

Managing Information Production in Teams*

Shota Ichihashi[†] Fei Li[‡] Dihan Zou[§]

May 1, 2024

Abstract

A principal has a stream of related decisions to make under imperfect information. He employs a finite group of agents to acquire information at a cost. The principal designs task allocation and payment schemes to robustly implement all agents engaging in information acquisition and truthful reporting. We characterize the optimal joint design of task allocation and payment scheme, which highlights a trade-off between task assignment diversification and peer monitoring efficiency. The optimal deterministic design features a chain structure of peer monitoring. Stochastic task allocation and payment scheme ease the tension between diversification and monitoring efficiency.

Keywords: team incentives, monitoring, information acquisition, full implementation

JEL Classification Codes: D82, D83, D86, L23, L25

*We are grateful to Yu Awaya, Arjada Bardhi, Gary Biglaiser, Nina Bobkova, Jaden Chen, Joyee Deb, Rahul Deb, Miaomiao Dong, Wei He, Ashwin Kambhampati, Jacob Kohlhepp, Maciej H. Kotowski, Margaret Meyer, Andrew Newman, Ichiro Obara, Marcin Peški, Peter Norman, Anna Sanktjohanser, Gan Tan, Can Tian, Huseyin Yildirim, Kun Zhang, and audiences at numerous conferences and seminars for their comments.

[†]Queen's University. s.ichihashi@queensu.ca

[‡]University of North Carolina at Chapel Hill. lifei@email.unc.edu

[§]University of North Carolina at Chapel Hill. dihan@email.unc.edu

Contents

- 1 Introduction** **1**

- 2 Model** **6**

- 3 Optimal Policy and Information Loss** **11**

- 4 Optimal Task Allocation and Contracting** **13**
 - 4.1 Least-Cost Compensation and Incentive Cost 13
 - 4.2 Optimal (Deterministic) Task Allocation 20
 - 4.3 Discussion 23

- 5 Random Assignment and Transparency** **25**

- 6 Conclusion** **29**

- A Appendix: Proofs** **29**

- B Appendix: Partial-Implementation Contract** **36**

- C Supplementary Material** **37**
 - C.1 Heterogeneous Costs of Information Acquisition 37
 - C.2 Circular Decision Environment 42
 - C.3 Imperfect Information Acquisition 44

1 Introduction

Modern organizations face numerous decision problems in their daily operations. Some problems are similar to each other and thus admit similar solutions, whereas other problems are more novel and require distinct approaches. Recognizing the heterogeneous similarity among decision problems is crucial for organizations to efficiently allocate and motivate their limited personnel resources for gathering decision-relevant information. Nevertheless, to the best of our knowledge, this aspect remains largely unexplored in the literature. To fill the gap, this paper offers a model of task allocation and contracting in team information production. The principal faces many decision problems. The solution to each problem is ex ante uncertain and correlated. The principal can task each agent to acquire information about one of the problems and discover its solution. However, there are more problems than agents, and the effort each agent puts into acquiring information is unobservable. The principal’s design problem consists of sampling and contracting—i.e., which problems to assign to agents and how to motivate them to produce high-quality information. We give some examples of this setting below.

Market research. A business owner wants to decide whether to launch an advertising campaign in each market. The returns to advertising are ex ante uncertain and correlated, that is, similar markets are more likely to have similar returns. The owner hires consultants, each of whom can investigate one market. The owner then needs to decide which markets to investigate and how to incentivize consultants to produce high-quality information.

Project evaluation. A granting agency wants to assess a project’s merits across several attributes. The agency can consult a limited number of reviewers, each possessing unique expertise in specific attributes. The agency needs to decide how to select reviewers (or attributes to focus on), incentivize their participation through monetary or non-monetary rewards, and consolidate their feedback.

Federated learning. A leading hospital wants to train an AI-based image-processing model for tumor diagnosis. The training requires a large amount of diverse patient data from hospitals nationwide. However, due to privacy concerns, the data must be processed within the hospital where the data originated.¹ The leading hospital needs to select participating hospitals,

¹E.g., *German Cancer Consortium’s Joint Imaging Platform* consists of 24 research institutes and hospitals across Germany. It aims at facilitating medical image processing with machine learning methods, while dealing with challenges of data protection requirements at participating sites. See [Rieke et al. \(2020\)](#) for more examples.

incentivize them for proper model training, and aggregate the training results.

We study an optimal contracting problem to study such issues. The principal needs to make decisions on a continuum of locations, each of which is indexed by $t \in [0, 1]$. Each location t has a binary uncertain state, or its local state, $x(t) \in \{A, B\}$. The local states are initially uncertain and correlated: If two locations are closer, the local states are more likely to take the same value. To capture such a correlation structure, we assume that the state $(x(t))_{t \in [0, 1]}$ is a realized path of a continuous time Markov chain.

To learn about the state, the principal hires n agents. The principal chooses a task allocation—which specifies the location each agent is assigned to—and a compensation scheme—which rewards agents based on their reports and the principal’s own information. Specifically, given a contract, each agent i privately chooses whether to work or shirk, where working is costly but enables agent i to privately observe the local state of her assigned location, $x(t_i) \in \{A, B\}$. Each agent then sends a report A or B to the principal. The principal also observes the realized local state $x(0)$ of location 0 at no cost, which captures his pre-existing knowledge. Based on the reports of the agents and his own signal, the principal pays agents according to the compensation scheme.

Based on the collected information, the principal chooses action A or B at each location. The principal incurs a loss whenever the action does not match the corresponding local state. We call the principal’s loss from this final decision stage an information loss.

The principal chooses a task allocation and a compensation scheme in order to minimize the sum of the information loss and payments, while *robustly* implementing work and truth-telling. That is, the principal designs a contract so that all agents work and report truthfully in a unique (correlated) rationalizable outcome, or equivalently, as a unique outcome that survives the iterative elimination of strictly dominated strategies.

We first characterize the task allocation that minimizes the information loss without taking into account payments to the agents. Such a task allocation maximally diversifies the agents’ locations. The principal learns about the entire set of local states well, including locations that are further away from the principal’s location (i.e., $t = 0$) and exhibit high ex ante uncertainty.

We then characterize the robustly optimal contract and show that it has two features. First, the optimal compensation scheme takes a *chain monitoring structure*, i.e., it pays each agent a bonus if and only if her report coincides with the report of her left neighbor. Intuitively, the agent closest to the principal, say agent 1, will be paid whenever her report coincides with the

principal’s signal, regardless of other agents’ reports. Thus agent 1 will be the first player who chooses “work” as a dominant strategy. Conditional on that agent 1 works, the right neighbor of agent 1, say agent 2, finds it dominant to work. This logic continues and ensures that all agents work and truthfully report their findings as a unique outcome. While other compensation schemes could also robustly implement the desired outcome, the chain monitoring structure does so in the cheapest way.

The robust compensation scheme contrasts with the optimal contract under partial implementation, which induces all agents to work in *an* equilibrium. Consistent with the informativeness principle, such a contract rewards each agent when her report coincides with the reports of her left and right neighbors (cf. [Holmström \(1982\)](#)). However, the contract also creates another equilibrium in which everyone shirks. The robustly optimal contract eliminates such strategic risk by underusing information for compensating agents.

The second feature of the robustly optimal contract is that, compared to the task allocation that minimizes information loss, the principal assigns agents closer to each other and closer to the principal. Such a task allocation renders monitoring more effective: When agents are assigned to closer locations, the report of each agent i and that of her left neighbor will be more strongly correlated. Then if agent i shirks, she would face a higher probability of sending an inconsistent report (and thus losing a bonus) than when her left neighbor is further away. As a result, the principal can induce effort at a relatively low bonus. However, the principal incurs large information losses at locations that are further away from the principal location, $t = 0$.

Finally, the principal benefits from keeping task allocation opaque. Specifically, suppose that the principal can commit to randomize task allocations and privately reveal a realized location to each agent. In such a case, the principal can virtually eliminate the trade-off between maximizing learning and minimizing agency cost. We construct a distribution over task allocations such that the first-best task allocation is realized with probability close to 1, but with small probabilities, one of the agents is secretly appointed to monitor another agent. The principal compensates each agent only when the task allocation differs from the first best, but each agent, not knowing whether she is monitored by another agent, will exert effort regardless of the realized task allocation. Such a contract attains the first-best information loss without sacrificing the effectiveness of monitoring. The result highlights the benefit for the principal of keeping task allocation uncertain and opaque from the agents’ perspectives.

Contributions. Our paper makes four contributions. First, we introduce a tractable model

for analyzing task allocation in team information production. The model parsimoniously captures the rich set of heterogeneous, interdependent tasks. Second, we study a novel trade-off between information diversification and peer monitoring efficiency. In addition to the standard underprovision of effort (which we do not study in this paper), the trade-off creates another distortion in motivating team information acquisition—the misallocation of tasks. Third, we highlight the importance of leveraging the principal’s existing knowledge in motivating agents to acquire information. For enhanced monitoring efficiency, the principal should assign agents to those closely aligned with the areas in which the principal holds information, relative to the efficient allocation. Finally, we propose a new rationale for maintaining confidentiality regarding task assignment and compensation.

Related Literature. The paper contributes to several strands of literature. First, we offer a new model to study task allocation in team production. Since the seminal work by [Holmstrom and Milgrom \(1991\)](#), there has been intensive discussion on allocating multiple tasks to a group of agents in team production settings. We recommend [Bolton and Dewatripont \(2004\)](#) for a comprehensive textbook treatment.² While the majority of the literature is not about information acquisition, [Bohren and Kravitz \(2019\)](#) consider a principal who employs agents to learn the realization of many independent draws. The principal can (i) assign multiple agents to the same task for peer monitoring and (ii) assign multiple tasks to the same agent and takes away the rewards from successful tasks if one failure is detected by her peer monitor. In their model, tasks are independent, so the trade-off between diversification and monitoring efficiency does not arise.

Second, we introduce task-design to the literature of information acquisition by multiple agents. To the best of our knowledge, this perspective has been neglected in the existing literature, with the exception of [Bohren and Kravitz \(2019\)](#) discussed above. Our optimal contract rewards agreement between agents due to correlated states. This property also arises in papers that study how to incentivize agents to costly learn some common state, such as [Pesendorfer and Wolinsky \(2003\)](#), [Miller et al. \(2005\)](#), [Gromb and Martimort \(2007\)](#), [Bohren and Kravitz \(2019\)](#), and [Azrieli \(2021, 2022\)](#). In contrast to these papers, we consider a principal who aims to implement information acquisition as a unique outcome. Also, the similarity of

²There is also a literature on task allocation initiated by [Garicano \(2000\)](#). This literature emphasizes the trade off between knowledge acquisition and communication cost, putting incentive concern aside. Also, papers in this literature typically assume tasks are unrelated.

tasks that agents investigate is endogenous. Rewarding consistency of reports also arises in other settings. For example, in [Deb et al. \(2018\)](#), rewarding consistency is useful for screening a strategic forecaster who observe signals about some persistent state.

Third, our paper complements the literature on spatial learning. The literature is started by [Jovanovic and Rob \(1990\)](#) and [Callander \(2011\)](#), who model the state as a realized path of Brownian motion. These papers use spatial learning to model the set of alternatives in the trial-and-error process.³ Recent papers model complex decision uncertainty using more general processes and discuss how it interacts with agency problems. See [Bardhi \(2022\)](#), [Bardhi and Bobkova \(2023\)](#), and [Dong and Mayskaya \(2023\)](#) for examples. Our paper differs from the above ones in several aspects. First, these papers model uncertainty as a realized path of a Gaussian process. In contrast, we consider a simple non-Gaussian setting. Second, these papers exclude monetary transfers, whereas optimal compensation is an important part of our study. Finally, none of these papers consider peer monitoring and robust implementation.

By focusing on robust implementation, we join the literature on robust implementation in teams. [Winter \(2004\)](#) and [Halac et al. \(2021\)](#) derive optimal contracts that induce all agents to work when the only verifiable information is the entire team’s success. [Camboni and Porcellacchia \(2023\)](#) show that even when the signal of an individual performance is also available, the optimal robust contract may not use it because it could introduce undesirable equilibria. These papers do not consider endogenous task allocation or monitoring structure. [Halac et al. \(2023\)](#) allow the principal to divide agents into groups, each delivering a signal of joint performance, but the number of groups is constrained. [Cusumano et al. \(2023\)](#) study a design of monitoring systems when the number of signals the principal can contract on is finite. In these papers, the principal directly designs the monitoring structure, which generates verifiable information about agents’ actions.⁴ In contrast, the principal in our model determines the peer monitoring structure through task allocation.

Finally, our paper is related to the literature on random monitoring in team production. The way in which the principal benefits from a stochastic task allocation is consistent with papers such as [Legros and Matthews \(1993\)](#), [Rahman and Obara \(2010\)](#), and [Rahman \(2012\)](#). These papers study contracts that induce some agents occasionally choose suboptimal actions

³Also see [Cetemen et al. \(2023\)](#) for a model of collective search based spatial learning.

⁴[Li and Yang \(2020\)](#) study the optimal joint design of incentive contract and monitoring technology. They focus on the trade-off between giving incentives to agents and saving the monitoring cost and apply their analysis to a team production setting.

to identify non-deviators.⁵ The literature also considers informational assumptions under which the principal can attain (nearly) efficient outcomes (see, e.g., Miller (1997) and Strausz (1999)). Our work differs from these papers in that we consider robust implementation, and the principal implements a random monitoring structure through task assignment.

Organization. The rest of the paper is organized as follows. Section 2 describes the model. Section 3 puts the agency problem aside and derives the principal’s optimal task allocation that minimizes information loss. Section 4 characterizes the optimal deterministic task allocation and contracting, and Section 5 allows the principal to assign tasks randomly and privately. Section 6 concludes.

2 Model

Decision Environment. The principal (“he”) faces a continuum of decision problems under uncertainty. Each problem is associated with an uncertain state, and the states for similar problems are more likely to take the same value. We model such a situation as spatial learning: There is a unit mass of locations $t \in [0, 1]$. A function

$$x : [0, 1] \rightarrow \{A, B\}$$

is called the *state*, and $x(t)$ is called the *local state* of location t . We model the state as a realized path of a continuous-time Markov chain. Specifically, $x(0)$ is a random draw from $\{A, B\}$ with equal probability. At each $t \in [0, 1]$, a shock arrives at rate $\lambda > 0$, which induces a change in local state ($x(t) \neq x(t^-) \triangleq \lim_{t' \uparrow t^-} x(t')$) with probability 0.5. If no shock arrives at t , $x(\cdot)$ is continuous at t . Under this specification, the distribution of local state $x(t)$ conditional on a shock at t is independent of other local states.⁶ Figure 1 presents an example of the realized state in which $x(0) = B$ and $x(t)$ changes value at t_1 and t_2 . At t_3 , a shock arrives, but the transition does not occur.

⁵Some recent development in principal-agent literature demonstrates that the opacity in incentive scheme (Ederer et al., 2018) and monitors’ identities (De Janvry et al., 2023) can also mitigate agents’ incentive to game the system. These papers do not consider moral hazard in team production.

⁶This independence would fail if the local state deterministically switched upon each shock, i.e., conditional on a shock at t , $x(t) \neq x(t^-)$ with probability 1. Furthermore, our analysis extends to the case in which (i) the set of local states is an arbitrary finite set, and upon the arrival of each shock, $x(t)$ is independently and uniformly drawn from the set, or (ii) when $x(t)$ ’s distribution is independent but non-uniform upon the arrival of each shock.




Figure 1: Illustration of a realized state where A, B are real numbers such that $A < B$.

The state realization process has several key features. First, the local state of one location is informative about another, and closer locations are more likely to have the same local state. To see this, take two locations, t and $t' > t$. Their local states differ (i.e., $x(t) \neq x(t')$) if and only if (i) at least one shock arrives in the interval $(t, t']$, which occurs with probability $\int_0^{t'-t} \lambda e^{-\lambda s} ds$, and (ii) the realized local state at the last shock is different from $x(t)$, which occurs with probability $1/2$. As a result, we have

$$\Pr[x(t) = x(t')] = 1 - \frac{1}{2} \int_0^{t'-t} \lambda e^{-\lambda s} ds = \frac{1}{2} + \frac{1}{2} e^{-\lambda(t'-t)}. \quad (1)$$

Thus, probability $\Pr[x(t) = x(t')]$ is greater than $1/2$ for any $t \leq t'$ and decreasing and convex in $|t' - t|$. Second, suppose that the principal observes the local states $\{x(t_i)\}_{i=1, \dots, m}$ of locations $t_1 < t_2 < \dots < t_m$. Then for any $j \in \{1, \dots, m-1\}$ and any location $t \in [t_j, t_{j+1}]$, we have

$$\Pr(x(t) = A | \{x(t_i)\}_{i=1, \dots, m}) = \Pr(x(t) = A | x(t_j), x(t_{j+1})), \quad (2)$$

because the state $x(\cdot)$ is a realized path of a Markov process.

For each location, the principal chooses action A or B and incurs a loss whenever the action does not match the local state. Formally, let

$$p : [0, 1] \rightarrow \{A, B\}$$

denote a *policy* of the principal, where $p(t)$ specifies his action for location $t \in [0, 1]$. Given policy p and state x that are piecewise constant (with at most finitely many jumps), the principal incurs a loss

$$\int_0^1 \mathbb{1}(p(t) \neq x(t)) dt,$$

where $\mathbb{1}(\cdot)$ denotes the indicator function.

The principal *privately* observes the local state $x(0)$ of location 0, but not the rest of the path. As a result, without additional information, the principal correctly predicts local state $x(t)$ with probability $\Pr[x(0) = x(t)]$. The probability $\Pr[x(0) = x(t)]$ is decreasing in t and λ , so we can view a decision problem (or a location) with a larger t as less understood or more novel from the principal’s perspective, and λ as the degree of uncertainty the principal faces.

Agents and Information Acquisition. To learn about the uncertain state, the principal hires $n \in \mathbb{N}$ agents (“she”) and assigns each agent i to one location, denoted by t_i . Without loss, we assume that $t_1 \leq t_2 \leq \dots \leq t_n$. The resulting (deterministic) *task allocation* is⁷

$$\tau = (t_1, \dots, t_n).$$

The distance between agent i and her left neighbor (i.e., agent $i - 1$) is denoted by

$$\Delta t_i = t_i - t_{i-1},$$

which captures the task similarity between the two agents. As Δt_i becomes smaller, probability $\Pr[x(t_i) = x(t_{i-1})]$ increases. As a result, the local states that agents i and $i - 1$ are tasked to learn are more likely to take the same value.

Given the task allocation, the agents make two choices. First, each agent simultaneously and privately chooses whether to acquire information ($a_i = 1$) or not ($a_i = 0$). If agent i chooses $a_i = 1$, she incurs a cost of c and privately observes the local state $x(t_i)$ of her assigned location t_i . If agent chooses $a_i = 0$, she incurs no cost but observes a pure noise (say \emptyset) regardless of the realized state. Second, each agent simultaneously sends a report $\hat{x}_i \in \{A, B\}$ to the principal. The profile of reports is denoted by

$$\chi = (\hat{x}_0, \hat{x}_1, \dots, \hat{x}_n).$$

To simplify exposition, we include the principal’s signal (i.e., local state $x(0)$) as a part of a report profile χ . We abstract away from the principal’s incentive to misreport and assume that $\hat{x}_0 = x(0)$ in any report profile.⁸ The profile of reports is the only contractible object.

In our model, the realized state is only partially indicative of an agent’s action: Even if the principal could ex post observe the state and use it to compensate an agent, the principal

⁷In Section 5, we consider a stochastic task allocation.

⁸To maintain the principal’s incentive, one would have to introduce money burning as in MacLeod (2003), making the principal’s payment independent of his report.

would still be unable to induce $a_i = 1$ at the first-best wage (which is c). Indeed, the agent can always shirk, randomly report A or B , and secure a probability 0.5 of matching her report with the true local state. The principal's inability to observe the state further exacerbates the agency problem, because the principal has to rely on agents' reports to compensate them.

The principal commits to a compensation scheme to incentivize agents to acquire information and truthfully report the local states. Let $w(\tau, \chi) \in \mathbb{R}_+^n$ denote the compensation profile, where its i th component $w_i(\tau, \chi)$ specifies the payment to agent i when the report profile is χ under task allocation τ . Agents are risk-neutral and have limited liability, so the payment to each agent must be nonnegative.⁹ The principal's payoff is not contractible. This assumption is justified when the principal's payoff is realized after a long time period or unverifiable.

Timing. The timeline of the game is as follows.

1. The principal chooses a task allocation τ and a compensation scheme w .
2. Nature chooses the state x .
3. Agents observe τ and w and then simultaneously choose whether to acquire information and what to report. The principal observes $x(0)$.
4. The profile of report χ becomes public. Agents receive payment according to the compensation scheme, and the principal chooses a policy p .

Joint Design of Task and Contract. A task allocation τ and a compensation scheme w together define a simultaneous-move incomplete information game between the agents, denoted by $\Gamma(\tau, w)$. In this game, the task allocation, compensation scheme, and distribution of states are common knowledge. Each agent i 's strategy specifies (i) whether to acquire information ($a_i \in \{0, 1\}$), and (ii) what to report conditional on the outcome of information acquisition ($b_i : \{\emptyset, A, B\} \rightarrow \{A, B\}$). The principal wishes to robustly implement the profile such that each agent acquires information and truthfully reports her discovery.

Formally, given a task allocation τ , a contract w *robustly implements work and truth-telling (RIWT)* if $a_i = 1, b_i(y) = y, \forall y = A, B$ is the unique outcome of (correlated) rationalizable

⁹We do not explicitly model an agent's decision to join the team. But one can assume that the agents' outside option is 0, which implies that an agent weakly prefers accepting the assigned task and contract and choosing $a_i = 0$ to taking the outside option.

profiles (Fudenberg and Tirole, 1991) of game $\Gamma(\tau, w)$.¹⁰ Let $\mathcal{W}(\tau)$ denote the set of RIWT contracts given task allocation τ . By the standard argument, we have $w \in \mathcal{W}(\tau)$ if and only if the action profile $a_i = 1$, $b_i(y) = y, \forall y = A, B$ is the only one that survives iterated elimination of strictly dominated strategies (IESDS) in game $\Gamma(\tau, w)$. This implementation is immune to strategic risk and only relies on agents' rationality and beliefs about their opponents' rationality, and so on (Bernheim, 1984; Pearce, 1984). It is also stronger than the usual Nash full implementation, which requires the desired action profile being the unique (Bayes) Nash equilibrium.¹¹

For each task allocation τ , the optimal contract design problem is formulated as

$$K(\tau) \triangleq \inf_{w \in \mathcal{W}(\tau)} \sum_{i=1}^n \sum_{\chi} w_i(\tau, \chi) \Pr(\chi|\tau), \quad (\text{P-C})$$

where $\Pr(\chi|\tau)$ is the probability of a truthful report profile $\chi \in \{A, B\}^{n+1}$ given task allocation τ . The compensation minimization problem (P-C) generally does not have a minimum. We call the infimum of the expected compensation, $K(\tau)$, as the *incentive cost* for task allocation τ .

Task allocation also affects the principal's decision quality. Fix a task allocation τ , and suppose that the principal observes the truthful report profile, $\chi = (x(0), x(t_1), \dots, x(t_n))$. We define the principal's *information loss* as

$$L(\tau) \triangleq \sum_{\chi} \left\{ \min_p \mathbb{E} \left[\int_0^1 \mathbb{1}(p(s) \neq x(s)) ds \mid \tau, \chi \right] \right\} \Pr(\chi|\tau). \quad (\text{P-I})$$

Here, the inner expectation operator $\mathbb{E}[\cdot|\tau, \chi]$ denotes the expectation over realized path $x(\cdot)$ conditional on the profile of truthful reports, χ . The minimization is over the set of admissible policies that consists of all piecewise-constant functions from $[0, 1]$ to $\{A, B\}$. The outer expectation is taken over truthful report profile χ .

In sum, the principal's task design problem is

$$\min_{\tau} L(\tau) + K(\tau),$$

¹⁰When $a_i = 0$, the off-path report $b_i(\emptyset)$ can be arbitrary for the implementation. Hence, we look for a unique outcome of rationalizable profiles.

¹¹For some pair τ, w , the induced game $\Gamma(\tau, w)$ may have a unique Nash equilibrium but multiple rationalizable action profiles ($\Gamma(\tau, w)$ is in general not supermodular). In this sense, our implementation criteria is harder to satisfy than Nash full implementation.

where for each task allocation τ , $K(\tau)$ is the incentive cost in program (P-C), and $L(\tau)$ is the information loss in program (P-I). As we will show, the principal faces a trade-off, because the task allocation that minimizes information loss $L(\cdot)$ is different from the one that minimizes incentive cost $K(\cdot)$.

3 Optimal Policy and Information Loss

We first characterize the principal's optimal policy and task allocation that solve problem (P-I), temporarily ignoring the agency problem. The problem is reduced to optimal sampling with size n to minimize the overall error of state estimation. Without incentive concerns, the principal make agents acquire information at sufficiently distant locations.

Proposition 1. *Fix a task allocation τ and a truthful report profile χ , the principal's optimal policy is given by*

$$p(t) = \begin{cases} x(0) & \text{if } t \in [0, \frac{t_1}{2}) \\ \hat{x}(t_i) & \text{if } t \in [\frac{t_{i-1}+t_i}{2}, \frac{t_i+t_{i+1}}{2}), \forall i = 1, 2, \dots, n-1, \\ \hat{x}(t_n) & \text{if } t \in [\frac{t_{n-1}+t_n}{2}, 1] \end{cases} \quad (3)$$

and the corresponding information loss is

$$L(\tau) = \frac{1}{2} - \frac{2n+1}{2\lambda} + \frac{1}{\lambda} \left[\sum_{i=1}^n e^{-\frac{\lambda(t_i-t_{i-1})}{2}} + \frac{1}{2} e^{-\lambda(1-t_n)} \right]. \quad (4)$$

Moreover, the information loss is minimized by the task allocation τ^\dagger such that

$$t_i^\dagger = \frac{2i}{2n+1}, \forall i = 1, 2, \dots, n. \quad (5)$$

The corresponding information loss is $L(\tau^\dagger) = \frac{1}{2} + \frac{2n+1}{2\lambda} \left(e^{-\frac{\lambda}{2n+1}} - 1 \right)$, which decreases in n and increases in λ .

We refer to τ^\dagger in expression (5) as the *efficient* or first-best task allocation. We will use τ^\dagger as the benchmark to evaluate the distortion in task allocation due to moral hazard.

We first discuss the optimality of policy (3). Armed with truthful reports χ , the principal makes perfectly informative decisions at locations t_0, t_1, \dots, t_n . At other locations $t \neq t_i, \forall i$, the principal's optimal action aligns with the nearest known local state, i.e.,

$$p(t) = \hat{x}(t_i^*), \text{ where } t_i^* = \arg \min_{t_i \in \{t_1, \dots, t_n\}} |t - t_i|.$$

That is, each report \hat{x}_i establishes a *range of illumination*, $[\frac{t_{i-1}+t_i}{2}, \frac{t_i+t_{i+1}}{2}]$. Within this range, the principal's optimal policy is guided by the report $\hat{x}(t_i)$. By this policy, the principal's *decision quality*, as measured by the likelihood of matching the local state $x(t)$ with the action $p(t)$, achieves its maximum at $t = t_i$ and diminishes as t diverges from t_i . At $\frac{t_i+t_{i+1}}{2}$, the principal is indifferent to match his action to \hat{x}_i and \hat{x}_{i+1} . Figure 2 depicts the decision quality at different locations and the corresponding information loss.

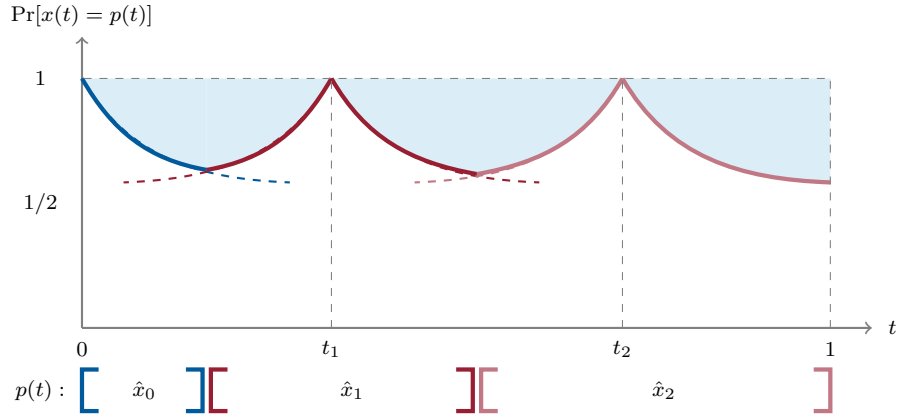


Figure 2: Illustrating decision quality under the efficient allocation with $n = 2$. The brackets correspond to each report's illumination range. The solid thick curve corresponds to the principal's decision quality under the optimal policy, and the grey area corresponds to the principal's information loss $L(\tau)$ under the optimal policy.

Proposition 1 also says that, without incentive concern, the optimal design ensures ample *diversification* in the tasks allocated to agents. Specifically, according to the formula (5), the




Figure 3: Illustrating the marginal impact of increasing t_i on decision quality for fixed t_{i-1} and t_{i+1} when t_i is on the left side of its illumination range. The marginal loss and the marginal gains are the corresponding areas delineated by four curves.

efficient task allocation equalizes the illumination range of each report \hat{x}_i , i.e., $\forall i = 1, \dots, n - 1$,

$$t_{i+1}^\dagger - t_i^\dagger = \Delta t^\dagger \triangleq \frac{2}{2n + 1}. \quad (6)$$

To see the intuition, fix t_{i-1} and t_{i+1} and consider the principal's choice of t_i . Suppose that task t_i is on the left-side of its illumination range (see Figure 3). The principal can then slightly move t_i to its right and decrease the information loss. Indeed, within the illumination range of t_i , the decision quality on the left side of t_i decreases but the decision quality on the right hand side increases. Because there are more tasks on the right hand side of the illumination range than the left hand side, the gain is greater than the loss. The symmetric argument holds when t_i is on the right side of its illumination range. As a result, each t_i must be in the middle of its illumination range, and task assignment (6) is the unique one with such a property.

With more agents (larger n), the illumination range of each agent's report under the efficient task allocation becomes narrower, which increases the principal's overall decision quality. As λ increases, the agent's report becomes less informative about the local states within its illumination range, which reduces the principal's overall decision quality.

4 Optimal Task Allocation and Contracting

We now study the optimal joint design of a task allocation and a compensation scheme that minimize the principal's total loss.

4.1 Least-Cost Compensation and Incentive Cost

We begin with the principal's optimal contract design problem (P-C).

Proposition 2. *Fix any task allocation τ . For each $\epsilon \geq 0$, define the compensation scheme w^ϵ as follows: For each $i = 1, \dots, n$ and report profile χ ,*

$$w_i^\epsilon(\chi) \triangleq \begin{cases} 2ce^{\lambda(t_i - t_{i-1})} + \epsilon & \text{if } \hat{x}_i = \hat{x}_{i-1} \\ 0 & \text{otherwise} \end{cases}. \quad (7)$$

For any $\epsilon > 0$, compensation scheme $w^\epsilon(\chi)$ robustly implements work and truth-telling (RIWT). The principal's incentive cost $K(\tau)$ coincides with the expected payment under $w^0(\chi)$, where

$$K(\tau) = c \sum_{i=1}^n [1 + e^{\lambda(t_i - t_{i-1})}]. \quad (8)$$

Moreover, for any task allocation τ , as the state uncertainty vanishes, we have

$$\lim_{\lambda \rightarrow 0} K(\tau) = K^\dagger \triangleq 2nc. \quad (9)$$

Finally, $K(\tau) \geq K^\dagger$ for any τ . We have $K(\tau) = K^\dagger$ if τ is such that for each agent i 's assignment $t_i > 0$, either $t_i = t_{i-1}$ or $t_i = t_{i+1}$.

Proposition 2 has three parts. First, it presents compensation schemes that robustly implement work and truth-telling. Under compensation scheme w^ϵ , an agent is paid if and only if her report is consistent with that of her "left neighbor." Thus the optimal contract induces a monitoring chain: agent 1 is monitored by the principal's report, and other agents are monitored by their peers. The distance $t_i - t_{i-1}$ between two adjacent agents determines the correlation of their truthful reports, thus defining the efficiency of peer monitoring. Second, such a compensation scheme is the cheapest way to robustly implement work and truth-telling, in the sense that the incentive cost is the infimum of the expected payments associated with w^ϵ across all ϵ . Third, the result points out two cases in which the principal can perfectly verify whether the agent's report matches the true state. One is when the local state never changes (i.e., $\lambda = 0$), in which case we have $x(0) = x(t_i)$ for any $t_i > 0$. The other one is when at least two agents are assigned to each location so that their reports must be perfectly correlated, maximizing the efficiency of peer monitoring.

The optimal contract in Proposition 2 contrasts with the informativeness principle (Holmström, 1979, 1982): Payment to each agent does not depend on the report of her right neighbor, although it is a part of the sufficient statistics for the agent's effort and report (see equation (2)). The reason is that in contrast to papers such as Holmström (1982), which focus on partial implementation, we focus on implementation as a unique (correlated) rationalizable profile.

To understand differences between the two approaches, consider two agents. Under partial implementation—where the principal can pick his favorite equilibrium—the optimal contract pays agent 1 if and only if $\hat{x}_0 = \hat{x}_1 = \hat{x}_2$ and agent 2 if and only if $\hat{x}_1 = \hat{x}_2$.¹² Moreover, to minimize the expected payment, the principal set payments so that each agent's incentive-compatibility constraint holds with equality. However, the induced game between the two agents also has an equilibrium in which no agent acquires information. Indeed, if agent 1 shirks, the

¹²The logic follows the standard argument. For each risk-neutral agent, it is optimal to pay her only in the event that maximizes the likelihood ratio of her information acquisition and truth-telling relative to no information acquisition, which is the event $\hat{x}_0 = \hat{x}_1 = \hat{x}_2$ for agent 1 and $\hat{x}_1 = \hat{x}_2$ for agent 2.

probability of $\hat{x}_1 = \hat{x}_2$ becomes independent of the effort and report of agent 2, so she will shirk as well. But if agent 2 shirks, the probability of $\hat{x}_0 = \hat{x}_1 = \hat{x}_2$ under agent 1’s effort and truthful reporting decreases, which violates agent 1’s incentive compatibility. The multiplicity of equilibria given the optimal contract under partial implementation stems from *monitoring externality*: One agent’s information acquisition and truth-telling enables the principal to better evaluate other agents. The optimal contract under partial implementation reflects mutual peer monitoring. It exploits the monitoring externality well but introduces strategic risk among agents due to its coordination feature, causing the multiplicity of equilibria.¹³

The optimal contract in Proposition 2 rules out the mutual peer monitoring to gain strategic robustness: Agent 1’s pay depends only on her own report and the principal’s, curbing her from the strategic risk induced by other agents’ choices. Once the incentive of agent 1 is established, the incentives of agents 2, 3, ..., n ’s can be pinned down in a similar manner.¹⁴ The optimality of the monitoring chain suggests a *robust informativeness principle*. The optimal contract isolates the signal generating process of $\{\hat{x}_0, \dots, \hat{x}_{i-1}\}$ from agent i ’s strategic influence, making these reports a reliable source of information to monitor agent i ’s action. Moreover, due to the Markov property of the state, it is sufficient to compare \hat{x}_i with \hat{x}_{i-1} in determining agent i ’s payment.

The rest of this section offers some intuition behind the optimality of the compensation scheme in equation (7). We begin with the case of a single agent and then the case of two agents. The full analysis is in the appendix.

Single-Agent Case. First, suppose that $n = 1$. For simplicity, we use “work” for the agent’s choice to acquire information and report truthfully and “shirk” for any strategy that does not acquire information. The event in which the agent’s report coincides with the principal’s signal (i.e., $x(0) = \hat{x}(t_1)$) maximizes the likelihood ratio of the agent’s working relative to shirking. Thus, the optimal contract pays the agent if and only if her report equals the principal’s private signal.

When the agent works, she incurs a cost c and receives compensation with a probability of $(1 + e^{-\lambda t_1})/2$, as indicated by expression (1). Conversely, the agent who shirks will avoid cost c , provide an uninformed report, and earn the compensation with probability $1/2$. Consequently, the agent will strictly prefer to work (i.e., the contract belongs to $\mathcal{W}(\tau)$) if her payment w_1

¹³See Appendix B for the optimal partial-implementation contract in general settings.

¹⁴Camboni and Porcellacchia (2023) recognize a similar information-waste result in a super-modular setting where the principal observes a team performance measure and a signal about each individual’s action.

given $\hat{x}_1 = x(0)$ satisfies the inequality

$$\frac{1}{2}(1 + e^{-\lambda t_1})w_1 - c > \frac{1}{2}w_1.$$

To minimize expected compensation, the principal sets $w_1 = 2ce^{\lambda t_1}$. The agent has the incentive to work because it increases the probability of matching her report with the principal's by $e^{-\lambda t_1}$, which is decreasing and convex. As t_1 increases, w_1 has to increase accordingly to maintain the agent's incentive compatibility. As $t_1 \rightarrow 0$, the payment converges to $w_1 = 2c$ rather than 0, because the agent can always make an uninformed report and secure an expected payoff $w_1/2$.

The expected incentive cost is the sum of the agent's individually rational payoff $w_1/2 = ce^{\lambda t_1}$ and information acquisition cost c , or

$$K(t_1) = c(1 + e^{\lambda t_1}). \quad (10)$$

The expected payment is linear in the agent's information acquisition cost and increases in t_1 . The coefficient's inverse, $1/(1 + e^{\lambda t_1})$, is called agent i 's (peer) *monitoring intensity*. The greater the monitoring intensity, the more closely the agent is monitored, and the less costly it becomes to induce her effort.

Two-Agent Case. Assume now that $n = 2$ and $t_1 \leq t_2$. For the desired outcome (i.e., both agents work) to be the unique (correlated) rationalizable profile, it has to be the unique outcome surviving IESDS in the induced game $\Gamma(w, \tau)$. The order of eliminating dominated strategies leads to two sets of contracts. Each set corresponds to a specific monitoring structure, as illustrated in Figure 4.

In the first case (i.e., the left panel of Figure 4), the first step of the IESDS procedure makes it strictly dominant for agent 1 to work. In the second step, the incentive for agent 2 is established. Within this set of contracts, we can determine the optimal compensation iteratively. In each of these iterations, the principal effectively deals with a single agent at a time. Specifically, to guarantee that agent 1 finds it strictly dominant to work, agent 1's contingent payment should only be based on her own and the principal's reports. Following the logic from the single-agent scenario, agent 1 receives payment if and only if $\hat{x}(t_1) = x(0)$. To solidify the incentive for information acquisition, the compensation must exceed $2ce^{\lambda t_1}$ slightly.

Once agent 1's incentive is guaranteed, the principal can rely on her report to monitor agent 2. By a parallel argument, to ensure that agent 2 finds it strictly dominant to work (conditional on agent 1 working), she should be paid if and only if $\hat{x}(t_2) = \hat{x}(t_1)$ and the payment should be




Figure 4: The left panel shows the chain monitoring structure in case 1, where the principal’s signal $x(0)$ is used to monitor agent 1 and agent 1’s report is used to monitor agent 2. The right panel shows the sandwich monitoring structure in case 2, where the principal’s signal is used to monitor agent 2 and agent 1 is monitored jointly by the principal’s signal and agent 2’s report. Each arrow points from the monitor to the monitored, with the associated roman numeral indicating the order of establishing information acquisition and truthful reporting as a strictly dominant strategy.

marginally above $2ce^{\lambda(t_2-t_1)}$. This case leads to a “chain” monitoring structure. The incentive cost under this chain monitoring structure is written as:

$$K^c(t_1, t_2) \triangleq c \sum_{i=1,2} \left[1 + e^{\lambda(t_i-t_{i-1})} \right], \quad (11)$$

which is the sum of the minimum expected payment to each agent $K_i^c(\tau) = c[1 + e^{\lambda(t_i-t_{i-1})}]$. Similar to the single-agent case, $K_i^c(\tau)$ is linear in c , and the coefficient’s inverse captures the monitoring intensity for agent i .

In the second case (i.e., the right panel of Figure 4), the IESDS procedure establishes agent 2’s dominance incentive in the first step, paying her if and only if $\hat{x}(t_2) = x(0)$. The payment must be slightly higher than $w_2 = 2ce^{\lambda t_2}$. Thus the incentive cost for agent 2 is $c(1 + e^{\lambda t_2})$.

In the second step, agent 1’s action can be jointly monitored by the principal’s private signal and agent 2’s report. The optimal contract pays agent 1 if and only if her report matches the reports of both the principal (her left neighbor) and agent 2 (her right neighbor), i.e.,

$$\hat{x}(t_1) = x(0) = \hat{x}(t_2).$$

In this case, the contract induces a “sandwich” monitoring structure.

To ensure agent 1’s strict incentive to work, the compensation w_1 , which is paid only when $x(t_1) = x(0) = \hat{x}(t_2)$, must satisfy

$$\frac{1}{4} \left(1 + e^{-\lambda t_1} \right) \left(1 + e^{-\lambda(t_2-t_1)} \right) w_1 - c > \frac{1}{4} \left(1 + e^{-\lambda t_2} \right) w_1. \quad (12)$$

The left-hand side is agent 1’s expected utility from working. By doing so, she incurs cost c and receives payment w_1 with probability $(1 + e^{-\lambda t_1})(1 + e^{-\lambda(t_2-t_1)})/4$. The right-hand side of

condition (12) is agent 1's payoff of shirking. In this case, she avoids incurring the cost and receives payment w_1 if her uninformed report matches the principal and agent 2's reports. As in the case of a single agent, inequality (12) gives us the infimum of w_1 that induces a strict incentive to work. The corresponding incentive cost for agent 1 is

$$c \left[1 + \frac{1 + e^{-\lambda t_2}}{e^{-\lambda t_1} + e^{-\lambda(t_2-t_1)}} \right]$$

To sum up, the infimum of the principal's expected payment under the sandwich monitoring structure is

$$K^s(t_1, t_2) = \underbrace{c \left[1 + \frac{1 + e^{-\lambda t_2}}{e^{-\lambda t_1} + e^{-\lambda(t_2-t_1)}} \right]}_{K_1^s(\tau)} + \underbrace{c \left[1 + e^{\lambda t_2} \right]}_{K_2^s(\tau)}, \quad (13)$$

where agent i 's monitoring intensity $c/K_i^s(\tau)$ differs from that under chain monitoring structures.

The principal's optimal contract design boils down to the comparison between $K^c(t_1, t_2)$ and $K^s(t_1, t_2)$. In what follows, we argue that the chain monitoring structure attains a lower incentive cost, i.e.,

$$K^s(t_1, t_2) \geq K^c(t_1, t_2), \forall t_1 \leq t_2.$$

Note that in both cases, fix t_2 , the expected compensation payment is symmetric in t_1 ; i.e.,

$$K^i(t, t_2) = K^i(t_2 - t, t_2), \forall t \in [0, t_2/2], i = c, s.$$

Also, the incentive costs in two cases are identical if $t_1 = 0$ or $t_1 = t_2$. We claim that at any $t_1 \in [0, t_2/2)$, $K^c(\cdot, t_2)$ is decreasing and $K^s(\cdot, t_2)$ is increasing. We plot K^c and K^s in Figure 5.

First, we differentiate $K^c(t_1, t_2) = c \sum_{i=1,2} \left[1 + e^{\lambda(t_i - t_{i-1})} \right]$ with respect to t_1 and obtain

$$\frac{\partial K^c}{\partial t_1} = c\lambda e^{\lambda t_1} - c\lambda e^{\lambda(t_2-t_1)}.$$

The first term of the left-hand side reflects the marginal impact on agent 1's expected compensation, and the second term reflects the marginal effect on agent 2's expected compensation. As t_1 increases, the principal's report becomes less informative in monitoring agent 1, resulting in an increase in the incentive cost for agent 1, but agent 1's truthful report becomes more informative in monitoring agent 2, resulting in a decrease in the incentive cost for agent 2. Overall, the total incentive cost decreases because the second effect dominates when $t_1 < t_2/2$ due to the convexity of the individual agent's expected compensation in her distance from her left neighbor. By symmetry, $K^c(\cdot, t_2)$ reaches its minimum at $t_1 = t_2/2$.




Figure 5: Illustrating expected payment functions $K^c(\cdot, t_2)$ (chain monitoring structure) and $K^s(\cdot, t_2)$ (sandwich monitoring structure).

Now consider the effect of increasing t_1 on $K^s(\cdot, t_2)$. For a given t_2 , t_1 affects $K^s(\cdot, t_2)$ only through $e^{-\lambda t_1} + e^{-\lambda(t_2-t_1)}$, which decreases in t_1 whenever $t_1 < t_2 - t_1$. As a result, $K^s(\cdot, t_2)$ increases in t_1 when $t_1 < t_2/2$. By symmetry, $K^s(\cdot, t_2)$ reaches its maximum at $t_1 = t_2/2$. The intuition is as follows. Agent 1 is paid if only if (i) the principal's and agent 2's reports are identical and (ii) they are aligned with agent 1's report. The first event occurs with probability $(1 + e^{-\lambda t_2})/2$, which is independent of t_1 . Conditional on the first event, agent 1 is more likely to be paid if her location is closer to either the principal's or agent 2's. Therefore, as t_1 moves from 0 toward $t_2/2$, agent 1's probability of being paid decreases, and so w_1 must increase accordingly to maintain her incentive compatibility.

In sum, when $n = 2$, the principal's incentive cost is $K(t_1, t_2) = K^c(t_1, t_2)$ for any $t_1 \leq t_2$. In the appendix, we extend the above argument to an arbitrary number of agents. To understand the intuition, consider the example of three agents in Figure 6. Consider a contract (τ, w) such that the order of IESDS is $t_2 \rightarrow t_1 \rightarrow t_3$, i.e., agent 2 is monitored by the principal's signal, agent 1 is monitored by the principal and agent 2's reports, and agent 3 is monitored by agent 2's report (the left panel of Figure 6). Notice that the principal, agent 1, and agent 2 form a local sandwich monitoring structure. The incentive cost can be written as

$$K^s(t_1, t_2) + c[1 + e^{\lambda(t_3-t_2)}],$$

where the last term is the minimum incentive-compatible expected payment to agent 3. Now, suppose we "locally" modify the order of IESDS as $t_1 \rightarrow t_2 \rightarrow t_3$. Then the principal, agents 1, and 2 form a chain monitoring structure (the right panel of Figure 6). The incentive cost is

$$K^c(t_1, t_2) + c[1 + e^{\lambda(t_3-t_2)}].$$



Figure 6: Illustrating the optimality of chain monitoring structure when $n = 3$.

By the previous argument, $K^c(t_1, t_2) \leq K^s(t_1, t_2)$. Because agent 3 is still monitored by agent 2's signal, agent 3's expected payment remains the same. As a result, the local modification is weakly profitable.

In general, whenever the order of IESDS induced by a contract is different from that of the chain monitoring structure, we can locally modify the contract to decrease the expected payments to the agents. As a result, the chain monitoring structure is robustly optimal for any number of agents.

Remark 1 (Robustness to collusion). The contract in Proposition 2 is susceptible to collusion: If the agents could privately send cheap-talk messages to each other before sending reports to the principal, agent 1 would reveal $x(t_1)$ to other agents, who will shirk and report $\hat{x}_i = x(t_1)$ to the principal. The resulting outcome gives strictly higher payoffs to agents $i = 2, \dots, n$ than the original equilibrium without changing agent 1's payoff. However, the principal can prevent this kind of collusion (which does not involve side payments) by modifying each $w_i^\epsilon(\chi)$ as $\hat{w}_i^\epsilon(\chi) \triangleq w_i^\epsilon(\chi) - \delta \mathbb{1}(\exists j \neq i, \hat{x}_j = \hat{x}_i)$, i.e., the principal penalizes agent i if her report matches the report of at least one agent.¹⁵ For any $\delta > 0$, no agent is willing to communicate their report truthfully to other agents.¹⁶ Moreover, so long as δ is sufficiently small relative to ϵ , contract $\hat{w}^\epsilon(\chi)$ robustly implements work and truth-telling.

4.2 Optimal (Deterministic) Task Allocation

Propositions 1 and 2 present a tension: On the one hand, the principal wishes to diversify agents' tasks to reduce the information loss. On the other hand, the principal wants to assign

¹⁵This modification would not prevent collusion with side payments, because the agents can still increase joint payoffs with the aforementioned information-sharing strategy.

¹⁶Formally, if the agents can send cheap-talk messages simultaneously before reporting to the principal, the only equilibrium of this cheap-talk game would be the babbling equilibrium.

agents similar tasks to reduce the incentive cost. This trade-off determines the solution to the principal's task allocation problem,

$$\min_{\tau} L(\tau) + K(\tau),$$

where the information loss $L(\tau)$ is given by equation (4) and the incentive cost $K(\tau)$ is given by equation (8). To state the next result, define the cutoff cost as follows:

$$\bar{c} \triangleq \frac{1 - e^{-\lambda}}{2\lambda}. \quad (14)$$

Recall that Δt_i^\dagger , defined by equation (6), is the distance between any two neighboring agents at the allocation that minimizes information loss. The following result characterizes the optimal task allocation.

Proposition 3. *Under the optimal task allocation, there is a unique $\Delta t^* \in [0, \Delta t^\dagger)$ such that*

$$t_i^* - t_{i-1}^* = \Delta t^*, \forall i = 1, 2, \dots, n,$$

with $t_0^* = 0$. If $c \geq \bar{c}$, we have $\Delta t^* = 0$. If $c < \bar{c}$, we have $\Delta t^* = \frac{2}{\lambda} \log x > 0$, where x solves

$$x^{2(n-1)} + 2\lambda e^\lambda c = \frac{e^\lambda}{x^3}. \quad (15)$$

Finally, Δt^* decreases in c and n , but increasing λ has an ambiguous effect on Δt^* .

Proposition 3 implies that the optimal task allocation equalizes the distance $t_i - t_{i-1}$ between any two neighboring agents to be Δt^* . The inequality $\Delta t^* < \Delta t^\dagger$ means that the agents work on more similar tasks than under the efficient allocation. Similar tasks generate more correlated outputs, enabling the principal to monitor the agents more effectively and reduce incentive costs. Meanwhile, so long as the cost c of information acquisition is below \bar{c} , the principal assigns agents to different tasks and learns nontrivial information about the state. We call

$$\left(n + \frac{1}{2}\right) \Delta t^* = \frac{\Delta t^*}{\Delta t^\dagger},$$

the principal's *knowledge frontier*. Within the frontier (i.e., $t \in [0, \Delta t^*/\Delta t^\dagger]$), the principal's decision quality is non-monotone, reaching maximum at each agent's assigned location t_i^* and minimum at the midpoints between adjacent agents' locations, $(t_i^* + t_{i-1}^*)/2$. However, at any location $t > \Delta t^*/\Delta t^\dagger$, the decision quality is lower than those within the knowledge frontier.

To understand the cutoff \bar{c} defined by (14), note that equations (4) and (8) imply

$$\frac{\partial L(0)}{\partial \Delta t} + \frac{\partial K(0)}{\partial \Delta t} = -n \left(\frac{1 - e^{-\lambda}}{2} \right) + nc\lambda.$$

The first term is negative and captures the benefit of diversifying the allocation of tasks. The second term is positive and equals the marginal increase in the incentive cost. If $c \geq \bar{c}$, the second effect dominates, so the principal chooses $\Delta t^* = 0$ and allocates all the agents to the same task.




Figure 7: Illustrating the solution to equation (15). The increasing dashed curve is the left-hand side and the decreasing solid curve is the right-hand side of equation (15).

Equation (15) corresponds to the principal's first-order condition with respect to Δt^* . The equation typically has no closed-form solution but enables us to establish the comparative statics in Proposition 3. Figure 7 depicts the left-hand and right-hand sides of equation (15).¹⁷ As the cost c of information acquisition falls, the left-hand side (i.e., the dashed curve) shifts downward, resulting in a larger x and thus Δt^* . At $c = 0$, equation (15) has solution $x = e^{\frac{\lambda}{2n+1}}$ and thus $\Delta t^* = \frac{2}{2n+1}$, i.e., the optimal task allocation coincides with the efficient allocation.

Second, as the team size grows (i.e., n increases), the dashed curve becomes steeper, which results in a smaller x and thus Δt^* . With a larger team, the principal receives more reports from the agents and attains a lower information loss. As a result, the principal effectively puts more weight on saving incentive costs. To do so, he assigns agents to work more closely with each other for more efficient peer monitoring.

¹⁷As we can see from the figure, equation (15) implies Δt^* exists and is unique. Indeed, the left-hand side of the equation strictly increases, whereas the right-hand strictly decreases from $+\infty$ (as $x \rightarrow 0$) to 0 (as $x \rightarrow \infty$).

Figure 7 also indicates that increasing λ has an ambiguous effect on Δt^* because it shifts both curves upwardly. Intuitively, increasing λ lowers the correlation between the local states assigned to agents, making it harder for the principal to monitor the agents. However, increasing λ also makes it more costly for the principal to deviate from the efficient allocation, because the local state of one location is less informative about the local state of another location. Depending on which effect dominates, a higher λ may increase or decrease Δt^* . One exception to this ambiguous comparative statics with respect to λ is when the principal hires only one agent ($n = 1$): equation (15) implies $t_1^* = \Delta t^* = -\frac{2}{3\lambda} \ln(e^{-\lambda} + 2\lambda c)$, i.e., t_1^* decreases as λ grows.

4.3 Discussion

Optimal Team Size. We have assumed that the principal hires an exogenous number of agents and induces every agent to acquire information. Suppose, instead, that the principal chooses the number of agents to hire and the action profile to implement. Then we can always assume that the principal induces all employed agents to acquire information, because inducing information acquisition by $n' < n$ agents is equivalent to hiring n' agents and inducing all of them to acquire information. The remaining question is how many agents to hire. The cost of hiring agents is a higher expected payment and the benefit is the reduction of information loss. Because the maximal information loss is bounded, the principal will optimally hire finitely many agents. Moreover, Proposition 3 implies that in a larger team, the agents work at closer locations, increasing the peer monitoring efficiency and lowering the incentive cost per agent.

Optimal Principal's Location. The benchmark model assumes that the principal observes the local state of location $t_0 = 0$. We obtain qualitatively the same results even when the principal can choose which local state $x(t_0)$ to observe, so long as the principal can commit to his location t_0 . In this case, there is an optimal task allocation $\tau = (t_0, t_1, \dots, t_n)$ such that: $t_0 \leq t_1 \leq \dots \leq t_n$; τ is symmetric around $1/2$; and any two neighboring locations are equidistant.¹⁸ Given such a task assignment, the principal adopts the chain monitoring structure, i.e., each agent $i \geq 2$ is monitored by agent $i - 1$'s signal, and agent 1 is monitored by the principal's signal. In particular, it is without loss for the principal to locate at the left of

¹⁸For any fixed t_0 and t_n , $t_{i+1} - t_i$ must be constant across $i = 0, \dots, n - 1$ by the same argument as in Figure 5. The optimal assignment also satisfies $t_0 = 1 - t_n$; for example, if $t_0 < 1 - t_n$, then the principal can reduce the information loss without changing the incentive cost by slightly shifting every t_i to the right by the same amount. Therefore, the optimal $\tau = (t_0, t_1, \dots, t_n)$ is symmetric around $1/2$ and equidistant.




Figure 8: The left panel shows a monotone optimal task allocation $0 < t_0^* < t_1^* < t_2^* < 1$. The right panel shows a non-monotone optimal task allocation $0 < t_1^{**} < t_0^{**} < t_2^{**} < 1$. Each arrow points from the monitor to the monitored, with the associated roman numeral indicating the order of establishing information acquisition and truthful reporting as a strictly dominant strategy.

agent 1, because swapping the principal and agent 1's locations affects neither the incentive cost nor the information loss. For example, suppose that the principal originally chooses $(t_0^{**}, t_1^{**}, t_2^{**})$ with $t_1^{**} < t_0^{**} < t_2^{**}$ (see Figure 8). The corresponding contract compares each agent's report to the principal's signal. However, the principal can attain the same incentive cost and information loss with $(t_0^*, t_1^*, t_2^*) = (t_1^{**}, t_0^{**}, t_2^{**})$ and the chain monitoring structure.

Imperfect Information Acquisition. Suppose that agents' effort does not perfectly reveal the local state. Concretely, suppose that each agent observes a signal $s_i \in \{A, B\}$ that equals $x(t_i)$ with probability $q \in (0.5, 1)$ if she acquires information. The principal also observes an imperfect signal of $x(0)$ with the same signal structure, and the signals are conditionally independent. Then the chain monitoring structure is not optimal, because the left neighbor's report is no longer a sufficient statistic for all information that efficiently identifies an agent's deviation. Instead, the principal should reward an agent when her report matches with all agents whose dominance incentives have been established in the previous rounds of IESDS. As a result, it becomes cheaper to induce effort of agents whose dominance incentives are established in later rounds of IESDS, because their reports are compared with more reports. For two agents, a variation of the chain monitoring structure is optimal: Each agent's report is compared with all information sources located at her left side. The optimal task allocation satisfies $t_2^* - t_1^* > t_1^*$, i.e., the agents are no longer equi-distant. See Online Appendix C.3 for the details.

On the Optimality of Chain Monitoring Structure. The optimality of the chain structure relies on at least two other assumptions. One is that all agents have the same cost of acquiring information. Otherwise, the optimal monitoring structure can be the sandwich monitoring structure (see, e.g., Figure 4). For intuition, recall the discussion below Proposition 2 with two agents: Compared with the chain monitoring, the sandwich monitoring structure saves incentive

cost for agent 1 but increases incentive cost for agent 2. When agent 1 faces a much higher information acquisition cost than agent 2, the cost saving outweighs the cost increment, and the sandwich monitoring structure becomes optimal (see Online Appendix C.1 for the details).

The other assumption is that the set of locations is a line. If we instead consider a circle, the optimal contract will entail a sandwich monitoring structure: At least, any agent is “sandwiched” by the principal on the circle. Due to symmetry of the circle, the optimal task allocation is not unique and the agents’ locations may not be equidistant (see Online Appendix C.2 for the details).

5 Random Assignment and Transparency

In this section, we show that if each agent privately observes her assigned location, the principal can randomize task assignment to simultaneously attain the minimal incentive cost and the first-best information loss.

Private task allocation is ubiquitous in real world. For instance, on crowdsourcing platforms such as Amazon Mechanical Turk, requestors can assign the same task to multiple workers without revealing the number of other workers assigned to the same task.¹⁹ A faculty member may hire research assistants from a large pool of students—who do not know each other—and secretly assign them the same dataset. Private task allocation would also be relevant to a company hiring external consultants or a granting agency hiring anonymous reviewers, who have limited opportunities to communicate with each other. In such a situation, the principal can substantively benefit from keeping task assignment private and uncertain.

Formally, stochastic task allocations are defined as follows:

Definition 1. *A (stochastic) task allocation policy is a pair (T, π) : T is a finite set of deterministic task allocations; $\pi \in \Delta(T)$ specifies the probability of each allocation being realized.*

Recall that $\tau^\dagger = (t_1^\dagger, \dots, t_n^\dagger)$ and $L(\tau^\dagger)$ are the first-best task allocation and corresponding information loss (Proposition 1). The minimal incentive cost under moral hazard is denoted by $K^\dagger := 2nc$, which arises when the principal can perfectly verify whether each agent’s report coincides with the local state of her assigned location (Proposition 2).

To highlight the intuition, we begin with the example in which the principal hires one agent, who does not observe the realized task assignment. For simplicity, we dispense with

¹⁹See <https://blog.mturk.com/cooking-tip-5-ask-multiple-workers-to-complete-a-hit-ec21c9fc0734>.

robust implementation and assume that the agent breaks ties in favor of the principal. Suppose that the principal assigns the agent to location $t_1 = 0$ or $t_1 = t_1^\dagger$ with probability $p \in (0, 1)$ or $1 - p$, respectively. The compensation $w(t_1, \chi)$ depends on the realized task allocation t_1 and the report profile $\chi = (x_0, \hat{x}_1)$. The same argument as the proof of Proposition 2 implies that the principal pays a positive amount only if the agent's report is aligned with the principal's signal $x(0)$ at each realized location. Hence, the principal's contract design becomes as follows:

$$\min_{w(0), w(t_1^\dagger) \geq 0} pw(0) + (1 - p)\frac{1}{2}(1 + e^{-\lambda t_1^\dagger})w(t_1^\dagger) \quad (16)$$

s.t.

$$pw(0) + (1 - p)\frac{1}{2}(1 + e^{-\lambda t_1^\dagger})w(t_1^\dagger) - c \geq \frac{1}{2}[pw(0) + (1 - p)w(t_1^\dagger)], \quad (17)$$

where $w(0)$ and $w(t_1^\dagger)$ denote the payments when the realized locations are 0 and t_1^\dagger , respectively. The constraint binds at the optimum, and we obtain the following solution:

$$w^*(0) = \frac{2c}{p}, w^*(t_1^\dagger) = 0. \quad (18)$$

In other words, the principal rewards the agent only if (i) she is perfectly monitored ($t_1 = 0$), and (ii) her report matches the true state $x(0)$. Even though the probability of a positive payment could be small, the compensation $w^*(0)$ is high enough to induce the agent to acquire information and tell the truth, regardless of the realized task allocation. The resulting expected payment is $2c$, which is the same as when the principal can perfectly verify whether the agent's report is aligned with the true local state. Taking $p \rightarrow 0$, the principal can implement the first-best task assignment t_1^\dagger with probability approaching 1.

Intuitively, the stochastic task allocation achieves conflicting goals through different realized assignments: Realization $t = t_1^\dagger$ minimizes information loss, while realization $t = 0$ enforces the most effective monitoring. Although the principal faces $t = t_1^\dagger$ most of the time, he will reward the agent only when the task realization is $t = 0$. The agent does not know the realized task assignment but understands that the probability of $t = 0$ is small. As a result, in order to satisfy the agent's incentive compatibility constraint in expectation, the reward must be inversely proportional to the probability of $t = 0$, leading to formula (18).

We now turn to our main setup, in which each agent observes her realized location but not other agents' locations. For each $p \in (0, 1)$, define task allocation policy (T^p, π^p) as follows:

$T^p = \{\tau^\dagger, \tau_1^\dagger, \dots, \tau_n^\dagger\}$, and

$$\pi^p(\tau) \triangleq \begin{cases} 1 - p & \text{if } \tau = \tau^\dagger \\ \frac{p}{n} & \text{if } \tau = \tau_i^\dagger \triangleq \{t_1^\dagger, \dots, t_{n-1}^\dagger, t_i^\dagger\}, i = 1, 2, \dots, n - 1 \\ \frac{p}{n} & \text{if } \tau = \tau_n^\dagger \triangleq \{t_1^\dagger, \dots, t_{n-2}^\dagger, t_n^\dagger, t_n^\dagger\}. \end{cases} \quad (19)$$

This task allocation policy assigns all agents to their respective first-best locations, $t_1^\dagger, \dots, t_n^\dagger$, with probability $1 - p$. With probability p/n each, the allocation policy chooses τ_i^\dagger , which continues to assign agents $1, \dots, n - 1$ to their first-best locations but now assigns agent n to t_i^\dagger so that her report is used to monitor agent i . With the remaining probability p/n , the policy draws τ_n^\dagger , which assigns all agents but $n - 1$ to their first-best locations and agent $n - 1$ to t_n^\dagger , so that agent $n - 1$'s report can be used to monitor agent n . As a result, when agent $i \in \{1, \dots, n\}$ learns that her own assignment is t_i^\dagger , she is uncertain whether the realized task allocation is τ^\dagger or τ_i^\dagger , i.e., whether another agent is assigned to the same location as i . When agent $i \in \{n - 1, n\}$ is assigned to location $t_j^\dagger \neq t_i^\dagger$, she knows that she is privately appointed to monitor agent j and her report will be compared to the principal's signal.

The following result presents a contract that, combined with the task allocation policy in (19), robustly implements the desired outcome.

Proposition 4. *Suppose that there are $n \geq 2$ agents. Assume that each agent observes the task allocation policy, the compensation scheme, and the realization of her assigned location, but not the realized locations of other agents. Take any $p \in (0, 1)$ and consider a stochastic task allocation policy (T^p, π^p) defined by (19). For any $\epsilon > 0$, the following contract robustly implements work and truth-telling (RIWT): For each agent $i = 1, \dots, n - 2$,*

$$w_i^\epsilon(\tau, \chi) = \begin{cases} \frac{2nc}{p} + \epsilon & \text{if } \tau = \tau_i^\dagger, \hat{x}_i = \hat{x}_n, \\ 0 & \text{otherwise;} \end{cases}, \quad (20)$$

for agent $n - 1$,

$$w_{n-1}^\epsilon(\tau, \chi) = \begin{cases} \frac{2nc}{p} + \epsilon & \text{if } \tau = \tau_{n-1}^\dagger, \hat{x}_{n-1} = \hat{x}_n, \\ 2ce^{\lambda(t_n^\dagger - t_{n-2}^\dagger)} + \epsilon & \text{if } \tau = \tau_n^\dagger, \hat{x}_{n-1} = \hat{x}_{n-2}, \\ 0 & \text{otherwise;} \end{cases}, \quad (21)$$

and for agent n ,

$$w_n^\epsilon(\tau, \chi) = \begin{cases} \frac{2nc}{p} + \epsilon & \text{if } \tau = \tau_n^\dagger, \hat{x}_n = \hat{x}_{n-1} \\ 2ce^{\lambda(t_i^\dagger - t_{i-1}^\dagger)} + \epsilon & \text{if } \tau = \tau_i^\dagger, \hat{x}_n = \hat{x}_{i-1}, i < n. \\ 0 & \text{otherwise.} \end{cases} \quad (22)$$

The principal's incentive cost is $K^\dagger + \frac{p}{n}(2ce^{2\lambda t_1^\dagger} + 2(n-1)ce^{\lambda t_1^\dagger} - 2nc)$, as if he can perfectly verify whether each agent's report matches her assigned location state, except for the events in which he needs to provide strict incentive to the monitor.

Corollary 1. *By using the contract described in Proposition 4 and taking $(\epsilon, p) \rightarrow (0, 0)$, the principal's total loss can arbitrarily approximate $L(\tau^\dagger) + K^\dagger$.*

Proposition 4 illustrates how the principal separates diversification and incentive provision using different realizations of task allocation. For example, suppose that agent $i < n$ learns that she is assigned to her first-best location, t_i^\dagger . Agent i is uncertain whether agent n is assigned to the same location. Moreover, agent i is paid only if (i) agent n 's location is t_i^\dagger (i.e., τ_i^\dagger is realized) and (ii) her report is aligned with agent n 's. As in the previous single-agent example, the combination of probabilistic monitoring and high reward makes agents act as if they are always monitored. To solve the “who-monitor-the-monitor” problem (Alchian and Demsetz, 1972), whenever agent n is assigned to location t_i^\dagger with $i < n$, she is offered a contingent payment that rewards her only if her report is aligned with \hat{x}_{i-1} . Finally, when the realized task allocation is τ_n^\dagger , agent $n-1$ is assigned to location t_n^\dagger secretly with probability p/n to verify whether agent n 's report is aligned with the state. We plot a two-agent case in Figure 9; agent 1 monitors agent 2 in one realized task allocation, and agent 2 monitors agent 1 in another.

For each $p \in (0, 1)$ and $\epsilon > 0$, the induced game is dominance solvable. Expression (22) implies that, if the realized allocation is τ_1^\dagger , agent n finds it strictly optimal to work regardless of what other agents do. Next, agent 1 believes that, with probability p/n , the realized task allocation is τ_1^\dagger and agent n will report $\hat{x}_n = x(t_1^\dagger)$. By expression (20), agent 1 finds it strictly optimal to work, which, in turn, ensures the incentive of agent n when she is assigned to t_2^\dagger . The IESDS argument continues, ensuring the only action profile survives the iterated process is that all agents acquire information and tell the truth in each realized allocation.

Corollary 1 states that keeping agents' assignment and compensation privacy in organizations can be beneficial. However, the contract that approximates $L(\tau^\dagger) + K^\dagger$ will have to pay




Figure 9: The left panel shows the chain monitoring structure when the realized task allocation is τ_1^\dagger , where the principal’s signal is used to monitor agent 2 and agent 2’s report is used to monitor agent 1. The right panel shows the chain monitoring structure when the realized task allocation is τ_2^\dagger , where the principal’s signal is used to monitor agent 1 and agent 1’s report is used to monitor agent 2. The filled blue circles correspond to locations assigned to agents, and the filled red circle corresponds to the principal’s location.

an arbitrarily large amount of reward with an arbitrarily small probability. Thus, the extent to which the principal can take advantage of stochastic and private design depends on unmodeled components such as the principal’s liquidity constraint and agents’ attitudes towards risks.

6 Conclusion

This paper studies a robust contracting problem in which the principal chooses a task allocation and a compensation scheme to incentivize agents to acquire information as a unique outcome. A task allocation plays a dual role—it determines the kind of information the agents generate and the intensity of peer monitoring each agent faces. Upon designing a contract, the principal faces the trade-off between information diversification and monitoring efficiency. The resulting optimal contract entails a new distortion whereby the principal gives up learning about states that are ex ante novel and less-understood in order to reduce the agency cost. Moreover, to provide robust incentives, the principal adopts a simple one-sided peer monitoring. When task allocation can be private, the principal can virtually eliminate this trade-off by adopting a stochastic task allocation. The result highlights the benefit of confidential task assignment.

A Appendix: Proofs

Proof of Proposition 1. Given τ, χ , the principal’s problem is

$$\min_p \mathbb{E} \left[\int_0^1 \mathbb{1}(p(s) \neq x(s)) ds \mid \tau, \chi \right]$$

Because the integrand $\mathbb{1}(p(s) \neq x(s))$ is bounded at each t , we can ignore the zero-measure policy shifts and rewrite the problem as pointwise optimization, i.e.,

$$\min_{p(t)} \mathbb{E}[\mathbb{1}(p(t) \neq x(t)) | \tau, \chi], \forall t \in T.$$

Note that $\mathbb{E}[\mathbb{1}(p(t) \neq x(t)) | \tau, \chi] = \Pr(p(t) \neq x(t) | \tau, \chi)$. The principal chooses each $p(t) = \hat{x}_k$ for some k to minimize the pointwise loss. At each $t \in [t_{i-1}, t_i]$, for some agents $i-1$ and i (consider principal as agent 0)

$$\Pr(x(t_j) \neq x(t) | \tau, \chi) \leq \Pr(x(t_k) \neq x(t) | \tau, \chi) \iff |t - t_{i-1}| \leq |t - t_i|.$$

Then the optimal policy $p(t) = \hat{x}(t_i^*)$, where $\hat{t}_i^* = \arg \min_{\tilde{t}} \{|\tilde{t} - t_i|\}_{i=0}^n, \forall t$. Note that $p(t)$ is a piecewise linear function of t . The information loss is computed as

$$\begin{aligned} L(\tau) &= \int_0^{\frac{t_1}{2}} \Pr(\hat{x}_0 \neq x(s)) ds + \sum_{i=2}^{n-1} \int_{\frac{t_{i-1}+t_i}{2}}^{\frac{t_i+t_{i+1}}{2}} \Pr(\hat{x}_{i-1} \neq x(s)) ds + \int_{\frac{t_{n-1}+t_n}{2}}^1 \Pr(\hat{x}_{n-1} \neq x(s)) ds \\ &= \frac{1}{2} - \frac{2n+1}{2\lambda} + \frac{1}{\lambda} \left[\sum_{i=1}^n e^{-\frac{\lambda(t_i-t_{i-1})}{2}} + \frac{1}{2} e^{-\lambda(1-t_n)} \right], \end{aligned}$$

which is strictly convex and symmetric in $\Delta t_i \equiv t_i - t_{i-1}, \forall i = 1, \dots, n$. Therefore, it attains a unique minimum at the symmetric solution $\Delta t_i^\dagger = \frac{2}{2n+1}, \forall i$, or

$$t_i^\dagger = \frac{2i}{2n+1}, \forall i = 1, 2, \dots, n.$$

Then, we are able to calculate the first-best information loss as

$$L(\tau^\dagger) = \frac{1}{2} + \frac{2n+1}{2\lambda} \left(e^{-\frac{\lambda}{2n+1}} - 1 \right).$$

Denote $y^\dagger \equiv e^{-\frac{\lambda}{2n+1}} < 1$. Using envelope theorem,

$$\frac{\partial L(\tau^\dagger)}{\partial n} = \frac{1}{\lambda} (y^\dagger - 1) < 0 \text{ and } \frac{\partial L(\tau^\dagger)}{\partial \lambda} = -\frac{2n+1}{2\lambda^2} (y^\dagger - 1) > 0.$$

□

Proof of Proposition 2. For convenience, we say an agent is k^{th} -ranked in dominance order (indexed by i_k) if the agent's strict dominance incentive to acquire information and truth-report ("work") is established in the k^{th} round of IESDS. We say agent i_j is *strategically risk-free* or *ranked higher* for agent i_k if $j < k$. Given τ , we proceed in three steps. The first two focus on

the optimal payment-contingent report profiles: The 1st-ranked agent (agent i_1) is paid if and only if $\hat{x}_0 = \hat{x}_{i_1}$, i.e., her report matches the principal's; then, agent i_k is paid if and only if her report matches with her nearest strategically risk-free neighbors. The last step shows that the identity permutation is uniquely optimal among all dominance orders, i.e., $t_k = t_{i_k}, \forall k$.

Step 1: Agent i_1 is paid if and only if $\hat{x}_0 = \hat{x}_{i_1}$. The incentive cost minimization is separable across agents. Since agent i_1 find it dominant to work, her wage scheme $w_{i_1}(\chi)$ solves

$$\inf_{w_{i_1}(\cdot)} \sum_{\chi} w_{i_1}(\chi) \Pr(\chi | a_{i_1} = 1, a_{-i_1} = \mathbf{1})$$

subject to

$$\sum_{\chi} w_{i_1}(\chi) \Pr(\chi | a_{i_1} = 1, a_{-i_1}) - c > \sum_{\chi} w_{i_1}(\chi) \Pr(\chi | a_{i_1} = 0, a_{-i_1}), \forall a_{-i_1} \quad (\text{IC})$$

The number of IC constraints is the number of permutations of $a_{-i_1} \in \{0, 1\}^{n-1}$, i.e., 2^{n-1} . In all IC constraints, we assume $\hat{x}_{i_1} = x_{i_1}$ whenever $a_{i_1} = 1$, and discuss the incentive of truth-reporting shortly. To find the contract that attains the infimum incentive cost, we focus on the limit problem with weak ICs. To satisfy the IC with $a_{-i_1} = \mathbf{0}$, the payment necessarily conditions on the report profile $\hat{x}_0 = \hat{x}_{i_1}$, where \hat{x}_0 is the principal's *truthful* report. It supports i_1 's dominant incentive of truth-reporting conditioned on acquiring information. We claim $\hat{x}_0 = \hat{x}_{i_1}$ is also sufficient for paying agent i_1 , so that the other agents' reports are irrelevant for her compensation. Then agent i_1 's contract is pinned down as if there were only one agent, and the infimum of the optimal wage is

$$w_{i_1}^0(\chi) = \begin{cases} 2ce^{\lambda t_{i_1}}, & \text{if } \hat{x}_0 = \hat{x}_{i_1} \\ 0, & \text{otherwise} \end{cases}$$

solved by the binding (IC) for the one-agent problem. The incentive cost for i_1 is then

$$K_{i_1}(\tau) = \sum_{\chi} w_{i_1}^0(\chi) \Pr(\chi | a_k = 1, \forall k) = c(1 + e^{\lambda t_{i_1}}).$$

To see $w_{i_1}^0$ is indeed the optimal wage scheme, suppose agent i_1 's contract conditions on \hat{x}_j , for some agent j . All possible payment-contingent report profiles can be divided into two categories: 1) $\hat{x}_j = \hat{x}_{i_1}$; 2) $\hat{x}_j \neq \hat{x}_{i_1}$. In both cases, only j 's action affects the probability of

payment-contingent report profiles of i_1 , so the relevant ICs reduce to

$$\sum_x w_{i_1}(\chi) \Pr(\chi|a_{i_1} = 1, a_j = 1) - c \geq \sum_x w_{i_1}(\chi) \Pr(\chi|a_{i_1} = 0, a_j = 1) \quad (\text{IC}_1)$$

$$\sum_x w_{i_1}(\chi) \Pr(\chi|a_{i_1} = 1, a_j = 0) - c \geq \sum_x w_{i_1}(\chi) \Pr(\chi|a_{i_1} = 0, a_j = 0) \quad (\text{IC}_0)$$

Suppose first that $\hat{x}_j = \hat{x}_{i_1}$ is a condition of the payment-contingent report profiles. Without loss, suppose agent j reports truthfully upon acquiring information. (IC_0) implies a lower bound for the wage: $w_{i_1}(\hat{x}_0 = \hat{x}_j = \hat{x}_{i_1}) \geq c / (\frac{1}{2}(1 + e^{-\lambda t_{i_1}}) \cdot \frac{1}{2} - \frac{1}{2} \cdot \frac{1}{2}) = 4ce^{\lambda t_{i_1}}$, incurring an expected compensation of $4ce^{\lambda t_{i_1}} \cdot \Pr(\hat{x}_0 = \hat{x}_j = \hat{x}_{i_1}|a_{i_1} = 1, a_j = 1) = c(1 + e^{\lambda t_{i_1}})(1 + e^{-\lambda(t_j - t_{i_1})}) > K_{i_1}(\tau)$. Therefore, adding $\hat{x}_j = \hat{x}_{i_1}$ into agent i_1 's contract increases cost while satisfying (IC_0) . Similarly, conditioning on $\hat{x}_j \neq \hat{x}_{i_1}$ increases cost while satisfying (IC_1) . Adding more agents' report profiles into the contract of i_1 increases cost in the same manner. Generally, if the payment-contingent profiles depend on $\hat{x}_{i_1} \neq \hat{x}_j$ (resp. $\hat{x}_{i_1} = \hat{x}_j$), the constraint involving $a_j = 1$ (resp. $a_j = 0$) requires a cost increase.

Step 2: Agent i_k is paid if and only if \hat{x}_{i_k} matches the reports of her nearest neighbors who are strategically risk-free. This is a generalization of Step 1. One could regard the principal as “agent i_0 ” ranked primarily in the dominance order. Agent i_1 is paid if and only if $\hat{x}_{i_1} = \hat{x}_0$, since the principal is the only “agent” that is strategically risk-free for agent i_1 . Now, consider agent i_k for a generic k , whose wage scheme is a solution to a problem identical to agent i_1 's, except that there are only 2^{n-k} constraints, for $k-1$ strategies have been eliminated in IESDS. Agent i_k may have strategically risk-free neighbors on both “left side” and “right side” (i.e., $\exists t_{i_j}, t_{i_m}$ s.t. $t_{i_j} \leq t_{i_k} \leq t_{i_m}$ and $j, m < k$), or only on the “left side” (i.e., $\nexists m$ s.t. $m < k$ and $t_{i_m} > t_{i_k}$). In the latter case, agent i_k is paid if and only if her report matches the nearest left side risk-free neighbor, say i_j , leading to the infimum of the optimal wage

$$w_{i_k}^0(\chi) = \begin{cases} 2ce^{\lambda(t_{i_k} - t_{i_j})}, & \text{if } \hat{x}_{i_j} = \hat{x}_{i_k} \\ 0, & \text{otherwise} \end{cases}$$

and incentive cost $K_{i_k}(\tau) = c(1 + e^{\lambda(t_{i_k} - t_{i_j})})$. Using the same logic as in Step 1, the wage scheme does not condition on lower-ranked agents' reports. As a result, *at least* $n - k$ agents' reports are excluded in agent i_k 's contract. Now what remains to show is the rationale to only condition on the *nearest* risk-free neighbor. Note that the agent's compensation is a solution to a linear programming with 2^{n-k} constraints, leading to at most 2^{n-k} payment-contingent

report profiles for agent i_k , which implies excludes *at most* $n - k$ agents' reports from i_k 's contract. Combining the arguments, agent i_k 's contract is contingent on *exactly* k agents' reports, i.e., agent i_k 's report should be compared with all k agents whose dominance incentive has been established before round k . However, the principal can further simplify the contract by conditioning only on the nearest risk-free neighbors' reports, which are the *sufficient statistics* for all informative reports. To illustrate, suppose agent i_k is paid if and only if $\hat{x}_0 = \hat{x}_{i_j} = \hat{x}_{i_k}$. The principal serves as the risk-free "agent" who is *not* the nearest neighbor to agent i_k . Then the expected compensation is $P(\hat{x}_0 = \hat{x}_{i_j} = \hat{x}_{i_k} | a = 1) \cdot 4ce^{\lambda(t_{i_k} - t_{i_j})} / (1 + e^{-\lambda t_j}) = c(1 + e^{\lambda(t_{i_k} - t_{i_j})})$, which is the same as $K_{i_k}(\tau)$. When i_k has nearest risk-free neighbors on both sides, the same logic follows through. We now characterize the optimal contract. Denote the nearest risk-free neighbors of i_k to be i_j and i_m , respectively. Then the relevant IC is

$$\sum_{\chi \in \mathcal{X}} w_{i_k}(\chi) [\Pr(\chi | a_{i_j} = 1, a_{i_k} = 1, a_{i_m} = 1) - \Pr(\chi | a_{i_j} = 1, a_{i_k} = 0, a_{i_m} = 1)] \geq c$$

where \mathcal{X} is the set of all combinations of the informative report profiles: $(\hat{x}_{i_j}, \hat{x}_{i_k}, \hat{x}_{i_m})$. The problem boils down to finding the report profiles that maximizes the likelihood ratio $\Pr(\cdot | a_{i_k} = 1) / \Pr(\cdot | a_{i_k} = 0)$, and the profile is $\hat{x}_{i_j} = \hat{x}_{i_k} = \hat{x}_{i_m}$. Both i_j and i_m have dominance incentive to work, which ensures i_k 's dominance incentive to truthfully report upon acquiring information. The binding IC implies

$$w_{i_k}^0(\chi) = \begin{cases} \frac{4c}{e^{-\lambda(t_{i_k} - t_{i_j})} + e^{-\lambda(t_{i_m} - t_{i_k})}} & \text{if } \hat{x}_{i_j} = \hat{x}_{i_k} = \hat{x}_{i_m} \\ 0 & \text{otherwise} \end{cases}$$

and the corresponding incentive cost is

$$K_{i_k}(\tau) = \frac{1}{4}(1 + e^{-\lambda(t_{i_k} - t_{i_j})})(1 + e^{-\lambda(t_{i_m} - t_{i_k})}) \cdot w_{i_k}^0(\hat{x}_{i_j} = \hat{x}_{i_k} = \hat{x}_{i_m}) = c \left[1 + \frac{1 + e^{-\lambda(t_{i_m} - t_{i_j})}}{e^{-\lambda(t_{i_k} - t_{i_j})} + e^{-\lambda(t_{i_m} - t_{i_k})}} \right].$$

Step 3: The identity permutation is optimal for dominance order. Given the task allocation τ , we are to show that the identity permutation (placing agent k as k^{th} -ranked) is the optimal dominance order. Starting from $n = 2$, there are only two permutations: agent 1 is 1st-ranked ("chain monitoring") or 2nd-ranked ("sandwich monitoring"). The incentive cost under the two monitoring structures are respectively:

$$K^c(t_1, t_2) = c[1 + e^{\lambda t_1}] + c[1 + e^{\lambda(t_2 - t_1)}], \quad K^s(t_1, t_2) = c[1 + e^{\lambda t_2}] + c \left[1 + \frac{1 + e^{-\lambda t_2}}{e^{-\lambda t_1} + e^{-\lambda(t_2 - t_1)}} \right].$$

The two incentive costs are the same when $t_1 = 0, t_2 > 0$. Moreover, with fixed t_2 , K^s increases as the distances among agents become more symmetric, driven by the denominator of the second term, while K^c decreases. Thus $K^s(t_1, t_2) \geq K^s(0, t_2) = K^c(0, t_2) \geq K^s(t_1, t_2), \forall t_1, t_2$. Therefore, chain monitoring always dominates sandwich monitoring in cost efficiency.

To apply the logic to $n \geq 3$ agents, we need to identify agent i_1 , who is 1st-ranked in the dominance order. Call the monitoring regime to be $\mathcal{M}^{c,n}$ “the general chain monitoring” if $i_1 = 1$, to be $\mathcal{M}^{s,n}$ “the general sandwich monitoring” if $i_1 = n$. Consider the case that $i_1 = n$. Similar to $n = 2$, we obtain $K^{s,n}(\tau) \geq K^{s,n}(t_n, \tau_{-n} = 0) = K^{c,n}(t_n, \tau_{-n} = 0) \geq K^{c,n}(\tau), \forall \tau$, which implies that $\mathcal{M}^{s,n}$ is dominated by $\mathcal{M}^{c,n}$. The inequalities on both sides are established by increasing $\Delta t_k \equiv t_k - t_{k-1}$ in the *dominance order*. E.g., starting from any $t_n > 0, \tau_{-n} = 0$, an increase in Δt_{i_2} reduces $K^{c,n}$ but increases $K^{s,n}$; then, fixing any $\Delta t_n > 0, \Delta t_{i_2} > 0$ and all other $\Delta t_k = 0$, an increase in Δt_{i_3} reduces $K^{c,n}$ but increases $K^{s,n}$ further. This process goes on until all possible τ are enumerated. If $i_1 < n$, we can split the agents into two parts: for $k \leq i_1$, we are back to \mathcal{M}^{s,i_1} , and \mathcal{M}^{c,i_1} is better; for $k \geq i_1$, the best monitoring structure is again $\mathcal{M}^{c,n-i_1+1}$. Hence, $\mathcal{M}^{c,n}$ is dominant for any n .

As a final remark, it is without loss to study the IESDS that only removes *one* agent’s strategy in each step. To see this, suppose both agent k_1 ’s and k_2 ’s dominant incentive is established (where $k_1 < k_2$) in some round k . Consider a “local improvement:” remove k_2 ’s strategy at round $k + 1$ instead of round k , and also postpone by one round for the strategy removal of all other agents after round k . There are two cases. If some strategically risk-free agent is located between k_1 and k_2 , then the incentive cost remains unchanged. Otherwise, the local improvement is strict because k_2 is more closely monitored (by k_1 and her right neighbor). Hence, removing multiple agents’ strategies in one round of IESDS may waste monitoring power.

In conclusion, (7) robustly implements work and truth-telling, for all $\epsilon > 0$. \square

Proof of Proposition 3. Thanks to Proposition 1 and 2, the task allocation problem is

$$\min_{\tau} \frac{1}{2} - \frac{2n+1}{2\lambda} + \frac{1}{\lambda} \left[\sum_{i=1}^n e^{-\frac{\lambda \Delta t_i}{2}} + \frac{1}{2} e^{-\lambda(1-\sum_{i=1}^n \Delta t_i)} \right] + c \sum_{i=1}^n [1 + e^{\lambda \Delta t_i}]$$

which is strictly convex and symmetric in the choice variables $\Delta t_i, \forall i$. Therefore, the objective attains minimum at the unique symmetric solution to the first-order conditions

$$x^{2(n-1)} + 2\lambda e^{\lambda} c = \frac{e^{\lambda}}{x^3}.$$

where $x \equiv e^{\frac{\lambda \Delta t^*}{2}} \geq 1$, $\Delta t^* = \Delta t_i$. Implicitly differentiating the optimal condition with respect to c and n , we find that Δt^* decreases in c and n when it is interior. However, the above characterization ignores the primal constraints $\tau \in \{(t_1, \dots, t_n) \in \mathbb{R}_+^n : t_1 \leq \dots \leq t_n\}$. At the boundary, $\Delta t^* = 0$ implies $c = \frac{1-e^{-\lambda}}{2\lambda} \equiv \bar{c}$. To sum up, if $c \geq \bar{c}$, $\Delta t^* = 0$. If $c < \bar{c}$, the solution is interior satisfying the above optimal condition, with $\Delta t^* = \frac{2}{\lambda} \log x > 0$. \square

Proof of Proposition 4. For agent i with $i = 1, 2, \dots, n-2$, her incentive cost is determined by the following linear programming.

$$\inf_{w_i(\tau_i^\dagger), w_i(\tau \neq \tau_i^\dagger)} \frac{p}{n} w_i(\tau_i^\dagger) + \left(1 - \frac{p}{n}\right) \frac{1}{2} (1 + e^{-\lambda(t_i - t_{i-1})}) w_i(\tau \neq \tau_i^\dagger)$$

subject to

$$\frac{p}{n} w_i(\tau_i^\dagger) + \left(1 - \frac{p}{n}\right) \frac{1}{2} (1 + e^{-\lambda(t_i - t_{i-1})}) w_i(\tau \neq \tau_i^\dagger) - c > \frac{p}{n} \cdot \frac{1}{2} w_i(\tau_i^\dagger) + \left(1 - \frac{p}{n}\right) \frac{1}{2} w_i(\tau \neq \tau_i^\dagger)$$

which leads to (20). In this construction, agent 1's dominance incentive is established in the *second* round of IESDS²⁰, after agent n with τ_1^\dagger . Moreover, each agent i 's dominance incentive is established in the $2i$ -th round, immediately after agent n with τ_i^\dagger , for all $i = 2, \dots, n-2$. Agent $n-1$ has two private types, the “monitor” type (when $\tau = \tau_n^\dagger$) and the “worker” type (when $\tau \neq \tau_n^\dagger$). By construction, when τ_n^\dagger is realized, agent n 's dominance incentive is established *after* $n-1$. Therefore, the monitor-type agent $n-1$ has a simple incentive cost under chain monitoring, and gets paid if and only if $\hat{x}_{n-1} = \hat{x}_{n-2}$ despite being located at t_n . The worker-type agent $n-1$ has the same incentive cost as agent $i = 1, 2, \dots, n-2$. Combining both types, agent $n-1$'s wage scheme satisfies (21). Finally, agent n has $n-1$ (private) monitor types (when $\tau \neq \tau_n^\dagger$) and one worker type (when $\tau = \tau_n^\dagger$). In each monitor type with $\tau = \tau_i^\dagger$, agent n is paid if and only if her report matches with agent $i-1$, whose dominance incentive is established one round before. Particularly, in the monitor type with $\tau = \tau_1^\dagger$, agent n 's report matches with the principal, establishing her dominance incentive in the *first* round of IESDS. In the worker type, agent n has the same incentive cost as agent $i = 1, 2, \dots, n-2$. Combining all types, agent n 's wage scheme satisfies (22). Thus, the total expected incentive cost is

$$\lim_{\epsilon \rightarrow 0} \sum_{i=1}^n \frac{p}{n} \cdot \left(\frac{2nc}{p} + \epsilon \right) - 2pc + \frac{p}{n} \cdot [2ce^{2\lambda t_1^\dagger} + \epsilon] + \frac{n-1}{n} p \cdot [2ce^{\lambda t_1^\dagger} + \epsilon] = K^\dagger + \frac{p}{n} (2ce^{2\lambda t_1^\dagger} + 2(n-1)ce^{\lambda t_1^\dagger} - 2nc)$$

The expression involving t_1^\dagger follows from $t_1^\dagger = t_i^\dagger - t_{i-1}^\dagger, \forall i$. \square

²⁰We consider the agent normal form in the IESDS process, i.e., different types of an agent are treated as many agents.

Proof of Corollary 1. The result follows directly from Proposition 4 by setting $p \approx 0$. \square

B Appendix: Partial-Implementation Contract

This section characterizes the incentive scheme under partial implementation, i.e., inducing all agents to work as a Nash equilibrium. The partial-implementation contract design problem is

$$\min_{\{w_i(\chi)\}_{i=1}^n} \sum_{i=1}^n \left[\sum_{\chi} w_i(\chi) \Pr(\chi | a_i = 1, a_{-i} = \mathbf{1}) \right]$$

subject to

$$\sum_{\chi} w_i(\chi) \Pr(\chi | a_i = 1, a_{-i} = \mathbf{1}) - c \geq \sum_{\chi} w_i(\chi) \Pr(\chi | a_i = 0, a_{-i} = \mathbf{1}) \quad (\text{P-IC})$$

for each $i = 1, \dots, n$. Constraint (P-IC) requires agent i has the incentive to acquire information and report truthfully if others will do so. Let

$$K_i^s(t_{i-1}, t_i, t_{i+1}) \equiv c \left[1 + \frac{1 + e^{-\lambda(t_{i+1}-t_{i-1})}}{e^{-\lambda(t_i-t_{i-1})} + e^{-\lambda(t_{i+1}-t_i)}} \right]$$

be the incentive cost of agent i sandwiched by her neighbors. Denote $K_n^s(t_{n-1}, t_n) \equiv c(1 + e^{\lambda(t_n-t_{n-1})})$ be the incentive cost of agent n .

Proposition B.1. *Fix a task allocation τ and $n \geq 2$, the incentive cost under partial implementation (with mutual monitoring) is*

$$K^{PI}(\tau) \equiv \sum_{i=1}^{n-1} K_i^s(t_{i-1}, t_i, t_{i+1}) + K_n^s(t_{n-1}, t_n), \quad (23)$$

attained by setting

$$w_j(\chi) = \begin{cases} \frac{4c}{e^{-\lambda(t_j-t_{j-1})} + e^{-\lambda(t_{j+1}-t_j)}} & \text{if } \hat{x}_{j-1} = \hat{x}_j = \hat{x}_{j+1}, \\ 0, & \text{otherwise} \end{cases},$$

for $j = 1, 2, \dots, n-1$, and

$$w_n(\chi) = \begin{cases} 2ce^{\lambda(t_n-t_{n-1})} & \text{if } \hat{x}_{n-1} = \hat{x}_n \\ 0, & \text{otherwise} \end{cases}.$$

Under the optimal partial-implementation contract, each agent $i < n$ is monitored by her left and right neighbors' reports and agent n is monitored by her left neighbor's report. Recall by the informativeness principle, an agent's payment should only be contingent on sufficient statistics, which is the truthful reports of her immediate neighbors. As a result, the optimal payment of each agent makes her indifferent to follow the desirable action (acquire information and tell the truth) given other agents do so. The result immediately follows.

C Supplementary Material

C.1 Heterogeneous Costs of Information Acquisition

The optimal monitoring structure will change when agents have different costs of information acquisition. To illustrate this point, we assume that there are two agents, where one agent has cost c^H and the other agent has cost $c^L \in (0, c^H)$. The principal's task allocation specifies (i) the locations of two tasks, $\tau = (t_1, t_2)$ with $t_1 \leq t_2$ and (ii) a who-does-what choice, i.e.,

$$\mathbf{c} \triangleq (c_1, c_2) \in \{(c^L, c^H), (c^H, c^L)\}.$$

Here, c_i denotes the cost of the agent assigned to location t_i . For example, $\mathbf{c} = (c^H, c^L)$ means that the high-cost agent works on task t_1 , which is closer to the principal.

The principal's problem is to choose a task allocation (τ, \mathbf{c}) to minimize his total loss, i.e.,

$$\min_{\tau, \mathbf{c}} L(\tau) + K(\tau, \mathbf{c}).$$

For each (τ, \mathbf{c}) , information loss $L(\tau)$ is still given by equation (4), but the incentive cost is

$$K(\tau, \mathbf{c}) = \min\{K^c(\tau, \mathbf{c}), K^s(\tau, \mathbf{c})\},$$

where

$$K^c(\tau, \mathbf{c}) = \sum_{i=1,2} c_i [1 + e^{\lambda(t_i - t_{i-1})}] \quad \text{and,}$$

$$K^s(\tau, \mathbf{c}) = c_1 [1 + e^{\lambda t_2}] + c_2 \left[1 + \frac{1 + e^{-\lambda t_2}}{e^{-\lambda t_1} + e^{-\lambda(t_2 - t_1)}} \right],$$

correspond to the incentives cost under the chain and sandwich monitoring structures, respectively (recall Figure 4).

To proceed further, we first prove that the principal can without loss of generality focus on task allocation such that the high-cost agent works on a task closer to the principal's location.

Lemma C.1. *The principal finds it weakly optimal to set $\mathbf{c} = \mathbf{c}^* = (c^H, c^L)$, i.e.,*

$$\min_{\tau} L(\tau) + K(\tau, (c^H, c^L)) \geq \min_{\tau} L(\tau) + K(\tau, (c^L, c^H)).$$

Moreover, the high-cost agent is more closely monitored than the low-cost agent in the sense of monitoring intensity, i.e.,

$$\frac{c^H}{K_1(\tau)} \geq \frac{c^L}{K_2(\tau)}, \forall \tau.$$

Proof of Lemma C.1. We show that it is always weakly optimal to set \mathbf{c}^* for each monitoring structure. Fix the chain monitoring structure, and fix any task allocation $\tau = (t_1, t_2)$. Suppose the arrangement is (c^L, c^H) , shifting the arrangement to (c^H, c^L) and $\tau' = (t_2 - t_1, t_2)$ keeps the principal's loss unchanged. Thus, it is always weakly optimal to set \mathbf{c}^* . Moreover, the high cost agent is more closely monitored. To see this, since the problem is sufficiently and necessarily characterized by the optimal conditions

$$\begin{aligned} e^{-\lambda(1-t_2)} - e^{-\frac{\lambda t_1}{2}} + 2\lambda c_1 e^{\lambda t_1} &= 0 \\ e^{-\lambda(1-t_2)} - e^{-\frac{\lambda(t_2-t_1)}{2}} + 2\lambda c_2 e^{\lambda(t_2-t_1)} &= 0 \end{aligned}$$

Cancelling out the term $e^{-\lambda(1-t_2)}$ in the optimal conditions yields

$$2\lambda(c_1 e^{\lambda t_1} - c_2 e^{\lambda(t_2-t_1)}) = e^{-\frac{\lambda t_1}{2}} - e^{-\frac{\lambda(t_2-t_1)}{2}}$$

Suppose, for the sake of contradiction, that agent 2 is more closely monitored, i.e., $t_1 > t_2 - t_1$. Then, the left-hand side of above equation is strictly positive, but right-hand side is weakly negative, a contradiction. Therefore, it must be the case that $t_1 \leq t_2 - t_1$, i.e., agent 1 is more closely monitored. Formally,

$$e^{\lambda t_1} \leq e^{\lambda(t_2-t_1)} \Rightarrow \frac{K_1^c(\tau)}{c^H} \leq \frac{K_2^c(\tau)}{c^L}$$

where the inequality only holds at the corner $t_1 = t_2 - t_1 = 0$.

On the other hand, fix the sandwich monitoring structure. The problem of who-does-what arrangement becomes

$$\inf_{(c_1, c_2) \in \{(c^H, c^L), (c^L, c^H)\}} L(\tau) + \left(1 + \frac{1 + e^{-\lambda t_2}}{e^{-\lambda t_1} + e^{-\lambda(t_2-t_1)}}\right) c_1 + (1 + e^{\lambda t_2}) c_2$$

Observe that

$$1 + \frac{1 + e^{-\lambda t_2}}{e^{-\lambda t_1} + e^{-\lambda(t_2-t_1)}} \leq 1 + e^{\lambda t_2},$$

and the inequality is strict if at least one of t_1 and $t_2 - t_1$ is strictly positive. Therefore, it is optimal to choose $(c_1, c_2) = \mathbf{c}^*$. This is true for all τ , and thus is true for the optimal τ . The high-cost agent is sandwiched, so that she is paid less per unit effort cost than the low-cost agent. Formally,

$$\frac{1 + e^{-\lambda t_2}}{e^{-\lambda t_1} + e^{-\lambda(t_2-t_1)}} \leq e^{\lambda t_2} \Rightarrow \frac{K_1^s(\tau)}{c^H} \leq \frac{K_2^s(\tau)}{c^L}$$

□

Lemma C.1 says that the principal always prefers keeping the high-cost agent “close” with greater monitoring intensity. This is intuitive. An agent’s minimum expected payment is the product of her information acquisition cost and the reciprocal of monitoring intensity. For each task allocation τ , to minimize the total incentive cost, it is optimal to assign the high-cost agent a location under more intense monitoring.

In what follows, we consider the effect of increasing the agents’ cost asymmetry while keeping the average cost. Specifically, fix $c > 0$ and let $c^H = c + \Delta c$ and $c^L = c - \Delta c$. A higher Δc means the higher cost asymmetry between the agents.

Proposition C.1. *As the agents’ cost disparity (Δc) increases, the principal is better off. Moreover, with sufficiently large cost disparity, the optimal design induces a sandwich monitoring structure, i.e., there exists $\widehat{\Delta c} < c$ such that for any $\Delta c > \widehat{\Delta c}$, we have*

$$\min_{\tau} L(\tau) + K^c(\tau, \mathbf{c}^*) \geq \min_{\tau} L(\tau) + K^s(\tau, \mathbf{c}^*),$$

and the inequality is strict if the optimal task allocation problem under the sandwich monitoring structure has an interior solution, i.e., $\exists \tau^s \in \arg \min_{\tau} L(\tau) + K^s(\tau, \mathbf{c}^*)$ and $t_2^s > t_1^s > 0$.

Proof of Proposition C.1. The first part of Proposition C.1 claims that the principal benefits from cost disparity. We show that this is true fixing each monitoring structure. First note that the loss function of the principal is

$$\mathcal{L} = \min_{\tau} L(\tau) + K(\tau)$$

By envelope theorem,

$$\frac{\partial \mathcal{L}}{\partial \Delta c} = \underbrace{\frac{\partial L(\tau^*)}{\partial \Delta c}}_{=0} + \frac{\partial K(\tau^*)}{\partial \Delta c} \quad (24)$$

First, fixing the chain monitoring structure, (24) becomes

$$\frac{\partial K^c(\tau^*)}{\partial \Delta c} = e^{\lambda t_1^*} - e^{\lambda(t_2^* - t_1^*)} \leq 0$$

by Lemma C.1. Hence, the loss decreases in Δc .

On the other hand, fixing the sandwich monitoring structure. Then (24) becomes

$$\frac{\partial K^s(\tau^*)}{\partial \Delta c} = \frac{1 + e^{-\lambda t_2^*}}{e^{-\lambda t_1^*} + e^{-\lambda(t_2^* - t_1^*)}} - e^{\lambda t_2^*} \leq 0,$$

as a result of Lemma C.1. Hence, the loss decreases in Δc . In conclusion, the principal's loss decreases as Δc increases even with flexible choice of monitoring structure.

To prove the second part of Proposition C.1, we characterize the condition for $K^c(\tau) \geq K^s(\tau), \forall \tau$, in which case the sandwich monitoring structure cost dominates the chain monitoring structure given any task allocation, implying the sandwich monitoring structure is optimal. To see this, with straightforward algebra, we can show that both $K^c(t_1; t_2)$ and $K^c(t_1; t_2) - K^s(t_1; t_2)$ are convex in t_1 . Note that $K^c(0; t_2) = K^s(0; t_2), \forall t_2$. Then it suffices to find a range of Δc such that $\frac{\partial K^c(0; t_2)}{\partial t_1} \geq \frac{\partial K^s(0; t_2)}{\partial t_1}, \forall t_2$, which boils down to

$$\Delta c \geq \widehat{\Delta c} \equiv \left(1 - \frac{4}{e^{\frac{8}{5}\lambda} + e^{\frac{4}{5}\lambda} + 2} \right) c.$$

Note that $\widehat{\Delta c}(\lambda)$ increases in λ . With smaller λ , we can guarantee a larger range of Δc for sandwich monitoring to be optimal. In particular, when $\lambda \rightarrow 0, \widehat{\Delta c}(\lambda) \rightarrow 0$. \square

Proposition C.1 says that the principal benefits from increasing agents' cost disparity. The logic is simple. Fix the task allocation (τ^*, \mathbf{c}^*) , consider a small increase in agents' cost disparity by $\epsilon > 0$. The corresponding impact on the principal's total loss can be approximated by

$$\epsilon \left[\frac{K_1(\tau^*)}{c^H} - \frac{K_2(\tau^*)}{c^L} \right] \leq 0.$$

The inequality is a consequence of Lemma C.1. Intuitively, each agent's optimal expected payment is her information acquisition cost divided by the monitoring intensity. Raising the cost disparity increases high-cost agent's expected payment but decreases the low-cost agent's payment, but the because high-cost agent's monitoring intensity is also higher than the low-cost agent, the total effect decreases the incentive cost.

The second part of Proposition C.1 says when the agents' cost dispersion is large enough, the chain monitoring structure can be dominated by the sandwich structure. The optimality of the sandwich structure can be demonstrated by Figure 10, which compares the incentive costs under the two monitoring structure by fixing t_2 and varying t_1 . With sandwich monitoring, the incentive cost $K^s(\cdot; t_2, \mathbf{c}^*)$ (blue curve) is still symmetric in t_1 as the homogeneous agent




Figure 10: Illustrating expected payment function $K^s(\cdot, t_2, \mathbf{c}^*)$ and $K^c(\cdot, t_2, \mathbf{c}^*)$.

case. This is because changing t_1 solely impacts agent 1's expected compensation, which under sandwich monitoring, depends on t_1 and $t_2 - t_1$ in a symmetric manner according to equation (13).

On the other hand, unlike the homogeneous agent case, the incentive cost with chain monitoring $K^c(\cdot; t_2, \mathbf{c}^*)$ (red curve) is asymmetric in t_1 . The reason is that varying t_1 impacts compensation for agent 1 and agent 2 differently: When fixing t_2 while increasing t_1 from 0 to t_2 , it becomes cheaper to compensate agent 2, and more costly to compensate agent 1. Crucial to the asymmetry, the effort cost of agent 1 is higher than that of agent 2. As t_1 gets close to t_2 , agent 1 is less closely monitored, leading an incentive cost ($\approx c^H(1 + e^{\lambda t_2}) + 2c^L$) to be higher than the case where t_1 is close to 0 ($\approx 2c^H + c^L(1 + e^{\lambda t_2})$). Moreover, note that the cost disparity (Δc) determines the speed of increment of $K^c(\cdot; t_2, \mathbf{c}^*)$. To see this, it is helpful to study the marginal effect of t_1 on incentive cost:

$$\frac{\partial K^c(t_1; t_2, \mathbf{c}^*)}{\partial t_1} = \lambda[(c + \Delta c)e^{\lambda t_1} - (c - \Delta c)e^{\lambda(t_2 - t_1)}]$$

which is increasing in Δc . Therefore, when the cost disparity (Δc) is large enough, the cost increase in agent 1 (the first term on the right-hand side) dominates the cost saving in agent 2 (the second term), making $K^c(\cdot, t_2, \mathbf{c}^*)$ strictly increases. It is intuitive that when K^s 's slope is sufficiently large, which is guaranteed when Δc is sufficiently large, K^s dominates K^c globally. In this case, suppose there is a minimal total loss achieved by the chain monitoring structure, say, $\min_{\tau} L(\tau) + K^c(\tau, \mathbf{c}^*)$. The principal can weakly benefit by shifting to the sandwich monitoring while keeping τ unchanged. This arrangement preserves the information loss, while lowering the incentive cost since $K^s(\tau, \mathbf{c}^*) \leq K^c(\tau, \mathbf{c}^*)$.

We generalize this result to n agents within a limited extent. When the cost dispersion is large enough, the optimal monitoring structure does not contain a local chain, i.e., all agents except for n are monitored by two peers. The formal result is upon request.

C.2 Circular Decision Environment

This subsection provides an alternative way to model the principal's decision environment. The main insight and the tractability of our baseline model are preserved.

We arrange locations uniformly on a circle with a radius of $1/(2\pi)$ rather than an interval. This ensures an even density of locations. We set the principal's location as the reference point, marked as 0, and label the other locations in a clockwise direction from this point. The distance between any two points, t and t' , is measured as the shortest path along the circle's circumference, i.e.,

$$d(t, t') = \min\{|t - t'|, |t + 1 - t'|\}.$$

The state is represented by a realized path of a continuous-time Markov chain as in the baseline model, with the condition that $x(0)$ equals $x(1)$. Specifically, $x(0)$ is randomly selected to be either state A or B . As t progresses, a shock arrives at rate λ , and when a shock occurs at location t , the local state $x(t)$ is determined by a new random draw, either A or B , randomly chosen. The only twist from the baseline model is that the local states at $t = 0$ and $t = 1$ must be the same. By simple algebra, the probability $\Pr[x(t) = x(0) | x(0) = x(1)]$ takes the following formula

$$\frac{1}{2} \left(1 + \frac{e^{-\lambda t} + e^{-\lambda(1-t)}}{1 + e^{-\lambda}} \right).$$

This probability is symmetric, meaning it is the same for t and $1 - t$, and it approaches zero as t nears 0 or 1. For any $t \in [0, 1/2]$, this probability is a decreasing and convex function. The midpoint, $t = 1/2$, represents the most "novel" or least understood location from the principal's perspective. The location circle is illustrated in Figure 11. The left panel illustrates information loss minimizing allocation, i.e., $t_1^\dagger = 1/3, t_2^\dagger = 2/3$. The right panel illustrates two optimal task allocations. One solution has $t_1^* \in (0, 1/2)$ and $t_2^* \in (t_1^*, 1/2 + t_1^*/2)$ and another solution has t_1^* and $t_2^{**} > 1/2 + t_1^*/2$ where $d(t_1^*, t_2^*) = d(0, t_2^{**})$.

In what follows, we assume $n = 2$ for simplicity. It is straightforward that given τ , the principal's optimal policy remains the same form as in Proposition 1, and the task allocation




Figure 11: The location circle and task allocation.

minimizing the information loss is such that

$$t_1^* = \frac{1}{3}, t_2^* = \frac{2}{3},$$

exhibiting ample diversification. Moreover, given τ , the optimal compensation scheme is to pay whoever has the closest distance to the principal if and only if her report is aligned with the principal's report. Without loss of generality, assume $t_1 < t_2, 1 - t_2$, and so agent 1 is paid by w_1 if and only if $x_0 = \hat{x}_1$ regardless of what agent 2 does. The optimal w_1 is set to make agent 1 indifferent. After establishing agent 1's incentive, agent 2 can be jointly monitored by the principal and agent 1's report, and agent 2 is paid by w_2 if and only if $x_0 = \hat{x}_1 = \hat{x}_2$, and w_2 is set to satisfy agent 2's indifference condition. The optimal task allocations is not unique. For instance, as illustrated by the right panel of Figure 11, if (t_1^*, t_2^*) is an optimal allocation, it is also optimal to assign (t_1^*, t_2^{**}) where

$$d(0, t_2^{**}) = d(t_1^*, t_2^*).$$

However, unlike in the benchmark model, the agents' locations may not be equidistant in general.²¹

To understand the intuition, notice that in the circle specification, an agent is always monitored by a sandwich structure. In the first-step of the IESDS, agent 1 is monitored by the principal but as both her left and right neighbors, and she is paid if and only if $x(0) = x(t_1) = x(1)$.

²¹It is difficult to solve optimal allocation in closed form, but numerical analysis reveals that in general, $d(0, t_1) = d(t_1, t_2)$ is suboptimal. The result is upon request.



Figure 12: Illustration of the multiplicity.

In the second step of IESDS, agent 2 is jointly monitored by the principal from the left-hand side and agent 1 from the right-hand side, and she is paid if and only if $\hat{x}_1 = \hat{x}_2 = x(0)$. This also specifies a sandwich structure. See Figure 12 for an illustration. Notice that whenever it is optimal to assign agent 2 to work at location t_2^* , it must be optimal to assign her to work at location t_2^{**} by symmetry.

C.3 Imperfect Information Acquisition

In this section, suppose each agent (conditioned on acquiring information) observes an imperfect binary signal from $\{A, B\}$, which matches the true state with probability $q \in (0.5, 1)$, i.i.d. across agents. For symmetry, assume the principal's information is imperfect: his signal also matches the true state with probability q .

Given any two *truthful* reports $\hat{x}(t')$, $\hat{x}(t)$ from two agents located at t' and t ($t' > t$), the critical probability that governs agents' incentive is

$$\begin{aligned} \Pr(\hat{x}(t') = \hat{x}(t)) &= \underbrace{\frac{1}{2}(1 + e^{-\lambda t_1})}_{\Pr(x(t')=x(t))} \cdot [q^2 + (1 - q)^2] + \underbrace{\frac{1}{2}(1 - e^{-\lambda t_1})}_{\Pr(x(t') \neq x(t))} \cdot 2q(1 - q) \\ &= \left(\frac{1}{2} - 2q(1 - q) \right) (1 + e^{-\lambda t_1}) + 2q(1 - q) \end{aligned}$$

On Sufficient Statistics. Note that the report of the nearest strategically risk-free neighbor is no longer a sufficient statistic monitoring an agent. Consider $n = 2$, suppose agent 1's dominance incentive is established in the first round of IESDS, and agent 2 the second (as in the chain monitoring). Suppose agent 2 is paid if and only if $\hat{x}_1 = \hat{x}_2$, then her binding IC constraint is

$$\left[\left(\frac{1}{2} - 2q(1 - q) \right) (1 + e^{-\lambda(t_2 - t_1)}) + 2q(1 - q) - \frac{1}{2} \right] w_2 = c,$$

which gives

$$w_2 = \frac{ce^{\lambda(t_2 - t_1)}}{\frac{1}{2} - 2q(1 - q)}.$$

The incentive cost is then

$$\begin{aligned} & \left[\left(\frac{1}{2} - 2q(1-q) \right) (1 + e^{-\lambda(t_2-t_1)}) + 2q(1-q) \right] \frac{ce^{\lambda(t_2-t_1)}}{\frac{1}{2} - 2q(1-q)} \\ & = c(1 + e^{\lambda(t_2-t_1)}) + \frac{2q(1-q)ce^{\lambda(t_2-t_1)}}{\frac{1}{2} - 2q(1-q)}, \end{aligned}$$

where the additional term, relative to the benchmark incentive cost, comes from the imperfect observation of agents. Now we show that utilizing an additional report \hat{x}_0 reduces the incentive cost. Suppose agent 2 is paid if and only if $\hat{x}_0 = \hat{x}_1 = \hat{x}_2$. Again, assume $a_1 = 1$. Then

$$\Pr(\hat{x}_0 = \hat{x}_1 = \hat{x}_2 | a_2 = 1) = \left[\frac{1}{4} - q(1-q) \right] (1 + e^{-\lambda t_1})(1 + e^{-\lambda(t_2-t_1)}) + q(1-q).$$

Using binding IC of agent 2, we get

$$w_2 = \frac{ce^{\lambda(t_2-t_1)}}{\left(\frac{1}{4} - q(1-q) \right) (1 + e^{-\lambda t_1})}.$$

The incentive cost is

$$K_2(\tau) = c(1 + e^{\lambda(t_2-t_1)}) + \frac{2q(1-q)ce^{\lambda(t_2-t_1)}}{\left(\frac{1}{2} - 2q(1-q) \right)} \frac{1}{(1 + e^{-\lambda t_1})},$$

which is less than the incentive cost based on (\hat{x}_1, \hat{x}_2) . When $q \rightarrow 1$, the second term vanishes and $K_2(\tau)$ converges to the benchmark incentive cost. Note that agent 2's compensation is determined by a linear programming problem of one constraint, which implies that there should be exactly one payment-contingent report. This report is $\hat{x}_0 = \hat{x}_1 = \hat{x}_2$, providing the maximal incentive to work for agent 2. This is different from our main analysis, because the nearest strategically risk-free neighbor's report is no longer a sufficient statistic.

Following the same logic, with generic n , an agent at incentive rank k must be paid if and only if her report match with those whose dominance incentive has been established in previous rounds of IESDS.

Proposition C.2. *Fixing any task allocation τ and imperfect learning parameter $q \in (0.5, 1)$, a contract that attains the minimal incentive cost must specify each agent to be paid if and only if her report matches with those whose dominance incentive has been established in the previous rounds of IESDS.*

Proof. We use a similar argument as the counterpart in the proof of Proposition 2 where $q = 1$. Given task allocation τ and a fixed dominance order, consider a generic agent whose dominance

incentive is established in round k of the IESDS. We first argue that this agent (denote by i_k) is paid upon at most 2^{n-k} report profiles, which is equivalent to excluding report profiles of $n - k$ agents from agent i_k 's contract. Then, we show the $n - k$ agents are exactly the agents whose dominance incentive has not been established before round k of the IESDS.

Step 1: Agent i_k is paid upon at most 2^{n-k} report profiles. Agent i_k 's incentive cost is determined by a linear programming with 2^{n-k} IC constraints:

$$\min_{w_1(\cdot)} \sum_{\chi} w_1(\chi) \Pr(\chi | a_{i_1} = 1, a_{-i_1} = 1)$$

subject to

$$\sum_{\chi} w_{i_k}(\chi) \Pr(\chi | a_{i_k} = 1, a_{-i_k}) - c \geq \sum_{\chi} w_{i_k}(\chi) \Pr(\chi | a_{i_k} = 0, a_{-i_k}), \forall a_{-i_k} \quad (\text{IC})$$

The number 2^{n-k} comes from the number of combinations of a_{-i_k} : at round k of IESDS, there are $k - 1$ agents whose dominance incentive has been established, and each of them only has a strategy $a = 1$; the remaining $n - k$ agents each has two strategies. By the Fundamental Theorem of Linear Programming, the above problem has a basic feasible solution with at most 2^{n-k} positive entries. In other words, there are at most 2^{n-k} report profiles generating positive payment for agent i_k , which requires at most $n - k$ other sources (agent or principal) of report to be excluded in the contract. To see this, consider an example with $k = 2$. There are at most 2^{n-2} report profiles generating positive payment for agent i_2 . Suppose agent i_2 is paid if and only if $\hat{x}_0 = \hat{x}_{i_2}$, then we have excluded $n - 1$ other sources of report from agent i_2 's contract. Then there are $2^{n-1} > 2^{n-2}$ report profiles generating positive payment, violating the previous claim. Indeed, there must be at most $n - 2$ other sources of report to be excluded, e.g., $\hat{x}_0 = \hat{x}_{i_1} = \hat{x}_{i_2}$.

To simplify the argument in the next step, denote

$$\mathcal{I}_{k-} \equiv \{i: i\text{'s dominance incentive is established before or at round } k \text{ of IESDS}\}$$

Then, the set of those whose incentive has not been established is $\mathcal{I}_{k+} \equiv \{1, 2, \dots, n\} \setminus \mathcal{I}_k$.

Step 2: Agent i_k 's payment is not contingent on the $n - k$ agents whose dominance incentive is established after round k . We prove by induction. First, suppose that agent i_k 's contract does not contain reports of any agent from \mathcal{I}_{k+} . By step 1, this implies *exactly* $n - k$ other sources of reports have been excluded in i_k 's contract, so i_k 's payment is dependent on all

agents in set \mathcal{I}_{k-} . In this case, denote the payment-contingent set of report profiles to be $X^* = \{\chi: \hat{x}_0 = \hat{x}_{i_1} = \hat{x}_{i_2}, \dots = \hat{x}_{i_k}\}$. Note that there is only one IC because the relevant $a_{-i_k} \equiv \mathbf{1}$. The binding IC solves for

$$w_{i_k}(X^*) = \frac{c}{\Pr(X^*|a_{i_k} = 1, a_{-i_k} = \mathbf{1}) - \Pr(X^*|a_{i_k} = 0, a_{-i_k} = \mathbf{1})}$$

and the incentive cost is

$$K_{i_k} = w_{i_k}(X^*) \Pr(X^*|a_{i_k} = 1, a_{-i_k} = \mathbf{1}).$$

Then, we claim that an inclusion of any report from agent in set \mathcal{I}_{k-} weakly increases the incentive cost for agent i_k . Denote a generic agent $i_j \in \mathcal{I}_{k+}$ ($j > k$). Without loss, suppose $t_{i_j} > t_{i_k}$, i.e., i_j is located to the “right” of agent i_k . There are two cases depending on the location of agent i_j .

Case 1: \nexists agent $m \in \mathcal{I}_{k-}$ such that $t_m > t_j$. In this case, the report \hat{x}_{i_j} is the “right-most” report that is included in agent i_k ’s contract. Denote the “second right-most” agent in i_k ’s contract to be agent ℓ (possibly $\ell = i_k$, which means there is no agent from \mathcal{I}_{k-} located between i_k and i_j). We discuss the case where the additional condition on the payment-contingent profiles by including \hat{x}_{i_j} is $\hat{x}_\ell = \hat{x}_{i_j}$, and the other case with $\hat{x}_\ell \neq \hat{x}_{i_j}$ is similar. Denote the new set of payment-contingent report profiles to be $X' = X^* \cap \{\hat{x}_\ell = \hat{x}_{i_j}\}$. Since $a_{i_j} \in \mathcal{I}_{k+}$, there are now 2 IC constraints for i_k , one with $a_{i_j} = 1$ and the other with $a_{i_j} = 0$. The optimal wage must satisfy both ICs. In particular, the IC with $a_{i_j} = 0$ provides a lower bound for the wage (we omit $a_{-(i_k, i_j)} = \mathbf{1}$ to simplify notation, as the agents in $\mathcal{I}_{k-} \setminus \{i_k\}$ always exert effort in i_k ’s incentive problem):

$$\begin{aligned} w_{i_k}(X') &\geq \frac{c}{\Pr(X'|a_{i_k} = 1, a_{i_j} = 0) - \Pr(X'|a_{i_k} = 0, a_{i_j} = 0)} \\ &= \frac{c}{\Pr(\hat{x}_{i_j} = \hat{x}_\ell | a_\ell = 1, a_{i_j} = 0) [\Pr(X^*|a_{i_k} = 1) - \Pr(X^*|a_{i_k} = 0)]} \\ &= \frac{1}{2} [\Pr(X^*|a_{i_k} = 1) - \Pr(X^*|a_{i_k} = 0)] \end{aligned}$$

The incentive cost changes to

$$\begin{aligned}
w_{i_k}(X') \Pr(X'|a_{i_k} = 1, a_{i_j} = 1) &= w_{i_k}(X') \Pr(X^*|a_{i_k} = 1) \Pr(\hat{x}_{i_j} = \hat{x}_\ell|a_\ell = 1, a_{i_j} = 1) \\
&= w_{i_k}(X') \Pr(X^*|a_{i_k} = 1) \left[\frac{1}{2} + e^{-\lambda(t_{i_j} - t_\ell)} \left(\frac{1}{2} - 2q(1 - q) \right) \right] \\
&\geq \frac{c}{\frac{1}{2}[\Pr(X^*|a_{i_k} = 1) - \Pr(X^*|a_{i_k} = 0)]} \Pr(X^*|a_{i_k} = 1) \left[\frac{1}{2} + e^{-\lambda(t_{i_j} - t_\ell)} \left(\frac{1}{2} - 2q(1 - q) \right) \right] \\
&= K_{i_k} \cdot (1 + e^{-\lambda(t_{i_j} - t_\ell)} (1 - 4q(1 - q))) > K_{i_k}
\end{aligned}$$

Similar approach shows that the incentive cost increases if the payment-contingent profiles depend on $\hat{x}_{i_k} \neq \hat{x}_j$, resulting from IC with $a_{i_j} = 1$. Therefore, it is suboptimal to condition agent i_k 's wage on agent i_j 's report in this case.

Case 2: \exists agent $m \in \mathcal{I}_{k-}$ such that $t_m > t_{i_j}$. Again use previous notation, we have three neighboring reports $\hat{x}_\ell, \hat{x}_{i_j}, \hat{x}_m$. We discuss the case where the additional condition on the payment-contingent profiles by including \hat{x}_{i_j} is $\hat{x}_\ell = \hat{x}_{i_j} = \hat{x}_m$, and the other case with $\hat{x}_\ell = \hat{x}_m \neq \hat{x}_{i_j}$ is similar. Denote the new set of payment-contingent report profiles to be $X'' = X^* \cap \{\hat{x}_\ell = \hat{x}_{i_j} = \hat{x}_m\}$. Again, there are two IC constraints for i_k as in Case 1. The IC with $a_{i_j} = 0$ provides a lower bound for the wage

$$\begin{aligned}
w_{i_k}(X'') &\geq \frac{c}{\Pr(X''|a_{i_k} = 1, a_{i_j} = 0) - \Pr(X''|a_{i_k} = 0, a_{i_j} = 0)} \\
&= \frac{c}{\frac{\Pr(\hat{x}_{i_j} = \hat{x}_m = \hat{x}_\ell|a_\ell = a_m = 1, a_{i_j} = 0)}{\Pr(\hat{x}_{i_j} = \hat{x}_\ell|a_\ell = a_m = 1)} [\Pr(X^*|a_{i_k} = 1) - \Pr(X^*|a_{i_k} = 0)]} \\
&= \frac{c}{\frac{1}{2}[\Pr(X^*|a_{i_k} = 1) - \Pr(X^*|a_{i_k} = 0)]}
\end{aligned}$$

The incentive cost changes to

$$\begin{aligned}
w_{i_k}(X'') \Pr(X''|a_{i_k} = 1, a_{i_j} = 1) &= w_{i_k}(X'') \Pr(X^*|a_{i_k} = 1) \frac{\Pr(\hat{x}_{i_j} = \hat{x}_m = \hat{x}_\ell|a_\ell = a_m = 1, a_{i_j} = 1)}{\Pr(\hat{x}_m = \hat{x}_\ell|a_\ell = a_m = 1)} \\
&= w_{i_k}(X'') \Pr(X^*|a_{i_k} = 1) \frac{[\frac{1}{4} - q(1 - q)] (1 + e^{-\lambda(t_m - t_{i_j})})(1 + e^{-\lambda(t_{i_j} - t_\ell)}) + q(1 - q)}{[\frac{1}{2} - 2q(1 - q)](1 + e^{-\lambda(t_m - t_\ell)}) + 2q(1 - q)} \\
&\geq \frac{c}{\frac{1}{2}[\Pr(X^*|a_{i_k} = 1) - \Pr(X^*|a_{i_k} = 0)]} \Pr(X^*|a_{i_k} = 1) \left[\frac{1}{2} + \frac{e^{-\lambda(t_{i_j} - t_\ell)} + e^{-\lambda(t_m - t_{i_j})}}{\Pr(\hat{x}_m = \hat{x}_\ell|a_\ell = a_m = 1)} \right] \\
&= K_{i_k} \cdot \left(1 + \frac{2e^{-\lambda(t_{i_j} - t_\ell)} + 2e^{-\lambda(t_m - t_{i_j})}}{\Pr(\hat{x}_m = \hat{x}_\ell|a_\ell = a_m = 1)} \right) > K_{i_k}
\end{aligned}$$

Again, similar approach shows that the incentive cost increases if the payment-contingent profiles depend on $\hat{x}_\ell = \hat{x}_m = \hat{x}_{i_j}$. We conclude that agent i_k 's wage does not depend on agent i_j 's report. Generally, if the payment-contingent profiles depend on $\hat{x}_{i_k} \neq \hat{x}_{i_j}$ (resp. $\hat{x}_{i_k} = \hat{x}_{i_j}$), then the constraint involving $a_{i_j} = 1$ (resp. $a_{i_j} = 0$) requires a cost increase. Using the above argument repeatedly, it is easy to see that it is suboptimal to contingent agent i_k 's wage on any reports of any combination of agents from set \mathcal{I}_{k+} . This proves the claim in Step 2.

Finally, combining Step 1 and 2, we conclude that agent i_k 's payment must be contingent on *exactly* k agents' reports, and the set of these agents is \mathcal{I}_{k-} . This finishes the proof. \square

On Information Loss. We claim that for $n = 2$, the information loss becomes

$$\begin{aligned} L^q(\tau) &= qL(\tau) + (1 - q) \left[2 \int_0^{\frac{t_1}{2}} \frac{1}{2} (1 + e^{-\lambda t}) dt + 2 \int_{t_1}^{\frac{t_1+t_2}{2}} \frac{1}{2} (1 + e^{-\lambda(t-t_1)}) dt + \int_{t_2}^1 \frac{1}{2} (1 + e^{-\lambda(t-t_2)}) dt \right] \\ &= \frac{1}{2} - \frac{5(2q-1)}{2\lambda} + \frac{2q-1}{\lambda} [e^{-\lambda\frac{t_2-t_1}{2}} + e^{-\lambda\frac{t_1}{2}} + \frac{1}{2}e^{-\lambda(1-t_2)}] \end{aligned}$$

The optimal policy still respects each report's illumination range. However, given each report \hat{x}_i , only with probability q this report matches the true state x_i , and the loss sources from the transition of state, which is captured in the term $qL(\tau)$. On the other hand, there is a probability $1 - q$ that $\hat{x}_i \neq x_i$, and the loss sources from *persistence* of x_i , which is captured with the term led by $1 - q$. As in the benchmark model, the information-loss-minimizing task allocation still allocate to each agent the same length of illumination range. This result naturally generalizes.

Proposition C.3. *Fixing imperfect learning parameter $q \in (0.5, 1)$, the information-loss-minimizing task allocation is τ^\dagger with*

$$t_i^\dagger = \frac{2i}{2n+1}, \forall i = 1, 2, \dots, n.$$

Optimal Task Allocation When $n = 2$. Next, we turn to the task allocation minimizing both incentive cost and information loss. The analysis of generic n is challenging, but we managed to grapple with the special case with $n = 2$. Again, we call the identity permutation of dominance order to be the chain monitoring structure. With $n = 2$, we illustrate that chain is still optimal. In the sandwich monitoring structure, agent 1 is paid if and only if $\hat{x}_0 = \hat{x}_1 = \hat{x}_2$. Her binding IC is

$$\left(\left[\frac{1}{4} - q(1-q) \right] (1 + e^{-\lambda t_1})(1 + e^{-\lambda(t_2-t_1)}) - \left(\frac{1}{4} - q(1-q) \right) (1 + e^{-\lambda t_2}) \right) w_1 = c$$

which gives

$$w_1 = \frac{c}{\left[\frac{1}{4} - q(1 - q)\right] (e^{-\lambda t_1} + e^{-\lambda(t_2 - t_1)})}$$

The incentive cost for sandwich monitoring is thus

$$K^s(t_1, t_2) = c \left(1 + \frac{1 + e^{-\lambda t_2}}{e^{-\lambda t_1} + e^{-\lambda(t_2 - t_1)}} + \frac{q(1 - q)}{\left[\frac{1}{4} - q(1 - q)\right] (e^{-\lambda t_1} + e^{-\lambda(t_2 - t_1)})} \right) + c(1 + e^{\lambda t_2}) + \frac{2q(1 - q)ce^{\lambda t_2}}{\frac{1}{2} - 2q(1 - q)}$$

Define the chain counterpart to be

$$K^c(t_1, t_2) = c(1 + e^{\lambda t_1}) + \frac{2q(1 - q)ce^{\lambda t_1}}{\frac{1}{2} - 2q(1 - q)} + c(1 + e^{\lambda(t_2 - t_1)}) + \frac{2q(1 - q)ce^{\lambda(t_2 - t_1)}}{\left(\frac{1}{2} - 2q(1 - q)\right)(1 + e^{-\lambda t_1})}$$

Fix any $t_2 > 0$, $K^s(\cdot, t_2)$ is symmetric and unimodal in $t_1 \in [0, t_2]$. On the other hand, simple algebra verifies $K^c(\cdot, t_2)$ is convex in $t_1 \in [0, t_2]$. Also, $K^c(t_2, t_2) = K^s(t_2, t_2)$ and $K^c(0, t_2) < K^s(0, t_2)$. Therefore, $K^c(t_1, t_2) \leq K^s(t_1, t_2)$, for all $t_1 \in [0, t_2]$. This holds for all $t_2 > 0$. Hence, the chain monitoring structure dominates the sandwich for all τ .

Note that the information loss is still symmetric in distances between agents, but the incentive cost is unbalanced, which leads to closer monitoring of agent 1 than agent 2. Indeed, define $\Delta t_1 = t_1 - 0$, $\Delta t_2 = t_2 - t_1$, then the task allocation problem is

$$\min_{\Delta t_1, \Delta t_2} L^p(\Delta t_1, \Delta t_2) + K^c(\Delta t_1, \Delta t_2)$$

has the following first-order conditions for an interior solution

$$\left(e^{-\lambda(1 - \Delta t_1 - \Delta t_2)} - e^{-\frac{\lambda \Delta t_1}{2}} \right) \frac{2q - 1}{2} + \lambda c e^{\lambda \Delta t_1} + \frac{2q(1 - q)c\lambda e^{\lambda \Delta t_1}}{\frac{1}{2} - 2q(1 - q)} + \frac{2q(1 - q)c\lambda e^{\lambda(\Delta t_2 - \Delta t_1)}}{\left[\frac{1}{2} - 2q(1 - q)\right](1 + e^{-\lambda \Delta t_1})^2} = 0 \quad (25)$$

$$\left(e^{-\lambda(1 - \Delta t_1 - \Delta t_2)} - e^{-\frac{\lambda \Delta t_2}{2}} \right) \frac{2q - 1}{2} + \lambda c e^{\lambda \Delta t_2} + \frac{2q(1 - q)c\lambda e^{\lambda \Delta t_2}}{\left[\frac{1}{2} - 2q(1 - q)\right](1 + e^{-\lambda \Delta t_1})} = 0 \quad (26)$$

Conditions (26) – (25) require

$$\left(e^{-\frac{\lambda \Delta t_1}{2}} - e^{-\frac{\lambda \Delta t_2}{2}} \right) \frac{2q - 1}{2} + \lambda c (e^{\lambda \Delta t_2} - e^{\lambda \Delta t_1}) + \frac{2q(1 - q)c\lambda}{\left[\frac{1}{2} - 2q(1 - q)\right](1 + e^{-\lambda \Delta t_1})^2} [e^{\lambda \Delta t_2} - e^{\lambda \Delta t_1} - 2 - e^{-\lambda \Delta t_1}] = 0$$

Suppose $\Delta t_1 \geq \Delta t_2$, then

$$e^{-\lambda(1-\Delta t_1-\Delta t_2)} - e^{-\frac{\lambda\Delta t_1}{2}} \leq 0, e^{\lambda\Delta t_2} - e^{\lambda\Delta t_1} \leq 0, e^{\lambda\Delta t_2} - e^{\lambda\Delta t_1} - 2 - e^{-\lambda\Delta t_1} < 0.$$

This implies that (26) – (25) < 0, a contradiction. Therefore, the interior solution τ must have the property $t_2 - t_1 > t_1$. The results are summarized in the following two propositions.

Proposition C.4. *Fixing $n = 2$, $\tau = (t_1, t_2)$, and the imperfect learning parameter $q \in (0.5, 1)$. The minimal incentive cost to robustly implement work and truth-reporting is*

$$K^c(t_1, t_2) = c(1 + e^{\lambda t_1}) + \frac{2q(1-q)ce^{\lambda t_1}}{\frac{1}{2} - 2q(1-q)} + c(1 + e^{\lambda(t_2-t_1)}) + \frac{2q(1-q)ce^{\lambda(t_2-t_1)}}{(\frac{1}{2} - 2q(1-q))(1 + e^{-\lambda t_1})}$$

which is attained by paying each agent if and only if her report matches all the reports to the left, i.e.,

$$w_1 = \begin{cases} \frac{ce^{\lambda t_1}}{\frac{1}{2} - 2q(1-q)}, & \text{if } \hat{x}_1 = \hat{x}_0 \\ 0, & \text{otherwise.} \end{cases}$$

$$w_2 = \begin{cases} \frac{ce^{\lambda(t_2-t_1)}}{(\frac{1}{4} - q(1-q))(1 + e^{-\lambda t_1})}, & \text{if } \hat{x}_2 = \hat{x}_1 = \hat{x}_0 \\ 0, & \text{otherwise.} \end{cases}$$

Proposition C.5. *Fixing $n = 2$ and the imperfect learning parameter $q \in (0.5, 1)$, the optimal (interior) task allocation exhibits $t_2 > t_2 - t_1$.*

References

- Alchian, A. A. and H. Demsetz (1972). Production, information costs, and economic organization. *American Economic Review* 62(5), 777–795.
- Azrieli, Y. (2021). Monitoring experts. *Theoretical Economics* 16(4), 1313–1350.
- Azrieli, Y. (2022). Delegated expertise: Implementability with peer-monitoring. *Games and Economic Behavior* 132, 240–254.
- Bardhi, A. (2022). Attributes: Selective learning and influence. *Available at SSRN 3468546*.
- Bardhi, A. and N. Bobkova (2023). Local evidence and diversity in minipublics. *Journal of Political Economy* 131(9), 2451–2508.

- Bernheim, B. D. (1984). Rationalizable strategic behavior. *Econometrica: Journal of the Econometric Society*, 1007–1028.
- Bohren, J. A. and T. Kravitz (2019). Peer monitoring with partial commitment.
- Bolton, P. and M. Dewatripont (2004). *Contract theory*. MIT press.
- Callander, S. (2011). Searching and learning by trial and error. *American Economic Review* 101(6), 2277–2308.
- Camboni, M. and M. Porcellacchia (2023). Monitoring team members: Information waste and the self-promotion trap.
- Cetemen, D., C. Urgan, and L. Yariv (2023). Collective progress: Dynamics of exit waves. *Journal of Political Economy* 131(9), 2402–2450.
- Cusumano, C., T. Gan, and F. Pieroth (2023). Misaligning incentives in teams: Unique implementation with optimal monitoring. *Available at SSRN 4533867*.
- De Janvry, A., G. He, E. Sadoulet, S. Wang, and Q. Zhang (2023). Subjective performance evaluation, influence activities, and bureaucratic work behavior: Evidence from china. *American Economic Review* 113(3), 766–799.
- Deb, R., M. M. Pai, and M. Said (2018). Evaluating strategic forecasters. *American Economic Review* 108(10), 3057–3103.
- Dong, M. and T. Mayskaya (2023). Does reducing communication barriers promote diversity? *Available at SSRN 4164162*.
- Ederer, F., R. Holden, and M. Meyer (2018). Gaming and strategic opacity in incentive provision. *The RAND Journal of Economics* 49(4), 819–854.
- Fudenberg, D. and J. Tirole (1991). *Game theory*. MIT press.
- Garicano, L. (2000). Hierarchies and the organization of knowledge in production. *Journal of political economy* 108(5), 874–904.
- Gromb, D. and D. Martimort (2007). Collusion and the organization of delegated expertise. *Journal of economic theory* 137(1), 271–299.

- Halac, M., I. Kremer, and E. Winter (2023). Monitoring teams. *American Economic Journal: Microeconomics*.
- Halac, M., E. Lipnowski, and D. Rappoport (2021). Rank uncertainty in organizations. *American Economic Review* 111(3), 757–786.
- Holmström, B. (1979). Moral hazard and observability. *The Bell journal of economics*, 74–91.
- Holmström, B. (1982). Moral hazard in teams. *The Bell journal of economics*, 324–340.
- Holmstrom, B. and P. Milgrom (1991). Multitask principal–agent analyses: Incentive contracts, asset ownership, and job design. *The Journal of Law, Economics, and Organization* 7(special_issue), 24–52.
- Jovanovic, B. and R. Rob (1990). Long waves and short waves: Growth through intensive and extensive search. *Econometrica: Journal of the Econometric Society*, 1391–1409.
- Legros, P. and S. A. Matthews (1993). Efficient and nearly-efficient partnerships. *The Review of Economic Studies* 60(3), 599–611.
- Li, A. and M. Yang (2020). Optimal incentive contract with endogenous monitoring technology. *Theoretical Economics* 15(3), 1135–1173.
- MacLeod, W. B. (2003). Optimal contracting with subjective evaluation. *American Economic Review* 93(1), 216–240.
- Miller, N., P. Resnick, and R. Zeckhauser (2005). Eliciting informative feedback: The peer-prediction method. *Management Science* 51(9), 1359–1373.
- Miller, N. H. (1997). Efficiency in partnerships with joint monitoring. *journal of economic theory* 77(2), 285–299.
- Pearce, D. G. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica: Journal of the Econometric Society*, 1029–1050.
- Pesendorfer, W. and A. Wolinsky (2003). Second opinions and price competition: Inefficiency in the market for expert advice. *The Review of Economic Studies* 70(2), 417–437.

Rahman, D. (2012). But who will monitor the monitor? *American Economic Review* 102(6), 2767–2797.

Rahman, D. and I. Obara (2010). Mediated partnerships. *Econometrica* 78(1), 285–308.

Rieke, N., J. Hancox, and W. Li *et al.* (2020). The future of digital health with federated learning. *npj Digital Medicine* 3, 119.

Strausz, R. (1999). Efficiency in sequential partnerships. *Journal of Economic Theory* 85(1), 140–156.

Winter, E. (2004). Incentives and discrimination. *American Economic Review* 94(3), 764–773.