This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

Toward a Theory of Systems Engineering Processes: A Principal–Agent Model of a One-Shot, Shallow Process

Salar Safarkhani, Ilias Bilionis, and Jitesh H. Panchal

Abstract—**Systems engineering processes (SEPs) coordinate the effort of different individuals to generate a product satisfying certain requirements. As the involved engineers are self-interested agents, the goals at different levels of the systems engineering hierarchy may deviate from the system-level goals, which may cause budget and schedule overruns. Therefore, there is a need of a systems engineering theory that accounts for the human behavior in systems design. As experience in the physical sciences shows, a lot of knowledge can be generated by studying simple hypothetical scenarios, which nevertheless retain some aspects of the original problem. To this end, the objective of this article is to study the simplest conceivable SEP, a principalagent model of a one-shot, shallow SEP. We assume that the systems engineer (SE) maximizes the expected utility of the system, while the subsystem engineers (sSE) seek to maximize their expected utilities. Furthermore, the SE is unable to monitor the effort of the sSE and may not have complete information about their types. However, the SE can incentivize the sSE by proposing specific contracts. To obtain an optimal incentive, we pose and solve numerically a bilevel optimization problem. Through extensive simulations, we study the optimal incentives arising from different system-level value functions under various combinations of effort costs, problem-solving skills, and task complexities. Our numerical examples show that, the passed-down requirements to the agents increase as the task complexity and uncertainty grow and they decrease with increasing the agents' costs.**

Index Terms—**Bilevel programming problem, complex systems, contract theory, expected utility, game theory, mechanism design, optimal incentives, principal–agent model, systems engineering theory, systems science.**

I. Introduction

COST and schedule overruns plague the majority of large systems engineering projects across multiple industry sectors including power [1], defense [2], and space [3]. As design mistakes are more expensive to correct during the production and operation phases, the design phase of the systems engineering process (SEP) has the largest potential impact on cost and schedule overruns. Collopy *et al.* [4] argued that requirements engineering (RE), which is a fundamental part of the design phase, is a major source of inefficiencies in systems engineering. In response, they developed value-driven design (VDD) [5], a systems design approach that starts with the identification of a system-level value function and guides the systems engineer (SE) to construct subsystem value functions that are aligned with the system goals. According to VDD, the subsystem engineers (sSE) and contractors should maximize the objective functions passed down by the SE instead of trying to meet requirements.

RE and VDD make the assumption that the goals of the human agents involved in the SEP are aligned with the SE goals. In particular, RE assumes that, agents attempt to maximize the probability of meeting the requirements, while VDD assumes that they will maximize the objective functions supplied by the SE. However, this assumption ignores the possibility that the design agents, as all humans, may have personal agendas that not necessarily aligned with the system-level goals.

Contrary to RE and VDD, it is more plausible that the design agents seek to maximize their own objectives. Indeed, there is experimental evidence that the quality of the outcome of a design task is strongly affected by the reward anticipated by the agent [6]–[8]. In other words, the agent decides how much effort and resources to devote to a design task after taking into account the potential reward. In the field, the reward could be explicitly implemented as an annual performance-based bonus, or, as it is the case most often, it could be implicitly encoded in expectations about job security, promotion, professional reputation, etc. To capture the human aspect in SEPs, one possible way is to follow a game-theoretic approach [9], [10]. Most generally, the SEP can be modeled as a dynamical hierarchical network game with incomplete information. Each layer of the hierarchy represents interactions among the SE and some sSEs, or the sSEs and other engineers or contractors. With the term "principal," we refer to any individual delegating a task, while we reserve the term "agent" for the individual carrying out the task. Note that an agent may simultaneously be the principal in a set of interactions down the network. For example, the sSE is the agent when considering their interaction with the SE (the principal), but the principal when considering their interaction with a contractor (the agent). At each time step, the principals pass down delegated tasks along with incentives, the agents choose the effort levels that maximize their expected utility, perform the task, and return the outcome to the principals.

Manuscript received March 29, 2019; revised September 20, 2019 and December 26, 2019; accepted December 26, 2019. This work was supported by the National Science Foundation under Grant 1728165. *(Corresponding author: Ilias Bilionis.)*

The authors are with the School of Mechanical Engineering, Purdue University, West Lafayette, IN 47907-2088 USA (e-mail: ssafarkh@purdue.edu; ibilion@purdue.edu; panchal@purdue.edu).

Digital Object Identifier 10.1109/JSYST.2020.2964668

The iterative and hierarchical nature of real SEPs makes them extremely difficult to model in their full generality. Given that our aim is to develop a theory of SEPs, we start from the simplest possible version of a SEP which retains, nevertheless, some of the important elements of the real process. Specifically, the objective of this article is to develop and analyze a principal–agent model of a one-shot, shallow SEP. The SEP is "one-shot" in the sense that decisions are made in one iteration and they are final. The term "shallow" refers to a one-layer-deep SEP hierarchy, i.e., only the SE (principal) and the sSEs (agents) are involved. The agents maximize their expected utility given the incentives provided by the principal, and the principal selects the incentive structure that maximizes the expected utility of the system. We pose this mechanism design problem [11] as a bilevel optimization problem and we solve it numerically.

A key component of our SEP model is the quality function of an agent. The quality function is a stochastic process that models the principal's beliefs about the outcome of the delegated design task given that the agent devotes a certain amount of effort. The quality function is affected by what the principal believes about the task complexity and the problem solving skills of the agent. Following our work [12], we model the design task as a maximization problem where the agent seeks the optimal solution. The principal expresses their prior beliefs about the task complexity by modeling the objective function as a random draw from a Gaussian process prior with a suitably selected covariance function.

As we showed in [12], conditioned on knowing the task complexity and the agent type, the quality function is well approximated by an increasing, concave function of effort with additive Gaussian noise. However, we will use a linear approximation for the quality function.

We study numerically two different scenarios. The first scenario assumes that the SE knows the agent types and the task complexity, but they do not observe the agent's effort. This situation is known in game theory as a *moral hazard* problem [13]. The most common way to solve a moral hazard problem is to use the first-order approach (FOA) [14]. In the FOA, the incentive compatibility constraint of the agent is replaced by its first-order necessary condition. However, the FOA depends on the convexity of the distribution function in effort which is not valid in our case. There have been several attempts to solve the principal–agent model, where the requirements of the FOA may fail, nonetheless they must still satisfy the monotone likelihood ratio property [15].

In the second scenario, we study the case of moral hazard with simultaneous *adverse selection* [16], i.e., the SE observes neither the effort nor the type of agents nor the task complexity. This is a Bayesian game with incomplete information.. In this case, the SE experiences additional loss in their expected utility, because the sSEs' can pretend to have different types. The revelation principle [17] guarantees that it suffices to search for the optimal mechanism within the set of incentive compatible mechanisms, i.e., within the set of mechanisms in which the sSEs are telling the truth about their types and technology maturity. In this article, we solve the optimization problem in the principal–agent model, numerically with making no assumptions about the quality function.

This article is organized as follows. In Section II, we will derive the mathematical model of the SEP and we will study the type-independent and type-dependent optimal contracts. We will also introduce the value and utility functions. In Section III, we perform an exhaustive numerical study and show the solutions for several case studies. Finally, we conclude in Section IV.

II. Modeling a One-Shot, Shallow SEP

A. Basic Definitions and Notation

As mentioned in the introduction, we develop a model of a one-shot (the game evolves in one iteration and the decisions are final), shallow (one-layer-deep hierarchy) SEP. The SE has decomposed the system into N subsystems and assigned an sSE to each one of them. We use $i = 1, \dots, N$ to label each subsystem. From now on, we refer to the SE as the principal and the sSEs as the agents. The principal delegates tasks to the agents along with incentives. The agents choose how much effort to devote on their task by maximizing their expected utility. The principal, anticipates this reaction and selects the incentives that maximize the system-level expected utility.

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space where, Ω is the sample space, $\mathcal{F}$ is a σ-algebra, and $\mathbb{P}$ is the probability measure. With $\omega \in \Omega$ we refer to the random state of nature. We use upper case letters for random variables (r.v.), bold upper case letters for their range, and lower case letters for their possible values. For example, the type of agent i is a r.v. Θ_i taking M_i discrete values θ_i in the set $\boldsymbol{\Theta}_i \equiv \Theta_i(\Omega) = \{1, \dots, M_i\}$. Collectively, we denote all types with the N-dimensional tuple $\Theta = (\Theta_1, \dots, \Theta_N)$ and we reserve Θ_{-i} to refer to the $(N-1)$-dimensional tuple containing all elements of Θ except Θ_i. This notation carries to any N-dimensional tuple. For example, θ and θ_{-i} are the type values for all agents and all agents except i, respectively. The range of Θ is $\boldsymbol{\Theta} = \times_{i=1}^{N} \boldsymbol{\Theta}_i$.

The principal believes that the agents types vary independently, i.e., they assign a probability mass function (p.m.f.) on Θ that factorizes over types as follows:

$$\mathbb{P}[\Theta = \theta] = \prod_{i=1}^{N} \mathbb{P}[\Theta_i = \theta_i] = \prod_{i=1}^{N} p_{i\theta_i} \tag{1}$$

for all θ in $\boldsymbol{\Theta}$, where $p_{ik} \geq 0$ is the probability that agent i has type k, for k in $\boldsymbol{\Theta}_i$. Of course, we must have $\sum_{k=1}^{M_j} p_{ik} = 1$, for all $i = 1, \dots, N$.

Each agent knows their type, but their state of knowledge about all other agents is the same as the principal's. That is, if agent i is of type $\Theta_i = \theta_i$, then their state of knowledge about everyone else is captured by the p.m.f.

$$\mathbb{P}[\Theta_{-i} = \theta_{-i} | \Theta_i = \theta_i] = \frac{\mathbb{P}[\Theta = \theta]}{\mathbb{P}[\Theta_i = \theta_i]} = \prod_{j \neq i} p_{j\theta_j}. \tag{2}$$

Agent i chooses a normalized effort level $e_i \in [0, 1]$ for his delegated task. We assume that this normalized effort is the percentage of an agent's maximum available effort. The units of the normalized effort depend on the nature of the agent's subsystem. If the principal and the agent are both part of same organization then the effort can be the time that the agent dedicates to the delegated task in a particular period of time, e.g., in a fiscal year. On the other hand, if the agent is a contractor, then the effort can be the percentage of the available yearly budget that the contractor spends on the assigned task. We represent the monetary cost of the ith agent's effort with the random process $C_i(e_i)$. In economic terms, $C_i(e_i)$ is the opportunity cost, i.e., the payoff of the best alternative project in agent could devote their effort. In general, we know that the process $C_i(e_i)$ should be an increasing function of the effort e_i. For simplicity, we assume that the cost of effort of the agents is quadratic

$$C_i(e_i) := c_{i\Theta_i} e_i^2 \tag{3}$$

with a type-dependent coefficient $c_{ik} > 0$ for all k in $\mathbf{\Theta}_i$.

The quality function of the ith agent is a real-valued random $Q_i(e_i) := Q_i(e_i)$ process paremeterized by the effort e_i. The quality function models everybody's beliefs about the design capabilities of agent i. The interpretation of the quality function is as follows. If agent i devotes to the task an effort of level e_i, then they produce a random outcome of quality $Q_i(e_i)$. In our previous work [12], we created a stochastic model for the quality function of a designer where we explicitly captured its dependence on the problem-solving skills of the designer and on the task complexity. In that work, we showed that $Q_i(e_i)$ has increasing and concave sample paths, that its mean function is increasing concave, and the standard deviation is decreasing with effort, albeit mildly, it is independent of the problem-solving skills of the designer, and it only increases mildly with increasing task complexity. Examining the spectral decomposition of the process for various cases, we observed that it can be well approximated by

$$Q_i(e_i) = q^0_{i\Theta_i}(e_i) + \sigma_{i\Theta_i}\Xi_i \tag{4}$$

where for k in $\mathbf{\Theta}_i$, $q^0_{ik}(e_i)$ is an increasing, concave, type-dependent mean quality function, $\sigma_{ik} > 0$ is a type-dependent standard deviation parameter capturing the aleatory uncertainty of the design process, and Ξ_i is a standard normal r.v. If we further assume that the time window for design is relatively small, then the $q^0_{ik}(e_i)$ term can be approximated as a linear function. Therefore, we will assume that the quality function is

$$Q_i(e_i) = \kappa_{i\Theta_i} e_i + \sigma_{i\Theta_i}\Xi_i \tag{5}$$

where κ is inversely proportional to the complexity of the problem. For instance, a large κ corresponds to a low-complexity task while a small κ corresponds to a high-complexity task. The standard deviation parameter σ captures the inherent uncertainty of the design process and depends on the maturity of the underlying technology. In summary, an agent's type is characterized by the triplet cost-complexity-uncertainty.

From the perspective of the principal, the r.v.'s Ξ_i are independent of the agents' types Θ_i as they represent the uncertain state of nature. A stronger assumption that we employ is that these Ξ_i's are also independent to each other. This assumption is strong because it essentially means that the qualities of the various subsystems are decoupled. Under these independence assumptions, the state of knowledge of the principal is captured by the following probability measure:

$$\mathbb{P}\left[\Theta = \theta, \Xi \in \times_{i=1}^N \mathbf{B}_i\right] = \prod_{i=1}^N \left[p_{i\theta_i}\int_{\mathbf{B}_i}\phi(\xi_i)d\xi_i\right] \tag{6}$$

for all $\theta \in \mathbf{\Theta}$ and all Borel-measurable $\mathbf{B}_i \subset \mathbb{R}$. Assuming that all these are common knowledge, the state of knowledge of agent i after they observe their type θ_i (but before they observe Ξ_i) is

$$\begin{aligned}
&\mathbb{P}\left[\Theta_{-i} = \theta_{-i}, \xi \in \times_{i=1}^N \mathbf{B}_i | \Theta_i = \theta_i\right] \\
&= \frac{\mathbb{P}\left[\Theta = \theta, \xi \in \times_{i=1}^N \mathbf{B}_i\right]}{\mathbb{P}[\Theta_i = \theta_i]} \\
&= \mathbb{P}[\Theta_{-i} = \theta_{-i}|\Theta_i = \theta_i]\prod_{i=1}^N \left[\int_{\mathbf{B}_i}\phi(\xi_i)d\xi_i\right].
\end{aligned} \tag{7}$$

Finally, we use $\mathbb{E}[\cdot]$ to denote the expectation of any quantity over the state of knowledge of the principal as characterized by the probability measure of (6). That is, the expectation of any function $f(\Theta, \Xi)$ of the agent types Θ and the state of nature Ξ is

$$\mathbb{E}[f(\Theta, \Xi)] = \sum_{\theta\in\mathbf{\Theta}}\int_{\mathbb{R}^N} f(\theta, \xi)\prod_{i=1}^N \left[p_{i\theta_i}\phi(\xi_i)\right]d\xi. \tag{8}$$

Similarly, we use the notation $\mathbb{E}_{i\theta_i}[\cdot]$ to denote the conditional expectation over the state of knowledge of an agent i who knows that their type is $\Theta_i = \theta_i$. This is the expectation $\mathbb{E}[\cdot|\Theta_i = \theta_i]$ with respect to the probability measure of (2) and we have

$$\mathbb{E}_{ik}[f(\Theta, \Xi)] = \sum_{\theta_{-i}\in\mathbf{\Theta}_{-i}}\int_{\mathbb{R}^N} f(\theta_i, \theta_{-i}, \xi)\frac{\prod_{j=1}^N\left[p_{j\theta_j}\phi(\xi_j)\right]}{p_{i\theta_i}}d\xi. \tag{9}$$

B. Type-Independent Optimal Contracts

We start by considering the case where the principal offers a single take-it-or-leave-it contract independent of the agent type. This is the situation usually encountered in contractual relationships between the SE and the sSEs within the same organization. The principal offers the contract and the agent decides whether or not to accept it. If the agent accepts, then they select their level of effort by maximizing their expected utility, they work on their design task, they return the outcome quality back to the principal, and they receive their reward. We show a schematic view of this type of contracts in Fig. 1(a). A contract is a monetary *transfer function* $t_i : \mathbb{R} \to \mathbb{R}$ that specifies the agent's *compensation* $t_i(q_i)$ contingent on the quality level q_i. Therefore, the payoff of the ith agent is the random process

$$\Pi_i(e_i) = t_i\left(Q_i\left(e_i\right)\right) - C_i(e_i). \tag{10}$$

We assume that the agent knows their type, but they choose the optimal effort level ex-ante, i.e., they choose the effort level before seeing the state of the nature Ξ_i. Denoting their monetary utility function by $U_i(\pi_i) = u_{i\Theta_i}(\pi_i)$, the ith agent selects an

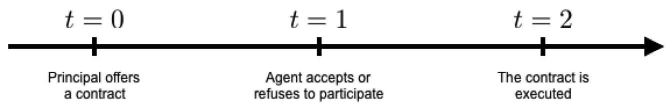

(a)

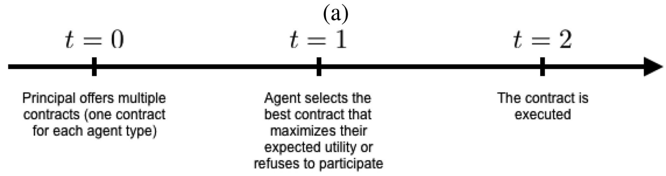

(b)

Fig. 1. (a) Timing of the contract for type-independent contracts. (b) Timing of the contract for type-dependent contracts.

effort level by solving

$$e^*_{i\theta_i} = \arg\max_{e_i\in[0,1]} \mathbb{E}_{i\theta_i}\left[U_i\left(\Pi_i(e_i)\right)\right]. \tag{11}$$

Let Q^*_i be the r.v. representing the quality function that the principal should expect from agent i if they act optimally, i.e.,

$$Q^*_i = Q_i(e^*_{i\Theta_i}). \tag{12}$$

Then, the system level value is a r.v. of the form

$$V = v(Q^*) \tag{13}$$

where $v : \mathbb{R}^N \to \mathbb{R}$ is a function of the subsystem outcomes Q^*. We introduce the form of the value function, $v(q)$, in Section III. Note that, even though in this article the r.v. V is assumed to be just a function of Q^*, in reality it may also depend on the random state of nature, e.g., future prices, demand for the system services. Consideration of the latter is problem-dependent and beyond the scope of this article.

Given the system value V and taking into account the transfers to the agents, the system-level payoff is the r.v.

$$\Pi_0 = V - \sum_{i=1}^{N} t_i(Q^*_i). \tag{14}$$

If the monetary utility of the principal is $u_0(\pi_0)$, then they should select the transfer functions $t(\cdot) = (t_1(\cdot),\ldots,t_N(\cdot))$ by solving

$$t^*(\cdot) = \arg\max_{t(\cdot)} \mathbb{E}\left[u_0\left(\Pi_0\right)\right]. \tag{15}$$

However, guarantee that they want to participate in the SEP, the expected utility of the sSEs must be greater than the expected utility they would enjoy if they participated in another project. Therefore, the SE must solve (15) subject to the *participation constraints*

$$\mathbb{E}_{i\theta_i}\left[U_i\left(\Pi_i\right)\right] \geq \bar{u}_{i\theta_i} \tag{16}$$

for all possible values of θ_i, and all $i = 1,\ldots,N$, where $\bar{u}_{i\theta_i}$ is known as the *reservation utility* of agent i.

C. Type-Dependent Optimal Contracts

By offering a single transfer function, the principal is unable to differentiate between the various agent types when adverse selection is an issue. That is, all agent types, independently of their cost, complexity, and uncertainty attributes, are offered exactly the same transfer function. In other words, with a single transfer function the principal is actually targeting the average agent. This necessarily leads to inefficiencies stemming from problems such as paying an agent involved in a low-complexity task more than a same cost and uncertainty agent involved in a high-complexity task.

The principal can gain in efficiency by offering different transfer functions (if any exist) that target specific agent types. For example, the principal could offer a transfer function that is suitable for cost-efficient, low-complexity, low-uncertainty agents, and one for cost-inefficient agents, low-complexity, low-uncertainty, etc., for any other combination that is supported by the principal's prior knowledge about the types of the agent population. To implement this strategy, the principal can employ the following extension to the mechanism of Section II-B. Prior to initiating work, the agents announce their types to the principal and they receive a contract that matches the announced type. In Fig. 1(b), we show how this type of contract evolves in time. Let us formulate this idea mathematically. The ith agent announces a type θ'_i in $\mathbf{\Theta}_i$ (not necessarily the same as their true type θ_i), and they receive the associated, type-specific, transfer function $t_{i\theta'_i}(\cdot)$. The payoff to agent i is now

$$\Pi_i(e_i,\theta'_i) = t_{i\theta'_i}(Q_i(e_i)) - C_i(e_i) \tag{17}$$

where all other quantities are like before. Given the announcement of a type θ'_i, the rational thing to do for agent i is to select a level of $e^*_i(\theta_i,\theta'_i)$ by maximizing their expected utility, i.e., by solving

$$e^*_{i\theta_i\theta'_i} = \arg\max_{e_i\in[0,1]}\mathbb{E}_{i\theta_i}[U_i(\Pi_i(e_i,\theta'_i))]. \tag{18}$$

Of course, the announcement of θ'_i is also a matter of choice and a rational agent should select also by maximizing their expected utility. The obvious issue here is that agents can lie about their type. For example, a cost-efficient agent (agent with low cost of effort) may pretend to be a cost-inefficient agent (agent with high cost of effort). Fortunately, the revelation principle [17] comes to the rescue and simplifies the situation. It guarantees that, among the optimal mechanisms, there is one that is incentive compatible. Thus, it will be sufficient if the principal constraints their contracts to over truth-telling mechanisms. Mathematically, to enforce truth-telling, the SE must satisfy the *incentive compatibility* constraints

$$\mathbb{E}_{i\theta_i}[U_i(\Pi_i(e^*_{i\Theta_i\theta_i},\theta_i))] \geq \mathbb{E}_{i\theta_i}[U_i(\Pi_i(e^*_{i\Theta_i\theta'_i},\theta'_i))] \tag{19}$$

for all $\theta_i \neq \theta'_i$ in $\mathbf{\Theta}_i$. Equation (19) expresses mathematically that "the expected payoff of agent i when they are telling the truth is always greater than or equal to the expected payoff they would enjoy if they lied."

Similar to the developments of Section II-B, the quality that the SE expects to receive is

$$Q^*_i = Q_i(e^*_{i\Theta_i\Theta_i}) \tag{20}$$

where we use the fact that the mechanism is incentive-compatible. The payoff of the SE becomes

$$\Pi_0 = V - \sum_{i=1}^{N} t_{i\Theta_i}(Q_i^*). \tag{21}$$

Therefore, to select the optimal transfer functions, the SE must solve

$$\max_{t(\cdot,\cdot)} \mathbb{E}\left[u_0(\Pi_0)\right] \tag{22}$$

subject to the incentive compatibility constraints of (19), and the participation constraints

$$\mathbb{E}_{i\theta_i}[U_i(\Pi_i(e^*_{i\Theta_i\Theta_i}, \Theta_i))] \geq \bar{u}_{i\theta_i} \tag{23}$$

for all $\theta_i \in \boldsymbol{\Theta}_i$, where we also assume that the incentive compatibility constrains hold.

D. Parameterization of the Transfer Functions

Transfer functions must be practically implementable. That is, they must be easily understood by the agent when expressed in the form of a contract. To be easily implementable, transfer functions should be easy to convey in the form of a table. To achieve this, we restrict our attention to functions that are made out of constants, step functions, linear functions, or combinations of these.

Despite the fact that including such functions would likely enhance the principal's payoff, we exclude transfer functions that encode penalties for poor agent performance, i.e., transfer functions that can take negative values. First, contracts with penalties may not be implementable if the principal and the agent reside within the same organization. Second, even when the agent is an external contractor penalties are not commonly encountered in practice. In particular, if the SE is a sensitive government office, e.g., the department of defense, national security may dictate that the contractors should be protected from bankruptcy. Third, we do not expect our theory to be empirically valid when penalties are included since, according to prospect theory [18], humans perceive losses differently. They are risk-seeking when the reference point starts at a loss and risk averse (RA) when the reference point starts at a gain.

To overcome these issues, we restrict our attention to transfer functions that include three simple additive terms: a constant term representing a participation payment, i.e., a payment received for accepting to be part of the project; a constant payment that is activated when a requirement is met; and a linear increasing part activated after meeting the requirement. The role of the latter two part is to incentivize the agent to meet and exceed the requirements.

We now describe this parameterization mathematically. The transfer function associated with type k in $\boldsymbol{\Theta_i}$ of agent i is parameterized by

$$\begin{aligned} t_{ik}\left(q_i\right) &= a_{ik,0} + a_{ik,1}\,\mathrm{H}\left(q_i - a_{ik,2}\right) \\ &\quad + a_{ik,3}\left(q_i - a_{ik,2}\right)\mathrm{H}\left(q_i - a_{ik,2}\right) \end{aligned} \tag{24}$$

where H is the Heaviside function ($H(x) = 1$ if $x \geq 0$ and 0 otherwise), and all the parameters $a_{ik,0}, \ldots, a_{ik,3}$ are nonnegative. In (24), $a_{ik,0}$ is the participation reward, $a_{ik,1}$ is the award for exceeding the passed-down requirement, $a_{ik,2}$ is the passed-down requirement, and $a_{ik,3}$ the payoff per unit quality exceeding the passed-down requirement. We will call these form of transfer functions the "requirement based plus incentive" (RPI) transfer function. In the case $a_{ik,3} = 0$, we call it the "requirement based" (RB) transfer function. At this point, it is worth mentioning that the passed-down requirement $a_{ik,2}$ is not necessarily the same as the true system requirement r_i, see our results in Section III. As we have shown in earlier work [10], the optimal passed-down requirement differs from the true system requirement. For notational convenience, we denote by $\mathbf{a}_{ik} \in \mathbb{R}_+^4$ ($\mathbb{R} = \{x \in \mathbb{R} : x \geq 0\}$) the transfer parameters pertaining to agent i of type $k \in \boldsymbol{\Theta}_i$, i.e.,

$$\mathbf{a}_{ik} = \left(a_{ik,0}, \ldots, a_{ik,3}\right). \tag{25}$$

Similarly, with $\mathbf{a}_i \in \mathbb{R}_+^{4M_i}$ we denote the transfer parameters pertaining to agent i for all types, i.e.,

$$\mathbf{a}_i = \left(\mathbf{a}_{i1}, \ldots, \mathbf{a}_{iM_i}\right) \tag{26}$$

and with $\mathbf{a} \in \mathbb{R}_+^{4\sum_{i=1}^{N} M_i}$ all the transfer parameters collectively, i.e.,

$$\mathbf{a} = \left(\mathbf{a}_1, \ldots, \mathbf{a}_N\right). \tag{27}$$

E. Numerical Solution of the Optimal Contract Problem

The optimal contract problem is an intractable bilevel, nonlinear programming problem. In particular, the SE's problem is for the case of type-dependent contracts is to maximize the expected system-level utility over the class of implementable contracts, i.e.,

$$\max_{\mathbf{a}} \mathbb{E}\left[u_0(\Pi_0)\right] \tag{28}$$

subject to

1) *Contract Implementability Constraints*

$$a_{ik,j} \geq 0 \tag{29}$$

for all $i = 1, \ldots, N, k = 1, \ldots, M_i, j = 0, \ldots, 3$;

2) *Individual Rationality Constraints*

$$e^*_{ikl} = \arg\max_{e_i \in [0,1]} \mathbb{E}_{ik}[U_i(\Pi_i(e_i, l))] \tag{30}$$

for all $i = 1, \ldots, N, k = 1, \ldots, M_i, l = 1, \ldots, M_i$;

3) *Participation Constraints*

$$\mathbb{E}_{ik}[U_i(\Pi_i(e^*_{\text{ikk}}, k))] \geq \bar{u}_{ik} \tag{31}$$

for all $i = 1, \ldots, N, k = 1, \ldots, M_i$; and

4) *Incentive Compatibility Constraints*

$$\mathbb{E}_{ik}[U_i(\Pi_i(e^*_{\text{ikk}}, k))] \geq \mathbb{E}_{ik}[U_i(\Pi_i(e^*_{ikl}, l))] \tag{32}$$

for all $i = 1, \ldots, N$ and $k \neq l$ in $\{1, \ldots, M_i\}$.

For the case of type-independent contracts, one adds the constraint $\mathbf{a}_{ik} = \mathbf{a}_{il}$ for all $i = 1, \ldots, N$ and $k \neq l$ in $\{1, \ldots, M_i\}$ and the incentive compatibility constraints are removed.

A common approach to solving bilevel programming problems is to replace the internal optimization with the corresponding Karush–Kuhn–Tucker condition. This approach is used when the internal problem is concave, i.e., when it has a unique maximum. However, in our case, concavity is not guaranteed, and we resort to nested optimization. We implement everything in Python using the Theano [19] symbolic computation package exploit automatic differentiation. We solve the follower problem using sequential least squares programming as implemented in the scipy package. We use simulated annealing to find the global optimum point of the leader problem. We first convert the constraint problem to the unconstrained problem using the penalty method such that

$$f\left(\mathbf{a}\right) = \mathbb{E}\left[u_0(\Pi_0)\right] + \sum_{i=1}^{N_c} \min\left(g_i\left(\mathbf{a}\right), 0\right) \tag{33}$$

where $g_i(\cdot)$'s are the constraints in (29)–(32). Maximizing the $f(\mathbf{a})$ in (33), is equivalent to finding the mode of the distribution

$$\pi_\gamma\left(\mathbf{a}\right) \propto \exp\left(\gamma f\left(\mathbf{a}\right)\right) \tag{34}$$

we use sequential Monte Carlo (SMC) [20] method to sample from this distribution by increasing γ from 0.001 to 50. To perform the SMC, we use the "pysmc" package [21]. To ensure the computational efficiency of our approach, we need to use a numerical quadrature rule to approximate the expectation over Ξ. This step is discussed in Appendix A. To guarantee the reproducibility of our results, we have published our code in an open source Github repository[1] with an MIT license.

F. Value Function and Risk Behavior

We assume two types of value functions, namely, the RB and RPI. Mathematically, we define these two value functions as

$$V_{\text{RB}} := v_0 \prod_{i=1}^{N} \left\{H(Q_i^* - 1)\right\} \tag{35}$$

and

$$V_{\text{RPI}} := v_0 \prod_{i=1}^{N} \left\{H(Q_i^* - 1)\right\}\left[1 + 0.2(Q_i^* - 1)\right] \tag{36}$$

respectively. In Fig. 2, we show these two value functions for one subsystem.

We consider two different risk behaviors for individuals, RA and risk neutral (RN). We use the utility function in (37), for the risk behavior of the agents and principal

$$u\left(\pi\left(\cdot\right)\right) = \begin{cases} a - be^{-c\pi(\cdot)}, & \text{for RA} \\ \pi\left(\cdot\right), & \text{for RN} \end{cases} \tag{37}$$

where $c = 2$ for an RA agent. The parameters a and b are

$$a = b = \frac{1}{1 - e^{-c}}.$$

We show these utility functions for the two different risk behaviors in Fig. 3.

[1] [Online]. Available: https://github.com/ebilionis/incentives

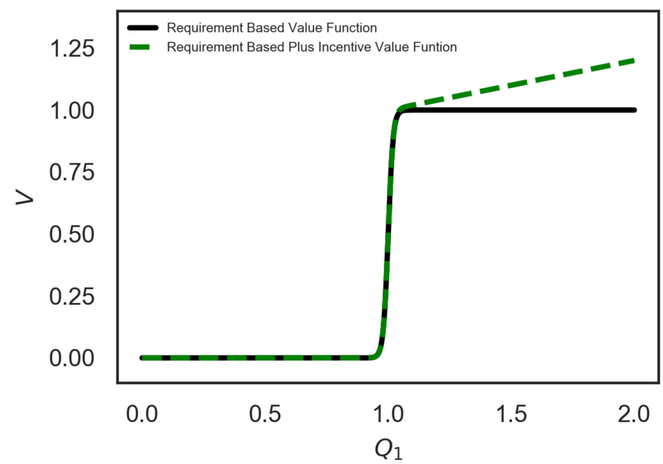

Fig. 2. RB value function (black solid line) and RPI value function (green dashed line).

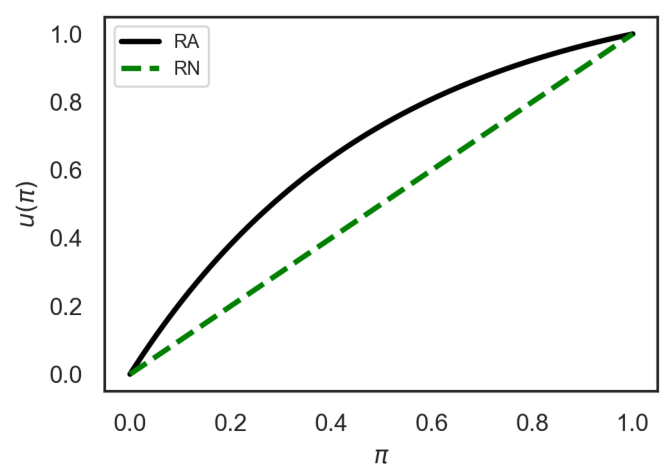

Fig. 3. Utility functions for RA (black solid line), RN (green dashed line).

III. Numerical Examples

In this section, we start by performing an exhaustive numerical investigation of the effects of task complexity, agent's cost of effort, uncertainty in the quality of the returned task, and adverse selection. In Section III-A1, we study the "moral hazard only" scenario with the RB transfer and value functions. In Section III-A2, we study the effect of the RPI transfer and value functions. We study the "moral hazard with adverse selection" in Section III-A3.

A. Numerical Investigation of the Proposed Model

In these numerical investigations, we consider a single RN principal and an RA agent. Each case study corresponds to a choice of task complexity [κ in (5)], cost of effort [c in (3)], and performance uncertainty [σ in (5)]. With regards to task complexity, we select $\kappa = 2.5$ for an easy task and $\kappa = 1.5$ for a hard task. For the cost of effort parameter, we associate $c = 0.1$ and $c = 0.4$ with the low- and high-cost agents, respectively. Finally, low- and high-uncertainty tasks are characterized by $\sigma = 0.1$ and $\sigma = 0.4$, respectively.

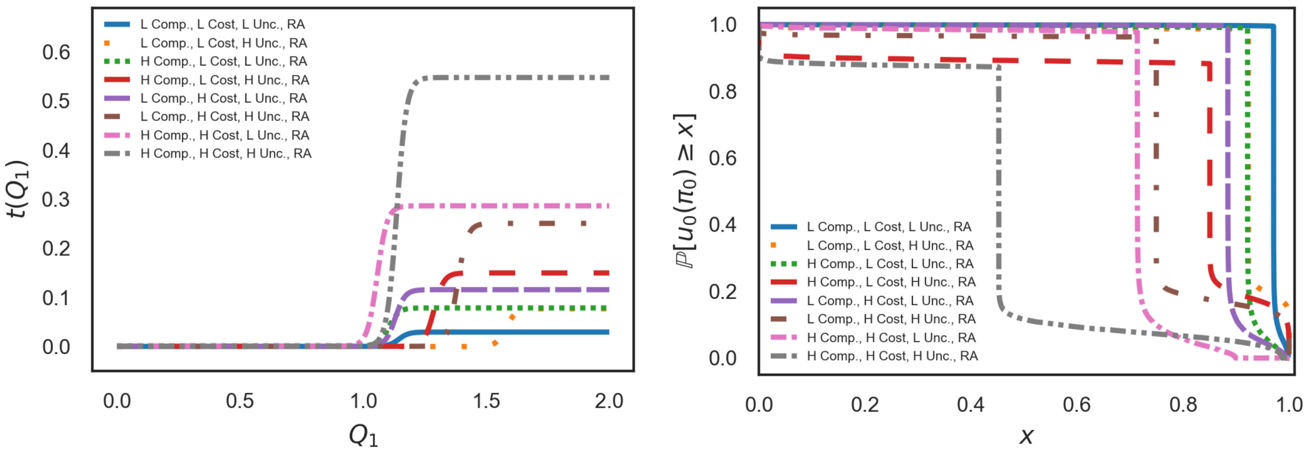

(a) (b)

Fig. 4. L and H stand for low and high, respectively, Comp. and Unc. stand for complexity and uncertainty, respectively. The low and high complexity denote the $\kappa_{11} = 2.5$ and $\kappa_{11} = 1.5$, respectively, low and high cost denote $c_{11} = 0.1$ and $c_{11} = 0.4$, respectively, low and high uncertainty denote $\sigma_{11} = 0.1$ and $\sigma_{11} = 0.4$, respectively, RA denotes the RA agent. (a) The RB transfer functions for several different agent types with respect to outcome of the subsystem (Q_1) for moral hazard scenario. (b) The exceedance for the moral hazard scenario using the RB transfer function.

Note that, the parameters $\kappa_{i\theta_i}$, $c_{i\theta_i}$, and $\sigma_{i\theta_i}$ have two indices. The first index i is the agent's (subsystem's) number and the second index is the type of the agent. We begin with a series of cases with a single agent with a known type denoted by 1 (moral-hazard-only case studies). In these cases, the parameters corresponding to complexity, cost, and uncertainty take the values κ_{11}, c_{11}, and σ_{11}, respectively. We end with a series of cases with a single agent but with an unknown type that can take two discrete, equally probable values 1 and 2 (moral-hazard-and-adverse-selection case studies). Consequently, κ_{11} denotes the effort coefficient of a type-1 agent 1, κ_{12} the same for a type-2 agent, and so on for all the other parameters.

To avoid numerical difficulties and singularities, we replace all Heaviside functions with a sigmoids, i.e.,

$$\hat{H}_\alpha(x) = \frac{1}{1 + e^{-\alpha x}} \tag{38}$$

where the parameter α controls the slope. We choose $\alpha = 50$ for the transfer functions and $\alpha = 100$ for the value function. We consider two types of value functions, RB and RPI value functions, see Section II-F. For the RB value function, we use the transfer function of (24) constrained so $a_{ik}, 3 = 0$ (RB transfer function). In other words, the agent is paid a constant amount if they achieve the requirement and there is no payment per quality exceeding the requirement. For the case of RPI value function, we remove this constraint.

1) Moral Hazard With RB Transfer and Value Functions: Consider the case of a single RA agent of known type and an RN principal with an RB value function. In Fig. 4(a), we show the transfer functions for several agent types covering all possible combinations of low/high complexity, low/high cost, and low/high task uncertainty. Fig. 4(b) depicts the probability that the principal's expected utility exceeds a given threshold for all these combinations. We refer to this curve as the *exceedance curve*. Finally, in Tables I and II, we report the expected utility

TABLE I
EXPECTED UTILITY OF THE PRINCIPAL FOR LOW-COST AGENT WITH RB VALUE FUNCTION

	Low Uncertainty	High Uncertainty
Low Complexity	0.97	0.93
High Complexity	0.92	0.79

TABLE II
EXPECTED UTILITY OF THE PRINCIPAL FOR HIGH-COST AGENT WITH RB VALUE FUNCTION

	Low Uncertainty	High Uncertainty
Low Complexity	0.89	0.77
High Complexity	0.72	0.45

of the principal for the low- and high-cost agents, respectively. We make the following observations.

1) For the same level of task complexity and uncertainty, but with increasing cost of effort.
   a) The optimal passed-down requirement decreases.
   b) The optimal payment for achieving the requirement increases.
   c) The principal's expected utility decreases.
   d) The exceedance curve shifts to the left.

   Intuitively, as the agent's cost of effort increases, the principal must make the contract more attractive to ensure that the participation constraints are satisfied. As a consequence, the probability that the principal's expected utility exceeds a given threshold decreases.
2) For the same level of task uncertainty and cost of effort, but with increasing complexity.
   a) The optimal passed-down requirement decreases.
   b) The optimal payment for achieving the requirement increases.
   c) The principal's expected utility decreases.
   d) The exceedance curve shifts to the left.

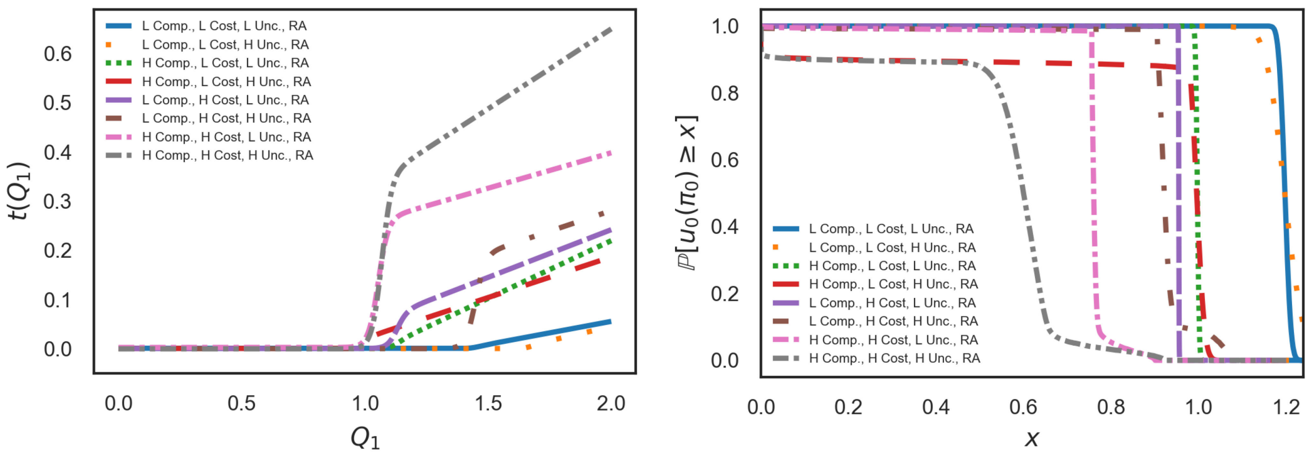

Fig. 5. L and H stand for low and high, respectively, Comp. and Unc. stand for complexity and uncertainty, respectively. The low and high complexity denote the $\kappa_{11} = 2.5$ and $\kappa_{11} = 1.5$, respectively, low and high cost denote $c_{11} = 0.1$ and $c_{11} = 0.4$, respectively, low and high uncertainty denote $\sigma_{11} = 0.1$ and $\sigma_{11} = 0.4$, respectively, RA denotes the RA agent. (a) The RPI transfer functions for several different agent types with respect to outcome of the subsystem (Q_1) for moral hazard scenario. (b) The exceedance curve for the moral hazard scenario using the RPI transfer function.

Thus, we see that an increase in task complexity has a similar effect as an increase in the agent's cost of effort. As in the previous case, to make sure that the agent wants to participate, the principal has to make the contract more attractive as task complexity increases.

3) For the same level of task complexity and cost of effort, but with increasing uncertainty.
   a) The optimal passed-down requirement increases.
   b) The optimal payment for achieving the requirement increases.
   c) The principals expected utility decreases.
   d) The exceedance curve shifts toward the bottom right.

   This case is the most interesting. Here, as the uncertainty of the task increases, the principal must increase the passed-down requirement to ensure that they are hedged against failure. At the same time, however, they must also increase the payment to ensure that the agent still has an incentive to participate.

4) For all cases considered, the optimal passed down requirement is greater than the true requirement (which is set to one). Note, however, this is not universally true. This article does not examine all possible combinations of cost, quality, and utility functions that could have been considered. Indeed, as we showed in our previous work [10], there are situations in which a smaller-than-the-true requirement can be optimal.

2) Moral Hazard With RPI Transfer and Value Functions: This case is identical to Section III-A1, albeit we use the RPI value function, see Section II-F, and the RPI transfer function [see (24)]. Fig 5(a), depicts the transfer functions for all combinations of agent types and task complexities. In Fig. 5(b), we show the exceedance curve using the RPI value and transfer functions. Finally, in Tables III and IV, we report the expected utility of the principal using the RPI transfer and value functions for the low- and high-cost agents, respectively. The results are qualitative similar to Section III-A1, with the following additional observations.

TABLE III
EXPECTED UTILITY OF THE PRINCIPAL FOR LOW-COST AGENT WITH RPI VALUE FUNCTION

	Low Uncertainty	High Uncertainty
Low Complexity	1.2	1.2
High Complexity	1.0	0.89

TABLE IV
EXPECTED UTILITY OF THE PRINCIPAL FOR HIGH-COST AGENT WITH RPI VALUE FUNCTION

	Low Uncertainty	High Uncertainty
Low Complexity	0.95	0.93
High Complexity	0.76	0.56

TABLE V
SUMMARY OF THE OBSERVATIONS

complexity	agent cost	uncertainty	requirement	payment
↑	fix	fix	↓	↑
fix	↑	fix	↓	↑
fix	fix	↑	↑	↑

1) For the same level of task complexity, uncertainty, and agent cost, the optimal reward for achieving the requirement decreases compared to the same cases in Section III-A1. Intuitively, as the principal has the option to reward the agent based on the quality exceeding the requirement, they prefer to pay less for fulfilling the requirement. Instead, the principal incentivizes the agent to improve the quality beyond the optimal passed-down requirement.
2) The slope of the transfer function beyond the passed-down requirement is almost identical to the slope of the value function.

In Table V, we summarize our observations for the results in Sections III-A1 and III-A2. In this table, we show how the passed-down requirement and payment change when we fix two parameters of the model (we denote it by "fix" in the table) and

TABLE VI
EXPECTED UTILITY OF THE AGENT WITH UNKNOWN COST FOR TWO DIFFERENT CONTRACTS

	$\mathbb{E}[u_1(\cdot)]$	$\mathbb{E}[u_0(\cdot)]$
Low Cost Agent (Type 1)	0.39	0.72
High Cost Agent (Type 2)	0	0.72

vary the third parameter. We denote increase by ↑ and decrease by ↓.

3) Moral Hazard With Adverse Selection: Consider the case of a single RA agent of unknown type, which takes two possible values, and an RN principal with an RB value function. We consider two possibilities for the unknown type as follows.

1) *Unknown Cost of Effort:* Here, we set $\kappa_{11} = \kappa_{12} = 1.5$ $(p(\kappa = \kappa_{11} = 1.5) = 1)$, $\sigma_{11} = \sigma_{12} = 0.1$ $(p(\sigma = \sigma_{11} = 0.1) = 1)$, and $p(c = c_{11} = 0.1) = 0.5$ and $p(c = c_{12} = 0.4) = 0.5$.
2) *Unknown Task Complexity:* For the unknown quality, we assume that $p(\kappa = \kappa_{11} = 2.5) = 0.5$ and $p(\kappa = \kappa_{12} = 1.5) = 0.5$, $\sigma_{11} = \sigma_{12} = 0.4$ $(p(\sigma = \sigma_{11} = 0.4) = 1)$, and $c_{11} = c_{12} = 0.4$ $(p(c = c_{11} = 0.4) = 1)$.

In this scenario, we maximize the expected utility of the principal subject to constraints in (29)–(32). The incentive compatibility constraint, (32), guarantees that the agent will choose the contract that is suitable for their true type. In other words, as there are two agent types' possibilities, the principal must offer two contracts, see Fig. 1(b). These two contracts must be designed in a way that there is no benefit for the agent to deviate from their true type, i.e., the contracts enforce the agent to be truth telling.

Solving the constraint optimization problem yields

$$\mathbf{a}_{11} = \mathbf{a}_{12} = (0, 0.29, 1.06)$$

i.e., the two contracts collapse into one. Note that the resulting contract is the same as the pure moral hazard case, Section III-A1, for an agent with type $\kappa_{11} = 1.5$, $\sigma_{11} = 0.1$, and $c_{11} = 0.4$. In other words, the principal must behave as if there was only a high-cost agent. That is, there are no contacts that can differentiate between a low- and a high-cost agent in this case.

A similar outcome occurs for unknown task complexity. The solution of the constraint optimization problem for this scenario is

$$\mathbf{a}_{11} = \mathbf{a}_{12} = (0, 0.08, 1.11)$$

which is the same as the optimum contract that is offered for the pure moral hazard case, Section III-A1, for an agent with type $\kappa_{11} = 1.5$, $\sigma_{11} = 0.4$, and $c_{11} = 0.4$. Therefore, in this case the principal must behave as if there the task is of high complexity.

Note that in both cases above, the collapse of the two contracts to one contract is not a generalizable property of our model. In particular, it may not happen if more flexible transfer functions are allowed, e.g., ones that allow performance penalties.

In Fig. 6, we show the transfer functions for the adverse selection scenarios with unknown cost and unknown quality. In Tables VI and VII, we show the expected utility of two types of agents and the principal using the optimum contract for unknown cost and unknown quality, respectively. The unknown cost and the unknown task complexity can be summarized as follows.

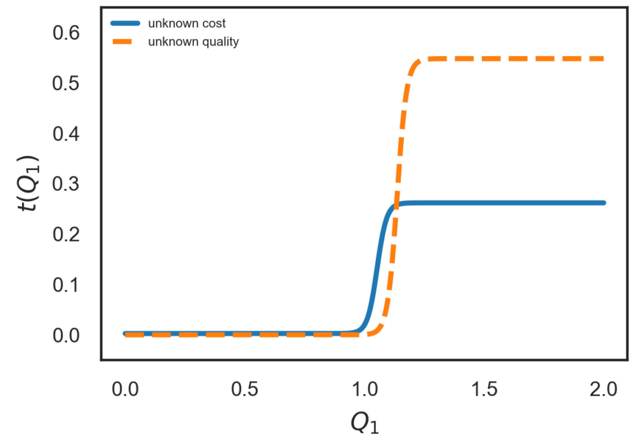

Fig. 6. Transfer function for the adverse selection scenarios with unknown cost (solid line) and unknown quality (dashed line), the agent is an RA. For unknown cost: $\kappa_{11} = \kappa_{12} = 1.5$ with probability 1, $\sigma_{11} = \sigma_{12} = 0.1$ with probability 1, and $c_{11} = 0.1$ with probability 0.5 and $c_{12} = 0.4$ with probability 0.5. For unknown quality: $\kappa_{11} = 2.5$ with probability 0.5 and $\kappa_{12} = 1.5$ with probability 0.5, $\sigma_{11} = \sigma_{12} = 0.4$ with probability 1, and $c_{11} = c_{12} = 0.4$ with probability 1.

TABLE VII
EXPECTED UTILITY OF THE AGENT WITH UNKNOWN QUALITY FOR TWO DIFFERENT CONTRACTS

	$\mathbb{E}[u_1(\cdot)]$	$\mathbb{E}[u_0(\cdot)]$
Low Complexity (Type 1)	0.52	0.45
High Complexity (Type 2)	0	0.45

1) The unknown cost comprises the following.
   a) The optimum transfer function for this problem is as same as that the principal would have offered for a single-type high-cost agent with $c_{11} = c_{12} = 0.4$ (moral hazard scenario with no adverse selection).
   b) The expected utility of the low-cost agent (efficient agent) is greater than that of the high-cost agent.

   In this case, the low-cost agent benefits because of information asymmetry. In other words, the principal must pay an information rent to the low-cost agent to reveal their type.
2) The unknown task complexity comprises the following.
   a) The optimum contract in this case is the contract that the principal would have offered for the single-type high-complexity task with $\kappa_{11} = \kappa_{12} = 1.5$.
   b) the expected utility of an agent dealing with a low-complexity task is greater than that of an agent dealing with a high-complexity task.

   Again, due to the information asymmetry, the agent benefits if the task complexity is low. The principal must pay an information rent to reveal the task complexity.

B. Satellite Design

In this section, we apply our method on a *simplified* satellite design. Typically a satellite consists of seven different subsystems [22], namely, electrical power subsystem, propulsion, attitude determination and control, on-board processing, telemetry, tracking and command, structures and thermal subsystems. We focus our attention on the propulsion subsystem ($N = 1$). To simplify the analysis, we assume that the design of these subsystems is assigned to a sSE in a one-shot fashion. Note that, the actual SEP of the satellite design is an iterative process and the information and results are exchanged back and forth in each iteration. Our model is a crude approximation of reality. The goal of the SE is to optimally incentivize the sSE to produce subsystem designs that meet the mission's requirements. Furthermore, we assume that the propulsion subsystem is decoupled from the other subsystems, i.e., there is no interactions between them, and that the SE knows the types of each sSE and therefore, there is no information asymmetry.

To extract the parameters of the model, i.e., $a_{11}, \sigma_{11}, c_{11}$, we will use available historical data. To this end, let I_1 be the cumulative, sector-wide investment on the propulsion subsystem and G_1 be the delivered specific impulse of solid propellants (I_{sp}). The specific impulse is defined as the ratio of thrust to weight flow rate of the propellant and is a measure of energy content of the propellants [22].

Historical data, say $\mathcal{D}_1 = \{(I_{1,i}, G_{1,i})\}_{i=1}^S$, of these quantities are readily available for many technologies. Of course, cumulative investment and best performance increase with time, i.e., $I_{1,i} \leq I_{1,i+1}$ and $G_{1,i} \leq G_{1,i+1}$. We model the relationship between G_1 and I_1 as

$$G_1 = G_{1,S} + A_1(I_1 - I_{1,S}) + \Sigma_1 \Xi_1 \tag{39}$$

where $G_{1,S}$ and $I_{1,S}$ are the current states of these variables, $\Xi_1 \sim \mathcal{N}(0, 1)$, and A_1 and Σ_1 are parameters to be estimated from the all available data, $\mathcal{D}_1$. We use a maximum likelihood estimator for A_1 and Σ_1. This is equivalent to a least squares estimate for A_i

$$\hat{A}_1 = \arg\min_{A_1} \sum_{i=1}^{S} [G_{1,S} + A_1(I_{1,i} - I_{1,S}) - G_{1,i}]^2 \tag{40}$$

and to setting Σ_1 equal to the mean residual square error

$$\hat{\Sigma}_1 = \frac{1}{S} \sum_{i=1}^{S} \left[G_{1,S} + \hat{A}_1(I_{1,i} - I_{1,S}) - G_{1,i}\right]^2. \tag{41}$$

Now, let G_1^r be the required quality for the propulsion subsystem in physical units. The scaled quality of a subsystem Q_1, can be defined as

$$Q_1 = \frac{G_1 - G_{1,S}}{G_1^r - G_{1,S}} \tag{42}$$

with this definition, we get $Q_1 = 0$ for the state-of-the-art, and $Q_1 = 1$ for the requirement. Substituting (39) in (42) and using the maximum likelihood estimates for A_1 and Σ_1, we obtain

$$Q_1 = \frac{\hat{A}_1}{G_1^r - G_{1,S}}(I_1 - I_{1,S}) + \frac{\hat{\Sigma}_1}{G_1^r - G_{1,S}}\Xi_1. \tag{43}$$

From this equation, we can identify the uncertainty σ_{11} in the quality function as

$$\sigma_{11} = \frac{\hat{\Sigma}_1}{G_1^r - G_{1,S}}. \tag{44}$$

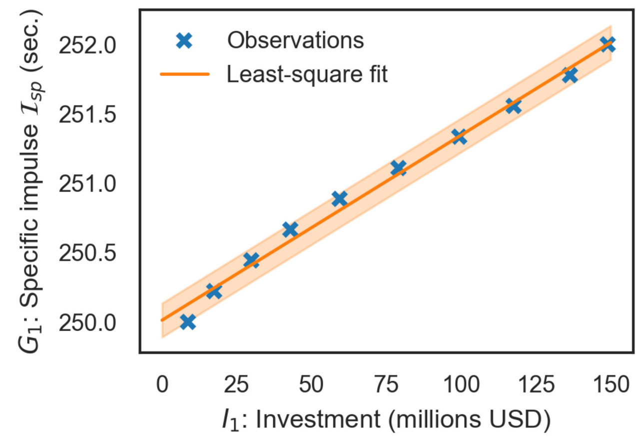

Fig. 7. Satellite case study (propulsion subsystem): Historical data (1979–1988) of specific impulse of solid mono-propellants versus cumulative investment per firm. The solid line and the shaded area correspond to the maximum likelihood fit of a linear regression model and the corresponding 95% prediction intervals, respectively.

Finally, we need to define effort. Let T_1 represents the time for which the propulsion engineer is to be hired. The cost of the agent per unit time is C_1. T_1 is just the duration of the SEP, we consider. The value C_1T_1 can be read from the balance sheets of publicly traded firms related to the technology. We can associate the effort variable e_1 with the additional investment required to buy the time of one engineer

$$e_1 = \frac{I_1 - I_{1,S}}{C_1 T_1} \tag{45}$$

that is, $e_1 = 1$ corresponds to the effort of one engineer for time T_1. Let us assume there are Z engineers work on the subsystem. Comparing this equation, (43) and (5), we get that the κ_{11} coefficient is given by

$$\kappa_{11} = \frac{Z C_1 T_1 \hat{A}_1}{G_1^r - G_{1,S}}. \tag{46}$$

To complete the picture, we need to talk about the value V_0 (in USD) of the system if the requirements are met. We can use this value to normalize all dollar quantities. That is, we set

$$v_0 = 1 \tag{47}$$

and for the cost per square effort of the agent we set

$$c_{11} = \frac{Z C_1 T_1}{V_0}. \tag{48}$$

Finally, we use some real data to fix some of the parameters. Trends in delivered I_{sp} [G_1 (s)] and investments by NASA [I_1 (millions USD)] in chemical propulsion technology with time are obtained from [23] and [24], respectively. The state-of-the-art solid propellant technology corresponds to a $G_{1,S}$ value of 252 s and $I_{1,S}$ value of 149.1 million USD. The maximum likelihood fit of the parameters results in a regression coefficient of $\hat{A}_1 = 0.0133$ s per million USD, and standard deviation $\hat{\Sigma}_1 = 0.12$ s. The corresponding data and the maximum likelihood fit are illustrated in Fig. 7. The value of C_1 is the median salary (per time) of a propulsion engineer which is approximately

TABLE VIII
MODEL PARAMETERS FOR TWO CASE STUDIES

G_1^r	V_0	κ_{11}	σ_{11}	c_{11}
252.2 s	50,000,000 USD	1.6	0.6	0.5
252.25 s	60,000,000 USD	1.28	0.48	0.4

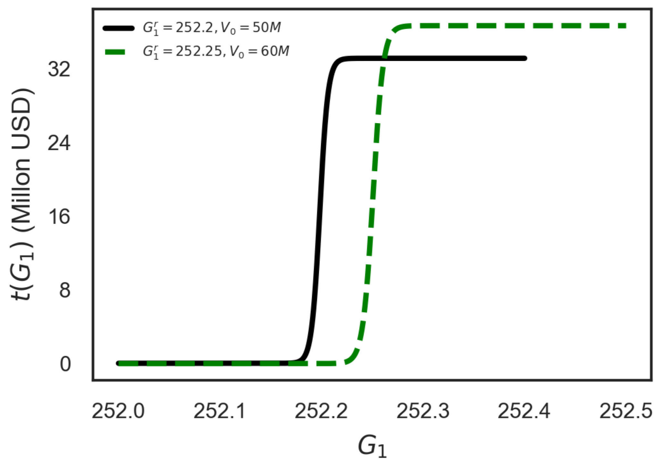

Fig. 8. Contracts for two case studies in satellite design.

USD120 000/year, according to the data obtained from [25]. For simplicity, also assume that $T_1 = 1$ year. Moreover, we assume that there are 200 engineers work on the subsystem, $Z = 200$. We will examine two case studies which is summarized in Table VIII.

Using RB value function, we depict the contracts for these two scenarios in Fig. 8.

IV. CONCLUSION

We developed a game-theoretic model for a one-shot shallow SEP. We posed and solved the problem of identifying the contract (transfer function) that maximizes the principal's expected utility. Our results show that, the optimum passed-down requirement is different from the real-system requirement. For the same level of task complexity and uncertainty, as the agent cost of effort increases, the passed-down requirement decreases and the award to achieving the requirement increases. In this way, the principal makes the contract more attractive to the high-cost agent and ensures that the participation constraint is satisfied. Similarly, for the same level of task uncertainty and cost of effort, increasing task complexity results in lower passed-down requirement and larger award for achieving the requirement. For the same level of task complexity and cost of effort, as the uncertainty increases both the passed-down requirement and the award for achieving the requirement increase. This is because the principal wants to make sure that the system requirements are achieved. Moreover, by increasing the task complexity, the task uncertainty, or the cost of effort, the principal earns less and the exceedance curve is shifted to the left. Using the RPI contracts, the principal pays smaller amount for achieving the requirement but, instead, they pay for per quality exceeding the requirement.

For the adverse selection scenario with RB value function, we observe that when the principal is maximally uncertain about the cost of the agent, the optimum contracts are equivalent to the contract designed for the high-cost agent in the single-type case with no adverse selection. The low-cost agent earns more expected utility than the high-cost agent. This is the information rent that the principal must pay to reveal the agents' types. Similarly, if the principal is maximally uncertain about the task complexity, the two optimum contracts for the unknown quality are equivalent to the contract that is offered to the high-complexity task where there is no adverse selection. Note that, the equivalence of the contracts in adverse selection scenario with the contract that is offered in absence of adverse selection is not universal. If the class of possible contracts is enlarged, e.g., to allow penalties, there may be a set of two contracts that differentiate types.

There are still many remaining questions in modeling SEPs using a game-theoretic approach. First, there is a need to study the hierarchical nature of SEPs with potentially coupled subsystems. Second, true SEPs are dynamic in nature with many iterations corresponding to exchange of information between the various agents. These are the topics of ongoing research toward a theoretical foundation of systems engineering design that accounts for human behavior.

APPENDIX A
NUMERICAL ESTIMATION OF THE REQUIRED EXPECTATIONS

For the numerical implementation of the suggested model, we need to be able to carry out expectations of the form of (9) a.k.a. (8) and (7). Since, we have at most two possible types in our case studies, the summation over the possible types is trivial. Focusing on expectations over Ξ, we evaluate them using a sparse grid quadrature rule [26]. In particular, any expectation of the form $\mathbb{E}[g(\Xi)]$ is approximated by

$$\mathbb{E}[g(\Xi)] \approx \sum_{s=1}^{N_s} w^{(s)} g\left(\xi^{(s)}\right) \tag{49}$$

where $w^{(s)}$ and $\xi^{(s)}$ are the $N_s = 127$ quadrature points of the level 6 sparse grid quadrature constructed by the Gauss-Hermite 1-D quadrature rule.

REFERENCES

[1] G. Locatelli, "Why are megaprojects, including nuclear power plants, delivered overbudget and late? Reasons and remedies," 2018, *arXiv:1802.07312*.

[2] U. S. G. A. Office, "Navy shipbuilding past performance provides valuable lessons for future investments," 2018.

[3] NASA, *Assessments of Major Projects*. GAO-19-262SP, 2018.

[4] I. Maddox, P. Collopy, and P. A. Farrington, "Value-based assessment of DoD acquisitions programs," *Procedia Comput. Sci.*, vol. 16, pp. 1161–1169, 2013.

[5] P. D. Collopy and P. M. Hollingsworth, "Value-driven design," *J. Aircr.*, vol. 48, no. 3, pp. 749–759, 2011.

[6] A. Bandura, "Social cognitive theory: An agentic perspective," *Annu. Rev. Psychol.*, vol. 52, no. 1, pp. 1–26, 2001.

[7] M. Shergadwala, I. Bilionis, K. N. Kannan, and J. H. Panchal, "Quantifying the impact of domain knowledge and problem framing on sequential decisions in engineering design," *J. Mech. Des.*, vol. 140, no. 10, 2018, Art. no. 101402.

[8] A. M. Chaudhari, Z. Sha, and J. H. Panchal, "Analyzing participant behaviors in design crowdsourcing contests using causal inference on field data," *J. Mech. Des.*, vol. 140, no. 9, 2018, Art. no. 091401.

[9] S. D. Vermillion and R. J. Malak, "Using a principal-agent model to investigate delegation in systems engineering," in *Proc. ASME Int. Des. Eng. Tech. Conf. Comput. Inf. Eng. Conf*, 2015, pp. V01BT02A046–V01BT02A046.

[10] S. Safarkhani, V. R. Kattakuri, I. Bilionis, and J. Panchal, "A principal-agent model of systems engineering processes with application to satellite design," 2019, *arXiv:1903.06979*.

[11] J. Watson, "Contract, mechanism design, and technological detail," *Econometrica*, vol. 75, no. 1, pp. 55–81, 2007.

[12] S. Safarkhani, I. Bilionis, and J. Panchal, "Understanding the effect of task complexity and problem-solving skills on the design performance of agents in systems engineering," in *Proc. ASME Int. Des. Eng. Tech. Conf. Comput. Inf. Eng. Conf.*, Aug. 2018, Paper DETC2018-85941.

[13] J. A. Mirrlees, "The theory of moral hazard and unobservable behaviour: Part I," *Rev. Econ. Stud.*, vol. 66, no. 1, pp. 3–21, 1999.

[14] W. P. Rogerson, "The first-order approach to principal-agent problems," *Econometrica*, vol. 53, no. 6, pp. 1357–1367, 1985.

[15] R. Ke and C. T. Ryan, "A general solution method for moral hazard problems," *Theor. Econ.*, vol. 13, no. 3, pp. 1425–1481, 2018.

[16] R. Myerson, *Game Theory: Analysis of Conflict*. Cambridge, MA, USA: Harvard Univ. Press, 1991.

[17] R. B. Myerson, "Optimal auction design," *Math. Oper. Res.*, vol. 6, no. 1, pp. 58–73, 1981.

[18] D. Kahneman and A. Tversky, "Prospect theory: An analysis of decision under risk," in *Handbook of the Fundamentals of Financial Decision Making: Part I*. Singapore: World Scientific, 2013, pp. 99–127.

[19] Theano Development Team, "Theano: A Python framework for fast computation of mathematical expressions," May 2016. [Online]. Available: http://arxiv.org/abs/1605.02688

[20] A. Doucet, S. Godsill, and C. Andrieu, "On sequential monte carlo sampling methods for Bayesian filtering," *Statist. Comput.*, vol. 10, no. 3, pp. 197–208, 2000.

[21] I. Bilionis, "Sequential Monte Carlo working on top of pymc," 2014. [Online]. Available: https://github.com/PredictiveScienceLab/pysmc

[22] J. R. Wertz, D. F. Everett, and J. J. Puschell, *Space Mission Engineering: The New SMAD*. Cleveland, OH, USA: Microcosm Press, 2011.

[23] "Solid," 1997. [Online]. Available: http://www.astronautix.com/s/solid.html

[24] "NASA Historical Data Books." [Online]. Available: https://history.nasa.gov/SP-4012/vol6/cover6.html

[25] "Propulsion Engineer Salaries," 2019. [Online]. Available: https://www.paysa.com/salaries/propulsion-engineer--t

[26] T. Gerstner and M. Griebel, "Numerical integration using sparse grids," *Numer. Algorithms*, vol. 18, no. 3, Jan. 1998, Art. no. 209.