Organizational Design in the Knowledge Economy*

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

April 13, 2025

Abstract

We study a principal-agent model of spatial learning. The principal must make decisions across a continuum of problems, but lacks complete knowledge of how each problem maps to its solution. This mapping is represented by the realized path of a Markov process. The principal, facing a fixed number of agents, assigns one problem to each agent, who can privately exert effort to solve the assigned problem. We characterize the optimal assignment of problems to agents and the compensation scheme that robustly implements effort by all agents. The model highlights a trade-off that shapes organizational boundaries in a knowledge economy.

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, Doruk Cetemen, Jaden Chen, Joyee Deb, Rahul Deb, Francesc Dilmé, Miaomiao Dong, Wei He, Ashwin Kambhampati, Jacob Kohlhepp, Maciej H. Kotowski, Jiangtao Li, Erik Madsen, Margaret Meyer, Andrew Newman, Ichiro Obara, Marcin Pęski, Peter Norman, Joao Ramos, Anna Sanktjohanser, Tan Gan, Can Tian, Huseyin Yildirim, Kun Zhang, and audiences at numerous conferences and seminars for their valuable comments. We are especially indebted to Steven Callander, whose constructive and detailed suggestions greatly improved the paper.

[†]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

1 Introduction

In the knowledge economy, organizations often confront many interdependent problems. To tackle them, an organization tasks agents with investigating some of the problems, learns from their discoveries, and extrapolates insights to other related problems. For example, a pharmaceutical company assigns researchers to test drug formulations for diseases, in which case test results convey information not only about tested combinations of formulations and diseases but also about untested ones. Informational interdependence across problems is also relevant in other organizations, such as when a tech firm assigns engineers to troubleshoot system failures or a policymaker deploys analysts to assess economic policies across regions.

These scenarios raise questions: How should an organization assign problems to a limited number of agents and incentivize them to solve the problems? What is the range of problems an organization of a given size can handle? These questions are fundamental to the study of organizations, from the internal structure to the scope of the organization itself.

We formalize such a situation as a principal-agent model of spatial learning. An organization consists of a principal and a fixed number of agents. The principal faces a continuum of problems, each requiring a decision. Each problem has a binary uncertain solution, which captures the principal's correct action for that problem. The set of problems is represented by the unit interval, and the realized state—i.e., the mapping of each problem to its solution—is represented by the realized path of a Markov process. Thus, solutions to different problems are correlated, and the principal can extrapolate from one problem's solution to others. The principal observes the solution to the status quo problem, located at the interval's left endpoint, but is uninformed about solutions to other problems.

The principal's design problem consists of two parts. First, the principal chooses a task allocation, which assigns one problem to each agent. Agents then decide whether to exert a costly effort. If an agent exerts effort, she generates a signal whose realization coincides with the solution of her assigned problem; otherwise, the signal is a state-independent draw. Second, the principal commits to a compensation scheme, which determines a (non-negative) payment to each agent as a function of realized signals. We solve for the optimal task allocation and compensation scheme that minimize the sum of payments to agents and the loss from decision problems, while robustly incentivizing all agents to exert effort as a unique outcome.

We begin with the first-best benchmark, in which effort choice is contractible. In this case, the principal arrays agents so that they are equally spaced and span the entire set of problems. The resulting task allocation reflects an equal treatment of problems and agents. Specifically, regions of the problem space are investigated equally well, and no parts of the problem space are prioritized over others. Each agent is responsible for an equally broad set of problems, so that each agent's *illumination range*—i.e., the set of problems for which the agent's discovery determines the principal's decision—is equalized across all agents.

We then turn to the second-best scenario, in which effort is not contractible. In contrast to the first-best case, the optimal task allocation no longer retains equal treatment of problems. Specifically, agents remain equally spaced, but compared with the first best, they are assigned to problems that are closer to each other and closer to the principal's status quo problem. As a result, problems on one side of the problem space are equally covered, whereas problems distant from the status quo remain relatively unexplored; consequently, the principal's decisions regarding these problems will be made with greater error.

The clustering of agents on one side also means that the agent on the flank faces a broader set of problems extending to her right. Her illumination range is much greater, like a lone lantern casting light into the darkness. However, the farther her light stretches, the dimmer and more diffuse it becomes, which makes it harder for the principal to infer solutions. As a result, although the agent's discovery provides some visibility over a wide area, the clarity and reliability of her insights diminish, which leads to lower decision quality.

The second-best task allocation reflects a novel trade-off: Assigning diverse problems maximizes the value of information generated by agents, but assigning similar problems reduces agency costs. Intuitively, the signals from agents who exert effort are strongly correlated when they investigate similar problems. In such a case, the principal can relatively cheaply incentivize agents' effort by implementing a cross-verification scheme that rewards agents for generating discoveries that are consistent with each other. As a result, the optimal balance between diversification and agency costs leads the principal to explore a narrow set of problems compared with the first best.

Our model opens up a new way of conceiving the scope and boundaries of organizations in a knowledge economy. For a fixed number of agents, an organization (or a firm) can be defined as the set of problems it handles, and its boundary as the boundary of that set. Our results show that—in the presence of moral hazard—the organization will focus on a narrow set of problems, and its boundary becomes closer to the problems the principal is familiar with. This result offers a novel explanation for a key idea in the management literature: that firms should focus on a limited set of "core competencies" (Prahalad and Hamel, 1990).

The second-best task allocation is associated with the optimal robust compensation scheme. We show that the compensation scheme exhibits a one-way chain structure: Each agent is rewarded if and only if her signal is consistent with the signal of her immediate left-hand neighbor, whose problem is similar to the agent's but closer to the principal's status quo problem. The compensation scheme functions as if peers monitor each other, with each agent monitoring the agent immediately to her right. Recall that the principal himself is informed about the status quo problem on the left endpoint of the interval and monitors the first agent to his right. That agent then monitors the next agent, and so on down the chain. This hierarchy is horizontal: No agent is more senior than another. However, creating a one-way peer evaluation system is the best way to robustly implement effort.

A key feature of the optimal compensation scheme is that it underuses information in compensating agents to robustly implement effort. The ordering of problems and agents means that a sufficient statistic for each agent's effort is the signal generated by the nearest neighbor on either side. In particular, except for the agent on the flank, each agent has the left and right immediate neighbors, and their signals are indicative of the agent's effort. However, our optimal contract discards half of that information. Indeed, robustness requires no circularity in the compensation structure—i.e., if compensation for agent i depends on the signal of agent j, then compensation for agent j never depends on the signal of agent i. A compensation scheme that uses more information than our solution might induce effort more cheaply in some equilibrium, but it also includes outcomes in which some agents shirk.

Finally, our model admits another interpretation, whereby a principal selects a fixed number of agents from a pool of heterogeneous applicants, each of whom has the expertise to solve exactly one problem. We can then interpret the optimal task allocation as the optimal team composition. According to this interpretation, our result suggests that the second-best hiring policy entails inefficiency, whereby the principal hires agents whose expertise is more similar to his own. By selecting agents with overlapping knowledge, the principal and agents can better evaluate others' performance. However, this narrows the scope of problems the organization can handle to the principal's area of expertise. Such a situation contrasts with the first-best hiring policy, which can be interpreted as maximizing the "diversity bonus" (Page, 2017).

Contribution to the literature First, our paper contributes to the literature on spatial learning by exploring the team incentive perspective. The literature begins with work by Jovanovic and Rob (1990) and Callander (2011), who model the state as a realized path of Brownian motion. They use spatial learning to model the set of alternatives in the trial-and-error process.[1] Recent papers extend the framework and study agency problems; see Bardhi (2022), Bardhi and Bobkova (2023), Dong and Mayskaya (2024), and Aybas and Callander (2023) for examples. Our paper differs from those studies in several respects. First, we use a binary-state Markov process instead of Brownian motion to simplify the analysis. Second, some of these papers also explore how to incentivize agents to communicate information in spatial-learning settings. However, unlike our study, they do not consider (i) monetary transfer between the principal and agents, (ii) task allocation, or (iii) robust implementation.

We propose a new framework to consider optimal task allocation in information-production organizations. It differs from the seminal paper by Garicano (2000) and his subsequent work (see Garicano and Rossi-Hansberg (2015) for a survey) in two key ways. First, Garicano (2000) focuses on *vertical* specialization, in which tasks are structured hierarchically. Agents are

[1] Various applications include law and economics (Callander and Clark, 2017); political economy (Callander and Harstad, 2015); and the incentives of researchers (Garfagnini and Strulovici, 2016; Carnehl and Schneider, 2025). Also see Cetemen et al. (2023) for a model of search-based collective spatial learning.

assigned to different levels based on their ability to handle problems, with lower-level agents managing routine tasks and escalating complex cases upward. In contrast, we study *horizontal* specialization, in which all agents operate at the same level and are responsible for solving a range of problems independently. Second, Garicano (2000) abstracts away from incentive issues, whereas we explicitly consider agency costs and monitoring as central to the trade-offs in task allocation.

We also contribute to the literature on information acquisition by multiple agents. To the best of our knowledge, the issue of optimal task allocation has largely been neglected in the literature. One exception is Bohren and Kravitz (2019), who study a model in which a principal employs agents to learn the realizations of many independent states. The principal can (i) assign multiple agents to the same task for peer monitoring and (ii) assign multiple tasks to the same agent and take 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 effectiveness does not arise.

More broadly, our optimal contract rewards agreement (i.e., generating the same signal realization) between agents. This property arises in other papers in which agents acquire information about some common state at cost, such as Gromb and Martimort (2007); Bohren and Kravitz (2019); and Azrieli (2021). In contrast to these papers, we consider a principal who aims to implement information acquisition as a unique outcome. Also, the similarity of tasks agents investigate is endogenous. Rewarding the consistency of signals also arises in other settings. For example, in Deb et al. (2018), rewarding consistency is useful for screening a strategic forecaster who observes signals about some persistent state.

Last, our paper adds to 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 individual performance is available, the optimal robust contract may not use it because it could introduce undesirable equilibria. Halac et al. (2024) allow the principal to divide agents into groups, with each delivering a signal of joint performance, but the number of groups is constrained. We recommend Halac (2025) for a recent survey. These papers do not consider endogenous task allocation. In contrast, the principal in our model determines the peer monitoring structure through task allocation.

Organization Section 2 describes the model. Section 3 derives the first-best allocation, which minimizes an information loss. Section 4 derives the optimal compensation scheme and task allocation, and Section 5 describes extensions. Omitted proofs for the main results are in the Appendix, and details of the extensions are in the Supplemental Appendix.

2 Model

We describe our model and the principal's design problem. Section 2.3 motivates and interprets some of the modeling assumptions.

2.1 Environment

We develop a principal-agent model of spatial learning. The principal faces a continuum of decision problems, each associated with an uncertain state. The principal assigns problems to agents, who can exert effort to uncover the state of the assigned problem. The principal assigns problems and compensates agents to minimize his overall loss.

The model consists of three stages: First, the principal posts a compensation scheme and assigns problems to agents. Second, agents choose whether to exert effort, which determines the quality of information they generate. Finally, the principal takes actions to minimize losses arising from incorrect decisions. Below, we first explain how the state is generated, then describe each stage.

Uncertain State The principal faces a set of interrelated decision problems, which we represent as a continuum of *locations* on the interval $[0, 1]$. Each location $t \in [0, 1]$ has a binary *local state* $x(t) \in \{A, B\}$, which represents the correct decision for the problem. These local states form the overall *state* of the world,

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

The state $x(\cdot)$ is the realized path of a Markov process, which is defined as follows. The initial state, $x(0)$, is a uniform random draw from $\{A, B\}$. At each $t \in [0, 1]$, a Poisson shock arrives at rate $\lambda > 0$. If no shock arrives, $x(\cdot)$ is continuous at t. If a shock arrives, then $x(t)$ is drawn independently and uniformly from $\{A, B\}$. This setup captures varying degrees of interdependence across problems—e.g., when $\lambda = 0$, all locations share the same state; as $\lambda \to \infty$, local states at different locations become independent (Section 2.3 summarizes properties of the state distribution we use). The state distribution is commonly known.

Figure 1 depicts a realized state such that $x(t) = A$ on $[t_1, t_2)$ and $x(t) = B$ otherwise. It corresponds to any realized path of the Markov process such that the initial state is $x(0) = B$ and state transitions occur at t_1 and t_2. In the figure, shocks arrive at t_1, t_2, and t_3, but state transitions occur only at t_1 and t_2.

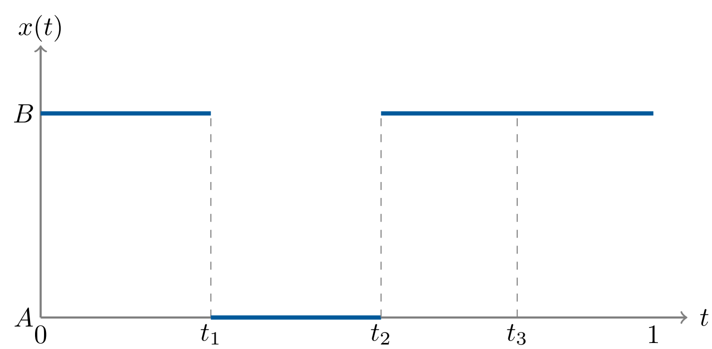

Figure 1: A state realization such that $x(t) = A$ if $t \in [t_1, t_2)$ and $x(t) = B$ otherwise. Here, A and B are real numbers such that $A < B$.

Task Allocation and Information Acquisition The principal faces n agents, where $n \in \mathbb{N}$. At the outset, the principal tasks agents to discover some of the local states. Formally, the principal chooses a *task allocation*, $\tau = (t_1, ..., t_n)$, where $t_i \in [0, 1]$ is the location assigned to agent i. Without loss of generality, we focus on task allocations such that $t_1 \leq \cdots \leq t_n$.

Given the task allocation, each agent i simultaneously decides whether to exert effort. An agent's choice determines the quality of the public signal s_i she produces. Specifically, if agent i exerts effort she incurs a cost of $c > 0$, and her signal equals the local state of her assigned location, i.e., $s_i = x(t_i)$. If agent i shirks—which is costless—her signal s_i is independent of the state and instead determined by a third party, called the *noisemaker*.

The noisemaker is a player who chooses a noise profile $(y_1, \dots, y_n) \in \{A, B\}^n$, where each y_i represents the realization of the signal produced by agent i if (and only if) shirking. The noisemaker is indifferent between all possible noise profiles in $\{A, B\}^n$ (i.e., its payoff equals 0 for any strategy profile) and moves simultaneously with agents.

The noisemaker is a modeling device. It allows us to capture a strategic ambiguity, whereby agents cannot predict or even form a belief about signals produced by shirkers. Also, it keeps the agent's strategy space binary, which facilitates the characterization of robust contracts. Section 2.3 provides a detailed discussion.

In addition to signals generated by agents, the principal observes the local state $x(0)$ for location 0 at no cost. Location 0 reflects the type of problems the organization can handle without being subject to agents' moral hazard. For example, it may reflect the principal's direct expertise or access to reliable information at that location without the help of agents (see Section 5 for further discussion). As a result, without agents, the principal correctly predicts local state $x(t)$ with probability $\Pr[x(0) = x(t)]$, which decreases in t and λ (see Section 2.3). Thus, the principal views a location with a larger t as a less understood or more novel problem, and a larger λ as a more uncertain decision environment.

Signal Profile and Compensation Scheme The realized state x, agents' actions, and the noisemaker's choice $(y_1, ..., y_n)$ determine a *signal profile*, denoted by $\mathbf{s} = (s_0, s_1, ..., s_n) \in \{A, B\}^{n+1}$. Here, s_0 denotes the principal's signal, which always equals the local state $x(0)$ at $t = 0$. Other signals are agents', where we have $s_i = x(t_i)$ if agent i exerts effort and $s_i = y_i$ if she shirks. A signal profile $\mathbf{s}$ is said to be *truthful* if $\mathbf{s}$ is a result of all agents exerting effort, i.e., $\mathbf{s} = (x(0), x(t_1), ..., x(t_n))$. Any truthful signal profile reveals the local states of at most

$n + 1$ locations.

The principal commits to how to compensate agents based on a realized signal profile. Formally, the principal chooses a *compensation scheme*, denoted by

$$w = (w_1(\mathbf{s}), ..., w_n(\mathbf{s}))_{\mathbf{s} \in \{A,B\}^{n+1}},$$

where agent i receives $w_i(\mathbf{s}) \geq 0$ when the signal profile is $\mathbf{s}$. We assume limited liability, so payments to agents must be nonnegative.[2] The signal profile is the only contractible object.

Principal's Decision Problem After observing the signal profile, the principal chooses action A or B for each location and incurs a loss whenever the action differs from the local state. The principal's *action policy* refers to any piecewise-constant function $p : [0, 1] \to \{A, B\}$, where $p(t)$ is the action for location $t \in [0, 1]$. Given action policy p and state x, the principal incurs a loss of

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

where $\mathbb{1}(\cdot)$ is the indicator function that takes value 1 or 0 if $p(t) \neq x(t)$ or $p(t) = x(t)$.

Payoffs Agents are risk-neutral: The payoff of an agent is equal to the payment from the principal minus effort cost c if she exerts effort; otherwise, the payoff is equal to the payment. The principal's objective is to minimize the sum of payments to agents and losses from the decision stage.

Timing The timing of the interaction is summarized as follows.

1. The principal publicly commits to a task allocation τ and a compensation scheme w.

2. Agents decide whether to exert effort, and simultaneously the noisemaker chooses $\mathbf{y} = (y_1, ..., y_n)$, which determines signal realizations for agents who shirk. Also, the state x is realized. Neither the principal nor agents directly observe $(x, \mathbf{y})$.

[2] We do not explicitly model the agent's participation decision—but assuming a zero outside option ensures that each agent weakly prefers accepting the assigned task and choosing $a_i = 0$ to opting out.

3. The signal profile $\mathbf{s} = (s_0, s_1, ..., s_n)$ becomes publicly observable. The principal chooses an action policy p, and payoffs are realized.

2.2 Principal's Design Problem

We now define the principal's design problem (i.e., Stage 1 above). For any pair of task allocation τ and compensation scheme w, Stages 2 and 3 in the above **Timing** define the following simultaneous-move incomplete information game between agents and the noisemaker, denoted by $\Gamma(\tau, w)$. Each agent i chooses between exerting effort $a_i = 1$ and shirking $a_i = 0$, and simultaneously the noisemaker chooses $(y_1, ..., y_n) \in \{A, B\}^n$, which applies to the signals of those who shirk. An agent's payoff is the expected payment from the principal—where the expectation is with respect to the realized state x—subtracted by, when applicable, effort cost c. The task allocation, compensation scheme, and state distribution are commonly known.

Fix any task allocation τ. We say that a contract w *robustly implements effort (RIE)* if all agents exert effort in every correlated rationalizable strategy profile of $\Gamma(\tau, w)$ (Fudenberg and Tirole, 1991). We call any such contract an *RIE contract* and write $\mathcal{W}(\tau)$ for the set of all RIE contracts. An RIE contract is immune to strategic risk and only relies on agents' rationality and beliefs about their opponents' rationality, and so on (Bernheim, 1984; Pearce, 1984). The notion is also stronger than requiring that all agents exert effort in any Bayesian Nash equilibrium.[3]

By the standard argument, a compensation scheme w belongs to $\mathcal{W}(\tau)$ if and only if all agents exert effort in any outcome that survives the iterated elimination of strictly dominated strategies (IESDS) in game $\Gamma(\tau, w)$. To facilitate analysis, we use this equivalent notion of implementation. The noisemaker is indifferent between any choices in $\{A, B\}^n$, so its strategies are never deleted during the IESDS.

[3]For some pair (τ, w), the induced game $\Gamma(\tau, w)$ may have a unique Nash equilibrium but multiple rationalizable action profiles, because $\Gamma(\tau, w)$ is in general not supermodular. In this sense, our implementation notion is more stringent than Nash full implementation.

Principal's Problem The principal's problem is to choose a task allocation τ and a compensation scheme w to minimize the sum of the expected payment and the loss in the decision stage, subject to the constraint that $w \in \mathcal{W}(\tau)$. We reformulate it as a two-stage problem.

First, for each (τ, w), denote the total expected payment to agents—when they all exert effort—as

$$\hat{K}(\tau, w) \triangleq \sum_{i=1}^{n} \sum_{\mathbf{s} \in \{A,B\}^{n+1}} w_i(\mathbf{s}) \Pr(\mathbf{s}|\tau),$$

where $\Pr(\mathbf{s}|\tau)$ is the probability of a truthful signal profile $\mathbf{s} \in \{A, B\}^{n+1}$ when the task allocation is τ and all agents exert effort.

We can then write the optimal robust contracting problem as

$$K(\tau) \triangleq \inf_{w \in \mathcal{W}(\tau)} \hat{K}(\tau, w). \tag{P-C}$$

The problem (P-C) generally does not have a minimum. We call the infimum of the expected payment, $K(\tau)$, the *incentive cost* for task allocation τ.

Second, we define the principal's *information loss* given task allocation τ as the minimized loss in the decision stage, i.e.,

$$L(\tau) \triangleq \sum_{\mathbf{s} \in \{A,B\}^{n+1}} \left\{ \min_{p} \mathbb{E}\left[\int_0^1 \mathbb{1}(p(t) \neq x(t)) \, \mathrm{d}t \,\middle|\, \tau, \mathbf{s} \right] \right\} \Pr(\mathbf{s}|\tau). \tag{P-I}$$

Here, the inner expectation $\mathbb{E}[\cdot|\tau, \mathbf{s}]$ is over the state x conditional on the truthful signal profile $\mathbf{s}$; the minimization is over the principal's action policies; and the outer expectation is over the truthful signal profile $\mathbf{s}$.

The principal's task allocation problem is then written as

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

where $K(\tau)$ is the incentive cost in program (P-C) and $L(\tau)$ is the information loss in program (P-I). The task allocation that minimizes incentive cost $K(\cdot)$ differs from one that minimizes information loss $L(\cdot)$. This is the principal's main trade-off.

2.3 Discussion of Assumptions

Before moving forward, we discuss some key assumptions of the model.

State Distribution Our Markov chain formulation of the state has several useful properties. First, the local state of one location is informative about other local states. To see this, take two locations, t and $t' > t$. We have $x(t) \neq x(t')$ if and only if, in the interval $(t, t']$, (i) at least one shock arrives, which occurs with probability $\int_0^{t'-t} \lambda e^{-\lambda s} ds$, and (ii) the realized local state at the last shock differs from $x(t)$, which occurs with probability 1/2. As a result, the two local states take the same value with probability

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

which is strictly greater than 1/2. The expression (1) also implies that the local states of closer locations are more likely to take the same value, i.e., $\Pr[x(t) = x(t')]$ is decreasing in $|t' - t|$. Thus, the local state of a location is more informative about nearby local states than distant ones.

Second, the local states of the nearest sampled locations are sufficient statistics for an unsampled location: Formally, suppose that the principal observes the local states of locations $t_1 < t_2 < \cdots < t_m$. The principal's belief over the local state at any unsampled location $t \in (t_j, t_{j+1})$ is given by

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

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

Third, we assume that if a shock arrives at t, then $x(t)$ is a uniform draw from $\{A, B\}$. Thus, conditional on a shock at t, the distribution of local state $x(t)$ is orthogonal to the local states of location $s < t$. This 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.

Finally, most spatial learning models—including the pioneering work by Callander (2011)—represent correlated states as the realized path of a Brownian motion to capture the complexity of the decision environment. In contrast, we use state correlation rather than complexity. By modeling states as the realization of a Markov chain, we simplify the contract space and obtain a closed-form solution for the optimal contract.

Robust Design and the Noisemaker We require a contract to induce all agents to exert effort as a unique outcome; furthermore, we use the noisemaker as a modeling device to capture this robustness requirement. We now motivate these assumptions.

First, when agents acquire information about correlated states, the principal can incentivize them through a cross-verification incentive scheme. Specifically, provided agent i exerts effort, agents i and j are more likely to generate the same signal when agent j also exerts effort than when she does not. Thus, the principal may reward agent j for generating the same signal as agent i. However, agent j's incentive depends on her belief about agent i's choice, which, in turn, depends on agent i's belief about the behavior of other agents. Managing this type of strategic uncertainty has been a focus of recent work on team production (see the related literature in the introduction), and we deem it essential for information production in an organization. Therefore, we study contracts that induce effort as a unique outcome.

In terms of modeling approach, we introduce the noisemaker in order to (i) capture a part of the robustness requirement—i.e., agents exert effort regardless of the signal realizations of shirking agents—and (ii) keep each agent's strategy space binary so that the IESDS is readily applicable. For example, an alternative approach, in which shirkers' signals are drawn from some known distribution, fails to capture the stronger robustness requirement in Part (i). At the same time, allowing shirkers to choose what (state-independent) signal to generate would expand the strategy space and complicate the characterization of robust contracts via the IESDS process. Thus introducing the noisemaker strikes a balance between the richness and tractability of the model.

3 First-best Task Allocation

We begin by characterizing the first-best task allocation, defined as the solution to the problem (P-I) of minimizing information loss without taking into account payments to agents. The principal would choose the first-best task allocation if effort choice is contractible.

The following lemma describes the principal's optimal action policy and the resulting information loss under any task allocation. When all agents exert effort, their signals and the principal's signal reveal the local states of at most $n+1$ locations. The principal then sets the action for each location equal to the realized local state of the nearest sampled location.

Lemma 1. *Fix any task allocation $\tau = (t_1, ..., t_n)$ with $t_1 \leq t_2 \leq \cdots \leq t_n$. Suppose that all agents exert effort. An action policy $p(\cdot)$ is optimal if and only if we have $p(t) = x(t^*)$ for almost every $t \in [0,1]$, where t^* is a nearest sampled location for location t, i.e., $t^* \in \arg\min_{t' \in \{0, t_1, ..., t_n\}} |t' - t|$. The resulting information loss is*

$$L(\tau) = \frac{1}{2} - \frac{2n+1}{2\lambda} + \frac{1}{\lambda}\Big[\sum_{i=1}^{n} e^{-\frac{\lambda(t_i - t_{i-1})}{2}} + \frac{1}{2}e^{-\lambda(1-t_n)}\Big]. \tag{3}$$

To gain some intuition, consider the principal's choice of an action for location t. Based on the local states revealed by the principal and agents, the principal infers whether $x(t)$ is more likely to be A or B, then matches the action $p(t)$ to the most likely realization. The optimal action for location t then equals the local state of the nearest sampled location, because a given location is more likely to have the same local state as location t when they are closer to each other.

The action policy in Lemma 1 determines the principal's *decision quality* at each location t—i.e., the probability with which the principal's action $p(t)$ matches the local state $x(t)$. Figure 2 depicts the principal's decision quality in the case of two agents. Decision quality achieves its maximum at $t = t_i$, decreases as t diverges from t_i, and attains a local minimum at $\frac{t_i + t_{i+1}}{2}$, where the principal is indifferent between matching his action to x_i and x_{i+1}. The information loss is represented by the light blue area. The first-best problem is to choose (t_1, t_2) to minimize this area.

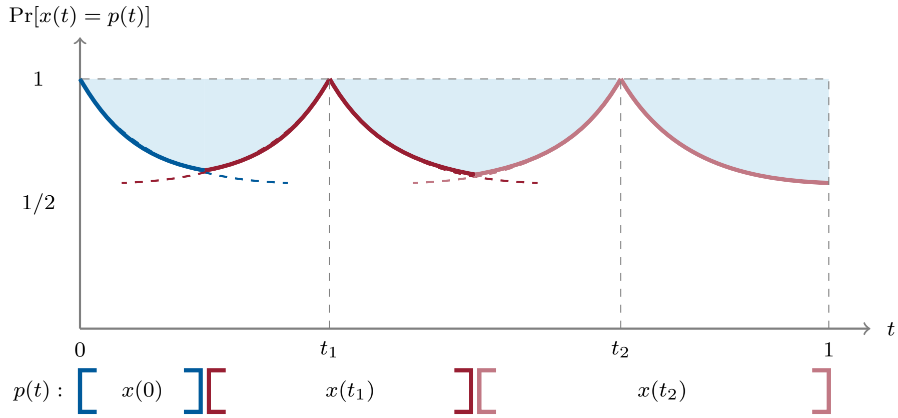

Figure 2: Solid thick curves represent the decision quality at each location under the optimal policy. The light blue area corresponds to the principal's information loss. Brackets represent agents' illumination ranges, whereby the local state in each bracket coincides with the principal's action.

The optimal action policy also determines each agent's *illumination range*, defined as the set of locations at which the principal's optimal actions coincide with the agent's signal realization (see the brackets in Figure 2). In other words, the illumination range represents the set of problems for which agent i's informational advantage compels the principal to defer to the agent's judgment, thereby granting the agent real authority (Aghion and Tirole, 1997). Whether the choice of an action is made by the principal or delegated to the agent is a matter of interpretation. According to Lemma 1, the illumination range of agent i equals the set of locations that are closer to t_i than to other agents or the principal, i.e., $[\frac{t_{i-1}+t_i}{2}, \frac{t_i+t_{i+1}}{2}]$.

In the first-best scenario, the principal assigns problems to agents so that their illumination ranges are equalized.

Proposition 1 (First-best Task Allocation)**.** *Information loss is minimized by the task allocation* $\tau^\dagger = (t_1^\dagger, ..., t_n^\dagger)$ *such that*

$$t_i^\dagger = \frac{2i}{2n+1}, \forall i = 1, 2, ..., n. \tag{4}$$

Task allocation $\tau^\dagger$ *equalizes the illumination range across all agents. The minimized information loss decreases in the number of agents,* n*, and increases in the shock-arrival rate,* λ*.*

To gain some intuition, fix t_{i-1} and t_{i+1} and consider the choice of agent i's location, $t_i \in [t_{i-1}, t_{i+1}]$. Suppose that t_i is closer to its left-hand neighbor (i.e., t_{i-1}) than its right-hand neighbor, or equivalently, that agent i is on the left side of her illumination range. In such a case, the principal can move agent i slightly to her right and decrease the information loss (see Figure 3). Indeed, within agent i's illumination range, the adjustment decreases decision quality on her left side but increases it on her right side. Because the right side of the illumination range has more tasks than the left side, the gain exceeds the loss. A symmetric argument holds when t_i is on the right side of her illumination range. As a result, each agent must be in the middle of her illumination range; this property pins down the first-best task allocation.

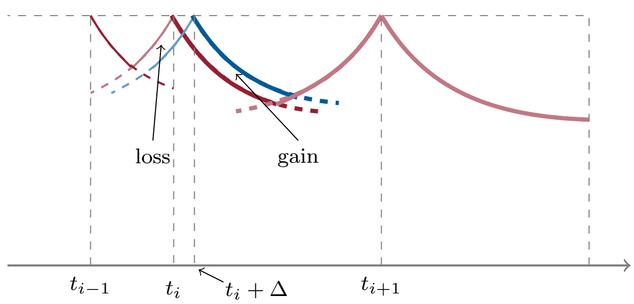

Figure 3: This figure illustrates the marginal impact of increasing t_i on decision quality for fixed t_{i-1} and t_{i+1} when t_i is on the left side of its illumination range. The marginal loss and the marginal gains are the corresponding areas delineated by four curves.

Proposition 1 implies that the first-best task allocation reflects an equal treatment of agents and problems. First, each agent is responsible for an equally broad set of problems, in that the illumination range is equalized across all agents. To see this, let $\Delta t^\dagger$ denote the distance between any two neighboring agents (or between agent 1 and the principal at $t = 0$) at the first best. The width of the illumination range of each agent i is also equal to $\Delta t^\dagger$ and independent of i, because each agent is exactly between her left and right neighbors.

Second, regions of the problem space are investigated equally well, and no parts of the problem space are prioritized over others. For example, consider the set of problems that has width $\Delta t^\dagger$, represented by interval $[t, t+\Delta t^\dagger]$. Then, the average decision quality in this interval is equalized across all possible t's. Thus, the principal learns about locations close to the status quo problem (i.e., those in interval $[0, \Delta t^\dagger]$) as accurately as those further from it (i.e., those in interval $[1-\Delta t^\dagger, 1]$).

Proposition 1 also provides comparative statics. When the principal faces more agents, the illumination range of each agent becomes narrower, which increases the principal's overall decision quality and decreases the information loss. Although the first-best task allocation does not depend on λ, a higher λ renders each agent's signal less informative about other local states, which increases the information loss.

4 Optimal Task Allocation and Compensation

We now turn to the main specification, in which agents' effort is not contractible. We begin with the principal's problem (P-C) of choosing a compensation scheme under a given task allocation, and characterize the contracts that attain the incentive cost. We then derive the second-best task allocation, which minimizes the sum of the incentive cost and information loss.

4.1 Least-cost Compensation and Incentive Cost

We define a class of compensation schemes that approximate the incentive cost. For any task allocation τ and any $\epsilon \geq 0$, define compensation scheme w^ϵ as follows: If the realized signal profile is $\mathbf{s} \in \{A, B\}^{n+1}$, each agent $i = 1, ..., n$ earns

$$w_i^\epsilon(\mathbf{s}) \triangleq \begin{cases} 2ce^{\lambda(t_i - t_{i-1})} + \epsilon & \text{if } s_i = s_{i-1} \\ 0 & \text{if } s_i \neq s_{i-1}, \end{cases} \tag{5}$$

where $t_0 = 0$. Compensation scheme w^ϵ rewards each agent i when i's signal coincides with the signal of her left-hand neighbor, who is agent $i-1$ for $i \geq 2$ or the principal for $i = 1$. The amount of the reward depends on the distance to the left-hand neighbor as well as effort cost c.

The following result states that the compensation schemes defined in (5) provide the cheapest way for the principal to robustly implement effort. Recall that given any task allocation τ, $K(\tau)$ is the incentive cost (P-C), and with an arbitrary compensation scheme w, $\hat{K}(\tau, w)$ is the total expected payment when all agents exert effort.

Proposition 2. *Fix any task allocation, τ. For any $\epsilon > 0$, compensation scheme w^ϵ defined by* (5) *robustly implements effort (i.e., $w^\epsilon \in \mathcal{W}(\tau)$). As $\epsilon \to 0$, it approximates the incentive cost, i.e.,*

$$K(\tau) = \lim_{\epsilon \to 0} \hat{K}(\tau, w^\epsilon) = \hat{K}(\tau, w^0) = c \sum_{i=1}^{n} \left[1 + e^{\lambda(t_i - t_{i-1})}\right]. \qquad (6)$$

Compensation scheme w^ϵ rewards an agent if and only if her signal matches her left-hand neighbor's. It induces a monitoring chain, whereby agent 1 is monitored by the principal, agent 2 is monitored by agent 1, and so on. Correspondingly, agent k is the k-th agent who eliminates shirking in the IESDS process. The distance $t_{i+1} - t_i$ between two adjacent agents determines the strength of correlation between their signals, and thus the effectiveness of monitoring for agent i.

The compensation scheme (5) has two key features. First, it underuses information. Specifically, agent i's payment does not depend on the signal of her right-hand neighbor, agent $i+1$, even though agent $i+1$'s signal is a part of the sufficient statistic for agent i's effort choice (Holmström, 1982).[4] To see why, consider the following situation: The principal faces two agents and pays each agent a bonus if her signal matches the other agent's signal. If the bonus is high enough, the induced game has an equilibrium in which both agents exert effort. However, the game has another equilibrium in which both agents shirk: If one agent shirks, her signal becomes independent of the state, which also discourages the other agent's effort. The compensation scheme (5) eliminates such coordination failure by ensuring that whenever agent i's pay depends on agent j's signal, agent j's pay never depends on agent i's signal.

Second, the compensation scheme (5) creates a one-directional, chain-like monitoring system. In this system, agents are ordered in a monitoring chain, with the principal at the top, followed by the agent monitored by the principal, and so on. The principal's signal, free from moral hazard, starts the monitoring chain and evaluates the agent who works on the task that is most similar to the principal's. Similarly, each agent i is assessed by the agent who works on the task that is most similar to agent i and is higher in the monitoring chain.

[4] Camboni and Porcellacchia (2023) recognize a similar information-waste result in a supermodular setting in which the principal observes a team performance measure and a signal about each individual's action.

A corollary of Proposition 2 is that the principal can reduce the incentive cost by assigning agents to locations that are closer to each other and to the principal:

Corollary 1. *The incentive cost decreases when agents are closer to each other and closer to the principal. Formally, take two task allocations,* $\tau = (t_1, ..., t_n)$ *and* $\tau' = (t'_1, ..., t'_n)$*, such that* $t_i - t_{i-1} \leq t'_i - t'_{i-1}$ *for* $i = 1, ..., n$ *with* $t_0 = t'_0 = 0$*. Then, we have* $K(\tau) \leq K(\tau')$*.*

Intuitively, if the principal assigns agents to closer locations, each agent who exerts effort faces a greater probability of generating the same signal as her left-hand neighbor—or equivalently, a greater probability of earning a bonus under compensation scheme (5). In contrast, the bonus probability is always 0.5 if agents shirk. As a result, the agents assigned to similar tasks face stronger incentives to work, which enables the principal to lower the bonus and save on the expected payment.

4.2 Proof Idea for Proposition 2

In this section, we sketch the proof of Proposition 2 by focusing on the case of two agents. The purpose is to clarify that (i) there are RIE contracts that are qualitatively different from the contracts in (5), but (ii) they cannot attain the incentive cost. Below, we say that an agent *works* if she exerts effort.

One Agent As a building block for the case of two agents, we begin with the case of one agent. Under any task allocation, the agent's signal and the principal's signal are more likely to take the same value (i.e., $s_0 = s_1 \in \{A, B\}$) when the agent works than when she shirks. Thus, the principal can focus on contracts that pay the agent some positive amount, denoted by w_1, if their signals coincide.

Given any such contract, consider the agent's incentive constraint. The principal's signal $s_0 = x(0)$ is A or B with equal probability. As a result, even if the agent shirks, she will still earn w_1 with probability 0.5 (for any choice of the noisemaker). But by working, the agent can increase the probability of the positive payment to $\Pr(x(0) = x(t_1)) = (1 + e^{-\lambda t_1})/2$ at cost c

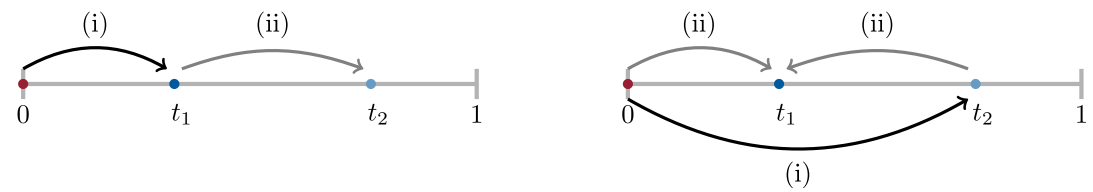

Figure 4: The left panel shows the chain monitoring structure, in which the principal's signal $x(0)$ is used to monitor agent 1 and agent 1's signal is used to monitor agent 2. The right panel shows the sandwich monitoring structure, in which the principal's signal is used to monitor agent 2, and agent 1 is monitored jointly by the principal's signal and agent 2's signal. Each arrow points from the monitor to the monitored, with the associated roman numeral indicating the order of elimination in the IESDS procedure.

(see expression (1)). Consequently, the agent strictly prefers to work (i.e., the contract belongs to $\mathcal{W}(\tau)$) if w_1 satisfies

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

The infimum of w_1's that satisfy this inequality is $w_1 = 2ce^{\lambda t_1}$. The corresponding expected payment, which is the incentive cost under a single agent, is

$$K(t_1) = c\big(1+e^{\lambda t_1}\big). \tag{7}$$

Two Agents We now turn to the case of two agents. Under any RIE contract, both agents work as a unique outcome that survives the IESDS. We can then categorize RIE contracts into two types: One represents RIE contracts such that agent 1 (i.e., the agent closer to the principal) is the first player who eliminates shirking. The other represents RIE contracts such that agent 2 is the first player who eliminates shirking. Each set of contracts 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 induces agent 1 to work as a dominant strategy. Agent 1's payment depends only on her own and the principal's signals. As in the case of one agent, agent 1 receives a positive payment if and only if her signal coincides with the principal's signal, and the infimum of the expected payment that induces working is $c\big(1+e^{\lambda t_1}\big)$. In the second step, the principal uses agent 1's signal to compensate agent 2. Again, as in the case of one agent, agent 2 receives a positive payment if and only if her signal coincides with agent 1's, and the infimum of the expected payment that induces working is $c\big(1+e^{\lambda(t_2-t_1)}\big)$. Thus, the infimum of the expected payments that induce both agents to work under this chain monitoring structure is

$$K^c(t_1,t_2) \triangleq c\sum_{i=1,2}\Big[1+e^{\lambda(t_i-t_{i-1})}\Big]. \tag{8}$$

In the second case (i.e., the right panel of Figure 4), the IESDS procedure first induces agent 2 to work as a dominant strategy; as before, the contract pays agent 2 when her signal matches the principal's signal. In the second step, the contract induces agent 1 to work by paying her if and only if her signal matches the signals of both the principal and agent 2. In this case, the contract induces a "sandwich" monitoring structure. Let $K^s(t_1,t_2)$ denote the infimum of the expected payments that induce both agents to work under the sandwich monitoring structure.

Compared with the chain monitoring structure, the sandwich monitoring structure increases the expected payment to agent 2, because her signal is compared with the signal of the principal, who is located further from agent 2. At the same time, the sandwich structure reduces the payment to agent 1, whose signal is compared not only with the principal's but also agent 2's.

We argue that the chain monitoring structure attains a lower total expected payment, i.e.,

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

To see this, we fix t_2 and view $K^s(t_1,t_2)$ and $K^c(t_1,t_2)$ as functions of t_1. The total expected payment $K^s(t_1,t_2)$ under the sandwich monitoring structure is minimized at $t_1=0$. Indeed, moving agent 1 from $t_1>0$ to $t_1=0$ will reduce the expected payment for agent 1 without changing the payment to agent 2. Under the chain monitoring structure, $K^c(t_1,t_2)$ is maximized at $t_1=0$. Indeed, moving agent 1 from $t_1>0$ to $t_1=0$ will decrease the payment to agent 1 and increase the payment to agent 2, but the total expected payment will increase because the expected payment to each agent i is convex in the distance Δt_i between agent i and her left-hand neighbor. Furthermore, if $t_1=0$, the expected payment is the same between the two

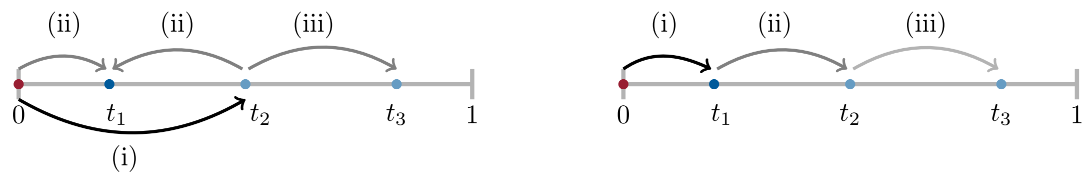

Figure 5: This figure illustrates the optimality of the chain monitoring structure when $n = 3$. Roman numerals denote the order of elimination in the IESDS process.

monitoring structures.[5] Therefore, we obtain

$$K^s(t_1, t_2) \geq K^s(0, t_2) = K^c(0, t_2) \geq K^c(t_1, t_2),$$

i.e., the chain monitoring structure attains a lower expected payment.

Finally, to see why the chain monitoring structure is optimal under any number of agents, consider the example of three agents in Figure 5. Suppose that under a given compensation scheme, the order of IESDS is agent 2, agent 1, and agent 3 (the left panel of Figure 5). Note that the principal, agent 1, and agent 2 form a sandwich monitoring structure. Suppose we exchange agents 1 and 2 in the order of IESDS. Then the principal and agents 1 and 2 form a chain monitoring structure (the right panel of Figure 5). By the previous argument, the modification and corresponding change in the compensation scheme reduce the sum of the expected payments to agents 1 and 2. The payment to agent 3 remains the same and equals $c\big(1 + e^{\lambda(t_3 - t_2)}\big)$. As a result, the modification weakly decreases the total expected payment. In general, whenever the order of IESDS induced by a contract differs from that of the chain monitoring structure, we can locally modify the contract to decrease the total expected payment.

4.3 Optimal Task Allocation

We now turn to the principal's task allocation problem, i.e.,

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

[5] If $t_1 = 0$, the principal does not benefit from using agent 2's signal to compensate agent 1. Therefore, the sandwich and chain monitoring structures attain the same infimum of the expected payments.

where the information loss $L(\tau)$ (in equation (3)) comes from the principal's optimal action policy and the incentive cost $K(\tau)$ (in equation (6)) comes from the chain compensation scheme.

The following result provides key qualitative properties of the principal's optimal task allocation. Recall that $\Delta t^\dagger = \frac{2}{2n+1}$ is the distance between any two neighboring agents under the first-best task allocation.

Proposition 3. *Compared with the first best, the optimal task allocation* $(t_1^*, ..., t_n^*)$ *locates agents closer to each other and closer to the principal. Formally, there is a unique* $\Delta t^* \in [0, \Delta t^\dagger)$ *such that* $t_i^* - t_{i-1}^* = \Delta t^*, \forall i = 1, 2, ..., n,$ *with* $t_0^* = 0$*. The distance* Δt^* *between each agent and her left-hand neighbor decreases in effort cost* c *and the number* n *of agents.*

Recall that the first-best task allocation reflects the equal treatment of problems and agents. In contrast, under moral hazard, the optimal task allocation no longer retains the equal treatment of problems, and the equal treatment of agents does not hold for one agent.

First, under the second-best allocation, agents are equally spaced but are assigned to closer locations and positioned nearer to the principal's location than in the first best. This means that the set of locations covered equally is on one side of the unit interval. Locations further away from $t = 0$ are left relatively unexplored, and as a result, the principal's decision quality declines. The intuition is present in Corollary 1: By clustering agents, the principal can make the cross-verification scheme (captured by (5)) more effective and reduce expected payments.

Second, the clustering of agents on one side means that agent n, who is furthest from the principal (i.e., $t = 0$), has a broader set of locations open on her right side. Thus, agent n's illumination range is much greater than that of other agents, although the light she casts on many of the problems is much weaker.

The task allocation in Proposition 3 optimally balances what the principal learns from agents' investigations against the cost of incentivizing them. To lower the incentive cost, the optimal balance motivates the principal to explore a narrower set of problems. In a sense, the principal abandons problems distant from his status quo problem and accepts that those decisions will be made with much less accuracy. Relative to the first best, therefore, the optimal task allocation for a given number of agents is tighter, with agents investigating problems that are more similar to each other and having information authority over a narrower set of problems (except for agent n). Simultaneously, parts of the problem space are essentially ignored.

The last part of Proposition 3 provides comparative statics.[6] First, if agents face a higher effort cost c, the principal prioritizes reducing the agency cost, and thereby assigns agents to closer locations.[7] Second, when the number n of agents increases, the principal receives more signals and attains a lower information loss. As a result, the principal puts more weight on saving the incentive cost. To do so, the principal assigns agents to work more closely with each other.

Remark 1 (Task Allocation as Team Composition)**.** We provide an alternative view of our model and results. Consider the following version of our model. The principal faces a unit mass of agents, represented by $[0, 1]$. Agents are heterogeneous and characterized by their expertise: Agent $t \in [0, 1]$ can discover the local state for location t at a cost, as in our baseline model, but she cannot do so for any other locations. The principal's problem is to choose which n agents to hire out of the pool and how to compensate them. A task allocation $(t_1, \ldots, t_n)$ is then viewed as the principal's decision to hire agents $t_1, ..., t_n$.

According to this interpretation, we can read our results as follows. At the first best, the principal maximizes the diversity of agents' expertise so that the organization can handle various kinds of problems equally well (Proposition 1). In contrast, in the presence of moral hazard, the principal prefers to hire agents who are similar to each other and similar to the principal, because they can easily monitor each other (Proposition 3). Otherwise, if one agent's expertise were very different from everyone else's, it would be expensive for the principal to incentivize such an agent, because no one in the organization would know what the right output (i.e., signal realization) should look like for that agent. Therefore, the second-best hiring policy suppresses the diversity of agents' expertise in order to mitigate the agency problem.

[6] Comparative statics with respect to λ is ambiguous: A higher λ lowers the correlation between local states, which increases the incentive cost but also makes it more attractive for the principal to diversify task allocation. A numerical example suggests that a higher λ may increase or decrease Δt^*.

[7] Proposition 3 admits a case in which the principal assigns all agents at $t = 0$; see Appendix A.5.

5 Discussion

In this section, we briefly discuss various extensions of our baseline model. Details are relegated to the Supplemental Appendix.

Stochastic Private Task Allocation In our baseline model, the principal publicly commits to a deterministic task allocation. In the Supplemental Appendix, we extend this framework by allowing the principal to commit to a probability distribution over task allocations and privately inform each agent of her realized task. The compensation scheme can then depend on both the realized task allocation and the signal profile. This stochastic design introduces uncertainty to the monitoring structure, which softens the tension between diversification benefits and the incentive costs. Specifically, the principal can simultaneously approximate (i) the first-best information loss and (ii) the minimal second-best expected payments that would emerge if all agents were at the same location. However, achieving such an outcome requires contracts that pay each agent arbitrarily large amounts with arbitrarily small probabilities.[8] Such contracts would be infeasible under practical conditions such as liquidity constraints for the principal or even slight risk aversion among agents.

Optimal Team Size We have assumed that the principal faces a fixed number of agents and induces every agent to exert effort. Suppose, instead, that the principal chooses the number of agents to hire and the action profile to implement. Then, without loss, we can assume that the principal induces all employed agents to exert effort, because inducing effort by $n' < n$ agents is equivalent to hiring n' agents and inducing all of them to exert effort. Hiring more agents leads to a higher expected payment and a lower information loss, which is akin to the familiar trade-off of inducing high versus low effort in the textbook moral hazard model (see,

[8]Our construction of stochastic allocation shares some similarity with papers such as Legros and Matthews (1993) and Rahman and Obara (2010). These papers study contracts that induce some agents to occasionally choose suboptimal actions in order to identify non-deviators. Our work differs from these papers in that we consider robust implementation and the principal implements a random monitoring structure through task assignment, rather than correlated equilibria, in a fixed game played by agents.

e.g., Proposition 14.B.3 in Mas-Colell et al. (1995)). Because the maximal information loss is bounded, the principal optimally hires finitely many agents.

Principal's Private Knowledge We assume that the principal observes $x(0)$. A direct consequence of this assumption is to allow the process of IESDS to take off, which makes it possible to consider a robust design.[9] This assumption reflects the idea that the principal is likely to either (i) investigate a fundamental problem by himself, (ii) be informed about some aspects of the project under his management from other sources, or (iii) be able to access outside experts who are not subject to moral hazard. The assumption that the principal receives a perfect signal of $x(0)$ is not essential; all results remain qualitatively the same if, for example, the principal observes a signal s_0 such that $\Pr(s_0 = A|x(0) = A) = \Pr(s_0 = B|x(0) = B) = q \in (0.5, 1)$.

Endogenous Location of the Principal The principal's location can be endogenous. Suppose that the principal can publicly commit to which local state $x(t_0)$ to observe. In this case, there exists an optimal task allocation $\tau = (t_0, t_1, ..., t_n)$ such that $t_0 \leq t_1 \leq \cdots \leq t_n$; τ is symmetric around 1/2; and any two neighboring locations are equidistant.[10] Given such a task assignment, the principal adopts the chain monitoring structure—i.e., agent 1 is monitored by the principal's signal, and each agent $i \geq 2$ is monitored by agent $i-1$'s signal. In particular, it is without loss for the principal to locate to the left of agent 1, because swapping the principal's and agent 1's locations affects neither the incentive cost nor the information loss. For example, consider task allocation $(t_0^{**}, t_1^{**}, t_2^{**})$ with $t_1^{**} < t_0^{**} < t_2^{**}$ (see Figure 6). The corresponding

[9] When the principal does not have any private signal, we can still consider an optimal partial-implementation contract and express the trade-off between diversification and monitoring power. The downside is losing the simple analytical characterization of optimal task allocation: The agent on the flank can only be monitored by agents on her right-hand side, which causes a complicated asymmetry in the design.

[10] For any fixed t_0 and t_n, $t_{i+1} - t_i$ must be constant across $i = 0, ..., n-1$ by the same argument as in Proposition 3. The optimal assignment also satisfies $t_0 = 1 - t_n$; for example, if $t_0 < 1 - t_n$, 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 $\tau = (t_0, t_1, ..., t_n)$ is symmetric around 1/2 and equidistant.

Figure 6: The left panel shows a monotone optimal task allocation $0 < t_0^* < t_1^* < t_2^* < 1$. The right panel shows a nonmonotone 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 as a strictly dominant strategy.

contract compares each agent's signal with the principal's signal. However, the principal can attain the same incentive cost and information loss with $(t_0^*, t_1^*, t_2^*) = (t_1^{**}, t_0^{**}, t_2^{**})$ under the chain monitoring structure.

Optimality of Other Compensation Schemes In the Supplemental Appendix, we focus on the case of two agents and consider the following three extensions: (i) agents' effort generates an imperfectly informative signal about the local state; (ii) agents face different effort costs; and (iii) the set of locations is represented by the unit circle. In each extension, the incentive cost is approximated by compensation schemes that differ from the chain-monitoring structure. However, the insights—that the optimal robust contract excludes mutual monitoring between agents, and that the agents are assigned to similar problems than the first best—continue to hold.

6 Conclusion

We study a robust contracting problem in which the principal chooses a task allocation and a compensation scheme to incentivize agents to acquire information. Task allocation plays a dual role: It determines the kind of information agents generate and the effectiveness of the cross-verification compensation scheme the principal employs. Consequently, upon designing a contract, the principal faces the trade-off between diversifying a task allocation and reducing agency costs. The resulting optimal contract entails a new distortion whereby the principal focuses on learning about relatively familiar problems and gives up learning about problems that are ex ante novel and less understood. Also, to provide robust incentives, the principal adopts a one-sided, chain-like compensation scheme. Our results imply that—in the presence of moral hazard—the organizational boundaries tend to be close to the principal's core expertise or, interpreted differently, the organization exhibits inefficiently low diversity of agents' expertise.

References

Aghion, P. and J. Tirole (1997). Formal and real authority in organizations. *Journal of Political Economy 105*(1), 1–29.

Aybas, Y. C. and S. Callander (2023). Efficient cheap talk in complex environments. Technical report, Stanford University.

Azrieli, Y. (2021). Monitoring experts. *Theoretical Economics 16*(4), 1313–1350.

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*, 1007–1028.

Bohren, J. A. and T. Kravitz (2019). Peer monitoring with partial commitment.

Callander, S. (2011). Searching and learning by trial and error. *American Economic Review 101*(6), 2277–2308.

Callander, S. and T. S. Clark (2017). Precedent and doctrine in a complicated world. *American Political Science Review 111*(1), 184–203.

Callander, S. and B. Harstad (2015). Experimentation in federal systems. *Quarterly Journal of Economics 130*(2), 951–1002.

Camboni, M. and M. Porcellacchia (2023). Monitoring team members: Information waste and the transparency trap. *Available at SSRN 4648556*.

Carnehl, C. and J. Schneider (2025). A quest for knowledge. *Econometrica 93*(2), 623–659.

Cetemen, D., C. Urgun, and L. Yariv (2023). Collective progress: Dynamics of exit waves. *Journal of Political Economy 131*(9), 2402–2450.

Deb, R., M. M. Pai, and M. Said (2018). Evaluating strategic forecasters. *American Economic Review 108*(10), 3057–3103.

Dong, M. and T. Mayskaya (2024). Does reducing communication barriers promote diversity? *Journal of Economic Theory 222*, 105932.

Fudenberg, D. and J. Tirole (1991). *Game Theory.* MIT Press.

Garfagnini, U. and B. Strulovici (2016). Social experimentation with interdependent and expanding technologies. *Review of Economic Studies 83*(4), 1579–1613.

Garicano, L. (2000). Hierarchies and the organization of knowledge in production. *Journal of Political Economy 108*(5), 874–904.

Garicano, L. and E. Rossi-Hansberg (2015). Knowledge-based hierarchies: Using organizations to understand the economy. *Annual Review of Economics 7*(1), 1–30.

Gromb, D. and D. Martimort (2007). Collusion and the organization of delegated expertise. *Journal of Economic Theory 137*(1), 271–299.

Halac, M. (2025). Contracting for coordination. *Journal of the European Economic Association*, jvaf015.

Halac, M., I. Kremer, and E. Winter (2024). Monitoring teams. *American Economic Journal: Microeconomics 16*(3), 134–161.

Halac, M., E. Lipnowski, and D. Rappoport (2021). Rank uncertainty in organizations. *American Economic Review 111*(3), 757–786.

Holmström, B. (1982). Moral hazard in teams. *Bell Journal of Economics*, 324–340.

Jovanovic, B. and R. Rob (1990). Long waves and short waves: Growth through intensive and extensive search. *Econometrica*, 1391–1409.

Legros, P. and S. A. Matthews (1993). Efficient and nearly-efficient partnerships. *Review of Economic Studies 60*(3), 599–611.

Mas-Colell, A., M. D. Whinston, and J. R. Green (1995). *Microeconomic theory.* Oxford University Press.

Page, S. E. (2017). *The Diversity Bonus: How Great Teams Pay Off in the Knowledge Economy*, Volume 2. Princeton University Press.

Pearce, D. G. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 1029–1050.

Prahalad, C. and G. Hamel (1990). The core competence of the corporation. *Harvard Business Review 68*(3), 79–91.

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

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

A Appendix: Omitted Proofs and Technical Details

A.1 Proof of Lemma 1

Take any task allocation $\tau = (t_1, ..., t_n)$, and suppose that all agents exert effort. Given a signal profile $\mathbf{s} = (x(0), x(t_1), ..., x(t_n))$, the principal's problem of choosing an action policy is

$$\min_p \mathbb{E}\left[\int_0^1 \mathbb{1}(p(t) \neq x(t)) \, \mathrm{d}t \mid \tau, \mathbf{s}\right].$$

The integrand $\mathbb{1}(p(t) \neq x(t))$ is bounded at each t. Thus, we can ignore the set of locations at which the principal's action changes—which has zero measure—and rewrite the problem as a pointwise optimization, i.e., $\min_{p(t)\in\{A,B\}} \mathbb{E}[\mathbb{1}(p(t) \neq x(t))|\tau, \mathbf{s}], \forall t \in [0, 1]$. Note that $\mathbb{E}[\mathbb{1}(p(t) \neq x(t))|\tau, \mathbf{s}] = \Pr(p(t) \neq x(t)|\tau, \mathbf{s})$.

Take any location $t \in [0, t_n]$, and let $i \in \{1, ..., n\}$ satisfy $t \in [t_{i-1}, t_i]$, where $t_0 = 0$. Due to the Markov property of the state distribution, probability $\Pr(p(t) \neq x(t)|\tau, \mathbf{s})$ depends only on $x(t_{i-1})$ and $x(t_i)$. Also, we have

$$\Pr(x(t) = x(t_i)|\tau, \mathbf{s}) \geq \Pr(x(t) = x(t_{i-1})|\tau, \mathbf{s}) \iff |t - t_i| \leq |t - t_{i-1}|.$$

Thus, the optimal action for location t is $p(t) = x(t^*)$, where $t^* \in \arg\min_{t' \in \{t_{i-1}, t_i\}} |t' - t| = \arg\min_{t' \in \{0, t_1, ..., t_n\}} |t' - t|$. This formula also applies to locations between t_n and 1, which do not have their right-hand neighbor.

We can write this optimal action policy explicitly as follows:

$$p(t) = \begin{cases} x(0) & \text{if} \quad t \in [0, \frac{t_1}{2}) \\ x(t_i) & \text{if} \quad t \in [\frac{t_{i-1}+t_i}{2}, \frac{t_i+t_{i+1}}{2}), \forall i = 1, 2, ..., n-1 \\ x(t_n) & \text{if} \quad t \in [\frac{t_{n-1}+t_n}{2}, 1] \end{cases} \tag{9}$$

with $t_0 = 0$. (To be precise, any optimal action policy is equal to (9) up to a measure-zero set of locations, which does not affect the information loss.)

The information loss from the optimal action policy is computed as

$$\begin{aligned} L(\tau) &= \int_0^{\frac{t_1}{2}} \Pr(x(0) \neq x(t))\ \mathrm{d}t + \sum_{i=1}^{n-1} \int_{\frac{t_{i-1}+t_i}{2}}^{\frac{t_i+t_{i+1}}{2}} \Pr(x(t_i) \neq x(t))\ \mathrm{d}t + \int_{\frac{t_{n-1}+t_n}{2}}^{1} \Pr(x(t_n) \neq x(t))\ \mathrm{d}t \\ &= \frac{1}{2} - \frac{2n+1}{2\lambda} + \frac{1}{\lambda}\left[\sum_{i=1}^{n} e^{-\frac{\lambda(t_i - t_{i-1})}{2}} + \frac{1}{2}e^{-\lambda(1-t_n)}\right]. \end{aligned}$$

Therefore, we obtain expression (3). □

A.2 Proof of Proposition 1

The information loss (3) in Lemma 1 is strictly convex and symmetric in $\Delta t_i = t_i - t_{i-1}, \forall i = 1, ..., n$. Thus, for a given t_n, we should choose $t_1, ..., t_{n-1}$ to equalize Δt_i. Let Δt be the equalized value, which satisfies $n\Delta t = t_n$. Plugging $t_n = n\Delta t$ and $t_i - t_{i-1} = \Delta t$ into the information loss and minimizing it with respect to Δt, we obtain $\Delta t^{\dagger} = \frac{2}{2n+1}$, or equivalently, $t_i^{\dagger} = \frac{2i}{2n+1}, \forall i = 1, 2, ..., n$. We can then 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 $y^{\dagger} \triangleq e^{-\frac{\lambda}{2n+1}} < 1$. Using the envelope theorem, we obtain $\frac{\partial L(\tau^{\dagger})}{\partial n} = \frac{1}{\lambda}\left(y^{\dagger} - 1\right) < 0$ and $\frac{\partial L(\tau^{\dagger})}{\partial \lambda} = -\frac{2n+1}{2\lambda^2}\left(y^{\dagger} - 1\right) > 0$. □

A.3 Proof of Proposition 2

Take any RIE contract. Without loss, assume that in each step of its iterated elimination of strictly dominated strategies (IESDS), exactly one agent eliminates shirking. We can then identify the order of elimination with a permutation of agents' names, generically written as $\mathbf{i} = (i_1, ..., i_n)$, i.e., agent i_k is the k-th agent to eliminate shirking. We call any such permutation $\mathbf{i}$ an *elimination order*. Given any elimination order $\mathbf{i}$, we say that an agent i is k^{th}*-ranked* if $i = i_k$. Agent i_j is *risk-free for agent* i_k if $j < k$, i.e., agent i_j eliminates shirking earlier than agent i_k according to order $\mathbf{i}$. Recall that without loss, we focus on task allocations τ such that $t_1 \leq t_2 \leq \cdots \leq t_n$.

The proof consists of four steps. In **Step 1**, we fix an arbitrary elimination order, $\mathbf{i} = (i_1, ..., i_n)$. Then, for each $k = 1, ..., n$, we identify a lower bound of expected payments for agent i_k across all RIE contracts that are consistent with elimination order $\mathbf{i}$. By adding up these lower bounds across agents, we obtain a lower bound of total expected payments across all RIE contracts consistent with elimination order $\mathbf{i}$. In **Step 2**, we minimize the lower bound of Step 1 across all elimination orders, which yields a lower bound of total payments across all RIE contracts. In **Step 3**, we show that, as $\epsilon \to 0$, compensation scheme w^{ϵ} defined by (5) attains the lower bound from Step 2. Finally, in **Step 4**, we show that for any $\epsilon > 0$, compensation scheme w^{ϵ} is an RIE contract.

Step 1: Lower bound of expected payment for a given elimination order Throughout this step, we fix an arbitrary elimination order, $\mathbf{i} = (i_1, ..., i_n)$. Unless otherwise stated, we use "lower bound for agent i_k" to mean a lower bound of expected payments for agent i_k across all

RIE contracts that are consistent with elimination order $\mathbf{i}$, where the expected payments are evaluated at all agents exerting effort.

This step consists of four substeps: **Step 1.1** derives a lower bound for agent i_1. **Step 1.2** restricts the class of contracts we need to focus on to derive a lower bound for other agents. **Step 1.3** derives a lower bound for agent i_2, and **Step 1.4** derives that for the rest of agents.

Step 1.1: Lower bound of payment for agent i_1 Let $\mathcal{W}(\tau, \mathbf{i})$ denote the set of all RIE contracts that admit elimination order $\mathbf{i}$. In the game $\Gamma(\tau, w)$ with any contract $w \in \mathcal{W}(\tau, \mathbf{i})$, agent i_1 strictly prefers to work regardless of the strategies of other agents and the noisemaker.

Specifically, because agent i_1 is the first to eliminate shirking, any other agent $j \neq i_1$ has both working $a_j = 1$ and shirking $a_j = 0$ uneliminated. For any agent j who works, her signal is the local state of her assigned location, i.e., $s_j = x(t_j)$. For any agent $j \neq i_1$ who shirks, the noisemaker's choice, $\mathbf{y} = (y_1, ..., y_n)$, determines her signal realization as $s_j = y_j \in \{A, B\}$.

Given any combination of other agents' strategies $\mathbf{a}_{-i_1} \in \mathcal{A}^*_{-i_1} \triangleq \{0, 1\}^{n-1}$ and the noisemaker's choice $\mathbf{y} \in \{A, B\}^n$, any contract $w \in \mathcal{W}(\tau, \mathbf{i})$ must be such that agent i_1 earns a strictly higher payoff by working than shirking. Thus, the compensation scheme must satisfy the following set of inequalities:

$$\sum_{\mathbf{s}} w_{i_1}(\mathbf{s}) \Pr(\mathbf{s}|a_{i_1} = 1, \mathbf{a}_{-i_1}, \mathbf{y}) - c > \sum_{\mathbf{s}} w_{i_1}(\mathbf{s}) \Pr(\mathbf{s}|a_{i_1} = 0, \mathbf{a}_{-i_1}, \mathbf{y}), \forall \mathbf{a}_{-i_1} \in \mathcal{A}^*_{-i_1}, \forall \mathbf{y} \in \{A, B\}^n, \tag{10}$$

where $\Pr(\mathbf{s}|a_{i_1}, \mathbf{a}_{-i_1}, \mathbf{y})$ is the probability of signal profile $\mathbf{s}$ given agents' strategies and the noisemaker's choice.

The goal of this substep is to find a lower bound of the expected payment for agent i_1 across all contracts that satisfy (10), where the expected payment is evaluated at all agents working. (Hereafter, whenever we say "payment" it means the payment evaluated at all agents working.)

Specifically, we find a lower bound by minimizing the expected payment for agent i_1 subject to weaker constraints that modify (10) as follows: (i) imposing only the constraints associated with all other agents shirking, i.e., $\mathbf{a}_{-i_1} = \mathbf{0}$; and (ii) replacing strict inequalities with weak ones. Meanwhile, it proves convenient to write the noisemaker's choice separately for agent i_1

as $y_{i_1} \in \{A, B\}$ and for other agents as $\mathbf{y}_{-i_1} \in \{A, B\}^{n-1}$.

The weaker constraints that reflect Points (i) and (ii) above are written as:

$$\sum_{\mathbf{s}} w_{i_1}(\mathbf{s}) \Pr(\mathbf{s}|a_{i_1} = 1, \mathbf{a}_{-i_1} = \mathbf{0}, y_{i_1} = A, \mathbf{y}_{-i_1}) - c \geq \sum_{\mathbf{s}} w_{i_1}(\mathbf{s}) \Pr(\mathbf{s}|a_{i_1} = 0, \mathbf{a}_{-i_1} = \mathbf{0}, y_{i_1} = A, \mathbf{y}_{-i_1}) \tag{11}$$

$$\sum_{\mathbf{s}} w_{i_1}(\mathbf{s}) \Pr(\mathbf{s}|a_{i_1} = 1, \mathbf{a}_{-i_1} = \mathbf{0}, y_{i_1} = B, \mathbf{y}_{-i_1}) - c \geq \sum_{\mathbf{s}} w_{i_1}(\mathbf{s}) \Pr(\mathbf{s}|a_{i_1} = 0, \mathbf{a}_{-i_1} = \mathbf{0}, y_{i_1} = B, \mathbf{y}_{-i_1}) \tag{12}$$

for all $\mathbf{y}_{-i_1} \in \{A, B\}^{n-1}$. In Appendix A.4, we show that, for any fixed $\mathbf{y}_{-i_1} \in \{A, B\}^{n-1}$, constraints (11) and (12) are equivalent to, respectively,

$$\frac{1}{4}(1 - e^{-\lambda t_{i_1}})[w_{i_1}(AB) - w_{i_1}(AA)] + \frac{1}{4}(1 + e^{-\lambda t_{i_1}})[w_{i_1}(BB) - w_{i_1}(BA)] \geq c; \quad \text{and} \quad (\text{ICA}_{\mathbf{y}_{-i_1}})$$
$$\frac{1}{4}(1 - e^{-\lambda t_{i_1}})[w_{i_1}(BA) - w_{i_1}(BB)] + \frac{1}{4}(1 + e^{-\lambda t_{i_1}})[w_{i_1}(AA) - w_{i_1}(AB)] \geq c, \quad (\text{ICB}_{\mathbf{y}_{-i_1}})$$

where we adopt the shorthand notation $w_{i_1}(s_0 s_{i_1}) := w_{i_1}(s_0 s_{i_1}; \mathbf{y}_{-i_1})$ for each $s_0 s_{i_1} \in \{AA, AB, BA, BB\}$.

We show that at the optimum—i.e., at the contract that minimizes the payment for agent i_1 given these constraints—the constraints must be binding. Suppose, for example, that $\text{ICA}_{\mathbf{y}_{-i_1}}$ is slack, then we have $w_{i_1}(AB) > 0$ or $w_{i_1}(BB) > 0$. If, e.g., $w_{i_1}(AB) > 0$, then the principal can decrease $w_{i_1}(AB)$ by some small $\epsilon > 0$, such that $\text{ICA}_{\mathbf{y}_{-i_1}}$ is still slack, $\text{ICB}_{\mathbf{y}_{-i_1}}$ is relaxed, and the expected payment for i_1 strictly decreases. A similar argument works for the case in which $\text{ICB}_{\mathbf{y}_{-i_1}}$ is slack.

Given that constraints $\text{ICA}_{\mathbf{y}_{-i_1}}$ and $\text{ICB}_{\mathbf{y}_{-i_1}}$ bind at the optimum, the system of equations uniquely solves $w_{i_1}(AA) - w_{i_1}(AB) = w_{i_1}(BB) - w_{i_1}(BA) = 2ce^{\lambda t_{i_1}}$. The principal should then reduce $w_{i_1}(AB)$ and $w_{i_1}(BA)$ as much as possible, until the limited liability constraint binds, i.e., $w_{i_1}(AB) = w_{i_1}(BA) = 0$. We then obtain $w_{i_1}(AA) = w_{i_1}(BB) = 2ce^{\lambda t_{i_1}}$.

The procedure works for each $\mathbf{y}_{-i_1}$. Note that although $\mathbf{y}_{-i_1}$ denotes the noisemaker's choice, the resulting payment rule, such as $w_{i_1}(s_0 s_{i_1}; \mathbf{y}_{-i_1})$, pins down the payment to agent i_1 when the realized signals for other agents are $\mathbf{s}_{-i_1} = \mathbf{y}_{-i_1} \in \{A, B\}^{n-1}$.

The resulting contract takes the following simple form: Regardless of the realized signals of other agents, agent i_1 earns $2ce^{\lambda t_{i_1}}$ if $s_0 = s_{i_1}$ and 0 otherwise. Formally, we obtain $w^*_{i_1}(AA;\mathbf{s}_{-i_1}) = w^*_{i_1}(BB;\mathbf{s}_{-i_1}) = 2ce^{\lambda t_{i_1}}$ and $w^*_{i_1}(AB;\mathbf{s}_{-i_1}) = w^*_{i_1}(BA;\mathbf{s}_{-i_1}) = 0$ for all $\mathbf{s}_{-i_1} \in \{A,B\}^{n-1}$. The corresponding lower bound of payment for agent i_1 is

$$\hat{K}^{\mathbf{i}}_{i_1}(\tau, w^*) = \sum_{\mathbf{s}\in\{A,B\}^{n+1}} w^*_{i_1}(\mathbf{s}) \Pr(\mathbf{s}|\tau, \mathbf{a}=\mathbf{1}) = 2ce^{\lambda t_{i_1}} \cdot \Pr(x_0 = x_{i_1}) = c(1+e^{\lambda t_{i_1}}).$$

Step 1.2: Simplifying contracts To characterize lower bounds for other agents, we first simplify the class of contracts we need to focus on. Recall that for a given $k \geq 2$, any agent i_j with $j < k$ is risk-free for agent i_k in that agent i_j eliminated shirking in an earlier step of the IESDS than agent i_k. Any agent i_j with $j > k$ is risky in that she has not eliminated shirking in the k-th round. Hereafter, we use "risky" and "risk-free" from the perspective of agent i_k.

We divide the signals of agents other than agent i_k as follows. Let $\mathbf{s}_{risky} \in \{A,B\}^{n-k}$ be the signal profile of risky agents. Next, we select at most two agents as agent i_k's *risk-free neighbors* in the following way:

(A) select one risk-free agent who is (i) located on the left side of agent i_k and (ii) closest to agent i_k compared to any other risk-free agents located on the left side of agent i_k.[11] If there is no agent on agent i_k's left side, select the principal as a risk-free neighbor of this category; and

(B) select one risk-free agent who is (i) not selected according to (A), (ii) located on the right-side of agent i_k, and (iii) closest to agent i_k compared to any other risk-free agents located on the right side of agent i_k.

The selection is arbitrary if multiple agents satisfy (A) or (B). Depending on the elimination order, agent i_k may not have a risk-free neighbor on her right side, in which case no agent satisfies (B). We call the players who satisfy (A) or (B) agent i_k's risk-free neighbors. Let $\mathbf{s}_{near}$ denote the signals of agent i_k's risk-free neighbors. Finally, let $\mathbf{s}_{far}$ be the signal profile of agent i_k's risk-free *non-neighbors*, i.e., agents who are risk-free but not risk-free neighbors.

[11] We use the "left side of agent i_k" to mean any location $t \leq t_{i_k}$. Similarly, the "right side of agent i_k" refers to any location $t \geq t_{i_k}$.

We show that upon finding a lower bound for agent i_k, we can focus on contracts that do not use the signals of risk-free non-neighbors. Recall that $a_{i_k} = 1$ denotes agent i_k's choice to work. With a slight abuse of notation, we write $a_{i_k} = A$ or $a_{i_k} = B$ for agent i_k's shirking and the noisemaker's choosing $y_{i_k} = A$ or $y_{i_k} = B$, respectively.

Take any compensation scheme for agent i_k, w_{i_k}. Then, for any signal profile, $(s_{i_k}, \mathbf{s}_{near}, \mathbf{s}_{far}, \mathbf{s}_{risky})$, define $\hat{w}_{i_k}$ as $\hat{w}_{i_k}(s_{i_k}, \mathbf{s}_{near}, \mathbf{s}_{far}; \mathbf{s}_{risky}) \triangleq \mathbb{E}_{\tilde{\mathbf{s}}_{far}}[w_{i_k}(s_{i_k}, \mathbf{s}_{near}, \tilde{\mathbf{s}}_{far}; \mathbf{s}_{risky})|\mathbf{s}_{near}]$. That is, we construct $\hat{w}_{i_k}$ by taking expectation with respect to $\mathbf{s}_{far}$ conditional on $\mathbf{s}_{near}$. The resulting $\hat{w}_{i_k}$ is independent of the signals $\mathbf{s}_{far}$ of agent i_k's risk-free non-neighbors.

Furthermore, suppose that (i) agent i_k chooses action a_{i_k}; (ii) all risk-free agents work; and (iii) all risky agents shirk and their realized signals (i.e., the noisemaker's choices) are $\mathbf{s}_{risky}$ with probability 1. Then, the expected payments for agent i_k under these two compensation schemes are equal, because

$$\begin{aligned}
\mathbb{E}[w_{i_k}(s_{i_k}, \mathbf{s}_{near}, \mathbf{s}_{far}; \mathbf{s}_{risky})|a_{i_k}] &= \mathbb{E}[\mathbb{E}_{\tilde{\mathbf{s}}_{far}}[w_{i_k}(s_{i_k}, \mathbf{s}_{near}, \tilde{\mathbf{s}}_{far}; \mathbf{s}_{risky})|a_{i_k}, \mathbf{s}_{near}]|a_{i_k}] \\
&= \mathbb{E}[\mathbb{E}_{\tilde{\mathbf{s}}_{far}}[w_{i_k}(s_{i_k}, \mathbf{s}_{near}, \tilde{\mathbf{s}}_{far}; \mathbf{s}_{risky})|\mathbf{s}_{near}]|a_{i_k}] \\
&= \mathbb{E}[\hat{w}_{i_k}(s_{i_k}, \mathbf{s}_{near}, \mathbf{s}_{far}; \mathbf{s}_{risky})|a_{i_k}].
\end{aligned}$$

where the first equality is the law of iterated expectation and the second equality follows from the Markov property. Therefore, it is without loss to focus on compensation schemes that do not use risk-free non-neighbor's signals, $\mathbf{s}_{far}$.

Step 1.3: Lower bound of payment for agent i_2 We now find a lower bound of payment to agent i_2. The relevant constraint is that agent i_2 must strictly prefer to work (i) regardless of the strategies of the noisemaker and the agents other than i_1 and i_2 but (ii) provided agent i_1 works.

As is in Step 1, we only consider the constraints such that $\mathbf{a}_{-(i_1,i_2)} = \mathbf{0} \in \mathcal{A}^*_{-(i_1,i_2)}$ (i.e., all agents other than i_1 and i_2 will shirk) and replace strict inequalities with weak ones. This modification relaxes constraints and reduces the lower bound for i_2. These weaker constraints are written as follows:

$$\sum_{\mathbf{s}} w_{i_2}(\mathbf{s}) \Pr(\mathbf{s}|a_{i_2}=1, a_{i_1}=1, \mathbf{a}_{-(i_1,i_2)}=\mathbf{0}, \mathbf{y}) - c \geq \sum_{\mathbf{s}} w_{i_2}(\mathbf{s}) \Pr(\mathbf{s}|a_{i_2}=0, a_{i_1}=1, \mathbf{a}_{-(i_1,i_2)}=\mathbf{0}, \mathbf{y})$$

for every $y_{i_2} \in \{A, B\}$ and every $\mathbf{y}_{-i_2} \in \{A, B\}^{n-1}$.

The same argument as that in **Step 1.1** ensures that all of these constraints bind at the payment-minimizing contract. Let $ICA_{\mathbf{y}_{-i_2}}$ denote this equality constraint with $y_{i_2} = A$ and $\mathbf{y}_{-i_2}$. Denote $ICB_{\mathbf{y}_{-i_2}}$ analogously.

Depending on the task allocation, agent i_2 has one or two risk-free neighbors. Correspondingly, we consider two cases: (1) $0 \leq t_{i_1} < t_{i_2}$ (one risk-free neighbor, "chain") and (2) $0 \leq t_{i_2} \leq t_{i_1}$ (two risk-free neighbors, "sandwich").

The chain case (1) reduces to **Step 1.1**, the derivation of i_1's lower bound. Indeed, the principal is a risk-free non-neighbor who is (weakly) further away from agent i_2 relative to agent i_1. By **Step 1.2**, it is without loss to use compensation schemes whose payment for agent i_2 does not depend on the principal's signal, s_0. The problem is then identical with **Step 1.1**, where agent i_1's signal in the current problem plays the role of the principal's signal in **Step 1.1**. As a result, the compensation scheme that attains the lower bound pays agent i_2 if and only if her signal $s_{i_2} = s_{i_1}$. The lower bound for agent i_2 is $\hat{K}_{i_2}(\tau, w^0) = c(1 + e^{\lambda(t_{i_2} - t_{i_1})})$.

We now consider the sandwich case ($0 \leq t_{i_2} \leq t_{i_1}$). Agent i_2 has risk-free neighbors on both sides: The principal on the left and agent i_1 on the right. In Appendix A.4, we show that, for any fixed $\mathbf{y}_{-i_2} \in \{A, B\}^{n-1}$, the two binding constraints, $\text{ICA}_{\mathbf{y}_{-i_2}}$ and $\text{ICB}_{\mathbf{y}_{-i_2}}$, can be written as

$$\begin{aligned}&\frac{1}{8}P_1^- P_2^- [w_{i_2}(ABA) - w_{i_2}(AAA)] + \frac{1}{8}P_1^- P_2^+ [w_{i_2}(ABB) - w_{i_2}(AAB)] \\ &\qquad + \frac{1}{8}P_1^+ P_2^- [w_{i_2}(BBA) - w_{i_2}(BAA)] + \frac{1}{8}P_1^+ P_2^+ [w_{i_2}(BBB) - w_{i_2}(BAB)] = c,\end{aligned}$$

$$\begin{aligned}&\frac{1}{8}P_1^+ P_2^+ [w_{i_2}(AAA) - w_{i_2}(ABA)] + \frac{1}{8}P_1^+ P_2^- [w_{i_2}(AAB) - w_{i_2}(ABB)] \\ &\qquad + \frac{1}{8}P_1^- P_2^+ [w_{i_2}(BAA) - w_{i_2}(BBA)] + \frac{1}{8}P_1^- P_2^- [w_{i_2}(BAB) - w_{i_2}(BBB)] = c\end{aligned}$$

where we use the shorthand notation, $w_{i_2}(s_0 s_{i_2} s_{i_1}) := w_{i_2}(s_0 s_{i_2} s_{i_1}; \mathbf{y}_{-i_2})$ as well as $P_1^+ = 1 + e^{-\lambda t_{i_2}}$, $P_1^- = 1 - e^{-\lambda t_{i_2}}$, $P_2^+ = 1 + e^{-\lambda(t_{i_1} - t_{i_2})}$, and $P_2^- = 1 - e^{-\lambda(t_{i_1} - t_{i_2})}$.

The above system of equations can be reduced to one equation by substituting $w_{i_2}(BAB) - w_{i_2}(BBB)$ in $\text{ICA}_{\mathbf{y}_{-i_2}}$ using $\text{ICB}_{\mathbf{y}_{-i_2}}$, which yields

$$\begin{aligned}
&\frac{1}{8}[(P_1^+ P_2^+)^2 - (P_1^- P_2^-)^2][w_{i_2}(AAA) - w_{i_2}(ABA)] \\
&\qquad + \frac{1}{8}[P_2^+ P_2^- ((P_1^+)^2 - (P_1^-)^2)][w_{i_2}(AAB) - w_{i_2}(ABB)] \\
&\qquad + \frac{1}{8}[P_1^+ P_1^- ((P_2^+)^2 - (P_2^-)^2)][w_{i_2}(BAA) - w_{i_2}(BBA)] = c[P_1^- P_2^- + P_1^+ P_2^+]
\end{aligned}$$

for any fixed $\mathbf{y}_{-i_2}$.

Recall that we are minimizing $\hat{K}_{i_2}(\tau, w) = \sum_{\mathbf{s} \in \{A,B\}^{n+1}} w_{i_2}(\mathbf{s}) \Pr(\mathbf{s}|\mathbf{a} = \mathbf{1})$ subject to the above constraint. This is a standard linear programming, whose optimal solution is pinned down by comparing the ratios between the coefficients of the variables in the objective (i.e., $\{\Pr(\mathbf{s}|\mathbf{a} = \mathbf{1})\}_{\mathbf{s} \in \{A,B\}^{n+1}}$) and the coefficients in the constraint.

First, it is optimal to set $w_{i_2}^*(ABA) = w_{i_2}^*(ABB) = w_{i_2}^*(BBA) = 0$ because reducing these wages allows us to reduce wages for other signal profiles.

Second, among the remaining three (i.e., $w_{i_2}(AAA)$, $w_{i_2}(AAB)$, and $w_{i_2}(BAA)$), $w_{i_2}(AAA)$ is the only signal profile that takes a positive value, because the relevant ratio (i.e., its coefficient in the objective to be minimized, divided by its coefficient in the constraint) is minimized:

$$\frac{\Pr(AAA|\mathbf{a} = \mathbf{1})}{\frac{1}{8}[(P_1^+ P_2^+)^2 - (P_1^- P_2^-)^2]} \leq \min\left\{ \frac{\Pr(AAB|\mathbf{a} = \mathbf{1})}{\frac{1}{8}[P_2^+ P_2^- ((P_1^+)^2 - (P_1^-)^2)]}, \frac{\Pr(BAA|\mathbf{a} = \mathbf{1})}{\frac{1}{8}[P_1^+ P_1^- ((P_2^+)^2 - (P_2^-)^2)]} \right\}.$$

Appendix A.4 proves this inequality based on direct calculation.

Hence, we conclude that $w_{i_2}^*(AAB) = w_{i_2}^*(BAA) = 0$, and $w_{i_2}^*(AAA) = \frac{4c}{e^{-\lambda t_{i_2}} + e^{-\lambda(t_{i_1} - t_{i_2})}}$. Combining the above solution with $\text{ICB}_{\mathbf{y}_{-i_2}}$, we get $w_{i_2}^*(BBB) = w_{i_2}^*(AAA)$, $w_{i_2}^*(BAB) = 0$. Also, the same argument applies to any $\mathbf{y}_{-i_2}$.

Step 1.4: Lower bound of payment for agent i_k with $k \geq 3$ Consider any agent i_k with $k \geq 3$. By **Step 1.2**, we can focus on contracts such that the payment to agent i_k does not depend on signals of her risk-free non-neighbors. Suppose that agent i_k has only one risk-free neighbor, agent k_-, on her left side. In this case, the problem of finding a lower bound of payment for agent i_k is exactly the same as the problem for agent i_2 with $0 \leq t_{i_1} < t_{i_2}$, where agent t_{i_1} plays the role of agent k_-. Thus, the lower bound for agent i_k is

$$\hat{K}^c_{i_k}(\tau, w^*) = c(1 + e^{\lambda(t_{i_k} - t_{k_-})}). \tag{13}$$

If i_k has two risk-free neighbors, agent k_- on her left and agent k_+ on her right, then the problem of finding a lower bound for agent i_k is the same as the problem for agent i_2 with $0 \leq t_{i_2} \leq t_{i_2}$, where the principal plays a role of agent k_- and agent i_1 plays a role of agent k_+. Thus, agent i_k's lower bound is

$$\hat{K}^s_{i_k}(\tau, w^*) = \frac{4c}{e^{-\lambda(t_{i_k} - t_{k_-})} + e^{-\lambda(t_{k_+} - t_{i_k})}} \cdot \Pr(x_{k_-} = x_{i_k} = x_{k_+}) = c\left(1 + \frac{1 + e^{-\lambda(t_{k_+} - t_{k_-})}}{e^{-\lambda(t_{i_k} - t_{k_-})} + e^{-\lambda(t_{k_+} - t_{i_k})}}\right).$$

Therefore, for any given elimination order, the lower bound of the total expected payment is $\hat{K}^{\mathbf{i}}(\tau, w^*) = \sum_{k=1}^{n} \hat{K}^{\mathbf{i}}_{i_k}(\tau, w^*)$.

Step 2: Identity permutation minimizes the lower bound. In **Step 1**, we have derived a lower bound of the total expected payment for any fixed elimination order $\mathbf{i} = (i_1, .., i_n)$. In this step, we show that for any task allocation τ, the identity permutation of the elimination order, i.e., $\mathbf{i}^c = (1, 2, ..., n)$, minimizes the lower bound.

First, suppose there are two agents, $n = 2$. There are only two permutations: Agent 1 is 1^{st}-ranked (elimination order $(1, 2)$, "chain") or 2^{nd}-ranked (elimination order $(2, 1)$, "sandwich"). The lower bound of expected payment under the two permutations are respectively:

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

For a fixed t_2, the two expected payments are equal when $t_1 = 0$ and $t_1 = t_2$. Also, $\hat{K}^s$ increases as t_1 approaches $t_2/2$ from left or right, while the same change decreases $\hat{K}^c$. Thus, for any $t_1, t_2 \in [0, 1]$ such that $t_1 \leq t_2$, we have $\hat{K}^s(t_1, t_2) \geq \hat{K}^s(0, t_2) = \hat{K}^c(0, t_2) \geq \hat{K}^c(t_1, t_2)$. Therefore, the chain elimination order attains a lower value of the lower bound than the sandwich elimination order, or equivalently, the identity permutation order minimizes the lower bound for $n = 2$.

Suppose now that there are $n \geq 3$ agents. Let $w^{\mathbf{i}}$ denote the contract that attains the lower bound of payment under elimination order $\mathbf{i}$. Suppose that elimination order $\mathbf{i}$ is different from the identity permutation order $\mathbf{i}^c$. Let k^* be the smallest step at which $\mathbf{i}$ deviates from $\mathbf{i}^c$, i.e., for any $j \leq k^* - 1$, $i_j = j$, but $i_{k^*} > k^*$. Then, we divide agents into three sets: $A = \{1, ..., k^* - 1\}$; $B = \{k^*, ..., i_{k^*}\}$; and $C = \{i_{k^*} + 1, ..., n\}$. First, we modify the elimination order $\mathbf{i}$ so that agents in A are the first $k^* - 1$ agents to eliminate shirking; agents in C are the last $n - i_{k^*}$ agents to eliminate shirking; and agents in $B = \{k^*, ..., i_{k^*}\}$ eliminate shirking after A but before C; also, within B, the elimination order is consistent with $\mathbf{i}$. This modification does not change the lower bound of payment, because the set of risk-free neighbors for each agent does not change at each step of elimination.

We show that the lower bound decreases if we modify payments and elimination order for some of the agents in set B. The modification will not affect agents in A or C. Thus, to simplify notation, we assume that sets A and C are empty, i.e., $k^* = 1$ and $i_{k^*} = n$. Define $k = i_2 < n$. Then, the principal ($t = 0$), agent k, and agent n form the sandwich case of $n = 2$. The principal can then reduce the lower bound of total payment by changing the order of elimination so that $i_1 = k$ and $i_2 = n$, with the rest of the elimination order following the original order. Thus, if a contract attains a lower bound for an elimination order that is different from the identify permutation, then we can modify the contract to strictly decrease the total payment. Therefore, the lower bound of the total payment is minimized by the elimination order $\mathbf{i}^c = (1, 2, ..., n)$. The corresponding payment for each agent is given by (13), so the lower bound of the total expected payment is $\hat{K}^{\mathbf{i}^c}(\tau, w^*) = c \sum_{i=1}^{n} (1 + e^{\lambda(t_i - t_{i-1})})$.

Step 3: w^ϵ in (5) attains the lower bound as $\epsilon \to 0$. By the definition of w^ϵ, we have

$$\hat{K}(\tau, w^\epsilon) = \sum_{i=1}^{n} \sum_{\mathbf{s} \in \{A,B\}^{n+1}} w_i^\epsilon(\mathbf{s}) \Pr(\mathbf{s}|\tau, \mathbf{a} = \mathbf{1}) = \sum_{i=1}^{n} c(1 + e^{\lambda(t_i - t_{i-1})}) + \epsilon \sum_{i=1}^{n} \Pr(x_i = x_{i-1})$$

Thus, we have $\lim_{\epsilon \to 0} \hat{K}(\tau, w^\epsilon) = \sum_{i=1}^{n} c(1 + e^{\lambda(t_i - t_{i-1})})$, i.e., the lower bound in **Step 2**.

Step 4: w^ϵ is an RIE Contract for any $\epsilon > 0$. We show that under compensation scheme w^ϵ with $\epsilon > 0$, agent k is the k-th player to eliminate shirking in the IESDS. First, agent 1 is the first player to eliminate shirking. Indeed, facing contract w^0, agent 1 weakly prefers to work regardless of the strategies of other agents and the noisemaker as shown in **Step 1.1**. Equivalently, by working instead of shirking, agent 1 can increase her expected compensation by more than her effort cost, c. Now, facing w^ϵ with $\epsilon > 0$, the increment of the expected compensation due to working is even higher, because the probability of extra bonus $\epsilon > 0$ is strictly higher when agent 1 works (note that regardless of the noisemaker's choice, the probability of $s_0 = s_1$ is 0.5 once agent 1 shirks). Thus, agent 1 strictly prefers to work regardless of the strategies of other agents and the noisemaker.

We can apply a similar argument to show that agent k is the k-th player to eliminate shirking. To prove this, for each $k = 2, ..., n$, we replace the principal and agent 1 in the previous paragraph with agent $k-1$ and agent k, respectively.

Summary We have shown that for any $\epsilon > 0$, contract w^ϵ is an RIE contract. Moreover, as $\epsilon \to 0$, the total expected payment under w^ϵ converges to the lower bound of total payments across all RIE contracts. Therefore, the limiting payment must be equal to the incentive cost, i.e., equations in (6) hold. □

A.4 Detail of Calculation for the Proof of Proposition 2

In this section, we fill in the detailed calculations omitted from Appendix A.3. In **Step 1.1**, constraints (11) and (12) for agent i_1 can be rewritten as

$$\sum_{\mathbf{s}} w_{i_1}(\mathbf{s})[\Pr(\mathbf{s}|a_{i_1} = 1, \mathbf{a}_{-i_1} = \mathbf{0}, y_{i_1} = A, \mathbf{y}_{-i_1}) - \Pr(\mathbf{s}|a_{i_1} = 0, \mathbf{a}_{-i_1} = \mathbf{0}, y_{i_1} = A, \mathbf{y}_{-i_1})] \geq c \tag{ICA$_{\mathbf{y}_{-i_1}}$}$$

$$\sum_{\mathbf{s}} w_{i_1}(\mathbf{s})[\Pr(\mathbf{s}|a_{i_1} = 1, \mathbf{a}_{-i_1} = \mathbf{0}, y_{i_1} = B, \mathbf{y}_{-i_1}) - \Pr(\mathbf{s}|a_{i_1} = 0, \mathbf{a}_{-i_1} = \mathbf{0}, y_{i_1} = B, \mathbf{y}_{-i_1})] \geq c \tag{ICB$_{\mathbf{y}_{-i_1}}$}$$

for all $\mathbf{y}_{-i_1} \in \{A, B\}^{n-1}$. Write $w_{i_1}(s_0 s_{i_1}) := w_{i_1}(s_0 s_{i_1}; \mathbf{y}_{-i_1})$ for every fixed $\mathbf{y}_{-i_1}$. Fixing $y_{i_1} = A$, we can enumerate $[\Pr(\mathbf{s}|a_{i_1} = 1, \mathbf{a}_{-i_1} = \mathbf{0}, \mathbf{y}) - \Pr(\mathbf{s}|a_{i_1} = 0, \mathbf{a}_{-i_1} = \mathbf{0}, \mathbf{y})]$ associated with all combinations of $s_0 s_{i_1}$ as the followings:

$$\Pr(AA;\mathbf{s}_{-i_1}|a_{i_1}=1,\mathbf{a}_{-i_1}=\mathbf{0},y_{i_1}=A,\mathbf{y}_{-i_1}) - \Pr(AA;\mathbf{s}_{-i_1}|a_{i_1}=0,\mathbf{a}_{-i_1}=\mathbf{0},y_{i_1}=A,\mathbf{y}_{-i_1})$$
$$= \underbrace{\frac{1}{2}}_{\Pr(s_0=A)} \times \underbrace{\frac{1}{2}(1+e^{-\lambda t_{i_1}})}_{\Pr(x(0)=x(t_{i_1}))} - \underbrace{\frac{1}{2}}_{\Pr(s_0=A)} \times \underbrace{1}_{\Pr(s_{i_1}=A|a_{i_1}=0,y_{i_1}=A)} = -\frac{1}{4}(1-e^{-\lambda t_{i_1}})$$

$$\Pr(AB;\mathbf{s}_{-i_1}|a_{i_1}=1,\mathbf{a}_{-i_1}=\mathbf{0},y_{i_1}=A,\mathbf{y}_{-i_1}) - \Pr(AB;\mathbf{s}_{-i_1}|a_{i_1}=0,\mathbf{a}_{-i_1}=\mathbf{0},y_{i_1}=A,\mathbf{y}_{-i_1})$$
$$= \frac{1}{2} \times \frac{1}{2}(1-e^{-\lambda t_{i_1}}) - \frac{1}{2} \times \underbrace{0}_{\Pr(s_{i_1}=B|a_{i_1}=0,y_{i_1}=A)} = \frac{1}{4}(1-e^{-\lambda t_{i_1}})$$

$$\Pr(BA;\mathbf{s}_{-i_1}|a_{i_1}=1,\mathbf{a}_{-i_1}=\mathbf{0},y_{i_1}=A,\mathbf{y}_{-i_1}) - \Pr(BA;\mathbf{s}_{-i_1}|a_{i_1}=0,\mathbf{a}_{-i_1}=\mathbf{0},y_{i_1}=A,\mathbf{y}_{-i_1})$$
$$= -\frac{1}{4}(1+e^{\lambda t_{i_1}})$$

$$\Pr(BB;\mathbf{s}_{-i_1}|a_{i_1}=1,\mathbf{a}_{-i_1}=\mathbf{0},y_{i_1}=A,\mathbf{y}_{-i_1}) - \Pr(BB;\mathbf{s}_{-i_1}|a_{i_1}=0,\mathbf{a}_{-i_1}=\mathbf{0},y_{i_1}=A,\mathbf{y}_{-i_1})$$
$$= \frac{1}{4}(1+e^{\lambda t_{i_1}})$$

Plugging into $\text{ICA}_{\mathbf{y}_{-i_1}}$ and rearranging, we get

$$\frac{1}{4}(1-e^{-\lambda t_{i_1}})[w_{i_1}(AB)-w_{i_1}(AA)] + \frac{1}{4}(1+e^{-\lambda t_{i_1}})[w_{i_1}(BB)-w_{i_1}(BA)] \geq c \qquad (\text{ICA}_{\mathbf{y}_{-i_1}})$$

The calculations for $\text{ICB}_{\mathbf{y}_{-i_1}}$ is symmetric.

In **Step 1.3**, agent i_2's compensation scheme must satisfy the following set of constraints:

$$\sum_{\mathbf{s}} w_{i_2}(\mathbf{s})[\Pr(\mathbf{s}|a_{i_2}=1,a_{i_1}=1,\mathbf{a}_{-(i_1,i_2)}=\mathbf{0},\mathbf{y}) - \Pr(\mathbf{s}|a_{i_2}=0,a_{i_1}=1,\mathbf{a}_{-(i_1,i_2)}=\mathbf{0},\mathbf{y})] \geq c$$

for every $y_{i_2} \in \{A,B\}$ and every $\mathbf{y}_{-i_2} \in \{A,B\}^{n-1}$. In the sandwich case ($0 \leq t_{i_2} \leq t_{i_1}$), given $y_{i_2}=A$ and a fixed $\mathbf{y}_{-i_2}$, we can enumerate $[\Pr(\mathbf{s}|a_{i_2}=1,a_{i_1}=1,\mathbf{a}_{-(i_1,i_2)}=\mathbf{0},\mathbf{y}) - \Pr(\mathbf{s}|a_{i_2}=0,a_{i_1}=1,\mathbf{a}_{-(i_1,i_2)}=\mathbf{0},\mathbf{y})]$ associated with all combinations of $s_0 s_{i_2} s_{i_1}$ as the followings (we henceforth omit $a_{i_1} = 1, \mathbf{a}_{-(i_1,i_2)} = \mathbf{0}$ to simplify notations):

$$\begin{aligned}
&\Pr(AAA|a_{i_2}=1, y_{i_2}=A, \mathbf{y}_{-i_2}) - \Pr(AAA|a_{i_2}=0, y_{i_2}=A, \mathbf{y}_{-i_2}) \\
&\quad = \underbrace{\frac{1}{2}}_{\Pr(s_0=A)} \times \underbrace{\frac{1}{2}(1+e^{-\lambda t_{i_2}})}_{\Pr(x(0)=x(t_{i_2}))} \times \underbrace{\frac{1}{2}(1+e^{-\lambda(t_{i_1}-t_{i_2})})}_{\Pr(x(t_{i_2})=x(t_{i_1}))} - \underbrace{\frac{1}{2}}_{\Pr(s_0=A)} \times \underbrace{1}_{\Pr(s_{i_2}=A|y_{i_2}=A)} \times \underbrace{\frac{1}{2}(1+e^{-\lambda t_{i_1}})}_{\Pr(x(0)=x(t_{i_1}))} = -\frac{1}{8}P_1^- P_2^- \\
&\Pr(ABA|a_{i_2}=1, y_{i_2}=A, \mathbf{y}_{-i_2}) - \Pr(ABA|a_{i_2}=0, y_{i_2}=A, \mathbf{y}_{-i_2}) \\
&\quad = \underbrace{\frac{1}{2}}_{\Pr(s_0=A)} \times \underbrace{\frac{1}{2}(1-e^{-\lambda t_{i_2}})}_{\Pr(x(0)\neq x(t_{i_2}))} \times \underbrace{\frac{1}{2}(1-e^{-\lambda(t_{i_1}-t_{i_2})})}_{\Pr(x(t_{i_2})\neq x(t_{i_1}))} - \underbrace{\frac{1}{2}}_{\Pr(s_0=A)} \times \underbrace{0}_{\Pr(s_{i_2}=B|y_{i_2}=A)} \times \underbrace{\frac{1}{2}(1+e^{-\lambda t_{i_1}})}_{\Pr(x(0)=x(t_{i_1}))} = \frac{1}{8}P_1^- P_2^- \\
&\Pr(ABB|a_{i_2}=1, y_{i_2}=A, \mathbf{y}_{-i_2}) - \Pr(ABB|a_{i_2}=0, y_{i_2}=A, \mathbf{y}_{-i_2}) = \frac{1}{8}P_1^- P_2^+ \\
&\Pr(AAB|a_{i_2}=1, y_{i_2}=A, \mathbf{y}_{-i_2}) - \Pr(AAB|a_{i_2}=0, y_{i_2}=A, \mathbf{y}_{-i_2}) = -\frac{1}{8}P_1^- P_2^+ \\
&\Pr(BBA|a_{i_2}=1, y_{i_2}=A, \mathbf{y}_{-i_2}) - \Pr(BBA|a_{i_2}=0, y_{i_2}=A, \mathbf{y}_{-i_2}) = \frac{1}{8}P_1^+ P_2^- \\
&\Pr(BAA|a_{i_2}=1, y_{i_2}=A, \mathbf{y}_{-i_2}) - \Pr(BAA|a_{i_2}=0, y_{i_2}=A, \mathbf{y}_{-i_2}) = -\frac{1}{8}P_1^+ P_2^- \\
&\Pr(BAB|a_{i_2}=1, y_{i_2}=A, \mathbf{y}_{-i_2}) - \Pr(BAB|a_{i_2}=0, y_{i_2}=A, \mathbf{y}_{-i_2}) = -\frac{1}{8}P_1^+ P_2^+ \\
&\Pr(BBB|a_{i_2}=1, y_{i_2}=A, \mathbf{y}_{-i_2}) - \Pr(BBB|a_{i_2}=0, y_{i_2}=A, \mathbf{y}_{-i_2}) = \frac{1}{8}P_1^+ P_2^+
\end{aligned}$$

where $P_1^+ = 1 + e^{-\lambda t_{i_2}}, P_1^- = 1 - e^{-\lambda t_{i_2}}, P_2^+ = 1 + e^{-\lambda(t_{i_1}-t_{i_2})}, P_2^- = 1 - e^{-\lambda(t_{i_1}-t_{i_2})}$. Plugging the above into the constraint with $y_{i_2} = A$, we get

$$\begin{aligned}
&\frac{1}{8}P_1^- P_2^-[w_{i_2}(ABA) - w_{i_2}(AAA)] + \frac{1}{8}P_1^- P_2^+[w_{i_2}(ABB) - w_{i_2}(AAB)] \\
&\qquad + \frac{1}{8}P_1^+ P_2^-[w_{i_2}(BBA) - w_{i_2}(BAA)] + \frac{1}{8}P_1^+ P_2^+[w_{i_2}(BBB) - w_{i_2}(BAB)] \geq c.
\end{aligned}$$

Similar calculations yield $\text{ICB}_{\mathbf{y}_{-i_2}}$.

Finally, we illustrate how to verify

$$\frac{\Pr(AAA|\mathbf{a}=\mathbf{1})}{\frac{1}{8}[(P_1^+P_2^+)^2 - (P_1^-P_2^-)^2]} \leq \min\left\{\frac{\Pr(AAB|\mathbf{a}=\mathbf{1})}{\frac{1}{8}[P_2^+P_2^-((P_1^+)^2 - (P_1^-)^2)]}, \frac{\Pr(BAA|\mathbf{a}=\mathbf{1})}{\frac{1}{8}[P_1^+P_1^-((P_2^+)^2 - (P_2^-)^2)]}\right\}$$

For the first term on the right-hand side, the inequality boils down to

$$\frac{\frac{1}{8}P_1^+P_2^+}{\frac{1}{2}(1+e^{-\lambda t_{i_1}})(e^{-\lambda(t_{i_1}-t_{i_2})} + e^{-\lambda t_{i_2}})} \leq \frac{\frac{1}{8}P_1^+P_2^-}{\frac{1}{2}P_2^+P_2^- e^{-\lambda t_{i_2}}},$$

This inequality simplifies to $2e^{-\lambda t_{i_1}} \leq e^{-\lambda(t_{i_1}-t_{i_2})} + e^{-\lambda(t_{i_1}+t_{i_2})}$, which is true as long as $t_{i_1} \geq t_{i_2}$, and the inequality is strict if $t_{i_1} > t_{i_2}$. This is because $e^{-\lambda t}$ is strictly convex. The other inequality is a mirror image of the first one and is obtained by similar calculations.

A.5 Proof of Proposition 3

Lemma 1 and Propositions 2 imply that the task allocation problem is

$$\min_{\tau} \frac{1}{2} - \frac{2n+1}{2\lambda} + \frac{1}{\lambda}\left[\sum_{i=1}^{n} e^{-\frac{\lambda \Delta t_i}{2}} + \frac{1}{2}e^{-\lambda(1-\sum_{i=1}^{n}\Delta t_i)}\right] + c\sum_{i=1}^{n}[1+e^{\lambda \Delta t_i}], \tag{14}$$

which is strictly convex and symmetric in the choice variables, $\Delta t_1, ..., \Delta t_n$. As a result, there is a unique $\Delta t \in [0, 1)$ such that $\Delta t_i = \Delta t$ for every i. Compared to the first best, which minimizes the information loss, the principal's objective has the incentive cost, $c\sum_{i=1}^{n}[1 + e^{\lambda \Delta t_i}] = cn[1 + e^{\lambda \Delta t}]$. This term is strictly increasing in Δt. The argument of monotone comparative statics implies that the optimal Δt, denoted by Δt^*, is strictly less than the first-best solution, $\Delta t^\dagger$.

To establish the comparative statics, we consider two cases. If $c < \bar{c} \triangleq \frac{1-e^{-\lambda}}{2\lambda}$, then the optimal Δt^* is strictly positive and uniquely solves the first-order condition $x^{2(n-1)} + 2\lambda e^{\lambda} c - \frac{e^{\lambda}}{x^3} = 0$, where $x \triangleq e^{\frac{\lambda \Delta t^*}{2}} \geq 1$. The left-hand side is strictly increasing in n and c, so to maintain the equation, the optimal Δt^* must decrease. If $c \geq \bar{c}$, then regardless of n, we have $\Delta t^* = 0$ at the optimum. Therefore, the optimal Δt^* is overall (weakly) decreasing in c and n. $\square$

If $c \geq \bar{c}$, the principal's loss is greater than when there are no agents, because the information generated by agents at $t = 0$ has no value to the principal, yet they still earn a rent of at least c.[12] Such a situation may arise because we assume that the principal must induce all agents to exert effort. In an extended model in which the principal could choose the number of agents to hire or the effort levels to induce, the principal would hire no agents or induce shirking by all agents whenever $c \geq \bar{c}$.

[12] To see this, suppose that all agents are at $t = 0$. If an agent faces (5), she is willing to work only if $2c + \epsilon - c \geq 0.5(2c + \epsilon)$, which implies $2c + \epsilon - c \geq c$. Thus, the agent's net payoff from exerting effort is at least c.

B Supplemental Appendix (Not for Publication)

B.1 Random Assignment and Transparency

In this section, we allow the principal to adopt a stochastic and private task allocation. Formally, we modify the model as follows: First, the principal publicly commits to a stochastic task allocation, which is a pair (T, π) such that $T \subseteq [0,1]^n$ is a finite set of deterministic task allocations and $\pi \in \Delta(T)$ is a probability distribution over T. The principal also commits to a compensation scheme, which specifies, for each realized task allocation $\tau \in T$ and a signal profile $\mathbf{s}$, the payment $w_i(\tau, \mathbf{s}) \geq 0$ for each agent i. Abusing notation, we continue using w for a compensation scheme of this model.

Once the principal chooses (T, π, w), Nature draws the state $x(\cdot)$ and the task allocation, $\tau = (t_1, ..., t_n) \sim \pi$. Each agent i observes (T, π, w) and t_i, and then decides whether to acquire information. Agents do not observe the realizations of other agents' locations. For example, if the principal randomizes over $(t_1, t_2) = (0.1, 0.3)$, $(0.1, 0.4)$, and $(0.2, 0.5)$ with equal probability, then agent 1 observes only whether $t_1 = 0.1$ or $t_1 = 0.2$, and after observing $t_1 = 0.1$, she infers that agent 2's location is 0.3 or 0.4 with equal probability.

Our notion of robust implementation extends to this setup. We apply the iterated elimination of strictly dominated strategies (hereafter, the IESDS) to the simultaneous-move incomplete information game $\Gamma(T, \pi, w)$, in which each agent privately observes her assigned location and then decides her action, and simultaneously, the noisemaker moves.

Even in this general setup, the minimal information loss is $L(\tau^\dagger)$ as in Proposition 1, and the minimal incentive cost is $K^\dagger = 2nc$, which, by Corollary 1, holds when all agents are at $t = 0$. The principal can achieve either of these outcomes (but not both) with a deterministic task allocation. However, the principal can now approximately attain the best possible total loss $L(\tau^\dagger) + K^\dagger$ by randomizing the task allocation and privately revealing each agent's realized location.

Proposition B.1. *Suppose that there are at least two agents. There is a sequence of stochastic task allocations and compensation schemes, $(T_k, \pi_k, w_k)_{k \in \mathbb{N}}$, such that for each k, compensation* *scheme w_k robustly implements effort, and as $k \to \infty$, the principal's total loss converges to $L(\tau^\dagger) + K^\dagger$.*

We first sketch the proof for the case of two agents and then prove the general case $n \geq 2$. Consider the following choice of the principal: With probability $1-p$, the principal chooses the first-best task allocation $\tau^\dagger = (t_1^\dagger, t_2^\dagger)$, which is $(\frac{2}{5}, \frac{4}{5})$ by Proposition 1. For the remaining event, the principal chooses $\tau_1^\dagger \triangleq (t_1^\dagger, t_1^\dagger)$ or $\tau_2^\dagger \triangleq (t_2^\dagger, t_2^\dagger)$ with probability $0.5p$. Task allocation $\tau_i^\dagger$ assigns both agents to agent i's first-best location, $t_i^\dagger$. The compensation scheme is as follows: If the realized task allocation is the first-best $\tau^\dagger$, there is no payment. Under task allocation $\tau_i^\dagger$, agent $j \neq i$ will receive $2ce^{\lambda t_i^\dagger} + \epsilon$ if her signal matches the principal's signal, $x(0)$, and agent i will receive $\frac{4c}{p} + \epsilon$ if her signal matches agent j's signal.

The key property of this policy is that even when the realized task allocation is the first-best, agent i still believes that the task allocation can be $\tau_i^\dagger$ and the other agent is assigned to the same location as agent i. While the probability of $\tau_i^\dagger$ can be small (as we will take $p \to 0$), the principal sets the bonus $\frac{4c}{p} + \epsilon$ for agent i large enough so that she finds it optimal to acquire information.

Indeed, for any $p \in (0,1)$ and $\epsilon > 0$, the IESDS uniquely selects both agents acquiring information: Suppose that agent 1 learns that her assigned location is agent 2's first-best location $t_2^\dagger$. Because the possible task allocations are $\tau^\dagger = (t_1^\dagger, t_2^\dagger)$, $\tau_1^\dagger = (t_1^\dagger, t_1^\dagger)$, and $\tau_2^\dagger = (t_2^\dagger, t_2^\dagger)$, agent 1 learns that the realized task allocation must be $\tau_2^\dagger$. As a result, agent 1 understands that she will receive $2ce^{\lambda t_2^\dagger} + \epsilon$ if her signal matches the principal's signal, $x(0)$, regardless of the other agent's action. By the same argument as the case of one agent (Section 4.2), we conclude that agent 1 strictly prefers to exert effort as a strictly dominant strategy. Symmetrically, agent 2 will work when the realized task allocation is $\tau_1^\dagger$.

Second, suppose that agent 1 learns that she is assigned to her first-best location $t_1^\dagger$. Agent 1 then infers that the task allocation is $\tau^\dagger$ or $\tau_1^\dagger$ with probability $\frac{1-p}{1-p+0.5p}$ or $\frac{0.5p}{1-p+0.5p}$, respectively. Agent 1 could possibly earn a positive payment only when the realization allocation $\tau_1^\dagger$, but agent 1 knows that, in such a case, agent 2 acquires information (because of the first step of the IESDS). Thus, by exerting effort, agent 1 can increase the probability of the positive payment

$\frac{4c}{p}+\epsilon$ by 0.5 in case the task allocation $\tau_1^\dagger$ is realized. As a result, the expected net gain for agent 1 of acquiring information after observing $t_1^\dagger$ is

$$\frac{0.5p}{1-p+0.5p}\cdot 0.5\cdot\left(\frac{4c}{p}+\epsilon\right)-c>0.5p\cdot 0.5\cdot\left(\frac{4c}{p}+\epsilon\right)-c=\epsilon>0.$$

Thus, agent 1 exerts effort as a strictly dominant strategy in the second step of the IESDS. By applying the symmetric argument to agent 2, we conclude that each agent 2 also chooses the desired strategy as a dominant strategy. Therefore, the IESDS selects the desired outcome for any realized task allocation.

Finally, by taking $(p,\epsilon)\to 0$, the principal will implement task allocation $\tau^\dagger$ with probability converging to 1, so the information loss converges to the first-best level. At the same time, the expected payment to each agent converges to $0.5p\cdot\frac{4c}{p}=2c$. Here, the payment regarding $2ce^{\lambda t_i^\dagger}+\epsilon$ will vanish as $p\to 0$, because it is independent of p and can arise only with probability $0.5p$. Therefore, the principal's total loss can be arbitrarily close to $L(\tau^\dagger)+K^\dagger$.

Proposition B.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 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.

Proof of Proposition B.1 For each $p\in(0,1)$, define task allocation policy (T^p,π^p) as follows: $T^p=\{\tau^\dagger,\tau_1^\dagger,...,\tau_n^\dagger\}$, and

$$\pi^p(\tau)\triangleq\begin{cases}1-p & \text{if } \tau=\tau^\dagger\\ \frac{p}{n} & \text{if } \tau=\tau_i^\dagger\triangleq\{t_1^\dagger,...,t_{n-1}^\dagger,t_i^\dagger\}, i=1,2,...,n-1\\ \frac{p}{n} & \text{if } \tau=\tau_n^\dagger\triangleq\{t_1^\dagger,...,t_{n-2}^\dagger,t_n^\dagger,t_n^\dagger\}.\end{cases}\tag{15}$$

This task allocation policy assigns all agents to their respective first-best locations, $t_1^\dagger,...,t_n^\dagger$, with probability $1-p$. With probability p/n each, the allocation policy chooses $\tau_i^\dagger$, which continues to assign agents $1,...,n-1$ to their first-best locations but now assigns agent n to $t_i^\dagger$

so that her signal is used to monitor agent i. With the remaining probability p/n, the policy draws $\tau_n^\dagger$, 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 signal can be used to monitor agent n. As a result, when agent $i \in \{1, ..., n\}$ learns that her own assignment is $t_i^\dagger$, she is uncertain whether the realized task allocation is $\tau^\dagger$ or $\tau_i^\dagger$, 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 signal will be compared to the principal's signal.

The following result presents a contract that, combined with the task allocation policy in (15), robustly implements the desired outcome.

Lemma B.1. *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 (15). For any $\epsilon > 0$, the following contract w^{p_k} robustly implements effort (RIE): For each agent $i = 1, ..., n-2$,*

$$w_i^{p,\epsilon}(\tau, \mathbf{s}) = \begin{cases} \frac{2nc}{p} + \epsilon & \textit{if } \tau = \tau_i^\dagger, s_i = s_n, \\ 0 & \textit{otherwise;} \end{cases}, \tag{16}$$

for agent $n-1$,

$$w_{n-1}^{p,\epsilon}(\tau, \mathbf{s}) = \begin{cases} \frac{2nc}{p} + \epsilon & \textit{if } \tau = \tau_{n-1}^\dagger, s_{n-1} = s_n, \\ 2ce^{\lambda(t_n^\dagger - t_{n-2}^\dagger)} + \epsilon & \textit{if } \tau = \tau_n^\dagger, s_{n-1} = s_{n-2} \\ 0 & \textit{otherwise;} \end{cases}, \tag{17}$$

and for agent n,

$$w_n^{p,\epsilon}(\tau, \mathbf{s}) = \begin{cases} \frac{2nc}{p} + \epsilon & \textit{if } \tau = \tau_n^\dagger, s_n = s_{n-1} \\ 2ce^{\lambda(t_i^\dagger - t_{i-1}^\dagger)} + \epsilon & \textit{if } \tau = \tau_i^\dagger, s_n = s_{i-1}, i < n \\ 0 & \textit{otherwise.} \end{cases} \cdot \tag{18}$$

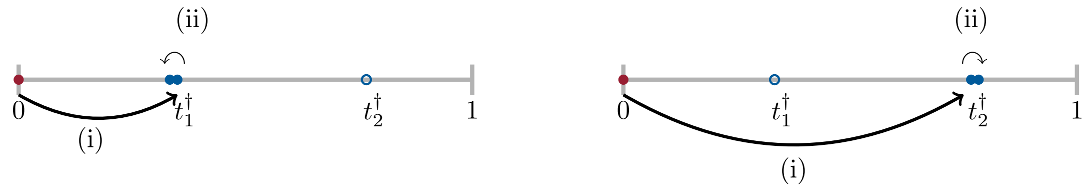

Figure 7: The left panel shows the chain monitoring structure when the realized task allocation is $\tau_1^\dagger$, where the principal's signal is used to monitor agent 2 and agent 2's signal is used to monitor agent 1. The right panel shows the chain monitoring structure when the realized task allocation is $\tau_2^\dagger$, where the principal's signal is used to monitor agent 1 and agent 1's signal 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.

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 signal matches her assigned location state, except for the events in which he needs to provide strict incentive to the monitor.*

Lemma B.1 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., $\tau_i^\dagger$ is realized) and (ii) her signal 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, 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 signal is aligned with s_{i-1}. Finally, when the realized task allocation is $\tau_n^\dagger$, agent $n-1$ is assigned to location $t_n^\dagger$ secretly with probability p/n to verify whether agent n's signal is aligned with the state. We plot a two-agent case in Figure 7; 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 compensation scheme RIE. Expression (18) implies that, if the realized allocation is $\tau_1^\dagger$, 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 $\tau_1^\dagger$

and agent n will generate $s_n = x(t_1^\dagger)$. By expression (16), 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 exert effort in each realized allocation.

Proof of Lemma B.1. For agent i with $i = 1, 2, ..., n-2$, 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 (16). In this construction, agent 1's dominance incentive is established in the *second* round of IESDS[13], after agent n with $\tau_1^\dagger$. Moreover, each agent i's dominance incentive is established in the $2i$-th round, immediately after agent n with $\tau_i^\dagger$, for all $i = 2, ..., n-2$. 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 $\tau_n^\dagger$ 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 $s_{n-1} = s_{n-2}$ despite being located at t_n. The worker-type agent $n-1$ has the same incentive cost as agent $i = 1, 2, ..., n-2$. Combining both types, agent $n-1$'s wage scheme satisfies (17). 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 signal 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 signal 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 $i = 1, 2, ..., n-2$. Combining all types, agent

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

n's wage scheme satisfies (18). Thus, the total expected incentive cost is

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

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

We now prove Proposition B.1. For an arbitrary constant $\delta \in (0, 1)$, define $p_k = \delta^k$, $k \in \mathbb{N}_+$. We obtain a sequence of probabilities $\{p_k\}_{k=1}^{\infty}$. Define (T^{p_k}, π^{p_k}) by (15). Applying Lemma B.1, for each k, we can find a compensation scheme w^{p_k} that RIE. The resulting total loss is

$$(1-p_k)L(\tau^\dagger) + \frac{n-1}{n} p_k [L(\tau^\dagger) + const] + \frac{1}{n} p_k [L(\tau^\dagger) + const] + K^\dagger + \frac{p_k}{n}(2ce^{2\lambda t_1^\dagger} + 2(n-1)ce^{\lambda t_1^\dagger} - 2nc).$$

As $k \to \infty$, we obtain the limit in Proposition B.1, i.e., $L(\tau^\dagger) + K^\dagger$. □

B.2 Imperfect Information Acquisition

We assume that each agent i who chooses to acquire information will generate signal s_i such that $\Pr(s_i = A | x(t_i) = A) = \Pr(s_i = B | x(t_i) = B) = q \in (0.5, 1)$. We also assume that the principal's signal about $x(0)$ follows the same signal structure. Conditional on the state, all the signals are independent.

Given any two signals $s(t'), s(t)$ from two agents located at t' and t ($t' > t$), conditional on both agents' exerting effort, the critical probability that governs agents' incentive is

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

The analysis of general n is challenging, but we managed to illustrate our point in the special case with $n = 2$, which we assume henceforth throughout this section.

On Sufficient Statistics. Note that the signal 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 $s_1 = s_2$, 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

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

where the additional term, relative to the benchmark incentive cost, comes from the imperfect observation of agents. Now we show that utilizing an additional signal s_0 reduces the incentive cost. Suppose agent 2 is paid if and only if $s_0 = s_1 = s_2$. Again, assume $a_1 = 1$. Then

$$\Pr(s_0 = s_1 = s_2 | a_2 = 1) = \left[\frac{1}{4} - q(1-q)\right](1 + e^{-\lambda t_1})(1 + e^{-\lambda(t_2-t_1)}) + q(1-q).$$

Using binding IC of agent 2, we get

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

The incentive cost is

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

which is less than the incentive cost based on (s_1, s_2). When $q \to 1$, 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 signal. This signal satisfies $s_0 = s_1 = s_2$, providing the maximal incentive to work for agent 2. This is different from our main analysis, because the nearest strategically risk-free neighbor's signal is no longer a sufficient statistic.

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

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

The optimal policy still respects each signal's illumination range. However, given each signal s_i, only with probability q this signal 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 $s_i \neq x_i$, 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 assigns each agent to the same length of illumination range.

Proposition B.2. *Fixing imperfect learning parameter $q \in (0.5, 1)$, the information-loss-minimizing task allocation is $\tau^\dagger$ with*

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

Optimal Task Allocation When $n = 2$. Next, we turn to the task allocation minimizing both incentive cost and information loss. 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 $s_0 = s_1 = s_2$. 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

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

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

$$\left(e^{-\lambda(1-\Delta t_1-\Delta t_2)} - e^{-\frac{\lambda\Delta t_1}{2}}\right)\frac{2q-1}{2} + \lambda c e^{\lambda\Delta t_1} + \frac{2q(1-q)c\lambda e^{\lambda\Delta t_1}}{\frac{1}{2}-2q(1-q)} + \frac{2q(1-q)c\lambda e^{\lambda(\Delta t_2-\Delta t_1)}}{[\frac{1}{2}-2q(1-q)](1+e^{-\lambda\Delta t_1})^2} = 0 \tag{19}$$

$$\left(e^{-\lambda(1-\Delta t_1-\Delta t_2)} - e^{-\frac{\lambda\Delta t_2}{2}}\right)\frac{2q-1}{2} + \lambda c e^{\lambda\Delta t_2} + \frac{2q(1-q)c\lambda e^{\lambda\Delta t_2}}{[\frac{1}{2}-2q(1-q)](1+e^{-\lambda\Delta t_1})} = 0 \tag{20}$$

Conditions (20) – (19) require

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

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 (20) − (19) < 0, a contradiction. Therefore, the interior solution τ must satisfy $t_2 - t_1 > t_1$. The results are summarized in the following two propositions.

Proposition B.3. *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 effort 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 signal matches all the signals to the left, i.e.,

$$w_1 = \begin{cases} \frac{ce^{\lambda t_1}}{\frac{1}{2} - 2q(1-q)}, & \text{if } s_1 = s_0 \\ 0, & \text{otherwise.} \end{cases}$$

$$w_2 = \begin{cases} \frac{ce^{\lambda(t_2 - t_1)}}{(\frac{1}{4} - q(1-q))(1 + e^{-\lambda t_1})}, & \text{if } s_2 = s_1 = s_0 \\ 0, & \text{otherwise.} \end{cases}$$

Proposition B.4. *Fixing $n = 2$ and the imperfect learning parameter $q \in (0.5, 1)$, the optimal (interior) task allocation exhibits $t_2 - t_1 > t_1$.*

B.3 Asymmetric Effort Cost

The optimal monitoring structure will change when agents have different costs of effort. 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 $(\tau, \mathbf{c})$ to minimize his total loss, i.e.,

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

For each $(\tau, \mathbf{c})$, information loss $L(\tau)$ is still given by equation (3), but the incentive cost is

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

where

$$K^c(\tau, \mathbf{c}) = \sum_{i=1,2} c_i\big[1 + e^{\lambda(t_i - t_{i-1})}\big] \tag{21}$$

$$K^s(\tau, \mathbf{c}) = c_1[1 + e^{\lambda t_2}] + c_2\Big[1 + \frac{1 + e^{-\lambda t_2}}{e^{-\lambda t_1} + e^{-\lambda(t_2 - t_1)}}\Big] \tag{22}$$

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 B.2. *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 B.2. 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

$$e^{-\lambda(1-t_2)} - e^{\frac{-\lambda t_1}{2}} + 2\lambda c_1 e^{\lambda t_1} = 0$$

$$e^{-\lambda(1-t_2)} - e^{\frac{-\lambda(t_2-t_1)}{2}} + 2\lambda c_2 e^{\lambda(t_2-t_1)} = 0$$

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 B.2 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 effort 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 B.5. *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., $\exists \tau^s \in \arg\min_{\tau} L(\tau) + K^s(\tau, c^*)$ *and* $t_2^s > t_1^s > 0$.

Proof of Proposition B.5. The first part of Proposition B.5 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,

$$\frac{\partial \mathcal{L}}{\partial \Delta c} = \underbrace{\frac{\partial L(\tau^*)}{\partial \Delta c}}_{=0} + \frac{\partial K(\tau^*)}{\partial \Delta c} \tag{23}$$

First, fixing the chain monitoring structure, (23) becomes

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

by Lemma B.2. Hence, the loss decreases in Δc.

On the other hand, fixing the sandwich monitoring structure. Then (23) 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 B.2. 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 B.5, 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

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

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 $\lambda \to 0$, $\widehat{\Delta c(\lambda)} \to 0$. $\square$

Proposition B.5 says that the principal benefits from increasing agents' cost disparity. The logic is simple. Fix the task allocation $(\tau^*, \mathbf{c}^*)$, consider a small increase in agents' cost disparity by $\epsilon > 0$. The corresponding impact on the principal's total loss can be approximated by

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

The inequality is a consequence of Lemma B.2. Intuitively, each agent's optimal expected payment is her effort 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.

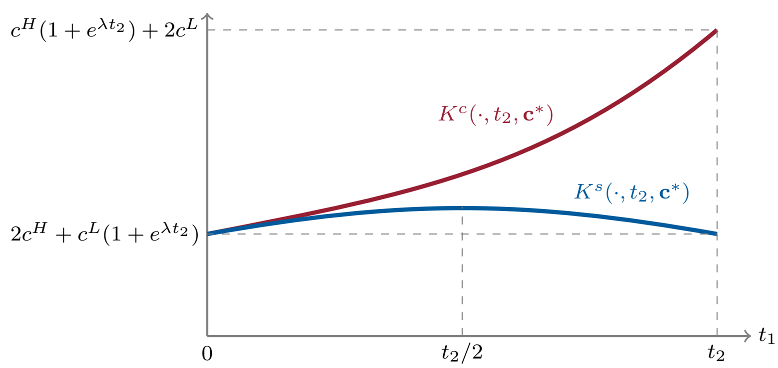

Figure 8: Illustrating expected payment function $K^s(\cdot, t_2, \mathbf{c}^*)$ and $K^c(\cdot, t_2, \mathbf{c}^*)$.

The second part of Proposition B.5 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 8, 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 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 (22).

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 slop 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.

B.4 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 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

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

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 9. 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 $t_2^{**} > 1/2 + t^*/2$ 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 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 signal is aligned with the

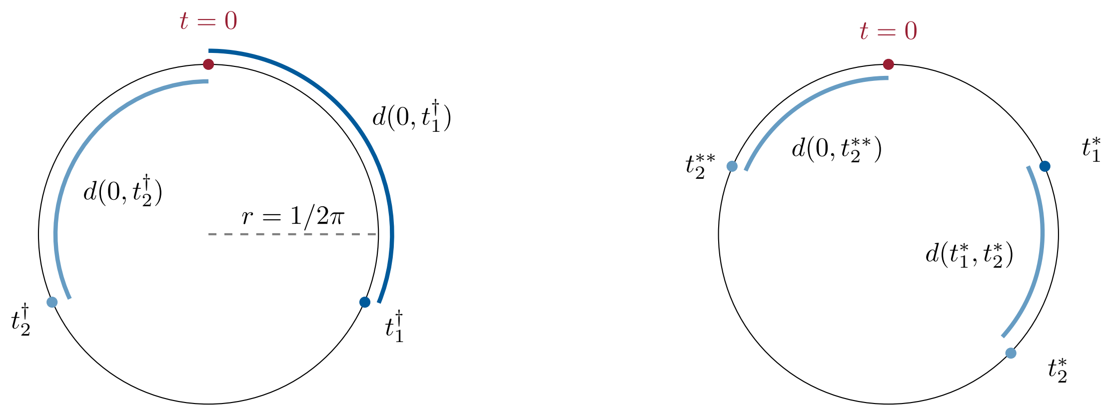

Figure 9: The location circle and task allocation.

principal's signal. 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 $s_0 = s_1$ 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 signal, and agent 2 is paid by w_2 if and only if $s_0 = s_1 = s_2$, 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 9, 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.[14]

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 $x(0) = x(t_1) = x(1)$. 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 $s_1 = s_2 = x(0)$. This also specifies a sandwich structure. See Figure 10 for an illustration. Notice that whenever it

[14]It 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 10: Illustration of the multiplicity.

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.