

The Cost of Informing Decision-Makers in Multi-Agent Maximum Coverage Problems With Random Resource Values

Bryce L. Ferguson¹, Student Member, IEEE, Dario Paccagnan², Member, IEEE,
and Jason R. Marden¹, Senior Member, IEEE

Abstract—The emergent behavior of a distributed system is conditioned by the information available to the local decision-makers. Therefore, one may expect that providing decision-makers with more information will improve system performance; in this letter, we find that this is not necessarily the case. In multi-agent maximum coverage problems, we find that even when agents' objectives are aligned with the global welfare, informing agents about the realization of the resource's random values can reduce equilibrium performance by a factor of 1/2. This affirms an important aspect of designing distributed systems: information need be shared carefully. We further this understanding by providing lower and upper bounds on the ratio of system welfare when information is (fully or partially) revealed and when it is not, termed the value-of-informing. We then identify a trade-off that emerges when optimizing the performance of the best-case and worst-case equilibrium.

Index Terms—Game theory, agent-based systems, uncertain systems, distributed algorithms.

I. INTRODUCTION

IN LARGE-SCALE systems, the prospect of distributing decision-making to local entities is becoming increasingly enticing as a method to reduce complexity while maintaining some level of performance. This can take the form of swarm control for robotic fleets [1], autonomous driving decisions in mobility services [2], local task assignment decisions [3], and many more. Taking a distributed approach entails assigning each agent a decision-making algorithm, such as maximizing an assigned local objective function [4], then analyzing the

Manuscript received 17 March 2023; revised 30 May 2023; accepted 21 June 2023. Date of publication 28 June 2023; date of current version 19 July 2023. This work was supported in part by the Office of Naval Research (ONR) under Grant N00014-20-1-2359, and in part by the Air Force Office of Scientific Research (AFOSR) under Grant FA95550-20-1-0054 and Grant FA9550-21-1-0203. Recommended by Senior Editor M. Guay. (Corresponding author: Bryce L. Ferguson.)

Bryce L. Ferguson and Jason R. Marden are with the Department of Electrical and Computer Engineering, University of California at Santa Barbara, Santa Barbara, CA 93106 USA (e-mail: blferguson@ece.ucsb.edu; jrmarden@ece.ucsb.edu).

Dario Paccagnan is with the Department of Computing, Imperial College London, SW7 2AZ London, U.K. (e-mail: d.paccagnan@imperial.ac.uk).

Digital Object Identifier 10.1109/LCSYS.2023.3290522

system's equilibria [5]. As agents need not possess full knowledge of the overall system; the local decisions (and, ultimately, global behavior) are dependent on the information communicated with and between the agents [6], [7]. In this letter, we address how available information affects system performance.

We focus on maximum coverage problems: a class of models in which each agent selects a set of resources from a ground set, with the objective of maximizing the value of covered resources. To solve this in a distributed fashion, each agent is given a utility function to evaluate what set of resources to cover; this forms a game played by the agents with their resource selection as their action and the evaluation of their assigned utility function as their payoff. Existing work has focused on how to design these utility rules and how well the resulting equilibria of the emergent game approximate the optimal welfare [8], [9], [10]. In this letter, we generalize this model to consider the case where agents have uncertainty about the resources' values. In this setting, we ask how revealing information to local decision-makers affects equilibrium performance. Interestingly, we find that *revealing truthful information about the system state can worsen system performance*. While this phenomenon has been observed before in social systems [11], here we find that similar conclusions hold even when the local decision-makers' objectives are *aligned* with the global welfare.

To study this, we consider a *Bayesian persuasion* framework, in which a well-informed system operator can strategically reveal information to agents using messages (or signals) which contain partial information [12]. In this letter, we study how this information revealing affects equilibrium welfare. To this end, we introduce a new performance metric termed the *value-of-informing*, which measures the ratio between the equilibrium welfare under an information-revealing policy and when no information is revealed. This measures the gain or loss in welfare from revealing information.

The framework of Bayesian persuasion has gained traction in the areas of economics, operations research, and engineering, but typically for settings concerned with the behavior and beliefs of human users (e.g., traffic routing apps [13], [14], [15], pricing/investing decisions [16], hybrid work policies [17], etc.). Results are typically restricted to a binary classification on whether revealing full information helps or

not [12], [18], or methods to compute optimal information revealing policies in limited settings [15], [19], [20], often with no guarantee on the magnitude of improvement. Here, we adapt the ideas of information provisioning to the setting of engineered systems, where designed decision-making components can improve their estimate of the system state by receiving relevant messages (e.g., a fleet of surveillance drones receiving live map updates).

Revealing information to local decision-makers can obviously improve the welfare of emergent system equilibria; however, in this letter, we find that this is not always the case. In fact, system performance may degrade by a factor of 1/2 when revealing truthful information to local decision-makers. The possible loss in system welfare from informing decision-makers comes from two sources (1) the multiplicity of equilibria and (2) the local objective assigned to each decision-maker. We study the aforementioned value-of-informing when considering the best-case and worst-case equilibria for different local objectives. Ultimately, we highlight a trade-off between the possible loss from revealing information to best-case and worst-case equilibrium guarantees when agents' local objectives are designed.

II. PROBLEM FORMULATION

A. Maximum Coverage Problem

Maximum coverage problems have been used to model resource allocation, sensor coverage, job scheduling, and more [21]. Consider the multi-agent maximum coverage problem, in which $\mathcal{R} = \{1, \dots, R\}$ is a finite set of resources. For each resource $r \in \mathcal{R}$, let $v_r \geq 0$ be the value of that resource; further, let $v \in \mathbb{R}_{\geq 0}^{|\mathcal{R}|}$ be the vector containing each resource value. Let $N = \{1, \dots, n\}$ be a set of agents, where each agent $i \in N$ can be assigned to cover a subset of resources $a_i \subseteq \mathcal{R}$. The set of allowable assignments for each agent is defined by $\mathcal{A}_i \subseteq 2^{\mathcal{R}}$. When each agent is assigned, an allocation of agents is denoted $a = (a_1, \dots, a_n) \in \mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_n$. Let (G, v) define a maximum coverage problem where $G = (N, \mathcal{R}, \mathcal{A})$.

In an allocation a , the system welfare is equal to the total value of resources covered by at least one agent, i.e., $W(a; v) = \sum_{r \in \cup_{i \in N} a_i} v_r$. However, finding an optimal allocation $a^{\text{opt}} \in \arg \max_{a \in \mathcal{A}} W(a; v)$ is NP-hard [9]. It is for this reason, we consider a distributed solution technique to approximate this optimal solution.

B. Distributed Decision Making

Let each agent $i \in N$ possess a local objective function

$$U_i(a; v) = \sum_{r \in a_i} v_r f(|a|_r),$$

which depends on their own action and the actions of every other agent. This local objective, or *utility function*, is parameterized by a local utility rule $f: \mathbb{N} \rightarrow \mathbb{R}$ that takes as argument $|a|_r$, the number of agents covering resource r in allocation a . The system operator can adopt a utility rule f without exact knowledge of the problem instance if needed.

When agents sequentially and repeatedly update their assignment to one that currently maximizes their utility




Fig. 1. Depiction of information signaling in maximum coverage problems. On the left is the support of a random state variable v . At right is a maximum coverage problem, to which we have assigned each agent a local objective. The agents in the coverage problem possess the prior distribution of the unknown state variable and receive some partial information π_k about the realization. The manner in which information is revealed will alter how agents evaluate their objectives and change the emergent behavior.

function, these best-response dynamics have fixed points that are the Nash equilibria of the underlying game. The convergence of best response dynamics in potential games to pure Nash equilibria are addressed in [5], [9]. Nash equilibrium allocations can be defined by

$$U_i(a^{\text{NE}}; v) \geq U_i(a'_i, a_{-i}^{\text{NE}}; v) \quad \forall a'_i \in \mathcal{A}_i, \quad i \in N, \quad (1)$$

where a_{-i} denotes the allocation of all agents but player i . Let $\text{NE}(G, v, f)$ denote the set of states satisfying (1). These states represent the possible solutions of the distributed dynamics, and we will consider their welfare as an approximation of the maximum coverage problem.

C. Uncertainty and Information Signaling

In this letter, we consider how uncertainty and information can affect the efficacy of distributed decision-making. We consider this uncertainty in the form of randomness for the resources values. Let $v \in \mathbb{R}_{\geq 0}^{|\mathcal{R}|}$ (the vector containing the value of each resource $r \in \mathcal{R}$) be a discrete random variable with prior distribution μ_0 . A realization of v determines each resource's value $[v_1, \dots, v_{|\mathcal{R}|}] = v$, i.e., the resource values may be correlated. Let the support of v be $\mathcal{V} := \text{supp}(v)$.

First, we consider the case where agents are *uninformed* about the system state, i.e., they know the prior distribution μ_0 but not the exact realization of v . In this setting, the agents optimize their expected utility, which we will denote $\bar{U}_i(a; \mu_0) = \mathbb{E}_{v \sim \mu_0} [\sum_{r \in a_i} v_r f(|a|_r)]$. Let $\text{NE}(G, \mu_0, f) = \text{NE}(G, \mathbb{E}_{v \sim \mu_0} [v], f)$. Additionally, the objective of the maximum coverage problem is to maximize the expected welfare, i.e., $\bar{W}(a; \mu_0) = \mathbb{E}_{\mu_0} [\sum_{r \in \cup_{i \in N} a_i} v_r]$.

As a means to try and improve the performance of the distributed decision-making agents, we may consider revealing information about the realization of the system state. One option is to reveal full information (or let agents know the realization exactly); however, either due to communication constraints or by design choice, it is often meaningful to reveal only partial information as well. In line with the broader Bayesian Persuasion framework, consider revealing information with an *information signaling policy* $\Pi = \{\pi_1, \dots, \pi_m\}$, where $\pi_k \subseteq \mathcal{V}$, $\pi_j \cap \pi_k = \emptyset$, and $\bigcup_{k=1}^m \pi_k = \mathcal{V}$. This signaling policy Π forms a partition over the support of our random state variable v . The signal

$\pi \in \Pi$ is revealed to the agents when $v \in \pi$.¹ The new system operates as follows: A system operator adopts a signaling policy Π and utility rule f . A state v is drawn from μ_0 , and all agents are informed of which element of Π the realization belongs to. The corresponding signal π is sent, and each player $i \in N$ computes the posterior belief on the realization $\mu_\pi(x) = \mathbb{P}[v = x | v \in \pi] = \mu_0(x) / (\sum_{v' \in \pi} \mu_0(v'))$ if $v \in \pi$ and zero otherwise. The agents then seek to maximize their expected utility with the posterior belief.

Agents may now condition their action on the received signal. Let $\alpha \in \mathcal{A}^\Pi$ denote a joint strategy, where an element $\alpha_i(\pi) \in \mathcal{A}$ captures the action agent i takes when they receive signal π . In a strategy profile α , agent i has an expected payoff of $\bar{U}_i(\alpha; \mu_0, \Pi) = \mathbb{E}_{v \sim \mu_0}[\sum_{r \in \alpha_i} v_r f(|\alpha|_r)]$. Note that α_i is implicitly a function of the received signal π , which itself is determined by the state variable v ; as such, each of α_i , π , and v are random variables. The expected welfare becomes $\bar{W}(\alpha; \mu_0, \Pi) = \mathbb{E}_{v \sim \mu_0}[\sum_{r \in \cup_{i \in N} \alpha_i} v_r]$.

When agents follow best-response dynamics, the set of fixed points becomes the set of *Bayes-Nash equilibria*, $\text{BNE}(G, \mu_0, \Pi, f)$. A strategy α^{BNE} in this set satisfies

$$\bar{U}_i(\alpha^{\text{BNE}}; \mu_0, \Pi) \geq \bar{U}_i(\alpha'_i, \alpha_{-i}^{\text{BNE}}; \mu_0, \Pi) \quad \forall \alpha'_i \in \mathcal{A}_i^\Pi. \quad (2)$$

The signaling policy Π will alter this equilibria set and thus the guarantees of our distributed solution to the maximum coverage problem.

Our motivation for this letter is understanding how giving agents information can affect system welfare. Due to the multiplicity of equilibria, we consider two perspectives: the *optimistic perspective* in which the system designer cares about the best attainable system performance, and the *pessimistic perspective* in which the system designer cares about the worst possible performance. For the optimistic perspective, the system designer evaluates a signaling policy Π by its effect on the system welfare in the best-case equilibrium; as such, let the optimistic *value-of-informing* with signaling policy Π be

$$\text{VoI}^+(G, \mu_0, \Pi, f) = \frac{\max_{\alpha \in \text{BNE}(G, \mu_0, \Pi, f)} W(\alpha; \mu_0, \Pi)}{\max_{a \in \text{NE}(G, \mu_0, f)} W(a; \mu_0)},$$

which measures the gain in optimistic welfare by using policy Π . Similarly, for the pessimistic perspective, the system operator evaluates a signaling policy Π by its effect on the system welfare in the worst-case equilibrium; let the pessimistic value-of-informing be

$$\text{VoI}^-(G, \mu_0, \Pi, f) = \frac{\min_{\alpha \in \text{BNE}(G, \mu_0, \Pi, f)} W(\alpha; \mu_0, \Pi)}{\min_{a \in \text{NE}(G, \mu_0, f)} W(a; \mu_0)},$$

which is the same ratio but now with the worst-case equilibrium strategy and allocation. These values inform a system operator of how revealing information will affect their equilibrium guarantees. They differ from the well-known price-of-anarchy/stability measures in that they relate two equilibrium performances rather than equilibrium to optimal.

¹In this letter, we consider signaling policies that are deterministic mappings from state to signal. In general, signal π could be drawn randomly based on state v . Many of the results easily generalize to this setting, but for ease of exposition and relevance to our problem setting of informing designed decision-makers, we present the results for deterministic signaling by treating Π as a partition of \mathcal{V} .

III. MAIN RESULTS

The main contribution of this letter is in lower and upper bounding the value-of-informing for best- and worst-case equilibria. These bounds depend on agents' local decision-making process. In Section III-A, we focus on the case where agents' payoffs are aligned with the system objective and find that revealing information improves the best-case equilibrium ($\text{VoI}^+ \geq 1$) but can worsen the worst-case equilibrium ($\text{VoI}^- \leq 1$). In Section III-B, we generalize these bounds to any local utility rule f . In Section III-C, we observe a trade-off in the lower bounds on VoI^+ and VoI^- ; Fig. 2 characterizes this trade-off and highlights the fact that altering the local objectives of agents affects the efficacy of revealing information.

A. Marginal Contribution

The first utility design we will consider is the marginal contribution, where each agent makes decisions that maximize their contribution to the system welfare, i.e., $U_i(a) = W(a) - W(a_{-i})$. This can be expressed by the local utility rule $f^{\text{mc}}(x) := \mathbb{1}[x = 1]$. When agents follow this utility rule, their preferences are aligned with global welfare. This utility rule has the property that it maximizes the best-case equilibrium guarantee, known as the price-of-stability ratio [10]. However, as of yet, the effect of revealing information to decision-makers using this utility function has not been addressed. In Theorem 1, we address this question by providing lower and upper bounds on the value-of-informing for the best- and worst-case equilibria.

Theorem 1: In a Bayesian Maximum Coverage problem (G, μ_0, Π) , with utility rule f^{mc} , the value-of-informing for the best-case equilibrium satisfies

$$1 \leq \text{VoI}^+(G, \mu_0, \Pi, f^{\text{mc}}) \leq |\Pi|, \quad (3a)$$

and for the worst-case equilibrium satisfies

$$1/2 \leq \text{VoI}^-(G, \mu_0, \Pi, f^{\text{mc}}) \leq 2|\Pi|. \quad (3b)$$

All of these bounds are tight, but the upper bounds on VoI^- .

Before proving the statement, we discuss the consequences of these results. We first see that revealing more information (increasing the cardinality of Π) provides significant opportunities for improvement in either perspective.² However, revealing information need not always have such a positive effect. In the optimistic perspective, revealing information can only improve performance ($\text{VoI}^+ \geq 1$); however, doing so does not come without consequence, as revealing information can reduce the quality of the worst-case equilibrium by a factor of 1/2 ($\text{VoI}^- = 1/2$). This fact affirms an important property of information in a multi-agent system: revealing information must be done carefully. The proof relies on the following lemma characterizing Bayes-Nash equilibria and the expected welfare.

Lemma 1: A joint strategy α is a Bayes-Nash equilibrium if and only if $(\alpha_1(\pi), \dots, \alpha_n(\pi)) \in \text{NE}(G, \mathbb{E}[v | \pi_k], f)$ for each $\pi \in \Pi$. Additionally, the expected welfare of a joint

²A more refined understanding of this improvement can be attained by considering the difference between the realizations and the prior mean.

strategy α is a weighted average of the welfare of the joint actions $\alpha(\pi)$ in the respective deterministic games, i.e.,

$$\bar{W}(\alpha; \mu_0, \Pi) = \sum_{k=1}^m p_k W(\alpha(\pi_k); \mathbb{E}[v|\pi_k]),$$

where $p_k = \sum_{v \in \pi_k} \mu_0(v)$.

Proof: Let $\alpha \in \mathcal{A}^\Pi$ denote a joint strategy. We show the first claim by observing the following transformation for any α ,

$$\bar{U}_i(\alpha; \mu_0, \Pi) = \mathbb{E} \left[\sum_{r \in \mathcal{R}} \mathbb{E}[v_r|\pi] f(|\alpha(\pi)|_r) \right] \quad (4a)$$

$$= \sum_{k=1}^m p_k U_i(\alpha(\pi_k); \mathbb{E}[v|\pi_k]), \quad (4b)$$

where (4a) holds from the law of total expectation. Because α can be any m -tuple of joint actions, (2) is satisfied if and only if $U_i(\alpha(\pi); \mathbb{E}[v|\pi]) \geq U_i(\alpha'_i, \alpha_{-i}(\pi); \mathbb{E}[v|\pi])$, for all $\alpha'_i \in \mathcal{A}_i$, $\pi \in \Pi$; or, that $\alpha(\pi)$ is a Nash equilibrium for the deterministic game G with values $\mathbb{E}[v|\pi]$ for each $\pi \in \Pi$. The second claim follows (4) with welfare in place of utility. ■

Proof of Theorem 1: Best-case equilibrium - We will make use of the function $W^*(v)$, which denotes the welfare of an optimal allocation in G when the values are v . We note that with the marginal contribution utility rule, each optimal allocation a^{opt} is an equilibrium; thus, the welfare of the best Nash equilibrium is the optimal welfare [10], or $W^*(v) = \max_{a \in \mathcal{A}} W(a; v) = \max_{a \in \text{NE}(G, v, f^{\text{mc}})} W(a^{\text{NE}})$. We first make several observations about the function W^* . Observe that $W^*(v) = \max_{a \in \mathcal{A}} \sum_{r \in \mathcal{R}} v_r \mathbb{1}[|a|_r > 0]$, is the point-wise maximum of a set of affine (and thus convex) functions of v , which is itself convex. Further, W^* is positively homogeneous, i.e., $W^*(\lambda v) = \lambda W^*(v)$ for all $\lambda \geq 0$ and $v \geq 0$, and W^* is monotone in v , i.e., $v \geq v' \Rightarrow W^*(v) \geq W^*(v')$ where “ \geq ” denotes the element-wise inequality.

Using the properties of W^* , we will prove the bounds on VoI^+ . First, the lower bound. Consider the Bayesian covering game (G, μ_0, Π) . Observe that

$$\max_{\alpha \in \text{NE}(G, \mu_0, f^{\text{mc}})} \bar{W}(\alpha; \mu_0) = W^* \left(\sum_{k=1}^m p_k \mathbb{E}[v|\pi_k] \right) \quad (5a)$$

$$\leq \sum_{k=1}^m p_k W^*(\mathbb{E}[v|\pi_k]) \quad (5b)$$

$$= \max_{\alpha \in \text{BNE}(G, \mu_0, \Pi, f^{\text{mc}})} \bar{W}(\alpha; \mu_0, \Pi), \quad (5c)$$

where (5a) holds from the fact the maximum-welfare Nash equilibrium is a system optimum with values $\mathbb{E}_{\mu_0}[v]$ and the law of total probability, (5b) holds from W^* convex and Jensen’s inequality, and (5c) holds from the second claim of Lemma 1 and the first claim of Lemma 1 with the fact the maximum-Nash is a system optimum. Rearranging terms gives the first inequality in (3a). It is tight when $|\mathcal{V}| = 1$. Now, we consider the upper bound on VoI^+ .

$$\max_{\alpha \in \text{BNE}(G, \mu_0, \Pi, f^{\text{mc}})} \bar{W}(\alpha; \mu_0, \Pi) = \sum_{k=1}^m p_k W^*(\mathbb{E}[v|\pi_k]) \quad (6a)$$

$$= \sum_{k=1}^m W^*(p_k \mathbb{E}[v|\pi_k]) \leq \sum_{k=1}^m W^*(\mathbb{E}_{\mu_0}[v]) \quad (6b)$$

$$= |\Pi| \left(\max_{a \in \text{NE}(G, \mu_0, f^{\text{mc}})} \bar{W}(a; \mu_0) \right), \quad (6c)$$

where (6a) holds from the fact a Bayes-Nash joint strategy α is an m -tuple of Nash equilibria, the maximum-welfare Nash equilibrium is optimal in $W(\cdot; v)$ and the second claim in Lemma 1. Eq. (6b) holds from W^* positive homogeneous and the monotonicity of W^* ; more specifically,

$$\begin{aligned} \mathbb{E}_{\mu_0}[v_r] &= \sum_{k \in [m]} p_k \mathbb{E}[v_r|\pi_k] \\ &= p_k \mathbb{E}[v_r|\pi_k] + \sum_{k' \in [m] \setminus k} p_{k'} \mathbb{E}[v_r|\pi_{k'}] \geq p_k \mathbb{E}[v_r|\pi_k], \end{aligned}$$

holds $\forall r \in \mathcal{R}, \pi_k \in \Pi$. Eq. (6c) holds from the definition W^* .

To see this is tight, consider a problem with R resources and one agent who can select a single one of them, i.e., $\mathcal{A}_1 = \{1, \dots, R\}$. Each resource can take on one of two values: 0 or 1. The prior μ_0 is that exactly one resource is ever of value 1 with equal probability; so the support of v has R elements, each of which occurs with probability $1/R$. When uninformed, the single agent is indifferent over their actions and cannot attain a payoff greater than $1/R$. When informed, the single agent can always select the resource of value 1. Thus $\text{VoI}^+ = R = |\mathcal{V}| = |\Pi|$.

Worst-case equilibrium - To focus on equilibrium strategies, let $a^{\text{NE}}(v)$ denote a Nash equilibrium joint action in the game G when the values are v . The following steps will hold for any Nash equilibrium, and the proof will be completed by considering $a^{\text{NE}}(v)$ as the welfare minimizing Nash equilibrium. Observe that the uninformed Nash welfare satisfies

$$\begin{aligned} W(a^{\text{NE}}(\mathbb{E}_{\mu_0}[v]); \mathbb{E}_{\mu_0}[v]) &\leq W^*(\mathbb{E}_{\mu_0}[v]) \\ &\leq \sum_{k=1}^m p_k W^*(\mathbb{E}[v|\pi_k]) \leq \sum_{k=1}^m p_k 2W(a^{\text{NE}}(\mathbb{E}[v|\pi_k]); \mathbb{E}[v|\pi_k]), \end{aligned}$$

where the first inequality holds from the definition of W^* , the second holds from properties of W^* shown in the first part of the proof, and the third holds from the price-of-anarchy bound of $1/2$ [22]. Letting $a^{\text{NE}}(v)$ be the worst-case Nash equilibrium when the values are v in each occurrence, the rightmost expression is the worst-case welfare in a Bayes-Nash joint strategy via Lemma 1. This gives the first inequality in (3a).

To see that this bound is tight, consider a resource allocation game with three resources, $\mathcal{R} = \{1, 2, 3\}$, and two players, $N = \{1, 2\}$ with two actions each: $\mathcal{A}_1 = \{r_1, r_2\}$ and $\mathcal{A}_2 = \{r_2, r_3\}$. Let resource r_1 have value $v_1 = 1$ and r_3 have value $v_3 = 0$. Let the value of resource r_2 be a random variable, where $v_2 = 1 - \varepsilon$ with probability $1 - p$ and $v_2 = 1 + \varepsilon(1 - p)$ with probability p . When the agents are not informed of the realization of v_2 , its expected value is $\mathbb{E}[v_2] = 1 - \varepsilon(1 - p)^2 < 1$ and there is a unique Nash equilibrium of $a = (r_1, r_2)$, providing a welfare of $W(a; \mu_0) = 2 - \varepsilon(1 - p)^2$. When the full revelation signaling policy $\Pi = \{\{r_1\}, \{r_2\}\}$ is used, then when $v_2 = 1 - \varepsilon$, there is a unique equilibrium of $\alpha^1 = (r_1, r_2)$, and when $v_2 = 1 + \varepsilon(1 - p)$ there are two equilibria, the worse of which is $\alpha^2 = (r_2, r_3)$. This gives an expected welfare of

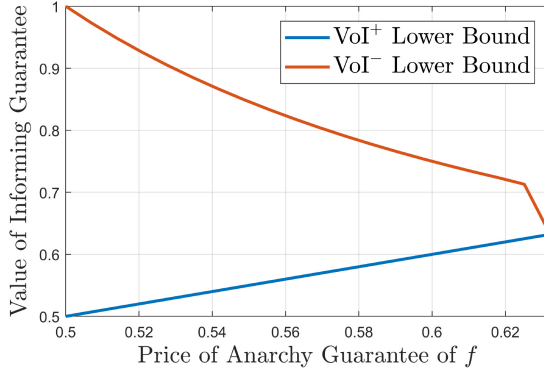


Fig. 2. Lower bounds on the value of revealing information for the best- and worst-case equilibria. If the utility rule f is designed to lessen the loss to worst-case equilibria (increasing Vol^-), then there is a greater possible loss to the worst-case equilibrium (decreasing Vol^+). This trade-off matches the tight bounds from Theorem 1 and Proposition 1, which appear as the endpoints of this plot. The bounds are generated by comparing the value-of-informing to the price-of-anarchy/stability via Theorem 2 and the price of stability to the price-of-anarchy via [10, Th. 4.1].

$W(\alpha; \mu_0) = (1-p)(2-\varepsilon) + p(1+\varepsilon(1-p))$ and a gain of informing agents of $\text{VoI}^-(\Pi, f^{\text{mc}}) = \frac{(1-p)(2-\varepsilon) + p(1+\varepsilon(1-p))}{2-\varepsilon(1-p)^2}$. Letting $p \rightarrow 1$ and $\varepsilon \rightarrow 0$ we get $\text{VoI}^-(\Pi, f^{\text{mc}}) \rightarrow 1/2$.

Finally, we show the upper bound on VoI^- . Observe that the welfare of a Bayes-Nash strategy $\alpha = \{a^{\text{NE}}(\mathbb{E}[v|\pi])\}_{\pi \in \Pi}$ under the signaling policy Π satisfies

$$\begin{aligned} \sum_{k=1}^m p_k W(a^{\text{NE}}(\mathbb{E}[v|\pi_k]); \mathbb{E}[v|\pi_k]) &\leq \sum_{k=1}^m p_k W^*(\mathbb{E}[v|\pi_k]) \\ &\leq |\Pi| \cdot W^*(\mathbb{E}_{\mu_0}[v]) \leq 2|\Pi| \cdot W(a^{\text{NE}}(\mathbb{E}_{\mu_0}[v]); \mathbb{E}_{\mu_0}[v]), \end{aligned}$$

where the first term is the expected payoff of α , and the first inequality holds from W^* being the optimal welfare, the second holds from the upper bound on VoI^+ , and the third holds from the price-of-anarchy bound.

B. Utility Design & Informing Efficacy

In Section III-A, we considered the special case in which agents' local objectives were aligned with the global objective using the marginal cost utility rule f^{mc} . In this section, we generalize this result to any utility rule by leveraging a connection to the *price-of-stability* and *price-of-anarchy*.

In the deterministic setting, the price-of-stability/anarchy are used to quantify how the best-case and worst-case equilibria approximate the system optimal. These metrics can be generalized to the Bayesian setting but need not be informative or insightful on the effects revealing information within a single problem instance [23] (i.e., failing to capture the benefits and consequences or comparing bounds derived from different problem instances). Let $\text{PoA}(G, v, f) = \frac{\min_{a \in \text{NE}(G, v, f)} W(a^{\text{NE}}; v)}{\max_{a \in \text{opt}_{\mathcal{A}}} W(a^{\text{opt}}; v)}$ denote the price-of-anarchy for a deterministic maximum coverage problem, and let $\text{PoS}(G, v, f) = \frac{\max_{a \in \text{NE}(G, v, f)} W(a^{\text{NE}}; v)}{\max_{a \in \text{opt}_{\mathcal{A}}} W(a^{\text{opt}}; v)}$ denote the price-of-stability.

Though not immediately apparent, we establish a connection between the price-of-stability/anarchy in deterministic covering games and the value-of-informing in Bayesian covering

games. In Theorem 2, we leverage this connection to generate bounds on VoI^+ and VoI^- for any utility design.

Theorem 2: Let $\psi := \inf_{v \in \text{conv}(\mathcal{V})} \text{PoS}(G, v, f)$ and $\rho := \inf_{v \in \text{conv}(\mathcal{V})} \text{PoA}(G, v, f)$, then the value of informing for the best-case equilibrium satisfies

$$\psi \leq \text{VoI}^+(G, \mu_0, \Pi, f) \leq \psi^{-1}|\Pi|, \quad (7a)$$

and for the worst-case equilibrium satisfies

$$\rho \leq \text{VoI}^-(G, \mu_0, \Pi, f) \leq \rho^{-1}|\Pi|. \quad (7b)$$

Proof of Theorem 2: The proof will rely on the function $W^*(v) = \max_{a \in \mathcal{A}} W(a; v)$ and several of its properties shown in the proof of Theorem 1. First, we prove the bounds on VoI^+ . Let $a \in \text{NE}(G, \mathbb{E}_{\mu_0}[v], f)$ be an arbitrary Nash equilibrium in the deterministic game G with values $\mathbb{E}_{\mu_0}[v]$ and utility rule f , and let $\alpha \in \text{BNE}(G, \mu_0, \Pi, f)$ be an arbitrary Bayes-Nash equilibrium in G with prior μ_0 on v and information signaling policy Π and utility function f . Observe that the expected welfare of α satisfies

$$\bar{W}(\alpha; \mu_0, \Pi) = \sum_{k=1}^m p_k W(\alpha(\pi_k); \mathbb{E}[v|\pi_k]) \quad (8a)$$

$$\leq \sum_{k=1}^m p_k W^*(\mathbb{E}[v|\pi_k]) \quad (8b)$$

$$\leq |\Pi| W^*(\mathbb{E}_{\mu_0}[v]) \quad (8c)$$

$$\leq \rho^{-1}|\Pi| \bar{W}(a; \mathbb{E}_{\mu_0}[v]), \quad (8d)$$

where (8a) holds from Lemma 1, (8b) holds from the definition of W^* , (8c) holds from the monotonicity and positive homogeneity of W^* (previously shown in (6b)-(6c)), and (8d) holds from the definition of ρ .

Similarly, we can show that the expected total welfare of the uninformed equilibrium a satisfies

$$\bar{W}(a; \mu_0) \leq W^*(v) \quad (9a)$$

$$\leq \sum_{k=1}^m p_k W^*(\mathbb{E}[v|\pi_k]) \quad (9b)$$

$$\leq \sum_{k=1}^m p_k \rho^{-1} W(\alpha(\pi); \mathbb{E}[v|\pi_k]) \quad (9c)$$

$$= \rho^{-1} W(\alpha; \mu_0, \Pi), \quad (9d)$$

where (9b) holds from the convexity of W^* , (9c) holds from the definition of ρ , and (9d) holds from Lemma 1. Because (8) and (9) hold for any $a \in \text{NE}(G, \mathbb{E}_{\mu_0}[v], f)$ and $\alpha \in \text{BNE}(G, \mu_0, \Pi, f)$, it holds for each being the respective welfare minimizing equilibria. If we consider only the case where a and α are the welfare maximizing equilibria, i.e., $a \in \arg \max_{a' \in \text{NE}(G, \mathbb{E}_{\mu_0}[v], f)} W(a'; \mathbb{E}_{\mu_0}[v])$ and $\alpha \in \arg \max_{\alpha' \in \text{BNE}(G, \mu_0, v, f)} \bar{W}(\alpha'; \mu_0, \Pi)$

C. Optimistic/Pessimistic Trade-Off

Theorem 2 highlighted the fact that altering the utility design will change the impact of information revealing; however, the given bounds need not be tight. It appears that using a utility design with a higher price-of-anarchy in the deterministic setting will lead to an improved lower bound on VoI^- .

As such, we will more closely consider the price-of-anarchy maximizing rule

$$f^g(x) := (x-1)! \frac{\frac{1}{(n-1)(n-1)!} + \sum_{i=x}^{n-1} \frac{1}{i!}}{\frac{1}{(n-1)(n-1)!} + \sum_{i=1}^{n-1} \frac{1}{i!}},$$

when $x > 0$, and $f^g(0) = 0$, proven optimal in [9]. In Proposition 1, we show tight lower bounds on VoI with f^g .

Proposition 1: While using the price-of-anarchy maximizing rule f^g , the value-of-informing for the best-case equilibrium satisfies

$$1 - \frac{1}{e} \leq \text{VoI}^+(G, \mu_0, \Pi, f^g), \quad (10a)$$

and for the worst-case equilibrium satisfies

$$1 - \frac{1}{e} \leq \text{VoI}^-(G, \mu_0, \Pi, f^g), \quad (10b)$$

Further, each of these lower bounds is tight.

Proof of Proposition 1: Theorem 2 can be used to show that each of (10a) and (10) are valid lower bounds. To see these bounds are tight, consider a maximum coverage problem with resource set $\mathcal{R} = \{r_{p_i}\}_{i=1}^{n-1} \cup \{r_{s_j}\}_{j=1}^z$ where $z = \lceil 1/(f^g(n) - \varepsilon) \rceil$. The first $n-1$ players can select a single resource from the public resources $\{r_{s_j}\}_{j=1}^z$ or their respective private resource r_{p_i} for player i , i.e., $\mathcal{A}_i = \{r_{s_1}, \dots, r_{s_z}, r_{p_i}\}$, $\forall i \in N \setminus \{n\}$. The final player has exactly one action to use all the shared resources simultaneously $\mathcal{A}_n = (r_{s_1}, \dots, r_{s_z})$. Each private resource r_{p_i} has value $v_{r_{p_i}} = f^g(n) - \varepsilon$ with probability one where $\varepsilon > 0$. The value of the shared resources are random; each takes on value 1 w.p. $1/z < f^g(n) - \varepsilon$ and 0 otherwise, and follow distribution μ_0 such that exactly one is ever the high value. When uninformed, each of the first $n-1$ players strictly prefers their private resource, giving a unique equilibrium welfare of $W(\alpha^{\text{NE}}; \mu_0) = 1 + (n-1)(f^g(n) - \varepsilon)$. Under the full information reveal policy Π , each of the $n-1$ agents strictly prefer to use the one shared resource of value 1, giving an expected welfare of $\bar{W}(\alpha^{\text{BNE}}; \mu_0, \Pi) = 1$. Because each of the informed and uninformed equilibria are unique, this gives $\text{VoI}^+(\mu_0, \Pi) = \text{VoI}^-(\mu_0, \Pi) = \frac{1}{1+(n-1)(f^g(n)-\varepsilon)} \rightarrow 1 - \frac{1}{e}$ as $\varepsilon \rightarrow 0$ and $n \rightarrow \infty$.

Comparing the lower bounds of Theorem 1 and Proposition 1 highlights a trade-off between revealing information in the optimistic and pessimistic perspectives. We further examine this trade-off in Fig. 2. Using Theorem 2 and recent results of [10], we can characterize lower bounds on VoI^+ and VoI^- for different utility rules.

IV. CONCLUSION AND FUTURE WORK

We addressed the possible benefit and consequences of revealing information to local decision-makers in a distributed system. By lower and upper bounding the value-of-informing, this letter (1) quantified the possible effects of information revealing and (2) identified a trade-off between the guarantees of revealing information in the optimistic and pessimistic perspective. Future work will answer the design question and develop methods to solve for information signaling policies that optimize expected system welfare.

REFERENCES

- [1] G. A. Korsah, A. Stentz, and M. B. Dias, "A comprehensive taxonomy for multi-robot task allocation," *Int. J. Robot. Res.*, vol. 32, no. 12, pp. 1495–1512, 2013.
- [2] S. Wollenstein-Betech, A. Houshmand, M. Salazar, M. Pavone, C. G. Cassandras, and I. C. Paschalidis, "Congestion-aware routing and rebalancing of autonomous mobility-on-demand systems in mixed traffic," in *Proc. IEEE Conf. Intell. Transp. Syst.*, 2020, pp. 1–7.
- [3] N. Mathew, S. L. Smith, and S. L. Waslander, "Planning paths for package delivery in heterogeneous multirobot teams," *IEEE Trans. Autom. Sci. Eng.*, vol. 12, no. 4, pp. 1298–1308, Oct. 2015.
- [4] J. R. Marden and J. S. Shamma, "Game theory and distributed control," in *Handbook of Game Theory*, vol. 4, H. P. Young and S. Zamir, Eds. Amsterdam, The Netherlands: Elsevier Sci., 2014.
- [5] B. Swenson, R. Murray, and S. Kar, "On best-response dynamics in potential games," *SIAM J. Control Optim.*, vol. 56, no. 4, pp. 2734–2767, Jul. 2018.
- [6] S. Rahili, J. Lu, W. Ren, and U. Al-Saggaf, "Distributed coverage control of mobile sensor networks in unknown environment using game theory: Algorithms and experiments," *IEEE Trans. Mobile Comput.*, vol. 17, no. 6, pp. 1303–1313, Jun. 2018.
- [7] B. L. Ferguson and J. R. Marden, "Robust utility design in distributed resource allocation problems with defective agents," *Dyn. Games Appl.*, vol. 13, pp. 208–230, Mar. 2023.
- [8] D. Paccagnan, R. Chandan, and J. R. Marden, "Utility design for distributed resource allocation—Part I: Characterizing and optimizing the exact price of anarchy," *IEEE Trans. Autom. Control*, vol. 65, no. 11, pp. 4616–4631, Nov. 2020.
- [9] M. Gairing, "Covering games: Approximation through non-cooperation," in *Int. Workshop Internet Netw. Econ.*, 2009, pp. 184–195.
- [10] V. Ramaswamy, D. Paccagnan, and J. R. Marden, "Multiagent maximum coverage problems: The tradeoff between anarchy and stability," *IEEE Trans. Autom. Control*, vol. 67, no. 4, pp. 1698–1712, Apr. 2022.
- [11] D. Acemoglu, A. Makhdoumi, A. Malekian, and A. Ozdaglar, "Informational Braess' paradox: The effect of information on traffic congestion," *Oper. Res.*, vol. 66, no. 4, pp. 893–917, 2018.
- [12] E. Kamenica and M. Gentzkow, "Bayesian persuasion," *Amer. Econ. Rev.*, vol. 101, no. 6, pp. 2590–2615, Oct. 2011.
- [13] M. Castiglioni, A. Celli, A. Marchesi, and N. Gatti, "Signaling in Bayesian network congestion games: The subtle power of symmetry," in *Proc. AAAI Conf. Artif. Intell.*, vol. 35, May 2021, pp. 5252–5259.
- [14] B. L. Ferguson, P. N. Brown, and J. R. Marden, "Avoiding unintended consequences: How incentives aid information provisioning in Bayesian congestion games," in *Proc. IEEE 61st Conf. Decis. Control (CDC)*, Dec. 2022, pp. 3781–3786.
- [15] H. Tavaafoghi and D. Teneketzis, "Informational incentives for congestion games," in *Proc. 55th Annu. Allerton Conf. Commun., Control, Comput., Allerton*, Jan. 2018, pp. 1285–1292.
- [16] Y. Chen and J. Zhang, "Signalling by Bayesian persuasion and pricing strategy," *Econ. J.*, vol. 130, no. 628, pp. 976–1007, 2020.
- [17] S. Shah, S. Amin, and P. Jaillet, "Optimal information provision for strategic hybrid workers," in *Proc. IEEE 61st Conf. Decis. Control (CDC)*, Dec. 2022, pp. 3807–3814.
- [18] F. Sezer, H. Khazaee, and C. Eksin, "Social welfare maximization and conformism via information design in linear-quadratic-Gaussian games," 2021, *arXiv:2102.13047*.
- [19] Y. Zhu and K. Savla, "Information design in nonatomic routing games with partial participation: Computation and properties," *IEEE Trans. Control Netw. Syst.*, vol. 9, no. 2, pp. 613–624, Jun. 2022.
- [20] M. Wu and S. Amin, "Information design for regulating traffic flows under uncertain network state," in *Proc. Annu. Allerton Conf. Commun., Control, Comput.*, 2019, pp. 671–678.
- [21] D. Paccagnan, R. Chandan, and J. R. Marden, "Utility and mechanism design in multi-agent systems: An overview," *Annu. Rev. Control*, vol. 53, pp. 315–328, Jan. 2022.
- [22] A. Vetta, "Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions," in *Proc. IEEE Symp. Found. Comput. Sci.*, 2002, pp. 416–425.
- [23] U. Bhaskar, Y. Cheng, Y. K. Ko, and C. Swamy, "Hardness results for signaling in Bayesian zero-sum and network routing games," in *Proc. ACM Conf. Econ. Comput.*, 2015, pp. 479–496.