

Information Elicitation Mechanisms for Statistical Estimation

Yuqing Kong,¹ Grant Schoenebeck,² Biaoshuai Tao,² Fang-Yi Yu^{2*}

¹CFCS, Computer Science Dept., Peking University, ²University of Michigan, Ann Arbor
yuqkong.kong@pku.edu.cn, {schoeneb, bstao, fayu}@umich.edu

Abstract

We study learning statistical properties from strategic agents with private information. In this problem, agents must be incentivized to truthfully reveal their information even when it cannot be directly verified. Moreover, the information reported by the agents must be aggregated into a statistical estimate. We study two fundamental statistical properties: estimating the mean of an unknown Gaussian, and linear regression with Gaussian error. The information of each agent is one point in a Euclidean space.

Our main results are two mechanisms for each of these problems which optimally aggregate the information of agents in the truth-telling equilibrium:

- A minimal (non-revelation) mechanism for large populations — agents only need to report one value, but that value need not be their point.
- A mechanism for small populations that is non-minimal — agents need to answer more than one question.

These mechanisms are “informed truthful” mechanisms where reporting unaltered data (truth-telling) 1) forms a strict Bayesian Nash equilibrium and 2) has strictly higher welfare than any oblivious equilibrium where agents’ strategies are independent of their private signals. We also show a minimal revelation mechanism (each agent only reports her signal) for a restricted setting and use an impossibility result to prove the necessity of this restriction.

We build upon the peer prediction literature in the single-question setting; however, most previous work in this area focuses on discrete signals, whereas our setting is inherently continuous, and we further simplify the agents’ reports.

1 Introduction

Traditional statistical estimation approaches assume inputs are given and produce an output. However, increasingly, inputs must be obtained by eliciting information from a diverse set of users. In this new environment, gathering inputs can be at least as important and difficult as the computation itself. In many settings, users must be rewarded for their

participation, and this is especially important when a representative sample is desired. However, rewarding users can create perverse incentives that lead to inaccurate reports, especially when the answers cannot be verified by the system. For example, agents may agree to participate due to the flat fee but then not take the time to report accurately or hide information due to privacy concerns.

This work develops systems to facilitate accurate statistical estimates by rewarding honest reporting—even when the information cannot be directly verified (peer prediction)—and then aggregating the results. In fact, we will see that often the aggregation and the reward go hand-in-hand: accurate rewards are produced by aggregating the information of other agents, but the truthfulness of the other agents relies on these accurate rewards. This creates a certain “full-pipeline” solution that integrates information elicitation and aggregation into a single process.

We consider two fundamental statistical estimation problems: mean estimation and linear regression. In the mean estimation problem, the signal space is \mathbb{R}^d . Each agent i can choose to access a signal $s_i \in \mathbb{R}^d$ which is drawn from an unknown Gaussian distribution. The mean of the Gaussian is unknown to the agents but believed to be drawn from a commonly known Gaussian prior distribution known to the agents, but not the mechanism. The mechanism’s goal is to estimate the mean of the distribution by motivating each agent i to report some information about s_i truthfully.

For the linear regression problem, each agent i has a feature/attribute vector $\phi(x_i) \in \mathbb{R}^{d+1}$ of a point x_i which is publicly known, and has access to a private signal $y_i \in \mathbb{R}$. The y_i point is from an unknown linear function applied to $\phi(x_i)$ plus some Gaussian error. While the linear function is not known, it is selected from some common prior. The mechanism’s goal is to estimate the linear function by motivating each agent i to report some information about y_i truthfully. Due to the space limit, all our results for the linear regression problem are deferred to the full version. However, the key technical ideas are all present in the mean estimation problem.

Challenges for continuous signal spaces The signal spaces in the above two problems are both continuous,

*Grant Schoenebeck, Biaoshuai Tao, and Fang-Yi Yu are pleased to acknowledge the support of the National Science Foundation AitF #1535912 and CAREER #1452915.
Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

which is different from most of the existing work where signal spaces are discrete. A key technical difficulty for continuous signals is to reduce the report space of each agent. Even if one could directly port the prior discrete mechanisms to the continuous setting, the resulting mechanisms would require agents to report their *posterior of others signals*. In our Gaussian setting, this would require agents to report the covariance matrices of their prediction of others’ signals. We do not believe it is realistic to ask agents for such information. For example, not everyone knows, or even has intuition about, covariance matrices.

In contrast, our mechanisms only ask agents for their signals or their expectation of another agent’s signal.

1.1 Our Results

Mechanism design goals We assume that the agents share a common prior, i.e. a common joint prior distribution over agents’ private signals, which is a typical assumption when agents are only assigned one task (Miller, Resnick, and Zeckhauser 2005; Prelec 2004). We seek to design information elicitation mechanisms that satisfy a number of useful properties including:

Informed truthful: (Shnayder et al. 2016) Truth-telling is a strict Bayesian Nash equilibrium which has the highest welfare (expected sum of agent payments) among all equilibria and a strictly higher welfare than any oblivious equilibrium where agents’ strategies are independent of their private signals.

Minimal: A mechanism is minimal if agents only need to report one value from the signal space. Additionally, we say a mechanism is a **minimal revelation** mechanism if the value it requests from each agent is the agent’s private information. If the one question the mechanism asks each agent may be about a different value (e.g. their best prediction for a peer’s value), we call it a **minimal non-revelation** mechanism.

Small population compatible: The mechanism works for a constant number of participants. In a large-group mechanism, we assume that the number of participants goes to infinity.

Detail-free: The mechanism need not know the common-prior of the agents.

Informative prior compatible: Agents may have some prior knowledge about the distribution of signals. In this case, the mechanism must be robust to an agent reporting a mixture of its signal and its prior when it is trying to predict the most likely signal of a peer. A common example of an “uninformative” prior is a Gaussian distribution with an infinite variance (e.g., a Jeffery prior defined in Definition 2.2).

Our mechanisms Our main result is to develop two informed truthful mechanisms for both mean estimation and linear regression which optimally aggregate the information of agents to estimate the statistical property in the truth-telling equilibrium: *Proxy BTS mechanism* and *Disagreement mechanism*. Our two mechanisms are based on

the same intuitions of the BTS mechanism (Prelec 2004)—reward surprisingly common signals—and the Disagreement mechanism (Kong and Schoenebeck 2018a; Radanovic and Faltings 2014)—which punishes agents for giving forecasts that appear to disagree with their reported signals. We also ask if there exists a minimal revelation mechanism. In general, we show that the answer is no. However, for the special case where agents’ knowledge of peer’s private information comes solely from their private information (not from the prior), we present the *Metric Mechanism*, a minimal revelation mechanism for small populations.

We summarize the properties of our mechanisms below:

mechanism	Metric	Proxy BTS	Disagreement
strict BNE	✓	✓	✓
informed truthful	✓	✓	✓
informative prior		✓	✓
small population	✓		✓
minimal	✓	✓	

Our mechanisms simplify agents’ reports by only collecting signals and/or posterior expectations without collecting the covariance matrices.

1.2 Related Work

Peer prediction Miller, Resnick, and Zeckhauser (2005) introduce the peer prediction mechanism which is the first mechanism that has truth-telling as a strict Bayesian Nash equilibrium and does not need verification. However, their mechanism requires the full knowledge of the common prior and there exist some equilibria that pay more than truth-telling. In particular, the oblivious equilibrium pays strictly more than truth-telling. Kong, Ligett, and Schoenebeck (2016) modify the original peer prediction mechanism such that truth-telling pays strictly better than any other equilibrium but still requires the full knowledge of the common prior. Prelec (2004) designs the first detail-free peer prediction mechanism—Bayesian truth serum (BTS). Moreover, BTS is informed truthful. However, BTS is non-minimal: each agent needs to report the forecast in BTS. In addition, BTS requires an infinite number of participants. A series of works (Radanovic and Faltings 2013; 2014; Witkowski and Parkes 2012)) relaxes the large population requirement but loses the informed truthfulness property. Kong and Schoenebeck (2018a) propose a mechanism that is detail-free, informed truthful and works for a small population—the disagreement mechanism. In the above peer prediction literature, agents are assigned a single task and it is typically assumed that the agents are homogeneous and share a common prior. Our work lies in the single-task setting and makes the common prior assumption. However, our study of the linear regression case allows agents to be heterogeneous.

Other works (Dasgupta and Ghosh 2013; Kong and Schoenebeck 2016; Shnayder et al. 2016) consider the multi-task setting, which requires many a priori similar tasks, but usually relaxes the assumptions of homogeneous agents and a common prior.

In contrast with the general peer prediction literature, we

consider a strategic statistical estimation setting where the private signals are from a continuous space.

Continuous signal space Goel and Faltings (2019) propose a mechanism that elicits continuous valued multi-attribute personal data (e.g. body measurements). However, unlike the present work, they consider a setting where agents report multiple, say d , attributes and the continuous valued attributes are modeled by a mixture multidimensional normal distribution with K components, and their results require that $d, K \geq 2$. They also require the number of agents to be infinite. We consider the setting where $d, K = 1$, and two of our proposed mechanisms are suitable for small populations.

Radanovic and Faltings (2014) also consider the continuous signal space. They apply a discretization approach and ask the agents to report their forecast over all the possible signals (non-minimal). However, the discretization approach leads either to large finite space size (when the cells are small), which renders the forecast report impractical; or loses a lot of information (when the cells are large). Moreover, a meaningful discretization requires the mechanism to have certain information about agents’ prior beliefs. Intuitively, without having any information on the prior distribution, the mechanism cannot decide what is a good “resolution” of discretization: no matter how small each “cell” is in a discretization, there exists a Gaussian distribution with sufficiently small variance such that the probability mass of this distribution is mostly in a single cell. Because of this, the mechanism is not detail-free. In contrast, we utilize the metric information such that we can design minimal mechanisms, and the forecast report is practical even in our non-minimal mechanism.

Kong and Schoenebeck (2018b) consider eliciting agents’ forecasts, which are also continuous values between 0 and 1. However, the strategic statistical estimation problems we study can come from high-dimensions and lack the structure which is exploited by Kong and Schoenebeck (2018b).

Strategic machine learning Cai, Daskalakis, and Papadimitriou (2014) study the statistical inference from strategic sources and this work is an inspiration for this paper. One key difference is that, in their model, each agent draws a signal with some noise (less noise is more costly), but then is assumed to truthfully report his received signal. We do not make the assumption that agents must report truthfully. This is especially significant in the case where agents have prior information about the signal that they could use to coordinate without procuring additional information. Their work focuses on incentivizing optimally cost-effective effort rather than truth-telling.

Liu and Chen (2018) use machine learning techniques to incentivize agents to report truthfully on heterogeneous tasks. Conceptually, our linear regression techniques are similar in that they use information aggregation to accomplish information elicitation. However, the work by Liu and Chen (2018) is rather different as it focuses on the discrete setting.

Chen et al. (2018b) also consider the problem of strategy-proof linear regression. However, their setting is very different from ours. In their model, the agents care about the outcome of the learner, and this is what incentivizes them to truthfully (or un-truthfully) report. In our setting, agents are motivated by monetary incentives and are indifferent to the outcome of the learner.

Finally, there is a series of work on procuring data when that data has different costs (Roth and Schoenebeck 2012; Cai, Daskalakis, and Papadimitriou 2014; Chen et al. 2018a; Zheng et al. 2017). Our model is better suited to study procuring high-quality data rather than cost efficiency.

2 Preliminaries

Throughout the paper, we use n to denote the total number of agents. Let Ω_S be the signal space, and each agent i obtains a signal $s_i \in \Omega_S$.

2.1 Prior and Posterior

Before obtaining a signal, each agent believes that the set of n signals (s_1, \dots, s_n) is sampled from a joint distribution \mathbf{P} over Ω_S^n called *the prior*, which is common knowledge. After agent i receives signal s_i , he updates his belief to the *posterior* $\mathbf{Q}_i(s_i)$, which is a distribution over the remaining $n-1$ signals $(s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n)$. We omit the input s_i and just write \mathbf{Q}_i when there is no confusion.

Let \mathcal{P} be a family of prior distributions. An important family of common priors we consider in this paper is the *Gaussian common prior*. First, we define a type of joint distribution called the *two-step Gaussian distribution*.

Definition 2.1. Given $m_0 \in \mathbb{R}^d$ and two positive definite matrices $\sigma^2, \tau^2 \in S_{++}^d \subset \mathbb{R}^{d \times d}$ where S_{++}^d is the set of d -dimensional positive definite matrices, a *two-step Gaussian distribution* with parameters $(n, m_0, \sigma^2, \tau^2)$ is a joint distribution $(X_1, \dots, X_n) \in \mathbb{R}^d$ defined as follows: a common mean (state) μ is sampled from $\mathcal{N}(m_0, \sigma^2)$; for all $i \in [n]$, $X_i \sim \mathcal{N}(\mu, \tau^2)$ are sampled independently and identically.

Definition 2.2. Consider the signal space $\Omega_S = \mathbb{R}^d$. A prior \mathbf{P} is called a *Gaussian common prior* if \mathbf{P} is a two-step Gaussian distribution with parameters $(n, m_0, \sigma^2, \tau^2)$ for certain $m_0 \in \mathbb{R}^d$ and positive definite matrices $\sigma^2, \tau^2 \in S_{++}^d$. We denote a Gaussian common prior by $G(n, m_0, \sigma^2, \tau^2)$. \mathbf{P} is called a *Jeffreys prior* if \mathbf{P} is a two-step Gaussian distribution with parameters (n, m_0, ∞, τ^2) .¹

We summarize some related properties of the Gaussian common prior below (see Bishop (2006)).

Proposition 2.3. *Under a Gaussian common prior $G(n, m_0, \sigma^2, \tau^2)$, the marginal distribution for each s_i is $\mathcal{N}(m_0, \sigma^2 + \tau^2)$, and after agent i receives signal s_i , the posterior \mathbf{Q}_i is a two-step Gaussian distribution with parameters $G(n-1, (\sigma^{-2} + \tau^{-2})^{-1}(\tau^{-2}s_i + \sigma^{-2}m_0), (\sigma^{-2} + \tau^{-2})^{-1}, \tau^2)$.*

¹The *Jeffreys prior* (Jeffreys 1946) is a special case of the Gaussian common prior where the distribution of the common mean is arbitrarily close to the uniform distribution. This property is captured by setting $\sigma = \infty$.

Consequently, under a Jeffreys prior $\mathbf{P} = G(n, m_0, \infty, \tau^2)$, after agent i receives signal s_i , the posterior \mathbf{Q}_i is a two-step Gaussian distribution $G(n-1, s_i, \tau^2, \tau^2)$.

2.2 Game Theory Basics

Informally, an *information elicitation mechanism* collects from each agent i a *report*, and rewards each agent i based on his report and the other agents' reports.

Definition 2.4 (Mechanism). An *information elicitation mechanism* $\mathcal{M} = (\Omega_R, \pi_R, M_1, \dots, M_n)$ specifies a space of allowed reports Ω_R , what kind of information should be contained in the report π_R , and rewards each agent i $M_i(\mathbf{r}) : \Omega_R^n \rightarrow \mathbb{R}_{\geq 0}$ upon receiving the report collection $\mathbf{r} = (r_1, \dots, r_n) \in \Omega_R^n$. We view π_R as a function $\pi_R : \mathcal{P} \times \Omega_S \rightarrow \Omega_R$ which maps a prior distribution and a signal to a report which truthfully contains the required information.²

Based on the nature of Ω_R and π_R , mechanisms can be classified into the following three types.

(Revelation) minimal $\Omega_R = \Omega_S$, and $\pi_R(\mathbf{P}, s_i) = s_i$. That is, the mechanism only collects each agent i 's private signal s_i ;

(Non-revelation) minimal $\Omega_R = \Omega_S$, and π_R can be arbitrary. That is, for each agent i , the mechanism collects a value $r \in \Omega_S$, which may be the private signal s_i , or anything else specified by the mechanism. For example, in the Gaussian common prior case, this may be the posterior of the common mean \mathbf{Q}_i (which is $(\sigma^2 + \tau^2)^{-1}(\sigma^2 s_i + \tau^2 m_0)$ as given in Proposition 2.3);

Non-minimal Ω_R and π_R can be arbitrary. For example, a non-minimal mechanism can collect from agent i his private signal s_i and the mean of his posterior belief \mathbf{Q}_i . In the case \mathbf{P} is a Gaussian common prior, we have $\Omega_R \in \mathbb{R}^d \times \mathbb{R}^d$ and $\pi_R(\mathbf{P}, s_i) = (s_i, \mu_i)$ where $\mu_i = (\sigma^{-2} + \tau^{-2})^{-1}(\tau^{-2} s_i + \sigma^{-2} m_0)$ (see Proposition 2.3). Our *disagreement mechanism* in Sect. 5 is an example of this.

The definitions of (*mixed*) *strategy* (denoted by $\theta_i : \mathcal{P} \times \Omega_S \rightarrow \Delta_{\Omega_R}$),³ (*expected*) *utility* (denoted by $u_i : \mathcal{P} \times \Delta_{\Omega_R} \rightarrow \mathbb{R}$), *social welfare* and *Bayes Nash equilibrium* follow their standard definitions in game theory, and the precise definitions are in the full version. A strategy profile $\boldsymbol{\theta} = (\theta_1, \theta_2, \dots, \theta_n)$ is *truth-telling* if each θ_i is a pure strategy and $\theta_i(\mathbf{P}, s_i) = \pi_R(\mathbf{P}, s_i)$. On the other hand, a strategy θ_i is *oblivious* if θ_i does not depend on the signals: for any $s_1, s_2 \in \Omega_S$ and any $\mathbf{P} \in \mathcal{P}$, we have $\theta_i(\mathbf{P}, s_1) = \theta_i(\mathbf{P}, s_2)$.

²Intuitively, a mechanism recommends π_R to agents which is a mapping from signals to reports, i.e., "what agents are supposed to report." This notion is necessary because a real number has countably many digits and a mechanism taking in a single real number as input can actually encode an arbitrarily long sequence of real numbers.

³In this work, the report spaces Ω_R are Euclidean spaces, and we only consider the agents' strategies have finite second moments: for each agent i and $s_i \in \Omega_S$, $\theta_i(\mathbf{P}, s_i)$, $\mathbb{E}[\theta_i(\mathbf{P}, s_i) \cdot \theta_i(\mathbf{P}, s_i)^\top]$ is finite. The restriction to finite second moment is largely a technicality as highly diffuse strategies are unlikely to pay well in our mechanisms.

2.3 Mechanism Design Goals

A mechanism is (*strictly*) *truthful* if the truth-telling strategy profile is a (strict) Bayes Nash equilibrium under every prior $\mathbf{P} \in \mathcal{P}$. A stronger goal is to design a mechanism that is (strictly) truthful and *the truth-telling profile has the maximum welfare*. The ultimate goal we want from a mechanism is the *informed truthfulness*.

Definition 2.5. A mechanism \mathcal{M} is *informed truthful* under a prior family \mathcal{P} if it is 1) strictly truthful, 2) the truth-telling profile has the maximum welfare,⁴ and 3) there is no oblivious Bayes Nash equilibrium or and for any $i \in [n]$ and $\mathbf{P} \in \mathcal{P}$, $u_i(\mathbf{P}, \boldsymbol{\theta}) > u_i(\mathbf{P}, \boldsymbol{\theta}')$, where $\boldsymbol{\theta}$ is the truth-telling profile and $\boldsymbol{\theta}'$ is any oblivious Bayes Nash equilibrium profile.

Definition 2.6. A mechanism \mathcal{M} is an *optimal estimator* of a statistic, if, in the truthful equilibrium, the collection of agents' report forms a sufficient statistic (Fisher 1922).

3 Minimal Revelation Mechanisms

In this section, we study minimal revelation mechanisms which ask agents to report private signals only. We first present a mechanism called *the metric mechanism* that is revelation minimal and informed truthful when agents have a Jeffery prior. However, in Sect. 3.2, we present a strong impossibility result showing that there is no revelation minimal mechanism that can achieve truthfulness in general Gaussian common priors.

3.1 Minimal Revelation Mechanism for Jeffreys Prior

In this subsection, we present *the metric mechanism* which is minimal revelation and informed truthful if agents have a Jeffreys prior as in Definition 2.2.

Mechanism 1 The metric mechanism $\mathcal{M}_{\text{metric}}$

- 1: Generate (predetermined) two disjoint groups A, B with size equal to $\lfloor \frac{n}{2} \rfloor$ and $n - \lfloor \frac{n}{2} \rfloor$ arbitrarily.
- 2: Each agent i reports a signal $\hat{s}_i \in \mathbb{R}^d$ where $\Omega_R = \mathbb{R}^d$ and $\pi_R(\mathbf{P}, s_i) = s_i$. If $|A| < |B|$ we randomly remove an agent in group B and give it 0 payment.
- 3: For each agent $i \in A$ we randomly choose a reference agent $j \in B$ and a competitor $k \in A$ (and vice versa for each agent in B) such that i, j, k are distinct. The payment to agent i is

$$M_i(\hat{\mathbf{s}}) = \begin{cases} -100 & \text{if } \hat{s}_i = \hat{s}_k \\ \|\hat{s}_j - \hat{s}_k\|_2^2 - \|\hat{s}_j - \hat{s}_i\|_2^2 & \text{otherwise.} \end{cases} \quad (1)$$

Theorem 3.1. *Given any Gaussian prior $G(n, m_0, \infty, \tau^2)$ defined in Definition 2.2 with $n \geq 4$, the minimal revelation mechanism $\mathcal{M}_{\text{metric}}$ is an optimal estimator, and is, in addition, informed truthful.*

⁴There is a subtle distinction. In Shnayder et al. (2016)'s definition, the truth-telling profile has the maximum expected payment to each agents. Our definition is equivalent to theirs when every agent uses the same strategy.

To prove this theorem, the following lemmas are required.

Proposition 3.2 (Bias variance decomposition). *Given a random vector $x \in \mathbb{R}^d$ and constant matrix $A \in \mathbb{R}^{d \times d}$,*

$$\mathbb{E}[x^\top A x] = \text{Tr}(A \text{Cov}(x)) + \mathbb{E}[x]^\top A \mathbb{E}[x],$$

where $\text{Cov}(x) = \mathbb{E}[(x - \mathbb{E}[x])^\top (x - \mathbb{E}[x])]$ is the covariance matrix of x .

Proposition 3.2 implies the following lemma, and the proof is in the full version.

Lemma 3.3. *Given a distribution $D \in \Delta_{\mathbb{R}^d}$ and an arbitrary constant positive definite matrix $A \in \mathbb{R}^{d \times d}$, define $U : \mathbb{R}^d \rightarrow \mathbb{R}$ as $U(x) = \mathbb{E}_{s \sim D} [(x - s)^\top A (x - s)]$. Then $U(\cdot)$ is continuous, and has unique minimizer $x = \mathbb{E}_{s \sim D}[s]$ with $\min U(x) = \text{Tr}(A \text{Cov}(s))$.*

Proof of Theorem 3.1. Strictly truthful: We first analyze the best response for agent i at truth-telling strategy profile θ . Note that if everyone tells the truth $\hat{s}_j = s_j$ which is a two-step Gaussian distribution, and the probability of $s_j = s_i$ for any $k \neq i$ is 0. So, a strategy θ_i is a best response if it minimizes the quadratic form

$$\mathbb{E}[\|\hat{s}_j - \hat{s}_i\|_2^2 \mid s_i] = \mathbb{E}[(s_j - \hat{s}_i)^\top \mathbb{I}(s_j - \hat{s}_i) \mid s_i] \quad (2)$$

Because identity matrix \mathbb{I} is positive definite, and s_j and \hat{s}_i are independent conditioned on s_i , by Lemma 3.3, it is minimized when $\mathbb{E}_{j \in B}[s_j - \hat{s}_i \mid s_i] = 0$ and $\mathbb{E}[\text{Tr}(\text{Cov}(\hat{s}_i)) \mid s_i] = 0$, so $\mathbb{E}[\theta_i(s_i)] = \mathbb{E}_{j \in B}[s_j \mid s_i] = s_i$ due to Jeffreys prior. As a result, truth-telling is a strict Bayesian Nash equilibrium.

Truth-telling has the maximum welfare: We want to show the social welfare of truth-telling equilibrium is better than or equal to all other non-oblivious strategy Bayesian Nash equilibrium.

Given any fixed report profile \hat{s} such that $\forall i \neq k \hat{s}_i \neq \hat{s}_k$, we consider the expected social welfare with respect to the randomness over the choice of reference agents j, k . If n is even,

$$\begin{aligned} \sum_i M_i(\mathbf{s}) &= \sum_i \mathbb{E}_{j,k} [\|\hat{s}_j - \hat{s}_k\|_2^2 - \|\hat{s}_j - \hat{s}_i\|_2^2] \\ &= 2 \sum_{i \in A, j \in B} \left(\frac{2\|\hat{s}_j - \hat{s}_k\|_2^2}{n} \right) - 2 \sum_{k \in A, j \in B} \left(\frac{2\|\hat{s}_j - \hat{s}_k\|_2^2}{n} \right) \\ &= 0. \end{aligned}$$

Informed truthful: Finally, we show any oblivious equilibrium has strictly less welfare. Suppose there is an oblivious equilibrium with strategy profile θ . $\theta_i(s_i)$ is independent of s_i , so \hat{s}_i are independent samples from random variables θ_i . Suppose there exists i, k in A , $\hat{s}_i \neq \hat{s}_k$, and, without loss of generality, $\hat{s}_i \neq \mathbb{E}_{j \in B}[\theta_j]$, it is not an equilibrium because there exists $s_i^* \triangleq (\hat{s}_i + \mathbb{E}_{j \in B}[\theta_j])/2$ which has larger expected payment by (2). Therefore for all i, k in A , $\hat{s}_i = \hat{s}_k$ and the average welfare is -100 . \square

Remark 3.4. There is a simple mechanism with 0-1 payments that achieves the same property. Instead of rewarding

each agent based on (1), we can reward i for being strictly “between” j and k :

$$M_i(\hat{\mathbf{s}}) = \begin{cases} 1, & \text{if } \|\hat{s}_j - \hat{s}_i\| < \|\hat{s}_j - \hat{s}_k\| \\ 0, & \text{otherwise.} \end{cases}$$

The same arguments in the proof of Theorem 3.1 show that this variant of the metric mechanism is informed truthful.

3.2 Impossibility Results

The theorem below, whose proof is deferred to the full version, shows that, minimal revelation mechanisms cannot achieve even the weakest truthful property—truth-telling as a strict Bayes Nash equilibrium, if no additional assumption is made about the prior distribution (other than it is a Gaussian common prior). Notice that having truth-telling as a (weak) Bayes Nash equilibrium is trivial: we can pay each agent a fixed amount regardless of what he reports.

Theorem 3.5. *For any number n of agents, there is no minimal revelation mechanism such that the truth-telling profile $\{\theta_i(\mathbf{P}, s_i) = s_i\}$ is a strict Bayes Nash equilibrium.*

In the full version, we also show that, if we consider the *nonzero-effort* setting where each agent i needs to spend a positive amount of effort $c_i > 0$ to obtain signal s_i and assume quasi-linear utility functions (i.e., for each agent i , his utility is given by his payment minus the effort c_i he spent), there does not even exist a mechanism that has truth-telling being a (weakly) Bayes Nash equilibrium.

4 Proxy Bayesian Truth Serum

In this section, we propose a minimal mechanism called *the proxy Bayesian truth serum* (the proxy BTS), which is informed truthful for Gaussian common priors. From each agent, the proxy BTS elicits either the agent’s private signal or the agent’s posterior expectation. We first present the mechanism in a way that collects *both* the private signal and posterior expectation from each agent, and then show how to make it minimal by collecting *either* the private signal *or* the posterior expectation.

Theorem 4.1. *Let $n \rightarrow \infty$ and assume the Gaussian common prior $G(n, m_0, \sigma^2, \tau^2)$ on \mathbb{R}^d defined in Definition 2.2. The mean estimation mechanism $\mathcal{M}_{\text{proxy}}$ in Mechanism 2 is an optimal estimator. Additionally, $\mathcal{M}_{\text{proxy}}$ is informed truthful.*

The main idea of Mechanism 2 is, like the Bayesian truth serum (BTS) mechanism in (Prelec 2004), to reward “surprisingly common” reports. The agents are asked to provide both the signals and the posterior beliefs of other agents’ signals, and their payments can be decomposed into two parts, *the prediction score* and *the information score*.

Prediction score is based on how *accurate* the reported prediction is. The first term $(-\hat{s}_j - \hat{t}_i)^\top L^{-1}(\hat{s}_j - \hat{t}_i)$ is larger when the agent’s prediction \hat{t}_i is closer to the reference agent’s reported signal s_j . Moreover, the value of the first term is essentially the log-likelihood of \hat{t}_i with respect to the Gaussian distribution $\mathcal{N}(\hat{s}_j, L)$. Thus, it is maximized from the perspective of agent i when the reported prediction

Mechanism 2 Proxy-BTS mechanism $\mathcal{M}_{\text{proxy}}$

- 1: Generate two disjoint groups A, B with sizes equal to $\lfloor \frac{n}{2} \rfloor$ and $n - \lfloor \frac{n}{2} \rfloor$ arbitrarily.
- 2: Each agent $i \in A$ reports a signal and a prediction of the posterior mean $r_i = (\hat{s}_i, \hat{t}_i)$. That is $\Omega_R = \mathbb{R}^d \times \mathbb{R}^d$ and $\pi_R(\mathbf{P}, s_i) = (s_i, \mathbb{E}[\mu|s_i])$. If $|A| < |B|$ randomly remove an agent in group B and give it 0 payment.
- 3: For each agent $i \in A$, randomly choose a reference agent $j \in B$ and all agents in $A_{-i} \triangleq A \setminus \{i\}$ as i 's competitors (and vice versa when $i \in B$). Calculate the sample means

$$\mu_s(i) \triangleq \frac{1}{|A_{-i}|} \sum_{k \in A_{-i}} \hat{s}_k, \quad \mu_t(i) \triangleq \frac{1}{|A_{-i}|} \sum_{k \in A_{-i}} \hat{t}_k$$

and sample covariances:

$$\frac{1}{|A_{-i}| - 1} \sum_{k \in A_{-i}} (\hat{s}_k - \mu_s(i))(\hat{s}_k - \mu_s(i))^\top \quad (3)$$

$$\frac{1}{|A_{-i}| - 1} \sum_{k \in A_{-i}} (\hat{t}_k - \mu_t(i))(\hat{t}_k - \mu_t(i))^\top \quad (4)$$

of the signal reports $\{\hat{s}_k\}_{k \in A_{-i}}$ and the prediction reports $\{\hat{t}_k\}_{k \in A_{-i}}$ respectively.

- 4: If (3) or (4) is not positive definite, skip to step 6. Otherwise, let $\Sigma_s(i), \Sigma_t(i)$ be their positive square roots and drop i when there is no ambiguity. Compute

$$K \triangleq \Sigma_s(\Sigma_s^{-1}\Sigma_t^2\Sigma_s^{-1})^{1/2}\Sigma_s^{-1}, \quad L \triangleq (K + \mathbb{I})\Sigma_s^2. \quad (5)$$

- 5: Depending on \hat{t}_i , the agent i 's prediction score $\text{PS}_i(\mathbf{r})$ is

$$-(\hat{s}_j - \hat{t}_i)^\top L^{-1}(\hat{s}_j - \hat{t}_i) + (\hat{s}_j - \mu_s)^\top \Sigma_s^{-2}(\hat{s}_j - \mu_s), \quad (6)$$

and depending on \hat{s}_i the information score $\text{IS}_i(\mathbf{r})$ is

$$-(\hat{s}_i - \mu_s)^\top \Sigma_s^{-2}(\hat{s}_i - \mu_s) + (\hat{s}_i - \hat{t}_j)^\top L^{-1}(\hat{s}_i - \hat{t}_j). \quad (7)$$

- 6: The reward for agent i , $M_i(\mathbf{r})$, is

$$\begin{cases} -100, & \text{if (3) or (4) are singular for some agent} \\ \text{IS}_i(\mathbf{r}) + \text{PS}_i(\mathbf{r}) & \text{otherwise} \end{cases}.$$

is the maximum likelihood estimator of the mean— $\mathbb{E}[s_j|s_i]$. The second term $(\hat{s}_j - \mu_s)^\top \Sigma_s^{-2}(\hat{s}_j - \mu_s)$ is independent of agent i 's report, and makes the social welfare, sum of IS and PS, equal to zero.

Information score is based on an agent's reported signal \hat{s}_i . A key observation for the information score is that the expectation of agent i on other agent's signal in two-step Gaussian prior is a linear combination of its private signal s_i and the prior mean m_0 . Therefore, if we have 1) agent i 's expectation on other agent's signal μ_s and 2) its expectation on other agent's prediction \hat{t}_j , we can recover agent i 's signal through a proper linear combination of those two values. Conceptually, we achieve this by rewarding agent i 's signal




Figure 1: Consider a Gaussian common prior $G(n, 0, 3, 1)$, and agent i receives $s_i = 4$. Agent i 's prediction score has two terms: The blue curve is the first term $-(\hat{s}_i - \mu_s)^\top \Sigma_s^{-2}(\hat{s}_i - \mu_s)$. The red curve is the second term $(\hat{s}_i - \hat{t}_j)^\top L^{-1}(\hat{s}_i - \hat{t}_j)$.

by two quadratic functions (Figure 1). The optimal point of the first term $-(\hat{s}_i - \mu_s)^\top \Sigma_s^{-2}(\hat{s}_i - \mu_s)$ is agent i 's expectation of μ_s , and the second term $(\hat{s}_i - \hat{t}_j)^\top L^{-1}(\hat{s}_i - \hat{t}_j)$ is centered on the agent's expectation of the reference agent's prediction \hat{t}_j . By picking L and Σ_s correctly, we can shift the optimal point to the proper linear combination of the above two values. Another way of seeing this is that the first term rewards the agent for giving a common answer while the second term punishes an agent for giving an expectedly common answer. Thus agents are rewarded for “surprisingly common” answers. While the mechanism is inspired by the intuition of Bayesian Truth Serum (Prelec 2004), the proof is not. It must be taken into consideration that agents only report posterior expectations instead of posterior beliefs in the original BTS.

Finally, if agents' strategy are oblivious, they have the same prediction on other's signal \hat{t} . Therefore (4) is singular, and everyone is punished with -100 . Contrarily, in the truth-telling strategy profile, the probability that (3) or (4) are singular is zero.

Remark 4.2. Note that the reward has two parts PS and IS. For agent i , the value PS only depends on \hat{t}_i , and IS only depends on \hat{s}_i . We can (randomly) ask half of the agents to report their signals, \hat{s}_i , and pay them by IS and the other half to report their prediction, \hat{t}_i , and pay them by PS. The resulting mechanism is still informed truthful, and it becomes (non-revelation) minimal.

5 Disagreement Mechanism

In this section, we present an informed truthful non-minimal mechanism for a small number of agents, called the *disagreement mechanism*. The intuition underpinning the mechanism is that agents' signal reports and forecast reports should “agree” because an agent's forecast is just the common prior updated by their signal. In particular, the mechanism will use other agents' signals to infer a mapping (presumably from the prior) of private signals to posterior predictions of the mean. Then, it will punish an agent based on how inconsistent his report is with this mapping. It borrows its name from the mechanism in Kong and Schoenebeck

(2018a) but the same idea was used in Radanovic and Faltings (2014). However, much additional effort is needed to adapt it into our continuous setting with smaller report spaces.

Mechanism 3 The disagreement mechanism $\mathcal{M}_{\text{disagree}}$

- 1: Each agent i reports a signal $\hat{s}_i \in \mathbb{R}^d$, the mean $\hat{\mu}_i$ of his posterior belief \mathbf{Q}_i , and an *amount of untruthfulness* $\xi_i \in \mathbb{R}_{\geq 0}$ where agent i confesses how untruthful his report is, with $\xi_i = 0$ being completely truthful. That is, $\Omega_R = \mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R}_{\geq 0}$ and $\pi_R(\mathbf{P}, s_i) = (s_i, \mu_i, 0)$.
- 2: Partition the agents into three groups A , B and C of sizes at least $d + 1$ arbitrarily. Let $\{a_0, a_1, \dots, a_d\}$, $\{b_0, b_1, \dots, b_d\}$ and $\{c_0, c_1, \dots, c_d\}$ be the first $d + 1$ agents in the three groups respectively.
- 3: Each agent i is assigned a *prediction score* defined as

$$\text{PS}_i(\mathbf{r}) = \begin{cases} -\|\hat{s}_{b_0} - \hat{\mu}_i\|^2 - \xi_i & \text{if } i \in A \\ -\|\hat{s}_{c_0} - \hat{\mu}_i\|^2 - \xi_i & \text{if } i \in B \\ -\|\hat{s}_{a_0} - \hat{\mu}_i\|^2 - \xi_i & \text{if } i \in C \end{cases}.$$

- 4: Each agent i is also assigned an *inconsistency score* computed as follows: for each $j \in \{1, \dots, d\}$, let $t_{a_j} = \hat{s}_{a_j} - \hat{s}_{a_0}$ and $T_a = [t_{a_1} t_{a_2} \dots t_{a_d}]$; define T_b, T_c similarly. Let $\nu_{a_j} = \hat{\mu}_{a_j} - \hat{\mu}_{a_0}$ and $U_a = [\nu_{a_1} \nu_{a_2} \dots \nu_{a_d}]$; define U_b, U_c similarly; The inconsistency score for each $i \in A$ is

$$\text{IS}_i(\mathbf{r}) = -\|U_b T_b^{-1}(\hat{s}_i - \hat{s}_{b_0}) - (\hat{\mu}_i - \hat{\mu}_{b_0})\|$$

if T_b is invertible and $\xi_i = 0$; otherwise $\text{IS}_i(\mathbf{r}) = 0$. Define the inconsistency score for agents in group B and C similarly (The report of each agent in B is compared to the first $d + 1$ reports in C , and the report of each agent in C is compared to the first $d + 1$ reports in A).

- 5: The *payment before normalization* for agent i is $\bar{M}_i(\mathbf{r}) = \text{PS}_i(\mathbf{r}) + \text{IS}_i(\mathbf{r})$.
- 6: Normalize the payments for each agent i as follows:

$$M_i(\mathbf{r}) = \begin{cases} \bar{M}_i(\mathbf{r}) - \frac{1}{|A|} \sum_{j \in B} \bar{M}_j(\mathbf{r}) & \text{if } i \in A \\ \bar{M}_i(\mathbf{r}) - \frac{1}{|B|} \sum_{j \in C} \bar{M}_j(\mathbf{r}) & \text{if } i \in B \\ \bar{M}_i(\mathbf{r}) - \frac{1}{|C|} \sum_{j \in A} \bar{M}_j(\mathbf{r}) & \text{if } i \in C \end{cases}.$$

- 7: If all the agents in group A report the same posterior mean $\hat{\mu}_i$, then update the score for each $i \in A$: $M_i(\mathbf{r}) \leftarrow M_i(\mathbf{r}) - 100$. Do the same for agents in group B and C .
-

The mechanism $\mathcal{M}_{\text{disagree}}$ is shown in Mechanism 3. The mechanism partitions agents into three groups A , B and C with sufficiently large sizes, and chooses $d + 1$ reference agents in each group. Specifically, all reports from group A (B , C) are compared to the $d + 1$ reference agents in group B (C , A). We say B, C, A are the reference groups of A, B, C respectively. Each agent is required to report a signal \hat{s}_i , the mean $\hat{\mu}_i$ of his posterior belief \mathbf{Q}_i , and an *amount of untruthfulness* $\xi_i \in \mathbb{R}_{\geq 0}$ where agent i confesses how untruthful his report is, with $\xi_i = 0$ being completely truthful. The prediction score measures how well $\hat{\mu}_i$ predicts the signal reported by the first reference agent in the reference group.

Specifically, we compute the squared distance between $\hat{\mu}_i$ and the reference agent's reported signal, with the amount of untruthfulness ξ_i punished (Step 3). The inconsistency score measures how consistent \hat{s}_i and $\hat{\mu}_i$ are (Step 4). In particular, to compute the inconsistency score for an agent $i \in A$, we first use the reports for the $d + 1$ reference agents from the reference group B to infer a bijection between private signals and posterior beliefs of the mean. The mechanism then checks whether an agent's posterior mean $\hat{\mu}_i$ and reported signal \hat{s}_i are indeed consistent with this bijection, and punishes agent i for the amount of inconsistency. Notice that we do not punish an agent i for inconsistency if he confesses a positive amount of untruthfulness $\xi_i > 0$. The payment to each agent i is then computed by the sum of the two scores (Step 5). At this stage, strict truthfulness is guaranteed: if all the other agents are truth-telling, an agent's unique best response is also truth-telling with zero confession of untruthfulness, as this is the unique way to simultaneously maximize the prediction score and the inconsistency score.

Next, we normalize the payments (without changing the set of all equilibria) such that the sum of all the n agents' payments is zero (Step 6). Finally, we punish all agents in each group if they report the same posterior (Step 7).⁵ Step 6 makes sure that the truth-telling equilibrium has the largest social welfare (which is 0). Step 7 makes sure there is no oblivious equilibrium at all. Intuitively, an agent i 's best strategy is to report $\hat{\mu}_i$ that best predicts the signal reported by the reference agent, because he can always avoid punishment from the inconsistency score by confessing a negligible $\xi_i > 0$. Since agent i can make his utility arbitrarily close to the best prediction score (if unable to make it equal to), in any oblivious equilibrium, all the agents in the same group should report the same best prediction $\hat{\mu}_i$. This is exactly the case in Step 7 where they are punished by -100 . Putting together, $\mathcal{M}_{\text{disagree}}$ is informed truthful. A rigorous proof of Theorem 5.1 below is deferred to the full version.

Theorem 5.1. *Assume a Gaussian prior $G(n, m_0, \sigma^2, \tau^2)$ in \mathbb{R}^d defined in Definition 2.2. When the number of agents $n \geq 3d + 3$, the mean estimation mechanism $\mathcal{M}_{\text{disagree}}$ is an optimal estimator. Additionally, $\mathcal{M}_{\text{disagree}}$ is informed truthful.*

Proof sketch. We first show that \mathcal{M} is strictly truthful. For an arbitrary agent i in, say, group A , upon receiving the private signal s_i , agent i believes the signal received by each of the other agents, including agent b_0 , is from the Gaussian distribution $\mathbf{q}_i(s_i) = \mathcal{N}(\mu_i, (\sigma^{-2} + \tau^{-2})^{-1} + \tau^2)$ where $\mu_i = (\sigma^{-2} + \tau^{-2})^{-1}(\tau^{-2}s_i + \sigma^{-2}m_0)$ (see Proposition 2.3). Firstly, we can show that $(\hat{\mu}_i, \xi_i) = (\mu_i, 0)$ is the unique maximizer to the prediction score by Proposition 3.3. Secondly, we can show that, fixing $(\hat{\mu}_i, \xi_i) = (\mu_i, 0)$, reporting

⁵An alternative design to the present mechanism would be to impose a sufficiently large enough punishment for each agent in group A (B , C) when all agents in group B (C , A) report the same *signal*—a more common design in the past literature. This also yields an informed truthful mechanism. However, this variant of the mechanism allows oblivious Nash equilibria, whereas the mechanism described in Mechanism 3 forbids all oblivious Nash equilibria, which is, in some sense, stronger.

$\hat{s}_i = s_i$ is the only way to make the inconsistency score 0. This again follows from Proposition 2.3. Therefore, fixing the other $n - 1$ agents' truthful report \mathbf{r}_{-i} , reporting $r_i = (s_i, \mu_i, 0)$ is the unique maximizer to $\bar{M}_i(\cdot, \mathbf{r}_{-i})$.

By the design of the mechanism, $\bar{M}_j(\mathbf{r})$ does not depend on r_i for each $j \in B$, a maximizer to $\bar{M}_i(\cdot, \mathbf{r}_{-i})$ is also a maximizer to $M_i(\cdot, \mathbf{r}_{-i})$. As a result, we conclude that \mathcal{M} is strictly truthful. Furthermore, Step 6 makes sure $\sum_{i=1}^n M_i(\mathbf{r}) = 0$, so the truth-telling profile has the maximum welfare 0.

Lastly, we can show that \mathcal{M} is informed truthful by showing that there is no oblivious Bayes Nash equilibrium. If there were an oblivious equilibrium θ such that all agents in a group do not report the same signal, then there must exist an agent whose reported posterior mean $\hat{\mu}_i$ does not predict the reference agent's signal as well as the others. In this case, i can improve his prediction while making the inconsistency score 0 by confessing a tiny amount of untruthfulness $\hat{\xi}_i > 0$. If there were an oblivious equilibrium θ such that all agents in a group reported the same signal, then an arbitrary agent in the group can slightly misreport his posterior mean $\hat{\mu}_i$ to avoid the -100 punishment. \square

6 Future Work

There are many attractive avenues available for future work. Within our model, the question of whether there exists a minimal non-revelation mechanism that applies in the small population setting remains open. Moving beyond our model, one could study priors that are more general than Gaussian distributions (e.g. exponential families), or learning problems beyond linear regression (e.g. SVMs). Additionally, one could study how to efficiently (with minimum cost) elicit data in a setting similar to ours. A final direction in full pipe-line learning, is to relax the assumptions on the agents being fully rational. For example, a fraction of the agents could be random, malicious, or naturally truth-telling.

References

- Bishop, C. M. 2006. *Pattern recognition and machine learning*. springer.
- Cai, Y.; Daskalakis, C.; and Papadimitriou, C. H. 2014. Optimum statistical estimation with strategic data sources. *arXiv preprint arXiv:1408.2539*.
- Chen, Y.; Immorlica, N.; Lucier, B.; Syrgkanis, V.; and Ziani, J. 2018a. Optimal data acquisition for statistical estimation. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, 27–44. ACM.
- Chen, Y.; Podimata, C.; Procaccia, A. D.; and Shah, N. 2018b. Strategyproof linear regression in high dimensions. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, 9–26. ACM.
- Dasgupta, A., and Ghosh, A. 2013. Crowdsourced judgement elicitation with endogenous proficiency. In *Proceedings of the 22nd international conference on World Wide Web*, 319–330. International World Wide Web Conferences Steering Committee.
- Fisher, R. A. 1922. On the mathematical foundations of theoretical statistics. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character* 222(594-604):309–368.
- Goel, N., and Faltings, B. 2019. Personalized peer truth serum for eliciting multi-attribute personal data. In *UAI*.
- Jeffreys, H. 1946. An invariant form for the prior probability in estimation problems. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences* 453–461.
- Kong, Y., and Schoenebeck, G. 2016. A Framework For Designing Information Elicitation Mechanisms That Reward Truth-telling. *ArXiv e-prints*.
- Kong, Y., and Schoenebeck, G. 2018a. Equilibrium selection in information elicitation without verification via information monotonicity. In *9th Innovations in Theoretical Computer Science Conference (ITCS 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- Kong, Y., and Schoenebeck, G. 2018b. Water from two rocks: Maximizing the mutual information. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, 177–194. ACM.
- Kong, Y.; Ligett, K.; and Schoenebeck, G. 2016. Putting peer prediction under the micro (economic) scope and making truth-telling focal. In *International Conference on Web and Internet Economics*, 251–264. Springer.
- Liu, Y., and Chen, Y. 2018. Surrogate scoring rules and a dominant truth serum for information elicitation. *CoRR* abs/1802.09158.
- Miller, N.; Resnick, P.; and Zeckhauser, R. 2005. Eliciting informative feedback: The peer-prediction method. *Management Science* 1359–1373.
- Prelec, D. 2004. A Bayesian Truth Serum for subjective data. *Science* 306(5695):462–466.
- Radanovic, G., and Faltings, B. 2013. A robust bayesian truth serum for non-binary signals. In *Twenty-Seventh AAAI Conference on Artificial Intelligence*.
- Radanovic, G., and Faltings, B. 2014. Incentives for truthful information elicitation of continuous signals. In *Twenty-Eighth AAAI Conference on Artificial Intelligence*.
- Roth, A., and Schoenebeck, G. 2012. Conducting truthful surveys, cheaply. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, 826–843. ACM.
- Shnayder, V.; Agarwal, A.; Frongillo, R.; and Parkes, D. C. 2016. Informed truthfulness in multi-task peer prediction. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, EC'16, 179–196. New York, NY, USA: ACM.
- Witkowski, J., and Parkes, D. C. 2012. Peer prediction without a common prior. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, 964–981. ACM.
- Zheng, S.; Waggoner, B.; Liu, Y.; and Chen, Y. 2017. Active information acquisition for linear optimization. *arXiv preprint arXiv:1709.10061*.