Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Efficient Resource Allocation Contracts to Reduce Adverse Events

Yong Liang, Peng Sun, Runyu Tang, Chong Zhang

informs
http://pubsonline.informs.org/journal/opre
OPERATIONS RESEARCH
Articles in Advance, pp. 1–19
ISSN 0030-364X (print), ISSN 1526-5463 (online)

Methods

Efficient Resource Allocation Contracts to Reduce Adverse Events

Yong Liang,[a] Peng Sun,[b] Runyu Tang,[c,*] Chong Zhang[d]

[a] Research Center for Contemporary Management, Key Research Institute of Humanities and Social Sciences at Universities, School of Economics and Management, Tsinghua University, Beijing 10084, China; [b] Fuqua School of Business, Duke University, Durham, North Carolina 27708; [c] School of Management, Xi'an Jiaotong University, Xi'an, Shaanxi 710049, China; [d] Department of Management, Tilburg University, 5037 AB Tilburg, Netherlands
*Corresponding author

Contact: liangyong@sem.tsinghua.edu.cn, https://orcid.org/0000-0002-7052-2248 (YL); peng.sun@duke.edu, https://orcid.org/0000-0002-3892-5936 (PS); tangry@xjtu.edu.cn, https://orcid.org/0000-0003-4353-665X (RT); c.zhang@tilburguniversity.edu, https://orcid.org/0000-0002-0848-5677 (CZ)

Received: September 1, 2020
Revised: May 30, 2021; December 16, 2021
Accepted: May 2, 2022
Published Online in Articles in Advance: July 1, 2022

Area of Review: Decision Analysis

https://doi.org/10.1287/opre.2022.2322

Abstract. Motivated by the allocation of online visits to product, service, and content suppliers in the platform economy, we consider a dynamic contract design problem in which a principal constantly determines the allocation of a resource (online visits) to multiple agents. Although agents are capable of running the business, they introduce adverse events, the frequency of which depends on each agent's effort level. We study continuous-time dynamic contracts that utilize resource allocation and monetary transfers to induce agents to exert effort and reduce the arrival rate of adverse events. In contrast to the single-agent case, in which efficiency is not achievable, we show that efficient and incentive-compatible contracts, which allocate all resources and induce agents to exert constant effort, generally exist with two or more agents. We devise an iterative algorithm that characterizes and calculates such contracts, and we specify the profit-maximizing contract for the principal. Furthermore, we provide efficient and incentive-compatible dynamic contracts that can be expressed in closed form and are therefore easy to understand and implement in practice.

Funding: Y. Liang acknowledges support from the National Key R&D Program of China [Grant 2020AAA0103801] and the National Natural Science Foundation of China [Grant 71872095].
Supplemental Material: The online appendix is available at https://doi.org/10.1287/opre.2022.2322.

Keywords: dynamic contract design • moral hazard • self-generating set • platform economy • stochastic optimal control

1. Introduction

Some of the most valuable companies by market capitalization, such as Amazon, Apple, and Alibaba, and many $1 billion unicorn start-ups, such as Airbnb, Bytedance, DiDi, and Uber, build online platforms to satisfy demand with supply from individual players. These digital platforms facilitate and reshape a wide range of businesses and social activities, from selling products (e.g., Amazon, Apple, and Alibaba) and providing services (e.g., DiDi, Uber, and Airbnb) to hosting news and information content (e.g., Bytedance). The success of the platform economy depends crucially on moral suppliers providing high-quality services, products, or information content. Adverse events related to low-quality supply, however, hurt a platform's reputation and society at large. Recent examples of adverse events on these platforms include brushing and fake reviews from third-party sellers on Amazon in 2020 (Dai and Tang 2020), fraudulent drivers on DiDi in 2018 (Feng 2019), inappropriate content on Reddit in 2016 (Marantz 2018), and tampering with food at Uber Eats in 2018 (Edelstein 2018). In many situations, product, service, and content suppliers on these platforms could exert effort to reduce the chance of adverse events. However, effort may be hard to verify, and adverse events may still occur despite the best effort.

In practice, some platforms create *resource allocation* incentives that adjust online visits allocated to different suppliers based on their performance. For example, Alibaba devised a point-based system on its retail platform. Each merchant is initially credited with a fixed amount of "points" that are deducted once the merchant breaches service quality promises or sells inferior products. Low-point merchants are penalized with low search visibility and limitations on marketing activities, or even termination (Alibaba Group 2015). Other platforms implemented similar point-based systems, as summarized in Table 1. Nonetheless, some platforms still struggle with frequent incidents (Mauldin 2019). From an operations perspective, the proper design of a resource allocation system to provide the right incentives for suppliers remains a challenge.

Table 1. Summary of Regulatory Systems for Well-Known Platforms

Platform	Category	Adverse event	Possible ramification
eBay Taobao	Online retailing	• Counterfeit goods • Breach of promises • Fake transactions	• Demotion of search weight • Limitations on promotion • Termination
DiDi Uber	Ride hailing	• Driving under influences • Safety issues • Fraud and theft	• Reduction in dispatch • Suspension • Termination
Ele.me Uber Eats	Food delivery	• Spoiled food • Delivery issues • Tampering with food	• Demotion of search weight • Reduction of subsidy • Suspension • Termination
Airbnb	Lodging rentals	• Inaccurate listings • Unclean spaces • Cancellations of confirmed reservations	• Loss of ("Superhost") visibility • Deactivation of listing
ByteDance Twitter YouTube	User-generated content	• Low-quality content • Fake news	• Restrictions on posting • Limitation on promotion

In this paper, we study dynamic contracts that motivate agents to exert effort to reduce the frequency of adverse events in a continuous-time setting over an infinite time horizon. Specifically, a risk-neutral principal (platform) with commitment power to design long-term contracts owns a fixed flow of resource (online visits) to be allocated among multiple agents (product, service, or content providers). Agents can use the resource to generate revenue, but they can also generate adverse events that are costly to the principal. Exerting effort allows agents to reduce the arrival rate of these adverse events. However, effort is costly to the agents and not observable to the principal. Following the usual limited liability assumption, the principal cannot align incentives by making agents bear the cost of adverse events. Therefore, the principal must rely on dynamic resource allocation and payment decisions that depend on past arrival times to induce effort from agents. Because most platforms have to direct all online visits to product, service, and content providers, we focus on *efficient* contracts that allocate all the resource to the agents. We also focus on *incentive-compatible* contracts that induce all agents to exert effort all the time. We investigate the existence of efficient and incentive-compatible (EIC) contracts and, if they exist, identify the one that maximizes the principal's utility. It is worth noting that an EIC contract may not be profit-maximizing. That is, a dynamic contract may yield a higher profit if it does not always allocate all resource to all agents, or does not always induce effort from all agents. However, such a more general class of contracts appears quite challenging to study and is left for future research.

From a theoretical perspective, our paper contributes to the continuous-time dynamic contracting literature with Poisson arrivals of adverse events by considering a multiagent setting. From a practical perspective, our results provide prescriptive guidance for practitioners to design easy-to-implement EIC contracts. Specifically, our main contributions are threefold. First, we extend the continuous-time dynamic moral hazard models of Biais et al. (2010) and Myerson (2015) to a multiagent setting, where a principal could leverage its resource allocation, besides payments, to incentivize agents to exert effort. In this context, we establish the existence of EIC contracts, despite the fact that they do not exist in the corresponding single-agent setting, as illustrated in Section 3. In particular, we describe the set of EIC contracts by characterizing the *self-generating set* (Abreu et al. 1990, Bernard and Frei 2016) of agents' total future utilities (also called *promised utilities*) that can be achieved using these contracts. Following an idea from Balseiro et al. (2019) of using support functions to represent the convex set of achievable promised utilities, we propose an iterative algorithm based on solving a sequence of (infinite-dimensional) linear optimization models to characterize the self-generating set. Intuitively, under the single-agent setting, due to the agent's limited liability, the principal has to withhold some resource from time to time in order to induce effort, and thus EIC contracts do not exist. In contrast, with multiple agents, our results suggest that the principal can keep allocating all the resource to agents and induce effort by redistributing the resource.

Second, we show that the solutions to the sequence of linear optimization models yield the optimal EIC contract that maximizes the principal's utility. The contract involves continuously adjusting all agents' allocations of resource and promised utilities between arrivals, and letting them take discrete jumps upon adverse arrivals. In particular, whenever an agent experiences an adverse event, this agent's allocation and promised utility take downward jumps, while other agents' allocations and promised utilities also take discrete jumps. Our analyses reveal that allowing such jumps in other agents' allocations and promised utilities is essential to the existence of EIC contracts.

Third, we design easy-to-implement dynamic EIC contracts in closed form that no longer require solving linear optimization problems. Although not necessarily optimal, these contracts are very easy to compute and implement. According to these contracts, each agent's income rate is proportional to the allocated resource. This feature is particularly relevant to online platform applications, where agents are individual players on the platform whose income is often proportional to the amount of online visits that they receive. Therefore, our easy-to-implement contracts allow online platforms to motivate players to maintain quality by adjusting online visit allocations in a straightforward manner.

Now we review the related literature. The study of moral hazard problems has gained growing attention since the early works of Holmström (1979) and Grossman and Hart (1983) on contract theory. In static settings, many studies examine bilateral contracting problems (see, e.g., Innes 1990, Baker 1992, Prendergast 2002). Early studies of dynamic moral hazard problems often consider discrete-time models (see, e.g., Rogerson 1985, Spear and Srivastava 1987, Gibbons and Murphy 1992, Holmström 1999). The main stochasticity in our setting is interarrival times or, equivalently, the frequency of adverse events. Therefore, we consider continuous-time models. There is a stream of literature on continuous-time moral hazard problems since Sannikov (2008), who uses a Brownian motion to capture uncertain outcomes following an agent's effort process. The methodology in Sannikov (2008) provides the analytical foundation for this stream of research.

Biais et al. (2010) extend the continuous-time optimal contracting framework to study firm dynamics based on a Poisson process of adverse events instead of a Brownian motion uncertainty process. Besides direct payments, an investor (the principal) can also change the size of a firm (the agent). In particular, they assume an upper bound to the speed at which the firm can scale up in size, while there is no constraint on downsizing. Our modeling framework is closely related to that of Biais et al. (2010), with some important differences. The most obvious one is that we consider contracting with multiple agents, while Biais et al. (2010) focus on the single-agent case. With a single agent, no incentive-compatible contract achieves efficiency in a setting with adverse events. With at least two agents, conversely, efficient contracts do exist, which is the focus of our study. A more nuanced difference is that while it is quite reasonable to assume that the speed at which a firm can scale up is upper-bounded, as in Biais et al. (2010), our principal can change the allocation at any time instantaneously, with no speed limit. For example, Airbnb is able to decrease one host's visibility while increasing another host's visibility by immediately redirecting lodging requests away from the former to the latter. This is in contrast to traditional manufacturing settings, in which ramping up the production rate takes time. Therefore, our modeling assumption is justified in the context of platforms allocating online visits. Furthermore, we assume a limited total amount of resource to be allocated, while the firm's size in Biais et al. (2010) can grow without bounds. These differences in our modeling assumption lead to different analyses and results from Biais et al. (2010).

From a modeling perspective, Myerson (2015) is closely related to Biais et al. (2010) and therefore also to our paper. Myerson (2015) studies a problem in political economy in which the principal can dynamically pay and/or replace an agent in order to motivate effort. Instead of replacing the agent, Chen et al. (2020) uses costly monitoring to resolve information asymmetry and dynamically schedules monitoring and payments to ensure an agent's effort. Both of these studies examine single-agent settings in which effort reduces the arrival rate of a Poisson process. Despite apparent connections, the specific dynamics of replacement (Myerson 2015), monitoring (Chen et al. 2020), and resource allocation (our paper) are different.

Analytically, our characterization of EIC contracts relies on the set of agents' promised utilities achievable by the contracts. Spear and Srivastava (1987) first proposed using promised utility as a state variable to study infinite-horizon contracts. Abreu et al. (1990) propose the concept of a *self-generating set* of promised utilities in a study of infinite-horizon repeated games among multiple agents with imperfect monitoring. The concept of self-generation is crucial for Fudenberg et al. (1994) to establish a folk theorem, which further implies that efficient mechanisms and contracts exist in dynamic adverse selection and moral hazard problems, respectively, when agents are infinitely patient (the discount factor approaches one). However, it is hard to deduce the corresponding dynamic mechanisms or contracts for operational purposes from their *existence* proofs. Balseiro et al. (2019) consider dynamic mechanism design without money and propose a mechanism that approaches efficiency as the discount factor approaches one. When there are only two agents, their mechanism's convergence rate to efficiency is optimal. Note that Balseiro et al. (2019) study an adverse selection problem in a discrete-time setting, whereas we consider a moral hazard problem in a continuous-time setting. Furthermore, their results rely on the time discount factor approaching one. In contrast, the existence of EIC contracts in our setting does not rely on agents being infinitely patient. An important analytical approach developed in Balseiro et al. (2019), which we adopt in our study, is to characterize the set of achievable promised utilities using its support functions. This approach allows us to compute the optimal EIC contract through iteratively solving a sequence of linear programs. Bernard and Frei (2016) extend the self-generating set concept from discrete-time settings to a continuous-time setting and establish the folk theorem for a general setup with Brownian motion uncertainties. Our linear program-based recursive algorithm may be extended to practically compute the support function of the self-generating set. Support functions are also used in static mechanism design (see, e.g., Goeree and Kushnir 2016, 2020).

Poisson arrivals need not be adverse events. In other settings, a principal may want to motivate agents' effort to increase the arrival rate of "good" arrivals. Shan (2017), for example, considers a principal hiring two agents to carry out a multistage project, whose successful outcomes follow a Poisson process with a rate jointly determined by the effort choices of both agents. The study considers free riding issues when the total effort levels determine the arrival rate, an aspect that we do not consider. Payment is the main contractual lever in that study, rather than resource allocation, which is another distinction with our study. Other recent works on good arrivals include Green and Taylor (2016) and Sun and Tian (2018), which study optimal contract design for single-agent settings and are more tangential to our study.

Recent operations management studies also examine incentive management issues using principal-agent models. For example, to mitigate product adulteration, Babich and Tang (2012) compare a deferred payment mechanism to a mechanism that combines deferred payment with inspection. Deferring payment until the principal is more certain about product quality provides a threat that induces an agent's good behavior. This line of work also includes the study of double-sided moral hazard problems in crowdfunding (Babich et al. 2021), irresponsible labor and environmental practices (Plambeck and Taylor 2016, Wang et al. 2016, Chen and Lee 2017, Cho et al. 2019), product adulteration (Rui and Lai 2015), sourcing and financing under suppliers' performance risks (Tang et al. 2018), and violations in health and safety regulations (Caro et al. 2018). Most of these papers consider static models and focus on classes of practically useful incentive mechanisms. Our study provides a dynamic modeling and analytical framework that allows the principal to leverage both dynamic resource allocation and payment to address incentive issues in some of these settings.

The rest of this paper is organized as follows. We first describe the model and introduce EIC contracts in Section 2. Section 3 introduces the self-generating set of achievable utilities and an iterative approach to construct the set of EIC contracts. Section 4 describes the optimal EIC contract that maximizes the principal's utility. Section 5 proposes two easy-to-implement dynamic EIC contracts with desirable properties that are relevant in practice. Section 6 further discusses a sufficient condition for the optimality of incentive-compatible contracts and extends the model to consider asymmetric agents. Finally, Section 7 concludes the paper and discusses further insights as well as future research directions. All the proofs are presented in the online appendix.

2. The Model

We consider a principal-agent model in a continuous-time setting. The principal has a resource, normalized to one per unit of time, to run a platform business, where the resource is the total visits to the platform. To run the business, the principal has to allocate the resource among multiple symmetric agents. Each agent generates some nonnegative revenue for the principal, and some for itself. In particular, the revenue rate for the principal is R^p per unit of resource and time, and the rate for the agent itself is R^a. We define $R := R^p + R^a$ as the total societal revenue rate. Both the principal and the agents are risk-neutral and discount future cash flows at the same rate of $\rho > 0$. We use $\mathcal{I} = \{1, \ldots, n\}$ to represent the set of agents. At each time epoch t, let $X_{i,t}$ denote the resource allocated to agent i. We assume that the allocation decisions satisfy the following condition, in which "FX" represents "feasible X":

$$X_{i,t} \geq 0, \text{ and } \sum_{i=1}^{n} X_{i,t} = 1, \ \forall t \geq 0, \ \forall i \in \mathcal{I}. \tag{FX}$$

That is, as mentioned in the introduction, we restrict contracts under consideration to allocate all the resource to agents all the time.

An agent's work may generate adverse events following a Poisson process. An adverse event from agent i arriving at time t causes a cost $C \cdot X_{i,t}$ to the principal, which is scaled with the resource allocated to the agent.[1] Each agent can reduce the arrival rate of the adverse events from $\bar{\lambda}$ to $\lambda = \bar{\lambda} - \Delta\lambda$ by exerting effort. Let $\boldsymbol{\Lambda} = \{\boldsymbol{\lambda}_t\}_{t\geq 0} = \{(\lambda_{i,t})_{i\in\mathcal{I}}\}_{t\geq 0}$ denote the agents' effort processes, where $\lambda_{i,t} \in \{\lambda, \bar{\lambda}\}$. Shirking brings an agent a benefit rate of $b \cdot X_{i,t}$, which is also proportional to the resource allocated to the agent. The principal observes only the history of adverse events but not the effort processes and, therefore, faces a dynamic moral hazard problem.

We assume that the principal has commitment power to issue a long-term contract, which is a contingency plan that both the principal and the agents are willing to follow through. It specifies both the payment and the resource allocation policies over time. We use an n-dimensional counting process $\{\boldsymbol{N}_t\}_{t\geq 0} = \{(N_{i,t})_{i\in\mathcal{I}}\}_{t\geq 0}$ to represent the number of adverse events induced by each agent up to time t. Define a filtration $\mathcal{F}^N = \{\mathcal{F}_t^N\}_{t\geq 0}$ generated by the counting process $\{\boldsymbol{N}_t\}_{t\geq 0}$, such that $\mathcal{F}_t^N$ captures the entire information history up to time t specified by the counting process $\{\boldsymbol{N}_t\}_{t\geq 0}$. The contract shall depend on the history of adverse events; that is, a contract consists of $\mathcal{F}^N$-predictable payment and resource allocation processes.

We follow the standard assumption that the agents have limited liability and are cash-constrained. That is, the principal can pay the agents but cannot be paid by them at any time.[2] Next, we use an n-dimensional $\mathcal{F}^N$-predictable process $\{\boldsymbol{L}_t\}_{t\geq 0} = \{(L_{i,t})_{i\in\mathcal{I}}\}_{t\geq 0}$ to represent the principal's cumulative payments to each agent $i \in \mathcal{I}$ up to time t. The limited liability and cash-constrained condition is

$$\mathrm{d}L_{i,t} \geq 0, \quad \forall t \geq 0, \ \forall i \in \mathcal{I}. \tag{LL}$$

Moreover, we consider $\mathrm{d}L_{i,t} = I_{i,t} + l_{i,t}\,\mathrm{d}t$, where $I_{i,t}$ and $l_{i,t}$ represent instantaneous payment and flow payment to agent i at time t, respectively. Following the same conventions as for $\boldsymbol{\lambda}_t$, we use n-dimensional nonnegative vectors $\boldsymbol{I}_t$ and $\boldsymbol{l}_t$ to represent the vectors of the instantaneous and flow payments, respectively.

Besides using payments, the principal can also dynamically allocate the resource among the agents to induce effort. Let $\{\boldsymbol{X}_t\}_{t\geq 0} = \{(X_{i,t})_{i\in\mathcal{I}}\}_{t\geq 0}$ be the resource allocation rule, where $\boldsymbol{X}_t$ represents the resource allocation at time t. Following (FX), vector $\boldsymbol{X}_t$ lies in the n-dimensional simplex. Formally, a contract $\Gamma = (\{\boldsymbol{L}_t\}_{t\geq 0}, \{\boldsymbol{X}_t\}_{t\geq 0})$ consists of $\mathcal{F}_t^N$-predictable payment and allocation processes, $\{\boldsymbol{L}_t\}_{t\geq 0}$ and $\{\boldsymbol{X}_t\}_{t\geq 0}$, respectively.

2.1. Agents' Utilities

Agents' utilities consist of the discounted total payments and potential benefits from shirking. Given contract Γ and agents' effort process $\boldsymbol{\Lambda} = \{\boldsymbol{\lambda}_t\}_{t\geq 0} = \{(\lambda_{i,t})_{i\in\mathcal{I}}\}_{t\geq 0}$, the *total expected utility* of agent i, denoted by $u_i(\Gamma, \boldsymbol{\Lambda})$, is

$$u_i(\Gamma, \boldsymbol{\Lambda}) := \mathbb{E}^{\boldsymbol{\Lambda}}\left[\int_0^\infty e^{-\rho t}\left[\mathrm{d}L_{i,t} + \left(R^a + b\,\mathbb{1}_{\{\lambda_{i,t}=\bar{\lambda}\}}\right)X_{i,t}\,\mathrm{d}t\right]\right], \quad \forall i \in \mathcal{I}, \tag{2.1}$$

where $\mathbb{E}^{\boldsymbol{\Lambda}}$ represents the expectation taken with respect to the probability measure induced by the arrival rate process $\boldsymbol{\Lambda}$ of Poisson processes and $\mathbb{1}_{\{\cdot\}}$ represents the indicator function.

It is standard and convenient to work with agents' *promised utilities* (see, e.g., Biais et al. 2010). The promised utility of agent i at time t is

$$W_{i,t}(\Gamma, \boldsymbol{\Lambda}) := \mathbb{E}^{\boldsymbol{\Lambda}}\left[\int_t^\infty e^{-\rho(s-t)}\left[\,\mathrm{d}L_{i,s} + \left(R^a + b\,\mathbb{1}_{\{\lambda_{i,t}=\bar{\lambda}\}}\right)X_{i,t}\mathrm{d}s\right] \,\middle|\, \mathcal{F}_t^N\right], \quad \forall t \geq 0, \ \forall i \in \mathcal{I}. \tag{2.2}$$

The promised utility $W_{i,t}(\Gamma, \boldsymbol{\Lambda})$ is a right-continuous process capturing the total discounted utility of agent i starting from time t. For convenience, we omit "$(\Gamma, \boldsymbol{\Lambda})$" when the context is clear and refer to $W_{i,t}(\Gamma, \boldsymbol{\Lambda})$ as $W_{i,t}$. Agents are free to walk away from the contract. Therefore, we require the following *individual rationality* (IR) constraint to ensure agents' participation:

$$W_{i,t} \geq 0, \quad \forall t \geq 0, \ \forall i \in \mathcal{I}. \tag{IR}$$

In the next lemma, we characterize the dynamics of agents' promised utilities $\boldsymbol{W}_t = (W_{i,t})_{i\in\mathcal{I}}$ under an arbitrary contract in terms of the stochastic integral with respect to the right-continuous $\{W_{i,t}\}_{t\geq 0}$ process. For this purpose, we define a left-continuous process $W_{i,t-} = \lim_{s\uparrow t} W_{i,s}$. We extend the definition of $W_{i,t}(\Gamma, \boldsymbol{\Lambda})$ such that $W_{i,0-}(\Gamma, \boldsymbol{\Lambda}) = u_i(\Gamma, \boldsymbol{\Lambda})$. Clearly, we have $W_{i,0}(\Gamma, \boldsymbol{\Lambda}) = W_{i,0-}(\Gamma, \boldsymbol{\Lambda}) = u_i(\Gamma, \boldsymbol{\Lambda})$. In summary, promised utility $W_{i,t}$ changes smoothly over time, except when there is an instantaneous payment to agent i or an arrival at any agent. We use the notation $H_{ij,t}$ to represent the jump in agent i's promised utility if there is an arrival at agent j at time t.

Lemma 2.1. *For any contract Γ and any effort process $\boldsymbol{\Lambda}$, there exist $\mathcal{F}_t^N$-predictable processes $\{\boldsymbol{H}_t = (H_{ij,t})_{i,j\in\mathcal{I}}\}_{t\geq 0}$ such that for any t_1 and t_2 with $0 \leq t_1 < t_2$, we have*

$$W_{i,t_2} = W_{i,t_1} + \int_{(t_1,t_2]} \mathrm{d}W_{i,s}, \quad \forall i \in \mathcal{I}, \tag{2.3}$$

in which

$$\mathrm{d}W_{i,t} = \left[\rho W_{i,t-} - \left(R^a + b\,\mathbb{1}_{\{\lambda_{i,t}=\bar{\lambda}\}}\right)X_{i,t} + \sum_{j\in\mathcal{I}} \lambda_{j,t}H_{ij,t}\right]\mathrm{d}t - \sum_{j\in\mathcal{I}} H_{ij,t}\,\mathrm{d}N_{j,t} - \mathrm{d}L_{i,t}, \quad \forall t \geq 0, \ \forall i \in \mathcal{I}, \tag{PK}$$

where the counting process $\{\boldsymbol{N}_t\}_{t\geq 0}$ is generated from the effort process $\boldsymbol{\Lambda}$. In addition, we need

$$H_{ij,t} \leq W_{i,t-}, \quad \forall t \geq 0, \forall i, j \in \mathcal{I}, \tag{2.4}$$

to satisfy (IR).

The *promise keeping* condition (PK) generalizes equation (13) of Biais et al. (2010) to the multiagent case, which ensures that $W_{i,t}$ is indeed agent i's continuation utility starting from time t. The term $H_{ij,t}$, if positive, represents a downward jump in agent i's promised utility caused by agent j's arrival at time t. If $H_{ij,t}$ is negative, then the jump is upward.

It is intuitive that if agent i experiences an adverse event, then its own promised utility takes a downward jump ($H_{ii,t} \geq 0$), which helps align the incentives between the principal and the agent. Allowing $H_{ij,t} \neq 0$ for $i \neq j$, however, is an important modeling construction. A natural extension from the single-agent model may include only the terms $H_{ii,t}$ without $H_{ij,t}$ for all $i \neq j$, and the corresponding dynamics of an agent's promised utilities would follow

$$\mathrm{d}W_{i,t} = \left[\rho W_{i,t-} - \left(R^a + b\,\mathbb{1}_{\{\lambda_{i,t}=\bar{\lambda}\}}\right)X_{i,t} + \lambda_{i,t}H_{ii,t}\right]\mathrm{d}t - H_{ii,t}\,\mathrm{d}N_{i,t} - \mathrm{d}L_{i,t}, \qquad \forall t \geq 0, \ \forall i \in \mathcal{I},$$

which reduces to the single-agent promise keeping condition when $|\mathcal{I}| = 1$. As we will see later in the paper, allowing $H_{ij,t}$ terms to be nonzero not only provides more flexibility in the contract design but is also essential for the existence of EIC contracts. Generally, the sign of $H_{ij,t}$ is unclear a priori, and we do not impose restrictions on it. In Section EC.7 of the online appendix, we explain the implications of allowing $H_{ij,t} \neq 0$ for $i \neq j$ using numerical examples.

2.2. Incentive Compatibility

We focus on contracts that always induce effort from agents, similar to Biais et al. (2010). In Section 6.1, we provide a sufficient condition such that this restriction does not affect optimality. In particular, a contract Γ is called *incentive-compatible* (IC), if it induces all agents to always exert effort and maintain a low arrival rate of adverse events. More precisely, this condition requires that, compared with any other strategy, agent i attains a higher total utility by always exerting effort, given that all other agents also always exert effort. That is,

$$\begin{aligned} &u_i(\Gamma, \bar{\mathbf{\Lambda}}) \geq u_i(\Gamma, \tilde{\mathbf{\Lambda}}_i), \quad \forall i \in \mathcal{I}, \forall \tilde{\mathbf{\Lambda}}_i, \\ \text{where} \quad &\bar{\mathbf{\Lambda}} := \{\lambda_{i,t} = \lambda, \ \forall i \in \mathcal{I}, \ \forall t \geq 0\}, \\ \text{and} \quad &\tilde{\mathbf{\Lambda}}_i := \{\lambda_{i,t} \text{ is } \mathcal{F}_t^N\text{-predictable}, \quad \text{and } \lambda_{j,t} = \lambda, \\ &\qquad \forall j \neq i, \forall t \geq 0\}. \end{aligned} \tag{2.5}$$

Let β denote the ratio between the shirking benefit b and the difference in arrival rates,

$$\beta := \frac{b}{\Delta\lambda}.$$

Intuitively, if the principal were to charge agent i an amount of $\beta X_{i,t}$ for each adverse arrival, then agent i would be indifferent between exerting effort and shirking. Heuristically, in a small time interval δ, agent i enjoys a shirking benefit $bX_{i,t}\delta$, which needs to be offset by a higher penalty cost, $\Delta\lambda\beta X_{i,t}$. Nonetheless, charging agents is not allowed in our setting, so the principal instead reduces agent i's promised utility by at least $\beta X_{i,t}$ for each arrival to induce effort. In summary, the value $\beta X_{i,t}$ is the minimum penalty on agent i if it introduces an adverse event. We formalize this result in the following proposition, which extends proposition 1 of Biais et al. (2010) to our multiagent setting.

Proposition 2.1. *Contract Γ satisfies the incentive-compatible condition (2.5) if and only if*

$$H_{ii,t} \geq \beta X_{i,t}, \quad \forall t \geq 0, \forall i \in \mathcal{I}. \tag{IC}$$

Proposition 2.1 along with condition (PK) imply that if agent i causes an adverse event, then the downward jump $H_{ii,t}$ of its promised utility is at least $\beta X_{i,t}$, in order to incentivize agent i to exert effort. Note that $H_{ij,t}$ for $i \neq j$ is not involved in the (IC) condition. Intuitively, the principal does not need to penalize the agent for an adverse event that is not associated with this agent.

Finally, we require a parameter, $\bar{w}$, as an upper bound to the agents' promised utilities. This parameter reflects the "commitment power" of the principal and is necessary to the model because of the "infinite back-loading" problem, where the principal always prefers to delay paying agents while promising to pay the corresponding interest (Myerson 2015). Without $\bar{w}$, the principal can indefinitely delay the payments such that the promised utilities grow to infinity without agents ever being paid. Therefore, we need the following *upper-bound* (UB) constraints:

$$W_{i,t} \leq \bar{w}, \quad \forall t \geq 0, \forall i \in \mathcal{I}. \tag{UB}$$

2.3. Profit Maximization Objective and EIC Contracts

Under an incentive-compatible contract Γ, the principal's profit is defined as

$$U(\Gamma) := \frac{R^p - \lambda C}{\rho} - \mathbb{E}^{\bar{\mathbf{\Lambda}}}\left[\int_0^\infty e^{-\rho t} \sum_{i \in \mathcal{I}} \mathrm{d}L_{i,t}\right], \tag{2.6}$$

which equals the total discounted revenue minus the cost of adverse events and payments to agents.

In this paper, we study the optimal contract that maximizes the principal's profit (2.6) among all contracts $\Gamma = \{\mathbf{L}_t, \mathbf{X}_t\}_{t\geq 0}$ that satisfy the following definition.

Definition 2.1. We call a contract $\Gamma = \{\mathbf{L}_t, \mathbf{X}_t\}_{t\geq 0}$ an *efficient and incentive-compatible contract*, or *EIC contract*, if there exists a promised utility process $\{\mathbf{W}_t\}_{t\geq 0}$ with $W_{i,0} = u_i(\Gamma, \bar{\mathbf{\Lambda}})$, such that $\mathbf{L}_t$, $\mathbf{X}_t$, and $\mathbf{W}_t$ satisfy (FX), (IC), (IR), (PK), (LL), and (UB).

With information asymmetry, it is not immediately clear whether an EIC contract exists. In fact, in the single-agent case, where effort reduces the arrival rate of adverse events, incentive-compatible contracts generally cannot achieve efficiency. Biais et al. (2010), for example, studies a single-agent problem similar to ours, in which the agent is a firm that must exert effort to reduce the arrival rate of adverse events. The principal (a financier, or society at large) dynamically adjusts the firm size in addition to using cash payments to induce effort. In their setting, the principal has to downsize the firm to yield credible threats when the agent's promised utility is below a certain threshold, which is economically inefficient.

The same logic behind inefficient allocation for a single-agent case works in our setting as well. Consider player i as the only agent. As long as the upper bound $\bar{w}$ is finite, no matter how high it is, for any positive arrival rate λ, there is always a positive probability such that a sequence of frequent arrivals pushes the promised utility below β. At this point, the (IC) condition implies that we cannot maintain incentive compatibility with $X_{i,t} = 1$ while still satisfying (IR). In this case, we must reduce $X_{i,t}$ to satisfy both

(IC) and (IR), which is inefficient. Tian et al. (2021) also discuss a similar intuition behind efficiency and incentive compatibility.

In contrast, and perhaps not a priori obvious, for the multiagent setting of our study, EIC contracts do exist. The fundamental rationale is that when one agent must be penalized by a reduction in its allocated resource, the principal could transfer the revoked resource to other agent(s). In the next section, we formally characterize the set of EIC contracts that form the basis of our characterization of the optimal contract in the following section.

3. Achievable Set of Promised Utilities and the Existence of EIC Contracts

In this section, we show that, under fairly general conditions, EIC contracts exist as long as there are at least two agents. For this purpose, it is helpful to consider the set of promised utilities that EIC contracts (if they exist) generate. Specifically, we define the *achievable set* of promised utilities by EIC contracts as

$$\mathcal{U} := \{\boldsymbol{w} = \{w_i\}_{i\in\mathcal{I}} \in [0,\bar{w}]^n \mid \exists \text{ EIC contract } \Gamma, \text{ such that } w_i = u_i(\Gamma,\bar{\Lambda})\}. \quad (3.1)$$

Therefore, the existence of an EIC contract is equivalent to the nonemptiness of the set of achievable utilities $\mathcal{U}$. Moreover, we claim, and will later show, that if an EIC contract Γ exists and yields a promised utility process $\{\boldsymbol{W}_t\}_{t\geq 0}$, then the definition of the achievable set $\mathcal{U}$ implies that for any time $t \geq 0$, we have $\boldsymbol{W}_t \in \mathcal{U}$. In other words, starting from an achievable set of promised utilities, all future promised utilities must also be achievable by some EIC contract and belong to this achievable set.

This evokes the *self-generating set* concept first introduced in the seminal paper by Abreu et al. (1990) for repeated games, and later used by Fudenberg et al. (1994) for imperfect public information repeated games and Balseiro et al. (2019) for dynamic mechanism design problems. All these papers study games in discrete-time settings and characterize the corresponding self-generating sets in a recursive manner: if a promised utility belongs to the set, then the next period's promised utility, according to the promised keeping constraint, must also belong to this set. One cannot directly adopt the definition of "self-generation" from the discrete-time setting because in a continuous-time setting, the notion of "next period" is not well defined. Bernard and Frei (2016) generalize the self-generating set concept to a continuous-time setting with Brownian motion uncertainties in the context of proving the folk theorem for continuous-time repeated games. In particular, definition 5 of that paper extends the self-generating set concept to the continuous-time setting, which is similar to the definition below for our setting.

Definition 3.1. A set $\mathcal{A} \subseteq [0,\bar{w}]^n$ is a *self-generating* set if, for any $\boldsymbol{W}_0 \in \mathcal{A}$, there exist $\mathcal{F}_t^N$-predictable processes $\boldsymbol{H}_t$, $\boldsymbol{X}_t$, $\boldsymbol{L}_t$, and an $\mathcal{F}_t^N$-adapted process $\{\boldsymbol{W}_t\}_{t\geq 0}$ starting from $\boldsymbol{W}_0$, that satisfy (FX), (IC), (IR), (PK), (LL), and (UB) such that $\boldsymbol{W}_t \in \mathcal{A}$ for all $t \geq 0$.

Definition 3.1 implies that should the agents start with promised utilities inside a self-generating set, their future promised utilities following an EIC contract would always stay in the same set.

Next, we draw the explicit connection between the self-generating set and the achievable set of promised utilities.

Proposition 3.1. *If set $\mathcal{A}$ is a self-generating set, then $\mathcal{A} \subseteq \mathcal{U}$.*

Proposition 3.1 states that every self-generating set is a subset of the achievable set. It implies that for any self-generating set $\mathcal{A}$ and a vector of promised utilities $\boldsymbol{W}_0 \in \mathcal{A}$, we have $\boldsymbol{W}_t \in \mathcal{A} \subseteq \mathcal{U}$ for any time $t \geq 0$. That is, there must exist an EIC contract Γ such that for any $i \in \mathcal{I}$, contract Γ delivers utility $W_{i,0}$ to agent i.

The next proposition further tightens the relationship between self-generating sets and the set of achievable utilities.

Proposition 3.2. *The achievable set $\mathcal{U}$ is a self-generating set.*

Propositions 3.1 and 3.2 imply that the achievable set $\mathcal{U}$ is the largest self-generating set, if it is not empty.

Next, we characterize the achievable set of promised utilities using its *support function*, which characterizes the boundary of a convex set by its normal vectors. In particular, we adopt the approach developed by Balseiro et al. (2019) for discrete-time dynamic mechanism design to our continuous-time moral hazard problem. Later, in Section 4, we further describe the optimal contract based on the characterization of the achievable set.

The general idea is an iterative approach that gradually shrinks the initial set $[0,\bar{w}]^n$ until reaching the largest self-generating set. In each iteration, based on the support function from the previous iteration, we solve a sequence of time-independent static optimization problems to obtain a new support function, which also defines a convex set. The final convex set, if not empty, is a self-generating set, and also the achievable set of promised utilities.

To this end, consider the following support function $\phi_{\mathcal{A}} : \mathbb{R}^n \to \mathbb{R}$ of any set $\mathcal{A} \subset \mathbb{R}^n$:

$$\phi_{\mathcal{A}}(\boldsymbol{\alpha}) := \inf_{\boldsymbol{w}\in\mathcal{A}} \boldsymbol{\alpha}^\top \boldsymbol{w}, \quad \forall \boldsymbol{\alpha} \in \mathbb{R}_+^n \text{ with } \|\boldsymbol{\alpha}\|_1 = 1.$$

The hyperplane $\{\boldsymbol{x} \mid \boldsymbol{\alpha}^\top \boldsymbol{x} = \phi_{\mathcal{A}}(\boldsymbol{\alpha})\}$ is a supporting hyperplane of set $\mathcal{A}$ with normal direction $\boldsymbol{\alpha}$. We focus on supporting hyperplanes with positive normal vectors $\boldsymbol{\alpha} \in \mathbb{R}^n_+$, because if a promised utility $\boldsymbol{w}$ is achievable by an EIC contract, then any promised utility that is component-wise greater than or equal to $\boldsymbol{w}$ is also achievable, as we show in the technical Lemma EC.3.1 in the online appendix. Moreover, from any support function ϕ, we can define a closed convex set as

$$\mathscr{G}(\phi) := \{\boldsymbol{w} \in [0, \bar{w}]^n \mid \boldsymbol{\alpha}^\top \boldsymbol{w} \geq \phi(\boldsymbol{\alpha}), \forall \boldsymbol{\alpha} \in \mathbb{R}^n_+, \|\boldsymbol{\alpha}\|_1 = 1\}, \tag{3.2}$$

referred to as the set characterized by support function ϕ. It is clear that for any set $\mathcal{A} \subseteq [0, \bar{w}]^n$, we must have

$$\mathcal{A} \subseteq \mathscr{G}(\phi_{\mathcal{A}}). \tag{3.3}$$

Next, we define an operator T, which maps from one support function to another, forming the foundation of our iterative approach. This operator is defined through the following linear program for any function $\phi : \mathbb{R}^n_+ \to \mathbb{R}$ and vector $\boldsymbol{\alpha} \in \mathbb{R}^n_+$:

$$[T\phi](\boldsymbol{\alpha}) := \inf_{\boldsymbol{w},\boldsymbol{x},\boldsymbol{y} \in \mathbb{R}^n; \boldsymbol{H},\boldsymbol{Z} \in \mathbb{R}^{n\times n}} \boldsymbol{\alpha}^\top \boldsymbol{w}$$

$$\text{subject to (s.t.)} \quad \sum_{i=1}^n x_i = 1, \quad x_i \geq 0, \quad \forall i \in \mathcal{I}, \tag{FXs}$$

$$H_{ii} \geq \beta x_i, \quad \forall i \in \mathcal{I}, \tag{ICs}$$

$$y_i = \rho w_i - R^a x_i + \lambda \sum_{j\in\mathcal{I}} H_{ij}, \quad \forall i \in \mathcal{I}, \tag{PKy}$$

$$Z_{ij} = w_i - H_{ij}, \quad \forall i, j \in \mathcal{I}, \tag{PKz}$$

$$\hat{\boldsymbol{\alpha}}^\top \boldsymbol{w} \geq \phi(\hat{\boldsymbol{\alpha}}), \quad \forall \hat{\boldsymbol{\alpha}} \in \mathbb{R}^n_+, \|\hat{\boldsymbol{\alpha}}_1\| = 1, \tag{SGw}$$

$$\boldsymbol{\alpha}^\top \boldsymbol{y} \geq 0, \tag{SGy}$$

$$\hat{\boldsymbol{\alpha}}^\top \boldsymbol{Z}_{\cdot j} \geq \phi(\hat{\boldsymbol{\alpha}}), \quad \forall j \in \mathcal{I}, \forall \hat{\boldsymbol{\alpha}} \in \mathbb{R}^n_+, \|\hat{\boldsymbol{\alpha}}\|_1 = 1, \tag{SGz}$$

$$w_i \geq 0, \quad \forall i \in \mathcal{I}, \tag{IRs}$$

$$w_i \leq \bar{w}, \quad \forall i \in \mathcal{I}, \tag{UBs}$$

where $\boldsymbol{Z}_{\cdot j}$ represents the vector $(Z_{ij})_{i\in\mathcal{I}}$. For any normal vector $\boldsymbol{\alpha}$, the linear program $[T\phi](\boldsymbol{\alpha})$ returns a real value. Therefore, $T\phi$ is also a function that maps $\mathbb{R}^n_+$ to $\mathbb{R}$, as is the case with ϕ.

In the optimization problem $[T\phi](\boldsymbol{\alpha})$, the decision variables $\boldsymbol{x}$ and $\boldsymbol{w}$ correspond to resource allocation and the current promised utilities, respectively. The decision variable H_{ij} in $\boldsymbol{H}$ corresponds to the jump to agent i's promised utility upon the arrival of an adverse event at agent j. It is easy to see that constraints (FXs), (ICs), (IRs), and (UBs) resemble (FX), (IC), (IR), and (UB) for an EIC contract, respectively. Note that this linear optimization model does not include payment decisions. In fact, we later show that it is sufficient to use this linear optimization to describe the achievable set, and adding decision variables representing payments does not help. As we mentioned earlier, this linear optimization also helps us construct a particular EIC contract. In such a construction, we can identify payments from the optimal solution of this linear program.

To explain variables $\boldsymbol{y}$ and $\boldsymbol{Z}$, it is helpful to rewrite condition (PK) without payments as

$$\mathrm{d}W_{i,t} = \left(\rho W_{i,t-} - R^a X_{i,t} + \lambda \sum_{j\in\mathcal{I}} H_{ij,t}\right) \mathrm{d}t - \sum_{j\in\mathcal{I}} H_{ij,t}\, \mathrm{d}N_{j,t}, \quad \forall t \geq 0, \ \forall i \in \mathcal{I}. \tag{3.4}$$

Therefore, constraint (PK_y) defines variable y_i as the smoothly changing term that multiplies dt in (3.4), and (PK_Z) implies that Z_{ij} represents the change in agent i's promised utility when an arrival occurs at agent j ($\mathrm{d}N_{j,t} = 1$ in (3.4)).

Finally, constraints (SGw), (SGy), and (SGz) capture the self-generating property. In particular, constraint (SGw) implies that the optimal $\boldsymbol{w}$ to this linear program must satisfy $[T\phi](\boldsymbol{\alpha}) = \boldsymbol{\alpha}^\top \boldsymbol{w} \geq \phi(\boldsymbol{\alpha})$ for all $\boldsymbol{\alpha} \in \mathbb{R}^n_+$ and $\|\boldsymbol{\alpha}\|_1 = 1$, which further indicates that for any function ϕ,

$$\mathscr{G}(T\phi) \subseteq \mathscr{G}(\phi). \tag{3.5}$$

This result further implies that iteratively applying operator T to a support function ϕ generates a sequence of ever-shrinking convex sets. The question is whether the limiting set in this sequence is desirable. The following result helps answer this question.

Lemma 3.1. *If $\mathcal{A}$ is a self-generating set, then we have*
1. *$\mathcal{A} \subseteq \mathscr{G}(T\phi_{\mathcal{A}})$, and*
2. *$\mathscr{G}(\phi_{\mathcal{A}})$ is a self-generating set.*

Conversely, if a convex set $\mathcal{A}$ satisfies $\mathcal{A} \subseteq \mathscr{G}(T\phi_{\mathcal{A}})$, then $\mathcal{A}$ is a self-generating set.

Lemma 3.1 implies that for a convex set $\mathscr{G}(\phi)$ to be self-generating, a necessary and sufficient condition is $\mathscr{G}(\phi) \subseteq \mathscr{G}(T\phi)$. In addition, considering (3.5), we know that a set $\mathscr{G}(\phi)$ being self-generating is equivalent to $\mathscr{G}(\phi) = \mathscr{G}(T\phi)$. Therefore, iteratively applying operator T to a support function, we obtain a self-generating set in the limit. The following theorem states that such a self-generating set is the achievable set.

Theorem 3.1. *Let $\mathcal{U}^0 = [0, \bar{w}]^n$, and define operator T^k such that $T^k\phi = T(T^{k-1}\phi)$ for all $k > 1$. Then, we have*

$$\lim_{k\to\infty} \mathscr{G}(T^k \phi_{\mathcal{U}^0}) = \mathcal{U} = \mathscr{G}(T\phi_{\mathcal{U}}) = \mathscr{G}(\phi_{\mathcal{U}}).$$

Theorem 3.1 further implies the following corollary.

Corollary 3.1. *The set of achievable utilities, $\mathcal{U}$, is closed and convex.*

If the achievable set $\mathcal{U}$ is nonempty, then we can construct EIC contracts and calculate the dynamics of agents' promised utilities from the feasible solutions to the linear program $[T\phi_{\mathcal{U}}](\boldsymbol{\alpha})$. We later exemplify this construction by describing an important EIC contract in Section 4. The next result addresses the condition under which the achievable set $\mathcal{U}$ exists.

Proposition 3.3. *For $n \geq 2$, there exists a threshold $\bar{\omega}$ that depends on model parameters n, ρ, b, λ, and $\bar{\lambda}$, such that the achievable set $\mathcal{U}$ is nonempty if and only if $\bar{w} \geq \bar{\omega}$. Furthermore, for $n = 1$, the achievable set $\mathcal{U}$ is empty.*

In the proof of Proposition 3.3, we also show that $\bar{\omega}$ is nonincreasing in n. Proposition 3.3 states that EIC contracts do not exist with a single agent in a continuous-time setting. In fact, EIC contracts exist in a discrete-time setting. For example, the principal can pay the agent β in each period if there is no adverse arrival and pay nothing if there is. At the cost of the principal, this contract mitigates the necessity of relying on resource allocation reduction to ensure incentive compatibility. However, if we shrink the length of the discrete-time period to approach the continuous-time model, then the total payment per unit of time would approach infinity. Such a contract is undesirable and violates the (UB) constraint.

Figure 1 illustrates the achievable sets found by the iterative approach for the two- and three-agent cases. Note that because the optimization problem $[T\phi](\boldsymbol{\alpha})$ is a semi-infinite linear program, we solve it approximately by considering a subset of constraints (SGw) and (SGz) only for $\hat{\boldsymbol{\alpha}}$ on a grid, such that $\hat{\boldsymbol{\alpha}} \geq 0$ and $\|\hat{\boldsymbol{\alpha}}\|_1 = 1$. Figure 1(a) also plots the supporting hyperplanes in solid lines, similar to figure 1(a) in the electronic companion of Balseiro et al. (2019). Notably, with two agents, the set $\mathcal{U}$ does not intersect with the axes, indicating that the promised utilities of both agents inside $\mathcal{U}$ are strictly positive. Equivalently, an EIC contract never terminates either of the two agents.

In Figure 1(b), however, the three hyperplanes defined by $w_1 = 0$, $w_2 = 0$, and $w_3 = 0$ do contribute to the boundary of the achievable set. That is, an EIC contract may terminate one of the agents. Furthermore, in this case, the intersection between the achievable set $\mathcal{U}$ and hyperplane $w_i = 0$ for $i = 1, 2$, or 3 is the achievable set in the corresponding two-agent cases without agent i. In other words, an EIC contract for three agents may reduce to an EIC contract for a two-agent case upon terminating one agent. This observation is intuitive. If the current promised utilities for the three agents is $\boldsymbol{w} = (w_1, w_2, w_3 = 0)$, then any EIC contract cannot increase w_3 and make it positive following (PK). Consequently, conditions $(w_1, w_2) \in \mathcal{U}(2)$ and $\boldsymbol{w} \in \mathcal{U}(3)$ are equivalent, where we use notation $\mathcal{U}(n)$ to highlight the achievable set in a setting with n agents. In general, the aforementioned logic suggests that on the boundary of the achievable set for n agents, if we restrict one agent's promised utility to zero, then we obtain the achievable set for $n-1$ agents.

The linear program $[T\phi_{\mathcal{U}}](\boldsymbol{\alpha})$ also directly demonstrates the following sensitivity result on how the model parameters affect the achievable set $\mathcal{U}$. Here, we use notation $\mathcal{U}(\theta)$ to highlight model parameter θ's impact on the achievable set, in which θ could be $n, b, \rho, \Delta\lambda$, or $\bar{w}$.

Proposition 3.4. *We have $\mathcal{U}(b_1) \supseteq \mathcal{U}(b_2)$ for $b_1 \leq b_2$, while keeping other model parameters the same. Similarly, we have $\mathcal{U}(\rho_1) \subseteq \mathcal{U}(\rho_2)$ for $\rho_1 \leq \rho_2$; $\mathcal{U}(\Delta\lambda_1) \subseteq \mathcal{U}(\Delta\lambda_2)$ for $\Delta\lambda_1 \leq \Delta\lambda_2$; and $\mathcal{U}(\bar{w}_1) \subseteq \mathcal{U}(\bar{w}_2)$ for $\bar{w}_1 \leq \bar{w}_2$.*

Finally, we note that allowing the terms $H_{ij,t}$ for $i \neq j$ to be nonzero is an essential condition for the existence of EIC contracts. Recall that $H_{ij,t}$ denotes the discrete jump in agent j's promised utility when agent i experiences an adverse event occurs at time t. The following result highlights the importance of allowing the promised utilities of the other agents to change when one agent experiences an adverse arrival.

Proposition 3.5. *Introduce constraints $H_{ij} = 0$ for all $i \neq j$ to the linear programs $[T\phi](\boldsymbol{\alpha})$ to obtain a new linear program $[\tilde{T}\phi](\boldsymbol{\alpha})$. We have $\lim_{k\to\infty} \mathcal{G}(\tilde{T}^k \phi_{\mathcal{U}_0}) = \emptyset$, which suggests that there does not exist an EIC contract with $H_{ij,t} = 0$ for all $t \geq 0$ and $i \neq j$.*

Before we close this section, it is worth reviewing the key results. In this section, we demonstrate the existence of EIC contracts by providing an iterative approach to obtain the achievable set of promised utilities, as shown in Theorem 3.1. Notably, the existence of our EIC contracts only requires that the upper bound $\bar{w}$ of the promised utilities is high enough and does not rely on agents being infinitely patient (Proposition 3.3). Proposition 3.5 further reveals that, generally, to achieve efficiency, whenever there is an adverse event arriving at one agent, all agents' allocations and promised utilities need to take discrete jumps.

4. Optimal EIC Contract

In this section we study the optimal contract in the set of EIC contracts that maximizes the principals' utility. In particular, the linear program $[T\phi_{\mathcal{U}}](\boldsymbol{\alpha})$ introduced in the last section not only provides the achievable set but also yields a set of EIC contracts. We call it the set of *boundary* EIC contracts, because the promised utilities remain on the boundary of the achievable set. After formally defining the boundary of the achievable set, we establish that the optimal EIC contract that maximizes the principal's utility is a boundary contract that sets all agents' initial promised utilities at the same value.

First, we formally define the *boundary* of an n-dimensional achievable set $\mathcal{U}(n) \subset [0, \bar{w}]^n$ as

$$\mathrm{bd}(\mathcal{U}(n)) := \mathrm{cl}\{\boldsymbol{w} \in \mathcal{U} \text{ and } \boldsymbol{w} > 0 \mid \exists \boldsymbol{\alpha} \in \mathbb{R}_+^n \text{ and } \|\boldsymbol{\alpha}\|_1 = 1 \text{ such that } \boldsymbol{\alpha}^\top \boldsymbol{w} = \phi_{\mathcal{U}}(\boldsymbol{\alpha})\}.$$

Therefore, it is the closure of the set of component-wise positive vectors on the boundary of the achievable set.

Figure 1. Illustration of Achievable Sets

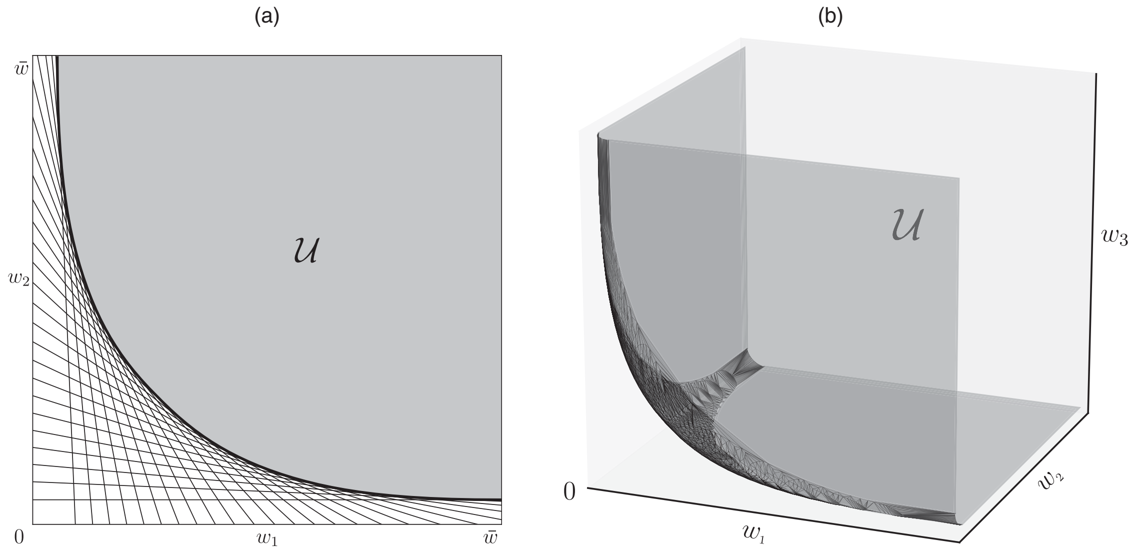

Note. (a) Two-agent case; (b) three-agent case.

We now present the key result of this section, which indicates that the optimal contract sets all agents' initial promised utilities to be the same on the boundary of the achievable set. Recall the definition of $u_i(\Gamma, \mathbf{\Lambda})$ from (2.1). With a slight abuse of notation, we use the following simplification to represent the agents' promised utilities from an EIC contract Γ under full effort:

$$\boldsymbol{u}(\Gamma) := \{u_i(\Gamma)\}_{i\in\mathcal{I}}, \text{ in which } u_i(\Gamma) := u_i(\Gamma, \bar{\Lambda}).$$

Clearly, for any EIC contract Γ, we have $\boldsymbol{u}(\Gamma) \in \mathcal{U}$.

Theorem 4.1. *Consider the EIC contracts $\hat{\Gamma}$ such that $\boldsymbol{u}(\hat{\Gamma}) \in \mathrm{bd}(\mathcal{U})$ and $u_i(\hat{\Gamma}) = u_j(\hat{\Gamma})$ for all $i \neq j$. We have, for any EIC contract Γ,*

$$U(\hat{\Gamma}) = \frac{R - \lambda C}{\rho} - \sum_i u_i(\hat{\Gamma}) \geq U(\Gamma).$$

Theorem 4.1 establishes that contract $\hat{\Gamma}$ is the optimal EIC contract for the principal. Theorem 4.1 and Proposition 3.4 further imply the following intuitive result.

Corollary 4.1. *The principal's total discounted utility under the optimal EIC contract increases in $\bar{w}$, R^p, n, and $\Delta\lambda$ (while keeping λ fixed), and decreases in b and C.*

Next, we characterize the complete dynamics of the optimal contract beyond the starting promised utilities. In particular, we establish two results. First, if the promised utility vector starts on the boundary of the achievable set, then it will stay on the boundary. Second, for any contract that starts the promised utilities on the boundary $\mathrm{bd}(\mathcal{U})$, the optimal solution to the linear program $[T\phi_{\mathcal{U}}](\boldsymbol{\alpha})$ identifies the dynamics of the promised utilities thereafter.

For this purpose, we first present the following technical result, which connects the linear program $[T\phi_{\mathcal{U}}](\boldsymbol{\alpha})$ introduced in the last section with the dynamics of promised utility.

Lemma 4.1. *At any optimal solution to the linear program $[T\phi_{\mathcal{U}}](\boldsymbol{\alpha})$, constraints (ICs) and (SGy) hold as equalities, and constraint (SGw) holds as an equality for $\hat{\boldsymbol{\alpha}} = \boldsymbol{\alpha}/\|\boldsymbol{\alpha}\|_1$. Similarly, for each $j \in \mathcal{I}$, there exists an $\hat{\boldsymbol{\alpha}}$ such that the corresponding (SGz) constraint holds as an equality.*

Lemma 4.1 reveals important geometric properties of the dynamics of promised utilities. In particular, constraints (SGy), (SGw), and (SGz) holding as equality at optimality implies that if we start the promised utilities on the boundary, then both the drift directions and the discrete jumps guarantee that the promised utilities stay on the boundary. Later in this section, we prove these claims formally in Proposition 4.1 based on Lemma 4.1. We require a few more new notations to fully characterize the dynamics of the promised utility. First, we need to characterize the normal vector of the boundary $\mathrm{bd}(\mathcal{U})$. To this end, for any $\boldsymbol{w} \in \mathbb{R}^n_+$, we define $\Pi(\boldsymbol{w}) \in \mathrm{bd}(\mathcal{U})$ as a projection of $\boldsymbol{w}$ onto the boundary $\mathrm{bd}(\mathcal{U})$, such that $\Pi(\boldsymbol{w})$ solves

$$\min_{\boldsymbol{\xi}\in\mathrm{bd}(\mathcal{U})} \|\boldsymbol{w} - \boldsymbol{\xi}\|_2. \tag{4.1}$$

Based on this projection, we define the mapping $\check{\boldsymbol{\alpha}}: \mathbb{R}^n_+ \to \mathbb{R}^n_+$, such that

$$\check{\boldsymbol{\alpha}}(\boldsymbol{w})^\top \Pi(\boldsymbol{w}) = \phi_{\mathcal{U}}(\check{\boldsymbol{\alpha}}(\boldsymbol{w})). \tag{4.2}$$

That is, for any $\boldsymbol{w} \in \mathcal{U}$, we have $\boldsymbol{w} - \Pi(\boldsymbol{w}) = c(\boldsymbol{w})\,\check{\boldsymbol{\alpha}}(\boldsymbol{w})$ for some scalar $c(\boldsymbol{w}) \in \mathbb{R}_+$ associated with $\boldsymbol{w}$. Therefore,

$\check{\boldsymbol{\alpha}}(\boldsymbol{w})$ is the normal vector of $\text{bd}(\mathcal{U})$ at any point $\boldsymbol{w} \in \text{bd}(\mathcal{U})$.

Furthermore, we need to consider the possibility of dropping an agent's promised utility to zero when there are more than two agents. For a rigorous presentation, we introduce the corresponding notation. When $n = 2$, any $\boldsymbol{w}$ on the boundary $\text{bd}(\mathcal{U})$ must be component-wise positive. For $n > 2$, however, up to $n - 2$ components of $\boldsymbol{w} \in \text{bd}(\mathcal{U})$ may be zero. For instance, as shown in Figure 1(b), the boundary $\text{bd}(\mathcal{U})$ corresponds to the dark meshed part of the set $\mathcal{U}$. The intersections of this boundary with the hyperplanes $w_i = 0$ for $i = 1, 2,$ or 3 reduce to the boundary sets for the achievable sets of respective two-agent settings. Therefore, we introduce the notation $\mathscr{C}(\boldsymbol{w})$ to represent collapsing the vector $\boldsymbol{w}$ into another vector that contains only its positive elements; that is,

$$\mathscr{C}(\boldsymbol{w}) := (w_i)_{i:w_i>0}.$$

It is clear that if the dimension of $\mathscr{C}(\boldsymbol{w})$ is $m < n$ for vector $\boldsymbol{w} \in \text{bd}(\mathcal{U}(n))$, then we must have $\mathscr{C}(\boldsymbol{w}) \in \text{bd}(\mathcal{U}(m))$.

Equipped with notations $\check{\boldsymbol{\alpha}}(\boldsymbol{w})$ and $\mathscr{C}(\boldsymbol{w})$, we are ready to present the dynamics of promised utilities. Starting from any vector of promised utilities $\boldsymbol{W}_0 \in \text{bd}(\mathcal{U})$, we define a promised utility process $\{\boldsymbol{W}_t\}_{t\geq 0}$ together with an $\mathcal{F}^N$-adapted process $\{\boldsymbol{\alpha}_t\}_{t\geq 0}$, a payment process $\{\boldsymbol{L}_t\}_{t\geq 0}$, an allocation process $\{\boldsymbol{X}_t\}_{t\geq 0}$, and a process of jumps $\{\boldsymbol{H}_t\}_{t\geq 0}$, such that we replace (2.3) with

$$\boldsymbol{W}_{t_2} = \mathscr{C}\left(\left(W_{i,t_1} + \int_{(t_1,t_2]} \mathrm{d}W_{i,s}\right)_{i:W_{i,t_1}>0}\right), \ \forall t_1 < t_2, \tag{4.3}$$

and define

$$\boldsymbol{\alpha}_t := \check{\boldsymbol{\alpha}}(\boldsymbol{W}_t), \tag{4.4}$$
$$\boldsymbol{X}_t := \boldsymbol{x}^*(\boldsymbol{\alpha}_{t-}), \tag{4.5}$$
$$\boldsymbol{H}_t := \boldsymbol{H}^*(\boldsymbol{\alpha}_{t-}), \tag{4.6}$$
$$\mathrm{d}L_{i,t} := (y_i^*(\boldsymbol{\alpha}_{t-}))^+ \, \mathbb{1}_{w_i^*(\boldsymbol{\alpha}_{t-})=\bar{w}} \, \mathrm{d}t, \ \text{and} \tag{4.7}$$
$$\mathrm{d}W_{i,t} := y_i^*(\boldsymbol{\alpha}_{t-})\,\mathrm{d}t + \sum_{j\in\mathcal{I}} \left(Z_{ij}^*(\boldsymbol{\alpha}_{t-}) - w_i^*(\boldsymbol{\alpha}_{t-})\right) \mathrm{d}N_{j,t} - \mathrm{d}L_{i,t}, \tag{4.8}$$

in which we use the notations $\boldsymbol{H}^*(\boldsymbol{\alpha})$, $\boldsymbol{w}^*(\boldsymbol{\alpha})$, $\boldsymbol{x}^*(\boldsymbol{\alpha})$, $\boldsymbol{y}^*(\boldsymbol{\alpha})$, and $\boldsymbol{Z}^*(\boldsymbol{\alpha})$ to represent the optimal decision variables from $[T\phi_{\mathcal{U}}](\boldsymbol{\alpha})$.

The next proposition indicates that there exists an EIC contract that yields a promised utility process following (4.3)–(4.8), and the promised utilities $\boldsymbol{W}_t$ always stay on the boundary $\text{bd}(\mathcal{U})$.

Proposition 4.1. *Starting from any $\boldsymbol{W}_0 \in \text{bd}(\mathcal{U})$, the processes $\{\boldsymbol{W}_t\}_{t\geq 0}$, $\{\boldsymbol{X}_t\}_{t\geq 0}$, $\{\boldsymbol{L}_t\}_{t\geq 0}$, and $\{\boldsymbol{H}_t\}_{t\geq 0}$ defined in* (4.3)–(4.8) *satisfy* (FX), (IC), (IR), (PK), (LL), *and* (UB). *Furthermore, $\boldsymbol{W}_t \in \text{bd}(\mathcal{U})$ for all $t \geq 0$.*

Remark 4.1. It is worth explaining the payment process described in (4.7). According to this expression, the specified EIC contract involves no instantaneous payment to an agent. This implies that upon an arrival at agent j, the upward jump $-H_{ij}^*$ to agent i's promised utility is upper-bounded, such that $w_i^* - H_{ij}^* \leq \bar{w}$. Otherwise, agent i would receive an instantaneous payment equal to the difference, $w_i^* - H_{ij}^* - \bar{w}$, in order to guarantee that agent i's promised utility does not exceed the upper bound $\bar{w}$. In addition, the flow payment occurs only when an agent's promised utility hits the upper bound $\bar{w}$. It is clear that the payment $\boldsymbol{L}_t$ according to (4.7) satisfies condition (LL). Therefore, although we do not explicitly include payment-related decision variables in the linear program $[T\phi_{\mathcal{U}}](\boldsymbol{\alpha})$, the optimal solution of this linear program allows us to construct an EIC contract, including both the promised utility and the payment processes. However, note that not all EIC contracts are defined according to Proposition 4.1. For instance, we cannot rule out the possibility that an EIC contract exists in which the payment is set as

$$\mathrm{d}L_{i,t} = \sum_j (w_i^* - H_{ij}^* - \bar{w})^+ \, \mathrm{d}N_{j,t} + (y_i^*)^+ \, \mathbb{1}_{w_i^*=\bar{w}} \, \mathrm{d}t,$$

where x^+ represents $\max\{x, 0\}$, and the instantaneous payment $(w_i^* - H_{ij}^* - \bar{w})^+$ is positive. □

The next result further states that following *any* EIC contract, if the vector of promised utilities is on the lower-left boundary of $\mathcal{U}$ at some point in time, then they always remain on the boundary thereafter.

Corollary 4.2. *For any EIC contract that yields a promised utility process $\{\boldsymbol{W}_t\}_{t\geq 0}$, if $\boldsymbol{W}_t \in \text{bd}(\mathcal{U})$ for some time $t \geq 0$, then $\boldsymbol{W}_{t'} \in \text{bd}(\mathcal{U})$ for all $t' \geq t$.*

Corollary 4.2 motivates us to focus on EIC contracts that keep the promised utility on the boundary $\mathcal{U}$ throughout the time horizon. We call these contracts *boundary EIC contracts*. In particular, the following corollary indicates that for any boundary EIC contract, the flow payment is zero unless one agent's promised utility is at $\bar{w}$.

Corollary 4.3. *For any boundary EIC contract, the flow payment $l_{i,t}$ is always zero, except when agent i's promised utility is at $\bar{w}$.*

Before closing this section, we summarize the computational procedure to obtain the optimal EIC contract. First, we obtain the achievable set $\mathcal{U}$ following the iterative procedure described in Theorem 3.1. Next, Theorem 4.1 demonstrates that we should start the promised utilities from the "midpoint" of the boundary of the achievable set, where the initial promised utility for each agent is the same. Finally, the support function $\phi_{\mathcal{U}}$ and the corresponding optimal solutions to the sequence of linear optimization problems $[T\phi_{\mathcal{U}}](\boldsymbol{\alpha})$

reveal the dynamics (4.3)–(4.8) that the optimal contract follows.

The optimal EIC contract has some nice features that make it easy to implement. For example, an agent is paid a constant flow payment only when the promised utility reaches the upper bound. Section EC.7 of the online appendix provides additional results and numerical examples illustrating the dynamics of the promised utilities for the two- and three-agent cases. Overall, however, the boundary contract can be quite complex to fully specify and accurately implement, because the promised utilities change constantly. This observation motivates us to propose two much simpler EIC contracts, which also possess further desirable properties.

5. Easy-to-Implement EIC Contracts

Although we can approximate the optimal EIC contract by solving a sequence of linear optimization models $[T\phi_{\mathcal{U}}](\boldsymbol{\alpha})$ with different $\boldsymbol{\alpha}$'s, such a procedure is complex from an operations perspective. In particular, at essentially any point in time, all agents' promised utilities keep moving. The movement, which is part of the solution to a linear program, is generally hard to characterize. Furthermore, in order to calculate the boundary of the achievable set accurately, the number of linear programs we must solve grows exponentially with the number of agents in the system. The size of each linear program $[T\phi_{\mathcal{U}}](\boldsymbol{\alpha})$ also grows with the number of agents due to constraints (SGw) and (SGz). Furthermore, the entire position and shape of the boundary of the achievable set $\mathcal{U}$, and therefore the optimal contract, are highly sensitive to the choice of the upper-bound $\bar{w}$ parameter. This upper bound captures the principal's commitment power in theory but may be hard to estimate in practice. In the spirit of finding solutions that are easy to calculate and implement, we propose two EIC contracts with closed-form expressions in this section.

We refer to the first easy-to-implement contract as the *simple EIC contract*. It has the following desirable properties in contrast to the boundary EIC contract. First, the agents' promised utilities and allocations do not change between arrivals. This property greatly simplifies implementation, because we only need to consider the jumps in promised utilities and allocations upon an adverse arrival. Second, our simple EIC contract relies only on allocation, and not payment, to induce effort. That is, at any time t, agent i's income is $X_{i,t}R^a$ and the principal only adjusts $X_{i,t}$ over time. This property is particularly desirable because, in an online platform setting, agents are often independent businesses operating on the platform, and the resource is the total online visits to the platform. In such a setting, an agent's income is often proportional to the total volume of visits to the agent's web page, and which agent's web page to show to the next online visit constitutes a key lever of the platform. Therefore, modeling agents' income as proportional to allocation is relevant in practice. Third, the simple EIC contract no longer relies on the exogenous parameter $\bar{w}$ as an input parameter. Instead, according to our simple EIC contract, the entire set of promised utilities achievable by the simple EIC contract lies on a simplex, such that $\sum_i W_{i,t} = \hat{w}$ for some $\hat{w}$, which is the total discounted revenue that agents would share according to the EIC contract. In particular, the total discounted revenue that agents receive is

$$\hat{w} := \frac{R^a}{\rho}. \tag{5.1}$$

Therefore, as long as the principal's commitment power, $\bar{w}$, is no less than $\hat{w} - (n-1)\check{w}$ or, equivalently, $\bar{w} \ge (R^a - \lambda\beta)/\rho$, the simple contract satisfies the (UB) constraint.

Further, let $\check{w}$ denote the lowest promised utility of an agent under the simple contract, when the agent's allocation is zero. In particular, we define

$$\check{w} := \frac{\lambda\beta}{(n-1)\rho}. \tag{5.2}$$

If all $n-1$ agents' promised utilities are at $\check{w}$, then the only agent that receives all the resource must have a promised utility of $\hat{w} - (n-1)\check{w}$. When this agent experiences an arrival, the promised utility must take a downward jump of at least $\beta \times 1$, following (FX) and (IC). Therefore, we require

$$\hat{w} - (n-1)\check{w} - \beta \ge \check{w},$$

which implies the following lower bound for the agents' revenue rate R^a to make our simple contract work:[3]

$$R^a \ge \left(\frac{n}{n-1}\lambda + \rho\right)\beta. \tag{5.3}$$

That is, the agents' promised utilities belong to the following set, which is a subset of an $(n-1)$-dimensional simplex:

$$\mathcal{U}_s := \left\{ \boldsymbol{w} \,\middle|\, \sum_{i\in\mathcal{I}} w_i = \hat{w},\ \check{w} \le w_i \le \hat{w} - (n-1)\check{w} \right\}.$$

With this setup, we formally define the *simple EIC contract*.

Definition 5.1. For any R^a that satisfies conditions (5.3), and $\boldsymbol{w} \in \mathcal{U}_s$, in which $\hat{w}$ and $\check{w}$ are defined as in (5.1) and (5.2), respectively, we define a *simple EIC contract* $\Gamma_s(\boldsymbol{w}; R^a)$ such that the payments are

$$l_{i,t} = 0 \text{ and } I_{i,t} = 0, \tag{5.4}$$

the allocations satisfy

$$X_{i,t} = \frac{W_{i,t-} - \check{w}}{\hat{w} - n\check{w}}, \tag{5.5}$$

and the promised utilities start from $\boldsymbol{W}_0 = \boldsymbol{w}$ and follow the dynamic

$$dW_{i,t} = -\sum_{j\in\mathcal{I}} H_{ij,t} dN_{j,t}, \tag{5.6}$$

in which

$$H_{ij,t} = \begin{cases} X_{i,t}\beta, & \text{for } i = j, \\ -\dfrac{X_{j,t}\beta}{n-1}, & \text{for } i \neq j. \end{cases} \tag{5.7}$$

It is worth highlighting some intuitions behind the closed-form expressions and conditions related to Definition 5.1. First, according to (5.5), if an agent's promised utility is at the lower bound $\check{w}$, then this agent receives zero resource. Conversely, if one agent's promised utility is at the upper bound, $\hat{w} - (n-1)\check{w}$, then all the resource must be allocated to this agent, and the definition of $\mathcal{U}_s$ implies that all other agents' promised utilities must be on the lower bound.

Next, the dynamic (5.6) implies that the promised utilities change only upon arrivals. As we mentioned earlier, this property significantly simplifies the contract implementation, especially compared with the boundary EIC contract. The change in promised utilities (5.7) further indicates that when one agent experiences an adverse arrival, all other agents' promised utilities increase by evenly splitting the promised utility loss of the focal agent. Expression (5.7) also ensures that $\sum_i H_{ij,t} = 0$, which further guarantees that $\sum_i W_{i,t}$ remains a constant, and therefore the promised utilities stay on a simplex. Furthermore, in the proof of the following theorem, we illustrate that in order to construct a contract with the desirable properties (5.4)–(5.7), we have to set $\check{w}$ according to (5.2).

Theorem 5.1. *Contract $\Gamma_s(\boldsymbol{w}; R^a)$ in Definition 5.1 satisfies (FX), (IC), (IR), (LL), (PK), and (UB) as long as $\bar{w} \geq \hat{w} - (n-1)\check{w}$ and is therefore an EIC contract. Furthermore, $\boldsymbol{W}_t \in \mathcal{U}_s$ for any $t \geq 0$ starting from $\boldsymbol{W}_0 = \boldsymbol{w}$ following (5.6). Therefore, $\mathcal{U}_s$ is a self-generating set.*

The proof of Theorem 5.1 critically depends on the lower-bound condition (5.3) for the agents' total income rate R^a. It is worth discussing the intuition behind why we need this lower bound. We need a lower bound on R^a because, in this setting, we assume that the only income that agents receive is proportional to the resource allocation, as specified in (5.4). If R^a is too low, then the set of promised utilities would be too small to allow a downward jump of at least $\beta X_{i,t}$ within the set. Consequently, the contract would no longer be both efficient and incentive-compatible. Nonetheless, it is worth noting that in case R^a is not large enough for condition (5.3) to hold, we can slightly adjust the flow payment to $l_{i,t} = \left[\left(\frac{n}{n-1}\lambda + \rho\right)\beta - R^a\right] X_{i,t}$ in Definition 5.1 for the simple contract to work. That is, the principal compensates the agents with flow payments proportional to the allocated resource, which is still quite easy to implement and therefore attractive in practice.

Although Theorem 5.1 establishes that the simple contract is an EIC contract, it may not be optimal for the principal. In particular, if the principal can promise an upper bound $\bar{w} > \hat{w} - (n-1)\check{w}$ while $R^a < \left(\frac{n}{n-1}\lambda + \rho\right)\beta$, then the simple contract is achievable by paying additional flow payments, as illustrated above, but is not optimal. If the upper bound is not high enough, on the other hand, for example, if $\bar{w} < \beta \left[\frac{\lambda}{\rho(n-1)} + 1\right]$, then the simple contract does not even exist. In these cases, it is necessary to resort to the optimal contract discussed in the previous section.

We use Figure 2 to illustrate the self-generating set $\mathcal{U}_s$ in a three-agent case, as well as how promised utilities change upon an arrival. First, the triangular shape in the figure represents the set $\mathcal{U}_s$, which is similar to a 2-simplex. At each extreme point, two of the three agents' promised utilities take the value $\check{w}$, while the other is $\hat{w} - (n-1)\check{w}$. In this figure, we consider a particular point o, representing the current promised utilities of the three agents. If an arrival occurs, then the promised utility may jump to point a, b, or c, depending on which agent suffers from the arrival. Note that the jump from o to a, b, or c is always perpendicular to one of the facets of this self-generating set. This is because as the point jumps toward one of the boundaries (the promised utility of the corresponding agent decreases), all other agents' promised utility increases equally, following (5.7). We show this result formally in Section EC.5.2 in the online appendix. This three-agent example also highlights another key difference between the simple contract Γ_s and the boundary EIC contract. Recall that a boundary EIC contract eventually terminates all but two agents. In contrast, under the simple contract, an agent's promised utility is lower-bounded by $\check{w}$ and no agent is

Figure 2. An Achievable Set Under Contract Γ_s and the Changes in Promised Utilities Upon an Adverse Arrival for a Three-Agent Case

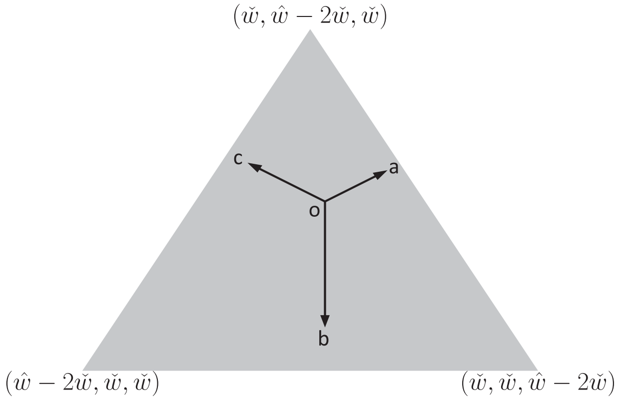

ever terminated. Even with zero resource and income at some point in time, an agent's promised utility remains positive because the next arrival at another agent with positive resource brings this agent's allocation and income flow back to a positive value.

We now present the second easy-to-implement EIC contract, and call it the *rotating contract*.[4] This contract also possesses the following desirable properties: (1) the agents' promised utilities and allocations do not change between arrivals, (2) the contract relies only on allocation and not payment to induce effort, and (3) the contract no longer relies on the exogenous parameter $\bar{w}$ as an input parameter. In summary, the contract allocates all the resource to one agent and rotates it to the next whenever an arrival occurs to the agent currently holding the resource. Under this contract, the agents' promised utilities take values from the set

$$\mathcal{W}_r = \left\{ w_{(i)} = \frac{R^a \tau^{i-1}}{(\lambda+\rho)(1-\tau^n)} \right\}_{i=1,\ldots,n}, \tag{5.8}$$

in which we define the notation

$$\tau = \frac{\lambda}{\lambda+\rho}.$$

The promised utility of the agent holding the resource is $w_{(1)}$, which will drop to $w_{(n)}$ upon the next arrival, and then climb back through the values in $\mathcal{W}_r$ until the agent holds the resource again. Such a contract is incentive-compatible if

$$R^a \geq \left(\frac{\rho}{1-\tau^{n-1}} + \lambda \right) \beta. \tag{5.9}$$

We define the rotating contract formally as follows.

Definition 5.2. For any R^a that satisfies (5.9), define notation $\hat{\beta} = \frac{R^a}{\frac{\rho}{1-\tau^{n-1}}+\lambda} \geq \beta$. A *rotating contract* is a contract such that the payments are

$$l_{i,t} = 0 \text{ and } I_{i,t} = 0,$$

the allocations satisfy

$$X_{i,t} = \begin{cases} 1, & \text{if } W_{i,t} = w_{(1)}, \\ 0, & \text{otherwise,} \end{cases} \tag{5.10}$$

and the promised utilities start from $W_{i,0} = w_{(i)}$ and follow the dynamic

$$\mathrm{d}W_{i,t} = -\sum_{j \in \mathcal{I}} H_{ij,t}\, \mathrm{d}N_{j,t},$$

in which

$$H_{ij,t} = \begin{cases} X_{i,t}\hat{\beta}, & \text{for } i = j, \\ X_{j,t} W_{i,t-}(1-1/\tau), & \text{for } i \neq j. \end{cases} \tag{5.11}$$

In other words, under the rotating contract, the resource is allocated to only one of the agents at any time. The income rate of the agent holding the resource is R^a. Upon the arrival of an adverse event to this agent, all the resource is allocated to the next agents, and the agent who just lost the resource must wait for $n-1$ adverse arrivals, one at each of the other agents who holds the resource, before regaining the resource. The next result formally establishes that the rotating contract is an EIC contract.

Theorem 5.2. *Under condition* (5.9), *the contract according to Definition* 5.2 *is an EIC contract. The corresponding self-generating set of promised utilities consists of vectors that are cyclic permutations of the sequence* $\mathcal{W}_r$.

Both lower bounds (5.3) and (5.9) decrease in n. This implies that these contracts may not exist for some values of R^a, unless we increase the number of agents. Comparing the lower bounds (5.3) and (5.9), it is easy to verify that for any $n \geq 2$,

$$\frac{\rho}{1-\tau^{n-1}} + \lambda \leq \frac{n}{n-1}\lambda + \rho,$$

with the equality holding at $n = 2$, and the two quantities approach each other when n approaches infinity. Therefore, the lower bound (5.9) for the rotating contract is generally lower than the one in (5.3) for the simple EIC contract. This result implies that the rotating contract is applicable to a wider range of values for R^a. On the other hand, the rotating contract concentrates all the resource to one agent at a time, while the simple EIC contract almost always spreads the resource among all agents. We show this result formally in Section EC.5.4 of the online appendix.

Finally, we remark that under either the simple contract or rotating contract, the (UB) constraint is automatically satisfied for an upper bound $\bar{w}$ high enough, because both of them are "cash-free" EIC contracts. That is, the principal does not need to pay any cash to the agents under these contracts, and the promised utilities are endogenously determined and upper-bounded.

6. Further Discussion and Extensions

In this section, we first provide a sufficient condition under which it is optimal for the principal to require agents to exert constant effort. Then, we extend the model to consider asymmetric agents.

6.1. Sufficient Condition for Incentive Compatibility

In the previous sections, we illustrated the profit-maximizing contract and easy-to-implement contracts, both under incentive compatibility constraints. When requiring effort from agents, the principal needs to pay the corresponding rent, either in the form of current or future payments. Therefore, in general, it may be better for the principal not to enforce the incentive compatibility constraints all the time. Optimal contract design that allows shirking is a very challenging problem in general, even in a single-agent setting. Therefore, our paper focuses on maximizing the principal's utility among EIC contracts.

A few recent papers on dynamic contracting deal with the shirking issue in single-agent settings, and the corresponding optimal dynamic contract structures are often intricate. For example, Zhu (2013) studies single-agent optimal contract design that allows shirking under Brownian motion uncertainty. The optimal contract involves controlling a "sticky" Brownian motion. Cao et al. (2020) consider a similar dynamic contract design problem under Poisson uncertainties. To avoid the "sticky" issue, they introduce a fixed cost when the effort is switched back on. Most other papers on dynamic contract design provide sufficient conditions under which optimal contracts already satisfy incentive compatibility (see, e.g., Biais et al. 2010, Varas et al. 2020, Tian et al. 2021). In this subsection, we also provide a sufficient condition under which enforcing incentive compatibility does not affect optimality.

First, we need to generalize the contract space under consideration, such that the principal may allow the agent not to exert effort from time to time. Specifically, in this subsection we define contract Γ to include not only payments and allocation processes but also additional $\mathcal{F}$-adapted effort processes $\mathcal{E} = \{\mathcal{E}_{i,t}\}_{i\in\mathcal{I},t\geq 0}$. In particular, the effort process $\mathcal{E}_{i,t} \in \{\lambda, \bar{\lambda}\}$ directs agent i to exert effort ($\mathcal{E}_{i,t} = \lambda$) or not ($\mathcal{E}_{i,t} = \bar{\lambda}$) at time t.

The contract needs to ensure that the agents are willing to comply with the directed effort process. That is, contract $\Gamma = (\{\boldsymbol{L}_t\}_{t\geq 0}, \{\boldsymbol{X}_t\}_{t\geq 0}, \mathcal{E})$ needs to satisfy the following *obedience* constraint:

$$u_i(\Gamma, \mathcal{E}) \geq u_i(\Gamma, \{\mathcal{E}_{-i,t}, \lambda_{i,t}\}_{t\geq 0}), \ \forall i \in \mathcal{I}, \quad \text{and } \mathcal{F}\text{-adapted effort process } \{\lambda_{i,t}\}_{t\geq 0}. \tag{6.1}$$

The following result expresses the obedience constraint (6.1) in a recursive form, which generalizes Proposition 2.1.

Proposition 6.1. *Following Lemma 2.1, consider the agents' promised utilities $\boldsymbol{W}_t$ and processes $\boldsymbol{H}_t$ specified by contract Γ and effort process $\boldsymbol{\Lambda} = \mathcal{E}$ according to Γ. Constraint (6.1) is equivalent to*

$$H_{ii,t} \geq \beta X_{i,t} \text{ if and only if } \mathcal{E}_{i,t} = \lambda, \ \forall t \geq 0, \forall i \in \mathcal{I}. \tag{OB}$$

We now generalize the principal's utility defined in (2.6) to the following definition,

$$\tilde{U}(\Gamma) = \frac{R^p}{\rho} - \mathbb{E}^{\mathcal{E}}\left[\int_0^\infty e^{-\rho t} \sum_{i\in\mathcal{I}} (CX_{i,t}\, \mathrm{d}N_{i,t} + \mathrm{d}L_{i,t})\right], \tag{6.2}$$

and we consider the contract design problem as $\max_\Gamma \tilde{U}(\Gamma)$ among all contracts Γ that satisfy (FX), (LL), (PK), (OB), (IR), and (UB). Next, we propose a sufficient condition under which a contract with full-effort process $\mathcal{E} = \bar{\boldsymbol{\Lambda}}$ solves this contract design optimization problem. The verification result is based on the following lemma, which establishes an upper bound on $\tilde{U}(\Gamma)$, similar to lemma EC.6 in Tian et al. (2021).

Lemma 6.1. *Suppose that $\hat{F}(\boldsymbol{w})$ is an upper-bounded, concave, and subdifferentiable function with $\partial \hat{F}(\boldsymbol{w})/\partial w_i \geq -1$ for all $i \in \mathcal{I}$. Consider any contract $\Gamma = (\{\boldsymbol{L}_t\}_{t\geq 0}, \{\boldsymbol{X}_t\}_{t\geq 0}, \mathcal{E})$ that satisfies (FX), (LL), (PK), (OB), (IR), and (UB), and yields the promised utility process $\{\boldsymbol{W}_t\}_{t\geq 0}$ starting from $\boldsymbol{W}_0 = \boldsymbol{u}(\Gamma, \mathcal{E})$. Define a stochastic process $\{\Psi_t\}_{t\geq 0}$ as*

$$\begin{aligned}\Psi_t := &\sum_{i\in\mathcal{I}} \frac{\partial \hat{F}(\boldsymbol{W}_t)}{\partial W_{i,t}} \left(\rho W_{i,t-} + \sum_{j\in\mathcal{I}} \mathcal{E}_{j,t} H_{ij,t} - R^a X_{i,t} - b X_{i,t} \mathbb{1}_{\{\mathcal{E}_{i,t} = \bar{\lambda}\}}\right) \\ &+ \sum_{i\in\mathcal{I}} \mathcal{E}_{i,t} \hat{F}(\boldsymbol{W}_t - \boldsymbol{H}_{\cdot i,t}) - \left(\sum_{i\in\mathcal{I}} \mathcal{E}_{i,t} + \rho\right) \hat{F}(\boldsymbol{W}_t) \\ &+ \sum_{i\in\mathcal{I}} (R^p - \mathcal{E}_{i,t} C) X_{i,t}.\end{aligned}$$

If the process $\{\Psi_t\}_{t\geq 0}$ is nonpositive almost surely, then $\hat{F}(\boldsymbol{u}(\Gamma, \mathcal{E})) \geq \tilde{U}(\Gamma)$.

We can now show that a sufficient condition for the principal to continuously induce effort from all agents is

$$C \geq \beta \text{ and } (5.9). \tag{6.3}$$

Note that the condition $C \geq \beta$ implies that the benefit of reducing arrival (C) is higher than the information rent (β). Therefore, under this condition, it is socially efficient to induce effort from agents. This condition often appears in the dynamic moral hazard literature (see, e.g., Biais et al. 2010).

To apply Lemma 6.1, we define the function $\hat{F}(\boldsymbol{w})$ for any $\boldsymbol{w} \in \mathbb{R}_+^n$ as

$$\hat{F}(\boldsymbol{w}) = \min\left(\frac{R^p - \lambda C}{\rho}, \ \frac{R^a + R^p - \lambda C}{\rho} - \sum_{i\in\mathcal{I}} w_i\right). \tag{6.4}$$

Proposition 6.2. *Under condition (6.3), for the function $\hat{F}$ defined in (6.4) and for any Γ that satisfies (FX), (LL), (PK), (OB), (IR), and (UB), and yields an initial promised utility $\boldsymbol{w} \in \mathbb{R}_+^n$, we have*

$$\hat{F}(\boldsymbol{w}) \geq \tilde{U}(\Gamma). \tag{6.5}$$

Furthermore, there exists an EIC contract Γ such that

$$U(\Gamma) = \frac{R^p - \lambda C}{\rho} = \max_{\boldsymbol{w}\in\mathbb{R}_+^n} \hat{F}(\boldsymbol{w}),$$

From Section 5, we know that under condition (5.9), the principal's utility under the rotating contract is

$$U(\Gamma_s(\boldsymbol{w}^*; R^a)) = \frac{R^p - \lambda C}{\rho}, \ \forall \boldsymbol{w}^* \in \mathcal{U}_s.$$

which implies that the upper bound $\max_{\boldsymbol{w}\in\mathbb{R}_+^n} \hat{F}(\boldsymbol{w})$ is achievable by an effort-inducing contract. Furthermore, the condition $C \geq \beta$ in (6.3) allows us to apply Lemma 6.1 to prove (6.5). Proposition 6.2 thus suggests that (6.3) is a sufficient condition for focusing on incentive-compatible contracts.

Before closing this subsection, we remark that if the principal does not have to allocate all the resource to agents all the time, withholding some resource at certain points may indeed yield higher profit for the principal. Also, as mentioned in the beginning of this subsection, it may be better for the principal to allow some agents to shirk from time to time. These general contract structures may yield higher profits than any EIC contract. Finding such contracts is generally challenging, and therefore we leave it to future research.

6.2. Asymmetric Agents

In this subsection, we consider an extension of the dynamic contract design problem with asymmetric agents. Assume that when the resource is given to agent i, the total revenue rate R is split between the principal and agent i into R_i^p and R_i^a, respectively. Further assume that agent i's shirking benefit rate is b_i, and let $\beta_i = b_i/\Delta\lambda$. The theoretical results presented in the previous sections hold for this asymmetric setting. We omit trivial repetitions for clarity of exposition and highlight only two important points.

First, we generalize the iterative approach presented in Section 3 to characterize the achievable set of promised utilities. Specifically, we define a new operator T_{asym} as

$$
\begin{aligned}
[T_{asym}\phi](\boldsymbol{\alpha}) :=& \inf_{\boldsymbol{w},\boldsymbol{x},\boldsymbol{y}\in\mathbb{R}^n;\boldsymbol{H},\boldsymbol{Z}\in\mathbb{R}^{n\times n}} \boldsymbol{\alpha}^\top \boldsymbol{w} \\
&\text{s.t.} \quad H_{ii} \geq \beta_i x_i, \quad \forall i \in \mathcal{I}, \qquad (\mathrm{IC_{sA}}) \\
&\qquad y_i = \rho w_i - R_i^a x_i + \lambda \sum_{j\in\mathcal{I}} H_{ij}, \ \forall i \in \mathcal{I}, \qquad (\mathrm{PK}_{yA}) \\
&\qquad (\mathrm{FX_s}),\ (\mathrm{PK_Z}),\ (\mathrm{SG_w}),\ (\mathrm{SG_y}),\ (\mathrm{SG_Z}),\ (\mathrm{IR_s}),\ \text{and}\ (\mathrm{UB_s}).
\end{aligned}
$$

The operator T_{asym} differs from T in its constraints $(\mathrm{IC_{sA}})$ and (PK_{yA}) due to asymmetry. Then, we can apply T_{asym} in the iterative approach proposed in Section 3 to obtain the corresponding asymmetric achievable set and the optimal contract. Let $\mathcal{U}_{asym}$ denote the achievable set with asymmetric agents. We have the following result, which generalizes Theorem 3.1 to this setting.

Theorem 6.1. *Let $\mathcal{U}^0 = [0, \bar{w}]^n$, and define operator T^k_{asym} such that $T^k_{asym}\phi = T_{asym}(T^{k-1}_{asym}\phi)$ for all $k > 1$. We have*

$$\lim_{k\to\infty} \mathcal{G}(T^k_{asym}\phi_{\mathcal{U}^0}) = \mathcal{U}_{asym}.$$

The proof follows directly from that of Theorem 3.1 by substituting β and R^a with β_i and R_i^a for each agent i, respectively, and hence is omitted for brevity. Similarly, the linear program $[T_{asym}\phi_{\mathcal{U}_{asym}}](\boldsymbol{\alpha})$ also yields a boundary EIC contract. Moreover, we have the following result, which generalizes Theorem 4.1.

Theorem 6.2. *Consider the boundary EIC contract $\hat{\Gamma} \in \mathrm{bd}(\mathcal{U}_{asym})$ such that*

$$\sum_i u_i(\hat{\Gamma}) \leq \sum_i u_i, \ \forall \boldsymbol{u} \in \mathcal{U}_{asym}.$$

We have $U(\hat{\Gamma}) \geq U(\Gamma)$ for any EIC contract Γ.

Theorem 6.2 suggests that the EIC contract that maximizes the principal's profit is the one that starts with agents' promised utilities at a point on the boundary of the achievable set $\mathcal{U}_{asym}$ with a normal vector $(1, 1, \ldots, 1)$.

Figure 3 depicts two numerical examples, both with two asymmetric agents. In particular, Figure 3(a) depicts the boundary of the achievable set as a solid curve when the two agents differ in their reserved revenue rates; that is, $R_1^a > R_2^a$. Figure 3(b) depicts the case when the two agents differ in their shirking benefits; that is, $b_2 > b_1$. As a benchmark, panels (a) and (b) of Figure 3 also plot the boundary of the achievable sets for the symmetric-agent setting, $\mathcal{U}$, where $R^a = R_2^a$ and $b = b_1$, respectively, as dashed curves. We first observe that, for both examples, the achievable sets satisfy $\mathcal{U}_{asym} \subsetneq \mathcal{U}$. Moreover, the minimum promised utilities of the agent with either a higher reserved revenue rate or shirking benefit are strictly higher. Consequently, the principal's maximum profit decreases in both cases with asymmetric agents compared with the benchmark case. These results confirm the intuition that the principal's utility decreases if some agents gain more market power or have stronger incentive to shirk. Interestingly, when only one agent's revenue rate (or shirking benefit) increases, the entire achievable set shrinks and becomes a subset of the original set. That is, the benefit due to one agent's improvement spills over to the other agents. Intuitively, this makes sense because the other agent becomes relatively more desirable to the principal compared with the base case.

7. Concluding Remarks and Further Discussion

Motivated by the emergence of platforms that allocate online visits among independent suppliers, this paper studies a dynamic moral hazard model, in which a principal leverages resource allocation and payment strategies to motivate multiple symmetric agents to reduce the frequency of adverse events. We demonstrate, by construction, the existence of EIC contracts that always allocate all available resource to agents and ensure that agents always exert effort to reduce the arrival rate of adverse events. Our construction is based on an iterative approach that solves a sequence

Figure 3. Comparisons Between the Symmetric- and Asymmetric-Agent Cases

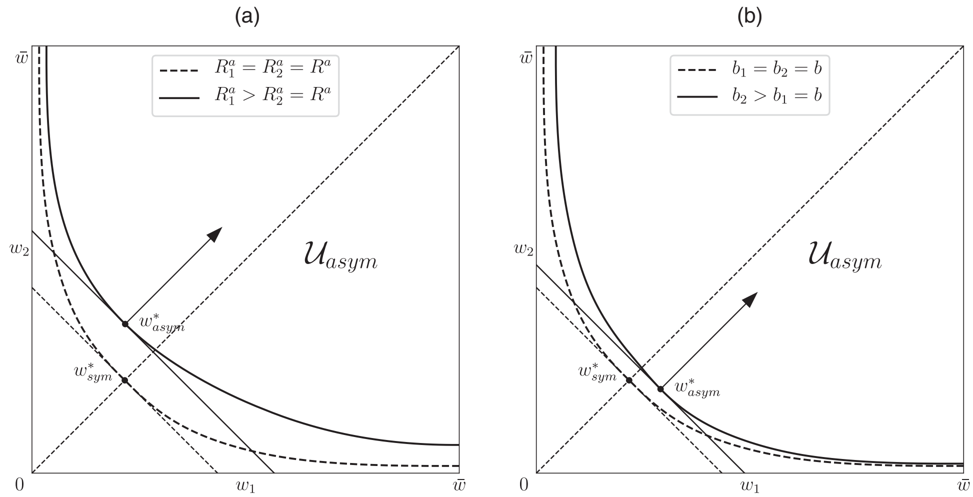

Note. (a) $R_1^a > R_2^a$; (b) $b_2 > b_1$.

of semi-infinite linear programs. We also specify the optimal EIC contract that maximizes the principal's utility. Nonetheless, from a practical perspective, the optimal contract can be cumbersome to characterize and implement. Therefore, we further provide a simple EIC contract and a rotating contract in closed form. These easy-to-implement EIC contracts possess desirable properties that make them relevant in practice. Our proposals thus provide prescriptive guidance for designing easy-to-implement EIC contracts that suit the need of practitioners.

Our analyses and results provide some interesting economic and managerial insights. For instance, in single-agent dynamic contracting settings similar to ours, one can perceive the decision of scaling firm size in Biais et al. (2010) or replacing a governor in Myerson (2015), analogously to our resource allocation decisions. In those settings, incentive-compatible contracts cannot achieve efficiency; that is, the firm has to be downsized and the governor has to be replaced within a finite time with probability one. This is because incentive compatibility requires the threat of reducing the agent's promised utility by a certain amount upon each arrival. In the multiagent setting, however, when we penalize an agent for an adverse event by reducing its promised utility and allocation, we can simultaneously increase other agents' allocations to maintain efficiency. For example, referring to Table 1, when Uber reduces the dispatching of riders to one driver, these riders should be dispatched among other drivers. Consequently, all other agents' promised utilities change due to one agent's adverse event. This flexibility allows us to design efficient contracts that induce agents to constantly exert effort. In fact, without such an option of potentially rewarding other agents upon the arrival at one agent, we show that EIC contracts no longer exist.

Notably, in a *discrete-time* version of the single-agent model where the agent's effort reduces the probability that an adverse arrival occurs in a period, EIC contracts may exist. For example, the principal pays the agent β at the end of each period when there is no arrival, and zero otherwise. Such a contract is naturally very costly to the principal and not desirable, as articulated in Holmström and Tirole (1997). If we take the usual approach of converging the discrete-time model to our continuous-time one by shrinking the time interval to zero while scaling other model parameters accordingly, β remains the same. It then becomes clear that paying the agent β per period results in an infinite payment per unit of time duration, which is not even feasible.

Furthermore, EIC contracts exist as long as there are two agents. One can put it as "a little bit of competition goes a long way." This result also implies that the principal can threaten agents with contract termination when there are more than two agents. In fact, we show that the optimal EIC contract that maximizes the principal's utility is the boundary contract starting from the point at which all agents' promised utilities are the same. This boundary contract terminates all but two agents within finite time.

We conclude this paper with some thoughts on potential future research directions. First, from a theoretical perspective, it may be interesting to investigate the optimal dynamic contracts in a setting where the principal is able to withhold some resource and/or agents are allowed to shirk. This appears a very challenging task, because even optimizing over the class of incentive-compatible (but not necessarily efficient) contracts involves optimality conditions in the form of a system of partial differential equations with delay. Second, in our model, we assume that agents do not collude. In particular, the EIC contracts in our setting leverage the competition among agents. It is interesting to investigate whether agents can suppress competition by collusion in multilateral dynamic moral hazard problems, and, if so, how to mitigat collusion in dynamic contract design. Third, in certain settings, rewarding one agent when there is an adverse event at another agent may create a perverse incentive for some agent to sabotage others. It would be interesting to investigate designs to mitigate this effect if it is a real concern. Finally, extending certain results of this paper into a discrete-time setting may be of interest for certain applications. In Section EC.8 of the online appendix, we provide the discrete-time versions of the simple and rotating contracts.

Acknowledgments

The authors gratefully thank the editor-in-chief and the area editor Professor John Birge, the anonymous associate editor, and reviewers for their insightful comments and suggestions. The authors contribute equally to this study and are ordered alphabetically.

Endnotes

[1] It is helpful to use the following example to explain this modeling choice. Consider adverse arrivals as low-quality or inappropriate content inadvertently posted by user-generated content providers (UGCs) on YouTube, Instagram, or even Coursera. The more visits that the platform allocates to this UGC, the greater the damage. Although we assume that the expected cost of an adverse event is proportional to the allocation, the arrival rate of adverse events is not. These modeling choices also capture adverse events in other settings. For example, outsourced parts without adequate quality control may trigger recalls, unmotivated representatives may cause customer complaints, and cutting corners in sanitation procedures may cause food poisoning episodes.

[2] This is a fair assumption in practice because the agents' liabilities are often hard to enforce for many practical reasons (Babich and Tang 2012). In theory, without this assumption, the principal may be able to effectively mitigate any misalignment of incentives, trivializing the contract design problem. Alternatively, one may assume a positive upper bound on the amount that the principal can charge the agent. This generalization does not affect our analysis and results in any fundamental way.

[3] From a practical point of view, if R^a is less than this lower bound, then the principal needs to pay a constant flow to make up the difference, as we illustrate later in Section 5.

[4] We acknowledge that this idea was proposed by one of our referees during the review process.

References

Abreu D, Pearce D, Stacchetti E (1990) Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica* 58(5): 1041–1063.

Alibaba Group (2015) Rules for enforcement action against non-compliance of transactions on alibaba.com. https://rule.alibaba.com/rule/detail/3310.htm.

Babich V, Tang CS (2012) Managing opportunistic supplier product adulteration: Deferred payments, inspection, and combined mechanisms. *Manufacturing Service Oper. Management* 14(2):301–314.

Babich V, Marinesi S, Tsoukalas G (2021) Does crowdfunding benefit entrepreneurs and venture capital investors? *Manufactuing Service Oper. Management* 23(2):508–524.

Baker GP (1992) Incentive contracts and performance measurement. *J. Political Econom.* 100(3):598–614.

Balseiro SR, Gurkan H, Sun P (2019) Multiagent mechanism design without money. *Oper. Res.* 67(5):1417–1436.

Bernard B, Frei C (2016) The folk theorem with imperfect public information in continuous time. *Theoret. Econom.* 11(2):411–453.

Biais B, Mariotti T, Rochet JC, Villeneuve S (2010) Large risks, limited liability, and dynamic moral hazard. *Econometrica* 78(1):73–118.

Cao P, Sun P, Tian F (2020) Punish underperformance with resting—Optimal dynamic contracts in the presence of switching cost. Working paper.

Caro F, Chintapalli P, Rajaram K, Tang CS (2018) Improving supplier compliance through joint and shared audits with collective penalty. *Manufacturing Service Oper. Management* 20(2):363–380.

Chen L, Lee HL (2017) Sourcing under supplier responsibility risk: The effects of certification, audit, and contingency payment. *Management Sci.* 63(9):2795–2812.

Chen M, Sun P, Xiao Y (2020) Optimal monitoring schedule in dynamic contracts. *Oper. Res.* 68(5):1285–1314.

Cho SH, Fang X, Tayur S, Xu Y (2019) Combating child labor: Incentives and information disclosure in global supply chains. *Manufacturing Service Oper. Management* 21(3):692–711.

Dai T, Tang SC (2020) Amazon has a trust problem. *Barron's* (August 7), https://www.barrons.com/articles/amazon-has-a-trust-problem-51596825474.

Edelstein S (2018) UberEats drivers are eating customers' food thanks to 'Loophole', report claims. *The Drive* (May 17), https://www.thedrive.com/news/25263/ubereats-drivers-are-eating-customers-food-thanks-to-loophole-report-claims.

Feng L (2019) Didi to relaunch the controversial hitch service? *Equal Ocean* (May 5), https://equalocean.com/analysis/201905051988/.

Fudenberg D, Levine D, Maskin E (1994) The folk theorem with imperfect public information. *Econometrica* 62(5):997–1039.

Gibbons R, Murphy KJ (1992) Optimal incentive contracts in the presence of career concerns: Theory and evidence. *J. Political Econom.* 100(3):468–505.

Goeree JK, Kushnir A (2016) Reduced form implementation for environments with value interdependencies. *Games Econom. Behav.* 99:250–256.

Goeree JK, Kushnir A (2020) A geometric approach to mechanism design. Working paper.

Green B, Taylor CR (2016) Breakthroughs, deadlines, and self-reported progress: Contracting for multistage projects. *Amer. Econom. Rev.* 106(12):3660–3699.

Grossman SJ, Hart OD (1983) An analysis of the principal-agent problem. *Econometrica* 51(1):7–45.

Holmström B (1979) Moral hazard and observability. *Bell J. Econom.* 10(1):74–91.

Holmström B (1999) Managerial incentive problems: A dynamic perspective. *Rev. Econom. Stud.* 66(1):169–182.

Holmström B, Tirole J (1997) Financial intermediation, loanable funds, and the real sector. *Quart. J. Econom.* 112(3):663–691.

Innes RD (1990) Limited liability and incentive contracting with ex-ante action choices. *J. Econom. Theory* 52(1):45–67.

Marantz A (2018) Reddit and the struggle to detoxify the Internet. *New Yorker* (March 19), https://www.newyorker.com/magazine/2018/03/19/reddit-and-the-struggle-to-detoxify-the-internet.

Mauldin W (2019) Alibaba's Taobao.com still a 'notorious market' despite improvements, U.S. says. *Wall Street Journal* (April 25), https://www.wsj.com/articles/alibabas-taobao-com-still-a-notorious-market-despite-improvements-u-s-says-11556210101.

Myerson RB (2015) Moral hazard in high office and the dynamics of aristocracy. *Econometrica* 83(6):2083–2126.

Plambeck EL, Taylor TA (2016) Supplier evasion of a buyer's audit: Implications for motivating supplier social and environmental responsibility. *Manufacturing Service Oper. Management* 18(2): 184–197.

Prendergast C (2002) The tenuous trade-off between risk and incentives. *J. Political Econom.* 110(5):1071–1102.

Rogerson WP (1985) Repeated moral hazard. *Econometrica* 53(1):69–76.

Rui H, Lai G (2015) Sourcing with deferred payment and inspection under supplier product adulteration risk. *Production Oper. Management* 24(6):934–946.

Sannikov Y (2008) A continuous-time version of the principal-agent problem. *Rev. Econom. Stud.* 75(3):957–984.

Shan Y (2017) Optimal contracts for research agents. *RAND J. Econom.* 48(1):94–124.

Spear SE, Srivastava S (1987) On repeated moral hazard with discounting. *Rev. Econom. Stud.* 54(4):599–617.

Sun P, Tian F (2018) Optimal contract to induce continued effort. *Management Sci.* 64(9):4193–4217.

Tang CS, Yang SA, Wu J (2018) Sourcing from suppliers with financial constraints and performance risk. *Manufacturing Service Oper. Management* 20(1):70–84.

Tian F, Sun P, Duenyas I (2021) Optimal contract for machine repair and maintenance. *Oper. Res.* 69(3):916–949.

Varas F, Marinovic I, Skrzypacz A (2020) Random inspections and periodic reviews: Optimal dynamic monitoring. *Rev. Econ. Stud.* 87(6):2893–2937.

Wang S, Sun P, de Véricourt F (2016) Inducing environmental disclosures: A dynamic mechanism design approach. *Oper. Res.* 64(2):371–389.

Zhu JY (2013) Optimal contracts with shirking. *Rev. Econ. Stud.* 80(2):812–839.

Yong Liang is an associate professor at the Research Center for Contemporary Management at the School of Economics and Management, Tsinghua University. His research interests include optimization theory and applications in operations management, supply chain management, and dynamic contract design.

Peng Sun is the J. B. Fuqua Professor of Decision Sciences at the Fuqua School of Business, Duke University. His research interests include optimization under uncertainty and optimal design of contracts and mechanisms in dynamic environments.

Runyu Tang is an assistant professor at the School of Management, Xi'an Jiaotong University. His research interests mainly lie in the area of data-driven optimization, game theory, supply chain management under emerging technologies, and dynamic mechanism design.

Chong Zhang is an assistant professor of supply chain management in the Department of Management at Tilburg University. She is broadly interested in marketing-operations interfaces, platform operations, information economics, game theory, and dynamic mechanism design. Her current research integrates both theoretical modeling and data analysis to study complex managerial problems with information asymmetry on digital platforms.