

Incentive Compatible Surveys via Posterior Probabilities

Jakša Cvitanić [†], Dražen Prelec [‡], Sonja Radas [§] and Hrvoje Šikić [¶]

April 1, 2019

Abstract

We consider the problem of eliciting truthful responses to a survey question, when the respondents share a common prior about which the survey planner is agnostic. The planner would therefore like to have a universal mechanism, which would induce honest answers for all possible priors. If the planner also requires a locality condition that ensures that the mechanism payoffs are

^{*}We are grateful to George Georgiadis, Matthew Elliott, Katrina Ligett, Patrick Ray, to conference audiences at the Southwest Economic Theory Conference 2014, 13th Conference on Research on Economic Theory & Econometrics, the 2017 Bayesian Crowd Conference, and to the anonymous referees for useful comments and suggestions. A previous version of this article circulated under the title "Mechanism design for an agnostic planner: universal mechanisms, logarithmic equilibrium payoffs and implementation".

[†]Division of the Humanities and Social Sciences, Caltech. E-mail: cvitanic@hss.caltech.edu. Research supported in part by NSF grant DMS-1613170.

[‡]MIT, Sloan School of Management, Department of Economics, Department of Brain and Cognitive Sciences. E-mail: dprelec@mit.edu. Supported by Intelligence Advanced Research Projects Activity (IARPA) via the Department of Interior National Business Center contract number D11PC20058. Disclaimer: The views and conclusions expressed herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DoI/NBC, or the U.S. Government.

[§]The Institute of Economics, Zagreb, and MIT, Sloan School of Management. E-mail: sradas@eizg.hr; sradas@mit.edu. This research was supported by a Marie Curie International Outgoing Fellowship within the 7th European Community Framework Programme, PIOF-GA-2013-622868 - BayInno.

[¶]University of Zagreb, Faculty of Science, Department of Mathematics. E-mail: hsikic@math.hr. Research supported in part by the MZOS grant 037-0372790-2799 of the Republic of Croatia and in part by Croatian Science Foundation under the project 3526.

determined by the respondents' posterior probabilities of the true state of nature, we prove that, under additional smoothness and sensitivity conditions, the payoff in the truth-telling equilibrium must be the logarithmic function of those posterior probabilities. Moreover, the respondents are necessarily ranked according to those probabilities. Finally, we discuss implementation issues.

Key words: proper scoring rules, robust/universal mechanisms, Bayesian Truth Serum, mechanism implementation, ranking experts

JEL codes: C11, D82, D83, M00

1 Introduction

Consider a problem of truthful elicitation of responses in a population of Bayesian agents who share a common prior.¹ We allow the survey planner to be "agnostic" about the prior (although the planner may have beliefs about the prior, she prefers to keep these beliefs private). The true outcome about which the responses are elicited is not verifiable. If the (lack of) respondent honesty or care is an issue, the survey planner may want to implement an incentive-compatible mechanism or 'scoring rule.' We say that a multi-person scoring rule is (strictly) incentive compatible, if, for each respondent π , when every other respondent responds truthfully, then providing the honest answer (strictly) maximizes that respondent's π expected score. Incentive compatible scoring rules play a major role in survey studies in various fields, most notably in economics and business. Thus, it would be of high value to both researchers and practitioners in numerous applications to characterize usable incentive-compatible algorithms. Are we to expect a variety of such algorithms, or is the incentive-compatibility a rather restrictive property? We show in this paper that, under a fairly natural "equilibrium-locality" condition, the property is more on the restrictive side, and guides us to an algorithm well-known in the literature. We now further elaborate more on the precise nature of our results.

Prelec (2004) introduced an algorithm based on two inputs by each respondent: a declaration of the respondent's type and his belief about the distribution of the type declarations by other players. BTS has two important properties (Prelec, 2004). First, it is strictly incentive compatible (IC), i.e., strict *truth-telling* is an equilibrium, so that the agents' "types" corresponding to their true responses, are fully revealed ². Second, the BTS equilibrium score of a respondent is, up to a linear transformation, equal to the logarithm of his posterior probability of the true state of nature , so that BTS *ranks respondents by posteriors*, henceforth called Posterior Ranking. We say that the BTS mechanism results in *logarithmic scoring*. BTS has been successfully applied in various fields, including new product adoption, Howie, Wang and Tsai (2011), economics and psychology, Kukla-Gryz, Tyrowicz, Krawczyk and Siwinski (2015), knowledge design, Miller, Bailey and Kirlik (2014) and criminology, Loughran, Paternoster and Thomas (2014).

The motivating questions for our paper are: Under which conditions do equilibrium payoffs gen-

¹For example, a market survey in which the respondents are asked whether they would buy a particular new product.

²In fact, we show that all strict equilibria in the BTS framework are either truth-telling or a types-permutation of truth-telling, and the scores are unique up to a linear transformation.

erated by a strictly Incentive Compatible scoring rule necessarily correspond to logarithmic scoring? Under which conditions must a strict Incentive Compatible equilibrium satisfy Posterior Locality? That is, under which conditions are equilibrium payoffs essentially equivalent to BTS payoffs?

Our main results are the following. We identify two conditions on equilibrium payoffs that we name "Posterior Locality", and a "Separation of Variables" property. The main theorem says

Incentive Compatible + Posterior Locality + Separation of Variables \longrightarrow logarithmic scoring.

The second result says

Incentive Compatible + Posterior Locality \longrightarrow Posterior Ranking.

To make these statements more precise, we describe now our setting in more detail. There is a state of nature drawn from a finite set, and an infinite population (for application purposes one can think of having a large population). There is a planner who asks each respondent to submit responses to a questionnaire. The respondents are players in a Bayesian game, with each player rewarded by the planner according to a score that depends on his responses and the responses of everyone else. Each player has a type, where types take values in a finite set. The variation in types can be interpreted as a consequence of players observing varying signals. (As was shown in Radanovic and Faltungs (2013), and in this paper using a different proof, reporting only types cannot lead to a truth-telling equilibrium.)³ The types are conditionally i.i.d. with respect to the state of nature, so that there is a single probability distribution, the "prior", that describes the joint distribution of the state of nature and types. We assume that the prior is common knowledge among players, but unknown to the planner ex-ante.

Let us now recall the notion of proper scoring rules in a framework with only one respondent. Consider a random variable Ω , taking values in $\{1, \dots, N\}$, $N > 1$, representing the state of nature. Based on the respondent's declarations, his belief about the distribution $\tilde{p} = (\tilde{p}^1, \dots, \tilde{p}^N)$ of Ω is implied. As a consequence, if the outcome $\Omega = i$ occurs, the respondent is paid $F_i(\tilde{p})$. In the

³We expect that our results would still hold approximately for finite, but large sample sizes. The exact theory for the finite case is very different and left for future research. For some elements of the finite case theory within the information-theoretic setup see Cvitanić, Prelec, Radas and Šikić (2019).

multi-respondent setting of this paper, each i will itself be a distribution, but over the beliefs of the respondents. In that case, if \tilde{p} is implied by respondent's answer and if the outcome $\Omega = i$ occurs, we will refer to \tilde{p}_i as his "local posterior".

With only one respondent, a family of functions $\{F_i\}_{i=1,\dots,N}$ is called a strictly proper scoring rule if it is incentive compatible for truth-telling, that is, the respondent's expected payoff is maximized at his true belief, the respondent's posterior, denoted by p . More precisely, for all probability vectors $\tilde{p} \neq p$, we have

$$\sum_{i=1}^N p^i F_i(p) > \sum_{i=1}^N p^i F_i(\tilde{p}). \quad (1.1)$$

There are many proper scoring rules. A general characterization with many examples is provided in Gneiting and Raftery (2007)⁴. An important special case arises if $F_i(p) = F_i(p_i)$ depends only on the local posterior, which is the probability p_i the respondent assigns to the outcome $\Omega = i$ that is actually realized, and does not depend on how probabilities are divided among the remaining counterfactual outcomes. In that case, the scoring rule is necessarily equal to a linear transformation of the logarithm of p_i (Savage (1971), Bernardo (1979)). Such a rule is a natural choice if the local posterior is interpreted as a measure of respondent's expertise, that is, if the quality of respondent's signal/type is measured by the probability assigned to the true state of nature.

As mentioned above, in the multi-player setting the values of Ω are interpreted as distributions of the types in the infinite population, and posterior as the probabilities that a type assigns to those values. We do not allow all distributions, but only finitely many. In practice, the respondents would be given a discretized choice of distributions, for example, "Do you think the percentage of votes for candidate A will fall in the range 0-10%, 10-20%, ...". Assuming that our game has a strictly type-separating equilibrium, our results depend only on the form of equilibrium payoffs, and not on the actual scoring rules that lead to those payoffs. The Posterior Locality condition posits that equilibrium payoffs F_i in the state $\Omega = i$ are functions of local posteriors, that is, of the posterior probabilities of the event $\Omega = i$.

Anticipating our results, we show now that, under mild smoothness conditions, if the equilibrium payoffs F_i satisfy a property analogous to (1.1), the difference in the state i scores of two respondents with local posteriors p^i and q^i , respectively, has to be approximately proportional to $\log(p^i) - \log(q^i)$ for

⁴Offerman, Sonnemans, Van De Kuilen and Wakker (2009) consider the case in which the respondents may have non-expected utilities, with only two possible states of nature.

$q \approx p$, up to the first order. To explain what we mean by that, assume now, for simplicity of notation, the case with two types, with local posteriors p^i and q^i . If a player implies, by his responses, local posterior p^i , he receives $F_i(p^i, q^i)$, and if he implies q^i , he receives $F_i(q^i, p^i)$. The incentive compatibility of the family $\{F\}_i$, as generalized from the one-player incentive compatibility (1.1) of proper scoring rules, means that the solution to the problem

$$\min_{q^i} \left\{ \sum_i p^i [F_i(p^i, q^i) - F_i(q^i, p^i)] + \lambda \sum_i q^i \right\}$$

is $q^i = p^i$, where λ is a Lagrange multiplier for the constraint $\sum_i q^i = 1$. The first order condition for the above problem is, denoting by ∂_x the partial derivative with respect to the x argument,

$$\partial_q F_i(p^i, p^i) - \partial_p F_i(p^i, p^i) = -\frac{\lambda}{p^i}$$

In the one player case, this reads $F'(p_i) = \lambda/p_i$, and results in the log function as the only proper scoring rule that satisfies locality. In our multi-player case, we fix p^i and expand the score difference up to the first order as a function of q^i around the point p^i , to get

$$F_i(p^i, q^i) - F_i(q^i, p^i) \approx [\partial_q F_i(p^i, p^i) - \partial_p F_i(p^i, p^i)](q^i - p^i)$$

Finally, combining the above equations and using that the first order Taylor expansion of the log function around $q_i = p_i$ is $\log(p_i) - \log(q_i) = \frac{1}{p_i}(p_i - q_i)$, we get

$$F_i(p^i, q^i) - F_i(q^i, p^i) \approx \lambda \left(1 - \frac{q^i}{p^i}\right) \approx \lambda(\log(p^i) - \log(q^i))$$

Building on this approximation, our first theorem says: if we add to incentive compatibility mild requirements on payoff smoothness and sensitivity on other players of the difference in equilibrium payoffs of two respondents, then the difference in incentive compatible scores of the two respondents is exactly proportional to the difference in logarithms of the implied local posteriors, rather than only approximately.

Our second theorem says that any incentive compatible equilibrium payoff $F_i(p_k^i, p_{-k}^i)$ of the player who implies, via his responses, probability p_k^i corresponding to type k , given that other types imply probabilities p_{-k}^i , is non-decreasing in the local posterior probability p_k^i . Consequently, the ranking of experts in equilibrium, if we consider p_k^i to be the measure of expertise, is the same given any incentive compatible mechanism, and it corresponds to the ranking by posteriors. The result is very general,

proven by purely algebraic methods. It is a generalization of the results in the literature in the case of one respondent, on the monotonicity being implied by incentive compatibility of proper scoring rules. See, e.g., McCarthy (1956), Savage (1971), Schervish (1989) and Schlag and van der Weele (2013).⁵

We also discuss implementation issues. Observe that in general, while a particular ex-post payoff of the form $F_i(p_k^i, p_{-k}^i)$ may arise in equilibrium in theory, it is not necessarily simple to implement it in practice. That is, the problem is how to implement the theoretically optimal payoff score using only the players' responses to a questionnaire designed by the agnostic planner, while having the questionnaire as simple as possible. Under an assumption somewhat stronger than Posterior Locality, but without assuming Separation of Variables, we show that the payoffs of all strictly-separating equilibria in our framework can be implemented by particular questionnaires, but the latter may be complex, except for the logarithmic, BTS case. In this context, let us recall that Prelec (2004) shows that promising the respondents the BTS scores ex-ante, results, in equilibrium, in the ex-post scores of the form $\log(p_k^i)$ (plus a term that does not vary with a respondent). We revisit this result and provide a detailed proof thereof. We also show that the budget-balanced strict equilibrium under BTS are necessarily separating.

Relationship to existing literature. Since Prelec (2004), proper scoring rules in the game-theoretic context, also called “information elicitation without verification”, have been studied extensively in recent years. In the case in which the planner knows the prior distribution of the player types, an early work is Miller, Resnick, and Zeckhauser (2005) who design a clever use of proper scoring rules to elicit truthful information. However, as the assumption that the planner knows the prior is unrealistic in practice, alternative approaches have been proposed. Jurca and Faltings (2007) use robust optimization to deal with small variations in belief models. Expanding on the framework of Prelec (2004) in which the planner does not know the prior, but the number of players is infinite, Witkowski and Parkes (2012) and Witkowski (2014) devise mechanisms, under the name Robust Bayesian Truth Serum (RBTS) that work for a finite number of respondents, but with only two types. Waggoner and Chen (2013) consider a general framework without assumptions on information structure. Radanovic and Faltings (2013, 2014) and Zhang and Chen (2014) develop mechanisms that are incentive compatible for any number of agents and non-binary player types. Dasgupta and Ghosh (2013) and

⁵See also Prelec, Seung and McCoy (2017) who define and test experimentally a broader class of algorithms to produce a ranking of experts according to their posteriors. Within this class, only BTS is known to be incentive-compatible.

Witkowski and Parkes (2013) relax the knowledge requirement even more by asking for more extensive reports by respondents in form of responses to multiple similar questions. Baillon (2016) implements truth-telling equilibria via a Bayesian Market, in the case of binary types. Cvitanić, Prelec, Riley and Tereick (2018) provide additional incentive-compatible mechanisms which are simple to explain to respondents. In Frongillo and Witkowski (2017), the authors characterize “minimal” peer-prediction mechanisms, that is, those in which the score depends only on the respondent’s reported type, and the reported type of another, suitably chosen “peer” respondent. We, on one hand, allow the scores to depend on reported types of all the respondents and on their predictions of those reports, and, on the other hand, we impose the assumption of locality on the scores. The preference for one or the other approach depends on whether minimality or locality is the preferred feature of the mechanism. The paper Kong and Schoenebeck (2019) develops an informational theoretic paradigm for designing incentive mechanisms, which includes, as special cases, many established mechanism, including BTS, and shows that properties of BTS can be proved in a simpler way by using a connection to Shannon mutual information. For a different connection between BTS and information theory, see Cvitanić, Prelec, Radas and Šikić (2019). Liu and Chen (2019) design a “uniform dominant” truth serum when there is a noisy signal of the ground truth, and there are sufficiently many agents and tasks. Their scoring depends on whether a report is informative or not, thus, in a sense, related to expertise of the respondent.

However, with the exception of Prelec (2004), all the papers mentioned above are concerned with incentive compatibility and not with ranking of respondents, so that the proposed mechanisms either do not satisfy our Posterior Locality condition, or they require the planner to know the common prior of the respondents. We focus on mechanisms that have all of these three properties: they allow IC equilibria, they can be implemented even when the planner is agnostic about the prior, and they rank the respondents by posteriors. In this case our main result says that, under relatively weak conditions, logarithmic scoring is the only possible equilibrium payoff form. When ranking by posteriors is not required from the mechanism, then the above papers provide many other ways to design IC mechanisms.

The problem we tackle in the paper can also be considered as one of mechanism design, since we seek to describe mechanisms that are both incentive-compatible and have attractive features for opinion elicitation applications. In one way, our approach is more general than typical mechanism

design models⁶ because we allow for both uncertainty regarding the players information (type), and uncertainty regarding the true state of nature, and those two may be correlated in a nontrivial way. It is exactly the joint distribution of the two that drives all the results. Our basic assumption is that the players have a common prior on this joint distribution, but that the prior is not used by the planner in designing the survey. We present this as a methodological rather than a substantive requirement: Although the planner may have some beliefs about the prior, she may prefer to keep these beliefs private and adopt the position of an agnostic/neutral outsider, not imposing her conjectures on the survey respondents. Thus, she is interested in an 'universal' mechanism, one that would work for all priors without any input from her side apart from the initial formulation of the multiple-choice question.⁷ In this sense, ours is a study of robust Bayesian mechanisms. On the other hand, our setup is less general in another way – the players do not choose actions other than reporting their responses, which is assumed to be costless. Thus, there is no modeling of utility/disutility drawn from actions, the only utility the players draw is from the expected payoff they attain. Moreover, our framework is less general than some models of robust mechanism design that, unlike ours, do not assume common knowledge of the prior distribution by all the players. (In our case only the planner may be ignorant.) We discuss in the conclusions section in what directions one could try to extend our results.

The rest of the paper is organized as follows: Section 2 introduces the model, Section 3 presents the main theoretical results, Section 4 discusses implementation issues, and we conclude in Section 5. The proofs are presented in the appendix.

2 Model, Definitions and Assumptions

In our model, a mechanism consists in giving scores to the players (respondents) of different types.⁸ Applications we have in mind are of the polling type: the respondents are asked to provide responses to queries assigned by a survey planner. The planner is interested in eliciting truthful opinions to a

⁶See, e.g., Maskin and Sjöström (2001), Bergemann and Morris (2012), Börgers (2013). We refer the reader to Bergemann and Morris (2012) for a detailed literature survey.

⁷In theory, using the "majority rule" mechanism that would ask for the common prior to be declared may result in an equilibrium that reveals the common prior; however, such a rule would not be implementable in practice, as we discuss in the paper.

⁸A negative of a score is usually called a transfer in the mechanism design literature; see, e.g., Börgers (2013).

multiple choice question, and ranking players according to the quality of their information, which in our framework, will mean according to their posterior probabilities of the true state of nature. For instance, the planner might be interested in the value of a certain wine bottle some years into the future, and asks experts to respond to appropriately designed questions. Broader applications include voting in elections, predicting political events, product market research, online product reviews, and any other application that involves a survey with a multiple choice question.⁹ We remark that it is not necessary to assume that the true response to the multiple choice question is verifiable.

2.1 The model

The players are indexed by $\pi \in R$, where R is infinite and countable.¹⁰ The state of nature is a random variable Ω , taking values in $\{1, \dots, N\}$, $N > 1$.¹¹ The players can be of $M > 1$ different types, that can be interpreted as random signals the players receive about the state of nature. Player π 's type is a random variable denoted T^π , and it takes values $t^\pi \in \{1, \dots, M\}$. We consider scoring mechanisms in which, for a given fixed positive integer K , player π submits as a response a K -dimensional value $a^\pi \in \mathbb{R}^K$ (a for “action”).

A response a^π would typically include a declaration of a respondent's type (choosing an answer to a multiple choice question), and it would also include responses to some other questions in order to be truth-inducing.¹² It could also include a declaration of the respondent's prior distribution of types and states of nature, as introduced below; that is, the respondent could be asked to state what his

⁹Prelec (2004) and Prelec, Seung and McCoy (2017) provide many more examples.

¹⁰We need the assumption that there are infinitely many players for several reasons: first, we don't want to impose assumptions on the form of the payoffs outside of equilibrium; for this, we will use the fact that, with infinite number of players, the form of equilibrium payoff does not change when a player of one type mimics the equilibrium strategy of another type; second, achieving truth-telling of types is much harder with finitely many players, and so is the implementation of equilibrium payoffs using practical inputs. We postpone to future research the analysis of the setup with finitely many players; finally, we need the infinite number of players because we invoke de Finetti's theorem in our model setup.

¹¹Strictly speaking, this is only an approximation for most applications, in which the state of nature could naturally have a continuous range of values. For instance, in the example about a wine bottle's value, the state of nature could be the percentage of experts who believe the bottle is worth more than one thousand dollars.

¹²In the section on implementation, we will see that another question might be about the percentage of other respondents choosing a specific choice from the multiple choice list.

prior is. We posit the following

Assumption 2.1 (i) *The family of signals $T^\pi, \pi \in R$, is a family of exchangeable random variables, and random variables $T^\pi, \pi \in R$, are i.i.d. conditional on the state of nature Ω .*

(ii) *If respondent π chooses response a^π , and the remaining responses are represented by $a^{-\pi}$, then his score is given by the scoring function $f(a^\pi, a^{-\pi})$, where the order of different respondents' responses in $a^{-\pi}$ does not matter, that is, f is symmetric in those.*

Condition (i) implies that the order in which we consider our players is irrelevant (from the point of view of the probability distribution of the entire sequence). Moreover, by de Finetti's theorem, the exchangeability assumption actually implies the second part of Assumption 2.1, that there exists a random variable Ω such that T^π 's are conditionally i.i.d. with respect to Ω ; see, e.g., Aldous (1985), or Chow and Teicher (1997).

The symmetry property in condition (ii) is a natural restriction considering that the planner does not make a distinction between different types, assumed exchangeable by condition (i).

From now on, we assume the players are risk-neutral, that is, each player maximizes his expected payoff.¹³

2.1.1 The prior and the posteriors

The joint distribution of types and states of nature is given by an $M \times N$ matrix $Q = [q_k^i]$, where

$$q_k^i = \Pr(T^\pi = k, \Omega = i).$$

Note that Q does not actually depend on π , a consequence of the exchangeability assumption.

We suppose that the matrix Q is common knowledge among the players, and refer to it as their “common prior”. However, Q is not used by the planner when designing the survey. In fact, the planner does not even need to know the number of the states of nature N . The only thing we assume

¹³Typically, mechanism design models consider only the types as being random, according to a prior which is known also to the planner. Our model is more general by considering random states of natures in addition to random types, with a non-degenerate correlation between the two. On the other hand, it is less general in that the players do not choose actions other than choosing what response to give.

that the planner uses is M . For example, M is needed for implementation using a multiple choice question – the planner has to offer exactly as many possible choices as there are types.¹⁴

Matrix Q determines the marginal probabilities of types, referred to as *type probabilities*, and the probabilities of states of nature given the type, referred to as *posteriors*. They are denoted

$$s_k = \Pr(T^\pi = k)$$

and

$$z_k^i = \Pr(\Omega = i \mid T^\pi = k).$$

We assume that the marginal probabilities of types and of states of nature are all strictly positive. The posteriors form a matrix $Z = [z_k^i]_{k=1, i=1}^{M, N}$. Note that z_k^i does not depend on π , that for every $k \in \{1, \dots, M\}$, we have $\sum_{i=1}^N z_k^i = 1$, and that any matrix with this property can be represented as a Z –matrix of posteriors for some joint distribution Q . We denote the vector (s_1, \dots, s_M) by S .

2.2 Equilibrium payoff and incentive compatibility

In the standard literature on scoring rules, there is only one respondent, asked to declare his posterior belief about the distribution of Ω , that is, to declare z^i ’s. If the outcome $\Omega = i$ occurs, the respondent is paid $F_i(z)$. A family of functions $\{F_i\}_{i=1, \dots, N}$ is called a strictly proper scoring rule if it is incentive compatible for truth-telling, that is, the respondent’s expected payoff is maximized at his true belief, meaning, for all probability vectors $\tilde{p} \neq p$, we have $\sum_{i=1}^N p^i F_i(p) > \sum_{i=1}^N p^i F_i(\tilde{p})$.

In our framework with infinitely many respondents, we consider only the payoff mechanisms that allow for a strictly separating Bayesian Nash equilibrium (SSNE), as defined below, in which the equilibrium payoffs are functions $F_i : (0, 1)^{2M} \rightarrow \mathbb{R}$, of the form $F_i(z_k^i, z_{-k}^i; s_k, s_{-k})$ where, for example, $z_{-k}^i = (z_1^i, \dots, z_{k-1}^i, z_{k+1}^i, \dots, z_M^i)$. We call this property *Posterior Locality*.

We now make this more precise. A pure strategy for player π is a map σ^π that maps a player’s type to his response choice a^π . We allow only pure strategies. The profile of all respondents’ pure strategies is denoted $\sigma(t)$, with entries $\sigma^\pi(t^\pi)$, and the profile excluding player π is denoted $\sigma^{-\pi}(t^{-\pi})$.

¹⁴To get around the problem of not knowing the common prior the planner could ask each player to state the whole prior distribution and harshly penalize the player who gives a response different from others. However, asking for the common prior is unlikely to work in practice - more likely than not, most responses would be different from each other, and the planner would have to penalize harshly most respondents.

The score for player π is given by $f(\sigma^\pi(t^\pi), \sigma^{-\pi}(t^{-\pi}))$, where f is a scoring function that takes the responses to the set of real numbers. Function $f(\cdot, \cdot)$ is of the same functional form for all N and Q .

We assume that the players maximize the expected score value. We will mostly restrict the payoff mechanisms to those which are “budget balanced”, that is, those for which the sum of the scores of all the players is equal to zero, with probability one.¹⁵

Here is the definition of equilibrium:

Definition 2.1

- (i) Given a prior matrix Q , we say that a scoring function f allows a Strict (Bayesian) Nash Equilibrium (SNE) if there exists a pure strategy profile $\sigma = \sigma_Q$ such that for all $\pi, t^\pi, t^{-\pi}, t^\gamma$, we have:

For an arbitrary response choice $a^\pi \neq \sigma^\pi(t^\pi)$, we have, with expectation taken with respect to the (conditional) distribution of Ω ,

$$E[f(a^\pi; \sigma^{-\pi}(t^{-\pi})) \mid T^\pi = t^\pi] < E[f(\sigma^\pi(t^\pi), \sigma^{-\pi}(t^{-\pi})) \mid T^\pi = t^\pi]$$

The strategy profile σ is called an SNE. If the equilibrium is also separating, that is, if, in addition to the above, we also have $\sigma^\pi(t^\pi) = \sigma^\gamma(t^\gamma) \Rightarrow t^\pi = t^\gamma$, we call σ a Strictly Separating (Bayesian) Nash Equilibrium (SSNE).

- (ii) We say that a scoring function f is an Universal Separating Scoring Rule (USSR), if for all Q it allows at least one budget-balanced SNE σ_Q , and if every budget-balanced SNE is an SSNE.

We will show later below that the budget-balanced logarithmic scoring can be implemented by an USSR (that is, by BTS), which also satisfies the assumptions below on the equilibrium payoffs.

From now on, we only consider USSR functions f , or non-budget-balanced versions thereof, so that there exists at least one SSNE. We use notation F_i for the state i ex-post payoff in an SSNE corresponding to f , budget-balanced or not. The following is the additional crucial assumption we impose, and it is an assumption on the ex-post, equilibrium payoffs F_i of an SSNE.

¹⁵It should be mentioned that in a budget-balanced game the players know they may receive negative “payments”, and some players may not be willing to participate. In practice, the “payments” will often not be monetary, but used as score points, and every respondent might be paid a non-negative amount, that may consist of a fixed fee and a variable fee that depends on the respondent’s score, or his ranking according to the scores. That is, what is used may not be a budget-balanced scoring rule, but a modification thereof.

Assumption 2.2

- **Posterior Locality.** $(\forall k \in \{1, \dots, M\})$ and $(\forall i \in \{1, \dots, N\})$, and $\forall j \neq k$, if $T^\pi = k$, and $\Omega = i$, the equilibrium score (but not necessarily the out-of-equilibrium scores) of player π has the representation, with $F_i : (0, 1)^{2M} \rightarrow \mathbb{R}$,

$$f(\sigma_Q^\pi(k), \sigma_Q^{-\pi}(t^{-\pi})) = F_i(z_k^i, z_{-k}^i; s_k, s_{-k}) .$$

We discuss this assumption in Section 2.2.1 below. We do not address uniqueness of equilibrium, and we study only the equilibria in which the realized equilibrium payoffs (but not necessarily the out-of-equilibrium payoffs) are of the above form. Moreover, we require that the realized equilibrium payoffs satisfy the conditions in the following definition.

Definition 2.2 *The family $\{F_i\}$ of functions of the form $F_i(z_k^i, z_{-k}^i; s_k, s_{-k})$ is called a Posterior-Local Equilibrium Payoff System (PLEPS) if the following is satisfied:*

- (i) **Symmetry:** $(\forall x, y \in (0, 1))$ $(\forall z_2, \dots, z_M, s_2, \dots, s_M \in (0, 1))$ $(\forall \text{permutation } \Pi \text{ of } \{2, \dots, M\})$, we have

$$F_i(x, z_2, z_3, \dots, z_M; y, s_2, \dots, s_M) = F_i(x, z_{\Pi(2)}, z_{\Pi(3)}, \dots, z_{\Pi(M)}; y, s_{\Pi(2)}, s_{\Pi(3)}, \dots, s_{\Pi(M)})$$

- (ii) **Incentive compatibility, strict separation inequality:**

$$(\forall Z - \text{matrix}) (\forall S - \text{vector}) (\forall k, j \in \{1, \dots, M\} \text{ such that } (z_k^1, \dots, z_k^N) \neq (z_j^1, \dots, z_j^N))$$

$$\sum_{i=1}^N z_k^i F_i(z_k^i, z_{-k}^i; s_k, s_{-k}) > \sum_{i=1}^N z_k^i F_i(z_j^i, z_{-j}^i; s_j, s_{-j}) \quad (2.1)$$

Assumption (i) on symmetry means that the equilibrium score of type k does not depend on the order of other types, and is consistent with Assumption 2.1 (ii) on the symmetry of scoring function f . Assumption (ii) implicitly assumes that the players are risk-neutral and maximize the expected score. By Proposition 2.1 below, it is automatically satisfied if F_i are the equilibrium payoffs in a truth-telling equilibrium.

We now elaborate more on the assumed form of equilibrium payoffs F_i .

Remark 2.1 The crucial assumption for the results of this paper is that the score of a player in equilibrium depends on the player's posterior z^i of the realized state of nature i , called local posterior. This is justified if the posterior is a good measure of a player's expertise. Note, however, that here we

think of expertise regarding the actual state of nature in this one particular survey, not about average accuracy over many surveys. It is the ex-post expertise, resulting from the signal a player receives, and not the ex-ante expertise of obtaining good signals.

There are cases in which the planner clearly wants to know about the distribution of types, such as elections or product market research, trying to estimate what percentage of population will vote for each candidate, or is likely to buy a product. In such cases, it is intuitive that a respondent with higher local posterior is a better expert – he has the highest probability of being right about the actual distribution of responses, reminiscent of the concept of maximum likelihood estimators that maximize the probability of the event that does actually occur. Moreover, if the survey study has more than one stage, for example, in market research, a mechanism that results in PLEPS payoffs could be used to identify experts in the first stage, and then only the experts could be used for further surveys, thus reducing the cost of the study. It is primarily these applications we have in mind. In other applications, such as, for example, surveying economists on whether this year’s inflation will be higher than a certain level, it is less clear that a higher local posterior on the distribution of types means a higher expertise. This is because in this example it is not necessarily the case that those that are better at estimating the percentage of their colleagues who will predict high inflation are also better at predicting the inflation. In such cases scoring rules other than those with ex-post PLEPS payoffs might be appropriate. In particular, if the planner is not concerned with identifying experts, but only with truth-telling, the assumption may exclude perfectly reasonable scoring rules, as in the papers mentioned in the literature survey in Introduction.

We also note that we look for the simplest possible equilibrium payoffs that describe players’ expertise, which is why the payoff F is not allowed to depend on other local probabilities that can be derived from the prior. On the other hand, the reason why we allow dependence on ex-ante type probabilities s_k, s_{-k} is because these, in implementation, translate to type frequencies, which may be used to make a mechanism budget-balanced. Actually, for budget balance, it is sufficient to have dependence on local conditional probabilities $s_k^i = \Pr(T^\pi = k \mid \Omega = i)$, but we allow dependence on s_k, s_{-k} for generality (except in the implementation section), as discussed next.

A natural question to ask is whether for any PLEPS F there exists a scoring rule f that implements it in equilibrium. In the implementation section below we argue that this is, indeed, the case, under the assumption that, instead of on possibly all s_k, s_{-k} , F_i depends only on $s_k^i = \Pr(T^\pi = k \mid \Omega = i)$.

It is also natural to ask if, for a given f , the equilibrium that implements F is unique. We show later below that this is essentially true for the benchmark example of the Bayesian Truth Serum scoring rule.

The following result is simple, but crucial for our results. It tells us what the score looks like for the type who mimics another type's equilibrium strategy. We emphasize that we need infinite number of players for this result.

Proposition 2.1 *Suppose there exists a strictly separating Bayesian Nash equilibrium for our game of respondents for which the ex-post payoffs are given by PLEPS $\{F_i\}$. Then, if a respondent of type k deviates from the equilibrium by using the strategy of type $j \neq k$, his deviation payoff is equal to $F_i(z_j^i, z_{-j}^i; s_j, s_{-j})$. That is, if a player of type k mimics the equilibrium strategy of type j , then his payoff is given by the equilibrium evaluation corresponding to type j .*

This holds because every type is represented by infinitely many players, the equilibrium payoffs are strictly separating, and the scoring function f is symmetric in their responses. The proof is in Appendix.

We have the following negative result, proved in Appendix, when the number of players is finite.¹⁶

Proposition 2.2 *Assume (only in this proposition) a finite number of players, and at least two players. Then, there exists no budget-balanced PLEPS.*

We also remark that even a non-budget balanced version of BTS is not incentive compatible when there are finitely many players.

2.2.1 Ex-ante vs. ex-post payoff: implementation

Even when identifying states of nature with possible empirical frequencies of responses, asking about posterior probabilities of state of nature is likely to be prohibitively complex in practice, because it would require respondents to provide a distribution over all possible empirical frequencies. Thus, in practice, the planner who wants the mechanism to result in ex-post payoffs F_i when the players play the truth-telling equilibrium, would like to find a way to induce those ex-post payoffs by promising to pay the players based on ex-ante scores that require much simpler inputs than the players' beliefs

¹⁶We leave for a future study a more thorough analysis of the case with the finite number of players.

about the distribution of the empirical frequencies. We will discuss this issue in the implementation section, and here we just mention the following. Our benchmark example of a PLEPS is the classical logarithmic scoring rule payoff

$$F_i(z_k^i, z_{-k}^i; s_k, s_{-k}) = \log(z_k^i) .$$

Prelec (2004) showed that the budget-balanced version of this payoff can be implemented by, in addition to asking (infinitely many) respondents to declare their own type – the multiple choice question – also asking them what they think is the percentage of other types in the population, that is, the empirical frequencies of each choice in the multiple choice question. It is much easier for the respondents to provide their estimates of empirical frequencies than their estimates of the probability distribution of the empirical frequencies. In the logarithmic case, this means that a respondent of type k is not asked for z_k^i 's and is not promised $\log(z_k^i)$, but he is asked for simpler inputs that determine his promised score via a specific function f (called Bayesian Truth Serum), and the value of the score will turn out to be equal to $\log(z_k^i)$ when the players play the truth-telling equilibrium.

3 Possible equilibrium payoffs

In this subsection we present examples of PLEPS's and address the question whether logarithmic equilibrium payoffs, or simple modifications thereof are the only possible PLEPS's.

3.1 Logarithmic equilibrium payoffs

3.1.1 The benchmark example – the logarithmic function

The canonical example of a PLEPS (ignoring budget-balancing) is the logarithmic function:

$$F_i(z_k^i, z_{-k}^i; s_k, s_{-k}) = \log(z_k^i)$$

More precisely, a player's equilibrium payoff is the logarithm of the posterior probability of the state of nature given his type. It is well known and straightforward to verify that this, indeed, satisfies the strict separation inequality (2.1). This is because of the well known *Gibbs inequality* which says that for a probability vector (p^1, \dots, p^N) , we have

$$0 = \min_{q^i \geq 0, \sum_i q^i = 1} \sum_{i=1}^N p^i [\log(p^i) - \log(q^i)] \tag{3.1}$$

This can be verified by noting that, with λ being a Lagrange multiplier for the constraint $\sum_i q^i = 1$, the first order conditions for the problem

$$0 = \min_{q^i} \left\{ \sum_{i=1}^N p^i [\log(p^i) - \log(q^i)] + \lambda \sum_i q^i \right\} \quad (3.2)$$

are $p^i/q^i = \lambda$, thus satisfied with $q^i = p^i$.

The question arises whether the log function is the only PLEPS (modulo budget balancing). The answer is negative in general, and we present a counterexample in what follows. Later below, we show that under mild additional conditions logarithmic equilibrium payoffs are, in fact, the only possible PLEPS's.

3.1.2 Other examples of PLEPS's

Let us first note that there are variations of the logarithmic equilibrium payoffs that produce equivalent scores when we require budget balance. For instance, if we set, for some function G symmetric in all the arguments, and some constant K , suppressing the dependence on the state of nature i ,

$$F(z_k, z_{-k}) = \log(z_k) - K \sum_{j \neq k} \log(z_j) + G(z_1, \dots, z_M)$$

then, function F corresponds to a PLEPS, as can be verified in the same way as for the problem (3.2). However, it is not really different from logarithmic equilibrium payoffs if we insist on budget balance, because, as is straightforward to check, if we add the constant term that makes it budget-balanced, we get the same equilibrium payoffs as for the budget-balanced logarithmic equilibrium payoffs.

We now present a PLEPS that has higher order terms that make it distinct from the logarithmic PLEPS, even if we make it budget-balanced.

Example 3.1 Consider the case with three types, $M = 3$, and denote

$$p^i = z_k^i, \quad (q^i, r^i) = z_{-k}^i$$

Define the following function:

$$F(p, q, r) = K \log(p) + p^4 - 2p^3(q + r) - 6p(qr^2 + q^2r)$$

It is straightforward to verify that, for large enough K , this function satisfies the strict separation inequality (2.1). This is because the first order conditions (FOC's) for the Lagrangian optimization problem

$$\min_{q^i} \left\{ \sum_i p^i [F(p^i, q^i, r^i) - F(q^i, p^i, r^i)] + \lambda \sum_i q^i \right\}$$

are, denoting with $\partial_x F$ the derivative with respect to x argument,

$$p^i [\partial_p F(q^i, p^i, r^i) - \partial_q F(p^i, q^i, r^i)] = \lambda \quad (3.3)$$

for some Lagrange multiplier λ , and these FOC's are satisfied for the above function with $q^i = p^i$. For large enough K , the FOC's are also sufficient conditions for optimality because the second order optimality conditions will also be satisfied, which implies that (2.1) is satisfied.

Remark 3.1 Even though there are “strange” PLEPS functions F_i as in the example above, as we explained in Introduction, for all of them the difference in two equilibrium payoffs is proportional to the difference of logarithmic payoffs, up to the first order. This is also true if F_i depends on type probabilities s_k , under the conditions of Lemma 3.1 below.

We next identify conditions under which there can be no second-order terms, and the budget-balanced logarithmic equilibrium payoff is the only budget-balanced PLEPS.

3.2 When are equilibrium payoffs logarithmic?

We assume in this section that $N \geq 3$.¹⁷ As we have just shown, the difference in the ex-post scores of two types is equal to the difference of the log scores up to the first order. We will now find conditions under which the higher order terms cannot appear, and under which any PLEPS is essentially a logarithmic equilibrium payoff.

We do the following:

- (i) we first state an assumption on the second order mixed derivative of the difference in equilibrium scores of two types;
- (ii) we then show that the assumption implies an additive representation of the equilibrium payoff of a given type – the equilibrium payoff is a sum of a term that does not depend on the posteriors of other types and a term that is symmetric in types.

¹⁷It is well known that there are quadratic scoring rules that are strictly separating when $N = 2$, for all priors.

- (iii) finally, we show that such additive representation is sufficient to imply logarithmic equilibrium payoffs, under a smoothness assumption.

For ease of notation we continue assuming $M = 3$, and use the above notation p^i, q^i, r^i for the local posteriors of the three types. Also denote by s_p, s_q, s_r the corresponding type probabilities. This is without loss of generality, the same proof works for more than three types.

The following is the assumption on “separation of variables” that we need; not surprisingly, in light of the first-order approximation above, it is an assumption on the second-order properties of the equilibrium payoffs. In particular, it is weaker than the assumption that the difference in equilibrium payoffs of two types does not depend on other types.

Assumption 3.1 *For all i , and all type probabilities s_p, s_q, s_r , the second mixed derivative (assumed to exist)*

$$\partial_{pq} [F_i(p^i, q^i, r^i; s_p, s_q, s_r) - F_i(q^i, p^i, r^i; s_q, s_p, s_r)]$$

of the difference in scores of two types with local posteriors p^i and q^i respectively, does not depend on other types’ local posteriors r^i .

The assumption says that the (mixed) sensitivity of the difference in equilibrium payoffs to the corresponding types is not affected by other types.

We now state the following additive representation result, proved in Appendix.

Proposition 3.1 *Consider a PLEPS $\{F_i\}$ such that Assumption 3.1 holds. Then, if, for some $p^0 \in (0, 1)$ and for any fixed type probabilities s_p, s_q, s_r the function $F_i(p^i, q^i, r^i; s_p, s_q, s_r)$ can be expanded as an infinite Taylor series around the point $(p^i, q^i, r^i) = (p^0, \dots, p^0) \in (0, 1)^M$, then, necessarily, the following Additive Representation (AR) holds:*

$$F_i(p^i, q^i, r^i; s_p, s_q, s_r) = G_i(p^i; s_p, s_q, s_r) + H_i(p^i, q^i, r^i; s_p, s_q, s_r) \quad (3.4)$$

where H_i is a function that is symmetric in all the pairs $(p^i, s_p), (q^i, s_q), (r^i, s_r)$, $i = 1, \dots, N$.

The main result of the section is the following:

Theorem 3.1 *Consider a PLEPS consisting of functions $F_i(p^i, q^i, r^i; s_p, s_q, s_r)$, $i = 1, 2, \dots, N$, that satisfy the assumptions of Proposition 3.1. Assume also that F_i is such that G_i is symmetric in all*

s_k variables, for every fixed p^i , $i = 1, \dots, N$. Then, we have, for some functions λ and B of type probabilities $S = (s_p, s_q, s_r)$,

$$G_i(p^i, s_p, s_q, s_r) = \lambda(S) \log p^i + B_i(S)$$

In particular, if the corresponding PLEPS is budget-balanced, the equilibrium payoff to the respondent with local posterior p_i is given by

$$F_i(p^i, q^i, r^i; s_p, s_q, s_r) = \lambda(S) \log p^i - \lambda(S) \sum_{t=p,q,r} s_t^i \log t^i \quad (3.5)$$

where s_o^i is the conditional probability of the type with posterior o in state i , $o = p, q, r$.

Remark 3.2 We emphasize again that this result is obtained by restricting only equilibrium properties of a scoring rule, without restrictions on the off-equilibrium properties.

Proof: Since F_i is a PLEPS, it satisfies separation property (2.1). By the stated symmetry of H_i , function G_i also satisfies the same type of inequality, which can be written as

$$0 = \min_{q^i} \left\{ \sum_i p^i G_i(p^i; s_p, s_q, s_r) - \sum_i p^i G_i(q^i; s_p, s_q, s_r) \right\}, \quad (3.6)$$

As shown in Savage (1971), this property implies that G_i is continuously differentiable in the p^i variable, $i = 1, \dots, N$. Then, by Lemma 3.1 below that identifies the first order condition for this minimization problem, there exists a Lagrange multiplier $\lambda(S)$ independent of p , such that, suppressing dependence on i ,

$$\lambda(S) \frac{1}{p^i} = \partial_p G(p^i; s_p, s_q, s_r)$$

The above implies the statement about the logarithmic form of G_i . Equation (3.5) is then straightforward to verify. ■

The following “Lagrange optimization” lemma is proved in Appendix. It gives the first order condition for the IC minimization problem in (3.7) below.

Lemma 3.1 Consider, in the above notation, functions $F_i(p^i, q^i, r^i; s_p, s_q, s_r)$, $i = 1, 2, \dots, N$, that are continuously differentiable in the p^i and q^i variables, and, for every fixed p^i, q^i, r^i , symmetric in all values of s variables. Recall the strict separation inequality (2.1), written in the form

$$0 = \min_{q^i} \left\{ \sum_i p^i F_i(p^i, q^i, r^i; s_p, s_q, s_r) - \sum_i p^i F_i(q^i, p^i, r^i; s_p, s_q, s_r) \right\}, \quad (3.7)$$

that is, the minimum over probabilities q^i is obtained at $q^i = p^i$. Then, there exists a function $\lambda(S) = \lambda(s_p, s_q, s_r)$ such that, for all $i, p^i, q^i, r^i, s_p, s_q, s_r$,

$$\lambda(S) = p^i [\partial_p F_i(p^i, p^i, r^i, s_p, s_q, s_r) - \partial_q F_i(p^i, p^i, r^i, s_p, s_q, s_r)] . \quad (3.8)$$

3.3 Ranking by (local) posteriors

We now show the following result: PLEPS payoffs necessarily rank the players according to the relative ranking of the corresponding local posteriors. That is, when using a scoring system resulting in an equilibrium with PLEPS payoffs, the planner will know which players are better experts than others, if she considers the level of the local posterior equivalent to the level of expertise.¹⁸ We emphasize that for this result it is crucial to assume that the equilibrium scores depend only on the local posteriors of the realized state of nature.

The main result of this section is

Theorem 3.2 PLEPS payoffs $\{F_i\}$ are strictly increasing in the posterior probabilities of the true state of nature. That is, functions F_i satisfy (for any prior distribution matrix Q),

$$\text{If } j, k \in \{1, \dots, M\} \text{ and } z_k^i > z_j^i, \text{ then } F_i(z_k^i, z_{-k}^i; s_k, s_{-k}) > F_i(z_j^i, z_{-j}^i; s_j, s_{-j}) \quad (3.9)$$

Put differently, if the planner wants to determine relative expertise of players receiving exchangeable signals, it is sufficient to design a scoring system which allows only for equilibria that are realized via a PLEPS. Thus, inequality (2.1) not only guarantees strict separation of types, but also has the posterior-based ranking as a direct consequence.

This theorem is a generalization of the results in the literature on the monotonicity being implied by incentive compatibility of proper scoring rules. See, e.g., McCarthy (1956), Savage (1971), Schervish

¹⁸If they were not ranked by their local posteriors, then, in the pre-game phase, they might want to avoid collecting information about the true state of nature, which is undesirable.

(1989) and Schlag and van der Weele (2013). Those papers consider only the non-game version of the problem with one respondent only. Moreover, they mostly use analytic methods to prove it, while our proof is completely algebraic.¹⁹

The intuition behind the result is that if type A 's posterior probability of a state was higher than the one of type B , but type A 's score in that state was lower, then, he would be better off pretending to be type B . To be more precise, consider the case with only two types, A and B , and two states of nature, 1 and 2. Denote by p_A and p_B the posterior probabilities of state 1, and suppose, without loss of generality, $p_A > p_B$. There are only two possible PLEPS scores in each state i , denoted $F_i(p_A, p_B)$ and $F_i(p_B, p_A)$ (suppressing dependence on S vector). Denote by D_i the difference in scores, $D_i^A = F_i(p_A, p_B) - F_i(p_B, p_A)$. The claim is that, in equilibrium, type A 's higher posterior probability of state i implies higher score in that state, that is, positive D_i^A . To argue this, note first that by the strict separation inequality, player A 's expected value of the differences in scores, that is, the weighted average of D_1^A and D_2^A with weights p_A and $1 - p_A$, is positive. By the same token, the weighted average of D_1^A and D_2^A with weights p_B and $1 - p_B$, is negative. The only way this can be possible when $p_A > p_B$ (thus also $1 - p_A < 1 - p_B$) is that $D_1^A > 0$ and $D_2^A < 0$. Thus, indeed, the type with higher posterior probability of a state receives higher score in that state. Or, put differently, if the type with higher posterior probability of a state does not receive higher score in that state, he would adopt the other type's strategy. In Appendix, we state and prove the above simple argument in a lemma, and extend it to any number of types and states.

4 Implementation

In this section we first show how to implement any PLEPS, up to an additional mild restriction, and then we elaborate on the Prelec (2004) result that the Bayesian Truth Serum algorithm provides a feasible implementation of budget-balanced logarithmic equilibrium payoffs of (3.5) (under standard Bayesian and rationality assumptions); see also Prelec, Seung and McCoy (2017). We also comment

¹⁹Theorem 3.1 is also in the spirit of the theorems that relate incentive compatibility to monotonicity in types, if we equate types with posterior probabilities; see Myerson (1981) for an early theorem of that type, and Vohra (2007) for a comprehensive treatment. However, our framework is different from the standard mechanism design framework, in that we have random states of nature, so that incentive compatibility is a property of a weighted sum (expected value conditional on type), not on the value itself. As a consequence, the methodology of those papers does not work.

on the uniqueness of equilibrium under the BTS scoring rule.

4.1 PLEPS implementation

By implementing an equilibrium payoff system F , we mean designing a questionnaire and a scoring rule such that the associated game allows an equilibrium with payoffs given by F .

We continue to assume there are infinitely many respondents. We consider the case in which the respondents are asked to choose the correct answer to a multiple choice question (to declare their type), and assume that the possible states of nature take values in the set of probability distributions of the responses to the multiple choice questions.

The following result shows that any PLEPS $\{F_i\}$ such that F_i depends only on local conditional type probabilities $s_k^i = \Pr(T^\pi = k \mid \Omega = i)$ can be implemented by using F_i to compute the value of the scoring function f in state i .

Proposition 4.1 *Consider a PLEPS with type k payoff in equilibrium state i given by $F_i(z_k^i, z_{-k}^i; s_k^i, s_{-k}^i)$; that is, in addition to z_k^j 's, the payoff depends only on local conditional type probabilities $s_k^i = \Pr(T^\pi = k \mid \Omega = i)$ instead of on possibly all ex-ante type probabilities contained in vector s_k . Then, this PLEPS can be implemented by the agnostic planner. More precisely, there exist questions that the planner can ask from which she can form estimates \hat{i} , \hat{z}_k^i and \hat{s}_k^i of the true state of nature i and the true probabilities z_k^i and s_k^i , and such that, if the planner announces that a player who declares type k will receive $F_i(\hat{z}_k^i, \hat{z}_{-k}^i; \hat{s}_k^i, \hat{s}_{-k}^i)$, then, truth-telling is an equilibrium.*

Proof: Suppose the planner asks the following from the respondents:

- (a) to choose the correct answer to the multiple choice question;
- (b) to state the possible states of nature, that is, to declare what the set of the possible distributions of the responses to (i) is, AND to state their perceived probability (z 's) for each of those distributions.

To guarantee that truth-telling is an equilibrium, the planner announces she will compute the scores F_i as follows. Her estimate \hat{i} of the true state of nature i will be given by the frequencies by which each particular answer to the multiple choice question has been chosen by the respondents. She will also make the estimates \hat{s}_k^i of the type probabilities equal to those frequencies. Having estimated state i , she will then choose the corresponding \hat{z}_k^i from all the z_k^j 's, $j = 1, \dots, N$, that a player provides

as the answer to (b). Having all of these, the planner will compute the corresponding values of F_i 's, by plugging \hat{z}_k^i 's and \hat{s}_k^i 's as the arguments of function F_i .

Suppose now that all players other than player π of type k play the truth-telling strategy. If player π also plays the truth-telling strategy, his payoff in state i is $F_i(z_k^i, z_{-k}^i; s_k^i, s_{-k}^i)$, because i , z 's and s 's are correctly estimated by the planner. If player π of type k declares type $j \neq k$, his payoff in state i is $F_i(z_j^i, z_{-j}^i; s_j^i, s_{-j}^i)$, because, with infinite number of players and all except player π being honest, i , z 's and s 's are again correctly estimated by the planner. By IC inequality (2.1), player's π expected value of the payoff when he is dishonest is less than the expected value of the payoff when he is honest, and he would not deviate.

■

Remark 4.1 The above implementation procedure is not robust – in practice, there will be more different outcomes of responses to question (b) than the number of types, and different respondents will consider different distributions of the responses to (a) as the possible outcomes for the states of nature. Thus, some approximate grouping of the responses would have to be done. Moreover, responding to (b) puts a large burden on the subjects, because they have to provide possible frequencies of the responses to (a) and distributions over those frequencies. For budget-balanced logarithmic equilibrium payoffs the story is different, as discussed in the next section: the Bayesian Truth Serum (BTS) scoring rule of Prelec (2004) implements budget-balanced logarithmic equilibrium payoffs using inputs that are simpler than those obtained from the responses to (b), and a procedure which is robust (that is, no grouping of similar responses is necessary).

4.2 Implementing logarithmic equilibrium payoffs by the Bayesian Truth Serum

We first recall the definition of the Bayesian Truth Serum (BTS). We specify the model in the notation of Section 2. We assume that there are infinitely (countably) many respondents, labeled $\pi \in R$. The truthful opinion of respondent π is represented by a pair of M -tuples $(X^\pi; Y^\pi) = ((X_1^\pi, \dots, X_M^\pi); (Y_1^\pi, \dots, Y_M^\pi))$ of random variables. Here, X_i^π 's take values zero or one, and only one is equal to one. This is interpreted as choosing an answer from a set of M possible answers. Random

variables Y_i^π 's take values in $[0, 1]$ and $\sum_{i=1}^M Y_i^\pi = 1$. The latter represent the declared opinion that respondent π has on what percentage of respondents will choose i as the correct answer.

As in Section 2, we assume that the infinite sequence $(X^\pi, \pi \in R)$ is exchangeable. Then, by de Finetti's theorem, there is an M -dimensional (potentially random) vector

$$\bar{X} = \lim_n \frac{1}{n} \sum_{r=1}^n X^\pi$$

taking values in $[0, 1]^M$, such that X^π 's are conditionally independent given \bar{X} . We interpret \bar{X} to be the true state of nature, denoted previously by Ω .

Denote by \bar{x}_j the sample mean of the declared values x_j^π of X_j^π over all respondents π , and by $\log \bar{y}_j$ the sample mean of all the declared values $\log y_j^\pi$ of $\log Y_j^\pi$ (so that \bar{y}_j is their geometric mean):

$$\log \bar{y}_j := \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{r=1}^n \log y_j^\pi$$

Definition 4.1 *The Bayesian Truth Serum (BTS) score function for respondent π is given by*

$$BTS^\pi = \sum_{j=1}^M x_j^\pi \log \frac{\bar{x}_j}{\bar{y}_j} + \sum_{j=1}^M \bar{x}_j \log \frac{y_j^\pi}{\bar{x}_j}$$

Prelec (2004) proved that BTS is an incentive compatible mechanism, in the sense that a respondent's payoff is maximized by declaring the true opinion, if everyone else does. Moreover, we can state a new “uniqueness” result, namely, that with the BTS mechanism any budget-balanced strict (Bayesian) Nash equilibrium is separating.

Remark 4.2 *It is a natural convention to define $\log(\bar{x}_j/\bar{y}_j) = 0$ if $\bar{x}_j = \bar{y}_j = 0$, as well as to define $\bar{x}_j \log(y_j^\pi/\bar{x}_j) = 0$ if $\bar{x}_j = 0$. Note that if $x_j^\pi = 0$ for all but a finite number of π 's, so that $\bar{x}_j = 0$, then it is optimal for every player π to correctly predict $y_j^\pi = 0$, so that \bar{y}_j is naturally defined to be zero²⁰. Under these conventions, the only possible budget-balanced SNE's are those which are separating, and in which the players of the same type have the same strategies. Indeed:*

-*(i) First, it is impossible to have an SNE in pure strategies in which two individuals of the same type choose different strategies and hence have different expected scores: suppose they have different strategies in this SNE. If player 1 switched to strategy 2, he would have a strictly lower value, by*

²⁰This is because increasing y_j does not change the score, while decreasing y_k for $k \neq j$ would lower the score, if $\bar{x}_k > 0$.

definition of "strict", and this value would be the same as player 2's value, because with infinite number of players, the value of one player is not affected by what another player does. For the same reason, if player 2 switched to strategy 1, his value would be equal to the original player 1's value, which we argued above is strictly larger. This means that player 2 was not playing an equilibrium strategy to start with. A contradiction.

- (ii) Second, two individuals of different types cannot have the same strategies in an SNE: if they did, by (i) all other players of their types also would choose the same strategy, which means that there would be a type k that nobody would "claim", that is a k such that $x_k^\pi = 0$ for all π . Because we assume budget balance, there is a player with a non-positive score. If that player deviated to type k , by above natural conventions his BTS score would be zero, which is weakly better than not deviating, so the equilibrium could not be strict.

Because of this, and since the truth-telling equilibrium is focal among strictly separating equilibria, from now on we consider x_i 's and y_i 's to be the truthful responses.

For the reader's convenience and to provide additional details, we recall the Prelec (2004) result that, in such a truth-telling equilibrium, the BTS score is equal to the budget-balanced logarithmic payoff, and we provide a detailed proof in Appendix.

Theorem 4.1 (Prelec 2004; Theorem 2) *Under the above assumptions, when the players play the truth-telling equilibrium, BTS scoring results in budget-balanced logarithmic equilibrium payoffs. More precisely, in the equilibrium we have*

$$BTS^\pi = \log \Pr(\bar{X} = \bar{x} | X^\pi = x^\pi) - \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{s=1}^n \log \Pr(\bar{X} = \bar{x} | X^\gamma = x^\gamma) \quad (4.1)$$

or, denoting $x^\pi = k$, $x^\gamma = j$, $\bar{x} = i$,

$$BTS^\pi = \log(\Pr(\Omega = i | T^\pi = k)) - \sum_{j=1}^M \Pr(T^\pi = j | \Omega = i) \log(\Pr(\Omega = i | T^\pi = j)) \quad (4.2)$$

Thus, the BTS score corresponds to the PLEPS function F_i that is logarithmic. Put differently, BTS implements budget-balanced logarithmic equilibrium payoffs by asking the players only two things: to choose an answer from the multiple choice list, and to predict what percentage of players will choose a particular answer.

To conclude, the main message of this section is the one confirming the appeal of BTS because of the following three properties: BTS always leads to a strictly separating equilibrium in which the players of the same type get the same score, it results in the benchmark, logarithmic ex-post scores, and it is easily implementable. We know of no other mechanism that has all of these properties.

5 Conclusions

We consider the problems of extracting true opinions from a large group of respondents and of ranking them according to their posteriors on the realized state of nature (local posteriors), in the case in which the planner is agnostic about the distribution of the states of nature and the respondents' types. Thus, the planner has to design a universal mechanism, that would work for all such distributions. One such mechanism is the one that is based on ex-post logarithmic payoffs. We prove the following results for equilibrium payoffs that are determined only by the local posteriors and type probabilities: (i) under assumptions on the sensitivity of score differences, the incentive compatible budget-balanced equilibria necessarily result in logarithmic payoffs; (ii) for arbitrary mechanisms, any incentive compatible equilibrium necessarily ranks the respondents according to the relative size of their posterior probabilities of the realized state of nature. We elaborate on the result from Prelec (2004) that the logarithmic equilibrium payoffs can be implemented using the BTS algorithm, and we note that other equilibrium payoff rules can also be implemented, but may require responses to more complex questions.

Our setup does not allow for players' actions other than costless expressing of their opinions. Thus, developing a more general analysis of robust mechanisms in our framework, in which the players also would draw utility from costly actions, is an unfinished task. In our model the experts have no reason to lie, but need positive incentive to tell the truth. One could envision a framework in which players have some reason to lie, for example they do not care about their own payoff, but want to manipulate the results so as to have some other type have the highest score. Or, a framework with known utilities and unknown correlation of types, in which the planner wants to elicit information about the correlations without disturbing the stated utilities; for example, the case in which the planner wants to ask players to predict what others will do, but she doesn't want the payoff they get for making these predictions to change any of the other incentives in the game. Finally, ours is a static game, while many applications are dynamic by nature.

6 Appendix

Proof of Proposition 2.1: We want to prove that, if a player of type k mimics the equilibrium strategy of type j , then his payoff is given by the equilibrium evaluation corresponding to type j .

First, recall that a pure strategy for player π is a map $\sigma^\pi(t^\pi)$, that maps a player's type to his response choice a^π . The profile of all respondents' pure strategies is denoted $\sigma(t)$, with entries $\sigma^\pi(t^\pi)$, and the profile excluding player π is denoted $\sigma^{-\pi}(t^{-\pi})$. The score for player π is given by $f(\sigma^\pi(t^\pi), \sigma^{-\pi}(t^{-\pi}))$. Let us denote by σ the equilibrium strategy profile of all the respondents, and define ρ to be the strategy profile that is identical to σ , except that a specific player π of type $k \neq j$ plays the strategy $\sigma^\pi(j)$ corresponding to type j . Let γ denote a player of type j . Then, we have that the payoff to the mimicry strategy, when π plays j is

$$\begin{aligned} f(\rho^\pi(k), \rho^{-\pi}(T^{-\pi})) &= f(\rho^\gamma(j), \rho^{-s}(T^{-s})) \\ &= f(\sigma^\gamma(j), \rho^{-s}(T^{-s})) \\ &= f(\sigma^\gamma(j), \sigma^{-s}(T^{-s})) \end{aligned}$$

because $\sigma^{-s}(T^{-s})$ and $\rho^{-s}(T^{-s})$ differ only in π 's response, and this does not matter with infinitely many players. This is because every type will be represented by infinitely many players, and f is symmetric in their responses. More precisely, to justify the last equality above, we argue as follows: We need to prove that if exactly one of $-s$ respondents deviates from the equilibrium response, then the score f of respondent γ does not change. Since σ corresponds to a strictly separating equilibrium, the sequence $\{\sigma^{-s}(T^{-s})\}$ consists of M different K -tuples, each being repeated infinitely many times. If a respondent of type k deviates to type j , that means that one repetition (among infinitely many) of K -tuples corresponding to type k becomes an additional repetition (among infinitely many) K -tuples corresponding to type j . We can then define a permutation of the sequence $\{\sigma^{-s}(T^{-s})\}$ which is equal to the deviation sequence $\{\rho^{-s}(T^{-s})\}$, and by symmetry of f , we prove the last equality in the equation above.

■

Proof of Proposition 2.2:

Recall that we want to prove that, with finitely many players there is no budget-balanced PLEPS.

For notational simplicity, we consider the case $M = 2$ with two types only, type 1 and type 2, with at least two players, and with $N = 3$, the states of nature 1 being $(2, 0)$ (two of type 1, zero of type 2), state 2 being $(1, 1)$, and state 3 being $(0, 2)$. The proof adjusts easily to the case when $M > 2$, since we can consider only those matrices Q in which the two types are isolated in a particular block.

We consider a Z matrix of the form

$$\begin{pmatrix} p & 1-p & 0 \\ 0 & q & 1-q \end{pmatrix}$$

where $0 < p, q < 1$ (notice that the rows correspond to types and columns to states of nature).

With finitely many players, any PLEPS functions F_i would depend on the posteriors based on the state of nature i corresponding to the declared types. For example, if the true state is $(2, 0)$, but one respondent declares herself as type 2, then the payoffs correspond to state $(1, 1)$.

The expected score of the truthful response for type 1 would be

$$pF_1(p, p) + (1-p)F_2(1-p, q)$$

If one respondent lies and declares his type 1 as type 2, then the expected value would be

$$pF_2(q, 1-p) + (1-p)F_3(1-q, 1-q)$$

Therefore, the separating inequality is

$$pF_1(p, p) + (1-p)F_2(1-p, q) > pF_2(q, 1-p) + (1-p)F_3(1-q, 1-q)$$

This becomes

$$p[F_1(p, p) - F_2(q, 1-p)] + (1-p)[F_2(1-p, q) - F_3(1-q, 1-q)] > 0$$

Similarly, when one type 2 respondent lies we get the following:

$$qF_2(q, 1-p) + (1-q)F_3(1-q, 1-q) > qF_1(p, p) + (1-q)F_2(1-p, q)$$

This becomes

$$q[F_2(q, 1-p) - F_1(p, p)] + (1-q)[F_3(1-q, 1-q) - F_2(1-p, q)] > 0$$

Suppose now that $p \neq q$. Without loss of generality we consider the case $p > q$, and apply Lemma 6.1 in Appendix on the inequalities above. We obtain

$$F_1(p, p) - F_2(q, 1 - p) > 0$$

$$F_2(1 - p, q) - F_3(1 - q, 1 - q) < 0$$

Assuming budget balance holds, we must have $F_1(p, p) = 0 = F_3(1 - q, 1 - q)$ and so

$$(1 - p)F_2(1 - p, q) + qF_2(q, 1 - p) = 0$$

Note that $F_1(p, p) = 0$ leads to $F_2(q, 1 - p) < 0$, while $F_3(1 - q, 1 - q) = 0$ leads to $F_2(1 - p, q) < 0$. This is in clear contradiction with the last equality.

■

Proof of Lemma 3.1:

Recall that we need to prove that (3.8) holds.

By the standard result on optimization under constraints (in our case the constraint being $\sum_i q^i = 1$), there exists a Lagrange multiplier function $\lambda(\vec{p}, \vec{r}, s_p, s_q, s_r)$, where, for example, $\vec{p} = (p^1, \dots, p^N)$, such that

$$p^i[\partial_p F_i(p^i, p^i, r^i, s_p, s_q, s_r) - \partial_q F_i(p^i, p^i, r^i, s_p, s_q, s_r)] = \lambda(\vec{p}, \vec{r}, s_p, s_q, s_r) \quad (6.1)$$

Fix an arbitrary value of i and p^i, r^i . Since $N > 2$, we can set $p^j = x, r^j = y$, for a fixed, but arbitrary $j \neq i$, for any $0 < x < 1 - p^i, 0 < y < 1 - r^i$. By the above equality we have that $\lambda(\vec{p}, \vec{r}, S)$ is a function $\lambda(p^i, r^i, S)$ of p^i, r^i, S , only, and we have

$$x[\partial_p F_j(x, x, y, S) - \partial_q F_j(x, x, y, S)] = \lambda(p^i, r^i, S),$$

for all $0 < x < 1 - p^i, 0 < y < 1 - r^i$. Since we can choose p^i, r^i arbitrarily small, we have then, for fixed S , that the left-hand side is constant across all values of x, y in $(0, 1)$, and because i is arbitrary we get that $\lambda(S)$ does not depend on any of the values $p^i, r^i, i = 1, \dots, N$.

■

Proof of Proposition 3.1:

We suppress dependence on i in this proof, and on s_p, s_q, s_r . We want to show that

$$F(p, q, r) = G(p) + H(p, q, r)$$

where H is symmetric in all the pairs $(p, s_p), (q, s_q), (r^j, s_{r^j})$.

For $p^0 \in (0, 1)$ denote

$$\bar{p} = p - p^0, \bar{q} = q - p^0, \bar{r} = r - p^0$$

From the smoothness and the symmetry property of F , we can write, for some functions a, b, c, d, e of the type probabilities, by Taylor's expansion,

$$F(p, q, r) = \sum_{n=0}^{\infty} a_n \bar{p}^n + \sum_{n=1}^{\infty} (b_n^q \bar{q}^n + b_n^{\pi} \bar{r}^n) + \sum_{m,n=1}^{\infty} \bar{p}^m (c_{m,n}^q \bar{q}^n + c_{m,n}^{\pi} \bar{r}^n) + \sum_{m,n=1}^{\infty} d_{m,n} \bar{q}^m \bar{r}^n + \sum_{l,m,n=1}^{\infty} e_{l,m,n} \bar{p}^l \bar{q}^m \bar{r}^n$$

where, by the symmetry property,

$$b_n^q(s_p, s_q, s_r) = b_n^{\pi}(s_p, s_r, s_q), \quad c_{m,n}^q(s_p, s_q, s_r) = c_{m,n}^{\pi}(s_p, s_r, s_q)$$

$$d_{m,n}(s_p, s_q, s_r) = d_{n,m}(s_p, s_r, s_q), \quad e_{l,m,n}(s_p, s_q, s_r) = e_{l,n,m}(s_p, s_r, s_q)$$

Note that it is sufficient to show that

$$c_{m,n}^{\pi} = d_{m,n}, \quad e_{l,m,n} = e_{m,l,n}$$

because then we can write

$$F(p, q, r) = \sum_{n=0}^{\infty} [a_n - b_n^q] \bar{p}^n + H(p, q, r)$$

where H is symmetric in all the pairs $(p^i, s_p), (q^i, s_q), (r_j^i, s_{r_j})$.

Let us consider the consequences of strict separation inequality (3.7), using Lemma 3.1. We have

$$\begin{aligned} & \partial_q F(p, p, r) - \partial_p F(p, p, r) \\ &= \sum_{n=1}^{\infty} n b_n^q \bar{p}^{n-1} + \sum_{m,n=1}^{\infty} c_{m,n}^q n \bar{p}^{m+n-1} + \sum_{m,n=1}^{\infty} d_{m,n} m \bar{p}^{m-1} \bar{r}^n + \sum_{l,m,n=1}^{\infty} e_{l,m,n} m \bar{p}^{l+m-1} \bar{r}^n \\ & \quad - \sum_{n=0}^{\infty} n a_n \bar{p}^{n-1} - \sum_{m,n=1}^{\infty} m \bar{p}^{m-1} (c_{m,n}^q \bar{p}^n + c_{m,n}^{\pi} \bar{r}^n) - \sum_{l,m,n=1}^{\infty} e_{l,m,n} l \bar{p}^{l+m-1} \bar{r}^n \end{aligned}$$

We can then write

$$p \partial_q F(p, p, r) - p \partial_p F(p, p, r)$$

$$\begin{aligned}
&= \sum_{n=1}^{\infty} nb_n^q \bar{p}^n + \sum_{m,n=1}^{\infty} c_{m,n}^q n \bar{p}^{m+n} + \sum_{m,n=1}^{\infty} d_{m,n} m \bar{p}^m \bar{r}^n + \sum_{l,m,n=1}^{\infty} e_{l,m,n} m \bar{p}^{l+m} \bar{r}^n \\
&\quad - \sum_{n=0}^{\infty} n a_n \bar{p}^n - \sum_{m,n=1}^{\infty} m \bar{p}^m (c_{m,n}^q \bar{p}^n + c_{m,n}^{\pi} \bar{r}^n) - \sum_{l,m,n=1}^{\infty} e_{l,m,n} l \bar{p}^{l+m} \bar{r}^n \\
&+ \sum_{n=1}^{\infty} nb_n^q p^0 \bar{p}^{n-1} + \sum_{m,n=1}^{\infty} n c_{m,n}^q p^0 \bar{p}^{m+n-1} + \sum_{m,n=1}^{\infty} d_{m,n} m p^0 \bar{p}^{m-1} \bar{r}^n + \sum_{l,m,n=1}^{\infty} e_{l,m,n} m p^0 \bar{p}^{l+m-1} \bar{r}^n \\
&\quad - \sum_{n=0}^{\infty} n a_n p^0 \bar{p}^{n-1} - \sum_{m,n=1}^{\infty} m p^0 \bar{p}^{m-1} (c_{m,n}^q \bar{p}^n + c_{m,n}^{\pi} \bar{r}^n) - \sum_{l,m,n=1}^{\infty} e_{l,m,n} l p^0 \bar{p}^{l+m-1} \bar{r}^n
\end{aligned}$$

By Lemma 3.1, to have a PLEPS this has to be equal to $(-\lambda)$ for all p, r , which is possible only if

- from \bar{r}^n terms:

$$c_{1,n}^{\pi} = d_{1,n} \quad (6.2)$$

- from $\bar{p}\bar{r}^n$ terms:

$$0 = c_{1,n}^{\pi} - d_{1,n} + c_{2,n}^{\pi} - d_{2,n} \quad (6.3)$$

- from $\bar{p}^2\bar{r}^n$ terms:

$$0 = 2(d_{2,n} - c_{2,n}^{\pi}) + 3p^0(d_{3,n} - c_{3,n}^{\pi}) + p^0(e_{1,2,n} - e_{2,1,n}) \quad (6.4)$$

- from $\bar{p}^3\bar{r}^n$ terms:

$$0 = 3(d_{3,n} - c_{3,n}^{\pi}) + (e_{1,2,n} - e_{2,1,n}) + 4p^0(d_{4,n} - c_{4,n}^{\pi}) + 2p^0(e_{1,3,n} - e_{3,1,n}) \quad (6.5)$$

And so on.

So, it is sufficient to show $e_{l,m,n} = e_{m,l,n}$. This follows directly from Assumption 3.1, because then the third mixed derivative of the difference $F(p, q, r) - F(q, p, r)$ in scores is zero for all p, q, r , that is,

$$0 = \sum_{l,m,n=1}^{\infty} l m n e_{l,m,n} \bar{p}^{l-1} \bar{q}^{m-1} \bar{r}^{n-1} - \sum_{l,m,n=1}^{\infty} l m n e_{l,m,n} \bar{q}^{l-1} \bar{p}^{m-1} \bar{r}^{n-1}$$

This completes the proof. ■

Then following lemma is the key ingredient in proving Theorem 3.2. It is a slight extension of Lemma A.1 in Schervish (1989).

Lemma 6.1 (Schervish 1989). *Let $0 < a \leq 1$, $p, q \in (0, a)$, and $p > q$. If A, B are real numbers such that*

$$pA + (a - p)B > 0$$

$$q(-A) + (a - q)(-B) > 0$$

then $A > 0$ and $B < 0$.

Proof:

Notice that $A \neq 0$. If not, then the two above inequalities become $(a - p)B > 0$ and $(a - q)(-B) > 0$, a contradiction. In order to prove the lemma we only need to prove that $A > 0$. Suppose to the contrary that $A < 0$. Then $B > 0$. From $(a - p)B > -pA$ it follows that $B > -\frac{p}{a-p}A > 0$. We then get $0 < q(-A) + (a - q)(-B) < q(-A) + \frac{a-q}{a-p}pA = Aa\frac{p-q}{a-p} < 0$, which is impossible. This contradiction proves $A > 0$.

■

Proof of Theorem 3.2:

Recall that we want to prove that PLEPS payoffs are strictly increasing in the posterior probabilities of the true state of nature.

We suppress the dependence on s_k 's in our notation. This is justified because fixing s_k 's does not restrict the choice of any two rows of the Z -matrix, because we can always define Q by $q_k^i = z_k^i s_k$.

We consider three cases separately according to the values of M and N .

Case 1: Assume $M = 2, N = 2$. The matrix Z can be written then as $\begin{bmatrix} z_1^1 & z_1^2 \\ z_2^1 & z_2^2 \end{bmatrix}$. If we denote $p := z_1^1, q := z_2^1$, then the matrix Z becomes $Z = \begin{bmatrix} p & 1-p \\ q & 1-q \end{bmatrix}$. Suppose $p > q$ (which is equivalent to $1-q > 1-p$). The IC property (2.1) of F_i implies

$$pF_1(p, q) + (1-p)F_2(1-p, 1-q) > pF_1(q, p) + (1-p)F_2(1-q, 1-p) \text{ and}$$

$$qF_1(q, p) + (1-q)F_2(1-q, 1-p) > qF_1(p, q) + (1-q)F_2(1-p, 1-q).$$

This leads to

$$p[F_1(p, q) - F_1(q, p)] + (1 - p)[F_2(1 - p, 1 - q) - F_2(1 - q, 1 - p)] > 0 \text{ and}$$

$$q[F_1(q, p) - F_1(p, q)] + (1 - q)[F_2(1 - q, 1 - p) - F_2(1 - p, 1 - q)] > 0.$$

We set $a = 1$, $A = F_1(p, q) - F_1(q, p)$ and $B = F_2(1 - p, 1 - q) - F_2(1 - q, 1 - p)$, and apply Lemma 6.1 in Appendix to the above equations. We obtain that $F_1(p, q) > F_1(q, p)$ and $F_2(1 - p, 1 - q) > F_2(1 - q, 1 - p)$, which proves the theorem in this case.

$$\begin{bmatrix} z_1^1 & z_1^2 \\ z_2^1 & z_2^2 \\ \vdots & \vdots \\ z_M^1 & z_M^2 \end{bmatrix}.$$

Case 2: Assume $M \geq 3, N = 2$. The matrix Z can be written as

The matrix entries satisfy $z_k^2 = 1 - z_k^1, k = 1, \dots, M$. Take any $k, j \in \{1, \dots, M\}$ such that $z_k^1 > z_j^1$ (which is equivalent to $z_j^2 > z_k^2$). Using the notation $p := z_k^1, q := z_j^1$, and the notation $z_{-j,k}^i$ for the $(N - 2)$ -tuple which consists of $\{z_1^i, \dots, z_M^i\} \setminus \{z_j^i, z_k^i\}$, from (2.1) we obtain that the following two equations hold:

$$\begin{aligned} pF_1(p, q, z_{-(j,k)}^1) + (1 - p)F_2(1 - p, 1 - q, z_{-(j,k)}^2) &> pF_1(q, p, z_{-(j,k)}^1) + (1 - p)F_2(1 - q, 1 - p, z_{-(j,k)}^2) \\ qF_1(q, p, z_{-(j,k)}^1) + (1 - q)F_2(1 - q, 1 - p, z_{-(j,k)}^2) &> qF_1(p, q, z_{-(j,k)}^1) + (1 - q)F_2(1 - p, 1 - q, z_{-(j,k)}^2). \end{aligned}$$

Hence, if we define A and B in the following way,

$$A = F_1(p, q, z_{-(j,k)}^1) - F_1(q, p, z_{-(j,k)}^1) = F_1(z_k^1, z_{-k}^1) - F_1(z_j^1, z_{-j}^1)$$

and

$$B = F_2(1 - p, 1 - q, z_{-(j,k)}^2) - F_2(1 - q, 1 - p, z_{-(j,k)}^2) = F_2(z_k^2, z_{-k}^2) - F_2(z_j^2, z_{-j}^2)$$

we are again within the framework of Lemma 6.1, and we conclude that $A > 0$ and $B < 0$, which proves (3.9) for both $i = 1$ and $i = 2$.

Case 3:

Assume $M \geq 2, N \geq 3$ and denote $Z = \begin{bmatrix} z_1^1 & z_1^2 & \dots & z_1^N \\ z_2^1 & z_2^2 & \dots & z_2^N \\ \vdots & \vdots & \vdots & \vdots \\ z_M^1 & z_M^2 & \dots & z_M^N \end{bmatrix}$. We will prove (3.9) for the i -th

column in the matrix Z . Choose any rows (types) $j, k \in \{1, \dots, M\}$ where $j \neq k$. Since the only requirement for the matrix Z is that its rows are non-degenerate probability distributions, and since the values of F_i depend only on the quantities in the i -th column, then in order to complete the proof

we need to prove only that for every $p := z_k^i$ and $q := z_j^i$, with $1 > p > q > 0$, for any choice of $z_{-(k,j)}^i \in (0, 1)^{M-2}$ (if $M = 2$ this last requirement is unnecessary), we have

$$F_i(p, q, z_{-(k,j)}^i) > F_i(q, p, z_{-(k,j)}^i) \quad (6.6)$$

To use (2.1) and apply Lemma 6.1 again, the idea now is to suitably modify the matrix Z , without changing its i -th column, by finding a matrix \tilde{Q} with the same type probabilities s_ℓ as the original matrix Q . As mentioned above, we can always do that by choosing $q_{\ell i} = z_\ell^i s_\ell$. Moreover, looking at the proof from Case 1 and Case 2, we see that we apply Lemma 6.1 to the payoff differences in two different states, corresponding to two different columns of matrix Z . Without loss of generality, we assume $i \neq 1$, and we change that column in a specific way. More precisely, instead of working with the original matrix Z , we work with the following Z -matrix, in which the i -th column is not changed:

Take $0 < \epsilon < 1$ and $a := 1 - \epsilon$; we work with the matrix \tilde{Z} given by

$$\tilde{Z}_l^t := \begin{cases} z_l^t, & \text{if } l \in \{1, \dots, M\} \setminus \{j, k\} \\ p, & \text{if } l = k, t = i \\ q, & \text{if } l = j, t = i \\ a - p, & \text{if } l = k, t = 1 \\ a - q, & \text{if } l = j, t = 1 \\ \frac{\epsilon}{N-2}, & \text{otherwise} \end{cases}$$

where p, q are arbitrary values in $(0, a)$ with $p > q$. Then for every choice of ϵ and p and q , we have that \tilde{Z} is a Z -matrix which differs from Z only in the j -th and k -th row and these rows are

$$\begin{bmatrix} a - p & \frac{\epsilon}{N-2} & \cdots & \frac{\epsilon}{N-2} & p & \frac{\epsilon}{N-2} & \cdots & \frac{\epsilon}{N-2} \\ a - q & \frac{\epsilon}{N-2} & \cdots & \frac{\epsilon}{N-2} & q & \frac{\epsilon}{N-2} & \cdots & \frac{\epsilon}{N-2} \end{bmatrix}$$

The IC property (2.1) applied now to j and k yields

$$\sum_{t=1}^N \tilde{z}_k^t F_t(\tilde{z}_k^t, \tilde{z}_{-k}^t) > \sum_{t=1}^N \tilde{z}_k^t F_t(\tilde{z}_j^t, \tilde{z}_{-j}^t)$$

$$\sum_{t=1}^N \tilde{z}_j^t F_t(\tilde{z}_j^t, \tilde{z}_{-j}^t) > \sum_{t=1}^N \tilde{z}_j^t F_t(\tilde{z}_k^t, \tilde{z}_{-k}^t)$$

Observe that for $t \in \{1, \dots, N\} \setminus \{1, i\}$ we have $\tilde{z}_j^t = \tilde{z}_k^t = \frac{\epsilon}{N-2}$. Hence on both sides of the above inequalities we have terms $\frac{\epsilon}{N-2} F_t(\frac{\epsilon}{N-2}, \frac{\epsilon}{N-2}, z_{-(j,k)}^t)$ and they cancel each other. Therefore, the inequalities take the following form:

$$(a-p)F_1(a-p, a-q, z_{-(j,k)}^1) + pF_i(p, q, z_{-(j,k)}^i) > (a-p)F_1(a-q, a-p, z_{-(j,k)}^1) + pF_i(q, p, z_{-(j,k)}^i)$$

$$(a-q)F_1(a-q, a-p, z_{-(j,k)}^1) + qF_i(q, p, z_{-(j,k)}^i) > (a-q)F_1(a-p, a-q, z_{-(j,k)}^1) + qF_i(p, q, z_{-(j,k)}^i)$$

If we set A and B as $A = F_i(p, q, z_{-(j,k)}^i) - F_i(q, p, z_{-(j,k)}^i)$ and $B = F_1(a-p, a-q, z_{-(j,k)}^1) - F_1(a-q, a-p, z_{-(j,k)}^1)$, we are again within the framework of Lemma 6.1. Therefore $A > 0$, which proves inequality (6.6) for $a > p > q > 0$. By letting $\epsilon \rightarrow 0$, we obtain (6.6) for $1 > p > q > 0$.

■

Proof of Theorem 4.1:

We want to derive the representation of the BTS score in terms of the logarithms of the local posteriors.

Let us denote

$$p_{ij} = \Pr(X_i^\pi = 1, X_j^\gamma = 1)$$

where we use the fact that, by exchangeability, the right-hand side does not depend on the choice of $\pi \neq \gamma$. Thus, we also have

$$\Pr(X^\pi = x^\pi | X^\gamma = x^\gamma) = \frac{p_{ij}}{\sum_{k=1}^M p_{kj}} \quad (6.7)$$

We will need the following three properties.

- Property I: $y_j^\pi = \sum_{i=1}^M x_i^\pi \frac{p_{ij}}{\sum_{k=1}^M p_{ki}}$
- Property II: $\log \Pr(X^\gamma = x^\gamma | X^\pi = x^\pi) = \sum_{j=1}^M x_j^\gamma \log y_j^\pi$, where conditioning indicates conditioning on the truthful response, hence on the signal.
- Property III: $\log \Pr(X^\pi = x^\pi | \bar{X} = \bar{x}) = \sum_{k=1}^M x_k^\pi \log \bar{x}_k$.

Property I is assumed because we assume a Bayesian game: the respondents compute conditional probabilities in a Bayesian fashion. Property II is a consequence of Property I and equation (6.7). For Property III, let ℓ be such that $x_\ell^\pi = 1$. De Finetti's theorem implies

$$Pr(X^\pi = x^\pi | \bar{X} = \bar{x}) = \bar{x}_\ell = \sum_{k=1}^M x_k^\pi \bar{x}_k .$$

The sum on the right always has only one term different from zero. Therefore, taking the log implies Property III.

Next, let x^γ be any values such that

$$\bar{x}_k = \lim_n \frac{1}{n} \sum_s x_k^\gamma$$

Note that we can use exchangeability to reorder the respondents so that $\pi = 1$ and $\gamma = 2, \dots, n+1$. For such choice of π and γ we have $\log Pr(X^\pi = x^\pi | X^\gamma = x^\gamma) = \sum_{j=1}^M x_j^\pi \log(y_j^\gamma)$. We may always omit those γ such that $Pr(X^\gamma = x^\gamma) = 0$. Thus, we actually have only finitely many choices for an M -tuple x^γ such that $0 < Pr(X^\gamma = x^\gamma) < 1$, and there is a lower bound A and an upper bound B such that $0 < A \leq Pr(X^\gamma = x^\gamma) \leq B < 1$. Then it follows that $A = \sqrt[n]{A^n} \leq \sqrt[n]{\prod_{s=1}^n Pr(X^\gamma = x^\gamma)} \leq \sqrt[n]{B^n} = B$. The log function is continuous, so $\log(\lim(f)) = \lim(\log(f))$ as long as f and $\lim(f)$ are both finite and strictly positive. We conclude that the limit $\lim_{n \rightarrow \infty} \prod_{s=1}^n Pr(X^\gamma = x^\gamma)$ exists, that it is not zero, and that we can take the log outside or inside the limit.

Next, using the above conclusion, from Properties I-III we get

$$\sum_{k=1}^M \bar{x}_k \log y_k^\pi = \lim_n \frac{1}{n} \sum_s \log Pr(X^\gamma = x^\gamma | X^\pi = x^\pi)$$

and

$$\sum_{k=1}^M x_k^\pi \log \bar{y}_k = \lim_n \frac{1}{n} \sum_s \log Pr(X^\pi = x^\pi | X^\gamma = x^\gamma)$$

and so, using Bayes rule,

$$\begin{aligned}
BTS^\pi &= \sum_{k=1}^M x_k^\pi \log \frac{\bar{x}_k}{\bar{y}_k} + \sum_{k=1}^M \bar{x}_k \log y_k^\pi \\
&= \log Pr(X^\pi = x^\pi | \bar{X} = \bar{x}) + \lim_n \frac{1}{n} \sum_s \log Pr(X^\gamma = x^\gamma | X^\pi = x^\pi) - \lim_n \frac{1}{n} \sum_s \log Pr(X^\pi = x^\pi | X^\gamma = x^\gamma) \\
&= \log \left(Pr(X^\pi = x^\pi | \bar{X} = \bar{x}) \lim_n \prod_{s=1}^n \frac{Pr^{1/n}(X^\gamma = x^\gamma | X^\pi = x^\pi)}{Pr^{1/n}(X^\pi = x^\pi | X^\gamma = x^\gamma)} \right) \\
&= \log \left(Pr(X^\pi = x^\pi | \bar{X} = \bar{x}) \frac{\lim_n \prod_{s=1}^n Pr^{1/n}(X^\gamma = x^\gamma)}{Pr(X^\pi = x^\pi)} \right) \\
&= \log Pr(\bar{X} = \bar{x} | X^\pi = x^\pi) - \log Pr(\bar{X} = \bar{x}) + \lim_n \frac{1}{n} \sum_s \log Pr(X^\gamma = x^\gamma)
\end{aligned}$$

Since the last two terms do not depend on π , and $\sum_r BTS^\pi = 0$, we get equation (4.1). Next, for fixed n and \bar{x} , denote by n_j the number of respondents who have type j , so that

$$\sum_j n_j = n$$

Then we can write equation (4.1) as

$$\begin{aligned}
BTS^\pi &= \log Pr(\bar{X} = \bar{x} | X^\pi = x^\pi) \\
&\quad - \lim_{n \rightarrow \infty} \frac{1}{n} \left[\sum_{s=1}^{n_1} \log Pr(\bar{X} = \bar{x} | X^\pi = x^1) + \dots + \sum_{s=n_{M-1}+1}^{n_M} \log Pr(\bar{X} = \bar{x} | X^\pi = x^M) \right] \\
&= \log Pr(\bar{X} = \bar{x} | X^\pi = x^\pi) - \lim_{n \rightarrow \infty} \left[\frac{n_1}{n} \log Pr(\bar{X} = \bar{x} | X^\pi = x^1) + \dots + \frac{n_M}{n} \log Pr(\bar{X} = \bar{x} | X^\pi = x^M) \right]
\end{aligned}$$

Since

$$\lim_{n \rightarrow \infty} \frac{n_j}{n} = Pr(T^\pi = j | \bar{X} = \bar{x})$$

we prove equation (4.2).

■

References

[1] Aldous, D.J. (1985) "Exchangeability and related topics". École d'été de probabilités de Saint-Flour, XIII - 1983, 1 - 198, Lecture Notes in Mathematics, 1117, Springer, Berlin.

[2] Bergemann, D. and Morris S. (2012) An Introduction to Robust Mechanism Design. In “Foundations and Trends in Microeconomics” 8, 169–230.

[3] Bernardo, J.M. (1979) ”Expected information as expected utility”. *The Annals of Statistics* 7, 686–690.

[4] Börgers, T. (with a chapter by Krähmer,D. and Strausz, R.) (2013) “An Introduction to the Theory of Mechanism Design”. Lecture Notes, U. of Michigan.

[5] Chow, Y.S. and Teicher, H. (1997) ”Probability Theory. Independence, Interchangeability, Martingales”. 3rd ed. Springer Texts in Statistics. Springer-Verlag, New York.

[6] Cvitanić, J., Prelec, D., Radas, S. and , Šikić, H. (2019) Game of Duels: Information-Theoretic Axiomatization of Scoring Rules. *IEEE Transactions on Information Theory*, 65, 530–537.

[7] Cvitanić, J., Prelec, D., Riley, B. and Tereick, B. (2018) Honesty via Choice-Matching. *American Economic Review: Insights*. Forthcoming.

[8] Dasgupta, A., and Ghosh, A. (2013) Crowdsourced Judgement Elicitation with Endogenous Proficiency. In Proceedings of the 22nd ACM International World Wide Web Conference (WWW13). 319–330.

[9] Feller, W. (1971). ”An Introduction to Probabilty Theory and its Applications”. Vol II, 2nd ed. John Wiley and Sons, Inc., New York.

[10] Frongillo, R. and Witkowski, J. (2017) A Geometric Perspective on Minimal Peer Prediction, ACM Transactions on Economics and Computation, Vol. V, No. N, Article A.

[11] Gneiting, T. and Raftery, A.E. (2007) Strictly Proper Scoring Rules, Prediction, and Estimation. *Journal of the American Statistical Association* 102, 359–378.

[12] Howie, P.J., Wang, Y. and Tsai, J. (2011) Predicting new product adoption using Bayesian truth serum, *Journal of Medical Marketing*, Vol. 11, Issue 1, 6-16.

[13] Jurca, R. and Faltungs, B. (2007) Robust Incentive-Compatible Feedback Payments. In Trust, Reputation and Security: Theories and Practice. LNAI, Vol. 4452. Springer-Verlag, 204–218.

[14] Kong, Y. and Schoenebeck, G. (2019) An Information Theoretic Framework For Designing Information Elicitation Mechanisms That Reward Truth-telling. *The ACM Transactions on Economics and Computation (TEAC)*, forthcoming.

[15] Kukla-Gryz, A., Tyrowicz, J., Krawczyk, M. and Siwinski, K. (2015) We All Do It, but Are We Willing to Admit? Incentivizing Digital Pirates' Confessions, *Applied Economics Letters*, v. 22, iss. 1-3, 184-88.

[16] Liu, Y. and Chen, Y. (2019) Surrogate Scoring Rules and a Uniform Dominant Truth Serum. Preprint, <https://arxiv.org/pdf/1802.09158>.

[17] Loughran, T., Paternoster, R. and Thomas, K. (2014) Incentivizing Responses to Self-report Questions in Perceptual Deterrence Studies: An Investigation of the Validity of Deterrence Theory Using Bayesian Truth Serum. *Journal of Quantitative Criminology*, Vol. 30, Issue 4, 677-707.

[18] Maskin, E. and Sjöström, T. (2002) Implementation theory. In “Handbook of Social Choice and Welfare”; K. J. Arrow, A. K. Sen and K. Suzumura (ed.) 1, 237–288, Elsevier.

[19] McCarthy, J. (1956) Measures of the value of information. *Proc Natl Acad Sci USA* 42, 654–655.

[20] Miller, S.R., Bailey, B.P. and Kirlik, A. (2014) Exploring the Utility of Bayesian Truth Serum for Assessing Design Knowledge, *Human-Computer Interaction*, vol.29, no.5-6, 487-515.

[21] Miller, N., Resnick, P. and Zeckhauser, R. (2005) Eliciting Informative Feedback: The Peer-Prediction Method. *Management Science* 51, 1359–1373.

[22] Myerson, R. B. (1981) Optimal Auction Design *Mathematics of Operations Research* 6, 58–73.

[23] Offerman, T., Sonnemans, J., Van De Kuilen, G. and Wakker, P.P. (2009) A Truth Serum for Non-Bayesians: Correcting Proper Scoring Rules for Risk Attitudes. *Review of Economic Studies* 76, 1461–1489.

[24] Prelec, D. (2004) A Bayesian Truth Serum for Subjective Data. *Science* 306, 462–466.

[25] Prelec, D., Seung, H.S. and McCoy, J. (2017) A solution to the single-question crowd wisdom problem. *Nature* 541, 532-535.

[26] Radanovic, G. and Faltings, B. (2013) A Robust Bayesian Truth Serum for Non-binary Signals. In Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI13). 833839.

[27] Savage, L.J. (1971) Elicitation of Personal Probabilities and Expectations. *Journal of the American Statistical Association* 66, 783–801

[28] Schervish, M. (1989) A General Method for Comparing Probability Assessors. *The Annals of Statistics* 17, 1856–1879.

[29] Schlag, K.H. and van der Weele, J.J. (2013) Eliciting Probabilities, Means, Medians, Variances and Covariances without assuming Risk Neutrality. *Theoretical Economics Letters* 2, 38–42.

[30] Vohra, R.V. (2007) Paths, Cycles and Mechanism Design. Preprint.

[31] Waggoner, B., and Chen, Y. (2013) Information Elicitation Sans Verification. In Proceedings of the 3rd Workshop on Social Computing and User Generated Content (SC13).

[32] Witkowski., J. (2014.) Robust Peer Prediction Mechanisms. Ph.D. Dissertation. Department of Computer Science, Albert-Ludwigs-Universitat Freiburg.

[33] Witkowski, J., and Parkes, D. (2012) A Robust Bayesian Truth Serum for Small Populations. In Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI12). 1492–1498.

[34] Witkowski, J., and Parkes, D. (2013) Learning the Prior in Minimal Peer Prediction. In Proceedings of the 3rd Workshop on Social Computing and User Generated Content (SC13)