Proceedings of the ASME 2018 International Design Engineering Technical Conferences & Computers and Information in Engineering Conference
IDETC/CIE 2018
August 26-29, 2018, Quebec City, Canada

DETC2018-85941

UNDERSTANDING THE EFFECT OF TASK COMPLEXITY AND PROBLEM-SOLVING SKILLS ON THE DESIGN PERFORMANCE OF AGENTS IN SYSTEMS ENGINEERING

Salar Safarkhani
School of Mechanical Engineering
Purdue University
West Lafayette, Indiana 47906
Email: ssafarkh@purdue.edu

Ilias Bilionis*
School of Mechanical Engineering
Purdue University
West Lafayette, Indiana 47906
Email: ibilion@purdue.edu

Jitesh H. Panchal
School of Mechanical Engineering
Purdue University
West Lafayette, Indiana 47906
Email: panchal@purdue.edu

ABSTRACT

Systems engineering processes coordinate the efforts of many individuals to design a complex system. However, the goals of the involved individuals do not necessarily align with the system-level goals. Everyone, including managers, systems engineers, subsystem engineers, component designers, and contractors, is self-interested. It is not currently understood how this discrepancy between organizational and personal goals affects the outcome of complex systems engineering processes. To answer this question, we need a systems engineering theory that accounts for human behavior. Such a theory can be ideally expressed as a dynamic hierarchical network game of incomplete information. The nodes of this network represent individual agents and the edges the transfer of information and incentives. All agents decide independently on how much effort they should devote to a delegated task by maximizing their expected utility; the expectation is over their beliefs about the actions of all other individuals and the moves of nature. An essential component of such a model is the quality function, defined as the map between an agent's effort and the quality of their job outcome. In the economics literature, the quality function is assumed to be a linear function of effort with additive Gaussian noise. This simplistic assumption ignores two critical factors relevant to systems engineering: (1) the complexity of the design task, and (2) the problem-solving skills of the agent. Systems engineers establish their beliefs about these two factors through years of job experience. In this paper, we encode these beliefs in clear mathematical statements about the form of the quality function. Our approach proceeds in two steps: (1) we construct a generative stochastic model of the delegated task, and (2) we develop a reduced order representation suitable for use in a more extensive game-theoretic model of a systems engineering process. Focusing on the early design stages of a systems engineering process, we model the design task as a function maximization problem and, thus, we associate the systems engineer's beliefs about the complexity of the task with their beliefs about the complexity of the function being maximized. Furthermore, we associate an agent's problem solving-skills with the strategy they use to solve the underlying function maximization problem. We identify two agent types: "naïve" (follows a random search strategy) and "skillful" (follows a Bayesian global optimization strategy). Through an extensive simulation study, we show that the assumption of the linear quality function is only valid for small effort levels. In general, the quality function is an increasing, concave function with derivative and curvature that depend on the problem complexity and agent's skills.

NOMENCLATURE

$A(x,\omega)$	Attribute function
CDF	Cumulative distribution function
E	Space of all possible effort levels
EI	Expected improvement

*Address all correspondence to this author.

GP	Gaussian process
KLE	Karhunen-Loéve expansion
PCA	Principal component analysis
PDF	Probability density function
SE	Systems engineer
SEP	Systems engineering process
$Q(e,\theta,\omega)$	Quality function
X	Space of candidate designs
e	Effort level (an element of E)
$m(x)$	Mean function of GP
$k(x,x')$	Covariance function of GP
sSE	Subsystem engineer
x	A candidate design (an element of X)
Θ	Space of all possible agent types
Ω	Sample space
θ	Type of an agent (an element of Θ)
ω	State of nature (an element of Ω)
$\mathscr{F}$	A σ-algebra of subsets of Ω
$\mathbb{P}$	Probability measure defined for all sets in $\mathscr{F}$

1 INTRODUCTION

Systems engineering processes (SEP) require the coordination of a large number of individuals, e.g., managers, systems engineers (SE), subsystem engineers (sSE), and contractors, to establish the design, deployment, operation, and retirement of complex systems [1, 2]. Naturally, these individuals are self-interested, i.e., they have their personal agendas which are not necessarily aligned with the system-level objectives. It has been postulated that this discrepancy between organizational and personal goals may be one of the leading factors behind the increasing cost overruns and delays in modern systems engineering [3]. However, to date, there is no comprehensive SEP theory which models the effects of human behavior.

To account for human behavior, we have to model SEPs within a game theoretic framework [4,5]. In particular, a SEP can be viewed as a dynamic hierarchical network game. Each layer of the hierarchy captures the interactions of a *principal*, e.g., the manager, the SE, or a sSE, with several *agents*, e.g., the SE, sSEs, component engineers, or contractors. At each iteration, the principal assigns tasks to the agents encoded via performance-based contracts. The agents, select how much effort to devote to their tasks by maximizing their expected utility. Finally, they report the product of their efforts back to the principal. Note that in this hierarchy, the agent of a layer may be the principal of a subsequent layer. The goal of the principal at the top of the hierarchy, presumably the owner of the system, is to select the contracts that maximize the expected system-level utility. In the economics literature, this is known as the *mechanism design* problem [6].

A critical component of any principal-agent model is the *quality function*, defined as the stochastic map between an agent's effort and the quality of the product they deliver. To be more precise, the quality function models the beliefs of the principal about the effort-to-quality map, i.e., it models what the principal *thinks* they can get if the agent decides to spend a given amount of effort. For analytical convenience, the quality function is usually taken to be a linear function of effort with some additive Gaussian noise [7]. This simplistic assumption is not sufficient for capturing the beliefs about the outcome of an assigned task within the context of a SEP. Focusing on the early design stages of a SEP, the outcome of design tasks depends predominately on beliefs about the complexity of the underlying problem and the problem-solving skills of the engaged agent.

The objective of this work is to mathematically model the dependence of the quality function on a principal's beliefs about the task complexity and the problem-solving skills of an agent, within the context of the early design stages of a SEP. We achieve this in two steps: (1) constructing a generative stochastic model of the delegated task, and (2) developing a reduced order representation suitable for use in an extensive game-theoretic framework. The generative model is essentially a random process labeled by effort. Each sample from this random process is a plausible effort-to-quality map. The reduced-order model is a mathematically convenient approximation of this random process constructed from multiple samples of effort-to-quality maps.

The details of the generative model are as follows. The design task assigned to an agent is modeled as a scalar function maximization problem. An agent's effort is measured in function evaluations used to solve this maximization problem. We assume that the principal encodes their beliefs about the complexity of the problem using a Gaussian process (GP) prior [8] over the space of possible scalar functions, e.g., by selecting a suitable mean and covariance function. We associate the problem-solving skills of the agent with the maximization strategy they choose to employ. We identify two agent types: "naïve" and "skillful". The naïve agent solves the maximization problem using random search. The skillful agent solves the maximization problem using Bayesian global optimization [9, 10]. That is, the naïve agent does not learn from past experience whereas the skillful agent does learn. The naïve approach has an alternative interpretation as a parallel search representative of a scenario where there is a team of engineers, all developing different ideas concurrently, and the team gets together and decides on the best solution at the end of the process. To obtain a plausible realization of the effort-to-quality map, we sample a scalar function from the GP prior and we simulate the behavior of the agent. Using extensive sampling datasets, we constructed a reduced order model by employing the Karhunen-Loève theorem [11].

The outline of the paper is as follows. In Sec. 2, we present our methodology starting with the definition of the quality function. In Sec. 2.1, we model a design task as a function maximization problem. In Sec. 2.2, we model the state of knowledge of the principal about the function that the agent is maximizing and we define the concept of task complexity. Sec. 2.3 discusses the modeling of the problem solving skills of an agent. The reduced order model is presented in Sec. 2.4. In Sec. 3, we present our numerical experiments and study the effect of complexity and problem solving skills on the quality function. We conclude in Sec. 4.

2 METHODOLOGY

Let $(\Omega, \mathscr{F}, \mathbb{P})$ be a probability space, where Ω is a sample space, $\mathscr{F}$ is a σ-algebra of subsets of Ω, and $\mathbb{P}$ is a probability measure. Elements of Ω are denoted by $\omega \in \Omega$. An $\omega \in \Omega$ corresponds to a specific realization of everything that is random in our models.

Consider the early design phase of a SEP and focus on an agent that resides at the leaves of the hierarchy, i.e., of a component engineer or a subcontractor. The behavior of these leaf agents depends only on inputs from the immediately higher level of the hierarchy and, thus, it is a natural starting point. The behavior of agents that lie at intermediate points of the SEP hierarchy should be defined recursively given inputs from even higher levels and assuming that subordinate agents act optimally. Such recursive constructions are the subject of on going research and will not be discussed in this paper. We will construct a stochastic model of the quality function of this leaf agent based on some generic beliefs about the complexity of the task that is assigned to them as well as their problem-solving skills. To this end, let $E = \{0, 1, 2, \dots\}$ be the space of possible effort levels that the agent may choose, denoting a specific effort level by $e \in E$. Similarly, let $\Theta = \{\text{"naïve"}, \text{"skillful"}\}$ be the set of possible agent types, denoting a specific type by $\theta \in \Theta$. For simplicity, let us measure the quality characterizing the outcome of a design task with a single real number, e.g., the utility/value of that design to the principal allocating the task.

Using these definitions, the *quality function* can be thought as a random process

$$Q : E \times \Theta \times \Omega \to \mathbb{R},$$

i.e., as a collection of Borel measurable functions $\{Q(e, \theta, \cdot)\}_{e \in E, \theta \in \Theta}$. Of course, not all such random processes yield reasonable quality functions. Any explicit model we construct should satisfy a minimum set of requirements:

1. $Q(e, \theta, \cdot)$ must have increasing paths in e, i.e., for all $e_1, e_2 \in E$ with $e_1 < e_2$, we must have:

$$Q(e_1, \theta, \omega) \leq Q(e_2, \theta, \omega),$$

   for all $\theta \in \Theta$, and for almost all $\omega \in \Omega$. This ensures that the more effort an agent spends the better the quality of the product of the task.

2. $Q(e, \theta, \omega)$ must be bounded above in probability, i.e., for all $\varepsilon > 0$, there exists an $M > 0$ such that

$$\mathbb{P}[\{\omega : Q(e, \theta, \omega) > M\}] < \varepsilon,$$

   for all $e \in E, \theta \in \Theta$. This ensures that the outcome quality cannot grow without bound no matter how much the effort of the agent is, e.g., because of physical limitations in the design.

We will now construct a generative model of $Q(e, \theta, \omega)$ that satisfies these properties and, furthermore, it makes explicit the dependence on task complexity and agent skills.

2.1 Modeling the Design Task as a Scalar Function Maximization

Let us assume that the agent's task is to maximize a scalar function over a set of candidate designs. For clarity, let the set of candidate designs be $X = [0, 1]$ (the ideas can easily be generalized to arbitrary sets.) The principal does not know exactly what the agent's objective function is. Nevertheless, let us assume that the principal believes that the scalar function the agent is maximizing is a sample from of a random process $A : X \times \Omega \to \mathbb{R}$. Let us refer to $A(x, \omega)$ as the *attribute function*. Thus, the principal believes that the agent is solving:

$$\max_{x \in X} A(x, \omega), \tag{1}$$

but they do not know exactly what $A(x, \omega)$ is.

2.2 Modeling the Beliefs of the Principal about the Attribute Function

To proceed, let us assume that the principal models the attribute function $A(x, \omega)$ as a GP, i.e.,

$$A \sim \mathrm{GP}(m, k), \tag{2}$$

where $m : X \to \mathbb{R}$ and $k : X \times X \to \mathbb{R}$ are the mean and the covariance function, respectively. The choices of GP priors is motivated by their successful application to human function learning by Griffiths et al. [9]. The beliefs of the principal about the plausible $A(x, \omega)$ are encoded in their choice of mean and covariance functions. Of course, any particular choice is context-dependent and the principal should make every effort to use any available data to estimate them. However, to advance our study, let us assume that principal's beliefs are reflected by the choice:

$$\begin{aligned} m(x) &= c \\ k(x, x') &= \sigma_s^2 \exp\left\{-\frac{(x - x')^2}{2l^2}\right\}, \end{aligned} \tag{3}$$

with signal strength σ_s and length-scale l. A constant mean encodes the principal's lack of knowledge about any particular trend of the attribute function. The regularity of the covariance function determines the regularity of sampled attribute functions, see [12]. Here, our choice of the *squared exponential* covariance function guarantees that the sampled attribute functions are infinitely differentiable with respect to x. The choice of the signal strength controls the principal's beliefs about the possible variations of the attribute function about its mean. Without loss of generality, we may set $c = 0$ and $\sigma_s = 1$, after a suitable affine transformation of the attribute function. Therefore, the only remaining parameter is the length-scale l. The length-scale of the covariance function is a measure of the problem complexity. On one hand, decreasing the length-scale, the fluctuations of the attribute function increase, making the task of the agent more difficult. On the other hand, increasing the length-scales yields smoother attribute functions thus making the underlying task easier. Of course, there are other aspects of problem complexity such as the number of dimensions, possible discontinuities, discrete choices, etc. Those are not considered in this paper.

2.3 Modeling the Problem-solving Skills of the Agent

The agent solves the problem of Eq. (1) by repeatedly evaluating the attribute function at design points of their choice. Each function evaluation counts as one unit of effort. Let $X_i : \Theta \times \Omega \to X$ be the random variables corresponding to the agent's query of the attribute function at effort level $i = 1, 2, \dots$, and $A_i : \Theta \times \Omega \to \mathbb{R}$ be the corresponding attributes they observe, i.e.,

$$A_i(\theta, \omega) := A(X_i(\theta, \omega), \omega), \tag{4}$$

for $\theta \in \Theta$.

The random variables X_i are not necessarily independent since at effort level $i+1$, the agent may use all observations $(X_1, A_1), \dots, (X_i, A_i)$ before they decide on X_{i+1}. Mathematically, this statement implies that the random variable X_{i+1} must be measurable with respect to the σ-algebra $\mathscr{F}_i$ generated by $(X_1, A_1), \dots, (X_i, A_i)$ [13]. In other words, the random process X_i must be *adapted* to the *filtration* $\{\mathscr{F}_i\}_{i \in E}$. The exact nature of this process depends on the beliefs of the principal about the skills of the agent, see Secs. 2.3.1 and 2.3.2 for specific choices corresponding to a naïve and a skillful agent.

In any case, we are now in the position to define mathematically the quality function $Q(e, \theta, \omega)$. It is:

$$Q(e, \theta, \omega) = \max_{1 \le i \le e} A_i(\theta, \omega). \tag{5}$$

Note that this definition does satisfy the two requirements for the quality function that we posed at the beginning of Sec. 2, namely that $Q(e, \theta, \omega)$ is an increasing function of e and that it is bounded above in probability. Furthermore, we are here operating under the assumption that the agent returns the best attribute they have found, i.e., that they are *honest*. Dishonest behavior, e.g., putting a design in the back-pocket for later use, is not modeled in this paper.

2.3.1 A Naïve Agent

The case $\theta =$ "naïve" corresponds to an agent that ignores past experience and simply chooses function evaluations at random. Mathematically,

$$X_i(\theta = \text{"naïve"}) \sim \mathscr{U}(X), \tag{6}$$

for all $i = 1, 2, \dots$, where $\mathscr{U}(X)$ is the uniform distribution over the space of feasible designs X.

2.3.2 A Skillful Agent

The case $\theta =$ "skillful" corresponds to an agent that learns from past experience and queries the function trying to exploit what they have learned. The problem of how individuals acquire new knowledge is known as *human function learning*. Griffiths et al. [9] model human function learning using a GP. Here we follow their approach. In particular, we assume that the agent's prior knowledge about the attribute function $A(x, \omega)$ is captured by the GP prior of Eq. (2). In economic terms, we assume that $A(x, \omega)$ is *common knowledge* for the principal and the agent. It is also possible to model the case in which the agent has *private knowledge*, but this is beyond the scope of this paper.

Now, let $i \in E$ and assume that the agent has already selected i designs,

$$X_{1:i} = (X_1, \dots, X_i), \tag{7}$$

and they have observed the corresponding i attributes:

$$A_{1:i} = (A_1, \dots, A_i). \tag{8}$$

We assume that the agent updates their state of knowledge about $A(x, \omega)$ by using Bayes rule to condition the prior GP of Eq. (2) on the observed data $X_{1:i}$ and $A_{1:i}$. The result is the *posterior GP*:

$$A | X_{1:i}, A_{1:i} \sim \mathrm{GP}(m_i, k_i), \tag{9}$$

where the posterior mean and covariance functions are given by:

$$m_i(x) = m(x) + k(x, X_{1:i}) k^{-1}(X_{1:i}, X_{1:i}) \left(A_{1:i} - m(X_{1:i})\right), \tag{10}$$

and

$$k_i(x,x') = k(x,x') - k(x,X_{1:i})k^{-1}(X_{1:i},X_{1:i})k(X_{1:i},x'), \quad (11)$$

respectively (see Ch. 2 of [8]). In these formulas, we have extended the definition of the mean and the covariance functions so that for any $X^1_{1:i^1}$ and $X^2_{1:i^2}$, $m(X^1_{1:i^1}) = \left(m\left(X^1_1\right), \ldots, m\left(X^1_{i^1}\right)\right)$, and $k\left(X^1_{1:i^1}, X^2_{1:i^2}\right)$ is the $\mathbb{R}^{i^1 \times i^2}$ matrix with (s,t)-element $k(X^1_s, X^2_t)$.

Following the experimental results in [10] and [14], we assume that a "skillful" agent selects X_{i+1} by maximizing the expected improvement in the attribute function. Suppose that the agent made a hypothetical query at $x \in X$ and they observed the attribute value $a \in \mathbb{R}$. The *improvement* they would have gotten over the observed attributes $A_{1:i}$ is

$$I_i(x,a) = \max\left\{0, a - \max_{1\le j\le i} A_j\right\}. \quad (12)$$

The *expected improvement* is obtained by taking the expectation of $I_i(x,a)$ over the agent's beliefs about a as captured by the posterior GP of Eq. (9), i.e.,

$$\mathrm{EI}_i(x) = \int I_i(x,a)\mathcal{N}(a|m_i(x), \sigma_i^2(x))da, \quad (13)$$

where $\sigma_i^2(x) = k_i(x,x)$, and $N(\cdot|\mu,\sigma^2)$ is the probability density function of a Gaussian random variable with mean μ and variance σ^2. It is actually possible to carry out the integration analytically [15] yielding:

$$\mathrm{EI}_i(x) = \left(m_i(x) - \max_{1\le j\le i} A_j\right)\Phi(Z_i(x)) + \sigma_i(x)\,\phi(Z_i(x)), \quad (14)$$

where

$$Z_i(x) = \frac{m_i(x) - \max_{1\le j\le i} A_j}{\sigma_i(x)},$$

and ϕ and Φ are the probability density function (PDF) and the cumulative distribution function (CDF) of standard normal, respectively. Therefore, the information acquisition strategy followed by the agent is assumed to be:

$$X_{i+1}(\theta = \text{"skillful"}) = \arg\max_{x\in X} \mathrm{EI}_i(x). \quad (15)$$

2.4 Constructing a Reduced Order Model

The quality function random process $Q(e,\theta,\omega)$ is not analytically available. Unfortunately, this makes its use in a game-theoretic context, e.g., for the study of Nash equilibria and optimal mechanisms of SEPs, extremely difficult. To remedy the situation, we propose to use samples from $Q(e,\theta,\omega)$ to construct a computationally efficient reduced order model. We outline this proposal below.

Let us consider a particular type of agent, $\theta \in \Theta$. According to the Karhunen-Loève theorem [11], if $Q(e,\theta,\omega)$ is square integrable for all $e \in E$, i.e., if $\mathbb{E}\left[Q^2(e,\theta,\omega)\right] < \infty$, then it admits the following representation:

$$Q(e,\theta,\omega) = Q^0(e,\theta) + \sum_{k=1}^{\infty}\sqrt{\lambda_k(\theta)}\xi_k(\omega)\phi_k(e,\theta), \quad (16)$$

where $Q^0(e,\theta)$ is the mean of the random field, $\lambda_k(\theta)$ and $\phi_k(e;\theta)$ are the eigenvalues and eigenvectors of its covariance function, respectively, and the ξ_k are uncorrelated zero mean and unit variance random variables.

The idea is to estimate all these quantities involved in Eq. (16) samples of the stochastic process $Q(e,\theta,\omega)$. This is achieved through the following algorithm: (i) sample a plausible attribute function from the GP specifying the beliefs of the principal (see Sec. 2.2); (ii) using the sampled attribute function as the underlying truth, simulate the behavior of an agent attempting to maximize it (see Sec. 2.3); and (iii) return the resulting realization of the effort vs quality function (see Eq. (5)). To get a practical model, we truncate Eq. (16) at M terms so that we capture at least 90% of the spectral energy of the random field. Since the effort levels are discrete, KLE is equivalent to principal component analysis (PCA), which we carry out using the Python package scikit-learn [16]. We approximate the probability density function of each ξ_k, i.e., $\mathrm{pdf}(\xi_k)$, using Gaussian mixture model [17].

3 NUMERICAL RESULTS

We run exhaustive numerical simulations in to understand the quality function dependence on task complexity and agent skill. In all simulations, the GP prior, which represents the principal's state of knowledge about the attribute function (Sec. 2.2), has a mean function $m(x) = 2$ and a signal strength $\sigma_s = 1$. We consider four levels of different problem complexity as captured by the covariance length-scale choices $l = 0.005, 0.01, 0.05$, and 0.1, along with the two agent skill levels (naïve and skillful). That is, the total number of cases we simulate is 4 (complexity lvels) $\times 2$ (skill levels) $= 8$ (cases). For each case, we take $50{,}000$ Monte Carlo (MC) samples from the corresponding random field $\{Q(e,\theta,\omega)\}_{1\le e\le 40}$. Using these $50{,}000$ sam-

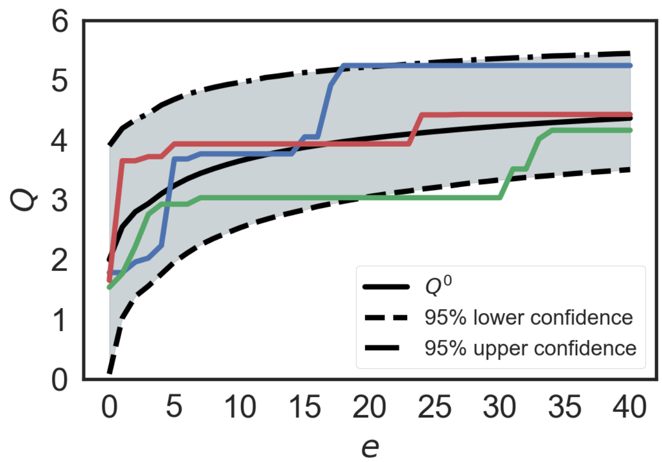

(a) Skillful, $l = 0.005$

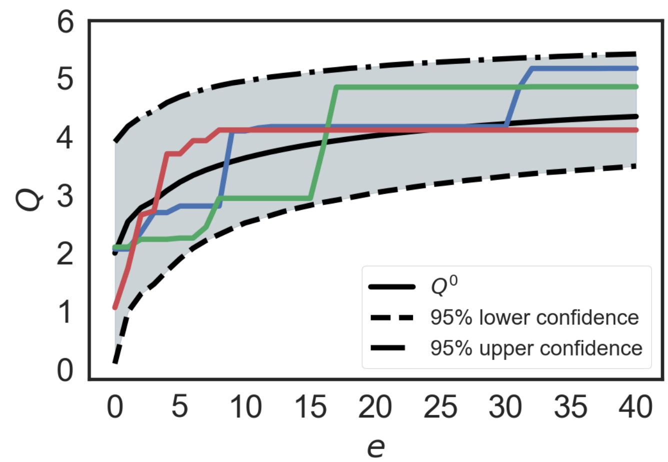

(b) Skillful, $l = 0.01$

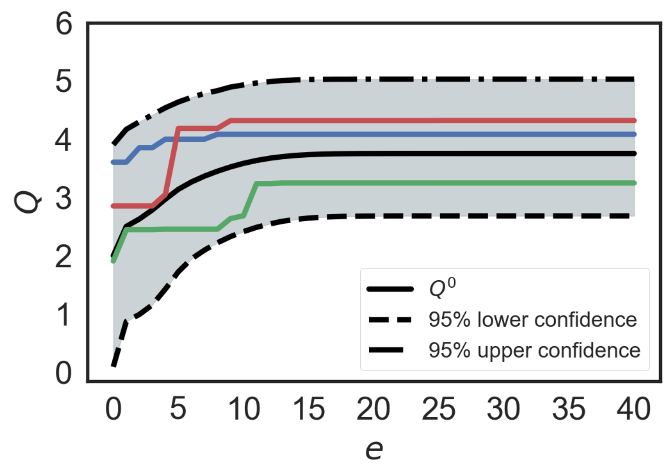

(c) Skillful, $l = 0.05$

FIGURE 1: Three random quality function samples (solid colored lines), the mean of all $50,000$ sampled quality functions (solid black line), and the 95% confidence levels (gray shaded area between the black dashed and dashed-dotted lines) for $\theta = $ skillful.

ples, we construct the reduced order model of Sec. 2.4 which we proceed to compare systematically with them.

In Figs 1 and 2, we depict three random quality function samples (solid colored lines), the mean of all $50,000$ sampled quality functions (solid black line), and the 95% confidence levels (gray shaded area between the black dashed and dashed-dotted lines) for $\theta = $ skillful and $\theta = $ naïve, respectively. The Figs 1a,1b, and 1c, and Figs 2a, 2b, and 2c are associated with decreasing complexity since they correspond to length-scale choices of $l = 0.005, 0.01$, and 0.05, respectively. We observe the following. First, all samples across every case are increasing piecewise constant and bounded from above functions of e. Second, the mean quality function ($Q^0(e,\theta)$) is increasing and concave showing a clear dependence on task complexity and agent skill which we study in the next paragraph. Third, the uncertainty is greater for small efforts, decreases mildly as the effort level increases, albeit it seems to have a definite non-zero lower bound. The latter is also discussed below.

Fig 3 depicts the mean quality function ($Q^0(e,\theta)$) for all 8 cases. First note that the upper bound to the $Q^0(e,\theta)$ increases with decreasing length-scale. This phenomenon is related to the distribution of the maximum of Gaussian random fields on compact domains [18]. Namely, the expectation of the maximum of a Gaussian random field on a compact domain increases as the length-scale decreases. To understand this phenomenon intuitively, take into account that the samples from GP priors with smaller length-scales have more opportunity to wiggle around the compact domain and reach extreme values. Therefore, com-

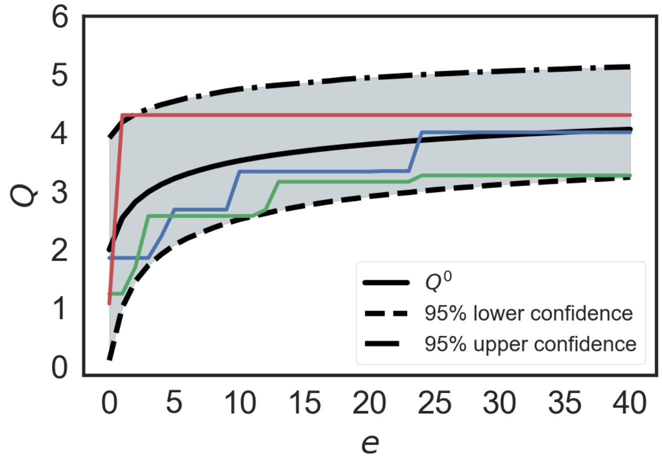

(a) Naïve, $l = 0.005$

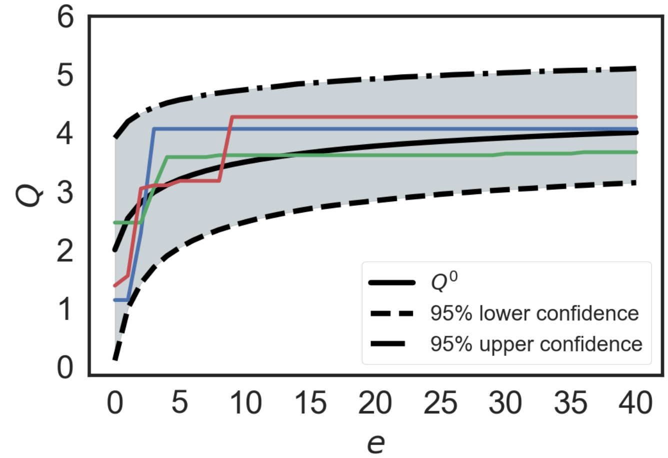

(b) Naïve, $l = 0.01$

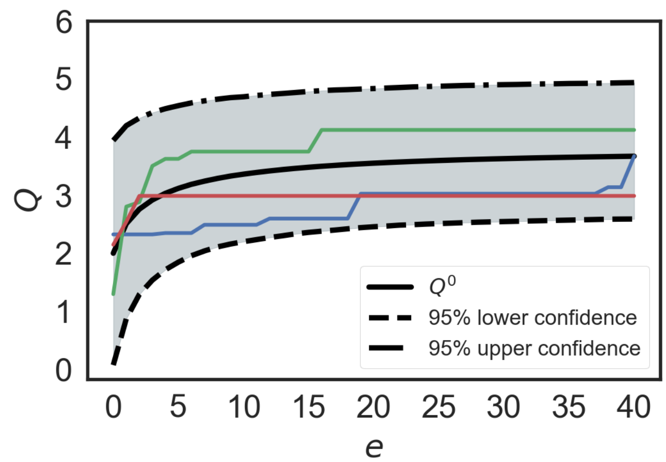

(c) Naïve, $l = 0.05$

FIGURE 2: Three random quality function samples (solid colored lines), the mean of all 50,000 sampled quality functions (solid black line), and the 95% confidence levels (gray shaded area between the black dashed and dashed-dotted lines) for $\theta = $ naïve.

paring the absolute values of $Q^0(e, \theta)$ across different length-scales is nonsensical. What is comparable across complexities is the amount of effort required to exceed a certain percentage of the maximum quality, e.g., the first effort level $e^* = e^*(\varepsilon)$ for which $Q^0(e^*(\varepsilon), \theta)/\sup_{e\in E} Q^0(e, \theta) > 1 - \varepsilon$ to become flat for some $\varepsilon > 0$. Naturally, in Fig 3 we observe that the maximum is reached later as complexity is increased. However, the mean quality function is directly comparable across agent skill levels. Comparing Fig. 3 (a) with (b), we observe that naïve agents require more effort to reach the same quality for the same complexity level.

Now, we focus on the reduced order model of Sec. 2.4. To construct it, we perform PCA on 80% of the MC samples. The remaining 20% of the MC samples are used for validation. We present our results for length-scales $l = 0.01$ and 0.05, which need 5 and 3 terms in the KLE to capture the 90% of spectral energy, respectively.

In Figs 4 and 5, we show the PDFs of retained ξ_k's for skillful and naïve agents, respectively. In all cases, the first KLE component ξ_1 is almost a perfect standard normal. However, for the higher order terms we start observing distinct non-Gaussian features. Also, the PDFs of the first three ξ_k's, for both length-scales and skill levels, are identical. It is only after the 4-th KLE component that we start observing differences in the PDFs.

In Fig 6, we show the eigenvalues (λ) and eigenfunctions (ϕ) of the reduced order model for the two length-scales. As expected, the eigenvalues decay faster for decreasing complexity. However, we do not observe any significant differences across

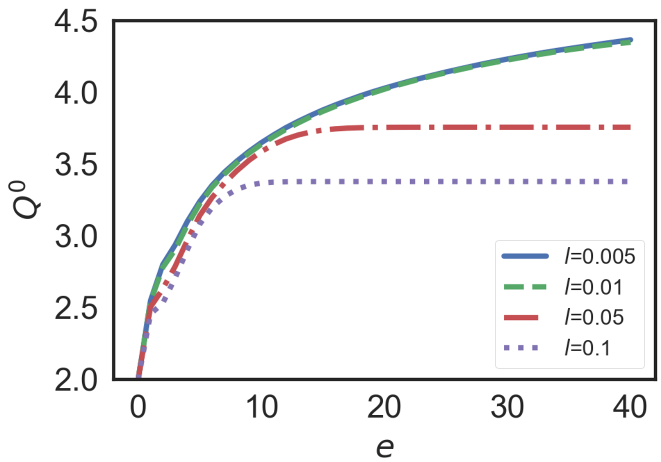

(a) Skillful

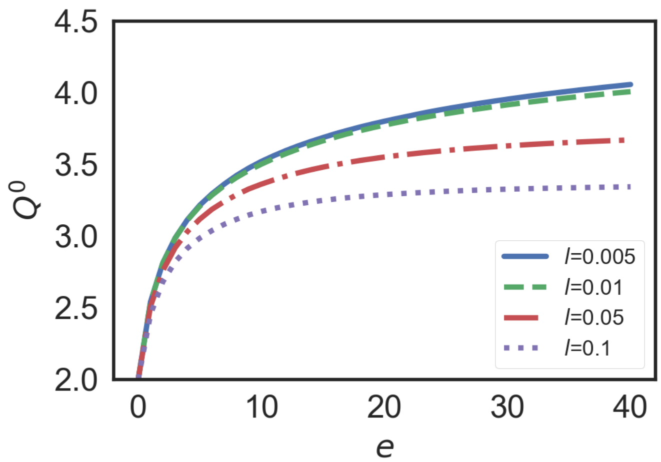

(b) Naïve

FIGURE 3: The effect of problem complexity on the mean of quality function ($Q^0(e,\theta)$) for skillful and naïve agents.

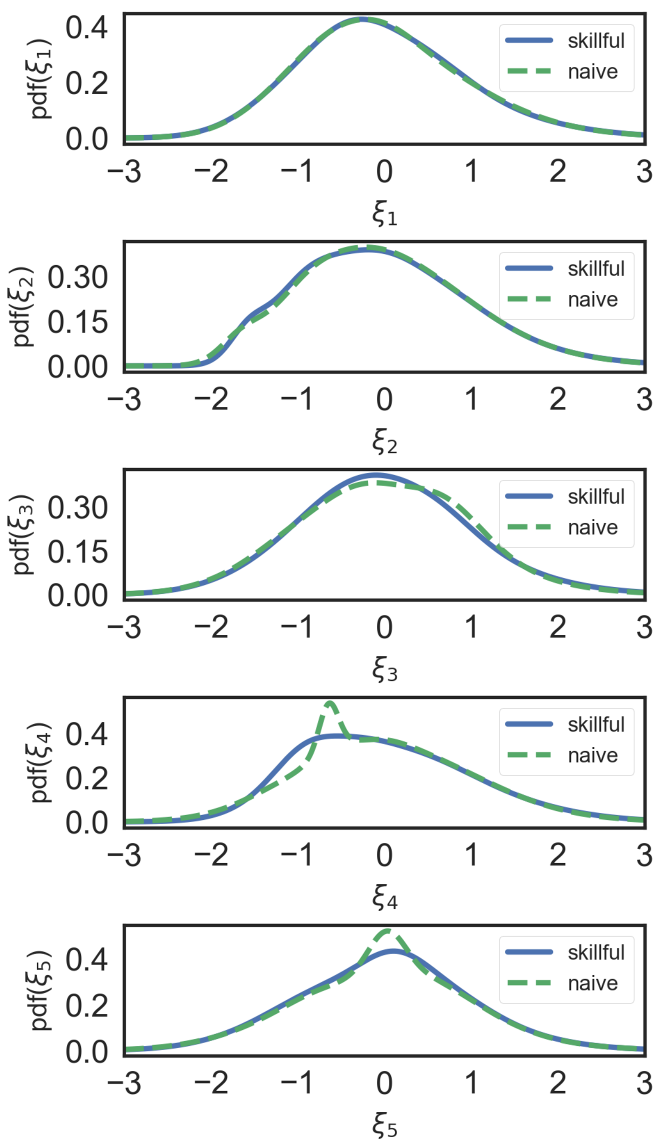

FIGURE 4: The PDF of random variables $\xi_1,\cdots,\xi_5$ of reduced order model for skillful and naïve agents with $l=0.01$.

skills. The eigenfunctions (especially at lower orders) seem almost independent of skill, but the higher order ones do exhibit a small variation as complexity changes. We outline and interpret intuitively these findings below.

The first eigenfunction for both agent types and all complexity levels is almost constant. That is, the first eigenfunction just adds a constant to the mean quality function. Therefore, the first eigenfunction captures the uncertainty in the maximum of the underlying attribute function. Furthermore, taking into account that the PDF of ξ_1 is almost a perfect standard normal, we see that the assumption of additive Gaussian noise is valid to first order.

The higher order eigenfunctions are non-constant. However, note that they have a bump at small efforts, but they converge to zero as the level of effort increases. This bump is an indication that these eigenfunctions capture uncertainties associated with the agent's search process. Furthermore, the higher the order of the eigenfunction, the more effort is needed for the bump to appear. That is, the high order eigenfunctions capture uncertainties in later stages of the search process. Consistent with this intuitive interpretation, notice that the eigenfunction bumps move to the left as complexity decreases. This is a reflection of the fact that in less complex tasks critical discoveries occur earlier.

Finally, in Fig 7, we compare the distribution of the $Q(e,\theta,\omega)$ at effort levels $e=20$ and 40 of reduced order model

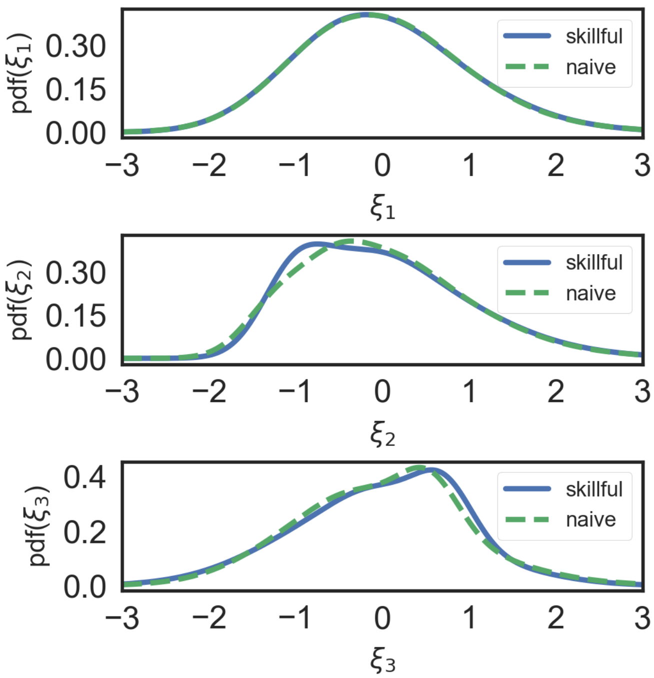

FIGURE 5: The PDF of random variables $\xi_1, \cdots, \xi_3$ of reduced order model for skillful and naïve agents with $l = 0.05$.

with those of the 20% MC samples that we set aside. The results obtained with the reduced order model match closely those obtained with the MC samples. As expected, the reduced order model slightly underestimates the variance. In Fig 8, we show the $Q^0(e, \theta)$ and 95% lower and upper confidence levels of $Q(e, \theta, \omega)$ alongside some sample functions with the reduced order random model, with $l = 0.01$ and 0.05. Our results matches with those from MC samples shown in Figs 1 and 2, albeit information about the "high frequency" behavior of the agent has been coarse grained.

4 CONCLUSIONS

We modeled the quality function of a leaf agent in the early design stages of the SEP hierarchy as a stochastic process. Our approach relies on the assumption that the design task assigned to an agent can be modeled as a scalar maximization problem. We explicitly captured the principal's beliefs about the task complexity and problem-solving skills of an agent. We studied two types of agents, a skillful agent who follows the Bayesian optimization algorithm in solving the maximization problem, and a naïve agent who searches randomly in the design space to solve the maximization problem. Finally, we constructed a reduced order model based on the KLE of the quality function that can be used in an extensive game-theoretic model of the SEP.

Note that our model does not aspire to predict the problem-solving behavior of real designers. Rather, we are investigating the mathematical implications that different information acquisition strategies have in the form of the quality function. However, assuming that real designers do maximize a well-defined mathematical function, we can conclude that the quality function corresponding to a naïve/skillful agent provides a lower/upper bound to the quality that should be expected by a real person lacking any domain-specific knowledge. In all other cases, the model is only a crude approximation of real agent behavior which may, nevertheless, be adequate for posing and solving the mechanism design problem.

We found that the common assumption that the quality function is linear with additive Gaussian noise is insufficient. We showed that, the quality function is an increasing concave function and that the derivative and curvature depend on the problem complexity and problem-solving skills of the agent. The derivative of this function is large at early stages of the effort, and it becomes smaller as the effort level increases for both skillful and naïve agents. The derivative at early stages of effort is lower for a naïve agent than that for a skillful agent. We also saw that the eigenfunctions of the reduced order model can be interpreted in the following way. The first eigenfunction is a constant that captures the uncertainty about the maximum possible quality value. The higher order eigenfunctions capture the uncertainty about the search process. We demonstrated that the statistical properties of the reduced order model match those of the MC samples of the full-blown stochastic process closely, albeit the fine details are coarse-grained. Therefore, we conclude that one may use the reduced order model in a principal-agent representation of the SEP.

There remain several open questions. First, we assumed that the agent starts the maximization problem from scratch. Usually, a designer may already know a lot about the attribute function and this information may or may not be available to the principal. However, we anticipate that our framework is easily adjusted to this case. Second, we assumed that the agent has only one information source, i.e., that there is no alternative way to gain information about the attribute. Such alternative sources of information could be simulations of varying complexity or building prototypes. Third, our account of complexity is quite restrictive as we have not covered all the possibilities such as varying design dimensions, covariance smoothness, and the existence of discrete choices. Finally, we did not discuss the quality function associated with design tasks requiring the discovery of the Pareto efficient frontier of multi-objective optimization problems. All these open questions are the subjects of ongoing research.

ACKNOWLEDGMENT

This work was supported by the National Science Foundation under Grant No. 1728165.

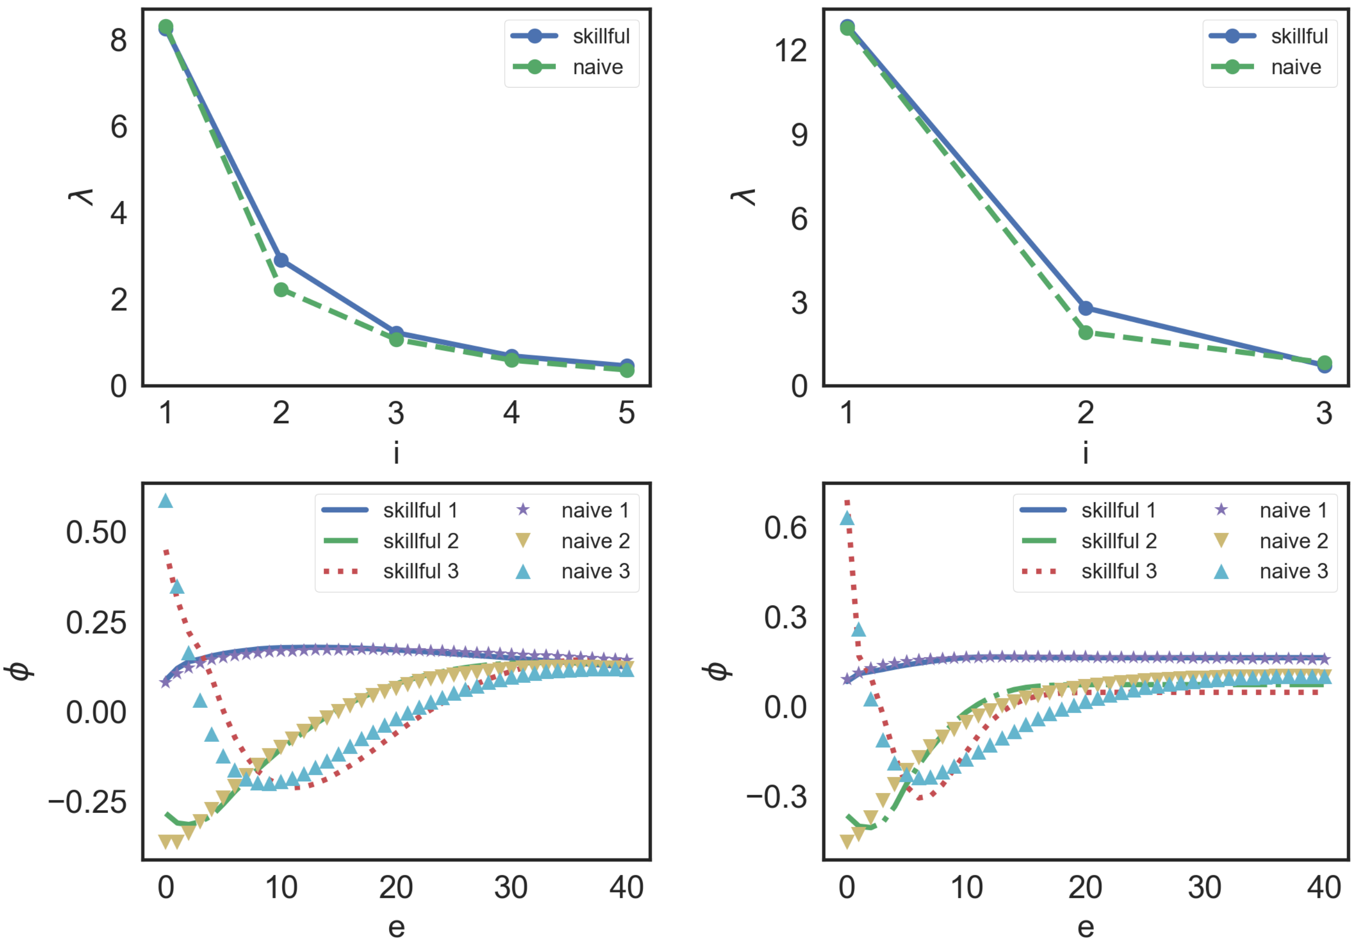

(a) $l = 0.01$ (b) $l = 0.05$

FIGURE 6: The eigenvalues (λ) and eigenfunction (ϕ) of the reduced order model that capture more than 90% of spectral energy of the random field for skillful and naïve agents with $l = 0.01$ and 0.05. Note, only first 3 eigenfunctions out of 5 are shown for $l = 0.01$.

REFERENCES

[1] Council on Systems Engineering, I., 2015. *INCOSE Systems Engineering Handbook: A Guide for System Life Cycle Processes and Activities*. 06.

[2] Nasa, Aeronautics, N., and Administration, S., 2008. *NASA Systems Engineering Handbook*. U.S. Government Printing Office.

[3] Maddox, I., Collopy, P., and Farrington, P. A., 2013. "Value-based assessment of dod acquisitions programs". *Procedia Computer Science,* ***16***, pp. 1161 – 1169. 2013 Conference on Systems Engineering Research.

[4] Shoham, Y., and Leyton-Brown, K., 2008. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, New York, NY, USA.

[5] Vermillion, S. D., and Malak, R. J., 2015. "Using a principal-agent model to investigate delegation in systems engineering". pp. V01BT02A046–.

[6] Fudenberg, D., and Tirole, J., 1991. *Game Theory*. MIT Press.

[7] Gibbons, Robert, ., [1992]. *Game theory for applied economists*. Princeton, N.J. : Princeton University Press, [1992] ©1992.

[8] Rasmussen, C. E., and Williams, C. K. I., 2005. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press.

[9] Griffiths, T. L., Lucas, C., Williams, J., and Kalish, M. L., 2009. "Modeling human function learning with gaussian processes". In *Advances in Neural Information Processing Systems 21*, D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, eds. Curran Associates, Inc., pp. 553–560.

[10] Borji, A., and Itti, L., 2013. "Bayesian optimization explains human active search". In *Advances in Neural Infor-*

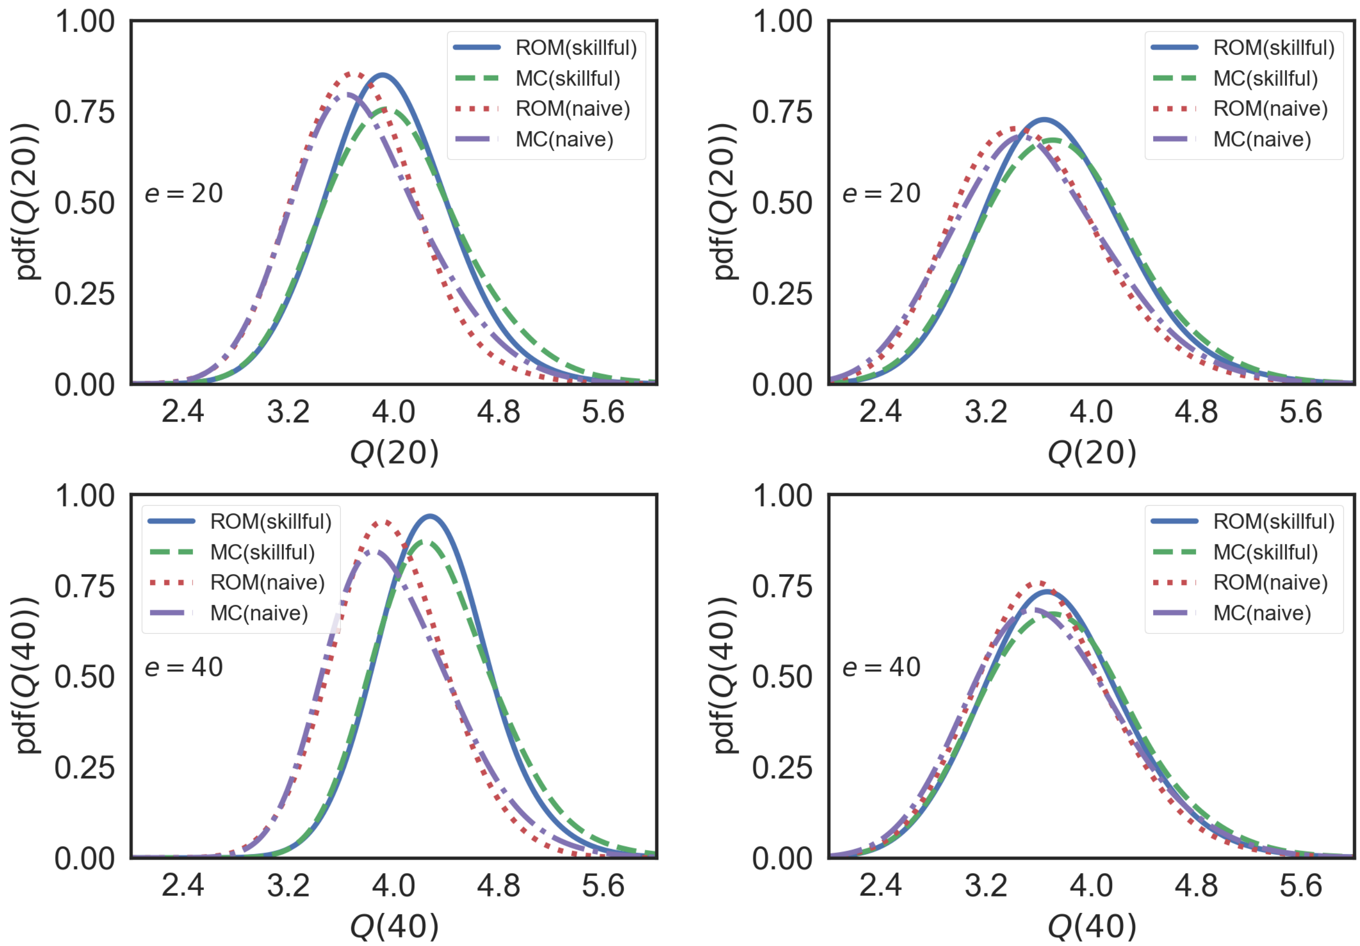

(a) $l = 0.01$ (b) $l = 0.05$

FIGURE 7: The distribution of the quality at effort levels $e = 20$ and 40 with reduced order model (ROM) and Monte Carlo (MC) samples for skillful and naïve agents with $l = 0.01$ and 0.05.

mation Processing Systems 26, C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, eds. Curran Associates, Inc., pp. 55–63.

[11] Ghanem, R., and Spanos, P., 2003. *Stochastic Finite Elements: A Spectral Approach*. Civil, Mechanical and Other Engineering Series. Dover Publications.

[12] Adler, R., 2010. *The Geometry of Random Fields*. Classics in Applied Mathematics. Society for Industrial and Applied Mathematics, Jan.

[13] Øksendal, B., 2003. *Stochastic Differential Equations: An Introduction with Applications*. Hochschultext / Universitext. Springer.

[14] Panchal, J. H., Sha, Z., and Kannan, K. N., 2017. "Understanding design decisions under competition using games with information acquisition and a behavioral experiment". *Journal of Mechanical Design*, ***139***(9), 07, pp. 091402–091402–12.

[15] Brochu, E., Cora, V. M., and de Freitas, N., 2010. "A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning". *CoRR*, ***abs/1012.2599***.

[16] Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E., 2011. "Scikit-learn: Machine learning in Python". *Journal of Machine Learning Research*, ***12***, pp. 2825–2830.

[17] Bishop, C. M., 2006. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA.

[18] Azas, J.-M., and Wschebor, M., 2008. "A general expression for the distribution of the maximum of a gaussian field

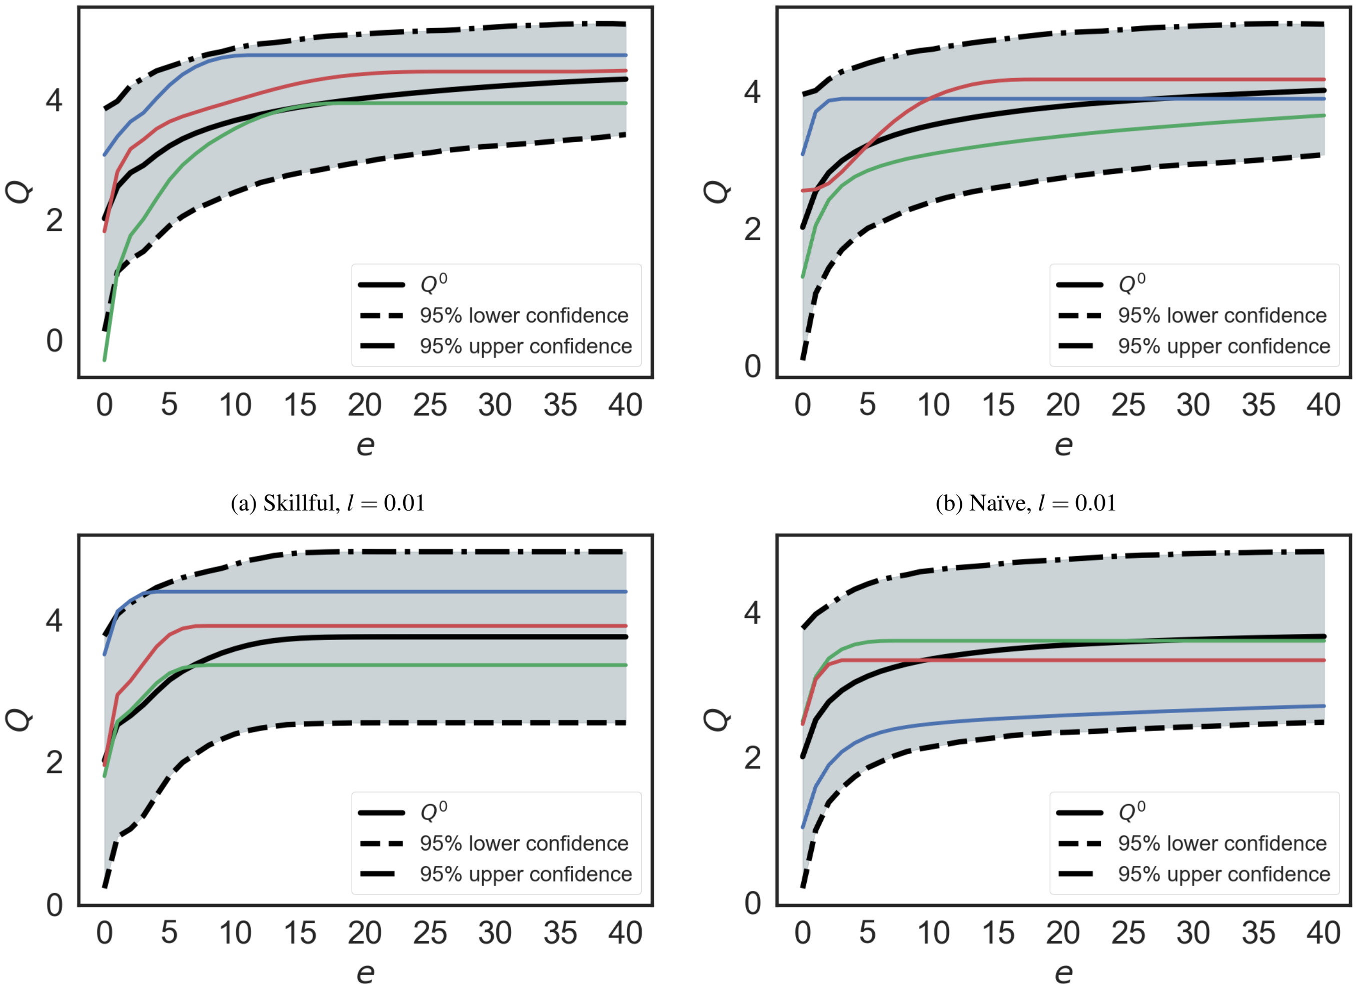

(a) Skillful, $l = 0.01$

(b) Naïve, $l = 0.01$

(c) Skillful, $l = 0.05$

(d) Naïve, $l = 0.05$

FIGURE 8: Three random quality function samples (solid colored lines), the mean of all 10,000 sampled quality functions from reduced order random model (solid black line), and the 95% confidence levels (gray shaded area between the black dashed and dashed-dotted lines). The first and the second columns correspond to $\theta =$ skillful and $\theta =$ naïve, respectively.

and the approximation of the tail". *Stochastic Processes and their Applications,* ***118***(7), pp. 1190 – 1218.