Program in Applied and Computational Mathematics

Project Report - Spring 2020

Will it Rain Tomorrow? Incentivizing Experts via Binary Scoring Rules

Georgy Noarov

supervised by
Matt Weinberg and Eric Neyman
Second reader:
Mark Braverman

Abstract

We consider a model that allows to compare binary scoring rules according to how well they elicit truthful, but also precise *reporting from experts. This is a new perspective on the incentive structure of scoring rules: traditionally, scoring rules are only required to be* proper*, i.e. make the expert reveal his true belief.*
Specifically, we define an incentivization index *that, in our model, lets compare any two strictly proper scoring rules by their power to entice the expert to perform costly effort in order to improve the precision (i.e. absolute error) of his report before submitting it. Furthermore, we find the unique proper scoring rule with the optimal incentivization index.*
In addition, we extend our results to the more general setting where precision is defined as the ℓ-th ($\ell \geq 1$) absolute moment of error.

1 Introduction

Note: This writeup and all results discussed therein (whether or not the respective definition, claim, lemma or theorem mentions it explicitly) are fully based on our paper [Neyman et al., 2020] coauthored with Eric Neyman and Matt Weinberg. The paper is quite general in its results, and the proofs are very involved analytically. Thus, our goal is to motivate the model we propose in the paper, as well as give a high-level view of its proofs. Thus, for instance, we consider here only the results obtained for a locally adaptive expert (see the definition below), while [Neyman et al., 2020] deals also with experts we call globally adaptive. Some other non-central results of the paper have been left out as well. The interested reader should refer to the full version of the mentioned paper for complete proofs.

Imagine an industrial company UmbrellaCo wanting to know, on a given day, what the chance of rain is for tomorrow. To this end, they contract a weather forecasting company, which naturally employs an expert in rain forecasts. Suppose that the expert can precisely simulate the next day, and see if it is raining or not on the simulated day. The simulation can be run as many times as the expert pleases. The total cost (resources + time spent running the computer test) of each simulation to the expert is c dollars, for some $c > 0$. The expert will estimate the chance of rain tomorrow as $\frac{h+1}{n+2}$ if he has run the simulation n times and saw rain in h of them. Equivalently, we model the event of rain tomorrow as a Bernoulli coin flip with some unknown probability $\mathfrak{p}$, and the expert starts with the prior value 0.5 on $\mathfrak{p}$ and performs a Bayesian update of it as he sees a sequence of coin flips drawn from that distribution.

Here is how the contract between UmbrellaCo and the expert works: UmbrellaCo chooses an increasing function $f : [0, 1] \to \mathbb{R}$ (which the expert knows). Suppose the expert has reported the chance of rain tomorrow to be $p \in [0, 1]$. On the next day, UmbrellaCo will observe whether it has rained or not. In the case it rains, it pays $f(p)$ dollars to the expert. In the case that it does not, it pays him $f(1 - p)$ dollars. In light of f being increasing, this is to represent that if it has indeed rained, then the company will pay more (closer to $f(1)$) for larger predicted probabilities, while if it has not rained, then the company will pay more for smaller predicted probabilities.

Of course, in order to game the system and receive a higher payment, the expert may report p while he actually believes that the chance of rain tomorrow is $q \neq p$. From his point of view, the expert would receive expected payout $qf(p) + (1 - q)f(1 - p)$ from the company. To make the expert always say the truth, the company therefore chooses f to be such that for every $q \in [0, 1]$, the $\arg\max_{p\in[0,1]}\{qf(p) + (1 - q)f(1 - p)\} = q$, and this argmax is unique.

Definition 1.1 (Good [1952])**.** *Such a function f is called a* (strictly) proper scoring rule. *If the argmax is allowed to not be unique, it is called a* weakly proper scoring rule. *The payoff the expert receives from reporting truthfully is then denoted*

$$R(p) := pf(p) + (1 - p)f(1 - p).$$

It may not be obvious that proper scoring rules exist, but they do – in fact, there is an infinite variety of such rules; Savage [1971] discusses some well-known ones and their general properties.

With f a proper scoring rule, the expert will definitely report his true belief to the company. Next, UmbrellaCo should really wish to know: how precise will the report be? Clearly, the more simulations the expert has performed, the closer his p will get to the true chance of rain. How many times will he want to simulate the rain? Let us model his behavior as *locally adaptive*: at any number of performed simulations, the expert performs another one if and only if his expected increase in payout from UmbrellaCo from doing so exceeds c, the cost of an extra simulation.

Definition 1.2. *The* locally adaptive *expert "flips the coin" (simulates the next day) for as long as the expected increase in payout for the next flip exceeds c:*

$$\frac{h+1}{n+2} R\left(\frac{h+2}{n+3}\right) + \left(1 - \frac{h+1}{n+2}\right) R\left(\frac{h+1}{n+3}\right) - R\left(\frac{h+1}{n+2}\right) > c.$$

How the expert simulates rain is now a stochastic process with a well-defined stopping rule. In addition to the above, the company UmbrellaCo supposes that the chance of rain tomorrow is itself distributed uniformly (UmbrellaCo's executive board consists of aliens from Mars and just does not know that due to the peculiarities of the Earth's weather, the chances of rain on various days throughout the year are definitely not uniformly distributed). The company is planning to use the expert's services for long enough that what matters to it is the expected expert's prediction error over many days. Mathematically, this is

Definition 1.3.

$$\mathit{Error}_f(c) := \mathbb{E}[|\hat{p} - p|],$$

where p is the random variable denoting the chance of rain, $\hat{p}$ the expert's estimate. Expectation is over the randomness of p and of the simulation process.

We may now ask the question: which proper scoring rule f minimizes this error, for fixed c? It turns out that such a rule does not exist! If f is a proper scoring rule, then any affine rescaling of this rule is also proper: $2f, 2f+10, 3f, 10f$, and so on. Clearly, $f+5$ or $3f$ pay more than f to the expert. So, he will simulate more on average. As one increases either $a \geq 1$ or $b \geq 0$, $af+b$ yields $\text{Error}_f(c) \to 0$.

Lemma 1.4. *As $c \to 0$, the expert's pay from UmbrellaCo tends to $\int_0^1 R(x)dx$.*

Proof. See Neyman et al. [2020]. □

Lemma 1.5. *The expert's reward $R(p)$ is minimized at $p = \frac{1}{2}$.*

Proof. $R = pf(p) + (1-p)f(1-p)$ is convex, and note it is symmetric around $\frac{1}{2}$ on $[0,1]$. □

Using the two lemmas together, we can rescale any proper scoring rule f so that its reward as $c \to 0$ is unit ($\int_0^1 R(x)dx = 1$), and the minimum pay the expert can receive from UmbrellaCo is $R(\frac{1}{2}) = 0$. Therefore, we may define normalized proper scoring rules as follows.

Definition 1.6 (Normalized proper scoring rule)**.** *A rule f which: 1) forces the expert to always report truthfully, 2) pays the expert at least 0 dollars, 3) as $c \to 0$, pays the expert 1 dollar on average.*

Given this definition, we may update our question from above, and we ask: how does UmbrellaCo choose among normalized proper scoring rules so as to minimize $\text{Error}_f(c)$? The result is provided by the following surprising theorem. For the meaning of 'analytically nice', please go to Section 3.

Theorem 1.7. *There exists a unique normalized proper scoring rule, g_{opt}, which is 'analytically nice' and which, when compared to any other 'analytically nice' f, satisfies*

$$\mathit{Error}_{g_{opt}}(c) < \mathit{Error}_f(c) \text{ for all } 0 < c < C_f, \text{ for some } C_f \text{ that depends on } f.$$

This rule is (where k is the rescaling constant):

$$g_{opt}(p) = \begin{cases} k\int_{\frac{1}{2}}^{p} \sqrt[5]{\frac{(1-t)^6}{t^7}}dt, & \text{if } p < 0.5 \\ k\int_{\frac{1}{2}}^{p} \sqrt[5]{\frac{t}{(1-t)^2}}dt, & \text{if } p \geq 0.5 \end{cases}$$

In the case UmbrellaCo does not seek an optimal scoring rule, but instead wants to compare two proper scoring rules of its own choice, it will be able to do so using this notion that we introduce in Neyman et al. [2020]:

Definition 1.8 (Incentivization index)**.** *For any 'analytically nice' normalized proper scoring rule f, we let its incentivization index be:*

$$Ind_f := \int_0^1 \sqrt[4]{\frac{x(1-x)}{R''(x)}}dx.$$

The incentivization index maps each nice enough scoring rule to a number. Its significance is that for any given nice enough scoring rules, f and g, we would be able to tell which one results in a smaller expert's error by simply computing and comparing the index for f and for g, which can be done due to the index's closed form.

Theorem 1.9. *For any two 'analytically nice' normalized proper scoring rules f, g,*

$$Ind_f < Ind_g \implies Error_f(c) < Error_g(c) \text{ for all } 0 < c < C_{f,g}, \text{ for some } C_{f,g}.$$

Here, again, the constant $C_{f,g}$ is only dependent on f and g.

Our main result in Neyman et al. [2020], which proves Theorem 1.7 and Theorem 1.9, is:

Theorem 1.10. *If f is an 'analytically nice' proper scoring rule, then as $c \to 0$,*

$$\sqrt[4]{\frac{1}{c}} \cdot Error_f(c) \to \sqrt{\frac{2}{\pi}} \cdot \sqrt[4]{2} \cdot Ind_f.$$

More generally, if $Error_f^\ell(c) := \mathbb{E}[|\hat{p} - p|^\ell]$ for $\ell \geq 1$, then as $c \to 0$:

$$c^{-\ell/4} \cdot Error_f^\ell(c) \to \mu_\ell \cdot 2^{\ell/4} \cdot Ind_f^\ell,$$

where

$$Ind_f^\ell := \int_0^1 \left(\frac{x(1-x)}{R''(x)}\right)^{\ell/4} dx$$

is the generalized incentivization index, and $\mu_\ell = \frac{2^{\ell/2}\Gamma(\frac{\ell+1}{2})}{\pi^{1/2}}$ is the ℓth moment of the standard Gaussian distribution.

1.1 Outline of this paper and the proofs

In Section 2, we briefly discuss some related work.

Section 3 contains a detailed outline of the proofs in our paper Neyman et al. [2020]. Specifically, in Section 3.1 we define how 'analytically nice' we require our scoring rules to be, and we mention that many well-known proper scoring rules fall into this group.

In Section 3.2, we discuss in a relatively detailed manner two simple results, which capture two essential characteristics of the expert's sampling process: First, we show a Taylor expansion-based formula on his expected increase in reward from round to round. Second, we demonstrate that as $c \to 0$, the expert will continue to flip for a guaranteed number of times, as an asymptotic function of c.

Section 3.3 contains a detailed description of the probability space that we set up so that the expert's coin-flipping trajectories are coupled for all possible values of the true rain probability p. Second, we define a sequence of events $\{\Omega_N\}$ that have probability $1-o(1)$, and guarantee that the expert's estimate will never stray from the true value of p by 'too much'. This is very important in the later course of the proof, since assuming Ω_N holds allows to show uniform convergence in p of many p-dependent random variables associated with the expert's coin-flipping process.

Section 3.4 first outlines the proof of how to capture the expert's stopping time quite precisely as a function of the random variable p and c, for all p except at the (vanishing as $c \to 0$) tails of the interval $[0, 1]$. It then concludes by showing how to obtain, from this bound on the stopping time, tight bounds on the value of the expert's error for any given p, c and, ultimately, how this results in the limit formula for the expert's expected error over all p declared in Theorem 1.10.

Note that much of Section 3 is just a detailed outline of the proofs in Neyman et al. [2020]. This is intentional: the proofs of the lemmas which we choose not to fully present here are very involved, and it is of more value to show the high-level logic of these proofs here; the interested reader should consult the original paper Neyman et al. [2020].

Finally, we give a brief conclusion, in which we highlight why our study is important in the context of incentivization via proper scoring rules.

2 Related Work

In a seminal paper, Brier [1950] defined the quadratic scoring rule (and described its applications in weather forecasting, like we do in this manuscript). Good [1952] formally defined the notion of scoring rule; he also brought in the logarithmic rule. The richness of the class of *proper* scoring rules, which incentivize truthful reporting by the expert, was explored soon after the introduction of scoring rules; see for instance Savage [1971]. Various properties of proper scoring rules, as applied to statistical prediction and estimation, are studied by Gneiting and Raftery [2007] and Dawid and Musio [2014]. Winkler et al. [1996] consider the "ex-ante" and the "ex-post" perspective on scoring rules; our view in this manuscript is ex-ante: the scoring rule is viewed as inducing a reward function, which the expert, in reporting his belief, takes into account.

We look to rank scoring rules by how well they incentivize the expert to acquire more precise knowledge of the distribution before they report their estimate. Interestingly, Tsakas [2019] considers the opposing perspective to ours. He recognizes that, for instance, in polling, it might be important to elicit the participants' pure, unaltered beliefs about a certain question. Therefore, he looks for scoring rules that incentivize the expert as much as possible to not perform any research and report his prior. Our model can be viewed as a restriction of Clemen [2002]'s model, where the expert is also propelled by the scoring rule to not just report his prior belief but do more research about the to-be-predicted distribution. In his model, somewhat more generally, the cost of each sample in the sampling sequence may be different to the expert. However, Clemen does not obtain rigorous results about how to compare different proper scoring rule by their incentivization power, which we do.

3 Proofs, Outlined

Here, we will refer to the expert's running estimate after n coin flips as Q_n. His stopping time will be denoted n_{stop}, and his marginal increase in expected reward at each flip will be called Δ_n.

3.1 Analytically nice proper scoring rules

Here, we shed light on which scoring rules yield to our analysis. The formal definition of this class of binary proper scoring rules, called *respectful* in Neyman et al. [2020], is as follows:

Definition 3.1. *A proper scoring rule f is analytically nice if its reward function R satisfies:*

1. *R is strongly convex on $(0, 1)$, namely, R'' is lower-bounded by some $a > 0$ on all of $(0, 1)$*
2. *R''' is Riemann integrable on any $[a, b] \subseteq (0, 1)$.*
3. *For any fixed ϵ, one can find constants t, C such that for all $0 < c < C$,*

$$|R'''(x)| \leq \frac{R''(x)}{c^{\frac{1}{6}-\epsilon}\sqrt{x(1-x)}}, \quad \text{for all } c^t \leq x \leq 1 - c^t.$$

An important question is: how rich is the defined class of scoring rules? Turns out, it is quite rich. First, we may see that the third condition is implied by an easier one: that $|R'''|$ is bounded on $(0, 1)$. With this in mind, we see that all polynomial scoring rules (i.e. scoring rules $f(x)$ that are polynomial in x) are 'analytically nice'. Indeed, the reward R corresponding to a polynomial f is also polynomial, and thus clearly satisfies Conditions 1 and 2, and also satisfies Condition 3 by virtue of its derivatives (and in particular R''') being bounded on $[0, 1]$. Therefore, in particular, the seminal Brier's quadratic scoring rule Brier [1950], which is expressed through its reward function $R(x) = x^2 + (1 - x)^2$, is also 'analytically nice'.

Somewhat more involved is to show that another famous proper scoring rule, the *logarithmic* rule given by $f(x) = \ln x$, is also 'analytically nice'. Its R''' is not bounded on $(0, 1)$, and so one cannot easily check the 3rd Condition of the definition above. However, the following (intricate) Claim shown in Neyman et al. [2020] ensures that also the logarithmic scoring rule, and a variety of other rules with diverging R''', are analytically nice.

Claim 1. *The 3rd condition of the definition above is implied, if R''' is bounded on any compact subset of $(0, 1)$ and moreover there are $k \neq 0$ and r such that $\lim_{x\to 0} x^r R'''(x) = k$.*

To conclude this discussion, we note that Neyman et al. [2020] proves a Weierstrass approximation theorem-based result showing that for a wide variety of proper scoring rules, there exists an arbitrary good approximation of those rules by polynomial ones (which are analytically good). Therefore, in that sense, a wide variety of proper scoring rules are approximately analytically good.

3.2 The probability space and its consequences

In order to simultaneously reason about the coin-flipping processes for all possible values of p in $[0, 1]$, we construct a common probability space for all these processes. Define the sequence $(\xi_n)_{n\geq 1}$ of iid. random variables each having uniform distribution on $[0, 1]$. Then for each $n \geq 1$ and for all $0 \leq p \leq 1$, we define the empirical process $Q_n(p) := \frac{1}{n+2}(1 + \sum_{i=1}^n 1_{\{\xi_i \leq p\}})$. It is easy to see that for any given value $p \in [0, 1]$, $Q_n(p)$ is precisely the running expert's estimate obtained from n Bernoulli distributed random variables with parameter p.

Now we will provide an asymptotic (as $n \to \infty$) upper bound on the fluctuations of the running estimate around the true value of p, and this bound will hold w.h.p. uniformly for all $p \in [1/n, 1 - 1/n]$. Before doing so, we define for every N the event Ω_N, by defining its complement as

$$\overline{\Omega}_N := \bigcup_{n=N}^{\infty} \bigcup_{j=0}^{n-1} \left\{ \left| Q_n\left(\frac{j}{n}\right) - \frac{j}{n} \right| > \frac{\sqrt{j(n-j)}}{2n^{1.49}} \right\}.$$

This definition of Ω_N is justified by the following lemma, which shows that given Ω_N, for all $p \in [1/n, 1-1/n]$ the fluctuations of the running estimate around the true value of p are uniformly bounded as $n \to \infty$.

Lemma 3.2. *There is a large enough N^* such that, for any $N \geq N^*$, conditioned on Ω_N, it holds for all $n \geq N$ that*

$$|Q_n(p) - p| \leq \frac{\sqrt{p(1-p)}}{n^{0.49}} \quad \forall n \geq N, \text{ and } \frac{1}{n} \leq p \leq 1 - \frac{1}{n}.$$

Proof. Take an arbitrary $n \geq N$ and an arbitrary $p \in [1/n, 1 - 1/n]$. Observe that there is an index j_p such that $1 \leq j_p \leq n - 2$ and $\frac{j_p}{n} \leq p \leq \frac{j_p+1}{n}$. Now, by monotonicity of $Q_n(\cdot)$ we have

$$\frac{j_p}{n} - \frac{\sqrt{j_p(n-j_p)}}{2n^{1.49}} \leq Q_n\left(\frac{j_p}{n}\right) \leq Q_n(p) \leq Q_n\left(\frac{j_p+1}{n}\right) \leq \frac{j_p+1}{n} + \frac{\sqrt{j_p(n-j_p)}}{2n^{1.49}},$$

implying

$$|Q_n(p) - p| \leq \frac{1}{n} + \frac{1}{2n^{1.49}} \max\{\sqrt{j_p(n-j_p)}, \sqrt{(j_p+1)(n-1-j_p)}\}.$$

Using that the right-hand side is maximized at p either $1/n$ or $1-1/n$, and that $\sqrt{p(n-p)}$ under the same conditions, we obtain after dividing through by $\sqrt{p(n-p)}$ and plugging say $p = 1/n$ in, that

$$\frac{|Q_n(p) - p|}{\sqrt{p(n-p)}} \leq \frac{1}{n\sqrt{p(1-p)}} + \frac{\sqrt{2}n}{2n^{1.49}} \leq \frac{1}{n^{0.49}}.$$

□

Furthermore, we demonstrate in Neyman et al. [2020] that the event Ω_N is extremely likely, the probability of its complement being exponentially small in N.

Lemma 3.3. *The probability of the event $\overline{\Omega}_N$ decreases exponentially fast in N:*

$$\Pr[\overline{\Omega}_N] = O(e^{-N^{0.01}}).$$

3.3 Stopping time estimation

This section begins with an asymptotic formula for the expected gain in reward from flipping one more time, as the overall number of flips increases. Then we determine how small c should be so that regardless of the true value of p and the flips' outcomes, the expert will not stop flipping until a set time n.

3.3.1 Expected increase in reward from flipping one more time

Lemma 3.4. *If R is twice differentiable at every point in $(0,1)$, then*

$$\Delta_{n+1} = \frac{Q_n\,(1-Q_n)}{2(n+3)^2}\left[Q_n R''(c_1) + (1-Q_n)\,R''(c_2)\right], \tag{1}$$

where $c_1 \in \left[\frac{S_n+1}{n+3}, Q_n\right]$ and $c_2 \in \left[Q_n, \frac{S_n+2}{n+3}\right]$ (in particular, each $c_i \in [Q_n - \frac{1}{n}, Q_n + \frac{1}{n}]$).

Proof. We use Taylor expansion with the remainder in Lagrange form. Observe that the proof below only requires that R be twice differentiable, and R' (resp. R'') be continuous, on the *open* interval $(0,1)$. This is because $\left[\frac{S_n+1}{n+3}, \frac{S_n+2}{n+3}\right] \subset (0,1)$ for all n.

To prove eq. (1), we specifically let $\delta = \frac{1}{n+3}$. Then, we note that the expected gain is equal to $Q_n R(\frac{S_n+2}{n+3}) + (1-Q_n)R(\frac{S_n+1}{n+3}) - R(Q_n)$, and expand both $R(\frac{S_n+2}{n+3})$ and $R(\frac{S_n+1}{n+3})$ around Q_n. As $\frac{S_n+2}{n+3} = Q_n + (1-Q_n)\delta$ and $\frac{S_n+1}{n+3} = Q_n - \delta Q_n$, which is checked by using the definitions of Q_n, δ, we have

$$\begin{aligned}
&Q_n R\left(\frac{S_n+2}{n+3}\right) + (1-Q_n) R\left(\frac{S_n+1}{n+3}\right) - R(Q_n) \\
&= Q_n \left(R(Q_n) + (1-Q_n)\delta R'(Q_n) + \frac{1}{2}(1-Q_n)^2\delta^2 R''(c_2)\right) \\
&+ (1-Q_n)\left(R(Q_n) - \delta Q_n R'(q) + \frac{1}{2}\delta^2 Q_n^2 R''(c_1)\right) - R(Q_n) \\
&= \frac{Q_n(1-Q_n)\delta^2}{2}\left(Q_n R''(c_1) + (1-Q_n)R''(c_2)\right).
\end{aligned}$$

Substitute back δ to obtain the desired statement. □

3.3.2 Guaranteeing a fixed number of coin flips

Lemma 3.5. *Assume R is twice differentiable and strongly convex on $(0,1)$. Then there exists a constant $C(R)$ which only depends on the reward function R (but not on p), such that for any given $n \geq 1$, $n_{stop} \geq n$ for any cost $c \leq \frac{C(R)}{n^3}$.*

Proof. Observe that by Equation (1), the expected gain in reward between steps $n, n+1$ equals $\Delta_{n+1} = \frac{1}{2(n+3)^2} Q_n (1-Q_n) [Q_n R''(c_1) + (1-Q_n) R''(c_2)]$ for some $c_1, c_2 \in [0,1]$. Writing $Q_n(1-Q_n) = \frac{(S_n+1)(n-S_n+1)}{(n+2)^2}$, we see that the numerator attains its minimum value $n+1$ when either $S_n = 0$ or $S_n = n$. Hence, $Q_n(1-Q_n) \geq \frac{n+1}{(n+2)^2} \geq \frac{1}{5n}$ for all $n \geq 1$.

Next, due to strong convexity of R, we may let $C_R := \min_{x\in(0,1)} R''(x) > 0$. Therefore, by the above

$$\Delta_{n+1} \geq \frac{1}{2(n+3)^2} \cdot \frac{1}{5n} [Q_n \cdot C_R + (1-Q_n) \cdot C_R] = \frac{C_R/10}{n(n+3)^2} \geq \frac{C_R/160}{n^3}.$$

So setting c less than or equal to $\frac{C_R}{160} n^{-3}$ guarantees that up until at least time n, the flipping will continue, because indeed for any time $m < n$, the expected gain from flipping once more is at least $\frac{C_R/160}{m^3} > \frac{C_R/160}{n^3} \geq c$. Q.e.d. □

3.4 A precise computation of bounds of the expert's stopping time

Now, roughly speaking, we show that as the *guaranteed* number of flips increases, $n_{\text{stop}}(p) \sim \frac{p(1-p)R''(p)}{2} c^{-1/2} (1 \pm O(c^{\frac{1}{300}}))$. This is easy to expect from the above Taylor-based formula for the expected marginal reward: indeed, one can say Δ_{n+1} is approximately $\frac{p(1-p)R''(p)}{2n^2}$, and equating this to the marginal cost c of an extra sample (the stopping condition for the expert) implies that $n_{stop} \approx \sqrt{\frac{p(1-p)R''(p)}{2c}}$. However, the formal proof of this is nontrivial.

Lemma 3.6. *Suppose Ω_N holds for a given N. Suppose the scoring rule is 'analytically nice', and recall the parameter t from its definition. Then, there are constants K, C such that for all $0 < c < C$,*

$$\sqrt{\frac{p(1-p)R''(p)}{2c}} \sqrt{1 - Kc^{1/300}} \leq n_{stop} \leq \sqrt{\frac{p(1-p)R''(p)}{2c}} \sqrt{1 + Kc^{1/300}} \text{ for all } 2c^t \leq p \leq 1 - 2c^t.$$

Proof. First, one ensures, by Lemma 3.5, that indeed at least N coins are flipped, so that we may indeed use that Ω_N holds. With a small technical assumption on c, one can see from Ω_N holding that

$$Q_n \in [p \pm \sqrt{p(1-p)} (\alpha c)^{0.49}],$$

for some constant α. Having this quite tight bound on Q_n, one then recalls the Taylor-based formula for the marginal expected reward,

$$\Delta_{n+1} = \frac{Q_n (1-Q_n)}{2(n+3)^2} \left[Q_n R''(c_1) + (1-Q_n) R''(c_2) \right]. \tag{2}$$

We would like to be able to show $R''(c_1) \approx R(p)$ and $R''(c_2) \approx R(p)$ in order to simplify this expression, and then elicit n_{stop} from equating the marginal reward to c, as discussed above. This is achieved via an intricate, but elementary, analytic lemma:

Claim 2.

$$|R''(p+\epsilon) - R''(p)| \leq R''(p)(e^{r|\epsilon|} - 1), \text{ for all } \epsilon \text{ such that } p, p+\epsilon \in [c^t, 1-c^t].$$

The proof of this claim relies on expressing the difference on the left-hand side as an integral of R''', and this is the main place where we use the most unintuitive, 3rd part of the definition of analytically nice scoring rules, which places a constraint on the rate of change of R'''.

As a result of this claim, we can show that for $i = 1, 2$, $|R''(c_i) - R''(p)| \leq 8\alpha^{0.49} R''(p) c^{1/300}$. Plugging this, as well as the above bound on Q_n that used Ω_N, into Lemma 3.4, one can see for K small enough that

$$\Delta_{n+1} \in [\frac{p(1-p)R''(p)(1 \pm Kc^{1/300})}{2n^2}],$$

and equating Δ_{n+1} with the marginal cost c, we get the declared bound on n_{stop}. □

Having tight guarantees on the magnitude of n_{stop}, depending on p, c, we are now in the position to tightly and uniformly (in p) bound the error of the expert over a range of true values of p. Formally, letting $err(p,c)$ be the random variable denoting the expected error of the expert for particular values of p, c, we have the following result:

Lemma 3.7. *For $\ell \geq 1$ and μ_ℓ the ℓth moment of a standard normal variable, and $N_c = (\alpha c)^{1/3}$,*

$$c^{-\ell/4}\, \mathbb{E}[(err(p,c))^\ell | \Omega_{N_c}] = (1 \pm o(1)) \mu_\ell \left(\frac{2p(1-p)}{R''(p)} \right)^{\ell/4} \quad \text{for all } 2c^t \leq p \leq 1 - 2c^t.$$

Here, the $o(1)$ term is with respect to c.

Proof. The crucial observation of the proof, which is quite long and involved, is that if we take n_0 to be $\sqrt{\frac{p(1-p)R''(p)}{2c}}\sqrt{1 - Kc^{1/300}}$, the lowest possible value of n_{stop} at given p, c according to the preceding lemma, then *the expert's guess will not change much at all* after the n_0th step. Thus, the final expert's error $err(p,c)$ will be approximable, in the limit, by $|Q_{n_0} - p|$.

The proof first shows the following claim:

$$\mathbb{E}[|Q_{n_0} - p|^\ell | \Omega_N] = \mu_\ell \left(\frac{p(1-p)}{n_0} \right)^{\ell/2} (1 + o(1)).$$

The proof of the claim relies on the Berry-Esseen theorem, a result which implies that the distribution function of the expert's guess becomes uniformly close to the CDF of a standard Gaussian variable.

After that, it is possible to show also that the contribution of all remaining steps from n_0 to n_{stop} to the expert's error vanishes in the limit: i.e., uniformly over all $p \in [2c^t, 1 - 2c^t]$, $\frac{\mathbb{E}[|Q_{n_0} - Q_{n_{stop}}|^\ell \Omega_N]}{\mathbb{E}[|Q_{n_0} - p|^\ell | \Omega_N]} \to 0$ as $c \to 0$. This claim and the above one give the desired conclusion of the lemma. □

Finally, we arrive at the proof of the main theorem 1.10.

Proof of Theorem 1.10. The proof essentially consists of integrating the preceding lemma over all $p \in [2c^t, 1 - 2c^t]$, then showing that the expert's error integrated over the tails $[0, 2c^t) \cup (1 - 2c^t, 1]$ is negligible, and finally, noting that one can remove the conditioning on Ω seen in the preceding lemma. The last point is due to Lemma 3.3, which shows that the complement of Ω_N has exponentially vanishing probability in N. □

Now, to conclude this section, we make a note about how we obtained the optimal scoring rule, as given in Theorem 1.7. This is a result of pointwise minimization of the integrand of the incentivization index formula, subject to the constraints that ensure a scoring rule is (weakly) proper. These constraints are derived e.g. in Neyman et al. [2020] and have a very simple algebraic form; see Section 4 of Neyman et al. [2020]. Thus, the optimization problem becomes solvable simply by writing down its Lagrangian that includes the weak-properness constraints. Finally, the paper shows that the obtained optimal rule is 'analytically good', so the incentivization index makes sense for it.

4 Conclusion

Proper scoring rules, which incentivize the expert to report truthfully, have been known since the 1950s. It is the only well-studied incentivization property of scoring rules. But then, assuming truthfulness, another very important incentivization issue is: how does the choice of a proper scoring rule motivate the expert to make his prediction as precise as possible – that is, do more research/simulations before reporting his belief? This is important: the expert might not refine his prediction and just make a baseless prediction of p that seems plausible to him at the moment. This would be an honest but useless prediction. Our model is the first one to rigorously study this question. The model is stylized, but nonetheless captures a variety of real-life scenarios, even with the assumption $c \to 0$ that we make.

References

Glenn W. Brier. Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78(1):1–3, 1950.

Robert T. Clemen. Incentive contracts and strictly proper scoring rules. *Test*, 11(1):167–189, Jun 2002. ISSN 1863-8260.

Alexander Philip Dawid and Monica Musio. Theory and applications of proper scoring rules. *METRON*, 72(2):169–183, Aug 2014. ISSN 2281-695X.

Tilmann Gneiting and Adrian E Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.

I. J. Good. Rational decisions. *Journal of the Royal Statistical Society: Series B (Methodological)*, 14(1):107–114, 1952.

Eric Neyman, Georgy Noarov, and S. Matthew Weinberg. Binary scoring rules that incentivize precision. *ArXiv*, abs/2002.10669, 2020.

Leonard J. Savage. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association*, 66(336):783–801, 1971. ISSN 01621459.

Elias Tsakas. Robust scoring rules. *SSRN*, Jul 2019.

R. L. Winkler, Javier Muñoz, José L. Cervera, José M. Bernardo, Gail Blattenberger, Joseph B. Kadane, Dennis V. Lindley, Allan H. Murphy, Robert M. Oliver, and David Ríos-Insua. Scoring rules and the evaluation of probabilities. *Test*, 5(1):1–60, Jun 1996. ISSN 1863-8260.