Peer prediction markets to elicit unverifiable information

Aurélien Baillon[1], Cem Peker*[2], and Sophie van der Zee[3]

[1,2,3]Erasmus School of Economics, Erasmus University Rotterdam

October, 2022

Abstract

This paper introduces an incentive mechanism to elicit subjective judgments. We consider binary questions where responses cannot be verified for accuracy. Self-reported information may not reflect the truth due to lack of cognitive effort or motives to lie. In our formal framework, agents choose whether to receive a costly private signal, which lead them to endorse either "yes" or "no" as an answer. Then, they either buy or sell a single unit of an asset, whose value is determined by the endorsement rate of "yes" answers. The price of the asset is set at the prior expectation of the endorsement rate. We obtain a separating equilibrium, where agents choose to receive a costly private signal and buy or sell the asset as a function of their signal. Trades reflect agents' true unverifiable information. Two experimental studies test the theoretical results. The first study shows that peer prediction markets motivate agents to seek costly information and reveal it in a simple prediction task. The second study implements an online field experiment to demonstrate feasibility in a natural setting. We find that agents in a peer prediction market are more likely to truthfully self-report socially stigmatized behavior that they can easily deny without being caught.

*Corresponding author: Cem Peker, acpeker@gmail.com

1 Introduction

"Have you stood less than 6 feet apart from another person in a queue yesterday?" Health surveys often require respondents to recollect past experiences. This experience can be seen as a private signal that a respondent acquire by exerting effort (recalling, to their mind, what they did a day earlier.) But how can we ensure that the respondents will, first, provide such effort and then, answer truthfully if there is no way to compare their answer to some truth?

Starting with Crémer and McLean (1988), the mechanism design literature has explored ways to reveal private signals. Miller et al. (2005), and more broadly the peer-prediction literature (Witkowski and Parkes, 2012b, 2013; Liu and Chen, 2017a), have proposed solutions exploiting the informativeness of a respondent's answer in predicting their peers' answers. For instance, imagine that we have some prior expectations about the rate of yes answers to the 6-feet-apart question. A respondent answering yes increases our expectations about the proportion of *other* people answering yes. Formally, this increase is a simple application of Bayesian updating when respondents draw a private signal (yes/no), with unknown probability p of yes signals: a yes signal makes higher values of p more likely than initially believed.[1] Intuitively, the yes answer to the 6-feet-apart question can suggest that others also had difficulty complying with a social distancing guidelines.

In this paper, we propose and implement a novel solution to incentivize private signals acquisition and revelation: a peer-prediction market (PPM). In a PPM, yes respondents are rewarded with the formula "yes answer rate - prior expectations of yes answer rate". Those who answer no get the opposite reward. If there are fewer yes answers than expected, yes respondents get a negative reward while no respondents get a positive one. Equivalently, a PPM can be presented as yes (no) respondents buying (selling) a single asset, the value of which is eventually determined by the proportion of yes answers. The price is set to the prior expectations. In a situation in which the yes-answer rate is expected to follow a random

[1] We assume here that signals are conditionally independent, i.e. independent given the probability of success. The probability of success is assumed to be itself drawn from a non-degenerate distribution over $(0, 1)$.

walk, a repeated PPM can be implemented in which the price at period t is the value of the asset at $t-1$.

First, we show that signal acquisition and truthful revelation is a Bayesian Nash equilibrium, providing a partial-implementation solution to the static problem. Our solution is minimal, in the sense that it does not ask respondents to provide more than their answer and it does not require the surveyor to share more than prior expectations with the respondents. We then extend our analysis to incorporate psychological costs, capturing the possible (mild) 'shame' of reporting an answer and potential deception costs.

Second, we test PPM in an online experiment closely following the theoretical model: respondents may exert an effort (i.e., complete a real-effort task borrowed from the experimental economics literature) to obtain a signal and report it; or they may simply answer randomly. We compare PPM with two benchmarks: flat fee (no incentives) and accuracy incentives (incentives when the signal generation process is observable). The latter is not applicable in surveys, where such process in unobservable but it provides a gauge for the effect of PPM. A flat fee decreases the effort rate by about 20 percentage points with respect to accuracy incentives. PPM allows us to recover half of this difference.

Third, we demonstrate feasibility in a natural setting. We implement the repeated PPM in the context of a health survey, involving questions of the 6-feet-apart type. The asset price is set to the previous week yes-rate. We hypothesize that people not exerting recollection efforts or feeling some slight shame for not complying with health guidelines are likely to deny having experienced such situation, and therefore that PPM will trigger higher rate of yes answers than a flat fee. We indeed obtain that more people admit experiencing situations in contradictions with health guidelines in the PPM treatment than in the flat fee treatment. This second study is, in nature, more exploratory and we cannot exclude alternative explanations. However, it shows that PPM can be applied to socially relevant questions, where psychological costs of reporting non-compliance may be present.

Related literature - PPMs offer a market-based solution to the problem of incentivizing effort in information elicitation without verification (Waggoner and Chen, 2013). Previous work introduced peer prediction mechanisms that consider the truthful elicitation problem only, and does not explicitly incorporate costly effort. The original peer prediction method (Miller et al., 2005) can be adjusted for costly effort via re-scaling of payments. However, the researcher has to know the full common prior belief of participants. Bayesian truth serum (Prelec, 2004) and its variants (Witkowski and Parkes, 2012a; Radanovic and Faltings, 2013, 2014) do not require any knowledge of priors. But, as a result, the researcher lacks information to scale rewards appropriately for costly effort. Other authors developed peer prediction mechanisms to incentivize effort in crowdsourcing tasks in which ground truth is unverifiable. Such mechanism rely on additional structure on agents' proficiency (Witkowski et al., 2013), multiple tasks (Dasgupta and Ghosh, 2013; Radanovic et al., 2016; Shnayder et al., 2016) or a dynamic framework (Liu and Chen, 2017b). Similar to the original peer prediction method, PPM is one-shot and 'minimal' (Witkowski and Parkes, 2013): agents complete a single task only. But, PPM requires less information on priors. Furthermore, PPM offers a simpler solution in binary problems compared to other peer prediction mechanisms with costly effort. The present paper is also the first of this stream of literature to include both cost of efforts and psychological costs in the model.

Closest to PPMs are Bayesian markets (Baillon, 2017), which provide a market solution to binary elicitation problem in a similar Bayesian setup to ours, except that information is not costly. Moreover, unlike PPM, an agent first reports her answer. She can later buy (sell) one unit of the asset only if she reported 'yes' ('no'). Price is determined randomly afterwards, so the agents decide on trade options before price is observed. In equilibrium, agents report their true judgments to be eligible for their desired trade. In the way they are set-up, PPMs aim to be closer to prediction markets than Bayesian markets are. Agents can trade freely, according to their private information, at a pre-specified price.

2 Theory

2.1 Agents and their information

A *center* (a researcher, a survey company) is interested in eliciting N *agents*' informed answers to a question Q, with possible answers $\{0,1\}$. Agents can answer randomly at no cost but they may also decide to provide an effort (thinking, remembering, looking for information, etc.) to obtain their informed answer. Formally, we model the informed answer as a *signal* $\tau_i \in \{0,1\}$, which agent $i \in \{1,\ldots,N\}$ can obtain by providing *effort* $e_i = 1$ at a cost $c_i > 0$ (expressed in monetary terms). The cost of no effort ($e_i = 0$) is 0. The probability of getting signal 1 is the same for all agents (hence, it is independent of the effort cost) but is unknown. We model it as a random variable ω over $[0,1]$. Denoting $\tau = (\tau_1,\ldots,\tau_N)$, a *state of nature* is thus a realization of ω and τ, with the *state space* being $\Omega = [0,1]\times\{0,1\}^N$.The probability space is (Ω,Σ,P), with Σ the Borel $\sigma-$algebra of Ω and we assume that P is countably additive. The next assumption describes the full signal technology.

Assumption 1 (Signal technology)**.** *The signal technology is such that for all* $i,j \in \{1\ldots,N\}$, $i \neq j$, *and* $o \in [0,1]$*:*

1. $P(\tau_i = 1|\omega = o) = o$*;*

2. $P(\tau_i = 1|\tau_j, \omega = o) = o$*;*

3. *and* $P(\omega)$ *is continuous over* $[0,1]$.

Part 1 of Assumption 1 states that the signal technology is anonymous, part 2 that it satisfies *conditional independence*, and part 3 that no value of ω has a probability mass. The latter excludes degenerate cases in which all agents could get the same signal for sure or in which ω would be known.

Let P_i represent the belief of agent i about the signal technology, and P_0 that of the center. It is common to assume $P_i = P_0 = P$ in peer prediction mechanisms.[2]. We allow agents to have different opinions on how likely various values of ω are but the following assumption restrict their belief in two ways.

Assumption 2 (Unbiased prior expectations)**.** *For all $i \in \{0, \ldots, N\}$, P_i satisfies properties 1-3 of Assumption 1 and $E_i(\omega) = E(\omega)$.*

Assumption 2 states that all agents and the center agree on the main properties of the signal technology and share the same prior expectation. It is a strong assumption, despite relaxing the often-used common prior assumption. Assumption 2 is plausible if (i) question Q is new and people have no reason to believe that answer 1 is more likely than answer 0, i.e., $E(\omega) = 0.5$; or (ii) signals of another group of agents have been publicly revealed (possibly with another mechanism); or (iii) the agents have no clue about ω but the center shares her prior expectation. In case (i), we do not need to assume uniform P_i over the possible values of ω; e.g., it can be bell-shaped for some agents. Case (ii) can correspond to situations in which question Q was asked in the past (to other agents) but the center and the (new) agents do not know whether the signal distribution will be exactly the same. For instance, imagine that, a month ago, it was published that 73% of people reported they could always stay 6 feet away from others. There are many reasons why this proportion might change but before agents try to remember their own experience, 73% is a good average guess about what others will answer. Let us denote $\bar{\omega} \equiv E(\omega)$, $\bar{\omega}_i^0 \equiv E_i(\omega|\tau_i = 0)$ and $\bar{\omega}_i^1 \equiv E_i(\omega|\tau_i = 1)$.

Lemma 1. *Under Assumptions 1 and 2, for all $i \in \{1, \ldots, N\}$, $0 < \bar{\omega}_i^0 < \bar{\omega} < \bar{\omega}_i^1 < 1$.*

Proof. First part 3 of Assumption 1 excludes $\bar{\omega} \in \{0, 1\}$.

Second, $P_i(\tau_i = 1) = \int_0^1 P_i(\tau_i = 1|\omega = o) \times P_i(\omega = o)do = \int_0^1 o \times P_i(\omega = o)do = E_i(\omega) = \bar{\omega}$. $\bar{\omega}_i^1 = \int_0^1 \frac{P_i(\tau_i=1|\omega=o)\times P_i(\omega=o)\times o}{P_i(\tau_i=1)}do = \int_0^1 \frac{o^2\times P_i(\omega=o)}{\bar{\omega}}do > \bar{\omega}$ because $\int_0^1 o^2 \times P_i(\omega = o) > \left(\int_0^1 o \times P_i(\omega = o)\right)^2 = \bar{\omega}^2$ by Jensen's inequality applied to the convex squared function and

[2]Or $P_i = P$ with no assumption on P_0 in the Bayesian truth-serum of (Prelec, 2004) or Bayesian markets of (Baillon, 2017)

the inequality is strict because we degenerate cases were excluded by Part 3 of Assumption 1, which also excludes a posterior expectation of 1. The proof of $0 < \bar{\omega}_i^0 < \bar{\omega}$ is symmetric. □

Lemma 1 shows that under our assumptions, all agents receiving signal 1 have higher expectations about ω than they had ex ante (and than the center) whereas agents with signal 0 decrease their expectations. Finally, we make the following assumption on agents' risk preferences:

Assumption 3 (Risk neutrality)**.** *Agents are risk neutral.*

Assumption 3 implies that agents maximize their expected payoffs. Section 2.2 introduces a market mechanism to exploit the difference in expectations established in Lemma 1. Assumption 3 suggests that agents' optimal strategy will not depend on a risk parameter.

2.2 The Market

The center implements a *peer-prediction market* for Q, in which an asset is traded whose value will be the proportion of agents reporting 1 as answer for Q multiplied by π, a scaling constant. If the currency is the dollar, $\pi = 10$ means that the asset is worth \$5 if 50% of the agents report 1.

Definition 1. *A peer-prediction market is defined by the following steps:*

1. *The center announces the asset price $\bar{\omega}\pi$.*

2. *Agents simultaneously choose a report $r_i \in \{0, 1\}$. Those who report 1 become buyers of the asset and those who report 0 become sellers.*

3. *The center computes the asset value $\bar{r}\pi = \frac{\pi}{N} \sum_{i=1}^{n} r_i$.*

4. *If $\bar{r} = 0$ or $\bar{r} = 1$, the market is stopped; no payment occurs.*

5. *Otherwise, buyers pay $\bar{\omega}\pi$ to the center in exchange of $\bar{r}\pi$ and sellers receive $\bar{\omega}\pi$ from the center in exchange of $\bar{r}\pi$.*

In a peer-prediction market, reporting a 1 answer ($r_i = 1$) is equivalent to betting that the proportion of 1 answers will be higher than $\bar{\omega}$, that is, buying the asset. Symmetrically, reporting a 0 answer is a bet on a proportion of 1 answers lower than $\bar{\omega}$. Step 5 specifies that all trades are made with the center, and not directly between agents. Direct trading would lead to complications such as the no-trade theorem (Milgrom and Stokey, 1982): knowing that someone wants to sell informs the buyer that someone received a 0 signal, and conversely. Ultimately, agents who report 1 get $(\bar{r} - \bar{\omega})\pi$ and those who report 0 get $(\bar{\omega} - \bar{r})\pi$. The center subsidies the market if need be. The agents subtract c_i from their earnings if they provided an effort.

2.3 Strategies and Equilibria

The agents' strategies in the peer-prediction market involve first deciding whether to exert an effort, and then what to report. We will consider mixed strategies only in reports, so agent i's strategy is given by (e_i, R_i, R_i^0, R_i^1) with R_i, R_i^0, and R_i^1 the probabilities of $r_i = 1$ if $e_i = 0$, if $e_i = 1$ and $\tau_i = 0$, and if $e_i = 1$ and $\tau_i = 1$ respectively. The strategy space is thus $\{0, 1\} \times [0, 1]^3$. The center is interested in situations in which agent i exerts an effort and answers truthfully, i.e., $e_i = 1$, $R_i^0 = 0$, and $R_i^1 = 1$. We need to make one final assumption, about what agents know about each others.

Assumption 4 (Common knowledge)**.** *The peer-prediction market functioning, the strategy space, the signal technology, the beliefs P_i, the costs c_i and agents' risk neutrality are common knowledge.*

Assumption 4 ensures that we have specified all the elements of a *Bayesian game*, as defined by (Osborne and Rubinstein, 1994, Definition 25.1). If beliefs and costs were not common knowledge, we would have to define higher-order beliefs, which would complicate the proofs. As we will see below the crucial part is not so much that agents know the exact beliefs of everyone, but rather that all agents know that Lemma 1 holds. Again for convenience, we let $N \to \infty$. It allows us to assimilate signal probability with signal proportion. It also allows us to neglect the impact of a single agent on the asset value.

Proposition 1. *Under Assumptions 1 to 4 and with N infinite, if $c_i > \pi$ for all $i \in \{1, \ldots, N\}$, then Nash equilibria are characterized by $e_i = 0$ and $R_i \in \{0, \bar{\omega}, 1\}$. Expected payoffs are 0.*

Proof. Possible earnings $(\bar{r} - \bar{\omega})\pi$ and $(\bar{\omega} - \bar{r})\pi$ are both strictly lower than π, and therefore than c_i if $c_i > \pi$. There are no incentives to provide efforts; hence, $e_i = 0$. Consider agent i and assume all other agents $j \neq i$ have the same probability to report 1 ($R_j = R$ for some $R \in [0, 1]$). Hence, with N infinite, the asset value $\bar{r}$ is R. Agent i hence expects to earn $[R_i \times (R - \bar{\omega}) + (1 - R_i) \times (\bar{\omega} - R)] \times \pi$. If $R \in (\bar{\omega}, 1]$, then $R_i = 1$ is optimal. If $R \in [0, \bar{\omega})$, then $R_i = 0$ is optimal. Finally, if $R = \bar{\omega}$, then any $R_i \in [0, 1]$ is optimal. Nash equilibria require $R_i = R$ such that no one has incentives to deviate. Hence, we must have either $R_i = 1$ for all i, or $R_i = 0$ for all i, or $R_i = \bar{\omega}$ for all i. In all these cases, earnings are 0 (remember that if $\bar{r} = 0$ or 1, no payoffs occur as specified in step 4 of Definition 1. □

Proposition 2. *Under Assumptions 1 to 4 and with N infinite, if for all $i \in \{1, \ldots, N\}$ $\frac{c_i}{\pi} < \bar{\omega} \times (\bar{\omega}_i^1 - \bar{\omega}) + (1 - \bar{\omega})\,(\bar{\omega} - \bar{\omega}_i^0)$, providing an effort and reporting truthfully ($e_i = 1$, $R_i^0 = 0$, and $R_i^1 = 1$) is a Nash equilibrium, and it strictly dominates the no-effort equilibria.*

Proof. Let us consider agent i's view point and assume $e_j = 1$, $R_j^0 = 0$, and $R_j^1 = 1$ for all $j \neq i$. Without any signal, agent i's expected earnings are $[R_i\,(E_i(\omega) - \bar{\omega}) + (1 - R_i)\,(\bar{\omega} - E_i(\omega))]$ $\times\, \pi = 0$ by Assumption 2.

With signal 1, agent i's expected earnings are $[R_i^1\,(\bar{\omega}_i^1 - \bar{\omega}) + (1 - R_i^1)\,(\bar{\omega} - \bar{\omega}_i^1)] \times \pi$. By Lemma 1, this is maximum for $R_i^1 = 1$, yielding $(\bar{\omega}_i^1 - \bar{\omega}) \times \pi > 0$.

With signal 0, agent i's expected earnings are $[R_i^0\,(\bar{\omega}_i^0 - \bar{\omega}) + (1 - R_i^0)\,(\bar{\omega} - \bar{\omega}_i^0)] \times \pi$. By Lemma 1 again, this is maximum for $R_i^0 = 0$, yielding $(\bar{\omega} - \bar{\omega}_i^0) \times \pi > 0$.

Before getting a signal, the expected gain is therefore,

$$\left[P_i(\tau_i = 1) \times \left(\bar{\omega}_i^1 - \bar{\omega}\right) + P_i(\tau_i = 0)\left(\bar{\omega} - \bar{\omega}_i^0\right)\right] \times \pi = \left[\bar{\omega} \times \left(\bar{\omega}_i^1 - \bar{\omega}\right) + (1-\bar{\omega})\left(\bar{\omega} - \bar{\omega}_i^0\right)\right] \times \pi.$$

This is strictly positive by construction and strictly more than c_i by assumption. Hence, the net earnings (once the costs are subtracted) are strictly positive and providing an effort is worth it. As a consequence, $e_i = 1$, $R_i^0 = 0$, and $R_i^1 = 1$ is a Nash equilibrium.

Finally, let us consider the case in which all agents but i provide no efforts and report 1 with probability R. The expected earnings are

$$\begin{cases} [R_i^1 \times (R - \bar{\omega}) + (1 - R_i^1) \times (\bar{\omega} - R)] \times \pi & \text{with signal 1} \\ [R_i^0 \times (R - \bar{\omega}) + (1 - R_i^0) \times (\bar{\omega} - R)] \times \pi & \text{with signal 0} \\ [R_i \times (R - \bar{\omega}) + (1 - R_i) \times (\bar{\omega} - R)] \times \pi & \text{with no signal.} \end{cases}$$

As in Proposition 1, the only equilibria must be of the form $R_i = R \in \{0, \omega, 1\}$, and by similar arguments $R_i^1 = R_i^0 = R \in \{0, \omega, 1\}$. The earnings are always 0 and the net earnings with effort are even strictly negative. Hence, $e_i = 0$, $R_i \in \{0, \omega, 1\}$ is also a Nash equilibrium (with $R_i^1 = R_i^0 = R_i$) but it is dominated by the equilibrium with effort and truthful reporting ($e_i = 1$, $R_i^0 = 0$, and $R_i^1 = 1$). $\square$

Proposition 3. *Under Assumptions 1 to 4 and with N infinite, if for $T \times 100\%$ of the agents $\frac{c_i}{\pi} > \bar{\omega} \times (T\bar{\omega} + (1-T)\bar{\omega}_i^1 - \bar{\omega}) + (1-\bar{\omega})\left(\bar{\omega} - T\bar{\omega} - (1-T)\bar{\omega}_i^0\right)$ and the inequality is reversed for the remaining agents, then there is a Nash equilibrium in which these $T \times 100\%$ will exert no efforts and report 1 with probability $\bar{\omega}$ and where the other agents exert efforts and report truthfully.*

Proof. First, let us assume that all agents but i play the strategy described in the proposition. With signal 1, agent i expects the asset value to be $T\bar{\omega} + (1-T)\omega_i^1$, and with signal 0 $T\bar{\omega} + (1-T)\omega_i^0$. By Lemma 1, $T\bar{\omega} + (1-T)\omega_i^0 < \bar{\omega} < T\bar{\omega} + (1-T)\omega_i^1$, and with the same argument as in the proof of Proposition 2, it is best to report truthfully $R_i^0 = 0$ and $R_i^1 = 1$. Ex ante, the expected benefit of exerting an effort is therefore $[\bar{\omega} \times (T\bar{\omega} + (1-T)\bar{\omega}_i^1 - \bar{\omega}) + (1-\bar{\omega})(\bar{\omega} - T\bar{\omega} - (1-T)\bar{\omega}_i^0))]\pi - c_i$.

If $\frac{c_i}{\pi} \leq \bar{\omega} \times (T\bar{\omega} + (1-T)\bar{\omega}_i^1 - \bar{\omega}) + (1-\bar{\omega})(\bar{\omega} - T\bar{\omega} - (1-T)\bar{\omega}_i^0)$ then $e_i = 1$ is optimal.

If $\frac{c_i}{\pi} > \bar{\omega} \times (T\bar{\omega} + (1-T)\bar{\omega}_i^1 - \bar{\omega}) + (1-\bar{\omega})(\bar{\omega} - T\bar{\omega} - (1-T)\bar{\omega}_i^0)$, an effort leads to negative net earnings, whereas exerting no efforts gives $[R_i \times (T\bar{\omega} + (1-T)E_i(\omega) - \bar{\omega}) + (1-R_i)(\bar{\omega} - T\bar{\omega} - (1-T)E_i(\omega))]\,\pi = 0$ because of the common prior expectations assumption. Hence, $e_i = 0$ and $R_i = \bar{\omega}$ is a best response in this case. □

In the equilibrium of Proposition 3, the $T\%$ of agents not providing an effort have negative externalities on others by decreasing the extent to which the asset value can differ from the prior expectations. This reduces the value of providing an effort for everyone.

2.4 Psychological costs

So far, we have only considered effort costs. In this subsection, two additional costs are considered:

- *Asymmetric reporting cost:* The cost $a_i \geq 0$ borne by agent i when reporting $r_i = 1$ per se, no matter whether the agent receives a signal and what this signal may be. We choose 1 arbitrarily, and without loss of generality. This cost can reflect a stigma associated with answer 1. As we will see in the theoretical results and later in the experimental applications, a_i should not be too high, thereby excluding major incentives to lie.

- *Deception cost:* The cost $d_i \geq 0$ of reporting $r_i = 0$ after receiving signal $\tau_i = 1$ or reporting $r_i = 1$ after receiving signal $\tau_i = 0$. This cost captures people's tendency to tell the truth. We assume that such costs can only occur when a signal has been received because cost for reporting an answer in spite of having no signal would be equivalent to decreasing the effort costs.

Assumption 5. *Agents bear asymmetric reporting costs* $a_i \geq 0$ *and deception costs* $d_i \geq 0$ *and these costs are common knowledge.*

Proposition 4. *Under Assumptions 1 to 5 and with* N *infinite, if for all* $i \in \{1, \ldots, N\}$ $\frac{c_i}{\pi} < \bar{\omega} \times \left(\bar{\omega}_i^1 - \bar{\omega} - \frac{a_i}{\pi}\right) + (1 - \bar{\omega})\left(\bar{\omega} - \bar{\omega}_i^0\right)$ *and* $\frac{a_i}{\pi} < \frac{d_i}{\pi} + 2\left(\bar{\omega}_i^1 - \bar{\omega}\right)$*, providing an effort and reporting truthfully (*$e_i = 1$*,* $R_i^0 = 0$*, and* $R_i^1 = 1$*) is a Nash equilibrium, and it strictly dominates the no-effort equilibrium.*

Proof. Let us consider agent i's view point and assume $e_j = 1$, $R_j^0 = 0$, and $R_j^1 = 1$ for all $j \neq i$. Without any signal, agent i's expected earnings are

$$\left[R_i\left(E_i(\omega) - \bar{\omega} - \frac{a_i}{\pi}\right) + (1 - R_i)\left(\bar{\omega} - E_i(\omega)\right)\right] \times \pi \leq 0.$$

With signal 1, agent i's expected earnings are

$$\left[R_i^1\left(\bar{\omega}_i^1 - \bar{\omega} - \frac{a_i}{\pi}\right) + (1 - R_i^1)\left(\bar{\omega} - \bar{\omega}_i^1 - \frac{d_i}{\pi}\right)\right] \times \pi - c_i.$$

This is maximum for $R_i^1 = 1$, because $\frac{a_i}{\pi} < \frac{d_i}{\pi} + 2\left(\bar{\omega}_i^1 - \bar{\omega}\right)$. With signal 0, agent i's expected earnings are

$$\left[R_i^0\left(\bar{\omega}_i^0 - \bar{\omega} - \frac{a_i}{\pi} - \frac{d_i}{\pi}\right) + (1 - R_i^0)\left(\bar{\omega} - \bar{\omega}_i^0\right)\right] \times \pi - c_i.$$

This is maximum for $R_i^0 = 0$. Before getting a signal, the expected payoff is therefore, $\left[\bar{\omega} \times \left(\bar{\omega}_i^1 - \bar{\omega} - \frac{a_i}{\pi}\right) + (1 - \bar{\omega})\left(\bar{\omega} - \bar{\omega}_i^0\right)\right] \times \pi - c_i$. This is strictly positive by assumption. Hence, providing an effort is worth it. As a consequence, $e_i = 1$, $R_i^0 = 0$, and $R_i^1 = 1$ is a Nash equilibrium.

Finally, let us consider the case in which all agents but i provide no efforts and report 0 (as in Proposition 1). The best agent i can do is to provide no effort and report 0 as well, yielding expected earnings 0, which is dominated by truth-telling. □

Proposition 4 establishes two sufficient conditions for the existence of a truth-telling equilibrium. The first one, as in Proposition 2, ensures that the expected payoffs with effort is higher than with no effort. The second one ensures that the cost of reporting the stigmatizing answer does not exceed the benefit of truth-telling. This benefit is twofold: the agent does not lie (so no deception costs d_i) and buys the asset instead of having to sell it. This leads to three remarks. First, costs of reporting a stigmatizing answers are moderated by the cost of lying. Second, if $\frac{a_i}{\pi} > \frac{d_i}{\pi} + 2\left(\bar{\omega}_i^1 - \bar{\omega}\right)$, the corresponding agent will anticipate to never report 1 anyhow and therefore, has no incentives to provide an effort. In other words, in our model, conscious lying has no reason to occur because agent will simply prefer not to get a signal and report the more acceptable answer. Third, a higher π is useful to both stimulate effort and reduce incentives to lie.

3 Experimental Evidence

Section 2 established the existence of an equilibrium agents in a PPM seek costly information and make informed trades. An agent's incentives in trading are based on her peers' behavior, as value of the asset is determined by other agents' trades. Are such peer-based incentives effective in eliciting effort in practice? This section presents evidence from two experimental studies. Section 3.1 provides a brief overview of the two studies and the findings. Sections 3.2 and 3.3 provide detailed information on the two studies and present the results.

3.1 Overview

We run two experimental studies to test if PPM elicit effort in judgment formation. Study 1 aims to test PPM in a controlled setting. We recruit participants for an online experiment where they are presented with pairs of virtual boxes, containing yellow and blue balls of unknown proportions. In each pair, one of the boxes is the 'actual box' with equal probability. Participants are asked to pick a box within each pair. Before making a pick, each participant could independently draw a single ball from the actual box if she completes a real effort task, which involves counting the number of zeroes in a binary matrix. In this design the actual box is known to the experimenter, implying that the information is verifiable. Testing the PPM in a verifiable task allows us to implement incentives for ex-post accuracy as a benchmark. Study 1 consists of three experimental conditions in which participants complete the same task. The control condition offers fixed rewards (a flat participation fee) while the two treatments implement PPM incentives and incentives for ex-post accuracy. Results suggest that the PPM elicit significantly more effort than fixed rewards while the effort is highest under incentives for ex-post accuracy. As discussed before, ex-post accuracy is not observable in practical elicitation problems of unverifiable information. The results of Study 1 suggest that the PPM are effective when ex-post rewards are not feasible.

Study 2 explores the feasibility of PPM in a practical problem of elicitation of unverifiable information. In response to the Covid-19 pandemic in 2020, governments around the world issued guidelines for social distancing and other safe practices. Policy makers would like to know if such guidance is followed by the public. When asked to self-report if they were following a safe practice, people may not recall instances where they failed to do so. In addition, people may be reluctant to admit unsafe practices due to the social stigma associated with such anti-social behavior. We implement the PPM in an online survey aimed at the residents of the UK. Participants are asked 8 questions, each involving an unsafe practice according to the Covid-19 guidance issued by the UK government. We find that under the PPM incentives, participants are more likely to admit not following the guidance and they took longer to respond on average. Study 2 allows us to test the PPM in a setup where psychological costs as well as effort costs are relevant.

3.2 Study 1 - PPM in a simple prediction task

3.2.1 Design and procedures

Tasks. Participants complete 10 *prediction tasks*. Each prediction task displays a pair of boxes as shown in Figure 1 below. There are 10 such pairs and each pair appears in a single prediction task only. One of the boxes in each pair is set as the 'actual box' via a coin flip prior to the experiment. Participants are informed that one of the boxes is the actual box, but they do not know which. In each task, participants are asked to pick one of the boxes, which may affect their rewards depending on the experimental condition.

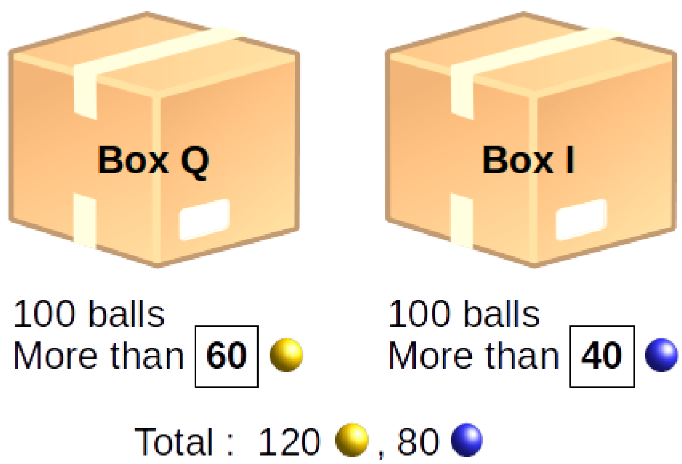

Figure 1: An example pair of boxes

In Figure 1, there are 120 yellow and 80 blue balls in total. Box Q contains more than 60 yellow balls while Box I contains more than 40 blue balls. The exact number of balls of each color are determined randomly according to the specifications. So, the number of yellow balls in Box Q is within $(60, 100]$. For example, if Box Q contains 80 yellow and 20 blue balls, Box Z contains 40 yellow and 60 blue balls. In the experiment, pairs of boxes are presented as shown in Figure 1. Thus, participants do not know the exact number of yellow and blue balls in a box. The boxes are constructed such that the left box (Box Q in Figure 1) always contains more than half of the total number of yellow balls. All 10 pairs are included in the supplemental material.

Before picking a box, each participant is offered a choice to observe a single draw from the actual box with replacement. Participants have to complete a *real effort task* to observe their draw. The effort task is counting the number of 0s in a matrix. Figure 2 shows one such matrix. There is a unique matrix for each effort task and there is a single effort task associated with each prediction task. The number of 0s in each matrix varies between 8 and 16.

$$\begin{array}{cccccc} 0 & 0 & 1 & 1 & 0 & 1 \\ 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 0 & 1 \end{array}$$

Figure 2: An example binary matrix

The sequence of events in each prediction task is as follows: First, participants are shown a pair of boxes and asked if they want to complete the effort task. If a participant skips the effort task, she is immediately asked to pick a box. Otherwise, she is presented the associated binary matrix and asked to report the number of 0s. The participant is required to report an accurate count to proceed. If the participant reports an inaccurate count, she is allowed unlimited number of retries until she reports an accurate count. Upon reporting the accurate count, the participant observes a random draw, which is either a blue or a yellow ball. Then, she proceeds to picking a box.

The prediction task is a representation of the binary question Q, where the two boxes in any pair correspond to the possible answers. The effort task corresponds to the costly signal in our framework. Participants are allowed to skip the effort task, in which case they make a pick without observing a draw. In any given pair, the total number of yellow (and blue) balls are known and boxes are a priori equally likely to be the actual box, which induces a common prior expectation on the number of yellow balls in the actual box. For example, the common prior expectation on yellow in Figure 1 is 60. If a participant draws a yellow (blue) ball, her posterior probability on left (right) box being the actual box is higher. An agent's best guess on the actual box matches with her draw and hence, corresponds to her type. Thus, a participant's draw is effectively the signal that fully determines her type.

Design. We set up three experimental conditions which differ only in reward structure. In the *flat* condition, participants receive a fixed reward of £3.25 for completing the experiment. In the *accuracy* treatment, participants receive a basis reward of £3.25. In addition, they earn £0.20 per accurate pick and lose £0.20 per inaccurate pick, where the accurate pick in a pair is picking the actual box. Thus, a participant's total reward is within $[£1.25, £5.25]$. The *PPM* treatment implements the PPM. Similar to the accuracy treatment the basis reward is £3.25. In addition, participants may earn a bonus from each pick which is determined by her peers' picks in the same pair and composition of the boxes. To illustrate, consider a participant who is asked to pick a box in the pair shown in Figure 1. Suppose, among all other participants, 82% picked Box Q and 18% picked Box I. Then, the participant earns $82 - 60 = 22$p if she picked Box Q, loses $40 - 18 = 22$p if she picked Box I. The number within the square below each box is serves as a threshold. The participant earns a positive bonus from her pick if the percentage of others who pick the same box in that pair exceeds the threshold of that box.

Rewards in the PPM treatment represent the incentives in a PPM. Consider the pair of boxes given in Figure 1. The actual box is either Box Q or Box I with equal probability. Prior expectation of a participant on the number of yellow balls is 60. Suppose the participant chooses to complete the effort task and draws a yellow ball. Her posterior probability on Box Q being the actual box is higher, which has two implications: i) her best guess on the actual box is Box Q, and ii) her posterior expectation on the number of yellow balls in the actual box is greater than 60. Then, the participant expects more than 60% of her peers to draw yellow and consider Box Q more likely as well. In a situation where all others pick the box they consider more likely, the participant expects more than 60% of her peers to pick

Box Q, resulting a positive expected bonus from picking Box Q herself. Vice versa holds for a participant who draws a blue ball. This setup is analogous to a PPM with $p = \omega_0 = 0.6$, where the differing best guesses of participants who draw different colors correspond to the types. Trades are represented by picks in the prediction task. Recall that the left (right) box in each pair contains more than the prior expectation on the number of yellow (blue) balls. The truthful strategy corresponds to a subject completing the effort task followed by picking the left (right) box if her draw is a yellow (blue) ball.

Participants in the flat condition have no direct financial incentives to complete the effort tasks as their reward does not depend on prediction accuracy. In contrast, rewards in the accuracy condition are determined by prediction accuracy. Thus, participants in the accuracy condition could be expected to complete effort tasks more frequently to maximize their accuracy. The PPM condition also provides incentives to complete effort tasks if, as predicted by the theory, participants consider their signal informative on others' picks. We could observe more effort task completion relative to the flat condition if the PPM incentives work in practice.

Participants. We recruit 210 subjects for an online experiment, implemented via Qualtrics. The subjects are recruited from Prolific, an online platform for conducting surveys. We restrict our subject pool to U.S. citizens who are students at the time of the experiment. Table B1 in Appendix B provides further information on the participants.

Procedure. The experiment was published on Prolific in May 2020. Subjects are randomly selected into one of the experimental conditions. They are first presented with instructions, which differ across the experimental conditions in rewards only. Then, subjects complete the prediction tasks. The order of the prediction tasks is randomized. Finally, subjects complete a short survey on demographics and their experience in the experiment.

3.2.2 Results

The primary question of interest is whether participants are more likely to seek costly information under the incentives provided by a PPM compared to fixed rewards. The effort task completion in control and PPM treatments allows us to test the effect of PPM incentives in eliciting effort. Furthermore in our prediction task, the ground truth (the actual box in any pair) is known to the experimenter. The accuracy treatment implements rewards for ex-post accuracy, which are not feasible in practice for elicitation without verification. We compare accuracy and PPM treatments to assess the effectiveness of PPM incentives relative to ex-post rewards.

We measure the frequency with which subjects completed the effort tasks across the experimental conditions. Figure 3 depicts the percentage of participants in each experimental condition who complete the associated effort task in each prediction task.

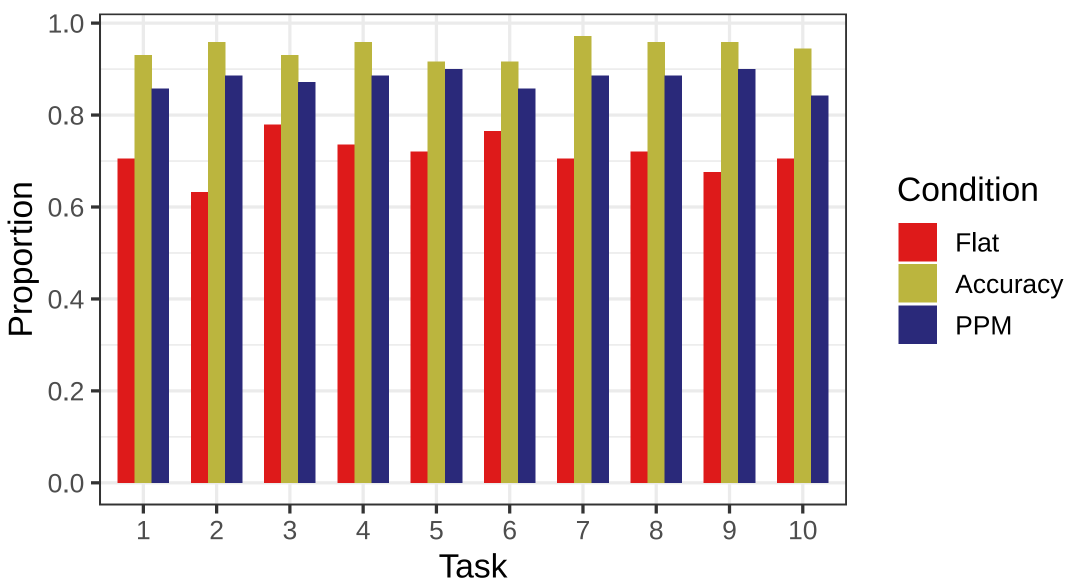

Figure 3: Proportion of participants who complete effort tasks in each prediction task.

The effort level is higher than zero, even in the control condition. Effort task completion is strictly higher in the PPM and accuracy treatments while the latter achieves the highest proportions. Figure 3 suggests that incentives provided by a PPM is effective in eliciting a higher proportion of informed judgments compared to a fixed reward. Incentives in the accuracy treatment are the most effective in eliciting effort.

We now investigate if subjects followed the truthful strategy, which also entails picking the left (right) box when a yellow (blue) ball is drawn. Given the simplicity of the predictions task, subjects do not have any external motives to make a non-truthful pick. However, deviations from the truthful strategy may occur due to confusion or errors. Figure 4 shows subjects' picks given their draw. The 3x3 grid depicts the three experimental conditions as well as the three possible situation after the effort task. A subject will receive a yellow or blue draw if she completes the effort task. Alternatively, the subject does not receive a draw if she skips the effort task. The bars show the number of picks in each task. Since picking the left (right) box when the draw is yellow (blue) is the truthful strategy, the number of left (right) picks are represented by yellow (blue) colored bars. The black dots show subjects' prior expectation on the number of yellow balls in the actual box, given that left and right boxes are equally likely to be the actual box. Table A1 in Appendix A provides the prior expectations on the number of yellow balls in each task.

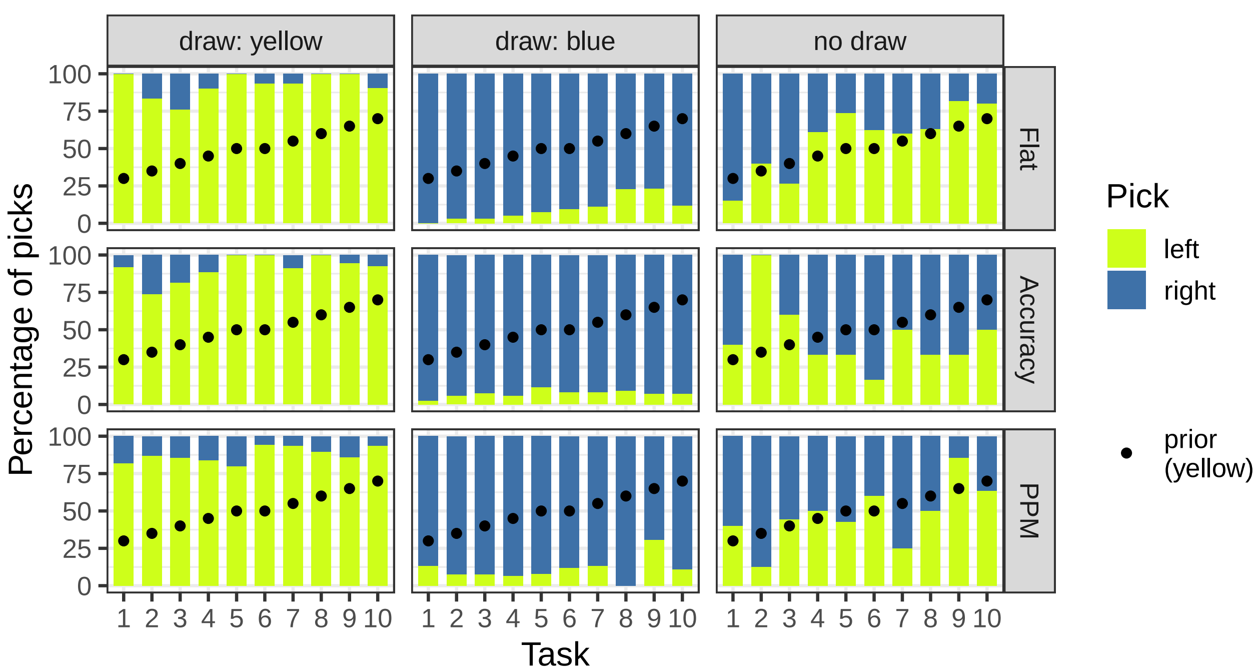

Figure 4: Subjects' picks

Figure 4 strongly suggests that the subjects pick according to the truthful strategy. Subjects who observe a yellow (blue) draw typically pick the left (right) box. The distribution of picks in PPM and Accuracy are very similar, so we can argue that the PPM incentives elicit subjects' true prediction. The same is true for the Flat condition as well. However, as shown in Figure 3, subjects in Flat are more likely to pick without completing the effort task. Thus, PPM is more effective in eliciting a complete truthful strategy. Also note that we do not observe the degenerate outcomes where all subjects coordinate on picking the same box. In contrast, subjects picks match with their signal as predicted by the truthful equilibrium.

The right-hand panel of Figure 4 illustrates the strategy subjects used if they did not draw. Interestingly, subjects in the PPM treatment (and in the Flat treatment) appeared to follow a mixed strategy, reporting left with a probability equal to the prior, as described in the equilibrium of Proposition 3. The probability to report left and the prior were correlated (Pearson: $\rho = 0,64$, $p = 0.048$) and not significantly different (t-test $t = -0.34$ $p = 0.739$) for PPM subjects who did not draw a ball, whereas they were uncorrelated and significantly different for those who drew a yellow ball or a blue ball (see Table C1 in the appendix).

For a statistical analysis on effort task completion, we estimate logistic regressions where probability of effort task completion is the dependent variable. Table 1 below shows the average marginal effects. The pooled data includes 2100 decisions to complete the effort task or not. We include binary indicators for the experimental conditions as dependent variables. The coefficient of 'PPM' in Table 1 measures the estimated difference from implementing PPM incentives instead of a flat fee on the likelihood of effort task completion in any task. The coefficient of 'Accuracy' measures the same for rewarding participants for ex-post accuracy. Models (1) and (2) use the whole sample of subjects. In (3) and (4), participants who gave an incorrect answer in the post-experimental quiz are excluded to construct a filtered sample. Specifications (2) and (4) also include various controls. The variables 'US citizen?' and 'Female?' are binary indicators for US residents and gender respectively while 'Age' is a numeric variable. In all models, standard errors are clustered at participant level. Tables C3

Dep. var.: P(effort task completed)				
	(whole sample)		*(filtered sample)*	
	(1)	(2)	(3)	(4)
PPM	0.10^{**}	0.09^{**}	0.10^{**}	0.08^{**}
	(0.03)	(0.03)	(0.03)	(0.03)
Accuracy	0.18^{***}	0.18^{***}	0.18^{***}	0.18^{***}
	(0.03)	(0.03)	(0.03)	(0.03)
Age		-0.00		-0.00
		(0.00)		(0.00)
Female?		0.04		0.04
		(0.03)		(0.03)
US resident?		-0.02		-0.02
		(0.06)		(0.06)
Num. obs.	2100	2070	2060	2030
Log Likelihood	-821.85	-768.69	-816.44	-763.58
Deviance	1643.70	1537.38	1632.88	1527.16
AIC	1649.70	1549.38	1638.88	1539.16
BIC	1666.65	1583.19	1655.77	1572.86

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$; $^{+}p < 0.1$

Table 1: Marginal effects, logistic regression (baseline category: flat)

and C4 in Appendix C present probit marginal effects and regression estimates from both logistic and probit models.

In all specifications, the marginal effects for PPM and accuracy treatments are positively significant. Based on model (1), we see that a participant in the PPM treatment is 10% more likely to complete the associated effort task in a given prediction task. Incentives provided by a PPM motivates agents to exert more effort compared to a fixed payment. For a comparison between Accuracy and PPM, Table C2 estimates the same logistic regression except that PPM is the baseline category. Incentives for ex-post accuracy is 9-11% more likely to elicit effort compared to a PPM. We can infer that incentives for ex-post accuracy is the most effective in effort elicitation, followed by PPM and flat payments. In the absence of verifiability, PPM provides an alternative for incentivizing effort and eliciting truthful judgments.

3.3 Study 2 - Eliciting Covid-19 experiences truthfully using PPM

Study 2 implements PPM incentives in measuring if the residents of the UK followed safety guidance during the Covid-19 pandemic. For most of the safe practices in the guidance, it is not feasible to monitor all individual behavior. Self-reported behavior is practically unverifiable. In an unincentivized or a flat-fee survey, participants may not exert the mental effort to recall and report their behavior truthfully. Furthermore, reporting costs can be asymmetric. Unsafe behavior is typically stigmatized and likely to be under-reported. We investigate if the PPM motivate participants to spend more time in answering questions and report their unsafe practices at a higher rate.

3.3.1 Design and procedures

Tasks. Participants are presented a survey consisting of 8 statements. Each statement describes a situation that was considered unsafe and inadvisable (if not prohibited) by the UK Covid-19 guidance at the time of this survey. For each statement, participants pick 'true' or 'false' to self-report if they have been in the described situation. Table 2 provides the list of questions:

	Statement
1.	I have been in an elevator with another person in it at least once in the last 7 days
2.	I may have stood less than 2 metres away from the person in front in a queue at least once in the last 7 days
3.	I was seated less than 2 metres away from someone who is not part of my household in a restaurant/cafe/bar at least once in the last 7 days
4.	I have been in a social gathering with more than 6 people who are not part of my household at least once in the last 7 days
5.	I have been in a busy shop/market with no restrictions on number of customers at least once in the last 7 days
6.	I participated in an indoor activity with more than 6 people who are not part of my household at least once in the last 7 days
7.	I have been in a shop/market where one or more of the staff did not wear a mask at least once in the last 7 days
8.	I had an interaction with someone experiencing high body temperature, persistent cough or loss of taste/smell at least once in the last 7 days

Table 2: Covid-19 survey questions

We ran this survey for two weeks with a new sample of participants every week. The two iterations of the survey are referred to as week 1 and week 2 surveys respectively. As we will introduce below, week 1 and week 2 surveys include experimental conditions that implement the PPM. We also run a week 0 survey to elicit information necessary to initialize the PPM. The week 0 survey uses the same questions, but they are presented in a slightly different way to elicit more information on the number of instances participants engaged in the described behavior. For example, question 1 in Table 2 is presented as 'In the last 7 days, I have been in an elevator with another person in it ...' and the participant picks one of the following answers: 'once or more', 'twice or more', '3 times or more', '4 times or more', '5 times or more'. Based on the results of the week 0 survey, we decided to implement two versions of each survey in weeks 1 and 2. Both versions ask the questions in Table 2, but in the second version ‘at least once’ is replaced with ‘at least twice’ in each question. We will provide more information on how week 0 survey is used in the design below.

Design. In week 0 survey, all participants receive a flat fee. In week 1 and 2 surveys, we manipulate incentives to create the control and treatment conditions. In the control, participants are rewarded with a flat fee for completing the survey while the treatment implements the PPM incentives. Figure 5 shows the experiment interface in the *PPM* condition.

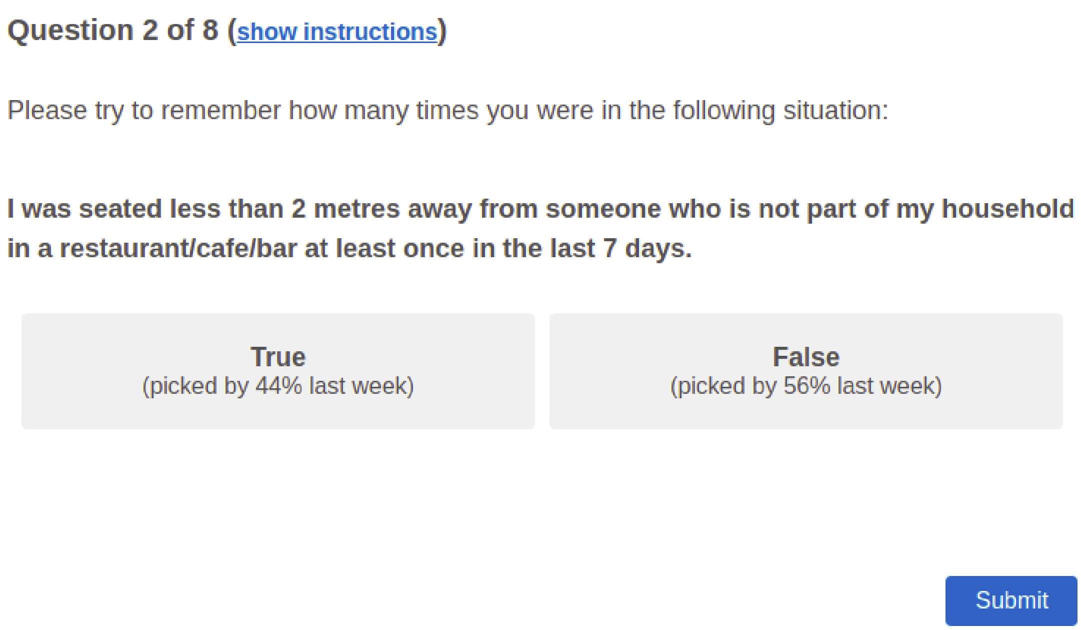

Figure 5: A screenshot from the treatment condition

The interface displays the statement and requires subjects to pick ‘true’ or ‘false’. The text below each alternative shows the percentage of participants who endorsed that alternative in the previous week’s survey. Recall that in our Bayesian setup, agents have a common prior expectation ω^0, which can be considered as the last realization of ω. The market maker sets $p = \omega^0$, which leads to the separating equilibrium. The endorsement rates of the previous iteration represents ω^0. Furthermore, participants’ bonus depends on the endorsement rates. In Figure 5, the endorsement rate of ‘true’ in the last iteration is 44%. A participant who picks ‘true’ in this iteration wins a positive (negative) bonus from this question if the realized endorsement rate in this iteration exceeds (falls below) 44%. The same holds for 'false', except that the threshold is 56%. Thus, the PPM condition essentially implements a repeated PPM where last iteration's realization determines the price for the current iteration. We will provide more information on the rewards below. The PPM incentives are expected to incentivize mental effort and/or overcome the psychological costs of reporting one's actual behavior. If PPM works as intended, we may expect decision times to be longer and endorsement rates for 'true' to be higher.

The control surveys are similar to the treatment surveys except that participants are rewarded with a flat fee. We implement two different types of control surveys. In the *control-1* condition, the survey interface does not present any information on previous iterations' endorsement rates. In contrast, the *control-2* survey shows the same screen as the PPM condition, as shown in Figure 5. So, the control-2 survey displays last week's endorsement rates. The rewards are fixed in both control-1 and control-2 surveys, thus the previous endorsement rates are irrelevant. Nevertheless, we included control-2 condition to check if merely presenting that information affects participants reports. If a PPM is effective, we could expect to see higher endorsement for 'true' and longer response times in the PPM condition compared to control-1. However, participants process additional information (previous endorsement rates) in the treatment condition, which might affect decision times. A significant difference between the PPM and the control-2 conditions would further suggest that the effect on reports and decision times is not simply due to the availability of previous endorsement rates.

Control-2 and PPM surveys present information on endorsement rates in the previous iteration. Week 0 survey is used to determine the previous endorsement rates presented in the control-2 and PPM surveys of week 1. Thus, week 0 data is used to initialize control-2 and PPM. Furthermore, the week 0 survey motivates our choice to run two versions where the statements include 'at least once' and 'at least twice' respectively. Table B2 in Appendix C provides the percentage of participants who pick 'true' in each question in the week 0 survey. For '3 times or more' and higher thresholds, the percentage of 'true' picks are close to 0. Then, participants in week 1 iteration of an 'at least 3 times' version may report 'true' simply because the threshold is very low and a few 'true' picks could easily bring the week 1 endorsement rates above the threshold. To avoid such cases, we only run two versions with 'at least once' and 'at least twice' respectively.

To summarize, we implement 6 surveys in a 3 (control-1,control-2,PPM) $\times$ 2 ('at least once', 'at least twice') design in each iteration. The week 0 survey is used to initialize the control-2 and PPM surveys in week 1 while week 2 surveys are initialized using week 1 results endorsement rates from the same survey.

Participants. Participants are recruited from Prolific, an online platform that provides subject pools for online experiments. We restrict our subject pool to students who currently reside in the UK. In total 692 participants completed our survey, 50 of which participate in week 0 survey while the remaining 642 participated in a week 1 or week 2 survey, assigned randomly in one of the 6 conditions explained above. One participant is excluded for being in a non-student status at the time of data collection. All surveys are implemented via Qualtrics. Table B3 in Appendix C provides further information on the participants.

Rewards. Control-1 and control-2 surveys pay a fixed reward of £1.75. In the PPM surveys, participants earn £0.75 for participation. In addition, they start with a bonus of £1. In each question, a participant's bonus changes according to the difference between the endorsement rate in the current survey versus the endorsement rate in the previous iteration. To illustrate, suppose a participant picked 'true' in a question in week 2 survey and endorsement rate of 'true' was 50% in week 1. If the realized endorsement rate of 'true' in week 2 at the same question is 70%, the subject wins $70 - 50 = 20$ pence. In contrast, if the endorsement rate in week is 30%, the subject loses $50 - 30 = 20$ pence. The previous week's endorsement rate serves as the price in a PPM while the current week's endorsement rate, unknown to the participant at the time of her decision, is analogous to realized value of the asset. For each participant in the PPM condition, we sum the gains and losses over all question to determine the net bonus.

Procedure. The experiment is conducted over three weeks and consists of week 0, 1 and 2 surveys that take place 7 days apart. The week 0 iteration is a single survey while in weeks 1 and 2, participants are randomly assigned to the different conditions. In each survey of each iteration, participants are first presented with instructions. Then they are asked to respond to the questions, which are presented in randomized order. Finally, participants complete a short survey on demographics and their experience in the experiment.

3.3.2 Results

Figure 6 shows the percentage of 'true' picks for each condition and version in the week 1 and week 2 surveys. Responses are pooled across questions and participants. Furthermore, we exclude 12 observations where the response time is longer than 60 seconds. Figure C1 in Appendix C suggest that these observations can be treated as outliers. Thus, they are excluded in all analyses below.

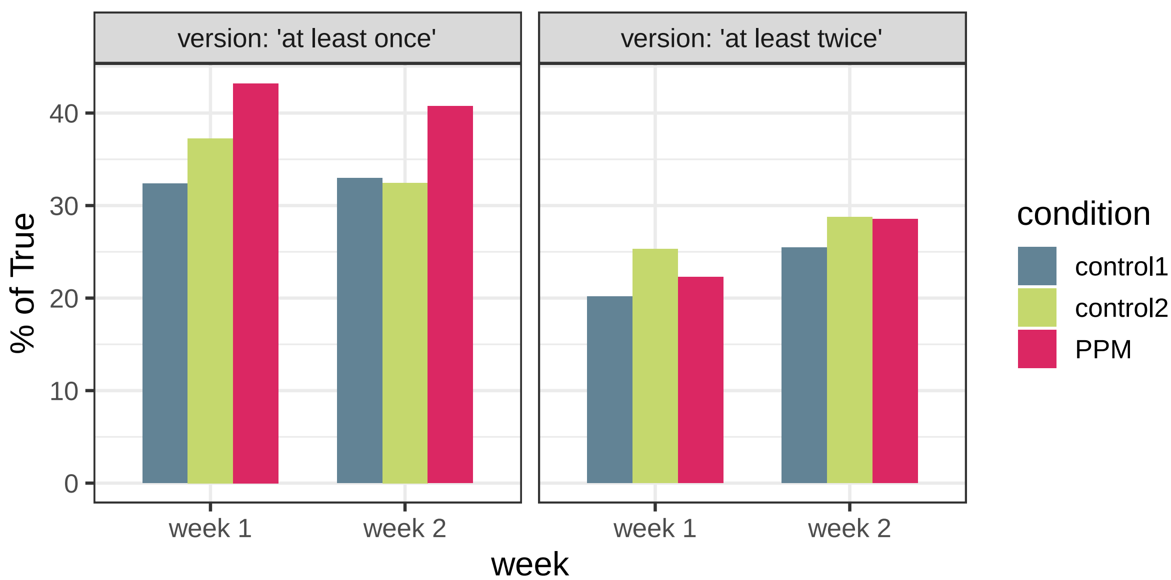

Figure 6: Percentage of 'true' picks in week 1 and 2 surveys.

In the 'at least once' surveys, the treatment elicits a higher percentage of 'true' responses compared to both controls. No such difference is observed in any iteration in the 'at least twice' version. Figure C2 in Appendix C shows a breakdown of percentage of 'true' across different questions. PPM elicits more 'true' in most questions in the 'at least once' version. Recall that week 1 surveys are initialized with the unincentivized week 0 survey (of a slightly different format) while week 2 surveys use data from week 1 survey of the corresponding condition. Since the prior has an effect on PPM, we will analyze the response data from weeks 1 and 2 separately.

Figure 7 depicts the response times for each version and week. We also categorize data according to the response type to see if response times differ.

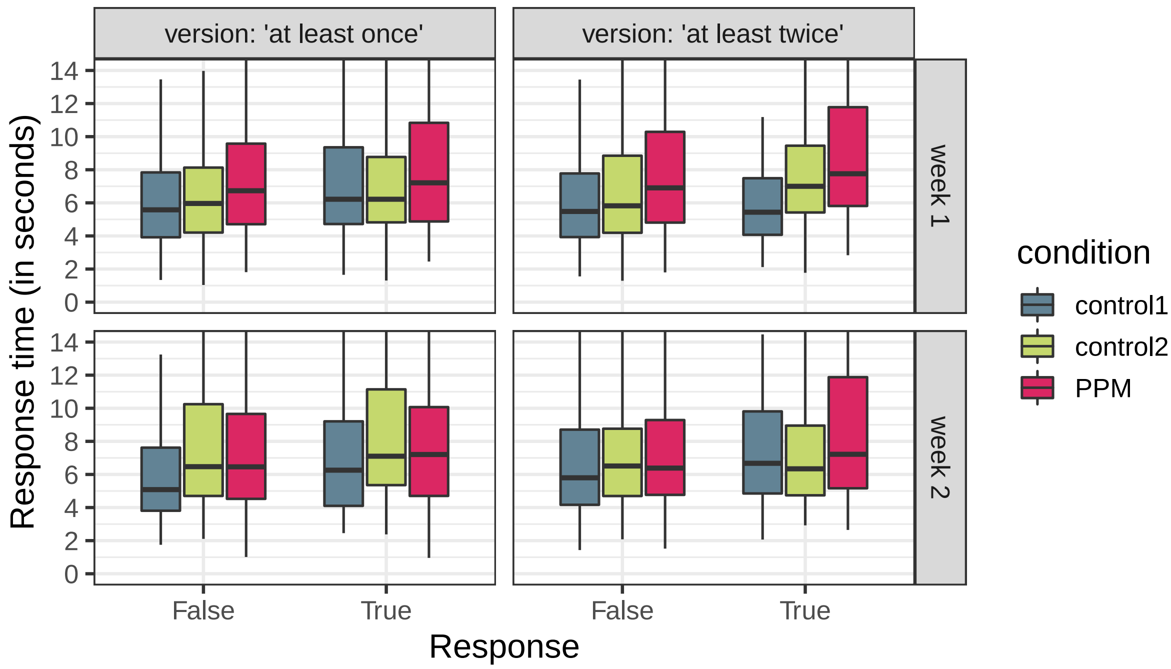

Figure 7: Response time of participants. The data points above 14 are included in calculations but not shown on the figure.

The median response time in the PPM condition is higher than the control-1 surveys in all iterations. The same is true for control-2 surveys in week 1. However, response times in control-2 and PPM are comparable in week 2 surveys. Interestingly, Figure 7 suggests that the response times are higher in the week 2 iteration of the 'at least once' control-2 survey than the week 1 iteration, which partially explains why response times are comparable to PPM. Figure C3 supports this observation. Since week 1 and 2 surveys are identical other than the percentage of 'true' in previous week's iteration, we pool response time data in the analysis below.

For a statistical analysis, we estimate two classes of regression models. Firstly, we estimate a logistic regression for participants' likelihood picking 'true' in any given question. Secondly, we estimate a linear regression model where response time is the dependent variable. In both models, control-1 is the baseline category and binary indicators for control-2 and PPM are variables of interest. We also include various demographic controls representing the age, gender and citizenship of participants. We focus on the 'at least once' versions of all iterations as Figure 6 suggested a possible difference for these versions only. Section C.2.3 in Appendix C performs the same analysis for 'at least twice' survey. As mentioned above, we pool response time data from weeks 1 and 2, but we estimate separate models week 1 and week 2 response data.

Table 3 presents the average marginal effects from the logistic regressions and the estimates from the response time regressions. Models (1) to (4) includes average marginal effects while (5) and (6) show the response time regressions with week 1 and 2 data pooled. Table C5 in Appendix C estimates models (5) and (6) separately for week 1 and 2 data. The intercept term in (5) and (6) represents the estimated response time in the control-1 condition. In all models, standard errors are clustered at the participant level.

	P(response = 'true'), marginal effects				*Response time*	
	(week 1)		*(week 2)*		*(pooled)*	
	(1)	(2)	(3)	(4)	(5)	(6)
(Intercept)					6.85^{***}	7.52^{***}
					(0.25)	(0.69)
Control-2	0.05	0.04	-0.01	-0.00	1.13^{**}	1.06^{**}
	(0.04)	(0.04)	(0.04)	(0.04)	(0.39)	(0.39)
PPM	0.11^{***}	0.10^{**}	0.08^{*}	0.08^{*}	1.74^{***}	1.69^{***}
	(0.03)	(0.03)	(0.04)	(0.04)	(0.44)	(0.44)
Age		-0.00		-0.00		0.00
		(0.00)		(0.00)		(0.02)
Female?		0.02		-0.02		0.28
		(0.03)		(0.03)		(0.36)
UK citizen?		-0.00		0.03		-1.05^{*}
		(0.03)		(0.04)		(0.41)
Num. obs.	1259	1259	1279	1279	2538	2538
Log Likelihood	-828.13	-826.36	-827.33	-825.89		
Deviance	1656.27	1652.72	1654.66	1651.78		
AIC	1662.27	1664.72	1660.66	1663.78		
BIC	1677.68	1695.55	1676.13	1694.70		
R^2					0.01	0.02
Adj. R^2					0.01	0.02
RMSE					5.87	5.85
N Clusters					318	318

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$; $^{+}p < 0.1$

Table 3: Logistic regression and linear regression on response times

The average marginal effects in Table 3 show that the PPM survey elicits a higher frequency of 'true' picks. According to model (1), a participant in the PPM condition of week 1 survey is 11 percentage points more likely to report 'true' for a given statement compared to a participant in the control-1 condition. In contrast, control-2 condition has no effect. A similar result holds for the week 2 survey where the marginal effect of the PPM condition is estimated to be 8%. Tables C6 and C7 in Appendix C show similar results in probit marginal effects and the logistic and probit regression estimates. Results support the equilibrium characterized in Proposition 4. PPMs motivate participants to declare unsafe practices at a higher rate, which suggest that such practices are under-reported in basic surveys. Higher rates of admitting an unsafe practice indicate that the PPM incentives dominate potential reporting costs associated with the stigmatized response. PPMs also encourage participants to exert more mental effort. The results of the response time regressions partially support this interpretation. In models (5) and (6), participants in the PPM survey spend significantly more time in their responses than the control-1 survey. However, the same effect is observed for the control-2 survey. The test two parameters (PPM vs control-2) in (6) results in an insignificant difference (mean difference = 0.62, $t = 1.359$, $p = 0.17$). Thus, higher decision times can also be the result of subjects processing more information in the form of last week's percentages.

4 Discussion

4.1 Related Literature

The original peer prediction method of Miller et al. (2005) asks agents to report an answer to a multiple choice question. It is assumed that the center knows the common prior (possibly from previous data). An agent's report is used to update the prior. The resulting posterior is used to predict what another agent reported. Accuracy of the posterior determines initial agent's reward through the use of a proper scoring rule. Our approach has the advantage to require slightly less information from the center.

Subsequent work extended peer prediction method to settings with weaker information requirements, at the cost of 'non-minimality' (Witkowski and Parkes, 2012b) or introducing a dynamic setup (Witkowski and Parkes, 2013; Zhang and Chen, 2014). In the Bayesian truth serum (Prelec, 2004, BTS), agents are assumed to have a common prior belief. But, the mechanism designer need not have access to that prior. So, BTS can be implemented without using previous data. In BTS, agents make two reports. In one, they respond to a multiple choice question. In the other, they predict the frequency of each possible response. Agents' prediction reports are scored based on accuracy. Private responses are scored according to actual vs predicted endorsement frequencies of responses, such that surprisingly common answers are scored higher. A Bayesian agent, who shares a common prior belief with others on population distribution of responses, expects her own response to be more common than average prediction of all agents. Thus, scoring incentivizes agents to report their true answer.

Both the original peer prediction method and BTS are solutions to the problem of truthful elicitation without verification. They do not incorporate costly effort. The peer prediction method can be adapted to costly effort by re-scaling payoffs, using the knowledge on common prior belief. Similar to the peer prediction method, PPM is one-shot and minimal. However, PPM does not require common prior, nor the complete knowledge of prior beliefs. The market maker is assumed to know the common prior expectation only. In binary elicitation problems, PPM provides a simpler and less information demanding alternative to the peer prediction method.

Recent work developed peer prediction mechanisms for effort elicitation in crowdsourcing problems with unverifiable tasks, such as peer grading, content classification etc. Witkowski et al. (2013) studied output agreement mechanisms, in which an agent receives positive payment if her report agrees with a peer agent. Simple output agreement mechanisms do not achieve truthful elicitation when an agent believes that she holds a minority opinion, which may also affect effort decision. Dasgupta and Ghosh (2013) use reports in multiple auxiliary questions to penalize agreement without effort in a binary question of interest. Given common prior expectation, PPM achieve the same binary elicitation while maintaining minimality (single task). Shnayder et al. (2016) generalize Dasgupta and Ghosh (2013) to obtain correlated agreement mechanism for non-binary questions. Correlated agreement uses multiple questions and requires knowledge of signs of individual correlations across questions. Peer truth serum for crowdsourcing is another peer agreement mechanism which uses agents' responses to multiple questions Radanovic et al. (2016). Liu and Chen (2017b) develop sequential peer prediction, in which agents submit answers sequentially and the mechanism learns the optimal reward for effort elicitation over time. Sequential peer prediction is minimal, but unlike PPM, requires a dynamic setup. In binary elicitation problems, PPM offers a simpler minimal alternative to other peer prediction mechanisms for effort elicitation.

Bayesian markets (Baillon, 2017) offer a market-based solution for truthful elicitation in binary questions. In a Bayesian market, agents report an answer to a binary of question of interest. There is a single asset, whose value is determined by the proportion of agents who report 'yes'. Agents receive a costless binary signal, which fully determine their type. Agents share a common prior belief on population distribution of types. As in our setup, agents update their beliefs using their own types. Belief updating is 'impersonal', agents with the same type have the same posterior beliefs. A Bayesian type-1 ('yes') agent expects a higher value of asset compared to a Bayesian type-0 ('no') agent. Agents who report yes (no) are allowed to only buy (sell) the asset, at a price drawn randomly from unit interval later. The market maker executes trades only when majority of agents in both sides of the market (yes and no) are willing to trade, which occurs when price is within posterior expectations of the two types. In this setup, both types are incentivized to report their true beliefs. Since type-1 agents have a higher posterior expectation, they prefer to become buyers when trade occurs. Vice versa for type-0 agents.

In binary truthful elicitation problems, Bayesian markets have an appeal over scoring-based methods: prediction reports and scoring are replaced by simple betting decisions and market payoffs. PPM follow a similar approach, but the elicitation procedure is simplified even further. Unlike Bayesian markets, participants in a PPM do not report an answer. They trade freely according to their private information. In equilibrium, participant's true judgments can be inferred from their trade. In a Bayesian market, trade is an auxiliary tool to incentivize truthful reports. If the randomly drawn price is not in the appropriate range, trade may not occur even in the truthful equilibrium. In a PPM, trade occurs at any price. PPM is more analogous to a prediction market as participants trade at a given price.

4.2 Theoretical limitations

PPM, like similar mechanisms, assume risk neutrality. Risk aversion could decrease the perceived incentives provided by the mechanism. When participation is compulsory however, the no effort strategy is also risky. In the presence of high risk aversion, a degenerate equilibrium with no-one providing effort and everyone reporting the same answer would dominate equilibria with efforts.

As illustrated by Propositions 1 to 3, there are several types of equilibria. To those should be added equilibria in which signal 1 agents report 0 and conversely. These latter equilibria did not occur in Study 1. Interestingly, at the aggregate level, subjects seemed to play the strategies of Proposition 3, and those who did not draw a signal played a mixed strategy (at the aggregate level) where the randomization probability was equal to the prior.

We considered a very simple model, binary in all dimensions. Effort could be continuous, signal informativeness could be a function of effort, and answers could be non-binary. We leave these refinements for future research.

4.3 Empirical limitations

Study 1 made use of tasks borrowed from the experimental literature, which allowed us to observe effort. The main drawback is that those tasks were artificial, and may have been seen as quite unnatural. To test whether PPM also elicits more effort and honest answers in a more realistic context, Study 2 was conducted. Results of Study 2 demonstrate the real-world validity of PPM.

The study was conducted online with participants from the Prolific platform. Participants from online platforms take part in experiments in an uncontrolled setting, for example, from home. This lack of experimental control has elicited concerns amongst researchers. However, experimental research has shown that this concerns is largely unfounded. Hauser and Schwarz (2016) demonstrated that participants from an online platform are more attentive than college students. Eyal et al. (2021) demonstrated that Prolific outperformed other participant platforms regarding data quality, supporting our decision to collect data through Prolific. To ensure high data quality in the current research, post-experimental quiz questions were included in Study 1, allowing to remove inattentive participants.

We initially planned to run Study 2 over four weeks, but we had to stop earlier when the pandemics amplified in the UK (second wave), making our questions less applicable. Fortunately, data collected during weeks 0, 1, and 2 already provide valuable insights on the effectiveness of PPM in a real-world context with unverifiable truths.

The questions we used were selected to ensure that the negative feeling of admitting such behavior would be limited. For instance, in most statements, non-compliance could have been due to behavior of others.

The final two empirical limitations concern the interpretation of the results. The flat fee condition suggests substantial intrinsic motivation. PPM clearly elicited further effort (Study 1) and honest responses (Study 2). Less clear is whether PPM elicited honest responses in Study 1 and whether PPM elicited increased effort in Study 2. PPM led to increased performance compared to the flat fee control on the prediction task in Study 1. In order to perform well on this task, participants should combine all available information - including information gained from completing the effort task - to form an educated guess. Including incorrect information, either through misunderstanding or dishonesty, would negatively influence their educated guess and subsequently, their performance on the task. Arguably, the need for honesty is not practically relevant in the prediction task, since there were no costs associated with honesty. In Study 1, honesty was the optimal choice, there was no trade-off to be made. To explicitly test whether PPM can elicit honest answers, Study 2 was designed. In Study 2, participants were asked about their violations of COVID guidelines. The discrepancy between the prevalence of self-reported lies (Debey et al., 2015) and lies told during experimental research (Feldman et al., 2002) demonstrates that people are reluctant to admit anti-social behavior. Since violations of COVID guidelines could negatively affect the health of both oneself and others, a violation of COVID guidelines can be seen as immoral behavior. Results of Study 2 demonstrate that participants in the PPM condition admitted more violations of COVID guidelines than participants in both control conditions. These results demonstrate that PPM can elicit more honest responses. PPM may have helped overcome the 'shame' of reporting non-compliance with health guidelines (a_i in the theory).

Unlike in Study 1, in Study 2 it is unclear whether PPM increased effort. Effort was operationalized as increased response time. While participants in the PPM condition took longer to respond than participants in the Control 1 condition, they did not take longer than participants in the Control 2 condition. The PPM and Control 2 condition contained additional information about the percentage of participants endorsing the COVID relation behaviour in the previous week's survey. The additional time it takes to read and process this information may better explain the response time differences between Control 1 on the one hand and Control 2 and PPM on the other hand. Since response time is a proxy rather than a direct measure of effort, this finding does not mean that PPM does not elicit additional effort, but the effect of PPM on effort in Study 2 remains inconclusive.

5 Conclusion

For events with ex-post verifiable outcomes, prediction markets are known to be effective in eliciting and aggregating informed judgments. However, prediction markets are not suitable for unverifiable judgments, as the outcome-based rewards are not feasible. Researchers and practitioners typically resort to simple surveys with fixed rewards, which do not provide incentives to acquire costly information. PPM provide a market mechanism that incentivize agents to seek information and trade truthfully on binary questions of unverifiable information. Experimental evidence suggests that incentives provided by a PPM motivates agents to seek costly information in judgment formation.

References

Baillon, A. (2017). Bayesian markets to elicit private information. *Proceedings of the National Academy of Sciences*.

Crémer, J. and McLean, R. P. (1988). Full extraction of the surplus in bayesian and dominant strategy auctions. *Econometrica: Journal of the Econometric Society*, pages 1247–1257.

Dasgupta, A. and Ghosh, A. (2013). Crowdsourced judgement elicitation with endogenous proficiency. In *Proceedings of the 22nd international conference on World Wide Web*, pages 319–330.

Debey, E., De Schryver, M., Logan, G. D., Suchotzki, K., and Verschuere, B. (2015). From junior to senior pinocchio: A cross-sectional lifespan investigation of deception. *Acta psychologica*, 160:58–68.

Eyal, P., David, R., Andrew, G., Zak, E., and Ekaterina, D. (2021). Data quality of platforms and panels for online behavioral research. *Behavior Research Methods*, pages 1–20.

Feldman, R. S., Forrest, J. A., and Happ, B. R. (2002). Self-presentation and verbal deception: Do self-presenters lie more? *Basic and applied social psychology*, 24(2):163–170.

Hauser, D. J. and Schwarz, N. (2016). Attentive turkers: Mturk participants perform better on online attention checks than do subject pool participants. *Behavior research methods*, 48(1):400–407.

Liu, Y. and Chen, Y. (2017a). Machine-learning aided peer prediction. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, EC '17, pages 63–80, New York, NY, USA. ACM.

Liu, Y. and Chen, Y. (2017b). Sequential peer prediction: Learning to elicit effort using posted prices. In *Thirty-First AAAI Conference on Artificial Intelligence*.

Milgrom, P. and Stokey, N. (1982). Information, trade and common knowledge. *Journal of economic theory*, 26(1):17–27.

Miller, N., Resnick, P., and Zeckhauser, R. (2005). Eliciting informative feedback: The peer-prediction method. *Management Science*, 51(9):1359–1373.

Osborne, M. J. and Rubinstein, A. (1994). *A course in game theory*. MIT press.

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.

Radanovic, G., Faltings, B., and Jurca, R. (2016). Incentives for effort in crowdsourcing using the peer truth serum. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 7(4):48.

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*, pages 179–196.

Waggoner, B. and Chen, Y. (2013). Information elicitation sans verification. In *Proceedings of the 3rd Workshop on Social Computing and User Generated Content, in conjunction with ACM EC'13*, volume 16.

Witkowski, J., Bachrach, Y., Key, P., and Parkes, D. C. (2013). Dwelling on the negative: Incentivizing effort in peer prediction. In *First AAAI Conference on Human Computation and Crowdsourcing*.

Witkowski, J. and Parkes, D. (2012a). A robust bayesian truth serum for small populations.

Witkowski, J. and Parkes, D. C. (2012b). Peer prediction without a common prior. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 964–981. ACM.

Witkowski, J. and Parkes, D. C. (2013). Learning the prior in minimal peer prediction. In *Proceedings of the 3rd Workshop on Social Computing and User Generated Content at the ACM Conference on Electronic Commerce*, volume 14.

Zhang, P. and Chen, Y. (2014). Elicitability and knowledge-free elicitation with peer prediction. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, pages 245–252. International Foundation for Autonomous Agents and Multiagent Systems.

A Additional experimental materials

Pair	Left box	Right box	Prior expectation on yellow in the actual box
1.	40 yellow, 60 blue	20 yellow, 80 blue	30
2.	40 yellow, 60 blue	30 yellow, 70 blue	35
3.	48 yellow, 52 blue	32 yellow, 68 blue	40
4.	56 yellow, 44 blue	34 yellow, 66 blue	45
5.	62 yellow, 38 blue	38 yellow, 62 blue	50
6.	57 yellow, 43 blue	43 yellow, 57 blue	50
7.	69 yellow, 31 blue	41 yellow, 59 blue	55
8.	69 yellow, 31 blue	51 yellow, 49 blue	60
9.	78 yellow, 22 blue	52 yellow, 48 blue	65
10.	77 yellow, 23 blue	63 yellow, 37 blue	70

Table A1: The content of boxes and prior expectation on yellow in each pair

B Summary statistics

Table B1: Summary statistics, Study 1

	Experimental Condition		
	Flat	Accuracy	PPM
Number of subjects	68	72	70
Female/Male	29/39	36/36	34/36
Average age	23.09	23.76	22.64
US resident	63	65	62
Average duration	8 min 59 sec	9 min 31 sec	9 min 8 sec
Average reward	£3.25	£3.50	£3.342
Correct answer in pre-experimental quiz	54	67	57
Correct answer in post-experimental quiz	68	72	66

Table B2: Study 2, Week 0 answers

	Percentage of 'true' picks				
Question	once or more	twice or more	3 times or more	4 times or more	5 times or more
1	18	12	6	4	4
2	76	50	20	6	2
3	58	22	8	4	2
4	16	8	0	0	0
5	70	34	14	4	2
6	24	10	8	4	2
7	54	24	8	2	2
8	12	4	2	2	2

Table B3: Summary statistics, Study 2

	Exp. Condition / version					
Week 1						
	Control-1 / 'once'	Control-2 / 'once'	Treatment / 'once'	Control-1 / 'twice'	Control-2 / 'twice'	Treatment / 'twice'
Number of subjects	53	53	52	54	54	53
Female/Male	36/17	36/17	33/19	36/18	25/29	33/20
Average age	24.85	23.53	22.73	23.11	23.57	25.17
UK/Non-UK citizen	42/11	36/17	40/12	44/10	45/9	37/16
Average duration	2 min 10 sec	2 min 38 sec	3 min 34 sec	2 min 14 sec	2 min 30 sec	3 min 38 sec
Average reward	£1.75	£1.75	£2.03	£1.75	£1.75	£1.81
Week 2						
Number of subjects	54	52	54	54	54	54
Female/Male	31/23	31/21	39/15	37/17	39/15	38/16
Average age	24.39	25.65	24.98	25.13	24.25	25.09
UK/Non-UK citizen	46/8	44/8	43/11	43/11	46/8	48/6
Average duration	2 min 14 sec	2 min 52 sec	3 min 44 sec	2 min 45 sec	2 min 25 sec	4 min 12 sec
Average bonus	£1.75	£1.75	£1.66	£1.75	£1.75	£1.73

C Additional results

C.1 Study 1

(a) Correlation tests

Draw	Pearson's C.C.	Spearman's C.C.
yellow	r = 0.53, p = 0.118	ρ = 0.52, p = 0.121
blue	r = 0.28, p = 0.425	ρ = 0.21, p = 0.555
no draw	r = 0.64, p = 0.048	ρ = 0.68, p = 0.032

(b) Two-sided t-test and Wilcoxon test

Draw	T-test	Wilcoxon test
yellow	t = 8.56, p < 0.001	W = 100, p < 0.001
blue	t = -8.12, p < 0.001	W = 1, p < 0.001
no draw	t = -0.34, p = 0.739	W = 44, p = 0.676

Table C1: Proportion of left picks vs prior expectation on the number of yellow balls in the actual box.

Dep. var.: P(effort task completed)				
	(whole sample)		*(filtered sample)*	
	(1)	(2)	(3)	(4)
Flat	-0.13^{*}	-0.11^{*}	-0.13^{*}	-0.10^{*}
	(0.05)	(0.05)	(0.05)	(0.05)
Accuracy	0.09^{*}	0.11^{**}	0.10^{*}	0.11^{**}
	(0.04)	(0.04)	(0.04)	(0.04)
Age		-0.00		-0.00
		(0.00)		(0.00)
Female?		0.04		0.03
		(0.03)		(0.03)
US resident		-0.02		-0.02
		(0.06)		(0.06)
Num. obs.	2100	2070	2060	2030
Log Likelihood	-821.85	-768.69	-816.44	-763.58
Deviance	1643.70	1537.38	1632.88	1527.16
AIC	1649.70	1549.38	1638.88	1539.16
BIC	1666.65	1583.19	1655.77	1572.86

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$; $^{+}p < 0.1$

Table C2: Marginal effects, logit regression (baseline category: Accuracy)

Dep. var.: P(effort task completed)				
	(whole sample)		*(filtered sample)*	
	(1)	(2)	(3)	(4)
PPM	0.11^{**}	0.10^{**}	0.11^{**}	0.09^{**}
	(0.04)	(0.03)	(0.04)	(0.03)
Accuracy	0.19^{***}	0.19^{***}	0.19^{***}	0.19^{***}
	(0.03)	(0.03)	(0.04)	(0.03)
Age		-0.00		-0.00
		(0.00)		(0.00)
Female?		0.04		0.03
		(0.03)		(0.04)
US resident		-0.03		-0.03
		(0.06)		(0.06)
Num. obs.	2100	2070	2060	2030
Log Likelihood	-821.85	-768.78	-816.44	-763.66
Deviance	1643.70	1537.56	1632.88	1527.33
AIC	1649.70	1549.56	1638.88	1539.33
BIC	1666.65	1583.37	1655.77	1573.02

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$; $^{+}p < 0.1$

Table C3: Marginal effects, probit regression (baseline category: Flat)

Dep. var.: P(effort task completed)								
	(logit)		*(logit, filtered)*		*(probit)*		*(probit, filtered)*	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(Intercept)	0.92^{***}	1.91^{*}	0.92^{***}	1.91^{*}	0.57^{***}	1.17^{*}	0.57^{***}	1.18^{*}
	(0.22)	(0.86)	(0.22)	(0.87)	(0.13)	(0.48)	(0.13)	(0.49)
Accuracy	1.91^{***}	2.15^{***}	1.91^{***}	2.15^{***}	1.03^{***}	1.13^{***}	1.03^{***}	1.13^{***}
	(0.43)	(0.41)	(0.43)	(0.41)	(0.22)	(0.20)	(0.22)	(0.20)
PPM	1.05^{***}	0.96^{*}	0.98^{**}	0.89^{*}	0.59^{**}	0.54^{**}	0.56^{**}	0.51^{*}
	(0.36)	(0.37)	(0.36)	(0.37)	(0.20)	(0.21)	(0.20)	(0.21)
Age		-0.04		-0.04		-0.02		-0.02
		(0.03)		(0.03)		(0.02)		(0.02)
Female?		0.37		0.33		0.19		0.17
		(0.33)		(0.33)		(0.18)		(0.18)
US resident?		-0.24		-0.19		-0.17		-0.14
		(0.65)		(0.65)		(0.33)		(0.34)

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$; $^{+}p < 0.1$

Table C4: Regression estimates (baseline: Flat)

821 ## C.2 Study 2

822 ### C.2.1 Figures on responses and response times

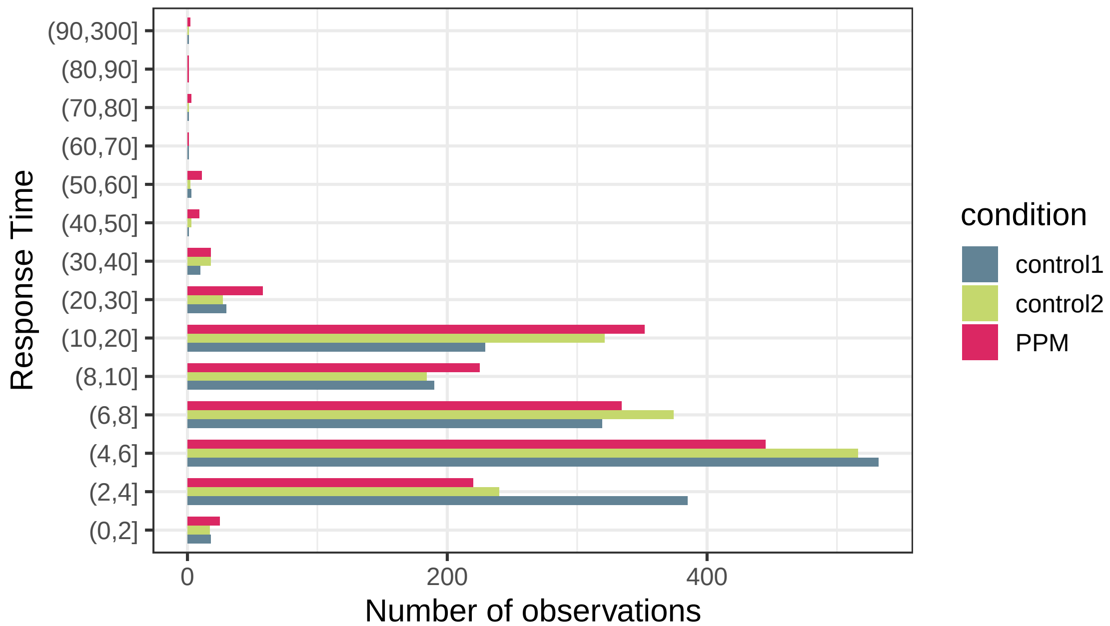

Figure C1: Response times

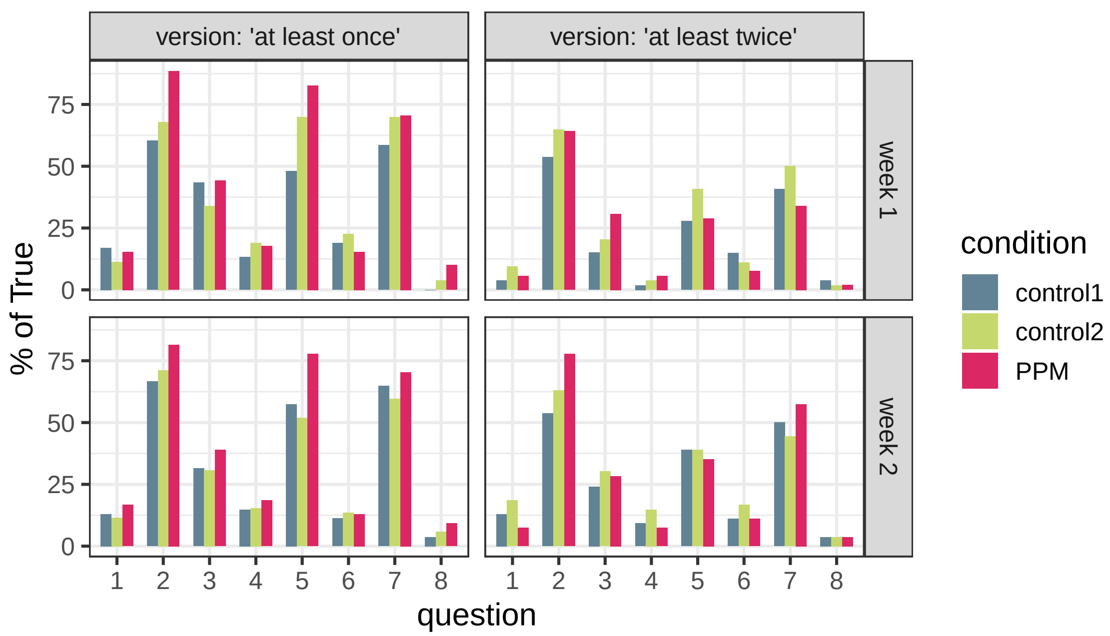

Figure C2: Proportion of participants who complete effort tasks in each prediction task.

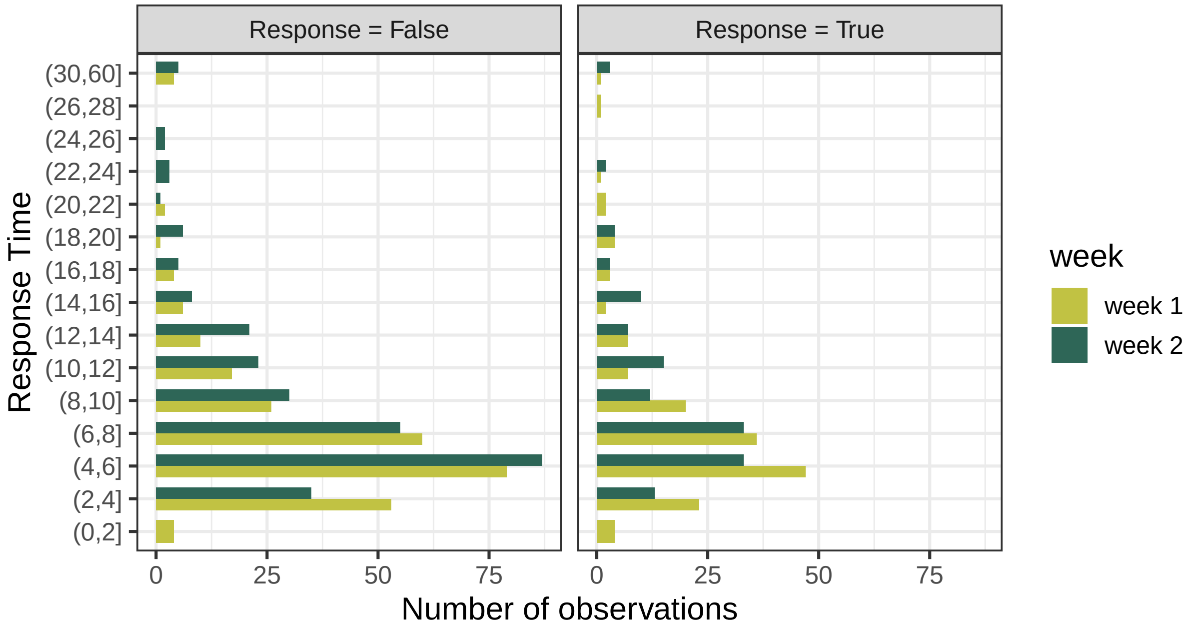

Figure C3: Response times in control-2 'at least once' survey, weeks 1 and 2

	(week 1)		*(week 2)*	
	(1)	(2)	(3)	(4)
(Intercept)	6.75^{***}	7.39^{***}	6.95^{***}	8.07^{***}
	(0.28)	(1.11)	(0.42)	(0.98)
Control-2	0.61	0.51	1.66^{**}	1.64^{**}
	(0.48)	(0.49)	(0.60)	(0.59)
PPM	2.37^{***}	2.35^{***}	1.14^{+}	0.99
	(0.62)	(0.61)	(0.62)	(0.63)
Age		-0.01		0.00
		(0.04)		(0.02)
Female?		0.28		0.41
		(0.50)		(0.51)
UK citizen?		-0.80		-1.65^{*}
		(0.51)		(0.64)
R^2	0.03	0.03	0.01	0.03
Adj. R^2	0.03	0.03	0.01	0.02
Num. obs.	1259	1259	1279	1279
RMSE	5.89	5.89	5.81	5.78
N Clusters	158	158	160	160

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$; $^{+}p < 0.1$

Table C5: Response time regressions, estimated separately for weeks 1 and 2.

C.2.2 Regression estimates, probit marginal effects

	(week 1)		*(week 2)*	
	(1)	(2)	(3)	(4)
Control-2	0.05	0.04	−0.01	−0.00
	(0.04)	(0.04)	(0.04)	(0.04)
PPM	0.11***	0.10**	0.08*	0.08*
	(0.03)	(0.03)	(0.04)	(0.04)
Age		−0.00		−0.00
		(0.00)		(0.00)
Female?		0.02		−0.02
		(0.03)		(0.03)
UK citizen?		−0.00		0.03
		(0.03)		(0.04)
Num. obs.	1259	1259	1279	1279
Log Likelihood	−828.13	−826.36	−827.33	−825.89
Deviance	1656.27	1652.71	1654.66	1651.78
AIC	1662.27	1664.71	1660.66	1663.78
BIC	1677.68	1695.54	1676.13	1694.71

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$; $^{+}p < 0.1$

Table C6: Probit marginal effects

	Logistic				*Probit*			
	(week 1)		*(week 2)*		*(week 1)*		*(week 2)*	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(Intercept)	−0.74***	−0.31	−0.71***	−0.56*	−0.46***	−0.20	−0.44***	−0.35*
	(0.10)	(0.33)	(0.11)	(0.28)	(0.06)	(0.20)	(0.06)	(0.17)
Control-2	0.22	0.19	−0.02	−0.01	0.13	0.12	−0.01	−0.01
	(0.16)	(0.16)	(0.16)	(0.16)	(0.10)	(0.10)	(0.09)	(0.09)
PPM	0.46***	0.43**	0.34*	0.36*	0.29***	0.26**	0.21*	0.22*
	(0.13)	(0.13)	(0.16)	(0.16)	(0.08)	(0.08)	(0.10)	(0.10)
Age		−0.02		−0.01		−0.01		−0.01
		(0.01)		(0.01)		(0.01)		(0.00)
Female		0.08		−0.09		0.05		−0.05
		(0.13)		(0.13)		(0.08)		(0.08)
UK citizen?		−0.01		0.14		−0.01		0.09
		(0.13)		(0.16)		(0.08)		(0.10)

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$; $^{\cdot}p < 0.1$

Table C7: Logistic and probit regression estimates (baseline: control-1)

C.2.3 Analysis on ‘at least twice’ survey data

	P(response = ‘true’), marginal effects				*Response time*	
	(week 1)		*(week 2)*		*(pooled)*	
	(1)	(2)	(3)	(4)	(5)	(6)
(Intercept)					6.76^{***}	9.11^{***}
					(0.36)	(1.00)
Control-2	0.05^{+}	0.05^{+}	0.03	0.04	0.94^{+}	1.18^{*}
	(0.03)	(0.03)	(0.04)	(0.04)	(0.53)	(0.50)
PPM	0.02	0.04	0.03	0.04	2.56^{***}	2.56^{***}
	(0.03)	(0.03)	(0.04)	(0.03)	(0.66)	(0.66)
Age		-0.00^{*}		-0.00^{+}		-0.07^{**}
		(0.00)		(0.00)		(0.03)
Female?		0.00		−0.02		0.84
		(0.02)		(0.03)		(0.55)
UK citizen?		0.07^{**}		−0.03		-1.65^{*}
		(0.03)		(0.04)		(0.72)
Num. obs.	1284	1276	1294	1286	1284	1276
Log Likelihood	−684.32	−674.50	−761.46	−754.97		
Deviance	1368.64	1349.01	1522.92	1509.94		
AIC	1374.64	1361.01	1528.92	1521.94		
BIC	1390.12	1391.91	1544.42	1552.90		
R^2					0.03	0.05
Adj. R^2					0.03	0.04
RMSE					6.06	6.02
N Clusters					161	160

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$; $^{+}p < 0.1$

Table C8: Logistic regression and linear regression on response times, ‘at least twice’ version