Scoring Rules for Adaptive Multi-stage Questions

Chun Kai Ling
Computer Science Department
Carnegie Mellon University
chunkail@andrew.cmu.edu

Gaurav Manek
Computer Science Department
Carnegie Mellon University
gmanek@andrew.cmu.edu

Abstract

Strictly Proper Scoring Rules (SPSRs) are typically designed to work with single questions. When applied to sequences of questions, one is restricted to deterministic question lists or strong assumptions of independence. We analyze the behavior of SPSRs when future questions are selected based on answers to present questions and show that in some cases they incentivize dishonesty. We propose a class of SPSRs called *Joint*-SPSRs which maximize the expected reward when all questions are answered honestly, even when future question selection depends on past answers. We propose one such rule, the *Multiplicative Quadratic Scoring Rule* and prove its correctness.

1 Introduction

Strictly Proper Scoring Rules (SPSRs) are scoring rules that are maximized in expectation when the reported belief is equal to the true belief. These incentives can be easily extended to a sequence of multiple independent questions; in particular, when the sequence of posed questions is independent of answers chosen.

We wish to establish similar results in the case where the *questions posed depend on the answers to past questions*. There are several scenarios in which this is necessary – for example, the question-setter can be more efficient by dynamically narrowing the scope of questions, or if a question requires follow-up questions.

1.1 Motivation

The following example shows why naively combining SPSR rewards does not work. Suppose we have 3 questions:

Q : the initial question, answered as a distribution over two options. The expert's response will be a vector of class probabilities $\boldsymbol{q} = (\boldsymbol{q}_1, \boldsymbol{q}_2)$, which we use to select between follow-up questions.
Example: *With what probability will NASDAQ do better than NYSE over the next week?*

Q_1 : a possible follow-up question, with any answer suitable for an SPSR. This question is asked with probability $\boldsymbol{q}_1$.
Example: *With what probability will AMZN do better than GOOGL over the next week?*

Q_2 : a possible follow-up question, with any answer suitable for an SPSR. This question is asked with probability $\boldsymbol{q}_2$.
Example: *With what probability will TWTR do better than BABA over the next week?*

We pose the *initial* question Q to an expert, eliciting a response of $\boldsymbol{q}$. We use this answer to determine a distribution over the possible *follow-up* questions. We can use any rule to construct this. For this example, we use the same distribution as $\boldsymbol{q}$. We sample the follow-up question Q_i from this

Final project report for course CMU 10709 "Fundamentals of Learning from the Crowd", Fall 2017. Instructor: Nihar B. Shah, TA: Ritesh Noothigattu. The formatting style file is borrowed from the NIPS conference.

distribution and pose it to the expert, who responds with $\boldsymbol{q}^{(i)}$. After we have both responses, we observe the event outcomes $a, a^{(i)}$, and pay the expert accordingly.

Without loss of generality, we use the Quadratic SPSR to score the initial and follow-up question. Consider an arbitrary respondent Bob, who is confident in his answer to Q_1 (where he expects to receive 0 reward), but not Q_2 (where he expects to receive -0.5 reward).

When answering the first question, Bob has an incentive to report a distribution skewed in favor of Q_1, since this will net him a greater *future* reward. This is an incentive to lie in the initial question, which cannot be accounted for in the traditional SPSR framework. Our project aim is to design a scoring rule that can elicit the truth even when each question posed may depend on previous answers.

2 Problem Formulation

We study simple questionnaires consisting of an *initial* question, which is always asked, and one *follow-up* question, which is selected based on the response to the initial question.

Questions Let the first and second questions be Q and Q_i respectively, where i is the index of the second question. Q_i is sampled from a finite question bank $\mathcal{Q} = \{Q_1, \dots Q_n\}$. The expert's response to $Q, Q^{(i)}$ is denoted by $\boldsymbol{q}, \boldsymbol{q}^{(i)}$ respectively. All answers are probability distributions.

Transitions The distribution of the index for the second question is given by $\Pr\{i|\boldsymbol{q}\}$. This is an arbitrary function of the respondent's previous answer $\boldsymbol{q}$. The respondent has full knowledge of this transition function. A possible transition function is the affine transition, elaborated in Appendix A.

Beliefs We denote the respondent's belief over the answers to $Q, Q^{(i)}$ as $\boldsymbol{p}, \boldsymbol{p}^{(i)}$ respectively. We wish to design a scoring rule that elicits $\boldsymbol{p}$ and $\boldsymbol{p}^{(i)}$, where i is the index of the second question asked.

Outcomes Respondents are paid after the outcome associated with Q and $Q^{(i)}$ are measured. These are denoted by a and $a^{(i)}$ respectively.

Our goal is to design a Jointly Strictly Proper Scoring Rule (J-SPSR) S, such that the respondent is encouraged to report his true beliefs $\boldsymbol{p}, \boldsymbol{p}^{(i)}$. Note that from the point of view of the respondent, the $a^{(i)}$ depends on the index of the second question i, which in turn depends on his response to the first question, $\boldsymbol{q}$. As in the example above, there may be an incentive to lie in $\boldsymbol{q}$ in order to obtain an advantage in the second question.

Definition 1. Joint Strictly Proper Scoring Rules (J-SPSR)
A J-SPSR is a multi-stage scoring rule $S(\boldsymbol{q}, \boldsymbol{q}^{(i)}, a, a^{(i)}, i)$ such that for all beliefs $\boldsymbol{p}, \boldsymbol{p}^{(i)}$, the expert is incentivised to answer the first question honestly:

$$\arg\max_{\boldsymbol{q}} \sum_{i\in[n]} \left(\Pr(i|\boldsymbol{q}) \cdot \max_{\boldsymbol{q}^{(i)}} \mathbb{E}_{a\sim p, a^{(i)}\sim p^{(i)}} \left[S(\boldsymbol{q}, \boldsymbol{q}^{(i)}, a, a^{(i)}, i) \right] \right) = \boldsymbol{p},$$

And such that the second question is also answered honestly when the first is honestly answered:

$$\arg\max_{\boldsymbol{q}^{(i)}} \mathbb{E}_{a\sim p, a^{(i)}\sim p^{(i)}} \left[S(\boldsymbol{p}, \boldsymbol{q}^{(i)}, a, a^{(i)}, i) \right] = \boldsymbol{p}^{(i)} \qquad \forall\, i \in [n] \text{ if } \Pr(i|\boldsymbol{q} = \boldsymbol{p}) > 0$$

Before continuing, we will assume the reader is familiar with the material in [1], in particular, the material in Section 2.1 and 3.1. Theorem 2 by Savage and McCarthy [1**?**] will be continually referenced throughout this report, and we refer to it as the *workhorse theorem*.

3 Additive Scoring Rules

The simplest proposal for a J-SPSR is the *additive J-SPSR* – one where S may be decomposed into the sum of 2 independent scoring rules, $S^{(0)}$ for the initial question, and $S^{(i)}$ for the follow-up question:

$$S(\boldsymbol{q}, \boldsymbol{q}^{(i)}, a, a^{(i)}, i) = S^{(0)}(\boldsymbol{q}, a) + S^{(i)}(\boldsymbol{q}^{(i)}, a^{(i)}). \tag{1}$$

We want to determine if it is possible for a SPSR to be additive. For simplicity, we assume that the distribution of follow-up questions $i \sim \boldsymbol{q}$ – that is, the probability of selecting follow-up question Q_i is the reported probability $\boldsymbol{q}_i$ (the i^{th} element of $\boldsymbol{q}$).

Theorem 1

If S is an additive J-SPSR then for all i, $S^{(i)}$ must be SPSR.

The proof is deferred to Appendix B.1.

Hence, we may assume that the reward from the second question is the *score* $G^{(i)}(\boldsymbol{p}^{(i)})$, where $G^{(i)}$ is the strictly convex function inducing $S^{(i)}$ [1]. However, the reward from answering the second question honestly may provide an incentive against answering the first question truthfully. For a fixed answer $\boldsymbol{q}$ to the initial question, different beliefs $\boldsymbol{p}^{(i)}$ for follow-up questions lead to different expected payoffs, assuming honest follow-up answers. The expected future reward from the follow-up question, as a function of $\boldsymbol{q}$ is

$$\mathbb{E}_{i\sim\boldsymbol{q}}\left[\mathbb{E}_{a^{(i)}\sim\boldsymbol{p}^{(i)}}[S^{(i)}(\boldsymbol{p}^{(i)},a^{(i)})]\right] = \mathbb{E}_{i\sim\boldsymbol{q}}\left[G^{(i)}(\boldsymbol{p}^{(i)})\right]$$

The maximum and minimum values of $S^{(i)}$ depends on the SPSR i.e., $G^{(i)}$. For convenience, let $\boldsymbol{x}$ be a vector where $\boldsymbol{x}_i = G^{(i)}(\boldsymbol{p}^{(i)})$, i.e. $\boldsymbol{x}$ contains the expected payoffs for each follow-up question. Then, the above expression simplifies to $\boldsymbol{x}^T q$.

Theorem 2

$S^{(0)}$ being an SPSR is not a sufficient condition to guarantee that an expert will answer honestly.

Proof. We provide an example in which $S^{(0)}$ is an SPSR but an expert will answer dishonestly. Consider the case from Section 1.1, where $\boldsymbol{p}$ is the expert's true belief and all of $S^{(i)}$ are quadratic scoring rules. The reported value is:

$$\arg\max_{\boldsymbol{q}} \mathbb{E}_{a,a^{(i)}}\left[S(\boldsymbol{q},\boldsymbol{q}^{(i)},a,a^{(i)},i)\right] = \arg\max_{\boldsymbol{q}} \mathbb{E}_a\left[S^{(0)}(\boldsymbol{q},a)\right] + \boldsymbol{x}^T\boldsymbol{q}$$

We know from [1] that the expectation term will be may be rewritten as a constant involving $\boldsymbol{p}$ (i.e. $G(\boldsymbol{p})$) minus the Bregman divergence between $\boldsymbol{q}$ and $\boldsymbol{p}$. For the quadratic, the divergence term is simply $||\boldsymbol{q}-\boldsymbol{p}||_2^2$. Hence,

$$\arg\max_{\boldsymbol{q}} \mathbb{E}_a\left[S^{(0)}(\boldsymbol{q},a)\right] + \boldsymbol{x}^T\boldsymbol{q} = \arg\max_{\boldsymbol{q}} -\sum_{i\in[n]}(\boldsymbol{p}_i-\boldsymbol{q}_i)^2 + (\boldsymbol{x}-\bar{\boldsymbol{x}})^T\boldsymbol{q}$$

where $\bar{\boldsymbol{x}}$ is the mean of $\boldsymbol{x}$, $\mathbf{1}\times\frac{\sum_j \boldsymbol{x}_j}{n}$, i.e., we perform normalization on $\boldsymbol{x}$. This normalization is equivalent to solving the constrained optimization problem $\boldsymbol{q}$ constrained to lie in the probability simplex. Differentiating and setting to 0 yields: $\boldsymbol{q}_i = \boldsymbol{p}_i + (\boldsymbol{x}_i - \bar{\boldsymbol{x}})/2$. Therefore, the error incurred is $|\boldsymbol{x}_i - \bar{\boldsymbol{x}}|/2$, which is non-zero unless $\boldsymbol{x}$ is equal in all entries – i.e. the expert is indifferent to all follow-up questions. □

3.1 Bounded slack

From our analysis of Theorem 2, we can suggest a mitigation strategy, assuming we are allowed to relax the strictness of our additive SPSR. This approach was inspired by work in [2].

We call an additive SPSR ξ-slack if it permits the expert to report $\boldsymbol{q}$ such that $\max(|\boldsymbol{q}-\boldsymbol{p}|) < \xi$, where $0 \le \xi \le \frac{1}{2}$ is a "slackness" bound. We observe that the incentive to lie in the first round can be reduced by offering a much larger reward for the first question than the second. In order to formalize this, we add a scaling term α to the first:

$$S(\boldsymbol{q},\boldsymbol{q}^{(i)},a,a^{(i)},i) = \alpha S^{(0)}(\boldsymbol{q},a) + S^{(i)}(\boldsymbol{q}^{(i)},a^{(i)}). \tag{2}$$

It is straightforward to apply Theorem 2 to show that the additive SPSR is $\frac{1}{2\alpha}$-slack when all scoring rules are quadratic. The key is noting that since the entries of $\boldsymbol{x}$ are bounded in the range $[-\frac{n-1}{n}, 0]$ (Appendix C.1), the maximum error is also bounded. Full derivation is in Appendix C.2.

The particular slackness bound is a result of the choice of SPSR. It is possible (though inconvenient) to derive the bound for any arbitrary SPSR that has a finite range.

4 The Multiplicative Quadratic Scoring Rule

In this section, we introduce the *Multiplicative Quadratic Scoring Rule* (MQSR), which satisfies the J-SPSR requirement precisely. (i.e. It yields the greatest expected score when the expert honestly answers both the first and second questions.)

4.1 Asymmetric Quadratic Scoring Rules

Our proposed method hinges upon the *Asymmetric Quadratic Scoring Rule* [3], a generalization of the Quadratic Scoring Rule originally invented to reward predictions based on their difficulty. We review some properties of this scoring rule in this section.

Let $G(\boldsymbol{p}) = \boldsymbol{p}^T M \boldsymbol{p}$, where M is a positive definite matrix. Then, $G(\boldsymbol{p})$ is convex in $\boldsymbol{p}$. Hence, we know it induces a SPSR. We do so explicitly using its Savage decomposition [1].

$$\begin{aligned} S(\boldsymbol{p}, i) &= G(\boldsymbol{p}) - < G'(\boldsymbol{p}), \boldsymbol{p} > + G'_i(\boldsymbol{p}) \\ &= 2(M\boldsymbol{p})_i - \boldsymbol{p}^T M \boldsymbol{p}, \end{aligned} \tag{3}$$

and S is a SPSR. In our case, we are interested in $S(\boldsymbol{p}, i)$ when M is a diagonal matrix with positive entries along the diagonal (implying M is positive definite).

Remark: We recover the commonly seen symmetric quadratic scoring rule (up to constants) by setting $M = I$.

Example 1

Let $M = \begin{bmatrix} 1000 & 0 \\ 0 & 1 \end{bmatrix}$. *Suppose the true beliefs are* $\mathbf{p}$*, and the worker reports* $\mathbf{q}$*. From the worker's perspective, his expected payoff is:*

$$\begin{aligned} Z(\boldsymbol{q}) &= \mathbb{E}_{i \sim \boldsymbol{p}}(2(M\boldsymbol{q})_i) - \boldsymbol{q}^T M \boldsymbol{q} \\ &= 2\boldsymbol{p}^T M \boldsymbol{q} - \boldsymbol{q}^T M \boldsymbol{q} \end{aligned} \tag{4}$$

This is a quadratic in q *and has a single maximum. Differentiating and setting to* 0*,*

$$\begin{aligned} &Z'(\boldsymbol{q}) = 2M\boldsymbol{p} - 2M\boldsymbol{q} = 2M(\boldsymbol{p} - \boldsymbol{q}) = 0 \\ &\boldsymbol{p} - \boldsymbol{q} = \boldsymbol{0}, \end{aligned}$$

where the inverse for M *exists by virtue of it being a positive definite. Since* $\mathbf{p} = \boldsymbol{q}$*,* S *is strictly proper. The key lesson here is that* the specific values along the diagonal do not matter *as long as they are positive.*

4.2 Eliciting Truthful Responses

In the above, we note that the asymmetric quadratic scoring rule possesses degrees of freedom in the matrix M (entries along the diagonal) while still retaining our desired properties. These degrees of freedom may be used to encode future expected payoffs – as long as these are positive, then *the worker is incentivised to respond truthfully, even if those entries are not known to the question-setter*! This remarkable property forms the basis of our proposed MQSR.

Let $S^{(i)}(\boldsymbol{q}^{(i)}, a^{(i)})$ be strictly proper scoring rules *whose ranges are non-negative*, i.e. $\forall i, S^{(i)}(\boldsymbol{q}^{(i)}, a^{(i)}) \geq 0$. Denote the expected payoff when the worker answers truthfully (i.e. the score) as $G^{(i)}(\boldsymbol{p}^{(i)})$. Note that $S^{(i)}$ only depends on the second question i, its response $\boldsymbol{q}^{(i)}$ and its observation $a^{(i)}$. The MQSR is defined as:

$$\begin{aligned} M &= I + \mathrm{diag}_i \left(S^{(i)}(\boldsymbol{q}^{(i)}, a^{(i)}) \right) \\ S(\boldsymbol{q}, \boldsymbol{q}^{(i)}, a, a^{(i)}, i) &= \mathbb{1}(i = a) \times (2M_{i,i}) - (q^T M)_i, \\ &= \begin{cases} (2 - q_i)\left(1 + S^{(i)}(\boldsymbol{q}^{(i)}, a^{(i)})\right), & a = i \\ -q_i \left(1 + S^{(i)}(\boldsymbol{q}^{(i)}, a^{(i)})\right), & a \neq i \end{cases} \end{aligned} \tag{5}$$

Eliciting Truthful Responses for the First Question

Looking carefully, our proposed scoring rule is merely the asymmetric quadratic scoring rule in disguise. To see why, fix the worker's policy for the second question and assume he always answers truthfully (this assumption will be lifted later, albeit with some effort). Because of this, M is a constant from the worker's perspective. We examine the expected payoff as a function of $\boldsymbol{q}$ (and implicitly $\boldsymbol{p}^{(i)}, a^{(i)}$) from their perspective.

$$\begin{aligned}
\mathbb{E}_{i\sim \boldsymbol{q}}\mathbb{E}_{a\sim \boldsymbol{p}}\left[\mathbb{1}(i=a)\times(2M_{i,i})-(\boldsymbol{q}^T M)_i\right] &= \sum_a\sum_i \boldsymbol{p}_a\boldsymbol{q}_i\times\left(\mathbb{1}(i=a)\times(2M_{i,i})-(\boldsymbol{q}^T M)_i\right)\\
&= \sum_i \boldsymbol{p}_i\boldsymbol{q}_i\times 2M_{i,i}-\sum_a \boldsymbol{p}_a\sum_i \boldsymbol{q}_i(\boldsymbol{q}^T M)_i\\
&= \sum_i \boldsymbol{p}_i\times 2(M\boldsymbol{q})_i-\boldsymbol{q}^T M\boldsymbol{q}\\
&= 2\boldsymbol{p}^T M\boldsymbol{q}-\boldsymbol{q}^T M\boldsymbol{q},
\end{aligned}$$

which is exactly of the expected reward (when reporting q with true beliefs p) for an Asymmetric Quadratic Scoring rule (equations (3). (4)). Thus, the expected payoff is maximized when $\boldsymbol{q}=\boldsymbol{p}$.

To elaborate, when the worker honestly answers the first question, he is paid exactly $\boldsymbol{p}^T M\boldsymbol{p}$. Since $S^{(i)}$ is non-negative, M is diagonal and strictly positive definite, by the Savage formulation, the worker is incentivised to report $\boldsymbol{q}=\boldsymbol{p}$. His expected reward is $\boldsymbol{p}^T M\boldsymbol{p}$. The beauty here is that the worker has an incentive to honestly answer the initial question because M is positive definite, which does not depend on their belief of the second question.

One possible misconception is that the asker needs to know M, which would require them to know all possible future answers. Examining equation 5 shows that the asker only needs the value for question i^{th} to reward the worker – M is merely a tool for analysis from the worker's perspective.

Removing the truthfulness assumption for the future
Now, we relax the assumption that the player's policy of answering the second question truthfully. In the previous paragraph, the assumption was the M was independent on the reported answer for the first question, i.e. $\boldsymbol{q}$. However, given i fixed, the best response to the second question, $\boldsymbol{q}^{(i)}$ will depend on $\boldsymbol{q}$. Consider the when the worker is faced with the second question, for arbitrary (possibly dishonest) $\boldsymbol{q}$ and i. To maximize expected payoff, the worker's response to the second question is

$$\arg\max_{\boldsymbol{q}^{(i)}}\sum_a\sum_{a^{(i)}}\boldsymbol{p}_a\boldsymbol{p}^{(i)}_{a^{(i)}}\left[\mathbb{1}(i=a)\times 2\left(1+S^{(i)}(\boldsymbol{q}^{(i)},a^{(i)})\right)-\boldsymbol{q}_i\left(1+S^{(i)}(\boldsymbol{q}^{(i)},a^{(i)})\right)\right] \tag{6}$$

$$=\arg\max_{\boldsymbol{q}^{(i)}}(2\boldsymbol{p}_i-\boldsymbol{q}_i)\sum_{a^{(i)}}\boldsymbol{p}^{(i)}_{a^{(i)}}\left(1+S^{(i)}(\boldsymbol{q}^{(i)},a^{(i)})\right) \tag{7}$$

$$=\begin{cases}\arg\max_{\boldsymbol{q}^{(i)}}\sum_{a^{(i)}}\boldsymbol{p}^{(i)}_{a^{(i)}}\left(1+S^{(i)}(\boldsymbol{q}^{(i)},a^{(i)})\right)=p^{(i)}, & 2\boldsymbol{p}_i>\boldsymbol{q}_i\\ \arg\min_{\boldsymbol{q}^{(i)}}\sum_{a^{(i)}}\boldsymbol{p}^{(i)}_{a^{(i)}}\left(1+S^{(i)}(\boldsymbol{q}^{(i)},a^{(i)})\right), & 2\boldsymbol{p}_i\le\boldsymbol{q}_i\end{cases}, \tag{8}$$

where in the first case, the maximum is unique due to $S^{(i)}$ being SPSR, and in the second, ties are broken arbitrarily.

The expression above is particularly interesting – this means that if the player was so dishonest in his first answer such that $2\boldsymbol{p}_i-\boldsymbol{q}_i$ is negative, and in turn causing the next question to be i, then he is better off answering the second question to *minimize* the expectation $\mathbb{E}_{q^{(i)}\sim p^{(i)}}\left[1+S^{(i)}(\boldsymbol{q}^{(i)},a^{(i)})\right]$, which is a term which we normally would maximize. That is, a respondent who has lied is incentivised to lie further. Applying the result to determine the worker's response to the first question, we obtain the following objective:

$$\arg\max_{\boldsymbol{q}}\sum_i\boldsymbol{q}_i(2\boldsymbol{p}_i-\boldsymbol{q}_i)\times\begin{cases}\max_{\boldsymbol{q}^{(i)}}\sum_{a^{(i)}}\boldsymbol{p}^{(i)}_{a^{(i)}}\left(1+S^{(i)}(\boldsymbol{q}^{(i)},a^{(i)})\right), & 2\boldsymbol{p}_i>\boldsymbol{q}_i\\ \min_{\boldsymbol{q}^{(i)}}\sum_{a^{(i)}}\boldsymbol{p}^{(i)}_{a^{(i)}}\left(1+S^{(i)}(\boldsymbol{q}^{(i)},a^{(i)})\right), & 2\boldsymbol{p}_i\le\boldsymbol{q}_i\end{cases} \tag{9}$$

Theorem 3
The expression in (9) *achieves a unique global maximum at* $\boldsymbol{q}=\boldsymbol{p}$.

Proof. Given $\boldsymbol{p}$, the set of vectors for $\boldsymbol{q}$ in the probability simplex may be partitioned into no more than 2^n (potentially empty) sets. Each of these sets is associated with a quadratic, based on which inequalities are satisfied in Equation (9). Each of these quadratics is maximized at $\boldsymbol{q}=\boldsymbol{p}$.

Furthermore, their values are equal where the cases in Equation (9) hold with equality (i.e. where $2\boldsymbol{p}_i = \boldsymbol{q}_i$ for some i), Hence the expression in Equation (9) is continuous.

Consider a ray being cast from $\boldsymbol{p}$ in any arbitrary direction $\boldsymbol{d}$, while remaining on the probability simplex, $\boldsymbol{q} = \boldsymbol{p} + t\boldsymbol{d}$. As the ray extends, it decreases strictly (due to strict concavity of Expression 9). At some point, for some i, $2\boldsymbol{p}_i = \boldsymbol{q}_i$. This causes the ray to transition to a new quadratic regime, with different shapes. However, the value of the function evaluated at $\boldsymbol{q}$ continues to decrease in this new regime, since the global maximum for the new quadratic remains $\boldsymbol{p}$. Moreover, this transition happens a finite number of times before exiting the probability simplex. Thus, Expression (9) is monotonically decreasing in every direction starting from $\boldsymbol{p}$, implying that it is maximized at $\boldsymbol{p}$. □

We numerically simulated this process to elicit expected payoff as a function of $\boldsymbol{q}$. The results are in Appendix D.

Eliciting Truthful Responses for the Second Question
This follows directly from earlier results. Substituting $\boldsymbol{p}_i = \boldsymbol{q}_i$ into Equation (8) gives us

$$\arg\max_{\boldsymbol{q}^{(i)}} \boldsymbol{p}_i \sum_{a^{(i)}} \boldsymbol{p}^{(i)}_{a^{(i)}} \left(1 + S^{(i)}(\boldsymbol{q}^{(i)}, a^{(i)})\right) = \arg\max_{\boldsymbol{q}^{(i)}} \sum_{a^{(i)}} \boldsymbol{p}^{(i)}_{a^{(i)}} \left(1 + S^{(i)}(\boldsymbol{q}^{(i)}, a^{(i)})\right) \tag{10}$$

which from the workhorse theorem is maximized at $\boldsymbol{q}^{(i)} = \boldsymbol{p}^{(i)}$.

5 A General Framework

At this stage, three natural questions arise: First, can we construct a J-SPSR which holds with arbitrary, nonlinear transitions? Second, can our method be generalized for other scoring rules other than the Quadratic? Third, can we extend our scheme to allow "nesting" follow-up questions?

We begin by tackling the first two questions:

5.1 General Scoring Rules with General Transitions

In this section, we will allow for arbitrary transition functions; the probability of going to question i after answering q is given by $\Pr(i|q)$, which we will assume is non-zero. (One such class of transition functions is the affine, which we discuss in Appendix A) More importantly, we show that it is possible to construct new J-SPSRs by *modulating* SPSR's with the quadratic scoring rule. (For ease of understanding, we present the same derivation with extra explanation on intermediate steps in Appendix E.)

Let $G(\boldsymbol{q}) = \boldsymbol{q}^T M \boldsymbol{q}$ for some positive definite diagonal M. Let $H(\boldsymbol{q})$ be any non-negative, non-decreasing, convex function defined over the probability simplex, such that $G(\boldsymbol{q})H(\boldsymbol{q})$ remains convex and induces a SPSR. For simplicity, we assume that H is at least once-differentiable. Denote this product by $\hat{G}(\boldsymbol{q})$, and similarly, $\hat{S}(\boldsymbol{q}, a)$ as the scoring rule induced by it. We have:

$$\hat{G}'(\boldsymbol{q}) = G'(\boldsymbol{q})H(\boldsymbol{q}) + H'(\boldsymbol{q})G(\boldsymbol{q}) = 2M\boldsymbol{q}H(\boldsymbol{q}) + H'(\boldsymbol{q})\boldsymbol{q}^T M\boldsymbol{q}$$
$$\hat{S}(\boldsymbol{q}, i) = -\boldsymbol{q}^T M\boldsymbol{q}H(\boldsymbol{q}) - \boldsymbol{q}^T M\boldsymbol{q}(H'(\boldsymbol{q}))^T\boldsymbol{q} + 2\boldsymbol{q}_i M_{ii} H(\boldsymbol{q}) + H'_i(\boldsymbol{q})\boldsymbol{q}^T M\boldsymbol{q}$$

Setting $H = 1$, yields the quadratic scoring rule. Applying the workhorse theorem,

$$\arg\max_{\boldsymbol{q}} \sum_i \boldsymbol{p}_i \hat{S}(\boldsymbol{q}, i) = \boldsymbol{p}$$

with uniqueness. As before, let $\{S^{(i)}\}$ be a collection of positive SPSR's we wish to apply to the second question. Re-defining $M = \mathrm{diag}_i(S^{(i)}(q^{(i)}, a^{(i)}))$, our proposed joint scoring rule is

$$S(\boldsymbol{q}, \boldsymbol{q}^{(i)}, a, a^{(i)} i) = \frac{\boldsymbol{q}_i}{\Pr(i|\boldsymbol{q})} \left(\mathbb{1}(i = a) 2M_{ii}H(\boldsymbol{q}) + \boldsymbol{q}_i M_{ii} H'_a(\boldsymbol{q}) - \boldsymbol{q}_i M_{ii} H(\boldsymbol{q}) - \boldsymbol{q}_i M_{ii} (H'(\boldsymbol{q}))^T\boldsymbol{q}\right) \tag{11}$$

Suppose the worker reported $\boldsymbol{q}$ and is asked question i. His expected payoff as a function of $\boldsymbol{q}^{(i)}$ is:

$$\frac{\boldsymbol{q}_i}{\Pr(i|\boldsymbol{q})}\sum_a\sum_{a^{(i)}}\boldsymbol{p}_a\boldsymbol{q}^{(i)}_{a(i)}\left[\mathbb{1}(i=a)(2M_{ii}H(\boldsymbol{q}))+\boldsymbol{q}_iM_{ii}H'_a(\boldsymbol{q})-\boldsymbol{q}_iM_{ii}H(\boldsymbol{q})-\boldsymbol{q}_iM_{ii}(H'(\boldsymbol{q}))^T\boldsymbol{q}\right]$$

$$=\frac{\boldsymbol{q}_i}{\Pr(i|\boldsymbol{q})}\left(\sum_{a^{(i)}}\boldsymbol{q}^{(i)}_{a^{(i)}}S^{(i)}(\boldsymbol{q}^{(i)},a^{(i)})\right)\left(2\boldsymbol{p}_iH(\boldsymbol{q})+\boldsymbol{q}_i\sum_a\boldsymbol{p}_aH'_a(\boldsymbol{q})-\boldsymbol{q}_iH(\boldsymbol{q})-\boldsymbol{q}_i(H'(\boldsymbol{q}))^T\boldsymbol{q}\right).$$

The worker seeks to maximize this quantity. If the factor on the right is positive, then the left factor (the usual SPSR) will be maximized. If it is negative, then it will be minimized. For brevity:

$$R(\boldsymbol{q},i)=\left(2\boldsymbol{p}_iH(\boldsymbol{q})+\boldsymbol{q}_i\sum_a\boldsymbol{p}_aH'_a(\boldsymbol{q})-\boldsymbol{q}_iH(q)-\boldsymbol{q}_i(H'(\boldsymbol{q}))^T\boldsymbol{q}\right)\tag{12}$$

$$K(\boldsymbol{q},i)=\begin{cases}\arg\max_a^{(i)}\left(\sum_{a^{(i)}}\boldsymbol{q}^{(i)}_{a^{(i)}}S^{(i)}(\boldsymbol{q}^{(i)},a^{(i)})\right), & R(\boldsymbol{q},i)>0\\ \arg\min_a^{(i)}\left(\sum_{a^{(i)}}\boldsymbol{q}^{(i)}_{a^{(i)}}S^{(i)}(\boldsymbol{q}^{(i)},a^{(i)})\right), & R(\boldsymbol{q},i)\le 0\end{cases}\tag{13}$$

As a sanity check, observe that if we set $H=1$, we obtain exactly the case described in Equation (8), where the 'cutoffs' are given by $2\boldsymbol{p}_i=\boldsymbol{q}_i$. Furthermore, if we assume $\boldsymbol{p}=\boldsymbol{q}$, $R(p,i)$ simplifies to $p_iH(p)$, which we have assumed to be positive during the construction of H; this result assures us that if the first question is answered truthfully, the second will be answered truthfully as well (the positive case in Equation (13)). For brevity, define the diagonal matrix

$$J(q)=\operatorname{diag}_i(K(\boldsymbol{q},i)),$$

and omit $\boldsymbol{q}$ for clarity. Now, we take argmax over q of expected payoff over all possible answers the worker gives for the first question, weighted by the probability of future transition to question i.

$$\arg\max_q\sum_i\Pr(i|q)\frac{q_i}{\Pr(i|q)}K(q,i)R(q,i)=\arg\max_q\sum_i q_iK(q,i)R(q,i)$$
$$=\arg\max_q\left(2p^TJqH(q)+q^TJq<H'(q),p>-q^TJqH(q)-q^TJq(H'(q))^Tq\right)$$

This is precisely the optimization equation from the scoring rule induced by $G(\boldsymbol{q})H(\boldsymbol{q})$ (see Equation (16)), and if J is fixed and positive definite, is minimized at $\boldsymbol{q}=\boldsymbol{p}$. Remember, however, that J derived from K, which is a function of $\boldsymbol{q}$. Assuming that for all i, the number of times R flips signs as we vary $\boldsymbol{q}$ along any direction is bounded by some constant c. Then the probability simplex may be partitioned by a finite number of regions, where within each region, J is constant, its value only depending on the *sign* of $R(\boldsymbol{q},i)$. For each region, we the function we want to minimize is given by a different SPSR, but each (when extended to the full probability simplex), achieves a maximum at $\boldsymbol{q}=\boldsymbol{p}$). The full piecewise function is also continuous.

As before, we argue that starting at $\boldsymbol{q}=\boldsymbol{p}$ and moving in an arbitrary direction along the probability simplex, the function will be strictly decreasing. Regardless of which partition we are in, consider the difference between the current value (with J chosen to be that which is associated with that partition) and the value at $\boldsymbol{q}=\boldsymbol{p}$. This is given by the Bregman divergence induced by $G(\boldsymbol{q})H(\boldsymbol{q})$ This value is always increasing as we move $\boldsymbol{q}$ away from $\boldsymbol{p}$, since this is a divergence (the rate of decrease may depend on the region). If we move from one region to another, there is no discontinuity, and the function remains strictly decreasing. Thus, the true value of $\boldsymbol{q}=\boldsymbol{p}$ is reported.

5.2 Generalization to Sequences of Questions

We can apply this technique recursively to obtain a J-SPSR for sequences of questions with length greater than 2. This may be done by replacing the diagonal matrix M on the first stage by a diagonal matrix containing *expected payoffs* from future questions assuming the second question is i (Equation (5)). For example, when the depth of the questionnaire is 3,

$$S(q,q^{(i)},q^{(i,j)},a,a^{(i)},a^{(i,j)},i,j)=\begin{cases}(2-q_i)\left(C+S(q^{(i)},q^{(i,j)},a^{(i)},a^{(i,j)},j)\right), & a=i\\ -q_i\left(C+S(q^{(i)},q^{(i,j)},a^{(i)},a^{(i,j)},j)\right), & a\ne i\end{cases}$$

where C is a constant sufficiently large to ensure that the function is positive, and hence M being positive definite. The derivation is straightforward and is hence omitted. The derivation proceeds as before, except that we minimizing or maximizing over the future reward for question 2, given by the nested J-SPSR of depth 2 (Equation (8)).

6 Discussion

As far as we know, MQSR works because of a specific property of the quadratic score which we have yet to identify. We have yet to generalize this to any other scoring rules, only to those constructed to be modulated by the quadratic.

MQSR is most easily compared to the bounded ξ-slackness (BS) approach, which is only an approximate solution to the J-SPSR problem. It has some advantages over the MQSR in practice:

1. Simplicity – MQSR is difficult to explain to a layman. This means that despite having nice theoretical properties, workers may not respond as we expect – especially when the depth, d is large.
2. It is 'subgame perfect'. At every point in the questionnaire, the worker has incentive to respond honestly. Under some circumstances, MQSR lacks this property – our proof that the worker is truthful in the second round relies on the premise that $\boldsymbol{q}$ is close enough to $\boldsymbol{p}$, i.e. the worker is already truthful in the first. MSQR incentivises honesty as long as $2\boldsymbol{p}_i - \boldsymbol{q}_i > 0$ (Equation (8)),which is an ample margin that this is unlikely to be an issue in practice.

However, MQSR has many strengths and desirable properties that BS lacks:

1. The bounds in BS are, in practice, extremely loose. In the earlier example a bound of 5% requires setting $\alpha \approx 100$. This means that the reward for the first question is a hundred times as much as the second.
2. BS suffers from exponentially vanishing rewards as the questionnaire depth d becomes large. If the smallest value we could pay was \$0.01, a blowup of 100 at each stage means that even a shallow questionnaire of depth 4 would cost us \$10,000 to administer once. This blowup renders BS unsuitable for any depth greater than 2.

   On the other hand, we hypothesize that the multi-stage version of MQSR has payoffs that scale linearly with depth, which makes it an attractive option!
3. The obvious advantage of MQSR is that it is exact. This makes it appealing from an academic perspective, as well as possible high-volume applications (such as in advertisement networks) where even small error margins can yield huge changes in revenue over time.

6.1 Future Work

Perhaps the greatest open question is the existence of other J-SPSRs. We have yet to identify the exact property of the quadratic scoring rule that allows us to construct the MSQR, and identifying this would be of considerable theoretical interest.

While constructing the MQSR, we discovered a technique of parameterizing SPSRs to yield this additional property. Our investigation was stymied by an inability to complete the final part of the proof. We present our construction, understanding, and the gap in our understanding in Appendix F

An important question about the MQSR is how it scales with depth. Based on our construction in Section 5.2, we hypothesize that its growth can be linear, but we have yet to prove that. This is an important result that, if proved, would make this tremendously efficient for questionnaires with deeply nested questions.

7 Conclusion

There are many fields in which having questionnaires change in response to answers is advantageous. Areas like crowdsourcing, imitation learning, adaptive quizzes or testing, and many more can apply the J-SPSR criterion formalized in this paper to construct reward schemes that incentivise honesty and improve the quality of their data.

To tackle this novel problem, we proposed an approximation based on ξ-slackness, an exact framework based on MQSRs (with several extensions), and provided comparative analysis between these options. There are clear areas for future investigation and promising open problems in the field.

References

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

[2] Nihar Shah and Dengyong Zhou. No oops, you won't do it again: Mechanisms for self-correction in crowdsourcing. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 1–10, New York, New York, USA, 20–22 Jun 2016. PMLR.

[3] Robert L. Winkler. Evaluating probabilities: Asymmetric scoring rules. *Management Science*, 40(11):1395–1405, 1994.

A Appendix: Modeling transitions

One simple transition model which can yield rich behavior while guaranteeing strictly positive transition property $\Pr(i|\boldsymbol{q})$ is the affine transform, which expresses the probability of selecting question q_k: $\Pr(Q^{(2)} = q_k|\boldsymbol{q}) = \frac{\boldsymbol{q}_k + \boldsymbol{\alpha}_k}{1+\sum_j \boldsymbol{\alpha}_j}$.

This may be interpreted as the weighted sum of the $\boldsymbol{q}$, regularized by $\boldsymbol{\alpha}$. The $\boldsymbol{\alpha}_k$'s may used to model the prior of the examiner, or the collective expected response from other respondents. Note that it is not necessary for $\boldsymbol{\alpha}$ to sum to 1: this allows us flexibility in weighing the distribution of question 2 as a Linear function of $\boldsymbol{q}$.

B Appendix: Proofs for Additive J-SPSR

B.1 Proof of Theorem 1 – Necessity of follow-up question being a SPSR

Proof. For clarity, we make the very mild technical condition that it is possible to transition to any $i \in [n]$ with positive probability – it may be easily verified that this assumption may be dropped.

Suppose S is an additive SPSR and the worker has answered the first question truthfully with p and transitioned to question i. At this stage, the reward from $S^{(0)}$ is beyond the worker's control, his response $\boldsymbol{q}^{(i)}$ may only affect the second component. It is easy to see that since $S^{(i)}$ only depends on $\boldsymbol{q}^{(i)}, a^{(i)}$, the worker will be incentivised to answer the second question truthfully. □

C Appendix: ξ-Slackness of the Quadratic Scoring Rule

C.1 Deriving the Range

The possible range of expected payoffs (assuming honesty) is simply the range of $G(\boldsymbol{q})$. For a quadratic, $G(\boldsymbol{q}) = ||\boldsymbol{q}||_2^2 - 1$. This achieves its maximum at the vertices, and has a minimum when it is uniform , giving a value of $-\frac{n-1}{n}$; which when $n = 2$, is $-\frac{1}{2}$.

C.2 Deriving the slackness bound:

From Theorem 2, we know that

$$\boldsymbol{q}_i = \boldsymbol{p}_i + (\boldsymbol{x}_i - \bar{\boldsymbol{x}})/2.$$

Observe that scaling the payoff of the first question by α is equivalent to scaling the payoff for the follow-up question by $\frac{1}{\alpha}$. After this scaling, the range of scores for the follow up questions become $[-\frac{n-1}{n\alpha}, 0]$. The largest possible deviation (coordinate-wise, i.e. max-norm) is when $|\boldsymbol{x}_i - \bar{\boldsymbol{x}}|$ is maximized. Observe that this is achieved when we set one entry in $\boldsymbol{x}$ to be the as low as possible, and every other entry to be as high as possible. In the case of scaled quadratics, this could be done by $\boldsymbol{x} = [-\frac{n-1}{n\alpha}, 0, ..., 0]^T$. A bound for the deviation is thus $\frac{1}{2\alpha}$ (which becomes closer to equality when $n \to \infty$).

D Appendix: Numerical Simulations of the Probability Simplex

We conducted some numerical simulations to test our scoring rule and illustrated expected payoffs. We simulated a small example with $n = 2$. The true belief is $\boldsymbol{p} = (0.75, 0.25)^T$, $p^{(1)} = (0.5, 0.5)$, $p^{(2)} = (0.1, 0.9)$. The spherical scoring rule was used for the follow-up question. Our implementation correctly identifies that the optimal policy is to report $\boldsymbol{q} = \boldsymbol{p}$.

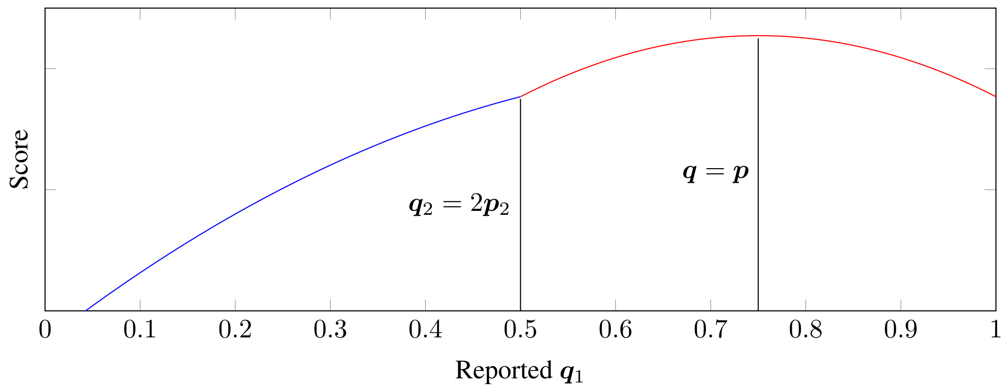

We move from one region to another in the probability simplex when $\boldsymbol{q}_i = 2\boldsymbol{p}_i$ for any i. In this chart, we observe the change in region at $\boldsymbol{q}_2 = 2\boldsymbol{p}_2 = 0.5$, exactly as predicted.

E Appendix: Full Derivation of General Scoring Rules with General Transitions

This appendix provides intermediate steps in the proof of the generalized version of the MQSR.

We begin by deriving $\hat{G}(\boldsymbol{q})$ and the induced scoring rule $\hat{S}(\boldsymbol{q}, a)$:

$$\begin{aligned}\hat{G}'(\boldsymbol{q}) &= G'(\boldsymbol{q})H(\boldsymbol{q}) + H'(\boldsymbol{q})G(\boldsymbol{q}) \\ &= 2M\boldsymbol{q}H(\boldsymbol{q}) + H'(\boldsymbol{q})\boldsymbol{q}^T M\boldsymbol{q} \qquad (14)\\ \hat{S}(\boldsymbol{q}, i) &= \boldsymbol{q}^T M\boldsymbol{q}H(\boldsymbol{q}) - < 2M\boldsymbol{q}H(\boldsymbol{q}) + H'(\boldsymbol{q})\boldsymbol{q}^T M\boldsymbol{q}, \boldsymbol{q} > + \left[2M\boldsymbol{q}H(\boldsymbol{q}) + H'(\boldsymbol{q})\boldsymbol{q}^T M\boldsymbol{q}\right]_i \\ &= -\boldsymbol{q}^T M\boldsymbol{q}H(\boldsymbol{q}) - \boldsymbol{q}^T M\boldsymbol{q}(H'(\boldsymbol{q}))^T\boldsymbol{q} + 2\boldsymbol{q}_i M_{ii}H(\boldsymbol{q}) + H'_i(\boldsymbol{q})\boldsymbol{q}^T M\boldsymbol{q} \qquad (15)\end{aligned}$$

Applying the workhorse theorem,

$$\begin{aligned}&\arg\max_q \sum_i p_i \hat{S}(\boldsymbol{q}, i) \\ &= \arg\max_q \sum_i \left(2\boldsymbol{p}_i\boldsymbol{q}_i M_{ii}H(\boldsymbol{q}) + \boldsymbol{q}^T M\boldsymbol{q}H'_i(\boldsymbol{q})p_i\right) - \boldsymbol{q}^T M\boldsymbol{q}H(\boldsymbol{q}) - \boldsymbol{q}^T M\boldsymbol{q}(H'(\boldsymbol{q}))^T\boldsymbol{q} \\ &= \arg\max_q \left(2\boldsymbol{p}^T M\boldsymbol{q}H(\boldsymbol{q}) + \boldsymbol{q}^T M\boldsymbol{q} < H'_i(\boldsymbol{q}), \boldsymbol{p} > -\boldsymbol{q}^T M\boldsymbol{q}H(\boldsymbol{q}) - \boldsymbol{q}^T M\boldsymbol{q}(H'(\boldsymbol{q}))^T\boldsymbol{q}\right) \qquad (16) \\ &= \boldsymbol{p}\end{aligned}$$

with uniqueness. As before, let $\{S^{(i)}\}$ be a collection of positive SPSR's we wish to apply to the second question. Re-defining $M = \mathrm{diag}_i(S^{(i)}(q^{(i)}, a^{(i)}))$, our proposed joint scoring rule is

$$S(q, q^{(i)}, a, a^{(i)}i) = \frac{q_i}{\Pr(i|q)}\left(\mathbb{1}(i = a)2M_{ii}H(q) + q_i M_{ii}H'_a(q) - q_i M_{ii}H(q) - q_i M_{ii}(H'(q))^T q\right) \qquad (17)$$

Suppose the worker has reported q and is now posed with question i. His expected payoff as a function of $q^{(i)}$ is

$$\frac{q_i}{\Pr(i|q)}\sum_a\sum_{a^{(i)}} p_a q^{(i)}_{a(i)}\left[\mathbb{1}(i=a)(2M_{ii}H(q))+q_iM_{ii}H'_a(q)-q_iM_{ii}H(q)-q_iM_{ii}(H'(q))^Tq\right]$$

$$=\frac{q_i}{\Pr(i|q)}\left(p_i\sum_{a^{(i)}}q^{(i)}_{a^{(i)}}2S^{(i)}(q^{(i)},a^{(i)})H(q)+\left(\sum_{a^{(i)}}q^{(i)}_{a^{(i)}}S^{(i)}(q^{(i)},a^{(i)})\right)\left(\sum_a p_aH'_a(q)q_i\right)\right.$$
$$\left.-\sum_{a^{(i)}}q^{(i)}_{a^{(i)}}\left(q_iS^{(i)}(q^{(i)},a^{(i)})H(q)+q_iS^{(i)}(q^{(i)},a^{(i)})(H'(q))^Tq\right)\right)$$
$$=\frac{q_i}{\Pr(i|q)}\left(\sum_{a^{(i)}}q^{(i)}_{a^{(i)}}S^{(i)}(q^{(i)},a^{(i)})\right)\left(2p_iH(q)+q_i\sum_a p_aH'_a(q)-q_iH(q)-q_i(H'(q))^Tq\right).$$

The worker seeks to maximize this quantity. If the factor on the right is positive, then the left factor (the usual SPSR) will be maximized. If it is negative, then it will be *minimized*. For brevity, we define:

$$R(q,i)=\left(2p_iH(q)+q_i\sum_a p_aH'_a(q)-q_iH(q)-q_i(H'(q))^Tq\right)\tag{18}$$

$$K(q,i)=\begin{cases}\arg\max_a^{(i)}\left(\sum_{a^{(i)}}q^{(i)}_{a^{(i)}}S^{(i)}(q^{(i)},a^{(i)})\right), & R(q,i)>0\\ \arg\min_a^{(i)}\left(\sum_{a^{(i)}}q^{(i)}_{a^{(i)}}S^{(i)}(q^{(i)},a^{(i)})\right), & R(q,i)\le 0\end{cases}\tag{19}$$

As a sanity check, observe that if we set $H=1$, we obtain exactly the case described in Equation (8), where the 'cutoffs' are given by $2\boldsymbol{p}_i=\boldsymbol{q}_i$. Furthermore, if we assume $\boldsymbol{p}=\boldsymbol{q}$, $R(p,i)$ simplifies to $p_iH(p)$, which we have assumed to be positive during the construction of H; this result assures us that if the first question is answered truthfully, the second will be answered truthfully as well (the positive case in Equation (19)). For brevity, define the diagonal matrix

$$J(q)=\mathrm{diag}_i(K(q,i)),$$

and we will omit the parameter q for clarity. Now, we take argmax over q of expected payoff over all possible answers the worker gives for the first question, weighted by the probability of future transition to question i.

$$\arg\max_q\sum_i\Pr(i|q)\frac{q_i}{\Pr(i|q)}K(q,i)R(q,i)$$
$$=\arg\max_q\sum_i q_iK(q,i)R(q,i)$$
$$=\arg\max_q\sum_i K(q,i)q_i\left(2p_iH(q)+q_i\sum_a p_aH'_a(q)-q_iH(q)-q_i(H'(q))^Tq\right)$$
$$=\arg\max_q\left(2p^TJqH(q)+\sum_i q_i^2J_{ii}\sum_a p_aH'_a(q)-q^TJqH(q)-q^TJq(H'(q))^Tq\right)$$
$$=\arg\max_q\left(2p^TJqH(q)+q^TJq<H'(q),p>-q^TJqH(q)-q^TJq(H'(q))^Tq\right)$$

This is precisely the optimization equation from the scoring rule induced by $G(\boldsymbol{q})H(\boldsymbol{q})$ (see Equation (16)), and if J is fixed and positive definite, is minimized at $\boldsymbol{q}=\boldsymbol{p}$.

We complete this derivation with the simplex traversal argument from Section 5.1

F Appendix: A Failed Approach

In this appendix, we discuss the Monotically Increasing Parameterized Scoring Rule (MIPSR), an initial attempt at constructing a J-SPSR that is promising but requires deeper mathematical insight. We provide the intuition behind this and clearly identify the remaining gap in the proof.

F.1 Monotically Increasing Parameterized Scoring Rules

Consider a parameterized variant of SPSRs, $\hat{S}(\boldsymbol{p}, a, \boldsymbol{\phi})$, where $\boldsymbol{\phi}$ is a vector of size n (equal to $\boldsymbol{p}$). We require $\hat{S}$ to be SPSR for all values of $\boldsymbol{\phi} \in \Phi$, where Φ is the set of permissible values of $\boldsymbol{\theta}$. Furthermore, we are interested where the payoffs (assuming honest responses) are strictly increasing as $\boldsymbol{\phi}$ increases.

$$\arg\max_{\boldsymbol{q}} \sum_i p_i \hat{S}(\boldsymbol{q}, a, \boldsymbol{\phi}) = \boldsymbol{p}, \qquad \forall \boldsymbol{\phi} \in \Phi, \boldsymbol{p} \tag{20}$$

$$\max_{\boldsymbol{q}} \sum_i p_i \hat{S}(\boldsymbol{q}, a, \boldsymbol{\phi}^{(1)}) < \max_{\boldsymbol{q}} \sum_i p_i \hat{S}(\boldsymbol{q}, a, \boldsymbol{\phi}^{(2)}), \qquad \forall \boldsymbol{\phi}^{(1)}, \boldsymbol{\phi}^{(2)} \in \Phi, \boldsymbol{\phi}^{(1)} \prec \boldsymbol{\phi}^{(2)}, \boldsymbol{p}, \tag{21}$$

and the arg-max is unique. Similarly, we define $\hat{G}(\boldsymbol{p}, \boldsymbol{\phi})$ as the score function inducing $\hat{S}$. As before, $\hat{G}$ is convex in $\boldsymbol{p}$ for all $\boldsymbol{\phi}$. Lastly, we will require that there exists a baseline $\tilde{\boldsymbol{\phi}}$ such that $\hat{S}(\boldsymbol{p}, a, \tilde{\boldsymbol{\phi}} + \lambda e_j)$ is known for all $\lambda \geq 0$, where $\boldsymbol{e}_j$ denotes the j-th elementary vector.

Example. Trivial additive parameterizations.
Let S be any SPSR. Then $\hat{S}(\boldsymbol{q}, a, \boldsymbol{\phi}) = S(\boldsymbol{q}, a) + \sum_j \boldsymbol{\phi}_j$ is a monotically increasing in $\boldsymbol{\phi}$. Setting $\tilde{\boldsymbol{\phi}} = \boldsymbol{0}$ gives a valid baseline.

Example. Trivial multiplicative parameterizations.
Let S be any strictly positive SPSR. Then $\hat{S}(\boldsymbol{q}, a, \boldsymbol{\phi}) = S(\boldsymbol{q}, a) \times (\sum_j \boldsymbol{\phi}_j)$ is a monotonically increasing in $\boldsymbol{\phi}$. Instead of summation, we could also use $\prod_j \boldsymbol{\phi}_j$ if $\boldsymbol{\phi} \in \mathbb{R}_+^n$. In both cases, setting $\tilde{\boldsymbol{\phi}} = \boldsymbol{1}$ gives a valid baseline.

Example. Non-trivial parameterizations of the Quadratic scoring rule.
Consider the Asymmetric Quadratic Scoring Rule. Then setting $M = I + \text{diag}(\boldsymbol{\phi})$ gives us a monotonically increasing rule for $\boldsymbol{\phi} \in \mathbb{R}_+^n$. Setting $\tilde{\boldsymbol{\phi}} = \boldsymbol{0}$ gives a valid baseline.

Example. Non-trivial parameterizations of the Spherical scoring rule.
Consider the generalized spherical scoring rule with the pseudospherical score. The expected payoff under truthful reporting is the α-norm $||\boldsymbol{p}||_\alpha$, $\alpha > 1$. Setting $\alpha = 1 + \frac{1}{\sum_j \boldsymbol{\phi}_j}$ gives us a monotonically increasing rule for $\boldsymbol{\phi}_j \succ \boldsymbol{0}$, since in $\mathbb{R}^n$, $||\boldsymbol{p}||_\alpha$ decreases as α increases. Setting $\tilde{\boldsymbol{\phi}} = \boldsymbol{0}$ gives a valid baseline.

F.2 Parameterized Scoring Rules with General Transitions

Let $\hat{S}(\boldsymbol{p}, a, \boldsymbol{\phi})$ and $\{S^{(i)}(\boldsymbol{p}^{(i)}, a^{(i)} | i \in [n]\}$ be the desired parameterized rule and set of rules for the first and second questions respectively, and $\hat{G}$, $\{G^{(i)} | i \in [n]\}$ be the score functions inducing them. Our 2-stage scoring rule is

$$S(\boldsymbol{q}, \boldsymbol{q}^{(i)}, a, a^{(i)}, i) = \mathbb{1}(a = i) \frac{\hat{G}_i(\boldsymbol{q}, \boldsymbol{\phi})}{\Pr(i|\boldsymbol{q})} + \hat{G}(\boldsymbol{q}, \boldsymbol{\phi}) - < \hat{G}'(\boldsymbol{q}, \boldsymbol{\phi}), \boldsymbol{q} > \tag{22}$$

$$\boldsymbol{\phi} = \tilde{\boldsymbol{\phi}} + S^{(i)}(\boldsymbol{q}^{(i)}, a^{(i)}) \boldsymbol{e}_i. \tag{23}$$

Note that the derivatives are partial, taken with respect to $\boldsymbol{p}$, and $\Pr(i|\boldsymbol{q})$ is any general transition function – previously $\Pr(i|\boldsymbol{q}) = q_i$. We compute the expected reward from the worker's perspective as a function of $\boldsymbol{q}$.

$$\begin{aligned} \mathbb{E}_{i \sim q} \mathbb{E}_{a \sim p} \left[S(\boldsymbol{q}, \boldsymbol{q}^{(i)}, a, a^{(i)}, i) \right] &= \sum_a \sum_i p_a \Pr(i|\boldsymbol{q}) \left(\mathbb{1}(a = i) \frac{\hat{G}_i(\boldsymbol{q}, \boldsymbol{\phi})}{\Pr(i|\boldsymbol{q})} + \hat{G}(\boldsymbol{q}, \boldsymbol{\phi}) - < \hat{G}'(\boldsymbol{q}, \boldsymbol{\phi}), \boldsymbol{q} > \right) \\ &= \sum_i p_i \left(\hat{G}_i(\boldsymbol{q}, \boldsymbol{\phi}) + \hat{G}(\boldsymbol{q}, \boldsymbol{\phi}) - < \hat{G}'(\boldsymbol{q}, \boldsymbol{\phi}), \boldsymbol{q} > \right). \end{aligned}$$

This is precisely the expected payoff with the parameterized scoring rule $\hat{S}(\boldsymbol{q}, i, \boldsymbol{\phi})$. Although $\boldsymbol{\phi}$ is unknown to the question setter, this expression achieves its maximum (regardless of $\boldsymbol{\phi}$) when $\boldsymbol{q} = \boldsymbol{p}$. However, we are not ready to proceed, since $\boldsymbol{\phi}$ can depend on $\boldsymbol{q}$.

F.3 The final piece of the puzzle

Since $\boldsymbol{\phi}$ affects the expected payoff in very non-linear ways, the expected payoff given $\boldsymbol{q}$ will vary smoothly as $\boldsymbol{q}$ changes; this change is not just due to G, but ϕ changing (due to the second question's response). We are unable to show the piecewise nature of the latter effect, unlike in MQSR, where the parameter was piecewise constant in $\boldsymbol{q}$. If the first question is answered truthfully, then there is incentive to answer truthfully for the second question. If the first question is answered truthfully, then there is an incentive to answer the second question truthfully. However, we are unable to prove if that the worker cannot obtain a better payoff by answering *both* questions dishonestly. It may be possible to perform further analysis by examining subgradients of the complicated function (which contains terms taking $\arg\max$ over $\boldsymbol{\phi}$), but it was beyond us.