Deep Bayesian Trust : A Dominant and Fair Incentive Mechanism for Crowd (Supplementary Material)

Naman Goel, Boi Faltings

In this supplement, we provide all the information that was omitted from the main text due to lack of space.

1 Mixed Strategies

In Section 3 of the paper, we claimed that mixed strategies can also be modeled using the pure strategy variables. In our model, we had two effort strategies (pure) $e_i = 0$ (no effort) or $e_i = 1$ (full effort). If the worker does not invest effort, her reporting strategy $\vec{S_i}$ is any probabilistic vector. If the worker does invest effort, her reporting strategy is a row stochastic matrix S_i.

A mixed effort strategy is the one in which a worker chooses to invest effort only with a probability q and doesn't invest effort with probability $1 - q$. A mixed reporting strategy is one in which worker can probabilistically choose among different reporting strategies (vectors $\vec{S_i}$ or matrices S_i). In either case, we can write the strategy of the worker as a convex combination of the pure strategies. For example, consider ternary answers space ($K = 3$) and a worker who is truthful with probability 0.5 and plays a heuristic strategy of always reporting 0 otherwise. The mixed reporting strategy of the worker is given by

$$S_i = 0.5 \cdot \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} + 0.5 \begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}$$

The uniform strategy space that we considered for multi-task settings includes these mixed strategies as long as the mixed strategy remains same on all tasks. This means that workers can probabilistically switch between pure strategies in multi-task settings but we don't consider the settings, in which workers swtich based on the task properties. This possibility is correctly ruled out in the literature by assuming that the tasks have similar properties so that workers can't choose a distinct strategy based on certain task property.

2 Solving Linear System of Equations

In Section 4 of the paper, we proposed a transitive method of calculating the unknown trustworthiness of worker. This required solving a linear system of Equations and required the coefficient matrix to have linearly independent rows.

As an example, consider the binary answer space ($K = 2$). If $T_j[1,1] + T_j[2,2] \neq 1$, solving the system of Equations gives the following,

$$T_i[1,1] = \frac{b \cdot \omega(y_i = 1|y_j = 1) - (1-a) \cdot \omega(y_i = 1|y_j = 2)}{ab - (1-a)(1-b)}$$

$$T_i[2,2] = 1 - \frac{a \cdot \omega(y_i = 1|y_j = 2) - (1-b) \cdot \omega(y_i = 1|y_j = 1)}{ab - (1-a)(1-b)}$$

where, $a = \frac{T_j[1,1] \cdot P(1)}{\omega(y_j=1)}$ and $b = \frac{T_j[2,2] \cdot P(2)}{\omega(y_j=2)}$.

$T_i[1,1]$ and $T_i[2,2]$ together define the trustworthiness matrix of i in binary answer space. For non-binary case also, similar closed form expression can be derived. In practice, many libraries are available to solve system of linear equations efficiently. We used the numpy.linalg library in Python for this purpose. It may be noted that the requirement for this estimation method to work in binary answer spaces ($K = 2$) is weaker than the general case ($K > 2$). It only requires that the reports of the peer j are not independent of the ground truth ($T_j[1,1] + T_j[2,2] \neq 1 \implies T_j[1,1] \neq T_j[2,1] \implies T_j[2,2] \neq T_j[1,2]$). As long as the reports of peer j have some correlation with the ground truth, we can estimate the trustworthiness matrix of worker i in binary answer space. For the general case, the requirement is explained in the paper (posterior distributions should not be identical for any two possible answers of the peer).

3 Missing Proofs

3.1 Proof of Lemma 1

Proof. Let's first write the expression for the probability $P(Y_j = y_j|Y_i = y_i)$ by applying Bayes' rule.

$$\begin{aligned} P(Y_j = y_j|Y_i = y_i) &= \sum_{g\in[K]} P(Y_j = y_j, G = g|Y_i = y_i) \\ &= \sum_{g\in[K]} P(Y_j = y_j|G = g, Y_i = y_i) \cdot P(G = g|Y_i = y_i) \end{aligned}$$

Since answers of i and j are conditionally independent given the ground truth, we have $P(Y_j = y_j|G = g, Y_i = y_i) = P(Y_j = y_j|G = g)$. This gives the following:

$$P(Y_j = y_j|Y_i = y_i) = \sum_{g\in[K]} P(Y_j = y_j|G = g) \cdot P(G = g|Y_i = y_i)$$

Now we apply the Bayes' rule again and expand the term $P(G = g|Y_j = y_j)$. This gives:

$$P(Y_j = y_j|Y_i = y_i) = \sum_{g\in[K]} P(Y_i = y_i|G = g) \cdot \frac{P(Y_j = y_j|G = g) \cdot P(G = g)}{P(Y_j = y_j)}$$

Note that $P(Y_i = y_i|G = g) = T_i[g, y_i]$ and $P(Y_j = y_j|G = g) = T_j[g, y_j]$. We thus get,

$$P(Y_j = y_j|Y_i = y_i) = \sum_{g\in[K]} T_i[g, y_i] \cdot \frac{T_j[g, y_j] \cdot P(G = g)}{P(Y_j = y_j)}$$

Assuming $|Q^i \cap Q^j| \to \infty$, we now use the law of large numbers and the continuous mapping theorem to replace $P(Y_j = y_j|Y_i = y_i)$ with empirical distribution $\omega(Y_j = y_j|Y_i = y_i)$ and $P(Y_j = y_j)$ with empirical distribution $\omega(Y_j = y_j)$. This finally gives,

$$\omega(Y_i = y_i|Y_j = y_j) = \sum_{k\in[K]} T_i[g, y_i] \cdot \left(\frac{T_j[k, y_j] \cdot P(k)}{\omega(Y_j = y_j)}\right)$$

□

3.2 Proof of Theorem 1

Proof. As $|Q^i \cap Q^j| \to \infty$, using lemma 1, the reward R_i' of any worker i in the Deep Bayesian Trust Mechanism is given by :

$$\begin{aligned} R_i' &= \beta\Big[\Big(\sum_{k\in[K]} T_i[k,k]\Big) - 1\Big] \\ &= \beta\Big[\Big(\sum_{k\in[K]}\sum_{m\in[K]} A_i[k,m]S_i[m,k]\Big) - 1\Big] \end{aligned}$$

(Using Proposition 1)

For binary answer space ($K = 2$), this can be expanded as :

$$R_i = \beta\Big[A_i[1,1]S_i[1,1] + A_i[1,2]S_i[2,1] + A_i[2,1]S_i[1,2] + A_i[2,2]S_i[2,2] - 1\Big]$$

Rearranging the terms,

$$R_i = \beta\Big[\Big[A_i[1,1]S_i[1,1] + A_i[2,1]S_i[1,2]\Big] + \Big[A_i[1,2]S_i[2,1] + A_i[2,2]S_i[2,2]\Big] - 1\Big]$$

Assuming $A_i[1,1] + A_i[2,2] > 1$s, we get that $A_i[1,1] > A_i[2,1]$ and $A_i[2,2] > A_i[1,2]$.

Now, note that $\big[A_i[1,1]S_i[1,1]+A_i[2,1]S_i[1,2]\big]$ is a convex combination of $A_i[1,1]$ and $A_i[1,2]$ with $S_i[1,1]$ and $S_i[2,1]$ being the convex coefficients. Since $A_i[1,1] > A_i[2,1]$, this convex sum is maximized by using $S_i[1,1] = 1$ and $S_i[1,2] = 0$. A similar argument follows for the second independent term in the reward. Thus, the total reward is maximized by the identity strategy matrix. The reward with $S_i = I$ is thus,

$$R_i = \beta\Big[A_i[1,1] + A_i[2,2] - 1\Big]$$

which is strictly positive.

The above analysis implies that whenever worker does solve the tasks, it is her best reporting strategy to report the answer as they are. Now we just need to ensure that investing effort is also the best effort strategy. If the worker doesn't invest effort and report heuristically, the value of the term $\Big(\sum_{k\in[K]} T_i[k,k]\Big) - 1$ is 0. This will be proved in the proof of Theorem 2. However, the worker saves the cost of effort too in this case and she neither earns anything nor loses anything. But when worker invest effort, she earns R_i' from the mechanism and loses C^E in the form of cost of effort. For truthful strategy ($e_i = 1, S_i = I$) to be the dominant uniform strategy, we need the following condition:

$$\beta\Big[A_i[1,1] + A_i[2,2] - 1\Big] - C^E > 0$$

This is true when

$$\beta > \frac{C^E}{A_i[1,1] + A_i[2,2] - 1}$$

For non-binary answer space ($K > 2$), the proof follows similarly assuming $A_i[k,k] > A_i[k',k], \ \forall k' \neq k$. The reward for truthful strategy is $R_i' = \beta\Big[\sum_{k\in[K]} A_i[k,k] - 1\Big]$ which is also strictly positive under the same assumption. □

3.3 Proof of Theorem 2

Proof. As $|Q^i \cap Q^j| \to \infty$, using lemma 1, the reward R_i of any worker i in the Deep Bayesian Trust Mechanism is given by :

$$R_i' = \beta\Big[\Big(\sum_{k\in[K]} T_i[k,k]\Big) - 1\Big]$$

We know that a worker's strategy is called heuristic if either $e_i = 0$ or $e_i = 1$ and S_i has identical rows. In both cases, it is easy to see (using Proposition 1) that the sum of diagonal entries of her trustworthiness matrix sum to 1, which implies,

$$\Big(\sum_{k\in[K]} T_i[k,k]\Big) - 1 = 1 - 1 = 0$$

$\square$

3.4 Proof of Theorem 3

Proof. The proof follows from Lemma 1, which ensures that reward of worker i converges in the limit to

$$\beta \cdot \Big(\sum_{k\in[K]} T_i[k,k]\Big) - 1$$

By definition, $T_i[k,k], \forall k \in [K]$ measure the accuracy of the answers reported by the worker. Thus, even though the mechanism has only estimates of the accuracy of the workers' answers and the estimates are indeed obtained using the answers of the peers and their trustworthiness but the consistency property of these estimates ensures the asymptotic fairness of the mechanism. $\square$

4 Simulation Results with Uniformly distributed Proficiencies

In another simulation setting, $A_i[k,k]\ \forall k \in [K]$ were uniformly distributed in $(\frac{1}{K}, 1]$. Figure 1a compares the reward distribution in the this setting. Observations similar to the former setting ($\beta(5,1)$ distributed proficiencies) are made in this setting as well. The difference is that for the truthful strategy, rewards in the former setting are slightly more skewed towards the positive side. It is an expected observation because the rewards of truthful workers are shown to be an increasing function of their proficiencies in the proof of Theorem 1 and the proficiency distribution in $\beta(5,1)$ is indeed skewed. Figure 1b shows that robustness of our mechanism with respect to the number of shared tasks when proficiencies are uniformly distributed. The observation here is also very similar to that discussed in the paper for $\beta(5,1)$ distributed proficiencies.

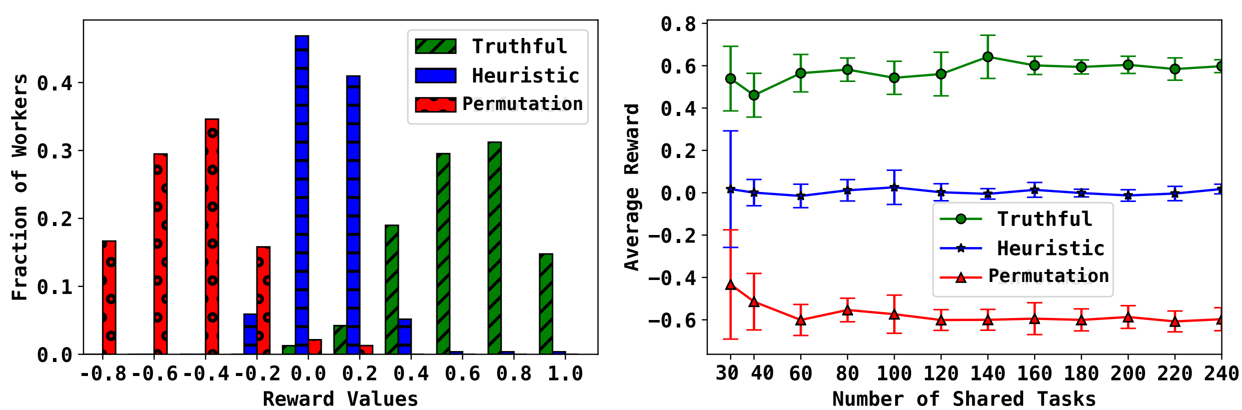

(a) Distribution of rewards for workers with different proficiencies playing different strategies

(b) Average reward of workers playing different strategies under different number of shared tasks

Figure 1: Simulation Results with uniformly distributed proficiencies

5 Amazon Mechanical Turk Study

Any conclusive study on the effects of reward mechanisms requires a large budget and needs to be conducted over a long period. Such a study is beyond the scope of this paper. We performed a limited scale study on Amazon Mechanical Turk to make some preliminary observations about the ease of implementing the mechanism and workers' response to the mechanism. We created some synthetic tasks which resemble tasks requiring human intelligence (natural language understanding) and elicited the answers of the crowd on MTurk with and without our reward mechanism in place. The advantage of using synthetically generated tasks was that we had access to ground truth and could judge the performance of workers with and without our mechanism objectively. The structure of our task (named *'story disentanglement task'*) was the following : we mixed a few sentences from different real news stories into one paragraph and asked the workers to count the number of news stories in the paragraph. Solving this task requires identifying the context of different sentences and whether they are related. The number of sentences in paragraphs was kept independent of the number of stories, making it harder to guess by just looking at the paragraph length. We asked workers to give a binary answer ('Yes' if the number of stories is less than 3 and 'No' otherwise). We also asked them to identify the sentences belonging to different stories. We will discuss only the binary answers of the workers. Each HIT corresponds to giving an answer for one such paragraph. We conducted the experiments under two settings.

- In the first setting, workers were told that they would be paid $0.03$$ per HIT and there would be additional performance based payments without discussing a specific reward rule. We will refer to this setting as the **unspecified** reward setting.
- In the second setting, each HIT was worth $0.03$$ and we explained our Deep Bayesian Trust mechanism to the workers in plain English with almost no use of mathematical language or notations. Figures 2a, 2b and 2c show screen shots of the instructions given to the crowdworkers. We will refer to this setting as the **DB Trust** setting.

In the DB Trust setting, batches of 80 HITs were designed such that each batch had 40 HITs in common with another batch to satisfy peer relationship. In both settings, we had 3 workers giving answers for each paragraph, giving us a total 3×480 HITs from 480 paragraphs. We thus collected a dataset of 1440 worker responses on these HITs, 720 in each setting. In total, 129 workers participated in the experiment. We judge the mechanism on two most important criteria. First, the ability to discourage workers from heuristic reporting and second, the ability to get more accurate answers from crowd.

Observations

1. Figure 3 compares the time workers spent on solving the tasks in the two settings. The fraction of HITs that were given very little time has significantly decreased with our mechanism and the fraction of HITS that were given more time has significantly increased (the green distribution with dots is more skewed towards the right side as compared to the red distribution with slashes, which is more skewed towards the left). This can be interpreted as a success in eliciting effort from the crowd and discouraging low quality/heuristic reporting. We used a browser based JavaScript solution to measure the actual time spent on solving tasks to get tight estimates of time spent in the DB Trust setting, without workers being aware of it. Amazon uses the difference between time of accepting and submitting a HIT as estimates of time spent, which (even after filtering very large values) tend to be highly inflated. As one can see, even with such tight estimates in the DB Trust setting, the time spent by workers is better.
2. The average accuracy of workers was found to increase from 70.86% in the unspecified setting to 79.17% in the DB Trust setting.
3. The average accuracy of all responses was also found to increase from 75.69% to 79.17% with our scheme.

HIT Preview

Story Disentanglement Instructions (Click to collapse)

Please read the instructions below carefully and attempt our task only if you commit to solving the complete task. If you don't read the instructions or spam our system, the task will be rejected and the worker ID will be permanently blocked for our future tasks.

Payment Scheme:

In addition to the fixed payment (80 x $0.03 = $2.4), there will be a performance based bonus. The bonus will be calculated in the following way :

We have prepared a set of batches each containing 80 micro-tasks. Among these batches, there are a few gold batches. The gold batches contain some micro-tasks for which we already know the correct answers. If you are assigned one such batch, your reward will depend on your accuracy on these tasks. The rest of the batches are designed in such a way that they have some tasks in common with one another. If you are assigned this kind of batch, we will determine your accuracy by looking at your answers and the answers obtained on a chain of batches starting from one of the gold batches. For example, if your batch is not a gold batch and has some tasks in common with a gold batch, then we will first determine the accuracy of the worker solving gold batch, then using this accuracy we will determine your accuracy through the answers given on common questions. Thus, irrespective of the type of batch you are assigned, your reward will depend only on your accuracy and nothing else. Your bonus will be proportional to difference of your accuracy and random guessing accuracy. For example, if your accuracy is 50% (which is one way of random guessing), your bonus will be 0. The bonus can also be negative, which means that if your accuracy is low, the negative bonus will reduce your fixed payment of $2.4 as well. It is important to note that accuracy is an aggregate over all the tasks you solve, hence a consistent effort throughout the batch will be required to maintain high accuracy.

The tasks are hosted on an external website. The link is given below in the task. Once you have submitted answers for the entire batch, you will get a unique token. You can submit this token as a solution on Amazon Mechanical Turk. This will let us know that you have solved our tasks and we will process your payment.

(a) Description of Payment Scheme

Sample Task:

In this job, you will be presented with 80 short text paragraphs. Each of these paragraphs have a mix of random parts of several news stories but we don't know exactly how many stories there are in each of the paragraphs. Help us in finding out.

Steps :

1. Read the given text carefully until very last sentence.
2. Determine how many stories are being told in the text.
3. Select 'YES' if you find that there are 1 or 2 stories in the text and select 'NO' if there are 3 stories.
4. Further, mention the last word of each story in the format described in example below. It is important to follow the format, otherwise we will not be able to use your answers.

Rules and Tips :

Each of the paragraphs contain many stories. Stories can be from the same field or different fields such as politics, business, sports, technology and entertainment. If they are from different fields, you may find it comparatively easier to count the stories but in general there are no guarantees that the stories will be from different fields. Any story can end abrubptly in the paragraph and a new story can begin. Any story in the paragraph can begin from any random sentence of the original story. Hence, reading every sentence with full attention is required. There is no correlation between the number of stories and the length of the paragraph. A paragraph with 2 stories can be of the same length as a paragraph with 3 stories.

(b) Task Instructions

Example :

Consider the following text paragraph :

```
The panel suggested different kinds of murders could be "graded" to recognise the seriousness of the offence. But he argued the minister shoul
d have volunteered a formal statement instead of having to be "dragged" to face MPs. Sound distribution of the cash could cut poverty levels t
o 36% from 53%, the government believes. Colombia has a population of about 44 million and half lives below poverty line. The seventh seed, wh
o has never gone beyond the quarter-finals in the year's first major and is lined up to meet Roddick in the last eight, is looking forward to
the match.
```

This paragraph contains parts of 3 different news stories as shown below. The first story is about some court case. The second story is about demographics in Colombia. The third story is about about an sports match. Hence, you should select 'No'.

```
The panel suggested different kinds of murders could be "graded" to recognise the seriousness of the offence. But he argued the minister shoul
d have volunteered a formal statement instead of having to be "dragged" to face MPs.

Sound distribution of the cash could cut poverty levels to 36% from 53%, the government believes. Colombia has a population of about 44 millio
n and half lives below poverty line.

The seventh seed, who has never gone beyond the quarter-finals in the year's first major and is lined up to meet Roddick in the last eight, is
 looking forward to the match.
```

The last word of each of the stories are : MPs, line, match. Hence, you should write :

```
MPs, line, match
```

in the text field. You can ignore any special characters such as " or . while writing the last word.

(c) Sample Task

Figure 2: Mechanical Turk Task

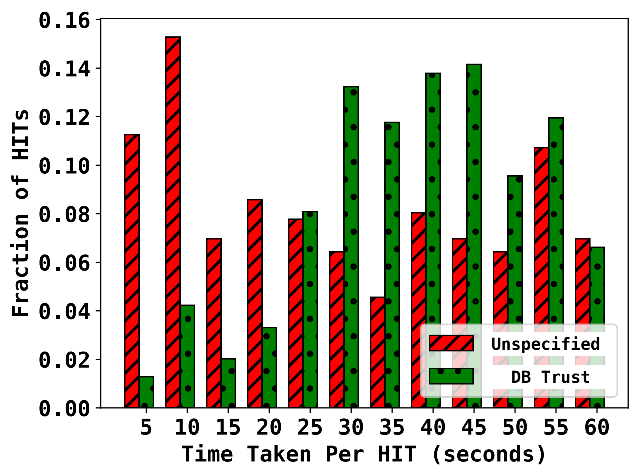

Figure 3: Time Spent on HITs