

Bayesian Ignorance

[Extended Abstract]

Noga Alon
Tel Aviv University and
Microsoft Israel R&D Center
nogaa@tau.ac.il

Yuval Emek
Microsoft Israel R&D Center
and Tel Aviv University
yuvale@eng.tau.ac.il

Michal Feldman
Hebrew University of
Jerusalem and Microsoft Israel
R&D Center
mfeldman@huji.ac.il

Moshe Tennenholtz
Microsoft Israel R&D Center
and Technion - Israel Institute
of Technology
moshet@ie.technion.ac.il

ABSTRACT

We quantify the effect of *Bayesian ignorance* by comparing the social cost obtained in a Bayesian game by agents with *local* views to the expected social cost of agents having *global* views. Both benevolent agents, whose goal is to minimize the social cost, and selfish agents, aiming at minimizing their own individual costs, are considered. When dealing with selfish agents, we consider both best and worst equilibria outcomes. While our model is general, most of our results concern the setting of *network cost sharing (NCS)* games. We provide tight asymptotic results on the effect of Bayesian ignorance in directed and undirected NCS games with benevolent and selfish agents. Among our findings we expose the counter-intuitive phenomenon that “ignorance is bliss”: Bayesian ignorance may substantially improve the social cost of selfish agents. We also prove that *public* random bits can replace the knowledge of the common prior in attempt to bound the effect of Bayesian ignorance in settings with benevolent agents. Together, our work initiates the study of the effects of local vs. global views on the social cost of agents in Bayesian contexts.

Categories and Subject Descriptors

F.2.3 [Theory of Computation]: Analysis of Algorithms and Problem Complexity—*tradeoffs among complexity measures*

General Terms

Theory

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

PODC'10, July 25–28, 2010, Zurich, Switzerland.

Copyright 2010 ACM 978-1-60558-888-9/10/07 ...\$10.00.

Keywords

Bayesian games, local vs. global view, network cost sharing

1. INTRODUCTION

It is very common that the participants of a distributed system are required to make decisions which are based on their own local views¹ rather than on a global view of the system. This lack of global view may have severe implications on the overall system’s performance. In this paper we introduce a novel approach for quantifying these implications. Our approach relies on the notion of Bayesian games that we now turn to describe.

Consider some k agents, where each agent i is associated with an *action space* A_i and a *type space* T_i . Every type profile $t \in T = T_1 \times \dots \times T_k$ (a.k.a. the *state* of the system) induces an *underlying* game G_t defined by a *cost function* $C_{i,t}$, $i \in [k]$, that maps each action profile $a \in A = A_1 \times \dots \times A_k$ to the cost incurred by agent i on a under t . A *Bayesian game* is merely a probability distribution p , referred to as the *common prior*, over the underlying games or, more formally, over the type profiles.

It is assumed that some type profile $t = (t_1, \dots, t_k) \in T$ is chosen with respect to the common prior p . The crux of the model is that although p is common knowledge, the agents have local views of the actual instantiation t so that each agent i knows only her own type t_i within the type profile t and should decide on her action based upon that local view alone. Therefore, a *strategy* of agent i is a function $s_i : T_i \rightarrow A_i$ that maps each type to one of her actions. The *social cost* incurred by a strategy profile is defined as the expected sum of the agents’ costs.

In light of the above, the main principle of Bayesian games is that the agents cannot coordinate their actions on the global state of the system as each agent’s action is dictated by her own local view. Obviously, the privilege of bearing global views could have dramatically affected the agent’s behavior and in particular, improve the social cost (assuming that this is the agent’s goal). We refer to this lack of global

¹The term local view in this context refers to an absence of information regarding the (global) system’s state. It should not be confused with a limited knowledge of the communication network.

view as *Bayesian ignorance* and our goal is to quantify its effect by comparing the social cost attained in the Bayesian game, i.e., under local views, to the expected social cost of agents having global views, where the expectation is taken with respect to the common prior distribution on the underlying games. In particular, we will focus on the ratio of the social cost of the optimal strategy profile in the Bayesian game to the expected optimal social cost in the underlying games.

The aforementioned discussion assumes that the agents are *benevolent* in the sense that their goal is to minimize the social cost. We will also consider *selfish* agents whose aim is to minimize their own individual costs. When dealing with selfish agents, we restrict attention to the set of *equilibrium profiles*². Specifically, we study the ratio of the social cost of a *best* equilibrium strategy profile in the Bayesian game to the expected social cost of the best equilibrium in the underlying games. Finally, we also study the same ratio with respect to a *worst* equilibrium. (Refer to Section 2 for a formal definition of the model.)

In our opinion, the effect of Bayesian ignorance must be taken into account when one attempts to design a distributed system with prior beliefs on the system’s global state. Specifically, if the social cost in the Bayesian setting is significantly higher than the expected social cost in the underlying complete-information setting, then the system’s performance still have much “room for improvement”. In such a scenario, the system’s designer may wish to invest into some sort of a correlation device that would reduce the effect of Bayesian ignorance, or alternatively, to redesign the system.

While our model and quality measures are general, most of the technical results established in the current paper concern the setting of *network cost sharing (NCS)* games. An NCS game is specified by a (directed or undirected) graph in which every edge is associated with a non-negative *cost*, and a set of k agents, each associated with a *source* vertex and a *destination* vertex. Each agent should *buy* a subset of the edges so as to connect her source to her destination. The cost of each edge is shared equally among all agents who bought it; the cost incurred by an agent is merely the sum of (partial) payments it made for the edges she bought. An NCS game is a congestion game [14], therefore it always admits an equilibrium in pure strategies. The social cost, which by definition, equals the sum of the agents’ costs, is just the total cost of the edges bought by all agents. In a Bayesian NCS game each agent knows her own source and destination, but not the sources and destinations of the others.

In Section 3 we provide tight asymptotic results on the effect of Bayesian ignorance in directed and undirected NCS games with benevolent and selfish agents. Among our results we expose two interesting phenomena. First, while allowing benevolent agents to bear global views is clearly socially beneficial, in selfish agent settings ignorance may be bliss. In particular, we present a Bayesian NCS game in which the social cost of the worst Bayesian equilibrium is asymptotically smaller than the expected social cost of the best equilibrium in the underlying games. In fact, in that Bayesian game the worst Bayesian equilibrium achieves the expected cost of the

²The assumption that selfish agents converge to an equilibrium profile (at least to a pure equilibrium profile when one exists) is among the most fundamental concepts in game theory.

globally optimal outcome. Second, in settings with benevolent agents we find that *public* random bits can replace the knowledge of the common prior in attempt to bound the effect of Bayesian ignorance (see Section 4).

Related work.

The effect of Bayesian ignorance is closely related to the notion of the *value of information* [10], defined as the amount a decision maker would be willing to pay for information prior to making a decision. This notion has been axiomatized by Gilboa and Lehrer [8]. Another work that is conceptually close to our study is that of Ashlagi *et al.* [5], which quantifies the loss or value that can be obtained due to lack of information about the number of agents in a resource selection game, alas not in a Bayesian setting. Similar in spirit to our observation, ignorance may improve the social welfare in their setting as well.

The network cost sharing game has been originally introduced by Anshelevich *et al.* [4] and has been extensively studied in a non-Bayesian setting in recent years. Within the context of NCS games, a great deal of attention has been given to the *price of anarchy* measure [12, 13], defined as the ratio of the cost of a worst Nash equilibrium to the social optimum and the *price of stability* measure [4], defined as the ratio of the cost of a best Nash equilibrium to the social optimum [6, 3, 1]. These measures can be thought of as quantifying the loss obtained due to selfish behavior. In contrast, the Bayesian ignorance measures introduced in the current paper quantify the loss (or gain) obtained due to local views (in either benevolent or selfish behaviors).

2. THE MODEL

Bayesian games.

A *Bayesian game* \mathcal{G} is a 5-tuple

$$\mathcal{G} = \langle k, \{A_i\}_{i \in [k]}, \{T_i\}_{i \in [k]}, \{C_{i,t}\}_{i \in [k], t \in T}, p \rangle,$$

where $k \in \mathbb{Z}_{>0}$ is the number of *agents*, A_i is the (finite) *action* space of agent i , T_i is the (finite) *type* space of agent i , $C_{i,t} : A \rightarrow \mathbb{R}$ is the *cost function* of agent i under the type profile $t \in T = T_1 \times \dots \times T_k$ that maps each action profile $a \in A = A_1 \times \dots \times A_k$ to the *cost* incurred by agent i from a under t , and $p \in \Delta(T)$ is a probability distribution over the type profiles T , referred to as the *common prior*.

It is assumed that a type profile $t = (t_1, \dots, t_k) \in T$ is chosen with probability $p(t)$. The fundamental principle of Bayesian games is that the common prior p (in fact, the whole 5-tuple \mathcal{G}) is common knowledge, but each agent i knows only her own type t_i out of the actual instantiation t and should decide on her action $a_i \in A_i$ based on that partial knowledge. A pure *strategy* of agent i in the Bayesian game \mathcal{G} is therefore a function $s_i : T_i \rightarrow A_i$ that maps her type to some action. We denote the strategy space of agent i by $S_i = A_i^{T_i}$ and the collection of all strategy profiles by $S = S_1 \times \dots \times S_k$.

Fix some strategy profile $s = (s_1, \dots, s_k) \in S$. Let $X_i(s)$ be the random variable (defined over the probability space p) that takes on the cost incurred by agent $i \in [k]$ from the strategy profile s , i.e., $X_i(s) = C_{i,t}(\{s_j(t_j)\}_{j \in [k]})$ with probability $p(t)$. The expected cost incurred by agent i from the strategy profile s (a.k.a. the *ex-ante cost*) is then defined

as

$$C_i(s) = \mathbb{E}[X_i(s)] = \sum_{t \in T} p(t) \cdot C_{i,t}(\{s_j(t_j)\}_{j \in [k]}) .$$

The strategy profile s is a (pure) *Bayesian equilibrium* of the Bayesian game \mathcal{G} if no agent gains on expectation from a unilateral deviation, that is, if for every $i \in [k]$ and for every $s'_i \in S_i$, $C_i(s) \leq C_i(s_{-i}, s'_i)$. Alternatively, s is a Bayesian equilibrium if for every $i \in [k]$, for every $t_i \in T_i$, and for every $s'_i \in S_i$ that agrees with s_i on all types in T_i except (maybe) t_i , $\mathbb{E}[X_i(s) \mid t_i] \leq \mathbb{E}[X_i(s_{-i}, s'_i) \mid t_i]$.

Potential functions.

We say that a Bayesian game \mathcal{G} is a *Bayesian potential game* if there exists a function $Q : S \rightarrow \mathbb{R}$ that satisfies $C_i(s) - C_i(s_{-i}, s'_i) = Q(s) - Q(s_{-i}, s'_i)$ for every $s \in S$, $i \in [k]$, and $s'_i \in S_i$. In this case Q is referred to as a *Bayesian potential function* for \mathcal{G} . Since the strategy profile space S is finite, there must exist some $s \in S$ such that the Bayesian potential function Q is minimized at s , i.e., $Q(s) \leq Q(s')$ for every $s' \in S$; by definition, this strategy profile s is a Bayesian equilibrium. A function $q_t : A \rightarrow \mathbb{R}$ is called a *potential function* for the type profile $t \in T$ if $C_{i,t}(a) - C_{i,t}(a_{-i}, a'_i) = q_t(a) - q_t(a_{-i}, a'_i)$ for every $i \in [k]$, $a = (a_1, \dots, a_k) \in A$, and $a'_i \in A_i$.

OBSERVATION 2.1. *If a potential function q_t exists for every type profile $t \in T$, then the function $Q : S \rightarrow \mathbb{R}$ defined by mapping each strategy profile $s = (s_1, \dots, s_k) \in S$ to*

$$Q(s) = \sum_{t \in T} p(t) \cdot q_t(\{s_j(t_j)\}_{j \in [k]})$$

is a Bayesian potential function for the Bayesian game \mathcal{G} .

PROOF. Consider some strategy profile $s = (s_1, \dots, s_k) \in S$, agent $i \in [k]$, and strategy $s'_i \in S_i$. We have

$$\begin{aligned} & C_i(s) - C_i(s_{-i}, s'_i) \\ &= \sum_{t \in T} p(t) \cdot [C_{i,t}(\{s_j(t_j)\}_{j \in [k]}) - C_{i,t}(\{s_j(t_j)\}_{j \in [k]-i}, s'_i(t_i))] \\ &= \sum_{t \in T} p(t) \cdot [q_t(\{s_j(t_j)\}_{j \in [k]}) - q_t(\{s_j(t_j)\}_{j \in [k]-i}, s'_i(t_i))] \\ &= Q(s) - Q(s_{-i}, s'_i) . \end{aligned}$$

The assertion follows. \square

Ignorance.

A Bayesian game $\mathcal{G} = \langle k, \{A_i\}_{i \in [k]}, \{T_i\}_{i \in [k]}, \{C_{i,t}\}_{i \in [k], t \in T}, p \rangle$ is a *partial-information game* in which every type profile $t \in T$ induces a *complete-information game* \mathcal{G}_t specified by the cost functions $C_{1,t}, \dots, C_{k,t}$. The complete-information games \mathcal{G}_t are referred to as the *underlying games*. We restrict our attention to Bayesian games admitting pure Bayesian equilibria in which all underlying games admit pure Nash equilibria (e.g., when the underlying games have potential functions).

Fixing some type profile $t \in T$, the *social cost* induced on t by the action profile $a \in A$ is defined as $K_t(a) = \sum_{i \in [k]} C_{i,t}(a)$. The *social cost* of the strategy profile $s \in S$ is then defined as

$$K(s) = \sum_{t \in T} p(t) \cdot K_t(\{s_i(t_i)\}_{i \in [k]}) = \sum_{i \in [k]} C_i(s) .$$

To avoid cumbersome notation, in what follows we denote $K_t(\{s_i(t_i)\}_{i \in [k]}) = K(s, t)$. We are interested in comparing the social cost of some strategy profiles in the partial-information Bayesian setting to the average social cost of some action profiles in the complete-information underlying games. More formally, we shall establish lower and upper bounds on ratios in which the numerator is either

- $\text{opt}^P(\mathcal{G}) = \min_{s \in S} K(s)$,
- $\text{best-eq}^P(\mathcal{G}) = \min_{\text{Bayesian equilibrium } s \text{ of } \mathcal{G}} K(s)$, or
- $\text{worst-eq}^P(\mathcal{G}) = \max_{\text{Bayesian equilibrium } s \text{ of } \mathcal{G}} K(s)$;

and the denominator is either

- $\text{opt}^C(\mathcal{G}) = \sum_{t \in T} p(t) \cdot \min_{a \in A} K_t(a)$,
- $\text{best-eq}^C(\mathcal{G}) = \sum_{t \in T} p(t) \cdot \min_{\text{Nash equilibrium } a \text{ of } \mathcal{G}_t} K_t(a)$, or
- $\text{worst-eq}^C(\mathcal{G}) = \sum_{t \in T} p(t) \cdot \max_{\text{Nash equilibrium } a \text{ of } \mathcal{G}_t} K_t(a)$.

These ratios can be thought of as reflecting the effect of *ignorance* on behalf of the agents: the numerator captures the social cost of a strategy profile in a partial-information setting where each agent knows only her own type (and the common prior p), while the denominator captures the average (with respect to p) social cost of action profiles in complete-information settings.

OBSERVATION 2.2. *Every Bayesian game \mathcal{G} admitting a pure Bayesian equilibrium satisfies*

$$\text{opt}^C(\mathcal{G}) \leq \text{opt}^P(\mathcal{G}) \leq \text{best-eq}^P(\mathcal{G}) \leq \text{worst-eq}^P(\mathcal{G}) .$$

Network cost sharing.

A *network cost sharing (NCS) game* is specified by a graph $G = (V, E)$ (may be directed or undirected), a non-negative real cost $c(e)$ associated with each edge $e \in E$, and a vertex pair $(x_i, y_i) \in V^2$ associated with each agent $i \in [k]$, where x_i (respectively, y_i) is referred to as the *source* (resp., *destination*) of agent i . The action space of each agent is 2^E ; it is convenient to think of action $a_i \subseteq E$ of agent i as if the agent *buys* the edges in a_i . Given some action profile $a = (a_1, \dots, a_k)$, the *payment* $\pi_i(e)$ of agent i for edge $e \in E$ is

$$\pi_i(e) = \begin{cases} \frac{c(e)}{|\{j \in [k] \mid e \in a_j\}|} & \text{if } e \in a_i; \\ 0 & \text{if } e \notin a_i. \end{cases}$$

The *total payment* of agent i for the action profile a is defined to be $\pi_i(a) = \sum_{e \in E} \pi_i(e)$. An NCS game is then defined by setting the cost incurred by agent i to be its total payment if a_i contains a path from x_i to y_i ; and ∞ otherwise.

NCS games fall into a more general family of games called *congestion games*. Rosenthal [14] shows that every congestion game admits a potential function. In the context of the NCS games, this potential function turns out to be

$$q(a) = \sum_{e \in E} c(e) \cdot H(|\{i \in [k] \mid e \in a_i\}|) ,$$

where $H(n) = 1 + 1/2 + \dots + 1/n$ is the n^{th} harmonic number (cf. [4]).

A *Bayesian NCS* game is depicted by a probability distribution p over NCS games with the same graph $G = (V, E)$ and edge costs $c(e)$. That is, the action space of each agent $i \in [k]$ is $A_i = 2^E$, the type space of each agent i is $T_i = V \times V$, and the cost functions $C_{i,t}$ are determined with respect to the complete-information NCS games induced by the type profiles $t \in T = (V \times V)^k$. Observation 2.1 implies that the function $Q : S \rightarrow \mathbb{R}$ defined by mapping each strategy profile $s \in S = \left((2^E)^{V \times V} \right)^k$ to

$$Q(s) = \sum_{t \in T} p(t) \cdot \sum_{e \in E} c(e) \cdot H(|\{i \in [k] \mid e \in s_i(t_i)\}|)$$

is a Bayesian potential function for that Bayesian NCS game.

3. BAYESIAN IGNORANCE IN NETWORK COST SHARING GAMES

In this section we establish various bounds regarding the effect of Bayesian ignorance on NCS games. These bounds are summarized in Table 1. Observe that the bounds we establish are asymptotically tight in all cases except for the $\frac{\text{best-eq}^P}{\text{best-eq}^C}$ ratio in undirected graphs for which logarithmic gaps still exist. It is quite easy to design a Bayesian NCS game over an undirected graph for which $\frac{\text{best-eq}^P}{\text{best-eq}^C} < 1$. However, designing a sequence of such games for which $\frac{\text{best-eq}^P}{\text{best-eq}^C} = o(1)$ (i.e., the ratio decreases asymptotically with the number of agents) seems to be a more difficult task; in particular, this would imply the existence of an NCS game with $o(1)$ price of stability. Whether or not such an NCS game exists is still an open question³.

It is also interesting to point out that the $\frac{\text{best-eq}^P}{\text{best-eq}^C} = O(1/\log k)$ existential bound for directed graphs is established via a k -agent Bayesian NCS game \mathcal{G} such that $\text{opt}^C(\mathcal{G}) = \text{worst-eq}^P(\mathcal{G}) = O(1)$, while $\text{best-eq}^C(\mathcal{G}) = \Omega(\log k)$. This demonstrates the potential usefulness of Bayesian ignorance to the benefit of the society as it means that in some scenarios the social cost of any equilibrium of selfish agents holding local views is (asymptotically) better than all equilibria of selfish agents with global views.

We now turn to establish the bounds exhibited in Table 1. We first observe that the universal lower bounds on the $\frac{\text{opt}^P}{\text{opt}^C}$ ratio in directed and undirected graphs simply follow from Observation 2.2. The matching existential bounds are trivial as every complete-information game is also a Bayesian game. We now turn to show that the effect of Bayesian ignorance cannot be too devastating.

LEMMA 3.1. *Every k -agent Bayesian NCS game \mathcal{G} satisfies $\text{worst-eq}^P(\mathcal{G})/\text{opt}^C(\mathcal{G}) \leq k$.*

PROOF. Consider some k -agent Bayesian NCS game \mathcal{G} on a (directed or undirected) graph $G = (V, E)$. For every type profile $t \in T$, an action profile $a \in A$ that minimizes $K_t(a)$ must contain a t_i -path for every $i \in [k]$. Hence, $K_t(a) \geq \max_{i \in [k]} \text{dist}_G(t_i)$ and

$$\text{opt}^C(\mathcal{G}) \geq \sum_{t \in T} p(t) \cdot \max_{i \in [k]} \text{dist}_G(t_i),$$

³Albers [1] shows such an example for a variant of the NCS game in which the agents have different weights. This variant is beyond the scope of the current paper.

where $\text{dist}_G(t_i)$ denotes the distance in G from the source to the destination dictated by t_i .

Let $s \in S$ be a strategy profile that realizes $\text{worst-eq}^P(\mathcal{G})$ and fix some agent $i \in [k]$ and type $t_i = (x_i, y_i) \in T_i$. Since s is a Bayesian equilibrium, we must have $\mathbb{E}[X_i(s) \mid t_i] \leq \text{dist}_G(t_i)$ as otherwise, agent i is better off deviating from her strategy s_i to a strategy $s'_i \in S_i$ that agrees with s_i on all types except t_i for which s'_i buys a shortest (x_i, y_i) -path and pays at most $\text{dist}_G(t_i)$. It follows that

$$\begin{aligned} & \text{worst-eq}^P(\mathcal{G}) \\ &= \sum_{i \in [k]} \mathbb{E}[X_i(s)] = \sum_{i \in [k]} \sum_{t_i \in T_i} \mathbb{P}(t_i) \cdot \mathbb{E}[X_i(s) \mid t_i] \\ &\leq \sum_{i \in [k]} \sum_{t_i \in T_i} \mathbb{P}(t_i) \cdot \text{dist}_G(t_i) = \sum_{i \in [k]} \sum_{t \in T} p(t) \cdot \text{dist}_G(t_i) \\ &= \sum_{t \in T} p(t) \cdot \sum_{i \in [k]} \text{dist}_G(t_i) \leq k \cdot \sum_{t \in T} p(t) \cdot \max_{i \in [k]} \text{dist}_G(t_i), \end{aligned}$$

which establishes the assertion. \square

Combined with Observation 2.2, Lemma 3.1 establishes the universal upper bounds on the $\frac{\text{opt}^P}{\text{opt}^C}$, $\frac{\text{best-eq}^P}{\text{best-eq}^C}$, and $\frac{\text{worst-eq}^P}{\text{worst-eq}^C}$ ratios in directed graphs and the universal upper bound on the $\frac{\text{worst-eq}^P}{\text{worst-eq}^C}$ ratio in undirected graphs. Since complete-information games are a special case of Bayesian games, Lemma 3.1 also establishes the universal lower bounds on the $\frac{\text{worst-eq}^P}{\text{worst-eq}^C}$ ratio in directed and undirected graphs.

The following lemma implies the existential lower bounds on the $\frac{\text{opt}^P}{\text{opt}^C}$ and $\frac{\text{best-eq}^P}{\text{best-eq}^C}$ ratios in directed graphs.

LEMMA 3.2. *For every k_0 , there exist some $k \geq k_0$ and a k -agent Bayesian NCS game \mathcal{G} on a directed $\Theta(k^2)$ -vertex graph such that $\text{opt}^P(\mathcal{G})/\text{worst-eq}^C(\mathcal{G}) = \Omega(k)$.*

PROOF. Let $m \geq k_0$ be a prime power and consider a finite *affine plane* (X, L) of order m : X is a set of m^2 points; $L \subseteq 2^X$ is a set of $m^2 + m$ lines. The affine plane (X, L) satisfies the following four properties.

- (1) Each line in L contains exactly m points.
- (2) Each point in X is contained in exactly $m + 1$ lines.
- (3) Given any two distinct points in X , there is exactly one line in L that contains both points.
- (4) Given any two lines in L , there is at most one point in X which is contained in both lines.

We construct a directed graph $G = (V, E)$ by setting $V = \{u\} \cup \{v_\ell \mid \ell \in L\} \cup \{w_p \mid p \in X\}$ and $E = \{(u, v_\ell) \mid \ell \in L\} \cup \{(v_\ell, w_p) \mid \ell \in L, p \in \ell\}$. That is, G consists of a *source* vertex u , $m^2 + m$ *intermediate* vertices v_ℓ indexed by the lines in L , and m^2 *sink* vertices w_p indexed by the points in X . There are edges connecting the source vertex to every intermediate vertex and edges connecting each intermediate vertex v_ℓ to every sink vertex w_p such that $p \in \ell$. The cost associated with the former type of edges is 1, while the cost associated with the latter is 0.

Fix $k = m + 1$. Our Bayesian NCS game \mathcal{G} is defined over the directed graph G and includes k agents. Each type profile $t \in T$ is characterized by some line $\ell \in L$ and by some permutation π of $[m]$. The source of all agents is always the source vertex u . Under the type profile $t(\ell, \pi)$, the destination of agent $i \in [m]$ is w_p so that p is the $\pi(i)$ th point in ℓ ; the destination of agent $k = m + 1$ is v_ℓ .

	directed graphs		undirected graphs	
	universal	existential	universal	existential
$\frac{\text{opt}^P}{\text{opt}^C}$	$O(k)$ ≥ 1	$\Omega(k), n = \Theta(k^2)$ $= 1$	$O(\log n)$ ≥ 1	$\Omega(\log n), k = \Theta(n)$ $= 1$
$\frac{\text{best-eq}^P}{\text{best-eq}^C}$	$O(k)$ $\Omega(1/\log k)$	$\Omega(k), n = \Theta(k^2)$ $O(1/\log k), n = \Theta(k)$	$O(\min\{k, \log k \log n\})$ $\Omega(1/\log k)$	$\Omega(\log n), k = \Theta(n)$ $< 1, n = O(1)$
$\frac{\text{worst-eq}^P}{\text{worst-eq}^C}$	$O(k)$ $\Omega(1/k)$	$\Omega(k), n = O(1)$ $O(1/k), n = O(1)$	$O(k)$ $\Omega(1/k)$	$\Omega(k), n = O(1)$ $O(1/k), n = O(1)$

Table 1: Bounds for k -agent Bayesian NCS games in n -vertex (directed or undirected) graphs. The universal columns correspond to absolute bounds holding for all Bayesian NCS games, while the existential columns correspond to the existence of some (infinitely many) Bayesian NCS games that satisfy the desired bounds. For example, the $\frac{\text{best-eq}^P}{\text{best-eq}^C}$ ratio in directed graphs is always at most $O(k)$ and at least $\Omega(1/\log k)$. These bounds are tight since there exists a Bayesian NCS game defined over a $\Theta(k^2)$ -vertex (respectively, $\Theta(k)$ -vertex) directed graph for which $\frac{\text{best-eq}^P}{\text{best-eq}^C} = \Omega(k)$ (resp., $\frac{\text{best-eq}^P}{\text{best-eq}^C} = O(1/\log k)$).

Now, consider some strategy profile $s \in S$ and fix some agent $i \in [m]$. By symmetry considerations, given some type profile $t \in T$, characterized by the line $\ell \in L$, agent i buys the 'right' edge (u, v_ℓ) with probability $1/m$; otherwise, she buys some edge $(u, v_{\ell'})$, $\ell' \neq \ell$, in which case she must be the only agent buying this edge since by definition of the affine plane, if her destination under t is the sink vertex u_p , then p is the unique point in the intersection of ℓ and ℓ' . It follows that $K(s) = m(1 - 1/m) = m - 1$.

On the other hand, we argue that for every type profile $t \in T$ characterized by the line $\ell \in L$, the unique Nash equilibrium in \mathcal{G}_t is the action profile $a \in A$ under which all agents buy the edge (u, v_ℓ) ; observe that $K_t(a) = 1$. It is easy to verify that a is indeed a Nash equilibrium. For every action profile $a' \in A$ such that $K_t(a') < \infty$, every edge $(u, v_{\ell'})$, $\ell' \neq \ell$, is bought by at most one agent i (this, once again, follows from the definition of the affine plane). Since the 'right' edge (u, v_ℓ) is already bought by agent k , such an agent i is better off deviating from her previous action, for which the cost she incurred was 1, and switch to buying the 'right' edge for which she would pay at most $1/2$. The assertion follows. \square

The existential upper bound on the $\frac{\text{best-eq}^P}{\text{best-eq}^C}$ ratio in directed graphs is established in the next lemma.

LEMMA 3.3. *For every k_0 , there exist some $k \geq k_0$ and a k -agent Bayesian NCS game \mathcal{G} on a directed $\Theta(k)$ -vertex graph such that $\text{worst-eq}^P(\mathcal{G})/\text{best-eq}^C(\mathcal{G}) = O(1/\log k)$.*

PROOF. Consider the directed graph G_k illustrated in Figure 1. This graph was first presented by Anshelevich *et al.* [4] in order to establish a lower bound on the price of stability of NCS games. We design a k -agent Bayesian NCS game \mathcal{G} over G_k . Vertex x serves as a common source for all agents. The destination of agent i is vertex y_i for $1 \leq i \leq k - 1$ with probability 1. Agent k has vertex z as her destination with probability $1/2$; otherwise her destination is vertex x (same as her source).

We argue that the strategy profile $s = (s_1, \dots, s_k) \in S$ under which agent i buys the edges (x, z) and (z, y_i) for every $1 \leq i \leq k - 1$ is the unique Bayesian equilibrium in \mathcal{G} . Indeed, since the edge (x, z) is bought by agent k with probability $1/2$, agent 1 prefers the strategy s_1 over buying the edge (x, y_1) . By induction on i , agent i prefers




Figure 1: The directed graph G_k . The cost of the edge (x, y_i) is $1/i$ for every $1 \leq i \leq k - 1$; the cost of the edge (x, z) is $1 + \epsilon$; all edges (z, y_i) cost 0.

the strategy s_i over buying the (x, y_i) . Therefore, $K(s) = 1 + \epsilon$.

In the underlying games on the other hand, when agent k 's destination is x , the unique Nash equilibrium is the action profile $a \in A$ under which each agent $1 \leq i \leq k - 1$ buys the edge (x, y_i) . It follows that $\text{best-eq}^C(\mathcal{G}) > H(k - 1)/2 = \Omega(\log k)$. The assertion follows. \square

Next, Lemma 3.4 yields the universal upper bound on the $\frac{\text{opt}^P}{\text{opt}^C}$ ratio in undirected graphs.

LEMMA 3.4. *Every Bayesian NCS game \mathcal{G} on an undirected n -vertex graph satisfies $\text{opt}^P(\mathcal{G})/\text{opt}^C(\mathcal{G}) \leq O(\log n)$.*

PROOF. Let $G = (V, E)$ be some n -vertex undirected graph with edge costs $c(e)$. We say that a weighted tree $\tau = (V_\tau, E_\tau)$ with $V_\tau \supseteq V$ dominates G if $\text{dist}_\tau(u, v) \geq \text{dist}_G(u, v)$ for every two vertices $u, v \in V$, where $\text{dist}_\tau(\cdot, \cdot)$ and $\text{dist}_G(\cdot, \cdot)$ stand for distances in τ and G , respectively. Fakcharoenphol *et al.* [7] prove that there exists a probability distribution \mathcal{T} over the set of dominating trees of G such that $\mathbb{E}_{\tau \in \mathcal{T}}[\text{dist}_\tau(u, v)] = O(\log n) \cdot \text{dist}_G(u, v)$ for every two vertices $u, v \in V$. By employing the technique of Gupta [9], we may remove from the dominating trees τ the vertices not in V without increasing the distortion by more than a constant factor, so in what follows we assume that $V_\tau = V$ for

every dominating tree τ in the support of \mathcal{T} (τ may still include edges that do not exist in G and edges that do exist in G but with arbitrary weights).

Let $\tau = (V, E_\tau)$ be a random dominating tree of G chosen according to \mathcal{T} . For each edge $e = (u, v) \in E_\tau$, let P_e be some (designated) shortest (u, v) -path in G . This is extended to every vertex pair $x, y \in V$ by setting

$$P_{(x,y)} = \bigcup_{e \in \tau(x,y)} P_e,$$

where $\tau(x, y)$ denotes the unique (x, y) -path in τ . Consider the strategy profile s that instructs each agent $i \in [k]$ to buy the edges in $P_{(x_i, y_i)}$ for the type $(x_i, y_i) \in T_i$. Note that in order to execute the strategy profile s , we have to assume that the agents have access to public random bits (all agents have to agree on the same random dominating tree τ); this assumption will be lifted soon.

Fix some type profile $t = ((x_1, y_1), \dots, (x_k, y_k)) \in (V \times V)^k$. Let F_t be the (unique) minimal E_τ subset that includes an (x_i, y_i) -path for every $i \in [k]$ and let F_t^G be an E subset of minimum total cost among all E subsets that include an (x_i, y_i) -path for every $i \in [k]$. By definition, an action profile $a_t \in A$ that buys F_t^G is optimal for the underlying game \mathcal{G}_t . Note that the action profile $\{s_i(t_i)\}_{i \in [k]}$ buys $\bigcup_{e \in F_t} P_e$.

We will soon argue that $\mathbb{E}_\mathcal{T}[K(s, t)] = O(\log n) \cdot K_t(a_t)$. This implies that for every common prior $p \in \Delta(T)$, we have

$$\begin{aligned} \mathbb{E}_\mathcal{T}[K(s)] &= \mathbb{E}_\mathcal{T} \left[\sum_{t \in T} p(t) \cdot K(s, t) \right] \\ &= \sum_{t \in T} p(t) \cdot \mathbb{E}_\mathcal{T}[K(s, t)] \\ &= O(\log n) \cdot \sum_{t \in T} p(t) \cdot K_t(a_t). \end{aligned}$$

The assertion follows since there must exist some dominating tree τ in the support of \mathcal{T} that meets the expectation and satisfies the desired bound. In other words, the assumption that the agents need access to public random bits is lifted as promised.

It remains to show that $\mathbb{E}_\mathcal{T}[K(s, t)] = O(\log n) \cdot K_t(a_t)$. The random choice of τ guarantees that $\mathbb{E}[\text{dist}_\tau(e)] = O(\log n) \cdot c(e)$ for every edge $e \in F_t^G$, thus

$$\begin{aligned} O(\log n) \cdot K_t(a_t) &\geq \sum_{e \in a_t} \mathbb{E}_\mathcal{T}[\text{dist}_\tau(e)] \\ &= \mathbb{E}_\mathcal{T} \left[\sum_{e \in a_t} \text{dist}_\tau(e) \right] \geq \mathbb{E}_\mathcal{T} \left[\sum_{e \in F_t} c(e) \right]. \end{aligned}$$

Since τ dominates G , it follows that $c(e) \geq \sum_{e' \in P_e} c(e')$ for every tree edge $e \in F_t$, hence

$$O(\log n) \cdot K_t(a_t) \geq \mathbb{E}_\mathcal{T} \left[\sum_{e \in F_t} \sum_{e' \in P_e} c(e') \right] \geq \mathbb{E}_\mathcal{T}[K(s, t)],$$

as promised. \square

The following lemma implies the existential lower bound on the $\frac{\text{opt}^P}{\text{opt}^C}$ ratio in undirected graphs.

LEMMA 3.5. *For every n_0 , there exist some $n \geq n_0$ and a $\Theta(n)$ -agent Bayesian NCS game \mathcal{G} on an undirected n -vertex graph such that $\text{opt}^P(\mathcal{G})/\text{opt}^C(\mathcal{G}) = \Omega(\log n)$.*

PROOF. Our lower bound is established via a reduction from (a probabilistic variant of) the *online Steiner tree* problem. An instance of the online Steiner tree problem consists of an undirected graph $G = (V, E)$ with edge costs $c(e)$, a *root* vertex v_0 , and an input sequence $\sigma = \langle v_1, \dots, v_{|\sigma|} \rangle \in V^{\leq n}$, where $n = |V|$ and $V^{\leq n} = \{\sigma \in V^* : |\sigma| \leq n\}$. An online algorithm **ALG** receives the sequence σ step by step. In each step $1 \leq i \leq |\sigma|$, the algorithm must react by connecting v_i to the root vertex. This is done by buying some edge subset $F_i \subseteq E$ so that $\bigcup_{1 \leq j \leq i} F_j$ includes a path connecting v_i to v_0 . The cost incurred by **ALG** on the input sequence σ is defined as $\text{ALG}(\sigma) = \sum_{1 \leq i \leq |\sigma|} \sum_{e \in F_i} c(e)$. A randomized online Steiner tree algorithm **ALG** is said to be α -competitive if $\max_{\sigma \in V^{\leq n}} \mathbb{E}[\text{ALG}(\sigma)]/\text{OPT}(\sigma) \leq \alpha$, where $\text{OPT}(\sigma)$ is the cost of a minimum Steiner tree spanning all vertices in σ .

Imase and Waxman [11] establish an $\Omega(\log n)$ lower bound on the competitiveness of deterministic online Steiner tree algorithms for the n -vertex *diamond graph* $G = (V, E)$. Their lower bound can be generalized to hold against randomized algorithms by designing a probability distribution $q \in \Delta(V^{\leq n})$ so that $\sum_{\sigma \in V^{\leq n}} q(\sigma) \cdot \text{ALG}(\sigma) = \Omega(\log n)$ for every deterministic online Steiner tree algorithm **ALG**, while $\text{OPT}(\sigma) = O(1)$ for every sequence σ in the support of q .

Consider some n -vertex undirected graph $G = (V, E)$ with edge costs $c(e)$. We argue that for every probability distribution $q \in \Delta(V^{\leq n})$, there exists an n -agent Bayesian NCS game \mathcal{G}^q defined over the graph G and a deterministic online Steiner tree algorithm **ALG** ^{q} such that

$$\frac{\text{opt}^P(\mathcal{G}^q)}{\text{opt}^C(\mathcal{G}^q)} \geq \frac{\sum_{\sigma \in V^{\leq n}} q(\sigma) \cdot \text{ALG}(\sigma)}{\sum_{\sigma \in V^{\leq n}} q(\sigma) \cdot \text{OPT}(\sigma)}.$$

By the aforementioned lower bound on the competitiveness of randomized online Steiner tree algorithms, we obtain an $\Omega(\log n)$ lower bound on the ratio $\text{opt}^P(\mathcal{G}^q)/\text{opt}^C(\mathcal{G}^q)$ for n -agent Bayesian NCS games defined over the n -vertex diamond graph.

Fix some probability distribution $q \in \Delta(V^{\leq n})$. We design the n -agent Bayesian NCS game \mathcal{G}^q with a common prior p as follows. Each vertex sequence $\sigma = \langle v_1, \dots, v_{|\sigma|} \rangle \in V^{\leq n}$ corresponds to some type profile $t^\sigma = (t_1^\sigma, \dots, t_n^\sigma)$ such that

$$t_i^\sigma = \begin{cases} (v_i, v_0) & \text{if } 1 \leq i \leq |\sigma|; \\ (v_0, v_0) & \text{if } |\sigma| < i \leq n. \end{cases}$$

The common prior of t^σ is set to be $p(t^\sigma) = q(\sigma)$.

Now, suppose that $\text{opt}^P(\mathcal{G}^q)/\text{opt}^C(\mathcal{G}^q) \leq \alpha$ and let $s = (s_1, \dots, s_n) \in S$ be the strategy profile that realizes $\text{opt}^P(\mathcal{G}^q)$. We design the deterministic online Steiner tree algorithm **ALG** ^{q} as follows: in step $1 \leq i \leq |\sigma|$, **ALG** ^{q} buys all edges in $s_i(v_i, v_0)$ that were not bought beforehand. It follows that

$$\begin{aligned} \sum_{\sigma \in V^{\leq n}} q(\sigma) \cdot \text{ALG}^q(\sigma) &= \sum_{\sigma \in V^{\leq n}} p(t^\sigma) \cdot K(s, t^\sigma) = K(s) \\ &\leq \alpha \cdot \text{opt}^C(\mathcal{G}^q) \\ &= \alpha \sum_{\sigma \in V^{\leq n}} p(t^\sigma) \cdot \min_{a \in A} K_{t^\sigma}(a) \\ &= \alpha \sum_{\sigma \in V^{\leq n}} q(\sigma) \cdot \text{OPT}(\sigma) \end{aligned}$$

which completes the proof. \square

In fact, it is not difficult to show that if \mathcal{G} is the Bayesian NCS game obtained by following the construction described

in the proof of Lemma 3.5 for the diamond graph, then the action profile that minimizes K_t is a Nash equilibrium of the underlying game \mathcal{G}_t for every type profile $t \in T$. Therefore, the existential lower bound on the $\frac{\text{best-eq}^P}{\text{best-eq}^C}$ ratio in undirected graphs is also established. Moreover, by applying the same line of arguments to the construction of Alon and Azar [2], we obtain an existential $\Omega(\log k / \log \log k)$ lower bound on the $\frac{\text{opt}^P}{\text{opt}^C}$ ratio of k -agent Bayesian NCS games in the Euclidean plane⁴.

The following two lemmata, whose proofs are deferred to the full version of this paper, yield the existential lower and upper bounds on the $\frac{\text{worst-eq}^P}{\text{worst-eq}^C}$ ratio in both directed and undirected graphs.

LEMMA 3.6. *For every k_0 , there exist some $k \geq k_0$ and a k -agent Bayesian NCS game \mathcal{G} on an (un)directed $O(1)$ -vertex graph such that $\text{worst-eq}^P(\mathcal{G})/\text{worst-eq}^C(\mathcal{G}) = \Omega(k)$.*

LEMMA 3.7. *For every k_0 , there exist some $k \geq k_0$ and a k -agent Bayesian NCS game \mathcal{G} on an (un)directed $O(1)$ -vertex graph such that $\text{worst-eq}^P(\mathcal{G})/\text{worst-eq}^C(\mathcal{G}) = O(1/k)$.*

Anshelevich *et al.* [4] prove a logarithmic upper bound on the price of stability in (complete-information) NCS games. Combined with Observation 2.2, this yields the universal lower bounds on the $\frac{\text{best-eq}^P}{\text{best-eq}^C}$ ratio in directed and undirected graphs. The next lemma extends the technique of Anshelevich *et al.* to Bayesian games. Combined with Lemmata 3.1 and 3.4, it establishes the universal upper bound on the $\frac{\text{best-eq}^P}{\text{best-eq}^C}$ ratio in undirected graphs.

LEMMA 3.8. *Every k -agent Bayesian NCS game \mathcal{G} satisfies $\text{best-eq}^P(\mathcal{G}) \leq H(k) \cdot \text{opt}^P(\mathcal{G})$.*

PROOF. Consider some k -agent Bayesian NCS game \mathcal{G} on a (directed or undirected) graph $G = (V, E)$. Recall that the potential of a strategy profile $s \in S$ is given by

$$Q(s) = \sum_{t \in T} p(t) \cdot \sum_{e \in E} c(e) \cdot H(|\{i \in [k] \mid e \in s_i(t_i)\}|),$$

while the social cost of s is

$$K(s) = \sum_{t \in T} p(t) \cdot \sum_{e \in E} c(e) \cdot 1(\exists i \in [k] \text{ s.t. } e \in s_i(t_i)).$$

Thus, $Q(s)/H(k) \leq K(s) \leq Q(s)$ for every strategy profile $s \in S$. Let s be a strategy profile that minimizes $Q(s)$ and let s^* be a strategy profile that realizes $\text{opt}^P(\mathcal{G})$. The assertion is established by recalling that s is a Bayesian equilibrium and observing that $K(s) \leq Q(s) \leq Q(s^*) \leq K(s^*) \cdot H(k)$. \square

4. PUBLIC RANDOM BITS AS A SUBSTITUTE FOR THE COMMON PRIOR

In this section we show that in the presence of public random bits, the agents can waive the knowledge of the common prior and still guarantee the same $\frac{\text{opt}^P}{\text{opt}^C}$ ratio. Formally, consider some 4-tuple $\phi = \langle k, \{A_i\}_{i \in [k]}, \{T_i\}_{i \in [k]}, \{C_{i,t}\}_{i \in [k], t \in T} \rangle$, where $C_{i,t}(a) \in \mathbb{R}_{>0}$

⁴Of course, the Euclidean plane is not a finite graph, so for the sake of formality, one can replace it with a very fine grid.

for every $i \in [k]$, $t \in T$, and $a \in A$. Every common prior $p \in \Delta(T)$ defines the Bayesian game $\mathcal{G}^p = \langle \phi, p \rangle$. We write $\text{RoE}(\phi) = r$ if for every $p \in \Delta(T)$, there exists some strategy profile $s \in S$ such that

$$\frac{\sum_{t \in T} p(t) \cdot K(s, t)}{\sum_{t \in T} p(t) \cdot \min_{s' \in S} K(s', t)} \leq r$$

and this fails to hold for any $r' < r$. (**RoE** stands for *ratio of expectations*.) In other words, $\text{RoE}(\phi) = r$ if r is the minimum real such that $\frac{\text{opt}^P(\mathcal{G}^p)}{\text{opt}^C(\mathcal{G}^p)} \leq r$ for every $p \in \Delta(T)$. We write $\text{EoR}(\phi) = r$ if for every $p \in \Delta(T)$, there exists some strategy profile $s \in S$ such that

$$\sum_{t \in T} p(t) \cdot \frac{K(s, t)}{\min_{s' \in S} K(s', t)} \leq r$$

and this fails to hold for any $r' < r$. (**EoR** stands for *expectation of ratios*.) We begin by showing that $\text{RoE}(\phi) = \text{EoR}(\phi)$; this turns out to be a special case of the following proposition.

PROPOSITION 4.1. *Consider some matrix $M \in \mathbb{R}_{>0}^{m \times n}$ and some vector $v \in \mathbb{R}_{>0}^n$. Let r^* be the minimum real r such that*

$$\forall p \in \Delta(n), \exists i \in [m] \text{ s.t. } \frac{\sum_{j \in [n]} p(j) \cdot M(i, j)}{\sum_{j \in [n]} p(j) \cdot v(j)} \leq r$$

and let r^{**} be the minimum real r such that

$$\forall p \in \Delta(n), \exists i \in [m] \text{ s.t. } \sum_{j \in [n]} p(j) \cdot \frac{M(i, j)}{v(j)} \leq r.$$

Then $r^* = r^{**}$.

PROOF. We show that r^{**} must be at least as large as r^* by designing a probability distribution $p' \in \Delta(n)$ such that

$$\forall i \in [m], \sum_{j \in [n]} p'(j) \frac{M(i, j)}{v(j)} \geq r^*.$$

By definition, there exists some probability distribution $p \in \Delta(n)$ such that

$$\forall i \in [m], \frac{\sum_{j \in [n]} p(j) \cdot M(i, j)}{\sum_{j \in [n]} p(j) \cdot v(j)} \geq r^*.$$

We define p' by setting $p'(j) = \frac{p(j) \cdot v(j)}{\alpha'}$ for every $j \in [n]$, where $\alpha' = \sum_{j \in [n]} p(j) \cdot v(j)$. On the other hand, we show that r^* must be at least as large as r^{**} by designing a probability distribution $p'' \in \Delta(n)$ such that

$$\forall i \in [m], \frac{\sum_{j \in [n]} p''(j) \cdot M(i, j)}{\sum_{j \in [n]} p''(j) \cdot v(j)} \geq r^{**}.$$

By definition, there exists some probability distribution $p \in \Delta(n)$ such that

$$\forall i \in [m], \sum_{j \in [n]} p(j) \frac{M(i, j)}{v(j)} \geq r^{**}.$$

We define p'' by setting $p''(j) = \frac{p(j)}{\alpha'' v(j)}$ for every $j \in [n]$, where $\alpha'' = \sum_{j \in [n]} p(j)/v(j)$. \square

We conclude that public random bits can replace the knowledge of the common prior when bounding the ratio $\frac{\text{opt}^P}{\text{opt}^C}$ by establishing the following lemma.

LEMMA 4.2. *There exists a probability distribution $q \in \Delta(S)$ such that for every $p \in \Delta(T)$,*

$$\frac{\sum_{t \in T} p(t) \sum_{s \in S} q(s) \cdot K(s, t)}{\sum_{t \in T} p(t) \cdot \min_{s' \in S} K(s', t)} \leq r = \text{RoE}(\phi)$$

and this fails to hold for any $r < \text{RoE}(\phi)$.

PROOF. We shall establish this lemma for $r = \text{EoR}(\phi)$; the assertion follows due to Proposition 4.1. Note first that it is sufficient to consider common priors $p \in \Delta(T)$ which are concentrated on single type profiles, namely, to prove that there exists a probability distribution $q \in \Delta(S)$ such that for every $t \in T$,

$$\frac{\sum_{s \in S} q(s) \cdot K(s, t)}{\min_{s' \in S} K(s', t)} \leq r = \text{EoR}(\phi)$$

and that this fails to hold for any $r < \text{EoR}(\phi)$. Fix $K'(s, t) = K(s, t) / \min_{s' \in S} K(s', t)$ for every $s \in S$ and $t \in T$. By the definition of $\text{EoR}(\phi)$, we know that for every $p \in \Delta(T)$, there exists some $s \in S$ such that

$$\sum_{t \in T} p(t) \cdot K'(s, t) \leq r = \text{EoR}(\phi)$$

and that this does not hold for any $r < \text{EoR}(\phi)$. The assertion is then established by von Neumann's minimax theorem that guarantees the existence of some $q \in \Delta(S)$ such that for every $t \in T$,

$$\sum_{s \in S} q(s) \cdot K'(s, t) \leq r = \text{EoR}(\phi),$$

where this fails to hold for any $r < \text{EoR}(\phi)$. \square

It can be easily shown that this conclusion also holds in the limit if $C_{i,t}(a) \rightarrow 0$ for some $i \in [k]$, $t \in T$, and $a \in A$. That is, a probabilistic strategy profile that guarantees the best possible bound on the $\frac{\text{opt}^P}{\text{opt}^C}$ ratio still exists if we allow some of the costs to be 0 and think of 0/0 as 1. Note that in this case, the $\frac{\text{opt}^P}{\text{opt}^C}$ ratio is not always defined, but if it is defined, then it can be guaranteed by a single probabilistic strategy profile.

5. CONCLUSIONS

A quantitative approach for analyzing the implications of the agent's local views in Bayesian distributed settings is introduced. Our approach is applied to Bayesian NCS games, yielding an almost complete picture of the various asymptotic bounds. It would be interesting to close the remaining gaps in the best equilibria undirected case (see Table 1) and in a larger scope, to apply the Bayesian ignorance measures to other Bayesian games. Among our findings, we show that in a benevolent agents setting, the knowledge of the common prior can be replaced by public random bits. In this context it would be interesting to understand what can be achieved with *private* random bits.

The effect of Bayesian ignorance resembles the extensively studied "price of ..." measures, however, to the best of our knowledge our approach is the first to use such measures in attempt to compare the social cost in a Bayesian game to the expected social cost in the underlying complete-information games. Moreover, unlike the "price of ..." measures, the effect of Bayesian ignorance makes sense outside the game theoretic (selfish agents) context.

6. REFERENCES

- [1] S. Albers. On the Value of Coordination in Network Design. *SIAM J. Comput. (SICOMP)*, 38(6):2273–2302, 2009.
- [2] N. Alon and Y. Azar. On-line Steiner trees in the Euclidean plane. In *Proc. of the 8th Symp. on Computational Geometry (SoCG)*, pages 337–343, 1992.
- [3] N. Andelman, M. Feldman, and Y. Mansour. Strong Price of Anarchy. In *Proc. ACM-SIAM Symp. on Discrete Algorithms (SODA)*, 2007.
- [4] E. Anshelevich, A. Dasgupta, J.M. Kleinberg, É. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. *SIAM J. Comput. (SICOMP)*, 38(4):1602–1623, 2008.
- [5] I. Ashlagi, D. Monderer, and M. Tennenholtz. Two-Terminal Routing Games with Unknown Active Players. *Artificial Intelligence Journal*, 173(15):1441–1455, 2009.
- [6] H.L. Chen, and T. Roughgarden. Network design with weighted players. In *Proc. ACM Symp. on Parallel Algorithms and Architecture (SPAA)*, pages 29–38, 2006.
- [7] J. Fakcharoenphol, S. Rao, and K. Talwar. A tight bound on approximating arbitrary metrics by tree metrics. *J. Comput. Syst. Sci. (JCSS)*, 69(3):485–497, 2004.
- [8] I. Gilboa, and E. Lehrer. The Value of Information - An Axiomatic Approach. *Journal of Mathematical Economics*, 20: 443–459, 1991.
- [9] A. Gupta. Steiner points in tree metrics don't (really) help. In *Proc. of the 12th ACM-SIAM Symp. on Discrete Algorithms (SODA)*, pages 220–227, 2001.
- [10] R.A. Howard. Information value theory. *IEEE Transactions on Systems Science and Cybernetics*, 2: 22–26, 1996.
- [11] M. Imase and B.M. Waxman. Dynamic Steiner tree problem. *SIAM J. Discrete Math. (SIDMA)*, 4(3):369–384, 1991.
- [12] E. Koutsoupias, and C.H. Papadimitriou. Worst-case equilibria. In *Proc. 16th Symp. Theoretical Aspects of Computer Science (STACS)*, 1999.
- [13] C.H. Papadimitriou. Algorithms, games, and the Internet. In *Proc. 33rd Symp. Theory of Computing (STOC)*, 749–753, 2001.
- [14] R.W. Rosenthal. The network equilibrium problem in integers. *Networks*, 3(1):53–59, 1973.