Incentive Engineering for Outsourced Computation in the Face of Collusion

Arman Khouzani, **Viet Pham**, *and Carlos Cid*

Information Security Group
Royal Holloway University of London

{arman.khouzani, viet.pham.2010, carlos.cid}@rhul.ac.uk

ITA-2014

Royal Holloway University of London

Information Security Group

Outsourced Computing

Research Problem

- Outsourcing of computational tasks:
  - Cryptographic solutions exit but can be an overkill if the parties are **not malicious but just lazy**: they may return *guessed* results just to save computational cost (and/or gain more reward given their capacity)
- Our Aim:
  - Designing optimal incentive schemes by the outsourcer (*principal*) combining *audits, redundancy, rewards, punishments and bounties* that guarantee participation and honest computation of the contractors (*agents*)
- Challenges:
  - Limited budget (for rewards and bounties), limited capacity for auditing, costly auditing, bounded enforceable fine, risk of "collusion" among participants

Royal Holloway University of London

Information Security Group

Problem Modeling

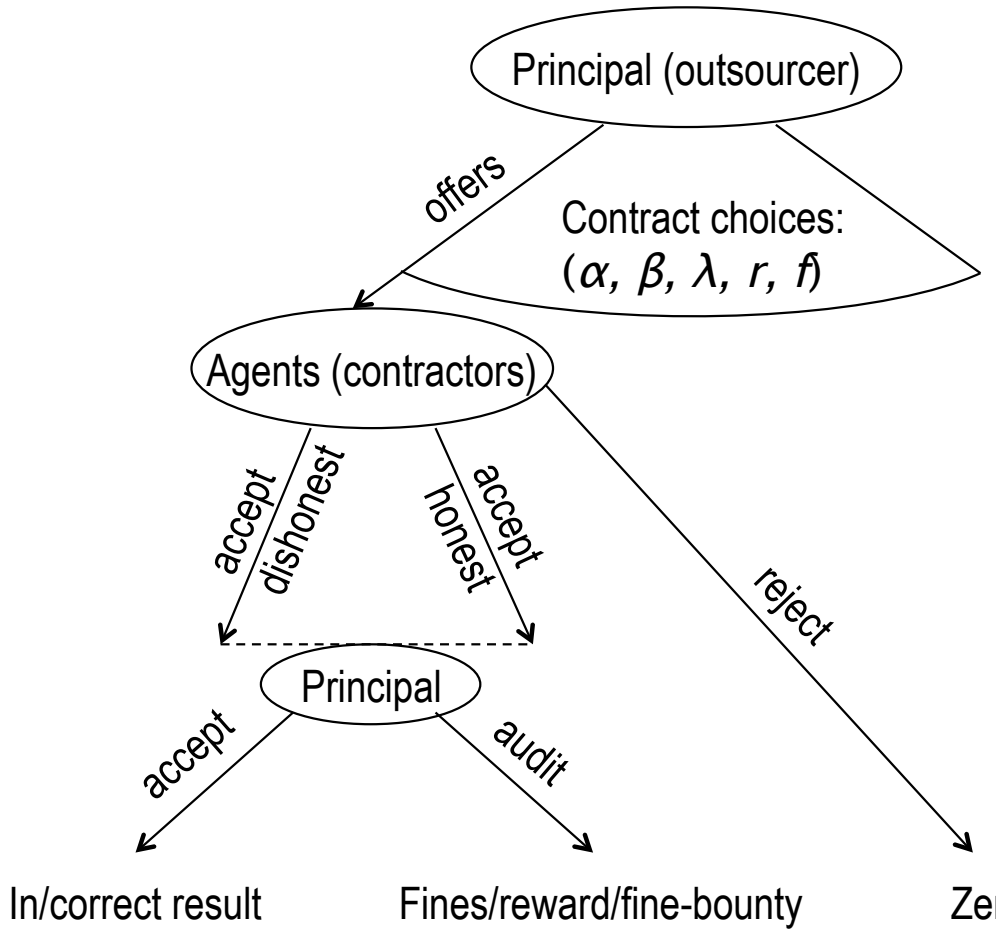

α	Prob. Redundancy
β	Ex-post prob. auditing two conflicted results ($\leq \Lambda$)
λ	Ex-ante prob. auditing single result ($\leq \Lambda$)
r	Reward ($\leq R$)
f	Fine ($\leq F$)

Summary of results

Previous work:

- Optimal contracts for single agent.
- Optimal contracts for one/two agents, given no collusion.

This work:

- Optimal contracts under information leakage.
- Optimal contracts under collusion.

Optimal Contract for a Single Agent

- The principal chooses the contract (auditing rate, reward and punishment) to maximize its utility ensuring *fully honest* computation.

$$\min_{r,f,\lambda} \mathcal{C} := r + \gamma\lambda$$

- Requiring full honesty translates to ensuring: $1 = \arg\max u_A(q)$. Following the Principal-Agent modeling in game theory, we will refer to this as the *incentive compatibility* constraint:

$$u_A(1) = r - c(1) \geq u_A(q_1) = [1 - (1 - q_1)\lambda]r - c(q_1) - (1 - q_1)\lambda f.$$

- The agent accepts the contract if its expected utility is larger than its *reserve utility*, $z \geq 0$. Hence, given incentive compatibility, this *participation constraint* is: $u_A(1) = r - c(1) \geq z.$
- This is a non-convex optimization, but satisfies (MFCQ), hence KKT.

Optimal Contract for a Single Agent

Proposition

The contract that enforces honest computation and is accepted by the agent, and minimizes the cost of the principal is by setting $f^* = F$ and choosing λ^*, r^* as given by the following:

$$\gamma \leq \frac{c}{\Lambda^2}: \begin{cases} [\frac{c}{\Lambda} - c]^+ \leq F: & \lambda^* = \frac{c}{c+F},\ r^* = c,\ \mathcal{C}^* = c + \frac{\gamma c}{c+F} \\ [\frac{c}{\Lambda} - R]^+ \leq F < [\frac{c}{\Lambda} - c]^+: & \lambda^* = \Lambda,\ r^* = \frac{c}{\Lambda} - F,\ \mathcal{C}^* = \frac{c}{\Lambda} + \gamma\Lambda - F \end{cases}$$

$$\gamma > \frac{c}{\Lambda^2}: \begin{cases} [\sqrt{c\gamma} - c]^+ \leq F: & \lambda^* = \frac{c}{c+F},\ r^* = c,\ \mathcal{C}^* = c + \frac{\gamma c}{c+F} \\ [\sqrt{c\gamma} - R]^+ \leq F < [\sqrt{c\gamma} - c]^+: & \lambda^* = \sqrt{\frac{c}{\gamma}}, r^* = \sqrt{c\gamma} - F, \mathcal{C}^* = 2\sqrt{c\gamma} - F \\ [\frac{c}{\Lambda} - R]^+ \leq F < [\sqrt{c\gamma} - R]^+: & \lambda^* = \frac{c}{R+F},\ r^* = R,\ \mathcal{C}^* = R + \frac{\gamma c}{R+F} \end{cases}$$

For $F < [\frac{c}{\Lambda} - R]^+$, the optimization is infeasible, i.e., there is no honesty-enforcing contract that is also accepted by the agent.

Proposition

Our optimal contracts stay feasible for any risk-averse agent as well.

Optimal Contract for a Single Agent

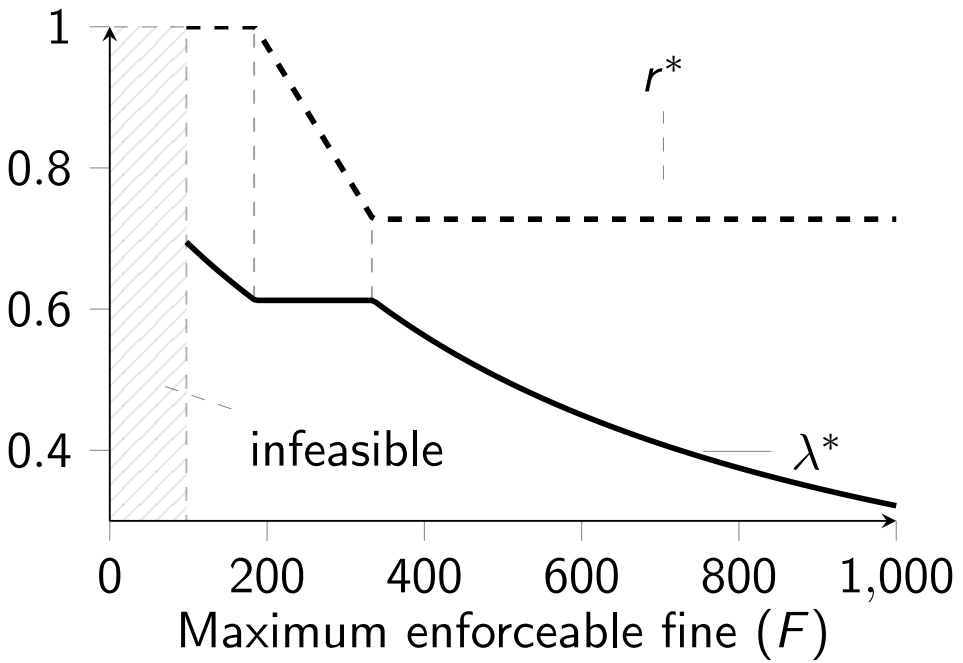

Figure: Example illustration of contract parameters r^*, λ^* w.r.t. the maximum enforceable fine F. Note that both r^* and λ^* are decreasing over F, however, r^* never falls below cost of honest computation c.

Optimal Contract for Two-Agent: **Baseline**

- A principal can use a **hybrid scheme** of sending the same job to multiple agents comparing the returned results (redundancy scheme), and to only one randomly selected agent and probabilistically audit it.
- Let $u_A(a_1, a_2)$: utility of agent 1, where $a_1, a_2 \in \{Honest, Cheat\}$.

$$u_A(H, H) = r - c, \quad u_A(C, H) = (1 - \alpha - \lambda)r/2 - (\alpha + \lambda/2)f.$$

- Principal's expected cost: $\mathcal{C} = 2r\alpha + \gamma\lambda + r(1 - \alpha) = (1 + \alpha)r + \gamma\lambda$.

$$\min_{r,f,\alpha,\lambda} r(1 + \alpha) + \gamma\lambda \text{ subject to:}$$

$$r \leq R,\ f \leq F,\ 0 \leq \lambda \leq \Lambda,\ \lambda \leq 1 - \alpha,\ \alpha \geq 0,\ \boxed{r \geq c},\ \boxed{r \geq \frac{c(1 + \alpha)}{\lambda + 2\alpha} - f}.$$

Royal Holloway University of London

Information Security Group

Optimal Contract for Two-Agent: **Baseline**

Proposition

Let $F_0 = c/\Lambda - c$ *and* $F_1 = c[c-\gamma]^+/[2\gamma - c]^+$.[a] *The optimal two-agent contract that guarantees participation and* (H, H) *as a Nash equilibrium is:*

$$\begin{cases} F_1 \leq F: & f^* = F, \alpha^* = \dfrac{c}{2F+c}, \lambda^* = 0, r^* = c,\ \mathcal{C}^* = c(1 + \dfrac{c}{2F+c}) \\ F_0 \leq F < F_1: & f^* = F, \alpha^* = 0, \lambda^* = \dfrac{c}{c+F}, r^* = c,\ \mathcal{C}^* = c(1 + \dfrac{\gamma}{F+c}) \\ F < \min(F_0, F_1): & f^* = F, \alpha^* = \dfrac{c - \Lambda(c+F)}{c+2F}, \lambda^* = \Lambda, r^* = c, \mathcal{C}^* = \dfrac{c(c+F)(2-\Lambda)}{c+2F} + \gamma\Lambda \end{cases}$$

For $\Lambda = 1$, (H, H) *is moreover the dominant Nash equilibrium.*

[a] We adopt the convention that $x/0 = +\infty$ for $x > 0$.

Optimal Contract for Two-Agent: Information Leakage

- Principal relied on agents' oblivion about when redundancy is used.
- Agents may be able to find out about task assignment of each other through a *side-channel* (hence the name information leakage). This lets them to *selectively* be honest.
- Hence, contract constraints must deal with two *information states*:
  - *Lone recipient*: $r - c \geq r(1-\rho) - f\rho$
  - *Redundancy*: $r - c \geq -f$

$$\min_{r,f,\alpha,\lambda} \mathcal{C} := r(1+\alpha) + \gamma\lambda \text{ subject to:}$$

$$f \leq F,\ 0 \leq \lambda \leq \Lambda,\ \lambda \leq 1-\alpha,\ \alpha \geq 0,\ \boxed{r \geq c},\ \boxed{r\lambda + f\lambda \geq c(1-\alpha)}.$$

Optimal Contract for Two-Agent: Information Leakage

Proposition

The optimal two-agent contract with information leakage, i.e., where the agents have access to the information of whether the same task is outsourced to the other agent or not, enforces honesty in that makes (H, H) a Nash equilibrium sets $f^* = F$, $r^* = c$, and:

$$\gamma \geq \frac{c}{\Lambda}: \begin{cases} F \geq [\gamma - c]^+ : & \lambda^* = \dfrac{c}{c+F},\ \alpha^* = 0,\ \mathcal{C}^* = c + \dfrac{\gamma c}{c+F} \\ F < [\gamma - c]^+ : & \lambda^* = 0,\ \alpha^* = 1,\ \mathcal{C}^* = 2c \end{cases}$$

$$\gamma < \frac{c}{\Lambda}: \begin{cases} F \geq [c/\Lambda - c]^+ : & \lambda^* = \dfrac{c}{c+F},\ \alpha^* = 0,\ \mathcal{C}^* = c + \dfrac{\gamma c}{c+F} \\ [\gamma - c]^+ \leq F < [c/\Lambda - c]^+ : & \lambda^* = \Lambda,\ \alpha^* = 1 - \Lambda\left(1 + \dfrac{F}{c}\right),\ \mathcal{C}^* = c\left(2 - \Lambda\left(1 + \dfrac{F}{c}\right)\right) + \gamma\Lambda \\ F < [\gamma - c]^+ : & \lambda^* = 0,\ \alpha^* = 1,\ \mathcal{C}^* = 2c \end{cases}$$

Royal Holloway University of London

Information Security Group

Optimal Contract for Two-Agent: Collusion

- The two agents may be able to coordinate their responses to report the same guessed result, saving computation cost without detection
- One way to discourage collusion: the returned results from the two agents can be audited by the principal with probability ν, (even) when they are the same.
- Incentive compatibility constraint: collusion should be a less attractive equilibrium, i.e., ensuring: $u_A(C, C) < u_A(H, H)$.
- With the introduction of ν, we have: $u_A(C, C) = r(1 - \nu) - F\nu$. Therefore, to make honesty a more attractive equilibrium than collusion, in the redundancy scheme information state, we must have:

$$r - c \geq r(1 - \nu) - F\nu$$

Royal Holloway University of London

Information Security Group

Optimal Contract for Two-Agent: Collusion

Proposition

The optimal contract that makes collusion a less attractive equilibrium than honest computation never uses the redundancy scheme at all.

- Intuitively, the principal can save the reward to the second agent by assigning the task to only one of them.
- We introduced bounty schemes, creating a prisoner's dilemma-like situation to undermine collusion: Make collusion a dis-equilibrium, i.e. $u_A(H, C) > u_A(C, C)$ – rather than a less desired equilibrium.
- When the returned results are different, the principal can randomly audit the task and reward the agent with the correct result (if any) with the "bounty" at largest credible promise, i.e., R.

Optimal Contract for Two-Agent: Collusion

- Let β be the probability of auditing by the principal if the task is assigned to two agents and the returned results are different
- Bounty Scheme One, Two and Hybrid: The difference between the schemes is how they treat the cases when the returned results are different AND not audited:
  - in *bounty scheme one*, both agents are punished at f;
  - in *bounty scheme two*, both agents are rewarded at r;
  - in the *hybrid bounty scheme*, the amount "paid" to the agents by the principal in such cases is x, a optimization variable with $-F \leq x \leq R/2$.

Optimal Contract for Two-Agent: Collusion

We derive partial closed-form solutions, establishing even in the presence of collusion, redundancy plus bounty schemes may still be optimal:

Corollary

For $F < [\gamma - c]^+$, in bounty scheme one if $\Lambda \geq 2c/R$, and in bounty scheme two if $\Lambda \geq c/\min(c+F, R-c)$, the optimal contract chooses redundancy $\alpha^* = 1$. The rest of the parameters for such a case are: $r^* = c$, $\lambda^* = \nu^* = 0$, $f^* = F$.

Corollary

For $F < [\gamma - c]^+$, if $\Lambda \geq \max\{2c/(R+F), (4c-R)/R\}$, the optimal hybrid bounty scheme contract chooses redundancy $\alpha^* = 1$. The rest of the parameters are: $r^* = c$, $\lambda^* = \nu^* = 0$, $f^* = F$ and $x^* = \min\{2cF/(R+F-2c), R/2\}$.

Optimal Contract for Two-Agent: Collusion

Example of optimal cost when $\gamma > c$:

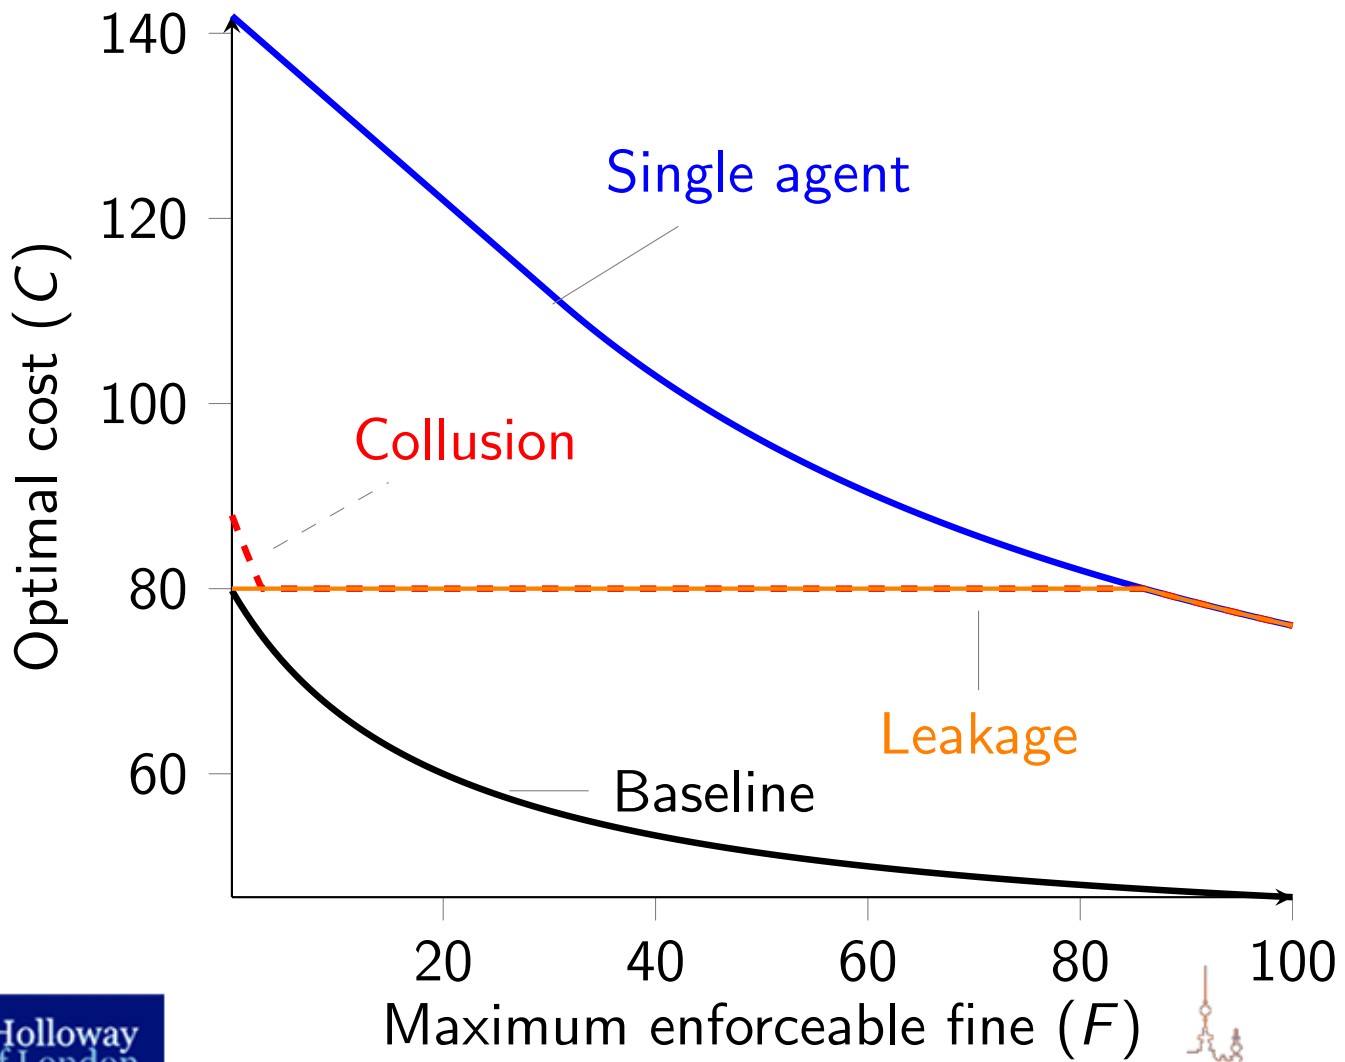

Royal Holloway University of London

Information Security Group

Generalizations

- **interactions among the agents**: the agents may be able to deceive their peers by giving them wrong signals about their state with the objective of winning the bounty.
- **enforceable commitments among colluding agents**: Assuming enforceable commitment, agents may agree to pass the honest result to one another, or intentionally plan for one of them to get the bounty, only to share it among themselves later.
- **global optimality of two-agent contracts**: In our previous work, we established that when agents are non-colluding and non-communicating, the optimal contracts developed assuming at most two agents per each task are in fact globally optimal among all contracts involving any number of agents per task. In the presence of information leakage and collusion, this is open.

Royal Holloway University of London

Information Security Group