

Malicious Bayesian Congestion Games^{*}

Martin Gairing

International Computer Science Institute, Berkeley, CA, USA.
gairing@icsi.berkeley.edu

Abstract. In this paper, we introduce *malicious Bayesian congestion games* as an extension to congestion games where players might act in a malicious way. In such a game each player has two *types*. Either the player is a rational player seeking to minimize her own delay, or – with a certain probability – the player is *malicious* in which case her only goal is to disturb the other players as much as possible.

We show that such games do in general not possess a Bayesian Nash equilibrium in pure strategies (i.e. a *pure Bayesian Nash equilibrium*). Moreover, given a game, we show that it is NP-complete to decide whether it admits a pure Bayesian Nash equilibrium. This result even holds when resource latency functions are linear, each player is malicious with the same probability, and all strategy sets consist of singleton sets of resources. For a slightly more restricted class of malicious Bayesian congestion games, we provide easy checkable properties that are necessary and sufficient for the existence of a pure Bayesian Nash equilibrium.

In the second part of the paper we study the impact of the malicious types on the overall performance of the system (i.e. the *social cost*). To measure this impact, we use the *Price of Malice*. We provide (tight) bounds on the Price of Malice for an interesting class of malicious Bayesian congestion games. Moreover, we show that for certain congestion games the advent of malicious types can also be beneficial to the system in the sense that the social cost of the worst case equilibrium decreases. We provide a tight bound on the maximum factor by which this happens.

1 Introduction

Motivation and Framework. Over the last decade, the study of strategic behavior in distributed systems has improved our understanding of modern computer artifacts such as the Internet. Normally, the users of such distributed systems are modeled as rational, utility optimizing players. However, in many real world scenarios, users do not necessarily act rationally, but rather *irrationally*. In this paper, we address one form of irrationality, namely, we allow that players act in a *malicious* way. In this case, the only goal of a malicious player is to disturb the (non-malicious) players as much as possible. The presence of *Denial of Service attacks* in the Internet is an example showing that such systems are quite realistic. In many such systems with malicious players, the players have only *incomplete information* about the set of malicious players. A standard approach for modeling games with incomplete information uses the Harsanyi transformation [14], which converts a game with incomplete information to a game where players have different *types*. The type of a player represents its private information that is not common knowledge to all players. In the resulting *Bayesian game*, each player’s uncertainty about each other’s type is described by a probability distribution.

One aspect of Game Theory that was studied extensively in recent years is the *Price of Anarchy* as introduced by Koutsoupias and Papadimitriou [16]. The Price of Anarchy is the worst case ratio between the value of the *social cost* in an equilibrium state of the system and that of some social optimum. Usually, the equilibrium state is defined as *Nash equilibrium* – a state in which no

^{*} This work was supported by a fellowship within the Postdoc-Programme of the German Academic Exchange Service (DAAD).

player can unilaterally improve her private objective function, also coined as *private cost*. A Nash equilibrium is *pure* if all players choose a pure strategy and *mixed* if players choose probability distributions over pure strategies.

While the celebrated result of Nash [20] guarantees the existence of a mixed Nash equilibrium for every finite game, pure Nash equilibria are not guaranteed to exist (see e.g. [9,12,17,18]). A natural question to ask, is whether a given game possesses a pure Nash equilibrium or not. We address this question by asking about the complexity of this decision problem.

A class of games that always possess pure Nash equilibria is the class of congestion games as introduced by Rosenthal [21]. Here, the strategy set of each player is a subset of the power set of given resources, the latency on each resource is described by a latency function in the number of players sharing this resource, and the private cost of each player is the sum of the latencies of its chosen resources. Milchtaich [18] considered weighted congestion games as an extension to congestion games in which the players have weights and thus different influence on the latency of the resources.

To measure the influence of malicious behavior, Moscibroda et al. [19] introduced the *Price of Byzantine Anarchy* as the worst case ratio between the social cost in an equilibrium state of the system under some assumption on the malicious players and the social cost of some social optimum without malicious players. They further define the *Price of Malice* as the ratio between the Price of Byzantine Anarchy and the Price of Anarchy. We will use a similar definition and define the equilibrium state as a Bayesian Nash equilibrium.

Contribution. In this paper, we introduce *malicious Bayesian congestion games* as an extension to congestion games where players might act in a malicious way. Following Harsanyi’s transformation [14], we allow each player to be of two *types*. Either the player is a rational player seeking to minimize her own delay, or – with a certain probability – the player is *malicious* in which case her only goal is to disturb the other players as much as possible. For such games we study the complexity of deciding whether a given game has a pure Bayesian Nash equilibrium. Moreover, we study the impact of the malicious types on the overall performance of the system (i.e. the *social cost*). To measure this impact, we use the *Price of Malice*, which we define similarly as Moscibroda et al. [19].

As our main result, we show that it is NP-complete to decide whether a given malicious Bayesian congestion game admits a pure Bayesian Nash equilibrium even if resource latency functions are linear and all strategy sets are singleton sets. We show that this result holds already for the very restricted case that each player is malicious with the same probability, and for the case that only one player is malicious with positive probability (Theorem 1). The same result even holds if we further restrict to the case that each player has at most four strategies and at most three players can be assigned to each resource (Theorem 2). For *symmetric* Bayesian congestion games with identical type probability, identical latency functions and strategy sets that consist only of singletons, we provide easy checkable properties that are necessary and sufficient for the existence of a pure Bayesian Nash equilibrium (Theorem 3).

We then shift gears and present results related to the Price of Malice. For general malicious Bayesian congestion games with linear latency functions, we show an upper bound on the Price of Byzantine Anarchy (Theorem 4). Moreover, we prove a lower bound on the same ratio that already holds for the case of identical type probabilities (Theorem 5). As a corollary, we get an asymptotic tight bound on the Price of Malice (Corollary 2). We close the paper with a tight lower bound on the maximum factor by which the social cost of a worst case (Bayesian) Nash equilibrium of a congestion game might decrease by introducing malicious types (Theorem 6).

Related Work. Congestion games and variants thereof have long been used to model non-cooperative resource sharing among selfish players. Rosenthal [21] showed that congestion games always possess pure Nash equilibria. The complexity of computing such a pure Nash equilibrium has been settled for arbitrary latency functions by Fabrikant et al. [8] and later for linear latency

functions by Ackermann et al. [1]. On the other hand, for weighted congestion games, Libman and Orda [17], Fotakis et al. [9] and Goemans et al. [12] provide examples that do not allow for a pure Nash equilibrium. Dunkel and Schulz [7] showed that it is NP-complete to decide the existence of a pure Nash equilibrium for a given weighted congestion game.

The Price of Anarchy for weighted congestion games has been studied extensively (see e.g. [3,2,5]). In case of linear latency functions, the Price of Anarchy is exactly $\frac{5}{2}$ for unweighted congestion games [5] and $1 + \Phi$ for weighted congestion games [3], where $\Phi = \frac{1+\sqrt{5}}{2}$ is the golden ratio. The exact value of the Price of Anarchy is also known for the case of polynomial latency functions [2]. For bounds on the Price of Anarchy of (weighted) congestion games with each strategy set being a singleton set of resources, we refer to [11] and references therein.

Several recent papers considered games allowing for malicious player behavior [4,15,19]. Moscibroda et al. [19] introduced the Price of Malice and gave bounds on the Price of Malice for a virus inoculation game where some of the players are malicious. In fact, our definition of Price of Malice is motivated by the corresponding definition from this paper. Karakostas et al. [15] and Babaioff et al. [4], study malicious player behavior in *non-atomic* congestion games. Here, each player from a continuum of infinitely many players controls only an infinitesimally small amount of weight and a fraction of those players is malicious. In contrast to those papers, our games are atomic, and thus have only finitely many players. This yields different results.

For general Bayesian games, questions concerning the complexity of deciding the existence of a pure Bayesian Nash equilibrium have been addressed in two recent works [6,13]. On the one hand, if the game is given in *standard normal form*, i.e. the utility functions and the type probability distribution are represented extensively as tables, then deciding the existence of a pure Bayesian Nash equilibrium is NP-complete [6]. On the other hand, if both – the utility functions and the type probability distribution – are succinctly encoded, then the problem becomes PP-complete [13]. In contrast to [6], malicious Bayesian congestion games are succinctly represented but they are more structured as the games considered by Gottlob et al. [13].

A certain class of Bayesian congestion game has been introduced in [10]. Here, players act completely rationally but they are uncertain about each other’s weight. Among other results, the authors show that such games always possess pure Bayesian Nash equilibria if latency functions are linear.

Roadmap. The rest of the paper is organized as follows. In Section 2, we introduce malicious Bayesian congestion games. In Section 3, we present our results on the complexity of deciding for pure Bayesian Nash equilibria, while Section 4 comprises our findings related to the Price of Malice. Due to lack of space, we omit some of the proofs. They can be found in the appendix.

2 Model

2.1 Congestion Games

Instance. A *congestion game* Γ is a tuple

$$\Gamma = (\mathcal{N}, E, (S_u)_{u \in \mathcal{N}}, (f_e)_{e \in E}).$$

Here, \mathcal{N} is the set of *players* and E is the finite set of *resources*. Throughout, we denote $n = |\mathcal{N}|$ and $r = |E|$ and assume $n \geq 2$ and $r \geq 2$. For every player $u \in \mathcal{N}$, $S_u \subseteq 2^E$ is the *strategy set* of player u . Denote $S = S_1 \times \dots \times S_n$. For every resource $e \in E$, the *latency function* $f_e : \mathbb{N} \rightarrow \mathbb{R}$ is a non-negative, non-decreasing function that describes the *latency* on resource e . For most of our results, we consider *affine latency functions* with non-negative coefficients, that is, for all resources $e \in E$, the latency function is of the form $f_e(\delta) = a_e \cdot \delta + b_e$ with $a_e, b_e \geq 0$. Affine latency functions

are *linear* if $b_e = 0$ for all $e \in E$. A congestion game is called *symmetric*, if $S_u = S'_u$ for any pair of players u, u' .

Strategies and Strategy Profiles. A *pure strategy* for player u is some specific strategy $s_u \in S_u$, while a *mixed strategy* $Q_u = (q(u, s_u))_{s_u \in S_u}$ is a probability distribution over S_u , where $q(u, s_u)$ denotes the probability that player u chooses the pure strategy s_u .

A *pure strategy profile* is an n -tuple $\mathbf{s} = (s_1, \dots, s_n)$ whereas a *mixed strategy profile* $\mathbf{Q} = (Q_1, \dots, Q_n)$ is represented by an n -tuple of mixed strategies. For a mixed strategy profile \mathbf{Q} , denote by $q(\mathbf{s}) = \prod_{u \in \mathcal{N}} q(u, s_u)$ the probability that the players choose the pure strategy profile \mathbf{s} .

Load and Private Cost. For a pure strategy profile \mathbf{s} , denote by $\delta_e(\mathbf{s}) = |\{u \in \mathcal{N} : e \in s_u\}|$ the *load* on resource $e \in [m]$, i.e. the number of players assigned to e . In the same way, for a partial strategy profile \mathbf{s}_{-i} , denote $\delta_e(\mathbf{s}_{-i}) = |\{u \in \mathcal{N} \setminus \{i\} : e \in s_u\}|$ the *load* on resource $e \in [m]$ without player i .

Fix a pure strategy profile \mathbf{s} . The *private cost* $\text{PC}_u(\mathbf{s})$ of player $u \in \mathcal{N}$ is defined by the *latency* of the chosen resources. Thus

$$\text{PC}_u(\mathbf{s}) = \sum_{e \in s_u} f_e(\delta_e(\mathbf{s})).$$

For a mixed strategy profile \mathbf{Q} , the *private cost* of player $u \in \mathcal{N}$ is

$$\text{PC}_u(\mathbf{Q}) = \sum_{\mathbf{s} \in S} q(\mathbf{s}) \cdot \text{PC}_u(\mathbf{s}).$$

Social Cost. Associated with a congestion game Γ and a mixed strategy profile \mathbf{Q} is the *social cost* $\text{SC}(\Gamma, \mathbf{Q})$ as a measure of social welfare. In particular we use the expected average latency. That is,

$$\begin{aligned} \text{SC}(\Gamma, \mathbf{Q}) &= \frac{1}{n} \sum_{u \in \mathcal{N}} \text{PC}_u(\mathbf{Q}) \\ &= \frac{1}{n} \sum_{u \in \mathcal{N}} \sum_{\mathbf{s} \in S} q(\mathbf{s}) \sum_{e \in s_u} f_e(\delta_e(\mathbf{s})) \\ &= \frac{1}{n} \sum_{\mathbf{s} \in S} q(\mathbf{s}) \sum_{e \in E} \delta_e(\mathbf{s}) \cdot f_e(\delta_e(\mathbf{s})). \end{aligned}$$

Observe, that this measure differs from the *total latency* [22] only by the factor n .

The *optimum* associated with a congestion game Γ is the least possible social cost, over all pure strategy profiles $\mathbf{s} \in S$. Thus,

$$\text{OPT}(\Gamma) = \min_{\mathbf{s} \in S} \text{SC}(\Gamma, \mathbf{s}).$$

Nash Equilibria. Given a congestion game and an associated mixed strategy profile \mathbf{Q} , player $u \in \mathcal{N}$ is *satisfied* if the player cannot improve its private cost by unilaterally changing its strategy. Otherwise, player u is *unsatisfied*. The mixed strategy profile \mathbf{Q} is a *Nash equilibrium* if and only if all players $u \in \mathcal{N}$ are satisfied, that is, $\text{PC}_u(\mathbf{Q}) \leq \text{PC}_u(\mathbf{Q}_{-u}, s_u)$ for all $u \in \mathcal{N}$ and $s_u \in S_u$.

Depending on the type of strategy profile we distinguish between *pure* and *mixed* Nash equilibria.

Price of Anarchy. Let \mathcal{G} be a class of congestion games. The *Price of Anarchy*, also called *coordination ratio* and denoted by PoA , is the supremum, over all instances $\Gamma \in \mathcal{G}$ and Nash equilibria \mathbf{Q} , of the ratio $\frac{\text{SC}(\Gamma, \mathbf{Q})}{\text{OPT}(\Gamma)}$. Thus,

$$\text{PoA} = \sup_{\Gamma \in \mathcal{G}, \mathbf{Q}} \frac{\text{SC}(\Gamma, \mathbf{Q})}{\text{OPT}(\Gamma)}.$$

2.2 Malicious Bayesian Congestion Games

Instance. A *malicious Bayesian congestion game* Ψ is an extension of congestion games, where each player is malicious with a certain probability. Following Harsanyi's approach, we model such a game with incomplete information as a Bayesian game, where each player $u \in \mathcal{N}$ can be of two types: Either u is *selfish* or *malicious*. For each type of player $u \in \mathcal{N}$ we introduce two independent type-agents u^s and u^m , denoting the *selfish* and *malicious type-agent* of player u , respectively.

Let p_u be the probability that player $u \in \mathcal{N}$ is malicious and call p_u the *type probability* of player u . Define the *type probability vector* $\mathbf{p} = (p_1, \dots, p_n)$ in the natural way. Denote $p_{\min} = \min_{u \in \mathcal{N}} p_u$. In the case of identical type probabilities $p_u = p$ for all player $u \in \mathcal{N}$. Define $\Delta = \sum_{u \in \mathcal{N}} p_u$ as the *expected number of malicious players*. Observe, that for identical type probabilities $\Delta = p \cdot n$. Denote by Γ_Ψ the congestion game that arises from the malicious Bayesian congestion game Ψ by setting $p_u = 0$ for all player $u \in \mathcal{N}$.

Summing up, a malicious Bayesian congestion game Ψ is given by a tuple

$$\Psi = (\mathcal{N}, E, (S_u)_{u \in \mathcal{N}}, (p_u)_{u \in \mathcal{N}}, (f_e)_{e \in E}).$$

Strategies and Strategy Profiles. A pure strategy σ_u for player $u \in \mathcal{N}$ is now a tuple $\sigma_u = (\sigma(u^s), \sigma(u^m)) \in S_u^2$, where $\sigma(u^s)$ and $\sigma(u^m)$ denote the strategy of the selfish type-agent and malicious type-agent of player u , respectively. Denote $\boldsymbol{\sigma} = (\sigma_1, \dots, \sigma_n)$. A *mixed strategy* Q_i is now a probability distribution over $S_i \times S_i$. Define \mathbf{Q} and $q(\boldsymbol{\sigma})$ as before.

Private Cost. For any type probability vector \mathbf{p} and pure strategy profile $\boldsymbol{\sigma}$, denote the *expected selfish load* on resource $e \in E$ by $\delta_e(\boldsymbol{\sigma}) = \sum_{u \in \mathcal{N}: e \in \sigma(u^s)} (1 - p_u)$ and the *expected malicious load* by $\kappa_e(\boldsymbol{\sigma}) = \sum_{u \in \mathcal{N}: e \in \sigma(u^m)} p_u$. For a partial assignment $\boldsymbol{\sigma}_{-u}$ define $\delta_e(\boldsymbol{\sigma}_{-u})$ and $\kappa_e(\boldsymbol{\sigma}_{-u})$ accordingly, by disregarding player u .

Fix any type probability vector \mathbf{p} and pure strategy profile $\boldsymbol{\sigma}$. The *private cost* $\text{PC}_u(\mathbf{p}, \boldsymbol{\sigma})$ of player $u \in \mathcal{N}$ is defined by

$$\text{PC}_u(\mathbf{p}, \boldsymbol{\sigma}) = \sum_{e \in \sigma(u^s)} f_e (\delta_e(\boldsymbol{\sigma}_{-u}) + \kappa_e(\boldsymbol{\sigma}_{-u}) + 1).$$

In other words $\text{PC}_u(\mathbf{p}, \boldsymbol{\sigma})$ is the expected latency that player u experiences if player u is selfish. For each player $u \in \mathcal{N}$, type-agent u^s aims to minimize $\text{PC}_u(\mathbf{p}, \boldsymbol{\sigma})$. Observe, that $\text{PC}_u(\mathbf{p}, \boldsymbol{\sigma})$ does not depend on $\sigma(u^m)$. For a mixed strategy profile \mathbf{Q} , define $\text{PC}_u(\mathbf{p}, \mathbf{Q})$ accordingly.

Social Cost. Let Ψ be a malicious Bayesian congestion game with type probability vector \mathbf{p} and let \mathbf{Q} be a mixed strategy profile for Ψ . We generalize the definition of *social cost* $\text{SC}(\Psi, \mathbf{Q})$ to the weighted average latency of the selfish type-agents. That is,

$$\text{SC}(\Psi, \mathbf{Q}) = \frac{\sum_{u \in \mathcal{N}} (1 - p_u) \cdot \text{PC}_u(\mathbf{p}, \mathbf{Q})}{n - \Delta}.$$

Bayesian Nash equilibria. A selfish type-agent is *satisfied* if she cannot unilaterally decrease her private cost, that is, $\text{PC}_u(\mathbf{Q}) \leq \text{PC}_u(\mathbf{Q}_{-u^s}, \sigma(u^s))$ for all $u \in \mathcal{N}$ and $\sigma(u^s) \in S_u$.

In contrast to the selfish type-agents, each malicious type-agent aims to maximize the social cost. So, a malicious type-agent is *satisfied* if she cannot increase the social cost by unilaterally changing her strategy.

For a malicious Bayesian congestion game, a mixed strategy profile \mathbf{Q} is a *Bayesian Nash equilibrium* if and only if both type-agents of all players $u \in \mathcal{N}$ are satisfied. Depending on the type of strategy profile we again differ between *pure* and *mixed* Bayesian Nash equilibria.

Price of Byzantine Anarchy and Price of Malice. For a fixed expected number of malicious players Δ , let $\mathcal{G}(\Delta)$ be the class of malicious Bayesian congestion games where $\sum_{u \in \mathcal{N}} p_u = \Delta$.

Similarly to [19], we define the *Price of Byzantine Anarchy*, denoted by PoB , as the supremum, over all instances $\Psi \in \mathcal{G}(\Delta)$ and Bayesian Nash equilibria \mathbf{Q} , of the ratio between the social cost in \mathbf{Q} and the optimum social cost of the corresponding congestion game Γ_Ψ . Thus,

$$\text{PoB}(\Delta) = \sup_{\Psi \in \mathcal{G}(\Delta), \mathbf{Q}} \frac{\text{SC}(\Psi, \mathbf{Q})}{\text{OPT}(\Gamma_\Psi)}.$$

Observe that for $\Delta = 0$, the Price of Byzantine Anarchy $\text{PoB}(0)$ reduces to the Price of Anarchy PoA as defined in Section 2.1.

Again similarly to [19], we define the *Price of Malice* by

$$\text{PoM}(\Delta) = \frac{\text{PoB}(\Delta)}{\text{PoB}(0)}.$$

3 Existence and Complexity of pure Bayesian Nash equilibria

In this section, we study the complexity of deciding whether a given malicious Bayesian congestion game possesses a pure Bayesian Nash equilibrium or not.

Theorem 1. *The problem of deciding whether a malicious Bayesian congestion game with linear latency functions possesses a pure Bayesian Nash equilibrium is NP-complete, even if all strategy sets consist of singletons and either of the following properties holds:*

- (a) *All players are malicious with the same probability p for any $0 < p < 1$.*
- (b) *Only one player is malicious with positive probability p for any $0 < p \leq 1$.*

Proof. Our proof uses a reduction from a restricted version of 3-SAT. Here, 3-SAT is restricted to instances where each clause is a disjunction of 2 or 3 variables and each variable occurs at most three times. Tovey [23] showed that it is NP-complete to decide the satisfiability of such instances. Consider an arbitrary instance of 3-SAT with set of variables $X = \{x_1, \dots, x_\ell\}$ and set of clauses $C = \{c_1, \dots, c_k\}$. Without loss of generality, we may assume that each variable occurs at most twice unnegated and at most twice negated.

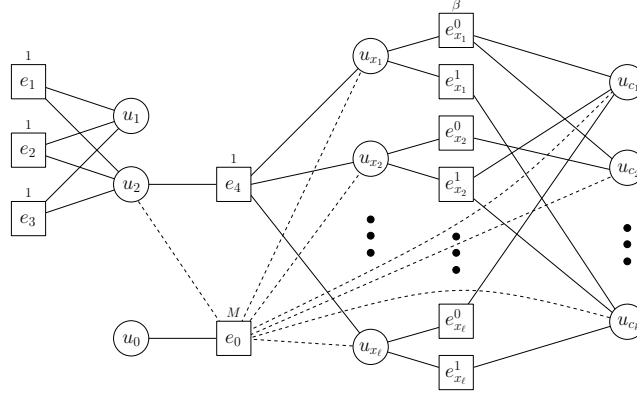


Fig. 1. Construction for the proof of Theorem 1

Part (a): We will construct a malicious Bayesian congestion game with singleton strategy sets and identical type probability p . Our construction imposes one player u_c for each clause $c \in C$,

one player u_x and two resources e_x^0, e_x^1 for each variable $x \in X$, 3 additional players u_0, u_1, u_2 , and 5 additional resources e_0, e_1, e_2, e_3, e_4 . Our construction is summarized in Figure 1. Resources are depicted as squares and players as circles and an edge (solid or dotted) between a resource e and a player u indicates that $\{e\}$ is in u 's strategy set. A number α above a resource e defines the slope of the corresponding linear latency function $f_e(\delta) = \alpha \cdot \delta$. Denote $E_v = \{e_{x_1}^0, e_{x_1}^1, \dots, e_{x_\ell}^0, e_{x_\ell}^1\}$. For the proof of part (a), let $\beta = 2 - p$. So, all resources $e \in E_v$ share the latency function $f_e(\delta) = (2 - p) \cdot \delta$.

Player u_0 can only be assigned to e_0 . Both u_0 and e_0 are used to collect the malicious type-agents of all players except player u_1 . Thus all those players have e_0 in their strategy set and M is chosen sufficiently large, such that for all those malicious type-agents e_0 is a dominant strategy and no selfish type other than u_0^s will ever prefer to choose e_0 . Choosing $M = \ell + 1$ suffices. Player u_1 and u_2 are connected to e_1, e_2 , and e_3 , while u_2 can also choose e_0 and e_4 . For each variable $x \in X$, the corresponding *variable player* u_x is connected to e_0, e_4, e_x^0 and e_x^1 . Assigning the selfish type-agent u_x^s to e_x^0 (resp. e_x^1) will be interpreted as setting x to true (resp. false). For each clause $c \in C$, the corresponding *clause player* u_c is connected to e_0 and to all resources e_x^0 (e_x^1) with $x \in X$ and x appears *negated* (*unnegated*) in c . For the example in Figure 1, $c_1 = (\bar{x}_1 \vee x_2 \vee \bar{x}_\ell)$, $c_2 = (\bar{x}_1 \vee \bar{x}_2)$, and $c_k = (x_1 \vee x_2 \vee x_\ell)$. Observe that by the structure of our 3-SAT instance, no more than two clause players are connected to each resource in E_v . This completes the construction of the malicious Bayesian congestion game.

We will first show that if the 3-SAT instance is satisfiable then the corresponding Bayesian congestion game possesses a pure Bayesian Nash equilibrium. Given a satisfying truth assignment, we define a strategy profile σ of the malicious Bayesian congestion game as follows:

- Both type-agents of player u_0 can only be assigned to e_0 .
- All malicious type-agents except u_1^m are assigned to resource e_0 . By the choice of M , none of those malicious type-agents can improve.
- Both type-agents of player u_1 are assigned to e_1 and no type-agent of any player is assigned to e_2 or e_3 . It is easy to see that neither u_1^m nor u_1^s have an incentive to switch.
- Type agent u_2^s is the only type-agent assigned to e_4 . So, u_2^s cannot improve.
- For each $x \in X$, the selfish type-agent u_x^s of variable player u_x is assigned to resource e_x^0 if $x = \text{true}$ in the satisfying truth assignment, and to e_x^1 otherwise. Each of these selfish type-agents is the only type-agent assigned to her resource. So, they all experience an expected latency of $\beta = 2 - p$ and changing to e_4 would yield the same expected latency. Thus, the selfish type-agents of all variable players are satisfied.
- Denote by E'_v the subset of resources from E_v to which no selfish type-agent of a variable player is assigned. Since we have a satisfying truth assignment, each clause player is connected to some resource from E'_v . For each $c \in C$, the selfish type-agent u_c^s is assigned to some resource in E'_v as follows:

Consider the sub-game that consists only of the selfish type-agents of the clause players u_c , $c \in C$ and the set of resources E'_v . Observe that this sub-game is a (non-malicious) congestion game and thus admits a pure Nash equilibrium [21]. Assign the selfish type-agents of each clause player according to this Nash equilibrium. So, none of these selfish type-agents can improve by changing to some other resource in E'_v . Moreover, at most two selfish type-agents are assigned to each resource in E'_v and there is exactly one selfish type-agent of a variable player assigned to each resource in $E_v \setminus E'_v$. Thus, the selfish type-agents of all clause players are satisfied.

Since no type-agent can improve by changing her strategy, it follows that σ is a pure Bayesian Nash equilibrium.

For the other direction observe that any pure Bayesian Nash equilibrium σ fulfills the following structural properties:

- (I) All malicious type-agents except u_1^m are assigned to resource e_0 and u_0^s is the only selfish type-agent assigned to e_0 .

(II) The selfish type-agent u_2^s is assigned to e_4 and no other type-agent is assigned to e_4 .

Property (I) follows immediately by the choice of M . We will now prove property (II).

By way of contradiction assume that u_2^s is assigned to a resource in $\{e_1, e_2, e_3\}$ in a pure Bayesian Nash equilibrium σ . In this case u_1^m will always choose the same resource as u_2^s . However, then there must be an empty resource in $\{e_1, e_2, e_3\}$ and u_2^s can improve by choosing this empty resource. This contradicts our assumption that σ is a pure Bayesian Nash equilibrium. Thus, u_2^s is assigned to e_4 . If some other type-agent is also assigned to e_4 , then u_2^s experiences an expected latency of at least $2 - p$ and u_2^s could decrease her expected latency to 1 by switching to the empty resource in $\{e_1, e_2, e_3\}$. Again a contradiction to σ being a pure Bayesian Nash equilibrium. It follows that u_2^s is the only type-agent assigned to e_4 in σ . This completes the proof of property (II).

Since u_2^s is the only type-agent assigned to e_4 it follows that for each variable $x \in X$ the corresponding selfish type-agent u_x^s is either assigned to e_x^0 or to e_x^1 . If u_x^s is not the only type-agent on that resource then her expected latency is at least $(2 - p)^2$ while changing to e_4 would improve her expected latency to $2 - p$, a contradiction to σ being a pure Bayesian Nash equilibrium. It follows that the selfish type-agents of all clause players are only assigned to resources in E_v to which no selfish type-agent of a variable player is assigned. This is only possible if the strategies of the selfish type-agents $u_x^s, x \in X$ correspond to a satisfying truth assignment. This completes the proof of part (a).

Part (b): To see that (b) holds we alter the construction depicted in Figure 1 slightly by deleting player u_0 and resource e_0 . Furthermore, in the new construction player u_1 is the only player that is malicious with positive probability p . For the slope of the latency functions of resources in E_v , let $\beta = \frac{3}{2}$ (in fact any $1 < \beta < 2$ would also do). The rest of the construction does not change. The proof now follows the same line of arguments as in part (a) with only minor changes. \square

Theorem 2. *The results from Theorem 1 hold, even if $|S_u| \leq 4$ for all players $u \in \mathcal{N}$ and for each resource $e \in E$ there are at most three players $u \in \mathcal{N}$ with $\{e\} \in S_u$.*

Proof (Sketch). We will slightly alter the construction from Figure 1. First observe that we have already $|S_u| \leq 4$ for all players $u \in \mathcal{N}$. Furthermore, the only resources that are in the strategy set of more than three players are e_4 and for part (a) also e_0 .




Fig. 2. Tree for $\ell = 3$

To resolve this for e_4 , disconnect all players from e_4 and replace the single resource e_4 with a binary tree of resources with root $e_{4,0}$ that has ℓ leaves $e_{4,x_1}, \dots, e_{4,x_\ell}$, all with depth $\lceil \log(\ell) \rceil$. For a resource e at level j the latency function is defined by $f_e(\delta) = \beta^j \cdot \delta$. So $f_{e_{4,0}}(\delta) = 1$ and $f_{e_{4,x}}(\delta) = \beta^{\lceil \log(\ell) \rceil} \cdot \delta$ for all leaves $x \in X$. For each pair of resources from two consecutive levels, we introduce a new player to connect them. Call those players *tree players*. Figure 2 shows the construction for $\ell = 3$. Player u_2 is connected to resource $e_{4,0}$ and each variable player $x \in X$ is connected to $e_{4,x}$. We also change the latency function of all resources $e \in E_v$ (cf. Theorem 1) to $f_e(\delta) = \beta^{\lceil \log(\ell) \rceil} \cdot \delta$.

Moreover, for part (a) we have to resolve that more than three players are connected to e_0 . To do so, we simply copy resource e_0 together with player u_0 multiple times and connect all players (including the tree players) except player u_1 to the new set of resources that evolve from e_0 . By having sufficiently many copies, this can be done, such that no more than three players are connected to each new resource. Again, M is chosen sufficiently large, e.g. $M = 2^{\lceil \log(\ell) \rceil + 1}$.

Observe that u_2^s will only selfishly choose $e_{4,0}$ if all tree players choose the strategy that is closer to the leaves. The rest of the proof now simply follows the proof of Theorem 1. \square

For the more restricted class of symmetric malicious Bayesian congestion game with singleton strategy sets, identical type probability p and identical latency functions we can easily decide whether a pure Bayesian Nash equilibrium exists or not.

Theorem 3. *A symmetric malicious Bayesian congestion game with singleton strategy sets, identical type probability p and identical (not necessarily linear) latency functions possesses a pure Bayesian Nash equilibrium if and only if either of the following properties holds:*

- (a) $p \leq \frac{1}{2}$ and $r = 2$
- (b) $p \leq \frac{1}{2}$ and $r|n$

Observe, that the proof of Theorem 3 is constructive. So, if the requirements for the existence of a pure Bayesian Nash equilibrium are fulfilled, then this equilibrium can also be easily constructed in linear time.

4 Price of Malice

We now shift gears and present our results that are related to the Price of Malice. We start with a general upper bound on the Price of Byzantine Anarchy. The proof of this upper bound uses a technique from [5] adapted to the model of malicious Bayesian congestion games.

Theorem 4. *Consider the class of malicious Bayesian congestion games $\mathcal{G}(\Delta)$ with affine latency functions. Then,*

$$PoB(\Delta) \leq \frac{n}{n - \Delta} (1 - p_{\min}) \left(\Delta + \frac{3 + \sqrt{5 + 4\Delta}}{2} \right).$$

For the case of identical type probabilities we can provide a better upper bound on the Price of Byzantine Anarchy. Observe that for identical type probabilities, $\Delta = p \cdot n$ and $p_{\min} = p$. As an immediate corollary to Theorem 4, we get:

Corollary 1. *Consider the class of malicious Bayesian congestion games $\mathcal{G}(\Delta)$ with affine latency functions and identical type probability p . Then,*

$$PoB(\Delta) \leq \Delta + \frac{3 + \sqrt{5 + 4\Delta}}{2}.$$

We proceed by introducing a malicious Bayesian congestion game that is parameterized by a parameter α . In the remainder of the paper, we will make use of this construction twice, each time with a different parameter α .

Example 1. Given some $\alpha > 0$, construct a malicious Bayesian congestion game $\Gamma(\alpha)$ with linear latency functions, $n \geq 3$ players and identical type probability p and $|E| = 2n$ as follows: Let $E = E_1 \cup E_2$ with $E_1 = \{g_1, \dots, g_n\}$ and $E_2 = \{h_1, \dots, h_n\}$. Each player $u \in \{1, \dots, n\}$ has three strategies in her strategy set. So, $S_u = \{s_u^1, s_u^2, s_u^3\}$ with $s_u^1 = \{g_u, h_u\}$, $s_u^2 = \{g_{u+1}, h_{u+1}, h_{u+2}\}$ and $s_u^3 = E_1 \cup E_2$, where $g_j = g_{j-n}$ and $h_j = h_{j-n}$ for $j > n$.

Each resource $e \in E_1$ has a latency function $f_e(\delta) = \alpha \cdot \delta$ whereas the resources $e \in E_2$ share the identity as their latency function, i.e. $f_e(\delta) = \delta$.

We make use of Example 1 to show a lower bound on the Price of Byzantine Anarchy:

Theorem 5. *Consider the class of malicious Bayesian congestion games $\mathcal{G}(\Delta)$ with linear latency functions and identical type probability p . Then, $\text{PoB}(\Delta) \geq \Delta + 2$.*

Proof. Consider the malicious Bayesian congestion game $\Psi = \Psi(\alpha)$ given in Example 1 with $\alpha = \frac{1+(n-1)p}{1-p}$. Observe that $\Delta = n \cdot p$.

Obviously, the optimum allocation \mathbf{s}^* for the corresponding non-malicious game Γ_Ψ is for each player $u \in \mathcal{N}$ to choose strategy s_u^1 . This yields $\text{SC}(\Gamma_\Psi, \mathbf{s}^*) = 1 + \alpha = \frac{2+(n-2)p}{1-p}$.

On the other hand, if $\sigma(u^m) = s_u^3$ and $\sigma(u^s) = s_u^2$ for all player $u \in \mathcal{N}$, then σ is a (pure) Bayesian Nash equilibrium for Ψ , with

$$\text{SC}(\Psi, \sigma) = 2(1 + (1-p) + (n-1)p) + (1 + (n-1)p)\alpha = \frac{2(1-p)(2 + (n-2)p) + (1 + (n-1)p)^2}{1-p}$$

It follows that

$$\begin{aligned} \frac{\text{SC}(\Psi, \sigma)}{\text{SC}(\Gamma_\Psi, \mathbf{s}^*)} &= 2(1-p) + \frac{(1 + (n-1)p)^2}{2 + (n-2)p} = 2(1-p) + \frac{1 + (n-1)p(2 + (n-1)p)}{2 + (n-2)p} \\ &> 2 - 3p + n \cdot p = \Delta + 2 - 3p. \end{aligned}$$

The Theorem follows for $p \rightarrow 0$, which implies $n \rightarrow \infty$. \square

Recall that the Price of Anarchy of (non-malicious) congestion games with affine latency functions is $\frac{5}{2}$ [5]. By combining this with Corollary 1 and Theorem 5 we get:

Corollary 2. *Consider the class of malicious Bayesian congestion games $\mathcal{G}(\Delta)$ with affine latency functions and identical type probability. Then, $\text{PoM}(\Delta) = \Theta(\Delta)$.*

For certain congestion games, introducing malicious types might also be beneficial to the system, in the sense that the social cost of the worst case equilibrium (one that maximizes social cost) decreases. To capture this, we define the *Windfall of Malice*. The term Windfall of Malice is due to [4]. For a malicious Bayesian congestion game Ψ , denote $\text{WoM}(\Psi)$ as the ratio between the social costs of the worst case Nash equilibrium of the corresponding congestion game Γ_Ψ and the worst case Bayesian Nash equilibrium of Ψ . We show,

Theorem 6. *For each $\epsilon > 0$ there is a malicious Bayesian congestion game Ψ with linear latency functions and identical type probability, such that $\text{WoM}(\Psi) \geq \frac{5}{2} - \epsilon$.*

Proof. Consider the malicious Bayesian congestion game $\Psi = \Psi(\alpha)$ given in Example 1 with $n = 3$ and $\alpha = 1$. This game (for $n \geq 3$) was already used in [5] to prove a lower bound on the Price of Anarchy for the corresponding non-malicious congestion games. For the congestion game Γ_Ψ that corresponds to Ψ , all players u choosing s_u^2 is a Nash equilibrium \mathbf{s} that maximizes social cost and $\text{SC}(\Gamma_\Psi, \mathbf{s}) = 5$.

Now, consider the malicious Bayesian congestion game Ψ , where $p > 0$. First observe that choosing s_u^3 is always a strictly dominant strategy for the malicious type-agent u^m for all $u \in \mathcal{N}$. Moreover, u^s will never choose s_u^3 . For $i \in \{2, 3\}$, let q_i be the probability that u_i^s chooses $s_{u_i}^1$. Then u_i^s chooses $s_{u_i}^2$ with probability $(1 - q_i)$. We will show that for all $p > 0$, the selfish type-agent u_1^s experiences a strictly lower expected latency, if she chooses $s_{u_1}^1$ and not $s_{u_1}^2$.

On the one hand, if u_1^s chooses $s_{u_1}^1$ then her expected latency is:

$$\underbrace{1 + 2p + (1 - q_2 + 1 - q_3)(1 - p)}_{h_1} + \underbrace{1 + 2p + (1 - q_2)(1 - p)}_{g_1} = 2 + 4p + (1 - p)(3 - 2q_2 - q_3)$$

On the other hand, if u_1^s chooses $s_{u_1}^2$ then her expected latency is:

$$\begin{aligned} & \underbrace{1 + 2p + (q_2 + 1 - q_3)(1 - p)}_{h_3} + \underbrace{1 + 2p + q_3(1 - p)}_{g_2} + \underbrace{1 + 2p + (1 - q_2 + q_3)(1 - p)}_{h_2} \\ & = 3 + 6p + (1 - p)(2 + q_3). \end{aligned}$$

However,

$$3 + 6p + (1 - p)(2 + q_3) - (2 + 4p + (1 - p)(3 - 2q_2 - q_3)) = 1 + 2p + (1 - p)(-1 + 2q_2 + 2q_3) \geq 3p.$$

So u_1^s is always better off by choosing $s_{u_1}^1$.

By symmetry it follows that for each $p > 0$ there is a unique (pure) Bayesian Nash equilibrium σ where $\sigma(u^s) = s_u^1$ and $\sigma(u^m) = s_u^3$ for all players $u \in \mathcal{N}$. For its social cost we get $\text{SC}(\Psi, \sigma) = 2 + 4p$.

So, for each $\epsilon > 0$ there is a sufficiently small p , such that

$$\text{WoM}(\Psi) = \frac{\text{SC}(\Gamma_\Psi, \mathbf{s})}{\text{SC}(\Psi, \sigma)} = \frac{5}{2 + 4p} \geq \frac{5}{2} - \epsilon.$$

This completes the proof of the theorem. □

This is actually a tight result, since for the considered class of malicious Bayesian games the Windfall of Malice cannot be larger than the Price of Anarchy of the corresponding class of congestion games which was shown to be $\frac{5}{2}$ in [5].

5 Conclusion and Open Problems

In this paper, we have introduced and studied a new extension to congestion games, that we call malicious Bayesian congestion games. More specifically, we have studied problems concerned with the complexity of deciding the existence of pure Bayesian Nash equilibria. Furthermore, we have presented results on the Price of Malice.

Although we were able to derive multiple interesting results, this work also gives rise to many interesting open problems. We conclude this paper by stating those, that we consider the most prominent ones.

- Our NP-completeness result in Theorem 1 holds even for linear latency functions, identical type probabilities, and if all strategy sets are singleton sets of resources. However, if such games are further restricted to symmetric games and identical linear latency functions, then deciding the existence of a pure Bayesian Nash equilibrium becomes a trivial task. We believe that this task can also be performed in polynomial time for *non-identical* linear latency functions and symmetric strategy sets.
- Although the upper bound in Corollary 1 and the corresponding lower bound in Theorem 5 are asymptotically tight, there is still potential to improve. We conjecture that in this case $\text{PoB}(\Delta) = \Delta + O(1)$.
- We believe that the concept of malicious Bayesian games is very interesting and deserves further study also in other scenarios. We hope, that our work will encourage others to study such malicious Bayesian games.

6 Acknowledgments

We are very grateful to Christos Papadimitriou and Andreas Maletti for many fruitful discussions on the topic. Moreover, we thank Florian Schoppmann for his helpful comments on an early version of this paper.

References

1. H. Ackermann, H. Röglin, and B. Vöcking. On the Impact of Combinatorial Structure on Congestion Games. In *Proceedings of the 47th Annual Symposium on Foundations of Computer Science (FOCS'06)*, pages 613–622, 2006.
2. S. Aland, D. Dumrauf, M. Gairing, B. Monien, and F. Schoppmann. Exact Price of Anarchy for Polynomial Congestion Games. In *Proceedings of the 23rd International Symposium on Theoretical Aspects of Computer Science (STACS'06)*, LNCS 3884, pages 218–229, 2006.
3. B. Awerbuch, Y. Azar, and A. Epstein. The Price of Routing Unsplittable Flow. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC'05)*, pages 57–66, 2005.
4. M. Babaioff, R. Kleinberg, and C. H. Papadimitriou. Congestion Games with Malicious Players. In *Proceedings of the 8th ACM Conference on Electronic Commerce (EC'07)*, pages 103–112, 2007.
5. G. Christodoulou and E. Koutsoupias. The Price of Anarchy of Finite Congestion Games. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC'05)*, pages 67–73, 2005.
6. V. Conitzer and T. Sandholm. New Complexity Results about Nash Equilibria. In *Proceedings of 18th International Joint Conference on Artificial Intelligence (IJCAI'03)*, pages 765–771, 2003.
7. J. Dunkel and A. S. Schulz. On the Complexity of Pure-Strategy Nash Equilibria in Congestion and Local-Effect Games. In *Proceedings of the 2nd International Workshop on Internet and Network Economics (WINE'06)*, LNCS 4286, pages 62–73, 2006.
8. A. Fabrikant, C. H. Papadimitriou, and K. Talwar. The Complexity of Pure Nash Equilibria. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC'04)*, pages 604–612, 2004.
9. D. Fotakis, S. Kontogiannis, and P. Spirakis. Symmetry in Network Congestion Games: Pure Equilibria and Anarchy Cost. In *Proceedings of the 3rd International Workshop on Approximation and Online Algorithms (WAOA'05)*, LNCS 3879, pages 151–160, 2005.
10. M. Gairing, B. Monien, and K. Tiemann. Selfish Routing with Incomplete Information. *Theory of Computing Systems*, 42(1):91–130, 2008.
11. M. Gairing and F. Schoppmann. Total Latency in Singleton Congestion Games. In *Proceedings of the 3rd International Workshop on Internet and Network Economics (WINE'07)*, LNCS 4858, pages 381–387, 2007.
12. M. X. Goemans, V. Mirrokni, and A. Vetta. Sink Equilibria and Convergence. In *Proceedings of the 46th Annual Symposium on Foundations of Computer Science (FOCS'05)*, pages 142–154, 2005.
13. G. Gottlob, G. Greco, and T. Mancini. Complexity of Pure Equilibria in Bayesian Games. In *Proceedings of 20th International Joint Conference on Artificial Intelligence (IJCAI'07)*, pages 1294–1299, 2007.
14. J. C. Harsanyi. Games with Incomplete Information Played by Bayesian Players, I, II, III. *Management Science*, 14:159–182, 320–332, 468–502, 1967.
15. G. Karakostas and A. Viglas. Equilibria for Networks with Malicious Users. In *Proceedings of the 14th International Symposium on Algorithms and Computation (ISAAC'03)*, LNCS 2906, pages 696–704, 2003.
16. E. Koutsoupias and C. H. Papadimitriou. Worst-Case Equilibria. In *Proceedings of the 16th International Symposium on Theoretical Aspects of Computer Science (STACS'99)*, LNCS 1563, pages 404–413, 1999.
17. L. Libman and A. Orda. Atomic Resource Sharing in Noncooperative Networks. *Telecommunication Systems*, 17(4):385–409, 2001.
18. I. Milchtaich. Congestion Games with Player-Specific Payoff Functions. *Games and Economic Behavior*, 13(1):111–124, 1996.
19. T. Moscibroda, S. Schmid, and R. Wattenhofer. When Selfish Meets Evil: Byzantine Players in a Virus Inoculation Game. In *Proceedings of the 25th Annual ACM Symposium on Principles of Distributed Computing (PODC'06)*, pages 35–44, 2006.
20. J. F. Nash. Non-Cooperative Games. *Annals of Mathematics*, 54(2):286–295, 1951.
21. R. W. Rosenthal. A Class of Games Possessing Pure-Strategy Nash Equilibria. *International Journal of Game Theory*, 2:65–67, 1973.
22. T. Roughgarden and É. Tardos. How Bad Is Selfish Routing? *Journal of the ACM*, 49(2):236–259, 2002.
23. C. A. Tovey. A Simplified NP-complete Satisfiability Problem. *Discrete Applied Mathematics*, 8:85–89, 1984.

A Appendix

A.1 Proof of Theorem 3

Proof. Assume that at least one of the properties holds. We will show that in each case this implies the existence of a pure Bayesian Nash equilibrium.

First assume that (a) holds: For each player u_i assign the selfish type-agent u_i^s alternately to the two resources. At each time, we assign the corresponding malicious type-agent to the other resource. It's not hard to see that the resulting strategy profile is a pure Bayesian Nash equilibrium.

Now assume that (b) holds: In this case assign $\frac{n}{r}$ selfish type-agents and $\frac{n}{r}$ malicious type-agents to each resource such that the selfish and malicious type-agent of each fixed player are not assigned to the same resource. Again it is easy to see that the resulting strategy profile is a pure Bayesian Nash equilibrium.

For the other direction we will show that if neither (a) nor (b) holds then the malicious Bayesian congestion game does not possess a pure Bayesian Nash equilibrium. By way of contradiction assume there exists a malicious Bayesian congestion game Γ satisfying neither (a) nor (b) but Γ admits a pure Bayesian Nash equilibrium σ . We consider 3 sub-cases:

Case 1: $p > \frac{1}{2}$ and $r = 2$

By way of contradiction assume σ assigns more than $\lceil \frac{n}{2} \rceil$ selfish type-agents to some resource e . If this is the case then σ will also assign all malicious type-agents to e . But then all the selfish type-agents on resource e can improve by switching to the other resource, a contradiction to σ being a pure Bayesian Nash equilibrium. It follows that at most $\lceil \frac{n}{2} \rceil$ selfish type-agents are assigned to each resource. This again implies that $\lceil \frac{n}{2} \rceil$ selfish type-agents are assigned to one resource (say e_1) and $\lfloor \frac{n}{2} \rfloor$ selfish type-agents are assigned to the other resource (say e_2). Denote $\mathcal{N}_1 = \{u \in \mathcal{N} \mid \sigma(u^s) = e_1\}$ the set of players with selfish type-agent assigned to resource e_1 and denote $\mathcal{N}_2 = \mathcal{N} \setminus \mathcal{N}_1$. For each player $u \in \mathcal{N}_2$ we have $\delta_{e_1}(\sigma_{-u}) > \delta_{e_2}(\sigma_{-u})$ which implies that $\sigma(u^m) = e_1$ for all $u \in \mathcal{N}_2$. Will now show that $\sigma(u^m) = e_2$ for all $u \in \mathcal{N}_1$. If n is even then this holds immediately by symmetry. So assume n is odd and $\exists u \in \mathcal{N}_1$ with $\sigma(u^m) = e_1$. Now consider an arbitrary player $u' \in \mathcal{N}_1$ with $u' \neq u$. Since n is odd it follows that $n \geq 3$ and thus such a player exists. Furthermore, it follows that $\lfloor \frac{n}{2} \rfloor = \lceil \frac{n}{2} \rceil - 1$. Then

$$\begin{aligned} \delta_{e_1}(\sigma_{-u'}) + \kappa_{e_1}(\sigma_{-u'}) &\geq (1-p) \cdot \left(\lceil \frac{n}{2} \rceil - 1 \right) + p \cdot \left(\lceil \frac{n}{2} \rceil + 1 \right) \\ &= \lceil \frac{n}{2} \rceil - 1 + p \end{aligned}$$

while

$$\begin{aligned} \delta_{e_2}(\sigma_{-u'}) + \kappa_{e_2}(\sigma_{-u'}) &\leq (1-p) \cdot \left\lfloor \frac{n}{2} \right\rfloor + p \cdot \left(\left\lfloor \frac{n}{2} \right\rfloor - 2 \right) \\ &= \left\lfloor \frac{n}{2} \right\rfloor - 1 - p \\ &< \delta_{e_1}(\sigma_{-u'}) + \kappa_{e_1}(\sigma_{-u'}), \end{aligned}$$

a contradiction to σ being a pure Bayesian Nash equilibrium. So, $\sigma(u^m) = e_2$ for all $u \in \mathcal{N}_1$. Summing up, for all $u \in \mathcal{N}_1$ we have $\sigma(u^s) = e_1$ and $\sigma(u^m) = e_2$ while for all $u \in \mathcal{N}_2$ we have $\sigma(u^s) = e_2$ and $\sigma(u^m) = e_1$. In such an assignment each selfish type-agent on resource e_2 can improve by switching to e_1 . This contradicts our initial assumption that σ is a pure Bayesian Nash equilibrium.

Case 2: $p > \frac{1}{2}$ and $r \mid n$

First assume by way of contradiction that there exists a resource $e \in E$ to which σ assigns more than $\frac{n}{r}$ selfish type-agents. It follows that there also exists some other resource e' to which σ

assigns less than $\frac{n}{r}$ selfish type-agents. Since σ is a pure Bayesian Nash equilibrium, no malicious type-agent is assigned to e' . However, then all selfish type-agents on e improve by switching to e' , a contradiction. It follows that σ assigns exactly $\frac{n}{r}$ selfish type-agents to each resource. This also implies that $\sigma(u^m) \neq \sigma(u^s)$ for all players $u \in \mathcal{N}$.

If more than $\frac{n}{r}$ malicious type-agents are assigned to some resource $e \in E$ then all selfish type-agents on e can improve since $p > \frac{1}{2}$. So, σ assigns exactly $\frac{n}{r}$ malicious type-agents to each resource. However, in such a pure strategy profile the selfish type-agent u^s of each player u can improve by switching to $\sigma(u^m)$. This contradicts our initial assumption that σ is a pure Bayesian Nash equilibrium.

Case 3: $r \geq 3$ and $\frac{n}{r} \notin \mathbb{N}$

First observe that if $n < r$ and the malicious type-agents are satisfied then there is always some resource $e \in E$ to which no type-agent is assigned. Each selfish type-agent can then improve by switching to e , a contradiction to σ being a pure Bayesian Nash equilibrium. So we may assume that $n > r$.

Let E^+ be the set of resources to which σ assigns at least $\lceil \frac{n}{r} \rceil$ selfish type-agents. Since $\frac{n}{r} \notin \mathbb{N}$ it follows that $1 \leq |E^+| \leq r - 1$. If $|E^+| \geq 2$ then σ assigns all malicious type-agents to a resource in E^+ . This implies that there exists a selfish type-agent assigned to some resource in E^+ that can improve by switching to some resource in $E \setminus E^+$. It follows that $|E^+| = 1$. Without loss of generality assume $E^+ = \{e_1\}$.

Since $|E^+| = 1$ it follows that σ assigns exactly $\lfloor \frac{n}{r} \rfloor = \lceil \frac{n}{r} \rceil - 1$ selfish type-agents to each resource $e \in E \setminus E^+$. Now, for all $u \in \mathcal{N}$, if $\sigma(u^s) \in E \setminus E^+$ then $\sigma(u^m) = e_1$. It follows that σ assigns at least $(r - 1) \cdot \lfloor \frac{n}{r} \rfloor$ malicious type-agents to e_1 and (by the pigeon hole principle) there exists a resource $e' \in E \setminus E^+$ to which σ assigns at most $\left\lfloor \frac{\lfloor \frac{n}{r} \rfloor}{r-1} \right\rfloor$ malicious type-agents. Since $r \geq 3$ and $n > r$ it follows that all selfish type-agents on resource e_1 can improve by switching to resource e' . This contradicts our initial assumption that σ is a pure Bayesian Nash equilibrium.

In each case we got a contradiction to our assumption that σ is a pure Bayesian Nash equilibrium, proving that Γ does not admit a pure Bayesian Nash equilibrium. This completes the proof of the theorem. \square

A.2 Proof of Theorem 4

Proof. The proof makes use of the following technical lemma.

Lemma 1. *For all $x, y \in \mathbb{R}$ and $c > 0$ we have $x \cdot y \leq c \cdot x^2 + \frac{1}{4c} \cdot y^2$*

Proof of Lemma 1. Observe that

$$\frac{1}{4c} \cdot y^2 - xy + cx^2 = \frac{1}{c} \cdot \left(\frac{y^2}{4} - cxy + c^2x^2 \right) = \frac{1}{c} \cdot \left(\frac{y}{2} - cx \right)^2 \geq 0.$$

The claim follows. \square

We proceed with the proof of Theorem 4. Let Ψ be an arbitrary malicious Bayesian congestion game from $\mathcal{G}(\Delta)$ and let Γ_Ψ be the corresponding (non-malicious) congestion game. Let \mathbf{Q} be an arbitrary Bayesian Nash equilibrium for Ψ . Furthermore, let \mathbf{s}^* be an optimum pure strategy

profile for Γ_{Ψ} . For each player $u \in \mathcal{N}$, we have

$$\begin{aligned}
\text{PC}_u(\mathbf{p}, \mathbf{Q}) &= \sum_{\boldsymbol{\sigma} \in \mathcal{S}^2} q(\boldsymbol{\sigma}) \cdot \text{PC}_u(\mathbf{p}, \boldsymbol{\sigma}) \\
&\leq \text{PC}_u(\mathbf{p}, (\mathbf{Q}_{-u^s}, s_u^*)) \\
&= \sum_{\boldsymbol{\sigma} \in \mathcal{S}^2} q(\boldsymbol{\sigma}) \cdot \text{PC}_u(\mathbf{p}, (\boldsymbol{\sigma}_{-u^s}, s_u^*)) \\
&= \sum_{\boldsymbol{\sigma} \in \mathcal{S}^2} q(\boldsymbol{\sigma}) \sum_{e \in s_u^*} f_e(\delta_e(\boldsymbol{\sigma}_{-u}) + \kappa_e(\boldsymbol{\sigma}_{-u}) + 1) \\
&\leq \sum_{\boldsymbol{\sigma} \in \mathcal{S}^2} q(\boldsymbol{\sigma}) \sum_{e \in s_u^*} f_e(\delta_e(\boldsymbol{\sigma}) + \Delta + 1),
\end{aligned}$$

where the first inequality follows since \mathbf{Q} is a Bayesian Nash equilibrium and the second inequality holds, since $\kappa_e(\boldsymbol{\sigma}_{-u}) \leq \kappa_e(\boldsymbol{\sigma}) \leq \Delta$ for all $e \in E$. So, we get

$$\begin{aligned}
\sum_{u \in \mathcal{N}} (1 - p_u) \cdot \text{PC}_u(\mathbf{p}, \mathbf{Q}) &\leq \sum_{\boldsymbol{\sigma} \in \mathcal{S}^2} q(\boldsymbol{\sigma}) \sum_{u \in \mathcal{N}} \sum_{e \in \sigma^*(u)} (1 - p_u) \cdot f_e(\delta_e(\boldsymbol{\sigma}) + \Delta + 1) \\
&\leq \sum_{\boldsymbol{\sigma} \in \mathcal{S}^2} q(\boldsymbol{\sigma}) \sum_{e \in E} (1 - p_{\min}) \cdot \delta_e(\mathbf{s}^*) \cdot f_e(\delta_e(\boldsymbol{\sigma}) + \Delta + 1) \\
&\leq (1 - p_{\min}) \cdot \sum_{\boldsymbol{\sigma} \in \mathcal{S}^2} q(\boldsymbol{\sigma}) \sum_{e \in E} \delta_e(\mathbf{s}^*) \cdot f_e(\Delta + 1) \\
&\quad + (1 - p_{\min}) \cdot \sum_{\boldsymbol{\sigma} \in \mathcal{S}^2} q(\boldsymbol{\sigma}) \sum_{e \in E} a_e \cdot \delta_e(\boldsymbol{\sigma}) \cdot \delta_e(\mathbf{s}^*)
\end{aligned}$$

Observe, that $\delta_e(\mathbf{s}^*) \geq 1$ and thus $\delta_e(\boldsymbol{\sigma}) \leq \delta_e(\mathbf{s}^*)^2$. Moreover, by applying Lemma 1 with $x = \delta_e(\boldsymbol{\sigma})$ and $y = \delta_e(\mathbf{s}^*)$, we get

$$\begin{aligned}
&\sum_{u \in \mathcal{N}} (1 - p_u) \cdot \text{PC}_u(\mathbf{p}, \mathbf{Q}) \\
&\leq (1 - p_{\min}) \cdot (\Delta + 1) \cdot \sum_{\boldsymbol{\sigma} \in \mathcal{S}^2} q(\boldsymbol{\sigma}) \sum_{e \in E} \delta_e(\mathbf{s}^*) \cdot f_e(\delta_e(\mathbf{s}^*)) \\
&\quad + (1 - p_{\min}) \cdot c \cdot \sum_{\boldsymbol{\sigma} \in \mathcal{S}^2} q(\boldsymbol{\sigma}) \sum_{e \in E} a_e \cdot \delta_e(\boldsymbol{\sigma})^2 \\
&\quad + (1 - p_{\min}) \cdot \frac{1}{4c} \cdot \sum_{\boldsymbol{\sigma} \in \mathcal{S}^2} q(\boldsymbol{\sigma}) \sum_{e \in E} a_e \cdot \delta_e(\mathbf{s}^*)^2 \\
&\leq (1 - p_{\min}) (\Delta + 1 + \frac{1}{4c}) \sum_{\boldsymbol{\sigma} \in \mathcal{S}^2} q(\boldsymbol{\sigma}) \sum_{e \in E} \delta_e(\mathbf{s}^*) \cdot f_e(\delta_e(\mathbf{s}^*)) \\
&\quad + c \cdot \sum_{\boldsymbol{\sigma} \in \mathcal{S}^2} q(\boldsymbol{\sigma}) \sum_{e \in E} \delta_e(\boldsymbol{\sigma}) f_e(\delta_e(\boldsymbol{\sigma})) \\
&= (1 - p_{\min}) (\Delta + 1 + \frac{1}{4c}) \cdot \sum_{u \in \mathcal{N}} \sum_{e \in s_u^*} f_e(\delta_e(\mathbf{s}^*)) \\
&\quad + c \cdot \sum_{\boldsymbol{\sigma} \in \mathcal{S}^2} q(\boldsymbol{\sigma}) \sum_{u \in \mathcal{N}} \sum_{e \in \sigma(u^s)} (1 - p_u) \cdot f_e(\delta_e(\boldsymbol{\sigma})) \\
&\leq (1 - p_{\min}) (\Delta + 1 + \frac{1}{4c}) \cdot \sum_{u \in \mathcal{N}} \sum_{e \in s_u^*} f_e(\delta_e(\mathbf{s}^*))
\end{aligned}$$

$$\begin{aligned}
& + c \cdot \sum_{\boldsymbol{\sigma} \in \mathcal{S}^2} q(\boldsymbol{\sigma}) \sum_{u \in \mathcal{N}} \sum_{e \in \sigma(u^s)} (1 - p_u) \cdot f_e(\delta_e(\boldsymbol{\sigma}_{-u}) + 1) \\
& \leq (1 - p_{\min})(\Delta + 1 + \frac{1}{4c}) \cdot \sum_{u \in \mathcal{N}} \text{PC}_u(\mathbf{s}^*) \\
& + c \cdot \sum_{u \in \mathcal{N}} (1 - p_u) \cdot \text{PC}_u(\mathbf{p}, \mathbf{Q})
\end{aligned}$$

It follows that

$$\begin{aligned}
\frac{\text{SC}(\mathbf{p}, \mathbf{Q})}{\text{SC}(\mathbf{s}^*)} &= \frac{n}{n - \Delta} \cdot \frac{\sum_{u \in \mathcal{N}} (1 - p_u) \cdot \text{PC}_u(\mathbf{Q})}{\sum_{u \in \mathcal{N}} \text{PC}_u(\mathbf{s}^*)} \\
&\leq \frac{n}{n - \Delta} (1 - p_{\min}) \frac{\Delta + 1 + \frac{1}{4c}}{1 - c}
\end{aligned} \tag{1}$$

Now, choosing $c = \frac{-1 + \sqrt{5 + 4\Delta}}{4(\Delta + 1)}$ yields

$$\begin{aligned}
\frac{\Delta + 1 + \frac{1}{4c}}{1 - c} &= \frac{\Delta + 1 + \frac{\Delta + 1}{-1 + \sqrt{5 + 4\Delta}}}{1 - \frac{-1 + \sqrt{5 + 4\Delta}}{4(\Delta + 1)}} \\
&= \frac{4(\Delta + 1)^2 (1 + \frac{1}{-1 + \sqrt{5 + 4\Delta}})}{5 + 4\Delta - \sqrt{5 + 4\Delta}} \\
&= \frac{4(\Delta + 1)^2 (1 + \frac{1 + \sqrt{5 + 4\Delta}}{4 + 4\Delta})}{5 + 4\Delta - \sqrt{5 + 4\Delta}} \\
&= \frac{(\Delta + 1)(5 + 4\Delta + \sqrt{5 + 4\Delta})}{5 + 4\Delta - \sqrt{5 + 4\Delta}} \\
&= \frac{(\Delta + 1)(\sqrt{5 + 4\Delta} + 1)}{\sqrt{5 + 4\Delta} - 1} \\
&= \frac{(\Delta + 1)(\sqrt{5 + 4\Delta} + 1)^2}{4 + 4\Delta} \\
&= \frac{1}{4}(4\Delta + 6 + 2\sqrt{5 + 4\Delta}) \\
&= \Delta + \frac{3 + \sqrt{5 + 4\Delta}}{2}
\end{aligned} \tag{2}$$

The theorem follows by combining (1) and (2) and since \mathbf{Q} is an arbitrary Bayesian Nash equilibrium. \square