction.

### 4.2. Enforce Single Prediction

Since multiple predictions are the source of the potential for manipulation, a mechanism that prevents that would restore general truthfulness.

**Proposition 9.** *Let  $S$  be a proper scoring rule. The mechanism that scores each prediction with the increment  $S(\mathbf{p}, x) - S(\mathbf{p}', x)$ , where  $\mathbf{p}$  is the predicted distribution and  $\mathbf{p}'$  the previous public prediction, and rewards each expert with the minimum score out of all her predictions, is truthful.*

PROOF. Once an expert has made a prediction, a further prediction may only lower her reward expectation and so is not optimal. Since incremental, proper scoring rules are truthful with a single prediction, the mechanism is truthful.  $\square$

While this mechanism is truthful, and loss-bounded (a consequence of Hanson’s loss-boundedness. See Section 1.2 (1)), it does not motivate promptness. Every expert needs to figure out the best time to make his single prediction, given other players’ strategies, resulting in a Bayes-Nash equilibrium. As Azar et al. (2016) show, this can be complex. Furthermore, full disclosure is not assured, since experts may choose to make a prediction before getting their final signal. A further caveat is the assumption that expert identity can be verified by the mechanism.

### 4.3. Discounting

The last, and, we will argue, the most successful mechanism we suggest for incentive compatibility is *discounting*.

Discounting essentially uses a proper market scoring rule. As explained in the Introduction, for a proper scoring rule  $S(\mathbf{p}, r)$ , the market scoring rule scores a prediction  $\mathbf{p}$   $S(\mathbf{p}, r) - S(\mathbf{p}', r)$  where  $\mathbf{p}'$  is the outcome distribution that was replaced by  $\mathbf{p}$ , and  $r$  is the outcome. Now any scoring rule can be scaled by an arbitrary constant  $k$  and remain proper. Furthermore,  $kS(\mathbf{p}, r) - k'S(\mathbf{p}', r)$  where  $k, k'$  may be different, is also proper. Generally, we can employ a time-varying scale factor  $k(t) \in \mathbb{R}_{>0}$ , and use the proper scoring rule  $k(t)S(\mathbf{p}, r)$ , where  $t$  is the time (which may be the elapsed time from some base, or an integer counter of events) of announcement of  $\mathbf{p}$ . Discounting means choosing  $k(t)$  that is weakly decreasing in  $t$ , whence we get a discounted scoring rule  $D(\mathbf{p}, r, t) := k(t)S(\mathbf{p}, r)$ , where  $t$  is the time of prediction  $\mathbf{p}$ . When  $t' < t$  is the time replaced prediction  $\mathbf{p}'$  was made, the discounted market scoring rule for prediction  $\mathbf{p}$  is

$$D(\mathbf{p}, r, t) - D(\mathbf{p}', r, t') = k(t)S(\mathbf{p}, r) - k(t')S(\mathbf{p}', r)$$

The idea of discounting is that, if without discounting Alice's optimal strategy is to hide or distort information on her first prediction, in order to reap a bigger benefit on her second prediction, her calculus will change if a sufficiently steep discount, motivating earlier predictions, is imposed on her reward.

We must be careful to use non-positive scoring rules with the discounting mechanism. Otherwise a possibility exists that the discounted scoring rule will have negative expectation for a prediction, creating a situation in which a player will prefer *not* predicting to making a truthful prediction. On the other hand, for non-positive scoring rules,  $S(\mathbf{p}, r) \leq 0$ , so that, whenever  $0 \leq k(t) \leq k(t')$

$$k(t)S(\mathbf{p}, r) - k(t')S(\mathbf{p}', r) \geq k'(t)[S(\mathbf{p}, r) - S(\mathbf{p}', r)]$$

So, if our original scoring rule had positive expectation, so does our discounted one.

Note that the logarithmic scoring rule  $\log p_r$  and quadratic scoring rule  $2p_r - \mathbf{p} \cdot \mathbf{p} - 1$ , which we have analyzed, are non-positive. Any scoring rule can be made non-positive by affine transform.

Being a market scoring rule, the discounting mechanism has bounded loss. More specifically

**Proposition 10.** *A non-negative discounted market scoring rule  $k(t)S(\mathbf{p}, r)$ , where  $t$  is the time of prediction  $\mathbf{p}$  and  $k(t)$  is a discounting function, has loss expectation bounded by  $-k(0)S(\boldsymbol{\pi}, \boldsymbol{\pi})$ , where  $\boldsymbol{\pi}$  is the prior ( $t = 0$  prediction).*

PROOF. The expectation of  $k(t)S(\mathbf{p}, r)$  for  $r \sim \mathbf{p}$  is  $k(t)S(\mathbf{p}, \mathbf{p})$ . Since the discounted scoring rule is a market scoring rule per Hanson (2003), its expected loss for  $r \sim \boldsymbol{\pi}$  and  $t = 0$  is, by (1) (Section 1.2), at most  $-k(0)S(\boldsymbol{\pi}, \boldsymbol{\pi})$ .

**Remark 1.** *Non-positive scoring rules suffer from an artefact that is the mirror image of positive ones: Experts have positive score expectation ( $[k(t) - k(t')]S(\mathbf{p}, r)$ ) for merely repeating the current prediction, and so will do so even if they have no information. In fact, the mechanism can offer them this “reward” automatically, sparing them the need to make an empty prediction. While this is ugly, its effect is minor, and the expected loss is still bounded by  $-k(0)S(\boldsymbol{\pi}, \boldsymbol{\pi})$  by Proposition 10. As in other cases, attempts to mend this would be counterproductive. E.g., if we deduct the “unearned”  $[k(t) - k(t')]S(\mathbf{p}, r)$  from every score, we will compromise truthfulness, as this reverts to the discounting mechanism proposed in Chen et al. (2010).*

**Proposition 11.** *Using the notation of Section 2.3, if  $S$  is a non-positive scoring rule, and there exists  $K$  such that*

$$K \geq \frac{\mathbb{E}_{b \sim B_1 | (A_1=a)} [S(h(a, b), h(a, b)) - S(h(\hat{a}, b), h(a, b))]}{S(g(a), g(a)) - S(g(\hat{a}), g(a))} \quad (9)$$

*for every  $A_1 = a$  and every possible  $\hat{a}$ , then the game is promptly truthful with a discount factor  $k(t)$  that satisfies  $k(t_1)/k(t_2) \geq K$ , where  $t_1$  is the time of Alice's first prediction and  $t_2$  is the time of the second.*

PROOF. If (9) is satisfied, scoring rule  $k(t)S(p, r)$  satisfies Corollary 2.  $\square$

From Proposition 11, we state sufficient conditions for discounting to succeed.

**Corollary 12.** *If*

1.  *$S$  is a non-positive strictly proper scoring rule, and*
2.  *$\mathbb{E}_{b \sim B_1 | (A_1=a)} [S(h(a, b), h(a, b)) - S(h(\hat{a}, b), h(a, b))]$  is bounded for every  $a, \hat{a}$* 
  - *In particular, if  $S$  is the quadratic scoring rule, or any other bounded scoring rule, and*
3. *The right-hand side of (9) is bounded when  $\hat{a} \rightarrow \pm\infty$ , and*
4.  *$\lim_{\hat{a} \rightarrow a} \frac{\frac{\partial^2}{\partial \hat{a}^2} S(h(\hat{a}, b), h(a, b))}{\frac{\partial^2}{\partial \hat{a}^2} S(g(\hat{a}), g(a))}$  exists for every  $a, b$ .*

*Then there exists a discount factor effective at restoring truthfulness.*

PROOF. As  $S$  is strictly proper, and the nominator of (9) is bounded, (9) can be unbounded only at infinity or at  $\hat{a} = a$ . At  $\hat{a} = a$ , (9) evaluates to  $\frac{0}{0}$ , so we invoke L'Hôpital's rule to find the limit. By the definition of a proper score rule, first derivatives again evaluate to  $\frac{0}{0}$ , so invoke L'Hôpital again for second derivatives.  $\square$

Discounting can restore truthfulness in the untruthful settings we found in Section 3.

**Proposition 13.** *Let Alice's, Bob's and the public signals in an Alice-Bob-Alice game belong to a jointly multivariate normal distribution, with covariance matrix given in (5). When using the discounted logarithmic scoring rule  $k(t) \log p_r$ , players are truthful and prompt iff signals are not perfectly correlated (i.e.  $|\rho| < 1$ ) and*

$$k(t_1)/k(t_2) \geq \frac{1 + \tau_C/\tau_A}{1 - \rho^2} \left/ \left[ 1 + \frac{(1 - \rho^2)(\tau_B + \tau_C)}{(\sqrt{\tau_A} - \rho\sqrt{\tau_B})^2} \right] \right. \quad (10)$$

*where  $t_1$  is the time of Alice's first prediction and  $t_2$  is the time of the second.*

PROOF. See Appendix.

Since Alice's signal is informative,  $\tau_A > 0$ , and Proposition 13 shows that discounting will be effective unless  $|\rho| = 1$ . Note that when the right-hand side of (10) evaluates to  $\leq 1$ , discounting is unnecessary. E.g. for  $\rho = 0$ , we have determined (Proposition 5) that truthfulness exists, and (10) evaluates to  $k(t_1)/k(t_2) \geq \frac{\tau_A + \tau_C}{\tau_A + \tau_B + \tau_C}$ . Thus setting  $k(t_1)/k(t_2) = 1$  is satisfactory.

**Proposition 14.** *Let Alice’s, Bob’s and the public signals in an Alice-Bob-Alice game belong to a jointly multivariate normal distribution, with covariance matrix given in (5). Then, iff signals are not perfectly correlated (i.e.  $|\rho| < 1$ ), there exists a discounted quadratic scoring rule  $k(t)[2p_r - \mathbf{p} \cdot \mathbf{p} - 1]$  for which players are truthful and prompt.*

PROOF. See Appendix.

Finally, we can formulate a proposition for the prompt truthfulness of a general prediction forum.

**Theorem 15.** *The public and a set of experts predict a variable  $\lambda$ . Let there be a fixed schedule in  $[0, T]$  specifying when experts and public receive signals, and experts may make a prediction. Agents may receive multiple signals, and experts may have multiple prediction opportunities.*

*Then the forum is generally truthful and prompt with a discount function  $k(t) \in \mathbb{R}_{>0}$  that makes all Alice-Bob-Alice subgames in the schedule truthful and prompt. These subgames are all occurrences where any expert, identified as “Alice”, has two consecutive prediction opportunities, and all experts who makes predictions in between are rolled into a single player identified as “Bob”.*

PROOF. See Appendix.

## 5. Discussion

### 5.1. Conclusions

We showed that using proper scoring rules for rewarding predictions often leads to reticence or deception. We formulated criteria for determining which prediction settings are truthful, and made a complete classification of the class of prediction settings with normal and lognormal signals, under the logarithmic and quadratic scoring rules. We suggested three new strategy-proof mechanisms, and identified one of them, discounting, as having all desirable properties, and proved the applicability and effectiveness of the discounting mechanism. Guidelines to implementing it are offered in Section 5.3 below.

### 5.2. A Market Maker for the Discounting Mechanism

Hanson (2003) has shown that every market scoring rule can be implemented by an automated market maker, who provides liquidity, and is willing to make any fair trade in shares of the form “pay \$1 if outcome is  $r$ ”, for every  $r$ . This applies to our discounting mechanism for the logarithmic scoring rule, which, as noted, takes the form of a market scoring rule for  $S(\mathbf{p}, r) = k(t) \log p_r$ .

Applying Hanson’s explanations to our case, if at  $t$  there is an inventory  $\mathbf{s}$  of shares  $s_r$  for every  $r$ , the instantaneous share price is

$$m(s_r) = e^{\frac{s_r}{k(t)}} / \int_{-\infty}^{\infty} e^{\frac{s_r}{k(t)}} dx$$

Assuming an infinitesimal trade path, this induces a cost function  $C(\mathbf{s}, t)$ . The cost of a trade changing the inventory from  $\mathbf{s}$  to  $\mathbf{s}'$  at  $t'$ , is

$$C(\mathbf{s}', t') - C(\mathbf{s}, t) := k(t') \log \int_{-\infty}^{\infty} e^{\frac{s'_r}{k(t')}} dx - k(t) \log \int_{-\infty}^{\infty} e^{\frac{s_r}{k(t)}} dx$$

Since  $k(t)$  is weakly decreasing in  $t$ , differentiating the above expression with respect to  $t'$  shows that the more any trade with the automated market maker is delayed, the higher its cost.

### 5.3. Implementation Notes

To implement the discounting mechanism  $k(t)S(p, r)$ , implement a standard LMSR prediction market with  $S(p, r) = \log p_r$ , adding a discounting scheme  $k(t)$ .

We offer some advice on how to select the discounting scheme, which is admittedly speculative or vague at times. Let us suppose that our prediction market is organized for the example which opens the paper: the estimation of company earnings by analysts. Instead of (or in addition to) issuing research reports, the analysts make predictions in the prediction market.

Proposition 11, in conjunction with Theorem 15, specifies a lower bound on the discount factor to be effective. The bound is dependent on the information model, which is typically known only in outline, and often only from empirical data. There is no upper bound, so imposing a discount factor so steep that it surely exceeds the required minimum is possible. This rough and simple method would work, technically, because the motivation provided by scoring rules is unaffected by scale. Practically, however, when translated into actual compensation, size does matter. To design more subtle, but still suitable discounting, we point out when discounting is most needed, and when it is not.

As is apparent from Example 1 and throughout the paper, withholding or distorting information makes sense only until that information becomes known anyway. Regardless of discounting, Alice will not distort information that Bob is likely to find out through his own devices. In most multi-player settings, like the analyst forum, private data and insights are temporary: What Alice knows now, Bob will figure out sooner or later, and vice versa. Deception is time-limited.

It is therefore probably safe to *not* discount during periods of prediction inactivity, or infrequent activity. It may even be safe to *reset* the discount factor to its original value at a quiet prediction period: Even if Alice knows about the reset in advance, and calculates that it is profitable to lie about an insight or information she has, her scheme will be foiled if her insight, or information, will become known to others through other means. If the reset is too far in the future, this becomes a losing bet. And, so long as the number of resets is bounded, the mechanism is still loss-bounded.

Of course, an *unannounced* reset of the discount factor is even better, as distortion in anticipation of a reset is not possible.

Discounting is most important when an information *shock* occurs. This is an event, or news, that reaches all analysts, and creates a scramble to update their previous now-obsolete estimate. An example is when the company issues an earnings warning, or new competition is announced. When this happens, all analysts suddenly have highly-correlated private information (by definition, all information not encompassed in the market price). Until a new market price stabilizes, there is great scope for deception. For example, the first analyst might try to deceive his peers about how severely he thinks the news will affect the earnings number. A strong discount is needed to put aside such thoughts.

It therefore seems appropriate to discount based on prediction *number* rather than prediction *time*, i.e.  $k(t)$  where  $t$  is a prediction counter rather than time-of-day, e.g.,  $k(t) = a^t$  where  $a \in (0, 1)$ . This has the merit of discounting heavily in periods of high activity, and lightly in periods of low activity. Resetting the discounting, by resetting  $t$  to 0, at unannounced random times may supplement this, to keep the discount factor from falling too low.



## References

## References

- Olivier Armantier and Nicolas Treich. 2013. Eliciting beliefs: Proper scoring rules, incentives, stakes and hedging. *European Economic Review* 62 (2013), 17–40.
- Yossi Azar, Amir Ban, and Yishay Mansour. 2016. When Should an Expert Make a Prediction?. In *Proceedings of the 2016 ACM Conference on Economics and Computation*. ACM, 125–142.
- G. W. Brier. 1950. Verification of forecasts expressed in terms of probability. *Weather Rev* 78 (1950), 1–3.
- Yiling Chen, Stanko Dimitrov, Rahul Sami, Daniel M. Reeves, David M. Pennock, Robin D. Hanson, Lance Fortnow, and Rica Gonen. 2010. Gaming Prediction Markets: Equilibrium Strategies with a Market Maker. *Algorithmica* 58, 4 (2010), 930–969.
- Yiling Chen, Daniel M Reeves, David M Pennock, Robin D Hanson, Lance Fortnow, and Rica Gonen. 2007. Bluffing and strategic reticence in prediction markets. In *International Workshop on Web and Internet Economics*. Springer, 70–81.
- Yiling Chen and Bo Waggoner. 2016. Informational substitutes. In *Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on*. IEEE, 239–247.
- Vincent Conitzer. 2009. Prediction markets, mechanism design, and cooperative game theory. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*. AUAI Press, 101–108.
- Bruno De Finetti. 1937. La prevision: Ses lois logiques, ses sources subjectives. *Ann. Inst. Henri Poincaré* 7 (1937), 1–68.
- Stanko Dimitrov and Rahul Sami. 2008. Non-myopic strategies in prediction markets. In *Proceedings of the 9th ACM conference on Electronic commerce*. ACM, 200–209.
- Xi Alice Gao, Jie Zhang, and Yiling Chen. 2013. What you jointly know determines how you act: Strategic interactions in prediction markets. In *Proceedings of the fourteenth ACM conference on Electronic commerce*. ACM, 489–506.
- Tilman Gneiting and Adrian E Raftery. 2007. Strictly Proper Scoring Rules, Prediction, and Estimation. *J. Amer. Statist. Assoc.* 102, 477 (2007), 359–378.
- I. J. Good. 1952. Rational Decisions. *Journal of the Royal Statistical Society. Series B (Methodological)* 14, 1 (1952), pp. 107–114.
- Robin Hanson. 2003. Combinatorial Information Market Design. *Information Systems Frontiers* 5, 1 (2003), 107–119.
- Yuqing Kong and Grant Schoenebeck. 2018. Optimizing Bayesian Information Revelation Strategy in Prediction Markets: the Alice Bob Alice Case. In *LIPICs-Leibniz International Proceedings in Informatics*, Vol. 94. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.

Theo Offerman, Joep Sonnemans, Gijs Van de Kuilen, and Peter P Wakker. 2009. A truth serum for non-bayesians: Correcting proper scoring rules for risk attitudes. *The Review of Economic Studies* 76, 4 (2009), 1461–1489.

Y.L. Tong. 2012. *The Multivariate Normal Distribution*. Springer Science & Business Media.

## Appendix A. Proof of Theorem 1

PROOF. Let  $A_1 = a, B_1 = b$ . Assume that Alice, on her first prediction, announces her prediction as  $\hat{a}$ , where possibly  $\hat{a} \neq a$ . Alice's belief at this time is  $g(a)$ , while all other agents' beliefs after Alice's announcement is  $g(\hat{a})$ . As discussed, Bob's prediction and Alice's second prediction are truthful. Therefore, after Bob's prediction the public belief is  $h(\hat{a}, b)$ . After Alice's second and last prediction, which is truthful, the public belief is  $h(a, b)$ .

Alice's total reward given outcome  $x$ ,  $\Pi_A(x)$  is

$$\Pi_A(x; a, \hat{a}, b) = \left\{ S(g(\hat{a}), x) - S(\pi, x) \right\} + \left\{ S(h(a, b), x) - S(h(\hat{a}, b), x) \right\} \quad (\text{A.1})$$

Bob's total reward given outcome  $x$ ,  $\Pi_B(x)$  is

$$\Pi_B(x; a, \hat{a}, b) = S(h(\hat{a}, b), x) - S(g(\hat{a}), x) \quad (\text{A.2})$$

The sum of Bob's and Alice's rewards is  $\Pi_A(x; a, \hat{a}, b) + \Pi_B(x; a, \hat{a}, b) = S(h(a, b), x) - S(\pi, x)$ , which does not depend on  $\hat{a}$ . Therefore the  $\hat{a}$  that maximizes Alice's reward is the one that minimizes Bob's reward. This is true for every  $x$  and  $b$ , and therefore it is also true for Alice's reward expectation at the time of her first prediction. This is before Alice learned that  $B_1 = b$ , but had a conditional distribution of it  $b \sim B_1 | (A_1 = a)$  based on her belief. In summary

$$\arg \max_{\hat{a}} \mathbb{E}_{b \sim B_1 | (A_1 = a)} \mathbb{E}_{\lambda \sim h(a, b)} \Pi_A(\lambda; a, \hat{a}, b) = \arg \min_{\hat{a}} \mathbb{E}_{b \sim B_1 | (A_1 = a)} \mathbb{E}_{\lambda \sim h(a, b)} \Pi_B(\lambda; a, \hat{a}, b)$$

□

## Appendix B. Proof of Proposition 4

PROOF. We know the player signals  $A_0, B_0$  from their prior predictions  $A_1, B_1$  by (4).

First, let us calculate the posterior distribution  $\pi_{AB} \sim N(\mu_{AB}, 1/\tau_{AB})$  from Alice's signal  $\pi_A \sim N(A_0, 1/\tau_A)$  and Bob's,  $\pi_B \sim N(B_0, 1/\tau_B)$ , alone.

Define an auxiliary signal  $\pi'_B = (1 - \alpha)\pi_A + \alpha\pi_B$  where  $\alpha := \frac{1}{1 - \rho\sqrt{\frac{\tau_A}{\tau_B}}}$ , whose mean is  $B'_0 = (1 - \alpha)A_0 + \alpha B_0$ . Its variance is

$$\begin{aligned} \sigma_B'^2 &= (1 - \alpha)^2 \text{Var} \pi_A + 2\alpha(1 - \alpha) \text{Cov}(\pi_A, \pi_B) + \alpha^2 \text{Var} \pi_B \\ &= (1 - \alpha)^2 \sigma_A^2 + 2\alpha(1 - \alpha) \rho \sigma_A \sigma_B + \alpha^2 \sigma_B^2 \end{aligned}$$

and accuracy  $\tau_B' = 1/\sigma_B'^2$ .

$\pi_A, \pi'_B$  are mutually independent, because

$$\text{Cov}(\pi_A, \pi'_B) = (1 - \alpha) \text{Var} \pi_A + \alpha \text{Cov}(\pi_A, \pi_B) = \frac{1 - \alpha}{\tau_A} + \frac{\alpha \rho}{\sqrt{\tau_A \tau_B}} = \frac{-\frac{\rho}{\sqrt{\tau_A \tau_B}} + \frac{\rho}{\sqrt{\tau_A \tau_B}}}{1 - \rho\sqrt{\frac{\tau_A}{\tau_B}}} = 0$$

Then  $\pi_{AB}$  is inferred from  $\pi_A, \pi'_B$  by the standard inference rule for independent normal observations.

$$\begin{aligned} \tau_{AB} &= \tau_A + \tau_B' \\ \mu_{AB} &= \frac{\tau_A A_0 + \tau_B' B'_0}{\tau_{AB}} \end{aligned}$$

Substituting for  $\alpha$  and simplifying, we get

$$\mu_{AB} = \frac{(\tau_A - \rho\sqrt{\tau_A\tau_B})A_0 + (\tau_B - \rho\sqrt{\tau_A\tau_B})B_0}{\tau_A - 2\rho\sqrt{\tau_A\tau_B} + \tau_B} \quad (\text{B.1})$$

$$\tau_{AB} = \frac{\tau_A - 2\rho\sqrt{\tau_A\tau_B} + \tau_B}{1 - \rho^2} \quad (\text{B.2})$$

Now, as the public signal  $\boldsymbol{\pi} \sim N(\boldsymbol{\pi}, 1/\tau_C)$  is independent of both  $\boldsymbol{\pi}_A$  and  $\boldsymbol{\pi}_B$ , it is independent of  $\boldsymbol{\pi}_{AB}$ . Applying again the inference rule for independent normal observations, we have

$$\tau_{ABC} = \tau_{AB} + \tau_C \quad (\text{B.3})$$

$$\mu_{ABC} = \frac{\tau_{AB}\mu_{AB} + \tau_C C_0}{\tau_{AB} + \tau_C} \quad (\text{B.4})$$

Substituting (B.1) and (B.2) in (B.4) and (B.3) we get (6) and (7).  $\square$

## Appendix C. Proof of Proposition 5

We first prove a lemma for divergences of the logarithmic scoring rule.

**Lemma 16.** *Let  $s$  be the logarithmic scoring rule, and let  $\mathbf{p} \sim N(\mu, 1/\tau)$  and  $\hat{\mathbf{p}} \sim N(\hat{\mu}, 1/\tau)$  be two normal distributions with the same accuracy  $\tau$ . Then*

$$S(\hat{\mathbf{p}}, \mathbf{p}) - S(\mathbf{p}, \mathbf{p}) = -\frac{\tau}{2}(\hat{\mu} - \mu)^2 \quad (\text{C.1})$$

PROOF.

$$\begin{aligned} S(\hat{\mathbf{p}}, \mathbf{p}) &= \frac{\tau}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{-\frac{\tau}{2}(x-\mu)^2} \log\left\{\frac{\tau}{\sqrt{2\pi}} e^{-\frac{\tau}{2}(x-\hat{\mu})^2}\right\} dx \\ &= \log \frac{\tau}{\sqrt{2\pi}} - \frac{\tau}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{-\frac{\tau}{2}(x-\mu)^2} \left\{\frac{\tau}{2}(x-\hat{\mu})^2\right\} dx \\ &= \log \frac{\tau}{\sqrt{2\pi}} - \frac{\tau}{2} \left[\frac{1}{\tau} + (\hat{\mu} - \mu)^2\right] \\ &= \log \frac{\tau}{\sqrt{2\pi}} - \frac{1}{2} - \frac{\tau}{2}(\hat{\mu} - \mu)^2 \end{aligned}$$

where we used the fact that the second moment of  $N(\mu, \sigma^2)$  is  $\sigma^2 + \mu^2$ .

Consequently (substitute  $\mathbf{p}$  for  $\hat{\mathbf{p}}$  above),

$$S(\mathbf{p}, \mathbf{p}) = \log \frac{\tau}{\sqrt{2\pi}} - \frac{1}{2}$$

and the lemma follows.  $\square$

The proof of the proposition follows.

PROOF. We use Corollary 3. Suppose that  $A_1 = a$ , and that Alice, when making her first prediction, pretends that  $A_1 = \hat{a} := a + c'$ , where  $c' = \frac{\tau_A}{\tau_A + \tau_C}c$ . I.e., by (4), Alice pretends her signal is  $A_0 + c$ .

By the Corollary, Alice is truthful if

$$\Delta(c; a, b) := [S(h(\hat{a}, b), h(a, b)) - S(h(a, b), h(a, b))] - [S(g(\hat{a}), g(a)) - S(g(a), g(a))] \geq 0$$

for every  $\hat{a}$  and  $b$ .  $g(a)$ , Alice's distribution at her first prediction, is given by (4). When  $B_1 = b$ , the inferred distribution from Alice's and Bob's predictions  $h(a, b)$  is given by Proposition 4. Substituting these, and using Lemma 16:

$$\begin{aligned}\Delta(c; a, b) &= \frac{\tau_{AC}}{2} \left( c \frac{\tau_A}{\tau_A + \tau_C} \right)^2 - \frac{\tau_{ABC}}{2} \left( c \frac{(\tau_A - \rho\sqrt{\tau_A\tau_B})}{\tau_A - 2\rho\sqrt{\tau_A\tau_B} + \tau_B + (1 - \rho^2)\tau_C} \right)^2 \\ &= c^2 \left\{ \frac{\tau_A^2}{\tau_A + \tau_C} - \frac{(\tau_A - \rho\sqrt{\tau_A\tau_B})^2}{(1 - \rho^2)(\tau_A - 2\rho\sqrt{\tau_A\tau_B} + \tau_B + (1 - \rho^2)\tau_C)} \right\} \\ &= c^2 \tau_A \left\{ \frac{\tau_A}{\tau_A + \tau_C} - \frac{(\sqrt{\tau_A} - \rho\sqrt{\tau_B})^2}{(1 - \rho^2)(\sqrt{\tau_A} - \rho\sqrt{\tau_B})^2 + (1 - \rho^2)^2(\tau_B + \tau_C)} \right\}\end{aligned}\tag{C.2}$$

(C.2) shows that  $c = 0$ , i.e., truthfulness, is either a global minimum or a global maximum of  $\Delta(c; a, b)$  depending on the sign of the factor multiplying  $c^2$ . Furthermore, this factor does not depend on  $b$  (nor on  $a$ ), and therefore Corollary 3 applies: Alice is promptly truthful iff

$$\frac{\tau_A}{\tau_A + \tau_C} - \frac{(\sqrt{\tau_A} - \rho\sqrt{\tau_B})^2}{(1 - \rho^2)(\sqrt{\tau_A} - \rho\sqrt{\tau_B})^2 + (1 - \rho^2)^2(\tau_B + \tau_C)} \geq 0\tag{C.3}$$

Rearranging, we get

$$1 - \rho^2 + (1 - \rho^2)^2 \frac{\tau_B + \tau_C}{(\sqrt{\tau_A} - \rho\sqrt{\tau_B})^2} \geq \frac{\tau_A + \tau_C}{\tau_A} \Rightarrow\tag{C.4}$$

$$(1 - \rho^2)^2(\tau_B + \tau_C) \geq (\rho^2 + \frac{\tau_C}{\tau_A})(\sqrt{\tau_A} - \rho\sqrt{\tau_B})^2 \Rightarrow\tag{C.5}$$

$$\left(1 - \rho^2\right)^2 \left(1 + \frac{\tau_C}{\tau_B}\right) \geq \left(\rho^2 + \frac{\tau_C}{\tau_A}\right) \left(\sqrt{\frac{\tau_A}{\tau_B}} - \rho\right)^2\tag{C.6}$$

as claimed.  $\square$

## Appendix D. Proof of Proposition 7

We first prove a lemma for divergences of the quadratic scoring rule.

**Lemma 17.** *Let  $s$  be the quadratic scoring rule, and let  $\mathbf{p} \sim N(\mu, 1/\tau)$  and  $\hat{\mathbf{p}} \sim N(\hat{\mu}, 1/\tau)$  be two normal distributions with the same accuracy  $\tau$ . Then*

$$S(\hat{\mathbf{p}}, \mathbf{p}) - S(\mathbf{p}, \mathbf{p}) = \frac{\tau}{\sqrt{2\pi}} [e^{-\frac{\tau}{4}(\mu - \hat{\mu})^2} - 1]\tag{D.1}$$

PROOF.

$$\begin{aligned}\hat{\mathbf{p}} \cdot \mathbf{p} &= \frac{\tau^2}{2\pi} \int_{-\infty}^{\infty} \exp\left[-\frac{\tau}{2}\left\{(x - \mu)^2 + (x - \hat{\mu})^2\right\}\right] dx \\ &= \frac{\tau^2}{2\pi} \int_{-\infty}^{\infty} \exp\left[-\frac{\tau}{2}\left\{2\left(x - \frac{\mu + \hat{\mu}}{2}\right)^2 + \frac{(\mu - \hat{\mu})^2}{2}\right\}\right] dx \\ &= \frac{\tau}{2\sqrt{2\pi}} e^{-\frac{\tau}{4}(\mu - \hat{\mu})^2} \frac{2\tau}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \exp\left[-\frac{2\tau}{2}\left(x - \frac{\mu + \hat{\mu}}{2}\right)^2\right] dx \\ &= \frac{\tau}{2\sqrt{2\pi}} e^{-\frac{\tau}{4}(\mu - \hat{\mu})^2}\end{aligned}$$

For the quadratic scoring rule

$$S(\hat{\mathbf{p}}, \mathbf{p}) - S(\mathbf{p}, \mathbf{p}) = \mathbb{E}_{r \sim \mathbf{p}} [S(\hat{\mathbf{p}}, r) - S(\mathbf{p}, r)] = 2\hat{\mathbf{p}} \cdot \mathbf{p} - \hat{\mathbf{p}} \cdot \hat{\mathbf{p}} - \mathbf{p} \cdot \mathbf{p}$$

Therefore

$$S(\hat{\mathbf{p}}, \mathbf{p}) - S(\mathbf{p}, \mathbf{p}) = \frac{\tau}{2\sqrt{2\pi}} [2e^{-\frac{\tau}{4}(\mu - \hat{\mu})^2} - 1 - 1] = \frac{\tau}{\sqrt{2\pi}} [e^{-\frac{\tau}{4}(\mu - \hat{\mu})^2} - 1]$$

□

Local truthfulness depends only on the sign of the second derivative of Alice's expectation of  $\Pi_B(x; a, \hat{a}, b)$ , as shown in the following lemma.

**Lemma 18.** *Let  $A_1 = a$ . Assume that Alice, on her first prediction, announces her prediction as  $\hat{a}$ , where possibly  $\hat{a} \neq a$ . Alice will be locally truthful if*

$$\frac{\partial^2}{\partial \hat{a}^2} \left[ \left\{ \mathbb{E}_{b \sim B_1 | (A_1 = a)} S(h(\hat{a}, b), h(a, b)) \right\} - S(g(\hat{a}), g(a)) \right] \Big|_{\hat{a}=a} > 0 \quad (\text{D.2})$$

PROOF. As  $s$  is proper, for every  $b$

$$\begin{aligned} \frac{\partial}{\partial \hat{a}} S(h(\hat{a}, b), h(a, b)) \Big|_{\hat{a}=a} &= 0 & \frac{\partial}{\partial \hat{a}} S(g(\hat{a}), g(a)) \Big|_{\hat{a}=a} &= 0 \\ \frac{\partial^2}{\partial \hat{a}^2} S(h(\hat{a}, b), h(a, b)) \Big|_{\hat{a}=a} &< 0 & \frac{\partial^2}{\partial \hat{a}^2} S(g(\hat{a}), g(a)) \Big|_{\hat{a}=a} &< 0 \end{aligned}$$

Since  $h(a, b)$  is a conditional distribution on  $g(a)$ , by the Law of Total Expectation we have

$$\mathbb{E}_{b \sim B_1 | (A_1 = a)} \mathbb{E}_{\lambda \sim h(a, b)} S(g(\hat{a}), \lambda) = \mathbb{E}_{\lambda \sim g(a)} S(g(\hat{a}), \lambda) = S(g(\hat{a}), g(a))$$

From the above and (A.2) we have that

$$\frac{\partial}{\partial \hat{a}} \mathbb{E}_{b \sim B_1 | (A_1 = a)} \mathbb{E}_{\lambda \sim h(a, b)} \Pi_B(\lambda; a, \hat{a}, b) \Big|_{\hat{a}=a} = 0$$

So  $\hat{a} = a$  is an extremum point for  $\mathbb{E}_{b \sim B_1 | (A_1 = a)} \mathbb{E}_{\lambda \sim h(a, b)} \Pi_B(\lambda; a, \hat{a}, b)$ , and so is a local minimum iff its second derivative is positive, or, equivalently, if (D.2) holds. □

The proof of the proposition follows.

PROOF. We use Corollary 2. Suppose that  $A_1 = a$ , and that Alice, when making her first prediction, pretends that  $A_1 = \hat{a} := a + c'$ , where  $c' = \frac{\tau_A}{\tau_A + \tau_C} c$ . I.e., by (4), Alice pretends her signal is  $A_0 + c$ .

By the Corollary, Alice is untruthful if

$$\Delta(c; a, b) := [S(h(\hat{a}, b), h(a, b)) - S(h(a, b), h(a, b))] - [S(g(\hat{a}), g(a)) - S(g(a), g(a))] < 0$$

for every  $a, b$  and  $c \neq 0$ .  $g(a)$ , Alice's distribution at her first prediction, is given by (4). When  $B_1 = b$ , the inferred distribution from Alice's and Bob's predictions  $h(a, b)$  is given by Proposition 4. Substituting these, and using Lemma 17:

$$\Delta(c; a, b) = \frac{1}{\sqrt{2\pi}} \left\{ \tau_{ABC} \left[ e^{-\frac{\tau_{ABC}}{4} \left( c \frac{(\tau_A - \rho \sqrt{\tau_A \tau_B})}{\tau_A - 2\rho \sqrt{\tau_A \tau_B} + \tau_B + (1 - \rho^2) \tau_C} \right)^2} - 1 \right] - \tau_{AC} \left[ e^{-\frac{\tau_{AC}}{4} \left( c \frac{\tau_A}{\tau_A + \tau_C} \right)^2} - 1 \right] \right\} \quad (\text{D.3})$$

Letting  $c \rightarrow \infty$  we deduce

$$\lim_{c \rightarrow \infty} \Delta(c; a, b) = \frac{1}{\sqrt{2\pi}} (-\tau_{ABC} + \tau_{AC}) < 0 \quad (\text{D.4})$$

because  $\tau_{ABC} > \tau_{AC}$ , since  $\pi_{ABC}$ 's distribution is conditional on  $\pi_{AC}$ , and Bob's signal is informative. As this does not depend on  $a$  nor on  $b$ , the setting is not truthful.

While not (globally) truthful, Alice may still be locally truthful. By Lemma 18, we need to find out when  $\frac{\partial}{\partial c^2} \Delta(c; a, b) \Big|_{c=0} > 0$ . From (D.3), noting that  $\frac{\partial}{\partial x^2} y e^{-zx^2} \Big|_{x=0} = -2yz$ .

$$\frac{\partial}{\partial c^2} \Delta(c; a, b) \Big|_{c=0} = -2\tau_{ABC} \frac{\tau_{ABC}}{4} \left( \frac{(\tau_A - \rho\sqrt{\tau_A\tau_B})}{\tau_A - 2\rho\sqrt{\tau_A\tau_B} + \tau_B + (1-\rho^2)\tau_C} \right)^2 + 2\tau_{AC} \frac{\tau_{AC}}{4} \left( \frac{\tau_A}{\tau_A + \tau_C} \right)^2 \quad (\text{D.5})$$

$$= -2\tau_{ABC} \frac{\tau_{ABC}}{4} \left( \frac{(\tau_A - \rho\sqrt{\tau_A\tau_B})}{\tau_{ABC}(1-\rho^2)} \right)^2 + 2\tau_{AC} \frac{\tau_{AC}}{4} \left( \frac{\tau_A}{\tau_{AC}} \right)^2 \quad (\text{D.6})$$

$$= \frac{1}{2}\tau_A^2 \left[ 1 - \left( \frac{(1-\rho\sqrt{\frac{\tau_B}{\tau_A}})}{(1-\rho^2)} \right)^2 \right] \quad (\text{D.7})$$

As this is true for every  $b$ , we deduce that local truthfulness holds for  $1 - \rho^2 > 1 - \rho\sqrt{\frac{\tau_B}{\tau_A}}$ . This entails  $0 < \rho < \frac{\sigma_A}{\sigma_B}$ , as claimed.  $\square$

## Appendix E. Proof of Proposition 8

PROOF. We show that truthfulness by all players is a Nash equilibrium, by showing that for each player, given truthfulness by every other player, if the player received her last signal, her best-response behavior is to make a truthful prediction. For suppose a player has received her last signal before the outcome is revealed. The player is aware of all her signals, as well as of previous predictions by her colleagues who have spoken.

The player then hypothesizes the following: First, that she makes a truthful prediction now, and second, she hypothesizes some sequence  $\mathcal{S}$  of predictions by other players, subject to the condition that all other players report truthfully their last signal (at least). The sequence ends in a last prediction  $\mathbf{p}(\mathcal{S})$ . Under the hypothesis,  $\mathbf{p}(\mathcal{S})$  is the player's subjective posterior distribution of the outcome from knowing all information in the (hypothetical) setting. Furthermore, it is the distribution that will be scored by the mechanism. That is to say, the player's score expectation for the scenario is  $S(\mathbf{p}, \mathbf{p})$ . Now, if the player deviates, by predicting untruthfully, or by failing to predict, the last prediction, still under the hypothesis, will be  $\mathbf{q}$ . Now since  $\mathcal{S}$  is proper,  $S(\mathbf{p}, \mathbf{p}) \geq S(\mathbf{q}, \mathbf{p})$  for any  $\mathbf{q}$ , and therefore there is no gain in deviating.

As this is true for *any* hypothetic future sequence  $\mathcal{S}$ , predicting truthfully at present is a best-response regardless of the future.  $\square$

## Appendix F. Proof of Proposition 13

PROOF. Referring to the proof of Proposition 5, from (C.3), clearly a discount factor  $k(t)$  that satisfies

$$k(t_1) \frac{\tau_A}{\tau_A + \tau_C} - k(t_2) \frac{(\sqrt{\tau_A} - \rho\sqrt{\tau_B})^2}{(1-\rho^2)(\sqrt{\tau_A} - \rho\sqrt{\tau_B})^2 + (1-\rho^2)^2(\tau_B + \tau_C)} \geq 0$$

will be effective. Rearranging this expression, the proposition follows.  $\square$

## Appendix G. Proof of Proposition 14

PROOF. The logarithmic scoring rule is bounded. Referring to the proof of Propostion 7, from (D.3) we see that

$$\lim_{\hat{a} \rightarrow \infty} \frac{\mathbb{E}_{b \sim B_1 | (A_1=a)} [S(h(a, b), h(a, b)) - S(h(\hat{a}, b), h(a, b))]}{S(g(a), g(a)) - S(g(\hat{a}), g(a))} = \frac{\tau_{ABC}}{\tau_{AC}}$$

is bounded. Furthermore, from (D.5),

$$\lim_{\hat{a} \rightarrow a} \frac{\frac{\partial^2}{\partial \hat{a}^2} S(h(\hat{a}, b), h(a, b))}{\frac{\partial^2}{\partial \hat{a}^2} S(g(\hat{a}), g(a))} = \left( \frac{(1 - \rho \sqrt{\frac{\tau_B}{\tau_A}})}{(1 - \rho^2)} \right)^2$$

which exists whenever  $|\rho| < 1$ . □

## Appendix H. Proof of Theorem 15

PROOF. We prove the proposition by backward induction. Suppose that a discount factor  $k(t)$  has been applied that would be sufficient to make all Alice-Bob-Alice interactions, with a “composite” Bob as described in the proposition, promptly truthful.

Assume that general prompt truthfulness is *not* in equilibrium. Call “Alice” the *last* player who is motivated to deviate from prompt truthfulness. This necessarily means that this is not Alice’s last prediction, and that all players who make predictions between Alice’s current prediction and her next one (call them *Bobbies*) are truthful.

Referring to Corollary 2 and its proof, it is readily seen that Alice maximizes her reward by minimizing the *aggregate* score of all Bobbies. Now, with a market scoring rule, in an interval where all players tell the truth, the aggregate score is the difference in score between final public information (public’s distribution before Alice speaks again) and initial public information (public’s distribution after Alice deviated). The sequence of events that led from initial information to final information has no effect on the aggregate score.

Any one of those Bobbies (call him Bob) has the same initial information, and, since all Bobbies speak the truth, the same final information. If we attribute all the signals of all Bobbies to Bob, we would again have the same initial information and same final information. By Corollary 2, Alice would tell the truth to minimize Bob’s expected aggregate score. She would therefore do the same to minimize the Bobbies’ aggregate score, contradicting our assumption that she deviates. This proves that no player deviates from prompt truthfulness. □