Online Appendices for "Privacy, Voting, and the Wisdom of Crowds"

Appendix C: Asymmetric Costs and Base Rates

In this section, we extend our core model to accommodate settings where the prior probability of the state of the world is not necessarily balanced, and the costs associated with different types of prediction errors are distinct. This generalization allows for a broader range of real-world scenarios where, for instance, false positives (e.g., launching a failed product) might be more or less costly than false negatives (e.g., shelving a potentially successful product), or where the baseline expectation for success deviates from 50%.

C.1. Extended Model Setup

We retain the fundamental structure described in Section 3 but introduce two key modifications. First, the prior probability of the state of the world being $Y=1$ (e.g., product success), denoted by $\theta = \Pr[Y=1]$, is allowed to be any value within the open interval $(0,1)$, rather than being fixed at $1/2$. Second, we differentiate the costs associated with the two possible types of errors the firm can make. We denote the cost of a false positive (Type I error: $a=1$ when $Y=0$) as $c_{\mathrm{FP}}>0$, and the cost of a false negative (Type II error: $a=0$ when $Y=1$) as $c_{\mathrm{FN}}>0$.

Consequently, the firm's cost function, representing its expected loss given its action a and the vector of public votes $\boldsymbol{x}_p$, is now expressed as:

$$C(a;\nu,\boldsymbol{x}_p) = \mathbb{E}_{Y|\boldsymbol{x}_p}[c_{\mathrm{FP}}\cdot \mathbb{I}(a=1,Y=0) + c_{\mathrm{FN}}\cdot \mathbb{I}(a=0,Y=1)]. \tag{11}$$

Here, the expectation is taken over the state of the world Y, conditional on the observed public votes $\boldsymbol{x}_p$. The firm seeks to choose action $a\in\{0,1\}$ to minimize this conditional expected cost.

Similarly, the expert's objective function must reflect these asymmetric costs, as their interests are partially aligned with the firm's success. The expected cost for expert i, given their private signal x_i and chosen vote $\tilde{x}_i$, becomes:

$$J_i(\tilde{x}_i;x_i) = \mathbb{E}_{Y,\boldsymbol{X}_p|x_i}\left[c_{\mathrm{FP}}\cdot \mathbb{I}(a=1,Y=0) + c_{\mathrm{FN}}\cdot \mathbb{I}(a=0,Y=1) + \omega\cdot w_{k,\alpha,\nu}(X_{p,i}) - B\cdot \frac{\mathbb{I}(X_{p,i}=Y)}{\sum_j \mathbb{I}(X_{p,j}=Y)}\right], \tag{12}$$

where the expectation incorporates uncertainty about the true state Y and the public votes $\boldsymbol{X}_p$ of all experts, given expert i's signal x_i. The terms represent the expected share of the firm's error cost, the expected privacy cost, and the expected accuracy bonus, respectively.

C.2. Firm's Optimal Decision Rule

With the asymmetric prior θ and error costs c_{FP} and c_{FN}, the firm's optimal decision rule deviates from the simple majority rule. The firm aims to minimize its expected cost conditional on the observed public votes $\boldsymbol{x}_p$. It will choose action $a=1$ if the expected cost of doing so is less than the expected cost of choosing $a=0$. That is, $a=1$ if $\mathbb{E}[C(1;\nu,\boldsymbol{x}_p)] < \mathbb{E}[C(0;\nu,\boldsymbol{x}_p)]$. Substituting from Equation 11, this condition becomes:

$$\Pr[Y=0\mid \boldsymbol{x}_p]\cdot c_{\mathrm{FP}} < \Pr[Y=1\mid \boldsymbol{x}_p]\cdot c_{\mathrm{FN}}.$$

Rearranging this inequality, we find that the firm chooses $a = 1$ if and only if the posterior probability of state $Y = 1$ exceeds a specific threshold:

$$\Pr[Y = 1 \mid \boldsymbol{x}_p] > \frac{c_{\mathrm{FP}}}{c_{\mathrm{FP}} + c_{\mathrm{FN}}}. \tag{13}$$

Let $\tau = \frac{c_{\mathrm{FP}}}{c_{\mathrm{FP}}+c_{\mathrm{FN}}}$ denote this decision threshold.

To make this rule operational, we express the posterior probability in terms of the observed data. Let $s = \sum_{i=1}^{2k+1} x_{p,i}$ be the number of public votes equal to 1 among the $2k+1$ experts. Assuming experts vote truthfully and their public votes $X_{p,i}$ are conditionally independent given Y, the likelihood of observing s votes for 1 given $Y = 1$ is $\binom{2k+1}{s}\phi^s(1-\phi)^{2k+1-s}$, and given $Y = 0$ is $\binom{2k+1}{s}(1-\phi)^s\phi^{2k+1-s}$, where $\phi = \Pr[X_{p,i} = 1|Y = 1] = \Pr[X_{p,i} = 0|Y = 0]$ is the accuracy of a public vote. Using Bayes' theorem, the posterior probability is:

$$\Pr[Y = 1 \mid \boldsymbol{x}_p] = \frac{\Pr[\boldsymbol{x}_p \mid Y = 1]\theta}{\Pr[\boldsymbol{x}_p \mid Y = 1]\theta + \Pr[\boldsymbol{x}_p \mid Y = 0](1-\theta)} = \frac{\theta}{\theta + (1-\theta)\frac{\Pr[\boldsymbol{x}_p|Y=0]}{\Pr[\boldsymbol{x}_p|Y=1]}}.$$

The likelihood ratio is $\frac{\Pr[\boldsymbol{x}_p|Y=0]}{\Pr[\boldsymbol{x}_p|Y=1]} = \frac{(1-\phi)^s\phi^{2k+1-s}}{\phi^s(1-\phi)^{2k+1-s}} = \left(\frac{1-\phi}{\phi}\right)^{2s-(2k+1)}$. Substituting this into the posterior expression gives:

$$\Pr[Y = 1 \mid \boldsymbol{x}_p] = \frac{\theta}{\theta + (1-\theta)\left(\frac{1-\phi}{\phi}\right)^{2s-(2k+1)}}. \tag{14}$$

Now, substituting this into the decision rule (13) and rearranging, we find that the firm chooses $a = 1$ if:

$$\left(\frac{\phi}{1-\phi}\right)^{2s-(2k+1)} > \frac{c_{\mathrm{FP}}(1-\theta)}{c_{\mathrm{FN}}\theta}.$$

Taking the natural logarithm of both sides (noting $\ln(\frac{\phi}{1-\phi}) > 0$ since $\phi > 1/2$) yields:

$$(2s - (2k+1))\ln\left(\frac{\phi}{1-\phi}\right) > \ln\left(\frac{c_{\mathrm{FP}}(1-\theta)}{c_{\mathrm{FN}}\theta}\right).$$

Solving for s, the firm chooses $a = 1$ if:

$$s > k + \frac{1}{2} + \frac{1}{2}\frac{\ln\left(\frac{c_{\mathrm{FP}}(1-\theta)}{c_{\mathrm{FN}}\theta}\right)}{\ln\left(\frac{\phi}{1-\phi}\right)}. \tag{15}$$

Let us define the adjustment term λ as:

$$\lambda = -\frac{1}{2}\frac{\ln\left(\frac{c_{\mathrm{FP}}(1-\theta)}{c_{\mathrm{FN}}\theta}\right)}{\ln\left(\frac{\phi}{1-\phi}\right)} = \frac{1}{2}\frac{\ln\left(\frac{c_{\mathrm{FN}}\theta}{c_{\mathrm{FP}}(1-\theta)}\right)}{\ln\left(\frac{\phi}{1-\phi}\right)}. \tag{16}$$

Then the condition (15) becomes $s > k + \frac{1}{2} - \lambda$. Since s must be an integer, the firm's optimal rule is to choose action $a = 1$ if $s \geq t$, and $a = 0$ if $s < t$, where the threshold t is given by:

$$t = \left\lceil k + \frac{1}{2} - \lambda \right\rceil. \tag{17}$$

Note that when the prior is symmetric ($\theta = 1/2$) and costs are equal ($c_{\mathrm{FP}} = c_{\mathrm{FN}}$), we have $\lambda = 0$, and the threshold becomes $t = \lceil k + 1/2 \rceil = k + 1$. This recovers the simple majority rule (at least $k+1$ votes for '1' result in action $a = 1$) from the symmetric baseline model. The term λ adjusts the threshold based on the relative prior likelihoods and error costs.

C.3. Expert's Truth-Telling Condition

The change in the firm's decision rule necessitates a recalculation of the conditions under which experts find it optimal to vote truthfully. As in the baseline model, the critical decision point arises for an expert who observes a signal $x_i = 0$. Such an expert compares the expected cost of voting truthfully ($\tilde{x}_i = 0$) versus voting untruthfully ($\tilde{x}_i = 1$).

First, let's determine the expert's posterior beliefs about the state of the world after observing $x_i = 0$. Using Bayes' rule:

$$p_1 = \Pr[Y = 1 \mid x_i = 0] = \frac{\Pr[x_i = 0 \mid Y = 1]\theta}{\Pr[x_i = 0]} = \frac{(1-\alpha)\theta}{(1-\alpha)\theta + \alpha(1-\theta)}, \tag{18}$$

$$p_0 = \Pr[Y = 0 \mid x_i = 0] = 1 - p_1 = \frac{\alpha(1-\theta)}{(1-\alpha)\theta + \alpha(1-\theta)}. \tag{19}$$

Note that $p_0 \geq p_1$ if and only if $\alpha(1-\theta) \geq (1-\alpha)\theta$, which simplifies to $\alpha \geq \theta$. If the expert's accuracy α is higher than the prior θ, then observing the signal $x_i = 0$ makes state $Y = 0$ more likely in the expert's posterior belief. We will maintain the focus on this standard case where the signal provides meaningful information relative to the prior (i.e., $\alpha > \theta$, implying $p_0 > p_1$).

An expert votes truthfully ($\tilde{x}_i = 0$ when $x_i = 0$) if $J_i(0;0) \leq J_i(1;0)$. The difference in expected cost, $\Delta J = J_i(0;0) - J_i(1;0)$, must be non-positive. By examining Eq. (12), this difference arises from three components: impact on the firm's decision error cost (when pivotal), privacy cost, and the accuracy bonus.

First, firm's prediction error difference: The expert's vote only influences the firm's action a when the number of "1" votes among the other $2k$ experts, denoted s_{-i}, is exactly $t-1$. Let $\xi = 1 - 0.5\nu$. The difference in the expected error cost, conditional on being pivotal and $x_i = 0$, is $\Delta E_{\text{pivotal}} = (2\xi - 1)(p_1 c_{\text{FN}} - p_0 c_{\text{FP}})$. The probability of being pivotal is $P_{\text{pivotal}} = \Pr[\text{Bin}(2k, \phi) = t-1]$. The total expected difference in the error cost component is $\Delta J_{\text{firm\_error}} = \Delta E_{\text{pivotal}} \cdot P_{\text{pivotal}}$.

Second, privacy cost difference: $\Delta J_{\text{privacy}} = \omega(v_{k,\alpha,\nu}(0,0) - v_{k,\alpha,\nu}(1,0)) = \omega \Delta v_{k,\alpha}(\tilde{\phi})$. Here $\tilde{\phi} = \phi - 1/2$.

Finally, accuracy bonus difference: Let $\gamma = \mathbb{E}_{N \sim \text{Bin}(2k,\phi)}[\frac{B}{1+N}]$. The difference in the expected bonus is $\Delta J_{\text{bonus}} = \mathbb{E}[\text{Bonus}|\tilde{x}_i = 1, x_i = 0] - \mathbb{E}[\text{Bonus}|\tilde{x}_i = 0, x_i = 0] = \gamma(\Pr[X_{p,i} = Y|0,0] - \Pr[X_{p,i} = Y|1,0]) = \gamma(2\xi - 1)(p_0 - p_1)$.

Combining these components, the truth-telling condition $J_i(0;0) \leq J_i(1;0) \iff \Delta J \leq 0$ becomes:

$$\omega \Delta v_{k,\alpha}(\tilde{\phi}) + (2\xi - 1)(p_1 c_{\text{FN}} - p_0 c_{\text{FP}}) P_{\text{pivotal}} + \gamma(2\xi - 1)(p_1 - p_0) \leq 0.$$

Rearranging to isolate ω:

$$\omega \Delta v_{k,\alpha}(\tilde{\phi}) \leq (2\xi - 1)[(p_0 c_{\text{FP}} - p_1 c_{\text{FN}}) P_{\text{pivotal}} + \gamma(p_0 - p_1)].$$

Using similar techniques, we derive a similar condition for truth-telling when $x_i = 1$: $J_i(0,1) - J_i(1,1) \geq 0$:

$$\omega(v_{k,\alpha,\nu}(0,1) - v_{k,\alpha,\nu}(1,1)) \geq (2\xi - 1)[(q_1 c_{\text{FN}} - q_0 c_{\text{FP}}) P_{\text{pivotal}} + \gamma(q_0 - q_1)],$$

where $q_1 = \Pr[Y = 1|x_i = 1]$ and $q_0 = \Pr[Y = 0|x_i = 1]$. Without loss of generality, we consider only the case when $q_1 \geq q_0$, which holds when $\alpha > 1 - \theta$. Further simplifying and using the relationship $2\xi - 1 = (1-\nu)(2\alpha - 1) = \frac{2\phi - 1}{2\alpha - 1}$ as well as that $v_{k,\alpha,\nu}(0,1) - v_{k,\alpha,\nu}(1,1) = v_{k,\alpha,\nu}(0,0) - v_{k,\alpha,\nu}(1,0) = \Delta v_{k,\alpha}(\tilde{\phi})$, we arrive at the following proposition.

PROPOSITION 5 **(Truth-Telling Condition with Asymmetric Parameters)**. *In the extended model with prior θ, error costs $c_{\mathrm{FP}}, c_{\mathrm{FN}}$, team size $2k+1$, and garbling level ν, experts vote truthfully (i.e., $\tilde{x}_i = x_i$ for all i) if and only if their privacy sensitivity ω satisfies both of the following conditions:*

$$\omega \le \frac{1}{2\alpha-1}\frac{2\tilde{\phi}}{\Delta v_{k,\alpha}(\tilde{\phi})}\left[\underbrace{(p_0(c_{\mathrm{FP}}+c_{\mathrm{FN}})-c_{\mathrm{FN}})\Pr[\mathrm{Bin}(2k,\phi)=t-1]}_{pivotality}+\underbrace{\gamma(2p_0-1)}_{bonus}\right] \tag{20}$$

$$\omega \ge \frac{1}{2\alpha-1}\frac{2\tilde{\phi}}{\Delta v_{k,\alpha}(\tilde{\phi})}\left[\underbrace{(c_{\mathrm{FP}}-q_1(c_{\mathrm{FP}}+c_{\mathrm{FN}}))\Pr[\mathrm{Bin}(2k,\phi)=t-1]}_{pivotality}-\underbrace{\gamma(2q_1-1)}_{bonus}\right] \tag{21}$$

where $\phi = \alpha - \frac{1}{2}(2\alpha-1)\nu$ is the public vote accuracy, $\tilde{\phi} = \phi - 1/2$, $\Delta v_{k,\alpha}(\tilde{\phi}) = v_{k,\alpha,\nu}(0,0) - v_{k,\alpha,\nu}(1,0)$ is the privacy cost difference.

This condition generalizes Proposition 1 to the asymmetric setting. Note that, under our assumptions $p_0 > p_1$ and $q_1 > q_0$ (or, equivalently, $\alpha > \max\{1-\theta,\theta\}$), the bonus term contributes positively to the RHS of the first condition and negatively to the RHS of the second condition. When $\theta = 1/2$ and $c_{\mathrm{FP}} = c_{\mathrm{FN}}$, the condition reduces to the symmetric case.

C.4. Perfect Learning with Asymmetric Parameters

We now explore the conditions necessary for achieving perfect learning in this extended model. Perfect learning occurs if the firm's expected cost converges to zero as the number of experts $(2k+1)$ approaches infinity. First, we generalize the relationship between the convergence rate of the public vote accuracy ϕ_k and the possibility of perfect learning.

LEMMA 2 **(Optimal Threshold Decision-Making and Perfect Learning)**. *Consider the extended model with prior $\theta \in (0,1)$ and costs $c_{\mathrm{FP}} > 0, c_{\mathrm{FN}} > 0$. Let $\{\phi_k\}_k$ be a sequence of probabilities (for each team size $2k+1$) of any expert's public vote being correct. Then if the firm uses the optimal decision threshold $t_k = \lceil k + 1/2 - \lambda_k \rceil$, where $\lambda_k = \frac{1}{2}\frac{\ln(c_{\mathrm{FN}}\theta/(c_{\mathrm{FP}}(1-\theta)))}{\ln(\phi_k/(1-\phi_k))}$, its expected error cost,*

$$\mathcal{C}_k = c_{\mathrm{FN}}\theta \cdot \Pr(\mathrm{Bin}(2k+1,\phi_k) < t_k) + c_{\mathrm{FP}}(1-\theta)\cdot \Pr(\mathrm{Bin}(2k+1, 1-\phi_k) \ge t_k),$$

converges to zero as $k \to \infty$ if and only if $\lim_{k\to\infty} k\tilde{\phi}_k^2 = +\infty$, where $\tilde{\phi}_k = \phi_k - 1/2$.

Proof: We seek the condition on the sequence $\{\tilde{\phi}_k\}_k$ such that $\lim_{k\to\infty}\mathcal{C}_k = 0$. This requires both the false negative probability $(\Pr(\mathrm{Bin}(2k+1,\phi_k) < t_k))$ and the false positive probability $(\Pr(\mathrm{Bin}(2k+1,1-\phi_k) \ge t_k))$ to converge to zero.

Let $C_{\mathrm{asym}} = \ln\left(\frac{c_{\mathrm{FN}}\theta}{c_{\mathrm{FP}}(1-\theta)}\right)$. Approximate the threshold t_k for large k where $\tilde{\phi}_k \to 0$ to obtain $\lambda_k \approx \frac{C_{\mathrm{asym}}}{8\tilde{\phi}_k}$ and $t_k \approx k + \frac{1}{2} - \frac{C_{\mathrm{asym}}}{8\tilde{\phi}_k}$. To see that, note that $\ln\left(\frac{\phi_k}{1-\phi_k}\right) = \ln\left(\frac{1+2\tilde{\phi}_k}{1-2\tilde{\phi}_k}\right) \approx \ln[(1+2\tilde{\phi}_k)^2] \approx 4\tilde{\phi}_k^2$ for small $\tilde{\phi}_k$.

We will first show that the probability of false negative is converging to zero under the condition of the Lemma. Denote the random variable $X_k^{(1)} \sim \mathrm{Bin}(2k+1,\phi_k)$, which captures the number of correct votes and has a mean $\mu_k^{(1)} = k + 1/2 + (2k+1)\tilde{\phi}_k$. Also denote $D_k^{(+)} = \mu_k^{(1)} - t_k \approx (2k+1)\tilde{\phi}_k + \frac{C_{\mathrm{asym}}}{8\tilde{\phi}_k} \ge 0$. We use Hoeffding's inequality to bound $\Pr[X_k^{(1)} < t_k]$—the probability of false negatives outcomes to happen:

$$\Pr[X_k^{(1)} < t_k] = \Pr[X_k^{(1)} - \mu_k^{(1)} < -D_k^{(+)}] \le \exp\left(-\frac{2(D_k^{(+)})^2}{2k+1}\right) \approx \exp\left(-\frac{2}{2k+1}\left((2k+1)\tilde{\phi}_k + \frac{C_{\mathrm{asym}}}{8\tilde{\phi}_k}\right)^2\right).$$

For the above expression to converge to 0, it must be that $\frac{1}{2k+1}\left((2k+1)\tilde{\phi}_k + \frac{C_{\text{asym}}}{8\tilde{\phi}_k}\right)^2 \to \infty$. Expanding this expression, we obtain $(2k+1)\tilde{\phi}_k^2 + \frac{C_{\text{asym}}}{4} + \frac{C_{\text{asym}}^2}{64(2k+1)\tilde{\phi}_k^2}$. It then must be that the dominant term $2k\tilde{\phi}_k^2 \to +\infty$ when $k \to \infty$ for the probability of false negative to converge to zero.

To show that the probability of false positive also converges to zero under the same condition of the Lemma, we can use similar techniques. In particular, we work with the random variable $Y_k(0) = (2k+1) - X_k^{(0)}$ and the threshold $t'_k = (2k+1) - t_k$. □

The above Lemma confirms that the fundamental scaling requirement for $\tilde{\phi}_k$ for perfect learning to occur is robust to the introduction of asymmetries.

Theorem 2 **(Achieving Perfect Learning with Asymmetric Parameters)**. *Consider the extended model satisfying Assumptions 1 and 2 regarding the privacy cost structure, and assume $\alpha > \theta$ (so $p_0 > p_1$).*

- *In the absence of a randomized response mechanism ($\nu = 0$, so $\tilde{\phi} = \alpha - 1/2$), the firm can achieve perfect learning ($\lim_{k\to\infty} \mathcal{C}_k = 0$) if and only if the firm provides accuracy bonuses ($B > 0$) and the safety-in-numbers effect is sufficiently strong ($\rho > 1$).*
- *When the firm uses the optimal randomized response mechanism (choosing ν_k^* to maximize ϕ_k subject to the truth-telling constraint (20)), perfect learning is achievable if and only if $B > 0$ and the safety-in-numbers effect satisfies $\rho > 1 - \beta/2$.*

When perfect learning is achieved, the optimal strategy involves employing an infinitely large team ($2k^* + 1 = \infty$ in the limit), and the probability of the firm choosing the correct action approaches 1.

Proof: The proof mirrors the logic of Theorem 1, using Proposition 5 and Lemma 2.

Necessity: Perfect learning requires $\lim_{k\to\infty} k\tilde{\phi}_k^2 = +\infty$ as per Lemma 2. The truth-telling conditions must also hold in the limit. The bonus terms dominate asymptotically the pivotality terms for large k. To see that, note that, as in the symmetric case, the bonus term decreases with k as $1/k$ while pivotality term decays exponentially with k. Note also that, under our assumptions, the bonus term is negative in the truth-telling condition under $x_i = 1$ (second inequality of Proposition 5). Hence, asymptotically, the second truth-telling condition of Proposition 5 is always satisfied. Thus, we only focus on the first truth-telling condition of Propostion 5. As per Assumption 2, the privacy cost term $\Delta v_{k,\alpha}(\tilde{\phi}_k)$ scales as $\tilde{\phi}_k^{1+\beta} k^{-\rho}$. Thus, for large k, the truth-telling condition requires (up to the constants):

$$\omega \lesssim \frac{2\tilde{\phi}_k}{\Delta v_{k,\alpha}(\tilde{\phi}_k)} \frac{1}{2\alpha - 1} \frac{B(p_0 - p_1)}{k} \approx B \frac{C_1 \tilde{\phi}_k / k}{C_2 \tilde{\phi}_k^{1+\beta} k^{-\rho}} = B \cdot \tilde{C} \frac{k^{\rho-1}}{\tilde{\phi}_k^{\beta}},$$

which can be rewritten as $\omega \tilde{\phi}_k^{\beta} k^{1-\rho} \lesssim B \cdot \tilde{C}$.

Now consider the two possibilities for achieving perfect learning:

- Case 1: $\nu = 0$, $\tilde{\phi}_k = \bar{\phi} > 0$. Condition for learning (i.e., $k\tilde{\phi}_k^2 \to \infty$) holds. Condition for truth-telling requires $\omega \bar{\phi}^{\beta} k^{1-\rho} \lesssim B \cdot \tilde{C}$. This holds for $k \to \infty$ only if $\rho > 1$ (note that $B > 0$ by setup).
- Case 2: Optimal ν_k^*, $\tilde{\phi}_k \to 0$. The optimality of $\tilde{\phi}_k$ implies that the truth-telling condition is asymptotically binding (otherwise, one could increase $\tilde{\phi}_k$ to improve learning) and $\lim_{k\to\infty} \omega \tilde{\phi}_k^{\beta} k^{1-\rho} = B \cdot \tilde{C}$. This suggests that $\tilde{\phi}_k$ scales asymptotically as $k^{-(1-\rho)/\beta}$. To show that this is indeed true, we need to show that $\tilde{\phi}_k / k^{-(1-\rho)/\beta} = \tilde{\phi}_k k^{(1-\rho)/\beta}$ converges to some constant. Consider the following three subcases:

— *Case 1a:* $\liminf_{k\to\infty} \tilde{\phi}_k k^{(1-\rho)/\beta} = 0$. Since $B \cdot \tilde{C} > 0$, this subcase contradicts the necessity of the truth-telling condition to be biding.

— *Case 1b:* $\liminf_{k\to\infty} \tilde{\phi}_k k^{(1-\rho)/\beta} = +\infty$. Since $B \cdot \tilde{C} < \infty$, this subcase contradicts the truth-telling condition.

— *Case 1c:* $\liminf_{k\to\infty} \tilde{\phi}_k k^{(1-\rho)/\beta} = C'$, *where* C' *is some constant.* This is the only viable subcase and it holds that $C' = \left(\frac{B\cdot\tilde{C}}{\omega}\right)^{1/\beta}$.

Hence, we showed, that indeed $\tilde{\phi}_k$ asymptotically scales as $k^{-(1-\rho)/\beta}$. Then, for the learning condition to hold, i.e., $k\tilde{\phi}_k^2 \to \infty$ to be true, we must have that $1 - \frac{2(1-\rho)}{\beta} > 0$, which implies that $\rho > 1 - \beta/2$.

To sum up: the condition $\rho > 1 - \beta/2$ (with $B > 0$) is necessary when randomized response us used, and $\rho > 1$ (with $B > 0$) is necessary without it.

Sufficiency: The proof follows that of Theorem 1—for each of the required parameter ranges, one can construct appropriate sequences $\tilde{\phi}_k$ that satisfy both truth-telling and the learning conditions. In particular, when $\rho > 1$ and $B > 0$, it is sufficient to pick $\tilde{\phi}_k = \alpha - \frac{1}{2} > 0$ for all k. When $1 - \beta/2 < \rho < 1$ and $B > 0$, one could use the sequence $\tilde{\phi}_k = C' k^{-(1-\rho)/\beta}$.

As a final remark, we note that the proof above relies on the assumption $\alpha > \max\{1-\theta, \theta\}$ ensuring the bonus term incentivizes truth-telling for $x_i = 0$. The specific values of $\theta, c_{\mathrm{FP}}, c_{\mathrm{FN}}$ affect constants but not the asymptotic feasibility condition $\rho > 1 - \beta/2$. □

This result demonstrates the robustness of our main findings regarding perfect learning. The critical interplay between the structural properties of privacy cost (safety in numbers ρ, adversarial restraint β) and the incentive mechanisms (accuracy bonus B, randomized response ν) dictates whether the firm can overcome privacy concerns to aggregate information effectively in large teams, even when facing asymmetric priors and error costs, provided the expert's signal is sufficiently informative relative to the prior ($\alpha > \max\{1-\theta, \theta\}$). The condition $\rho > 1 - \beta/2$ remains the key threshold for achieving the wisdom of crowds in the presence of privacy-conscious experts when both randomized response and accuracy bonuses are employed.

Appendix D: Unknown Expert Abilities

We now extend our model to consider the case where experts' abilities are uncertain. Specifically, each expert's ability is drawn from a two-point distribution, with only the expert knowing their own ability with certainty.

D.1. Model Modification for Unknown Expert Abilities

In the baseline model, we assumed that each expert has a known ability α that determines the quality of their signal. We now modify this assumption as follows:

1. Each expert $i \in \mathcal{I}$ has an ability parameter $\alpha_i \in \{\alpha^L, \alpha^H\}$, where $\alpha_L < \alpha_H$ and both are greater than 0.5.

2. The probability that expert i has high ability ($\alpha_i = \alpha^H$) is $\lambda_A \in (0,1)$, and the probability of low ability ($\alpha_i = \alpha^L$) is $1 - \lambda_A$.

3. Each expert knows their own ability with certainty, but neither the firm nor other experts can observe this ability.

4. The distribution of abilities is common knowledge among all players.

This modification captures the realistic scenario where experts vary in their ability to predict the state of the world, but their exact level of expertise is private information. Let each agent be one of the two types, corresponding to their ability level (α^L or α^H). We also define $\bar{\alpha} = \lambda_A \alpha^H + (1-\lambda_A)\alpha^L$ as the expected ability of a randomly chosen expert.

Note that the firm's optimal decision rule with unknown expert abilities is identical to the one with known homogeneous abilities. This occurs because, without observing individual expert types, the firm treats each expert as having the average ability $\bar{\alpha}$. However, as we will see, the experts' strategic behavior changes because of their private knowledge of their own abilities.

D.2. The Expert's Truth-Telling Condition

We now derive the condition under which an expert will vote truthfully, taking into account their private knowledge of their own ability.

Proposition 6 **(Truth-Telling Condition with Unknown Abilities)**. *An expert with ability $\alpha_i \in \{\alpha^L, \alpha^H\}$ votes truthfully if and only if their privacy sensitivity ω is sufficiently low:*

$$\omega \le \frac{2\tilde{\phi}^z}{\Delta v_{k,\alpha_i}(\tilde{\phi}^z)} \cdot \left[\frac{1}{4^k} \cdot \binom{2k}{k}(1-4\tilde{\phi}^2)^k + B \cdot \frac{1-2^{-(2k+1)} \cdot (1-2\tilde{\phi})^{2k+1}}{(\tilde{\phi}+\frac{1}{2})(2k+1)}\right], \tag{22}$$

where $\tilde{\phi}^z = \phi^z - \frac{1}{2}$, $\phi_i = \alpha^z - \frac{1}{2}(2\alpha^z - 1)\nu$ is the probability that expert i's with type $z \in \{L, H\}$ public vote is correct, $\tilde{\phi} = \bar{\alpha} - \frac{1}{2}(2\bar{\alpha} - 1)\nu$ is the expected probability that the expert with unknown ability is correct, and $\Delta v_{k,\alpha^z}(\tilde{\phi}^z) = v_{k,\alpha^z,\nu}(0,0) - v_{k,\alpha^z,\nu}(1,0)$ is the difference in expected privacy costs.

Proof: Following the same approach as in Proposition 1, an expert votes truthfully if the expected cost of doing so is lower than that of voting untruthfully. The only modification here is that the expert uses his own correctness probability to weigh the payoffs in different action, and the average probability of correctness to decide on the pivotality and probability of winning the bonus.

We can see that the structure of the truth-telling condition remains the same, and that if a low-ability expert is truthful, then the high-ability expert is truthful as well. The asymptotic results for the perfect learning regime therefore remain the same, with perfect learning being possible in the same range of parameters.

Appendix E: Unknown Privacy Costs

In this section, we extend our model to consider the case where privacy costs are unknown to experts. This captures scenarios where experts face uncertainty about the reputational or personal consequences of their signals. We modify the general model presented in Section 3.1 by introducing uncertainty about privacy costs. Let (ω_0, ω_1) be a tuple of privacy parameters, where:

- ω_0 represents the privacy cost incurred when the adversary believes the expert's signal is 0
- ω_1 represents the privacy cost incurred when the adversary believes the expert's signal is 1

In the original setting, we had $\omega_1 = 0$, meaning experts only incurred privacy costs for signals indicating "bad news." In this extended setting, ω_1 can be greater than zero, allowing for privacy costs associated with "good news" as well.

Furthermore, we assume that this tuple (ω_0, ω_1) follows a two-point distribution: the tuple takes values (ω_0^A, ω_1^A) with probability p_A and (ω_0^B, ω_1^B) with probability $p_B = 1 - p_A$.

The expert's cost function becomes:

$$J_i(\tilde{x}i; x_i) = \mathbb{E}_{Y, \boldsymbol{X}_p|x_i}\left[|Y-a| + \mathbb{E}_{(\omega_0,\omega_1)}\left[\sum_{l\in 0,1} \omega_l w_{k,\alpha,\nu}(l)\mathbb{I}(X_{p,i}=l)\right] - B\cdot\frac{\mathbb{I}(X_{p,i}=Y)}{\sum_j \mathbb{I}(X_{p,j}=Y)}\right], \tag{23}$$

where $\mathbb{E}_{(\omega_0,\omega_1)}$ denotes expectation over the distribution of privacy parameters.

Note that the firm's optimal action rule, the expert's pivotality condition, and the condition for perfect remain the same—the only thing that changes is the expert's truth-telling condition.

E.1. Expert's Truth-Telling Conditions

We now derive the conditions under which experts vote truthfully for both $x_i = 0$ and $x_i = 1$.

Let us define the expected privacy cost for an expert with signal x_i who votes $\tilde{x}_i$ as:

$$v_{k,\alpha,\nu}(\tilde{x}_i, x_i) = \mathbb{E}_{X_{p,i}|\tilde{x}_i,x_i}\left[\mathbb{E}_{(\omega_0,\omega_1)}[\omega_0\cdot w_{k,\alpha,\nu}(0)\cdot\mathbb{I}(X_{p,i}=0) + \omega_1\cdot w_{k,\alpha,\nu}(1)\cdot\mathbb{I}(X_{p,i}=1)]\right] \tag{24}$$

For an expert with signal $x_i = 0$, the expected cost of voting truthfully is:

$$J_i(0;0) = v_{k,\alpha,\nu}(0,0) + (1-\phi)\binom{2k}{k}\phi^k(1-\phi)^k + \Pr[\text{Bin}(2k,\phi) < k] - \phi\gamma, \tag{25}$$

The expected cost of voting untruthfully is:

$$J_i(1;0) = v_{k,\alpha,\nu}(1,0) + \phi\binom{2k}{k}\phi^k(1-\phi)^k + \Pr[\text{Bin}(2k,\phi) < k] - (1-\phi)\gamma, \tag{26}$$

Similarly, for an expert with signal $x_i = 1$, the expected cost of voting truthfully is:

$$J_i(1;1) = v_{k,\alpha,\nu}(1,1) + (1-\phi)\binom{2k}{k}\phi^k(1-\phi)^k + \Pr[\text{Bin}(2k,\phi) < k] - \phi\gamma, \tag{27}$$

The expected cost of voting untruthfully is:

$$J_i(0;1) = v_{k,\alpha,\nu}(0,1) + \phi\binom{2k}{k}\phi^k(1-\phi)^k + \Pr[\text{Bin}(2k,\phi) < k] - (1-\phi)\gamma, \tag{28}$$

Define $\Delta v^0_{k,\alpha}(\tilde{\phi}) = v_{k,\alpha,\nu}(0,0) - v_{k,\alpha,\nu}(1,0)$ as the difference in expected privacy costs between voting truthfully and untruthfully when the signal is $x_i = 0$. Similarly, define $\Delta v^1_{k,\alpha}(\tilde{\phi}) = v_{k,\alpha,\nu}(1,1) - v_{k,\alpha,\nu}(0,1)$ for signal $x_i = 1$.

An expert votes truthfully when signal is $x_i = 0$ if and only if $J_i(0;0) \le J_i(1;0)$, which gives:

$$\Delta v^0_{k,\alpha}(\tilde{\phi}) \le 2\tilde{\phi}\cdot\left[\frac{1}{4^k}\cdot\binom{2k}{k}(1-4\tilde{\phi}^2)^k + B\cdot\frac{1-2^{-(2k+1)}\cdot(1-2\tilde{\phi})^{2k+1}}{(\tilde{\phi}+\frac{1}{2})(2k+1)}\right] \tag{29}$$

Similarly, an expert votes truthfully when signal is $x_i = 1$ if and only if $J_i(1;1) \le J_i(0;1)$, which gives:

$$\Delta v^1_{k,\alpha}(\tilde{\phi}) \le 2\tilde{\phi}\cdot\left[\frac{1}{4^k}\cdot\binom{2k}{k}(1-4\tilde{\phi}^2)^k + B\cdot\frac{1-2^{-(2k+1)}\cdot(1-2\tilde{\phi})^{2k+1}}{(\tilde{\phi}+\frac{1}{2})(2k+1)}\right] \tag{30}$$

This yields the following proposition:

Proposition 7. *If the firm employs $2k+1$ experts and sets garbling level ν, then experts vote truthfully (i.e., $\tilde{x}_i = x_i$ for every i) if and only if both conditions hold:*

$$\Delta v^0_{k,\alpha}(\tilde{\phi}) \le 2\tilde{\phi} \cdot \left[\frac{1}{4^k} \cdot \binom{2k}{k}(1-4\tilde{\phi}^2)^k + B \cdot \frac{1 - 2^{-(2k+1)} \cdot (1-2\tilde{\phi})^{2k+1}}{(\tilde{\phi}+\frac{1}{2})(2k+1)}\right] \tag{31}$$

$$\Delta v^1_{k,\alpha}(\tilde{\phi}) \le 2\tilde{\phi} \cdot \left[\frac{1}{4^k} \cdot \binom{2k}{k}(1-4\tilde{\phi}^2)^k + B \cdot \frac{1 - 2^{-(2k+1)} \cdot (1-2\tilde{\phi})^{2k+1}}{(\tilde{\phi}+\frac{1}{2})(2k+1)}\right] \tag{32}$$

where $\tilde{\phi} = \phi - 1/2$ and $\phi = \alpha - \frac{1}{2}(2\alpha - 1)\nu$ is the probability that any expert's public vote is correct.

This proposition generalizes the truth-telling condition from the original model. While the original model only needed to ensure truthful voting when the expert's signal was $x_i = 0$ (since truthful voting was always optimal when $x_i = 1$ due to $\omega_1 = 0$), our extended model requires conditions for both signal values.

To distinguish between perfect and imperfect learning regimes in our extended model with unknown privacy costs, we make the following assumption about the privacy loss functions $\Delta v^0_{k,\alpha}(\tilde{\phi})$ and $\Delta v^1_{k,\alpha}(\tilde{\phi})$:

Assumption 3 **(Limiting Behavior of Privacy Loss Functions)**. *Let β and ρ be strictly positive constants. For $j \in \{0,1\}$ and any k:*

$$\lim_{\tilde{\phi}\to\frac{1}{2}} \frac{\Delta v^j_{k,\alpha}(\tilde{\phi})}{\tilde{\phi}^{1+\beta}} = \frac{C_j}{k^\rho} + o\left(\frac{1}{k^\rho}\right) \tag{33}$$

For any $\tilde{\phi}$:

$$\lim_{k\to\infty} \frac{\Delta v^j_{k,\alpha}(\tilde{\phi})}{k^{-\rho}} = C_j \tilde{\phi}^{1+\beta} + o\left(\tilde{\phi}^{1+\beta}\right) \tag{34}$$

where C_j are positive constants.

This assumption generalizes the one in the original model, ensuring that both privacy loss functions (for signals 0 and 1) increase superlinearly with informativeness and decrease with team size due to the "safety in numbers" effect. Under this assumption, the truth-telling conditions in Proposition 7 adopt the same structure as in the base model, implying that asymptotically, the results of Theorem 1 carry over to the extended model.

Appendix F: Quadratic Loss Function

In this section, we extend our base model to a setting in which the firm chooses a continuous action $a(\boldsymbol{x}_p) \in [0,1]$ and seeks to minimize a quadratic loss function, $\ell(a,Y) = (Y-a)^2$. We demonstrate that our core findings regarding the conditions for perfect learning (Theorem 1) are robust to this change. The proof proceeds in two main parts. First, we establish that the mathematical benchmark for perfect learning—the necessary rate of convergence for public vote accuracy—is identical to the one in Lemma 1. Second, we prove that the strategic conditions required to achieve this benchmark are asymptotically the same as in our main analysis.

F.1. The Condition for Perfect Learning with Quadratic Loss

In the quadratic loss setting, the firm's optimal action for a given public vote profile is the posterior mean of the state, $a^*(\boldsymbol{x}_p) = \mu(s) = \Pr(Y = 1 \mid s)$, where s is the sum of '1'-votes. The firm's minimized expected cost for a given vote profile is the posterior variance of the state, $g(s) = \mu(s)(1 - \mu(s))$. The firm's total ex-ante expected cost is therefore $\mathcal{C}_{quad} = \mathbb{E}_{\boldsymbol{X}_p}[g(s(\boldsymbol{X}_p))]$. We will compare it to the firm's cost under absolute loss that is used in our main model, $\mathcal{C}_{abs} = \mathbb{E}_{\boldsymbol{X}_p}[\min(\mu(s), 1 - \mu(s))]$.

LEMMA 3 **(Equivalence of Learning Conditions)**. *The firm's expected cost under quadratic loss, $\mathcal{C}_{quad}$, converges to zero as $k \to \infty$ if and only if its expected cost under absolute loss, $\mathcal{C}_{abs}$, converges to zero.*

Proof: The proof relies on a "sandwich" inequality that relates the two cost functions for any given posterior $\mu \in [0, 1]$:

$$\frac{1}{2}\min(\mu, 1-\mu) \le \mu(1-\mu) \le \min(\mu, 1-\mu). \tag{35}$$

To prove this inequality, consider two cases.

- If $\mu \le 1/2$, then $\min(\mu, 1-\mu) = \mu$. The inequality becomes $\frac{1}{2}\mu \le \mu(1-\mu) \le \mu$. The right side holds since $1 - \mu \le 1$. The left side holds since $\mu \le 1/2$ implies $1/2 \le 1 - \mu$.
- If $\mu > 1/2$, then $\min(\mu, 1-\mu) = 1 - \mu$. The inequality becomes $\frac{1}{2}(1-\mu) \le \mu(1-\mu) \le 1 - \mu$. The right side holds since $\mu \le 1$. The left side holds since $\mu > 1/2$ implies $1/2 < \mu$.

Since the inequality holds for any posterior μ, we can take the expectation over all possible vote profiles $\boldsymbol{X}_p$:

$$\frac{1}{2}\mathbb{E}_{\boldsymbol{X}_p}[\min(\mu(s), 1-\mu(s))] \le \mathbb{E}_{\boldsymbol{X}_p}[\mu(s)(1-\mu(s))] \le \mathbb{E}_{\boldsymbol{X}_p}[\min(\mu(s), 1-\mu(s))].$$

This simplifies to $\frac{1}{2}\mathcal{C}_{abs} \le \mathcal{C}_{quad} \le \mathcal{C}_{abs}$. This relationship shows that $\mathcal{C}_{quad} \to 0$ if and only if $\mathcal{C}_{abs} \to 0$. From Lemma 1 in the main text, we know that $\mathcal{C}_{abs} \to 0$ if and only if $\lim_{k\to\infty} k\tilde{\phi}_k^2 = \infty$. Therefore, this same condition is necessary and sufficient for perfect learning under a quadratic loss function. □

F.2. Achieving the Perfect Learning Condition

Having established the benchmark for perfect learning, we now analyze the expert's strategic incentives to determine the conditions under which this benchmark is achievable. The expert's truth-telling condition, $J_i(0;0) \le J_i(1;0)$, now includes the marginal gain from reducing the firm's quadratic loss instead of the binary pivotality term. The marginal gain from a truthful vote is proportional to the expected reduction in squared error. This is given by $(2\xi - 1)\mathbb{E}_{Y,S_{-i}|X_i=0}[(\mu(S_{-i}) - Y)^2 - (\mu(S_{-i}+1) - Y)^2]$. Denote $S_0 \sim \mathrm{Bin}(2k, 1-\phi_k)$ and $S_1 \sim \mathrm{Bin}(2k, \phi_k)$. Then, using the law of iterated expectations (conditioning on $Y = 0$ and $Y = 1$) and exploiting the symmetry of the posterior $\mu(s)$ and the distributions of votes, this term simplifies:

$$\begin{aligned}\text{Marginal Gain} &\propto (2\xi - 1)\left(\alpha\mathbb{E}_{S_0}[\mu(S_0+1)^2 - \mu(S_0)^2] + (1-\alpha)\mathbb{E}_{S_1}[(1-\mu(S_1+1))^2 - (1-\mu(S_1))^2]\right)\\ &= (2\xi-1)(2\alpha-1)\mathbb{E}_{S_0}\left[\mu(S_0+1)^2 - \mu(S_0)^2\right].\end{aligned}$$

Since $(2\xi - 1)(2\alpha - 1) = 2\phi - 1$, the final expression is:

$$\text{Marginal Gain} = (2\phi - 1)\mathbb{E}_{S_0}\left[\mu(S_0+1)^2 - \mu(S_0)^2\right],$$

where S_0 is the number of incorrect votes among the other $2k$ experts when the true state is $Y = 0$. For our asymptotic results to hold, we must prove that this marginal gain term decays exponentially in k.

LEMMA 4 **(Exponential Decay of Marginal Gain)**. *The expected marginal gain,* $(2\phi_k - 1)\mathbb{E}_{S_0}[\mu(S_0 + 1)^2 - \mu(S_0)^2]$, *is bounded by a function that decays exponentially in k under the perfect learning condition.*

Proof: It suffices to show that $\mathbb{E}_{S_0}[\mu(S_0+1)^2 - \mu(S_0)^2]$ decays exponentially. First, we establish an upper bound:

$$\begin{aligned}\mathbb{E}_{S_0}[\mu(S_0+1)^2 - \mu(S_0)^2] &= \mathbb{E}_{S_0}[(\mu(S_0+1) - \mu(S_0))(\mu(S_0+1) + \mu(S_0))] \\ &\leq 2 \cdot \mathbb{E}_{S_0}[\mu(S_0+1) - \mu(S_0)],\end{aligned}$$

where the inequality follows from $0 < \mu(s) < 1$. We now bound the term $\mathbb{E}_{S_0}[\mu(S_0+1) - \mu(S_0)]$. The expectation is a sum over all possible values of S_0:

$$\sum_{s_0=0}^{2k} p(s_0)(\mu(s_0+1) - \mu(s_0)),$$

where $p(s_0) = \Pr(S_0 = s_0)$. We split this sum into two regions at a cutoff point s_c. A strategic choice for the cutoff is the midpoint between the mean of S_0 and the point of maximum uncertainty (k). In terms of $\tilde{\phi}_k = \phi_k - 1/2$, we have $\mathbb{E}S_0 = k(1-2\tilde{\phi}_k)$, and we choose the cutoff $s_c = k(1-\tilde{\phi}_k)$.

Region 1: $S_0 \geq s_c$ **(The "unlikely" zone).** In this region, we bound the contribution by the tail probability of the event:

$$\sum_{s_0=s_c}^{2k} p(s_0)(\mu(s_0+1) - \mu(s_0)) \leq \sum_{s_0=s_c}^{2k} p(s_0) = \Pr(S_0 \geq s_c),$$

since $\mu(s_0+1) - \mu(s_0) < 1$. We apply Hoeffding's inequality, $\Pr(S_0 - \mathbb{E}S_0 \geq t) \leq \exp(-2t^2/n)$, where $n = 2k$ and the deviation is $t = s_c - \mathbb{E}S_0 = k\tilde{\phi}_k$. This gives:

$$\Pr(S_0 \geq s_c) \leq \exp\left(-\frac{2(k\tilde{\phi}_k)^2}{2k}\right) = \exp(-k\tilde{\phi}_k^2).$$

The contribution from this region is thus exponentially small.

Region 2: $S_0 < s_c$ **(The "low-impact" zone).** In this region, we bound the impact term, $\mu(s_0+1) - \mu(s_0)$, by its largest value, which occurs at the boundary s_c:

$$\sum_{s_0=0}^{s_c-1} p(s_0)(\mu(s_0+1) - \mu(s_0)) \leq \mu(s_c)\sum_{s_0=0}^{s_c-1} p(s_0) \leq \mu(s_c).$$

We now bound $\mu(s_c)$. Substituting $s_c = k(1-\tilde{\phi}_k)$ into the posterior function yields:

$$\mu(s_c) = \frac{1}{1 + \left(\frac{1+2\tilde{\phi}_k}{1-2\tilde{\phi}_k}\right)^{2k\tilde{\phi}_k+1}} \leq \left(\frac{1-2\tilde{\phi}_k}{1+2\tilde{\phi}_k}\right)^{2k\tilde{\phi}_k+1}.$$

To analyze the asymptotic behavior, we take the logarithm of the bound. Using the Taylor expansion $\ln(\frac{1-u}{1+u}) \approx -4u$ for small u:

$$\ln(\text{bound}) \approx (2k\tilde{\phi}_k + 1) \cdot (-4\tilde{\phi}_k) \approx -8k\tilde{\phi}_k^2.$$

Thus, the upper bound $\mu(s_c)$ is approximately $\exp(-8k\tilde{\phi}_k^2)$, which is also exponentially small. The total expected marginal gain is bounded by a sum of two exponentially decaying terms. □

F.3. Conclusion

The preceding lemmas establish two key facts. First, the mathematical condition for achieving perfect learning remains $\lim_{k\to\infty} k\tilde{\phi}_k^2 = \infty$. Second, the expert's marginal incentive to contribute to the firm's accuracy decays exponentially, just as the pivotality term did in the main model. Because this term vanishes, the accuracy bonus remains the dominant incentive for truth-telling in large teams, as the rest of the truth-telling condition remains unchanged. The asymptotic analysis of the truth-telling condition is therefore structurally identical to that in the proof of Theorem 1. Consequently, the conditions for achieving perfect learning, namely $\rho > 1 - \beta/2$ when $B > 0$ and randomized response is used, are fully robust to the change to a quadratic loss function.

Appendix G: Experts with Dependent Signals

In this section, we extend the model of Section 3 to incorporate dependence among experts. Specifically, we replace the assumption that each expert receives a signal directly related to the state of the world Y with a more sophisticated information structure where expert signals are drawn from an intermediate binary random variable that introduces correlation.

G.1. Model with Dependent Experts

We maintain most of the setup from Section 3.1 but introduce the following modifications. The state of the world $Y \in \{0,1\}$ generates an intermediate binary variable $W \in \{0,1\}$ that represents a noisy measurement of Y. Specifically, W equals Y with probability $\mu \in (0.5, 1]$, i.e., $\Pr[W = Y] = \mu$. The parameter μ can be interpreted as the measurement accuracy, with $\mu = 1$ corresponding to perfect measurement.

Instead of observing signals directly related to Y, each expert i privately observes a noisy but informative signal $X_i \in \{0,1\}$ about the intermediate variable W. The quality of this signal is defined by the expert's *ability* α such that $\Pr[X_i = W] = \alpha$. As in the original model, we assume that $\alpha > 0.5$ and that α is common knowledge and consistent across all experts. While the experts' signals X_i remain conditionally independent given W, they are now conditionally dependent given Y via their common relationship with W. This creates a more realistic scenario where experts' assessments exhibit correlation—even experts with entirely independent information sources may reach similar conclusions because they are observing manifestations of the same underlying phenomenon. The rest of the model—including the expert voting process, the firm's randomized response mechanism, the firm's decision-making, and the cost functions—remains structurally the same as in Section 3.1.

The firm's optimal decision-making function remains unchanged—it is still the majority rule, as the signals remain exchangeable. The pivotality condition stays the same as well.

G.2. Expert's Truth-Telling Condition

We now derive the condition under which experts will vote truthfully under the extended model.

PROPOSITION 8 **(Expert's Truth-Telling Condition)**. *Under the extended model with dependent experts, experts vote truthfully (i.e., $\tilde{x}_i = x_i$ for every i) if and only if their privacy sensitivity ω is sufficiently low:*

$$\omega \leq \underbrace{\overbrace{\frac{2(2\mu-1)\tilde{\phi}_k}{\Delta v_{k,\alpha,\mu}(\tilde{\phi}_k)}}^{accuracy}}_{privacy} \cdot \left[\underbrace{\frac{1}{4^k} \cdot \binom{2k}{k} (1-4\tilde{\phi}_k^2)^k}_{pivotality} + \underbrace{B \cdot \frac{1-(1-\psi_k)^{2k+1}}{(2k+1)\psi_k}}_{bonus} \right], \tag{36}$$

where $\tilde{\phi}_k = \phi_k - 1/2$, and ψ_k is the probability that an expert's public vote matches the true state Y given that their signal is $x_i = 0$, equal to $1-\mu+(2\mu-1)\phi_k$.

Proof: Following a similar approach as in the original model, we need to compare the expected costs of voting truthfully versus untruthfully when the expert's signal is $x_i = 0$. Let $J_i(0;0)$ be the expected cost of voting truthfully and $J_i(1;0)$ be the expected cost of voting untruthfully.

For an expert with signal $x_i = 0$:

$$J_i(0;0) = \omega v_{k,\alpha,\mu,\nu}(0,0) + (1-\psi_k)\binom{2k}{k}\phi_k^k(1-\phi_k)^k + \Pr[\mathrm{Bin}(2k,\psi_k) < k] - \psi_k\gamma, \tag{37}$$

$$J_i(1;0) = \omega v_{k,\alpha,\mu,\nu}(1,0) + \psi_k\binom{2k}{k}\phi_k^k(1-\phi_k)^k + \Pr[\mathrm{Bin}(2k,\psi_k) < k] - (1-\psi_k)\gamma, \tag{38}$$

where $\gamma = B \cdot \frac{1-(1-\psi_k)^{2k+1}}{(2k+1)\psi_k}$.

The expert votes truthfully when $J_i(0;0) \leq J_i(1;0)$. After taking the difference and rearranging the terms, we obtain the truth-telling condition stated in the proposition. (Note that $2\psi_k - 1 = (2\mu-1)(2\phi_k-1)$, delivering the $2\mu-1$ factor in our new truth-telling condition.) $\square$

G.3. Randomized Response Convergence and Perfect Learning

We now explore the rate of randomized response convergence necessary for perfect learning under the extended model. Because of the dependency across signals, it is no longer possible to attain perfect learning, as the probability of the crowd being correct is now bounded from above by μ, the precision of the common signal. The expected cost is now written as:

$$(1-\mu) + (2\mu-1)\Pr(\mathrm{Bin}(2k+1,\phi_k) < k).$$

We can now define the *almost-perfect* learning regime, meaning that the expected cost converges to $1-\mu$ asymptotically. Since our new expected cost function is the linear transformation of our original expected cost function, the almost-perfect learning regime is attained under the same condition: $\lim_{k\to\infty} k\tilde{\phi}_k^2 = +\infty$. Also, since the truth-telling condition has the same structure as in our original setting, the almost-perfect learning regime is now attainable under the same conditions that the perfect learning regime was attainable in the original setting.

Appendix H: Unknown Outcomes

In the base model presented in Section 3, experts are rewarded based on whether their public votes match the eventual outcome Y. This assumes that the outcome is perfectly observed at the time of reward allocation. However, in many practical scenarios, the true outcome may not be immediately observable. For instance, the success of a product or policy might take years to fully manifest, while rewards need to be allocated much sooner. In this extension, we modify the model by making the outcome unknown at the time experts are rewarded.

H.1. Extended Model with Noisy Outcome Proxy

We extend the model by introducing a binary random variable $W \in \{0,1\}$ that serves as a noisy proxy for the true outcome Y. The proxy W is related to the true outcome through a parameter μ such that:

$$\Pr[Y = w | W = w] = \mu \tag{39}$$

for all $w \in \{0,1\}$. We assume that $\mu > \frac{1}{2}$, meaning that the proxy provides some information about the true outcome, but is not perfectly accurate. A higher value of μ indicates a more informative proxy signal.

The key modification to the model is that experts now receive accuracy rewards based on whether their public votes match the proxy W rather than the true outcome Y. However, the firm's ultimate utility still depends on the accuracy of the prediction of the actual outcome Y, and the experts' accuracy-based payoffs are still based on the actual outcome Y.

The expert's expected cost function becomes:

$$J_i(\tilde{x}_i; x_i) = \mathbb{E}_{Y,W,\boldsymbol{X}_p|x_i}\left[|Y-a| + \omega \cdot w_{k,\alpha,\nu}(X_{p,i}) - B \cdot \frac{\mathbb{I}(X_{p,i} = W)}{\sum_j \mathbb{I}(X_{p,j} = W)}\right] \tag{40}$$

Note that only the reward term has changed compared to Equation (1), with W replacing Y in the indicator function. The firm's cost function remains unchanged:

$$C(a; \nu, \boldsymbol{x}_p) = \mathbb{E}_{Y|\boldsymbol{x}_p}|Y - a| \tag{41}$$

Also, since the experts are still exchangeable, the firms optimal decision is still given by the majority rule.

H.2. Expert's Pivotality and Truth-telling Conditions

We now derive the expert's pivotality and truth-telling conditions under the modified reward structure.

Proposition 9 **(Expert's Truth-telling Condition with Noisy Proxy)**. *If the firm employs $2k+1$ experts and sets garbling level ν, then experts vote truthfully (i.e., $\tilde{x}_i = x_i$ for every i) if and only if their privacy sensitivity ω is sufficiently low:*

$$\omega \leq \frac{\overbrace{2(2\mu-1)\tilde{\phi}}^{accuracy}}{\underbrace{\Delta v_{k,\alpha}(\tilde{\phi})}_{privacy}} \cdot \left[\underbrace{\frac{1}{4^k} \cdot \binom{2k}{k}(1-4\tilde{\phi}^2)^k}_{pivotality} + \underbrace{B \cdot \frac{1-(1-\psi)^{2k+1}}{(2k+1)\psi}}_{bonus}\right] \tag{42}$$

where $\phi = \alpha - \frac{1}{2}(2\alpha - 1)\nu$, $\tilde{\phi} = \phi - 1/2$, and $\psi = \mu\phi + (1-\mu)(1-\phi)$ is the probability that any expert's public vote is correct.

Proof: Following the approach used in the base model, we compare the expected costs of voting truthfully versus untruthfully for an expert with signal $x_i = 0$. The expert votes truthfully if $J_i(0;0) \leq J_i(1;0)$.

For $J_i(0;0)$, the expected cost of voting truthfully:

$$J_i(0;0) = \omega v_{k,\alpha,\nu}(0,0) + (1-\psi)\binom{2k}{k}\phi^k(1-\phi)^k + \Pr[\mathrm{Bin}(2k,\phi) < k] - \psi\gamma \tag{43}$$

For $J_i(1;0)$, the expected cost of voting untruthfully:

$$J_i(1;0) = \omega v_{k,\alpha,\nu}(1,0) + \psi\binom{2k}{k}\phi^k(1-\phi)^k + \Pr[\mathrm{Bin}(2k,\phi) < k] - (1-\psi)\gamma \tag{44}$$

where $\gamma = \mathbb{E}_{N\sim\text{Bin}(2k,\psi)}[\frac{B}{1+N}]$ is the expected accuracy reward.

The difference is:

$$J_i(0;0) - J_i(1;0) = \omega[v_{k,\alpha,\nu}(0,0) - v_{k,\alpha,\nu}(1,0)] - (2\psi-1)\binom{2k}{k}\phi^k(1-\phi)^k - (2\psi-1)\gamma \tag{45}$$

By simplifying the expression and noting that $(2\psi-1) = (2\mu-1)(2\phi-1)$, we obtain the truth-telling condition in Equation (42). □

The key difference between this truth-telling condition and the one in the base model is the factor $(2\mu-1)$ in the truth telling condition. This factor represents the informativeness of the proxy W about the true outcome Y. When $\mu = 1$ (perfect proxy), we recover the original truth-telling condition. When $\mu = 0.5$ (uninformative proxy), the incentives for truth-telling disappear entirely.

H.3. Rate of Randomized Response Convergence and Perfect Learning

We now revisit the concept of perfect learning in the presence of a noisy proxy. Similarly to the model with dependency across signals, it is no longer possible to attain perfect learning, as the probability of the crowd being correct is now bounded from above by μ, the quality of the outcome proxy. The expected cost is now again written as:

$$(1-\mu) + (2\mu-1)\Pr(\text{Bin}(2k+1,\phi_k) < k).$$

We can now define the *almost-perfect* learning regime, meaning that the expected cost asymptotically converges to $1-\mu$. Since our new expected cost function is a linear transformation of our original expected cost function, the almost-perfect learning regime is attained under the same condition: $\lim_{k\to\infty} k\tilde{\phi}_k^2 = +\infty$. Also, since the truth-telling condition has the same structure as in our original setting, the almost-perfect learning regime is now attainable under the same conditions that the perfect learning regime was attainable in the original setting.

Appendix I: Private Vote Rewards

We now extend the model presented in Section 3 to a setting where expert rewards are based on the accuracy of private votes rather than public votes. This modification reflects scenarios where private votes are revealed after privacy concerns are no longer relevant, allowing for more accurate performance evaluation.

The basic structure of the model remains the same as in Section 3, with a firm and a set $\mathcal{I}$ of experts who observe noisy signals about a binary state of the world $Y \in \{0,1\}$. As before, each expert i observes a private signal X_i with ability α such that $\Pr[X_i = Y] = \alpha$.

The key modification is in how expert rewards are determined. While in the original model the rewards were based on whether the public votes $x_{p,i}$ matched the realized outcome, in our extended model, rewards are based on whether the private votes $\tilde{x}_i$ matched the realized outcome. This represents a scenario where, after the decision is made and the true state is revealed, the firm can observe or infer the experts' private votes for reward allocation purposes.

Formally, we modify the expert's expected cost function from Equation (1) as follows:

$$J_i(\tilde{x}_i; x_i) = \mathbb{E}_{Y,\boldsymbol{X}_p|x_i}\left[|Y-a| + \omega\cdot w_{k,\alpha,\nu}(X_{p,i}) - B\cdot\frac{\mathbb{I}(\tilde{x}_i = Y)}{\sum_j \mathbb{I}(\tilde{x}_j = Y)}\right], \tag{46}$$

The key difference is in the reward term, where $\mathbb{I}(X_{p,i} = Y)$ has been replaced with $\mathbb{I}(\tilde{x}_i = Y)$. This means that experts are rewarded for the accuracy of their private votes rather than their public votes. The firm's decision-making function remains unchanged from the original model:

$$C(a; \nu, \boldsymbol{x}_p) = \mathbb{E}_{Y|\boldsymbol{x}_p}|Y - a|, \tag{47}$$

Thus, the firm still bases its decision on the public votes, which means that the experts' incentives to influence the firm's decision remain tied to the public voting mechanism. Since the experts' public votes remain exchangeable, the firm's optimal aggregation function remains to be the majority rule.

We will now derive the expert's truth-telling condition under this modified reward structure. As in the original model, we need to compare the expected costs of voting truthfully versus untruthfully. First, we calculate the expected cost when an expert with signal $x_i = 0$ votes truthfully ($\tilde{x}_i = 0$):

$$J_i(0;0) = \omega v_{k,\alpha,\nu}(0,0) + (1-\phi)\binom{2k}{k}\phi^k(1-\phi)^k + \Pr[\text{Bin}(2k,\phi) < k] - \alpha\gamma, \tag{48}$$

where $\gamma = \mathbb{E}_{N\sim\text{Bin}(2k,\alpha)}[\frac{B}{1+N}] = B \cdot \frac{1-(1-\alpha)^{2k+1}}{(2k+1)\alpha}$ represents the expected reward based on the number of other experts whose private votes correctly match the state (which happens with probability α when experts vote truthfully).

Similarly, the expected cost when the expert votes untruthfully ($\tilde{x}_i = 1$) given signal $x_i = 0$ is:

$$J_i(1;0) = \omega v_{k,\alpha,\nu}(1,0) + \phi\binom{2k}{k}\phi^k(1-\phi)^k + \Pr[\text{Bin}(2k,\phi) < k] - (1-\alpha)\gamma. \tag{49}$$

The expert votes truthfully when $J_i(0;0) \leq J_i(1;0)$, which leads to the following proposition:

Proposition 10 **(Truth-Telling Condition with Private Vote Rewards)**. *If the firm employs $2k+1$ experts and sets garbling level ν, then experts vote truthfully (i.e., $\tilde{x}_i = x_i$ for every i) if and only if their privacy sensitivity ω is sufficiently low:*

$$\omega \leq \underbrace{\frac{1}{\Delta v_{k,\alpha}(\tilde{\phi})}}_{privacy} \cdot \left[\underbrace{\frac{2\tilde{\phi}}{4^k}\cdot\binom{2k}{k}(1-4\tilde{\phi}^2)^k}_{pivotality} + \underbrace{B\cdot\frac{(2\alpha-1)(1-(1-\alpha)^{2k+1})}{\alpha(2k+1)}}_{bonus}\right], \tag{50}$$

where $\tilde{\phi} = \phi - 1/2$ and $\phi = \alpha - \frac{1}{2}(2\alpha-1)\nu$ is the probability that any expert's public vote is correct.

Note the key difference in the truth-telling condition compared to the original model: the bonus term now depends directly on $(2\alpha - 1)$ rather than on how the public votes correlate with the state. This reflects that rewards are now based on private votes, which match the state with probability α when truthful, regardless of the garbling level ν.

The aggregation function is still the majority rule, and the condition for the perfect learning remains the same. However, the truth-telling condition is now slightly modified. As a result, the asymptotics on $\tilde{\phi}_k$ is now $k^{-\frac{1-\rho}{1+\beta}}$, resulting in a looser condition on perfect learning: $\rho > \frac{1-\beta}{2}$. Eventually rewarding the experts based on private votes expands the parameter region in which perfect learning is possible.

Appendix J: Finite-Sample Upper Bound on the Expected Cost

As a more refined result, we can obtain a finite-simple upper bound on the expected error in the imperfect learning regime when there is no performance-based reward. From the proof of the Lemma, we may obtain that the expected error is bounded from above by $\frac{1}{2}\exp(-k\tilde{\phi}^2)$. This suggests that it is enough to obtain a lower bound on k and $\tilde{\phi}$ to obtain the upper bound on the expected error. To do so, we will a) find a feasible crowd size k that supports truth-telling; and b) for this k, determine a lower bound on $\tilde{\phi}$ that still supports truth-telling.

To get analytically tractable results, let us fix the privacy cost function to be given by a Cobb-Douglas function $k^{-\rho}\tilde{\phi}^{1+\beta}$. The truth-telling condition is then given by:

$$\omega = k^{\rho}\tilde{\phi}^{-\beta}\frac{1}{4^k}\binom{2k}{k}(1-4\tilde{\phi}^2)^k.$$

We look at the binding condition as it should yield the optimal crowd size and optimal $\tilde{\phi}$—otherwise we can increase $\tilde{\phi}$ or k to improve our performance.

For any fixed k, we can make the truth-telling condition hold with equality by picking a relatively small $\tilde{\phi}$ as the right hand side tends to infinity as $\tilde{\phi}$ goes to 0. Let us obtain a lower bound on such $\tilde{\phi}$ by establishing a lower bound on the right hand side function. First, we know that $\frac{1}{4^k}\binom{2k}{k}$ is bounded from below by $\frac{1}{2}k^{-1/2}$ by a well-known combinatorial result. Second, we know that $\tilde{\phi} < \alpha - 1/2 = \tilde{\alpha}$. Combining the two bounds, we obtain a new equation, the solution to which would be a lower bound on the solution of the original equation:

$$\omega = \frac{1}{2}k^{\rho-1/2}\tilde{\phi}^{-\beta}(1-4\tilde{\alpha}^2)^k$$

By simplifying this equation, we can obtain the following family of lower bounds on $\tilde{\phi}$, indexed by k:

$$\tilde{\phi} \geq \underline{\phi} = \min\left[\left(\frac{k^{1/2-\rho}(1-4\tilde{\alpha}^2)^k}{2\omega}\right)^{\frac{1}{\beta}}, \tilde{\alpha}\right].$$

Appendix K: Experts of Heterogeneous Ability

So far, we assumed that all experts have the same ability—which in our setting corresponds to the accuracy of their binary signal being $\alpha = \Pr(X_i = Y)$. In practice, though, the experts may be heterogeneous in terms of their ability. In this case, the firm can potentially improve the accuracy of its decisions by (i) weighing the experts' votes according to their ability and (ii) assigning personalized garbling probabilities to each expert.

We represent this heterogeneity by assuming that all experts belong to one of the two sets: the low-ability set $\mathcal{I}_L$ of size n_L and the high-ability set $\mathcal{I}_H$ of size n_H. All experts in $\mathcal{I}_L$ have ability α_L and those in $\mathcal{I}_H$ have ability α_H, with $\alpha_L < \alpha_H$. The firm selects the garbling levels ν_L and ν_H for these sets of experts to optimize its expected mismatch cost. As in the preceding section, we denote the probabilities of correct votes for each group by ϕ_L and ϕ_H.[18] Given these probabilities and the vector $\boldsymbol{x}_p$ of each expert's public vote, the firm's optimal action function $a^*(\boldsymbol{x}_p)$ can be derived through a *weighted majority rule*:

$$a^*(\boldsymbol{x}_p) = \mathbb{I}\left[\sum_{i_1\in\mathcal{I}_L}\psi_L(1-2x_{p,i_1}) + \sum_{i_2\in\mathcal{I}_H}\psi_H(1-2x_{p,i_2}) > 0\right],$$

[18] For example, ϕ_L is given by $\frac{1}{2}(1+(2\alpha_L-1)(1-\nu_L))$.

where, the "0" and "1" votes are represented numerically as -1 and 1 respectively and also, the low-(high-)ability votes are assigned the weights $\psi_L = \ln \frac{\phi_L}{1-\phi_L}$ and $\psi_H = \ln \frac{\phi_H}{1-\phi_H}$. This optimal weighting scheme was independently derived in social choice (Nitzan and Paroush 1982) and computer science literature (Minsky and Papert 1969, Chapter 12.4); see also Berend and Kontorovich (2014) for a connection to boosting algorithms.

The voting profiles could be summarized as tuples (m_L, m_H) where m_L (and m_H) is the number of correct low-(high-)ability votes. The probability of any such voting profile is given by

$$p(m_L, m_H; \phi_L, \phi_H) = \binom{n_H}{m_H}\binom{n_L}{m_L}\phi_H^{m_H}(1-\phi_H^{m_H})^{n_H-m_H}\phi_H^{m_L}(1-\phi_L^{m_L})^{n_L-m_L}.$$

To formalize the firm's expected error, we will introduce binary variables $z(m_L, m_H)$ constrained to be equal to 1 if the firm makes an incorrect decision given the voting profile (m_L, m_H) and to 0 otherwise. For each profile (m_L, m_H), these constraints can be encoded using a big-M formulation (Bradley et al. 1977, Chapter 9):

$$\psi_L(2m_L - n_L) + \psi_H(2m_H - n_H) \geq \underline{M} z[m_L, m_H] + \varepsilon,$$

$$\psi_L(2m_L - n_L) + \psi_H(2m_H - n_H) \leq \bar{M} z[m_L, m_H] - \varepsilon,$$

where $\underline{M}$ is a negative number with a large absolute value, $\bar{M}$ is a large positive number, and ε is a small positive number. The firm's expected cost function can then be written as

$$\sum_{0\leq m_L\leq n_L}\sum_{0\leq m_H\leq n_H} p(m_L, m_H; \phi_L, \phi_H) z(m_L, m_H).$$

Now, we will derive the constraints for truthful voting for the low-ability experts (the high-ability case can be derived analogously). We will omit bonus for accuracy for simplicity. We will also introduce the variable $z_L(m_L, m_H)$, which encodes that a low-ability expert is pivotal given that the rest of the votes are described by the profile (m_L, m_H) (see Section 4.1 for the definition of expert's pivotality). Then, the pivotality constraints can be written using a big-M formulation:

$$\psi_L(2m_L - n_L + 1) + \psi_H(2m_H - n_H) \geq \underline{M} \cdot (1 - z_L(m_L, m_H)),$$

$$\psi_L(2m_L - n_L + 1) + \psi_H(2m_H - n_H) \leq \bar{M} \cdot (1 - z_L(m_L, m_H)),$$

In these expressions, $0 \leq m_L \leq n_L - 1$ (because we're focusing on one of the low-ability experts) and $0 \leq m_H \leq n_H$. Given the variables $z_L(\cdot,\cdot)$, we can then write down the truthful voting (alt., incentive compatibility) constraint as:

$$\omega \leq \frac{2\phi_L - 1}{\Delta v_{k,\alpha}(\phi_L - 1/2)} \cdot \left(\sum_{m_L=0}^{n_L-1}\sum_{m_H=0}^{n_H} p(m_L, m_H; \phi_L, \phi_H) \cdot \left(\frac{n_L - (m_L - 1)}{n_L(1-\phi_L)}\right) \cdot z_L(m_L, m_H)\right).$$

The resulting problem is a complicated mixed-integer nonlinear programming that is unlikely to have an analytically tractable solution. Therefore, we investigate this problem numerically. In particular, we consider two scenarios: (i) no personalized garbling, implying that an additional constraint $\nu_L = \nu_H$ is imposed, and (ii) full anonymization, meaning that $\omega = 0$ (since there are no privacy risks as experts are anonymous) but $\psi_L = \psi_H$ (because there is no way to weigh the experts heterogeneously if they are anonymous).

Notice that if the garbling is personalized (i.e., if we can assign different garbling levels to different-ability experts), then increasing the ability of any given expert cannot reduce the expected error. Indeed, we can increase the garbling level for said expert so that the correctness probability—and therefore the truth-telling constraints—stay the same, meaning that we're guaranteed at least the same payoff. However, if the garbling level is *not* personalized, the firm may encounter situations in which improving the ability of one expert reduces the overall performance of the team. For example, consider the case with $\rho = 0$, $\beta = 1$, and $\omega = 0.1$. Suppose that initially there are 5 experts, all with ability $\alpha = 0.6$. Now assume that the ability of one of these experts increases to 0.8. It turns out that after this change, the expected error increases from 0.36 to 0.42. This happens because now the firm has to combat the low-ability experts' free-riding by providing better privacy protection and, therefore, increased garbling.

Our extended model also provides an additional argument *against* anonymization. In particular, we find numerically by testing a wide range of sets of model parameters that under certain conditions, anonymization may be detrimental to the firm's performance (as measured by the expected cost that the firm incurs). This may be the case because with vote anonymization, the firm is not able to assign a higher weight to the team's most capable expert and thus, the firm loses accuracy. For an illustration, consider the following setting: $\rho = 0$, $\beta = 1$, $\omega = 0.1$, $n_L = 4$, $n_L = 1$, $\alpha_L = 0.6$ and $\alpha_H = 0.8$ (assume that the garbling is now personalized). We can show numerically that with anonymization, the expected error is equal to 0.24 while without anonymization it is just 0.19, which signifies a better performance of the firm.

Appendix L: Experts Exerting Costly Effort to Increase Ability

The platform aims at incentivizing experts to exert high effort, which improves their ability from α_L to $\alpha_H > \alpha_L$ at a personal cost F. If an expert exerts such costly effort with probability ψ their expected ability is $\bar{\alpha} = \psi\alpha_H + (1-\psi)\alpha_L$. Denote $\phi = \frac{1}{2}(1 + (2\bar{\alpha} - 1)(1-\nu))$, $\phi_L = \frac{1}{2}(1 + (2\alpha_L - 1)(1-\nu))$, and $\phi_H = \frac{1}{2}(1 + (2\alpha_H - 1)(1-\nu))$.

The firm optimizes over the common garbling level ν. The firm's optimization problem can be formulated as a mixed integer nonlinear program with the objective of minimizing the expected cost.

$$\sum_{l=0}^{k} \binom{2k+1}{l} \phi^l (1-\phi)^{2k+1-l}.$$

We encode the incentive compatibility on the effort choice in the following constraints:

$$\psi \left[(\phi_H - \phi_L) \left(\binom{2k}{k} \phi^k (1-\phi)^k + \gamma \right) - F \right] \geq 0,$$

$$(1-\psi) \left[(\phi_H - \phi_L) \left(\binom{2k}{k} \phi^k (1-\phi)^k + \gamma \right) - F \right] \leq 0.$$

Incentive compatibility on the truthful voting choice is encoded as follows:

$$\omega \leq \frac{2\phi - 1}{\Delta v_{k,\alpha}(\phi - 1/2)} \left(\binom{2k}{k} \phi^k (1-\phi)^k + \gamma \right),$$

where, γ is defined as in our original setting: $\gamma = \mathbb{E}_{N\sim\mathrm{Bin}(2k,\phi)}[\frac{B}{1+N}] = B \cdot \frac{1-(1-\phi)^{2k+1}}{(2k+1)\phi}$. Similarly to the previous section, we also use this model to test numerically on a number of sets of model parameters that there exist scenarios under which a randomized response combined with the accuracy bonuses may lead to a *lower* firm's expected cost when compared to a fully anonymized case (in which $\omega = 0$ as there are no privacy concerns and also $\gamma = 0$ as there is no way for the firm to award personalized bonuses when identities of the experts are anonymized). For illustration, consider the following set of model parameters: $\rho = 0, \beta = 1, k = 5, \alpha_L = 0.6, \alpha_H = 0.7$ and $F = 0.02$. In the fully anonymized case, $\omega = 0$, as there are no privacy concerns. At the same time, there is no way for the firm to reward the accuracy of the experts, and hence $B = 0$ in this case too. Without anonymization, we assume that $\omega = B = 0.01$. We obtain that in the fully anonymized case, the expected error is approximately 0.245 versus the expected error without anonymization: 0.24. The observed better performance of the firm in the non-anonymized case can be explained by the increased effort of the experts: $\psi = 0.03$ as opposed to 0.007 with full anonymization.