On the Price of Mediation

Milan Bradonjić
Los Alamos National Laboratory
Theoretical Division, and Center for Nonlinear Studies
milan@lanl.gov

Gunes Ercal-Ozkaya
Kansas University
Department of Electrical Engineering and Computer Science
gunes@ku.edu

Adam Meyerson
University of California, Los Angeles
Computer Science Department
awm@cs.ucla.edu

Alan Roytman
University of California, Los Angeles
Computer Science Department
alanr@cs.ucla.edu

ABSTRACT

We study the relationship between the social cost of correlated equilibria and the social cost of Nash equilibria. In contrast to previous work focusing on the possible benefits of a benevolent mediator, we define and bound the Price of Mediation (PoM): the ratio of the cost of the worst correlated equilibrium to the cost of the worst Nash. We observe that in practice, the heuristics used for mediation are frequently non-optimal, and from an economic perspective mediators may be inept or self-interested. Recent results on computation of equilibria also motivate our work.

We consider the Price of Mediation for general games with small numbers of players and pure strategies. For games with two players each having two pure strategies we prove a tight bound of two on the PoM. For larger games (either more players, or more pure strategies per player, or both) we show that the PoM can be unbounded.

Most of our results focus on symmetric congestion games (also known as load balancing games). We show that for general convex cost functions, the PoM can grow exponentially in the number of players. We prove that PoM is one for linear costs and at most a small constant (but can be larger than one) for concave costs. For polynomial cost functions, we prove bounds on the PoM which are exponential in the degree.

1. INTRODUCTION

We consider games played by independent self-interested agents, and study the effect of adding a *mediator* who suggests strategies without any power to enforce them. This yields a correlated equilibrium [2, 3] instead of a Nash [10].

Correlated equilibria are interesting for a number of reasons. In many cases they are tractable to compute [11, 12] whereas computing Nash equilibria is PPAD-complete even for 2-player matrix games [4, 6]. The multi-agent machine learning community has observed that for *no-regret* algorithms, play converges to a correlated equilibrium [7], whereas only under additional restrictions can convergence to a Nash equilibrium be guaranteed [8].

This paper measures the potential social cost imposed by inept or malicious mediation. We observe that in many cases it is non-trivial to produce an optimum mediatior, requiring global information or substantial computing time. In many cases system designers impose heuristic mediation which may not be optimum. There are also many social science examples of inept or even self-interested mediators.

We define a measure called the **Price of Mediation** (PoM) which is the ratio of the social cost of the worst correlated equilibrium solution to the social cost of the worst Nash equilibrium solution. We show that even for two-player matrix games with only two strategies per player, the PoM can be as large as a factor of two. For larger matrix games, we show that the PoM can be arbitrarily large. Thus the costs imposed by inept or malicious mediation can in fact be quite severe.

The main results of our paper bound the PoM for symmetric congestion games (also known as load balancing games). These games have been well-studied in previous work from the standpoint of Price of Anarchy [9], and are a natural place to start in measuring the PoM. We show that for symmetric congestion games with linear cost functions the PoM is one. For games with concave cost functions, we give an example where the PoM is greater than one, but show that it can be no larger than four. For games with convex costs, we show that in the case of arbitrary or exponential cost functions the PoM can be arbitrarily bad (although it is bounded above by $O(m^{n-1})$ where n is the number of players and m is the number of strategies). For polynomial costs, we show that the PoM can grow exponentially in the degree, and cannot be substantially worse than this.

1.1 Related Work

The paper of Ashlagi et al [1] measures the social benefit of optimum mediation in some simple classes of games. The paper considers games with positive utilities, whereas our results focus on games with positive costs (negative utilities). Major results in [1] include proving that mediation helps by at most a factor of 4/3 in a matrix game with two players and two strategies, and proving that mediation can help by an unbounded factor even in congestion games with few players and strategies. We show that in contrast, the Price of Mediation in a two player and two strategy game is tightly bounded by 2. Similar examples to those in [1] exist to show that the Price of Mediation is unbounded in

EC'09, July 6–10, 2009, Stanford, California, USA.

small congestion games with positive utilities. Under the more standard model of positive costs, we are able to bound the PoM in terms of the number of strategies and players, or alternatively based upon the type of cost function. It may be interesting to revisit the results of [1] from the standpoint of positive costs.

A recent result of Roughgarden [14] compares the worst correlated equilibrium to the social optimum. He shows that for some classes of games (including congestion games), this Extended Price of Anarchy is bounded identically to the Price of Anarchy (ratio of worst Nash to social optimum). This does not tightly bound the PoM, which could be as large as the extended price of anarchy or as small as one. Our results complement [14] by bounding the ratio of the worst correlated and Nash for *particular* games rather than taking the extreme points for each over an entire class of games.

2. DEFINITIONS

A game is defined by a set of players i, a strategy set S_i for each player, and a cost (or utility) function c_i for each player. The *outcome* of a game is a selection of strategy for each player (thus a member of the product space of strategy sets) and the cost function maps from each outcome to a real number. In this paper we will deal with games having non-negative costs (or equivalently non-positive utilities).

Consider a probability distribution over outcomes where $p(\vec{s})$ is the probability of the outcome $\vec{s}$. For player i, the expected cost when playing strategy s_i^* is given by:

$$\Sigma_{\vec{s}:s_i=s_i^*} c_i(\vec{s})p(\vec{s})$$

If player i were to defect from strategy s_i^* to s_i', then if $\vec{s'}$ represents $\vec{s}$ with coordinate s_i modified to be s_i', then the expected cost would be:

$$\Sigma_{\vec{s}:s_i=s_i^*} c_i(\vec{s'})p(\vec{s})$$

A distribution is a **correlated equilibrium** if for any player i and any pair s_i^*, s_i' of strategies for player i, defecting from s_i^* to s_i' does not decrease the expected cost. We imagine a mediator selecting the outcome of the game; correlated equilibria guarantee that obeying the mediator will minimize the expected cost for each player, and thus that it is in the best interest of the players to follow the outcome selected through mediation.

A **Nash equilibrium** can be defined in a similar way, with the additional constraint that the probability distribution over outcomes must correspond to independent random selections by the various players. Thus we have a probability $p_i(s_i)$ for each player i and each $s_i \in S_i$ and require that $p(s_1, s_2, ..., s_n) = p_1(s_1)p_2(s_2)...p_n(s_n)$ for all outcomes.

The social cost of a distribution is $C(p) = \Sigma_{\vec{s}} p(\vec{s})\Sigma_i c_i(\vec{s})$. We will consider the **Price of Mediation** (PoM) which is the ratio of the maximum $C(p)$ where p is a correlated equilibrium divided by the maximum $C(p)$ where p is a Nash equilibrium. The PoM is a property of the particular game, and since every Nash is also a correlated equilibrium we will always have PoM at least one. We will sometimes refer to the PoM for a *class* of games, which is the maximum value of the PoM for any game in that class.

3. SMALL GAMES

We begin by showing that even for games of only two players and two strategies, the PoM can be as large as two. For larger games, we show that the PoM can be unbounded.

Note that adding a scalar to all costs will not change the set of Nash or Correlated Equilibrium solutions, which implies that since it is possible to design a game where the PoM is not equal to one, then by subtracting a suitable scalar from all costs we can obtain a situation where the worst Nash has negative cost and the worst Correlated has positive cost. For this reason we consider exclusively games with non-negative costs in this paper (it is also interesting to consider games with non-negative benefits or payoffs, but allowing either sign to be present in a single game matrix immediately yields infinite Price of Mediation).

3.1 Two-by-Two Games

Consider games with two players, each of whom has two possible strategies. We will assume that every outcome of the game assigns a non-negative cost to each player. Perhaps surprisingly, even in this simple case, it is possible for the Price of Mediation to asymptotically approach two.

LEMMA 3.1. *For any $\delta > 0$, there exists a two-by-two game Γ with non-negative costs such that $PoM(\Gamma) > 2 - \delta$.*

PROOF. Consider the following game matrix Γ_ϵ:

$$\begin{pmatrix} 0,1 & \frac{1}{\epsilon^2}, (2-2\epsilon) \\ (1-2\epsilon), \epsilon & 0,0 \end{pmatrix}$$

There are three Nash equilibria for this game (two pure and one mixed); the worst social cost belongs to the pure Nash which simply plays strategy one for each player and has cost 1.

The worst correlated equilibrium for this game is given by V_C:

$$V_C = \frac{1}{1+\epsilon-3\epsilon^2+2\epsilon^3}\begin{pmatrix} 1 & \epsilon^2(1-2\epsilon) \\ 0 & \epsilon(1-2\epsilon)^2 \end{pmatrix} \quad (1)$$

It is straightforward to verify that this is a correlated equilibrium by considering the various cases and computing that defecting from the strategy suggested by the mediator will never decrease the cost to the players. We observe that $1+\epsilon-3\epsilon^2+2\epsilon^3 \le 1+\epsilon$.

$$|c(V_C)| \ge \frac{1}{1+\epsilon}(1+\epsilon^2(1-2\epsilon)(1/\epsilon^2)) = 2(\frac{1-\epsilon}{1+\epsilon})$$

It's clear that by setting ϵ to be sufficiently small we can get this to be arbitrarily close to two, completing the proof. □

In fact this bound is tight; there is no such game where the Price of Mediation actually reaches two. The proof of the following lemma involves a great deal of algebra and case analysis, and we will defer it to Appendix A.

LEMMA 3.2. *For any two-by-two game with non-negative costs Γ, $PoM(\Gamma) < 2$.*

3.2 Two Player, Three Strategy Games

We now prove that a mediator can be arbitrarily bad for three-by-three games with two players. Consider the following cost matrix:

$$\begin{pmatrix} 10,1 & 0,0 & 0,1 \\ 1,2 & 5,3 & 1,4 \\ 0,5 & 2,2 & 10,0 \end{pmatrix}$$

The only and therefore worst Nash equilibrium of this game belongs to player one selecting strategy one and player two selecting strategy two, for a total cost of 0. The following is a correlated equilibrium for this game, with a social cost of $\frac{18}{5}$:

$$\begin{pmatrix} 0 & \frac{1}{15} & 0 \\ \frac{3}{5} & 0 & \frac{1}{15} \\ \frac{1}{5} & \frac{1}{30} & \frac{1}{30} \end{pmatrix}$$

We will defer the reader to Appendix B for proof that this is in fact a correlated equilibrium and the worst Nash equilibrium has a social cost of 0. We can then increase the values in the cost matrix by some small value ϵ such that the cost of the correlated equilibrium given is only slightly perturbed while the cost of the worst Nash will look like 2ϵ. Thus, the Price of Mediation for this game is at least $\frac{9}{5\epsilon}$, which can be made arbitrarily bad as ϵ approaches 0. Moreover, this proves that the Price of Mediation of games with two players who both have at least three strategies is unbounded. This follows since we can "embed" this cost matrix into any matrix which has at least three rows and at least three columns, filling the rest of the entries with sufficiently large values M. We get the following theorem:

THEOREM 3.3. *The Price of Mediation for the class of two player games where each player has at least three strategies is unbounded.*

3.3 Three Player, Two Strategy Games

We now consider three-player games, proving the following theorem:

THEOREM 3.4. *The price of mediation for the class of three player, two strategy games is unbounded.*

PROOF. Consider the following two-by-two game Γ:

$$\begin{pmatrix} 1+\frac{1}{\epsilon}, 1+\frac{1}{\epsilon} & 0,1 \\ 1,0 & \frac{1}{1-\epsilon}, \frac{1}{1-\epsilon} \end{pmatrix}$$

There are only three Nash equilibria for this game. These include two pure equilibria where one player chooses the first strategy and the other chooses the second strategy, and a mixed Nash where each player chooses the first strategy independently with probability ϵ. Note that in all cases the chance of the outcome where both choose the first strategy is at most ϵ^2. The social cost of Nash (the mixed) will be $2(1+\epsilon)$.

There exists the following correlated equilibrium. Both players select strategy one with probability $\frac{\epsilon}{2-\epsilon}$. With probability $\frac{1-\epsilon}{2-\epsilon}$ player one selects strategy one and player two selects strategy two. With probability $\frac{1-\epsilon}{2-\epsilon}$ player one selects strategy two and player two selects strategy one.

If player one selects strategy two, then he knows his opponent plays strategy one. He is getting a cost of 1 but if he defects he gets $1+\frac{1}{\epsilon}$ which is worse. If player one selects strategy one, then his expected payoff is the same whether he defects or not. So it's a correlated equilibrium.

Of course, the cost of this correlated equilibrium looks like $\frac{2-2\epsilon}{2-\epsilon} + \frac{\epsilon}{2-\epsilon}(2+\frac{2}{\epsilon}) = \frac{4}{2-\epsilon}$ which is very close to the cost of the worst Nash.

Now consider adding a third player. This player does not really have a choice of strategy: he has a cost of M whenever *both* the first two players select strategy one, and a cost of 0 otherwise. Let us assume that $M\epsilon^2 >> 2$. Now for any Nash, the probability that both the first two players select strategy one is at most ϵ^2 so the total cost will look like $M\epsilon^2 + 2(1+\epsilon)$ at most. But for the correlated equilibrium, we get a cost of $\frac{4}{2-\epsilon} + \frac{M\epsilon}{2-\epsilon}$. This is a gap of $\frac{2}{\epsilon}$, which we can make arbitrarily bad by decreasing ϵ. $\square$

4. SYMMETRIC CONGESTION GAMES

We have established that there exist games (even with only two players) where the Price of Mediation is arbitrarily bad. We now constrain ourselves to congestion games on parallel links (also called symmetric congestion games or load balancing games). Each of n players selects one of m strategies, then pays a *cost* which is an increasing function of the number of players selecting the same strategy. These games have been well-studied from the viewpoint of Price of Anarchy [9], and are a version of the classic model of [15] where the players have finite size.

4.1 Small Symmetric Congestion Games

We will first bound the Price of Mediation in terms of the number of players and strategies. This result can be contrasted with our theorem 3.4 for more general games.

THEOREM 4.1. *For any symmetric congestion game with n players and m strategies per player, the PoM is at most m^{n-1}. There is a game with PoM at least $\Omega(\frac{1}{n}m^{n-2})$.*

PROOF. For a game with m strategies, let the cost function of the ith strategy be $f_i(x)$. Let the strategies be such that $f_1(n) \leq f_2(n) \leq \cdots \leq f_m(n)$. First, the correlated equilibrium has cost $\leq nf_1(n)$, otherwise players would defect to strategy one regardless of the mediator's instructions. The game is symmetric with respect to players so there exists a symmetric Nash equilibrium, where the players have the same probability distribution [5]. Let p_i be the probability of playing the strategy i in the Nash equilibrium. Then the cost of the Nash is at least $\Sigma_i p_i^n n f_i(n)$ (this essentially assumes that $f_i(x) = 0$ for $x < n$; otherwise the cost of the Nash is only greater). Using the fact that $f_1(n) \leq f_i(n)$, the price of mediation must satisfy:

$$PoM \leq \frac{1}{\sum_{i=1}^{m} p_i^n} \leq \frac{1}{m(1/m)^n} = m^{n-1}$$

which is obtained by maximizing the sum of p_i^n subject to a sum of one.

For the other direction, we consider a congestion game where the costs are $f(n) = 1$ and otherwise zero. Since this is convex, the worst Nash will be uniformly at random with expected cost $\frac{n}{m^{n-1}}$. On the other hand, we construct a correlated equilibrium which assigns all players to the same randomly chosen strategy with probability $\frac{1}{mn+1}$ and assigns one randomly chosen player to one strategy and the rest to another with probability $\frac{mn}{mn+1}$. In this equilibrium, each player pays a cost of $\frac{1}{mn+1}$. However, defecting leads to paying a cost of 1 with strictly larger probability showing that this is a correlated equilibrium. □

These results are not quite tight, as there is a gap of nm between upper and lower bound. For the special case of *totally symmetric congestion games* in which all strategies have identical cost functions we have obtained tight bounds for two and three players; these bounds are deferred to appendix C. Nonetheless, we observe that these bounds increase very quickly with n (the number of players) and that interesting instances of congestion games often involve a large number of players (for example drivers on a highway or packets in a network). We will attempt to prove better bounds on the Price of Mediation by restricting the cost functions in our congestion game.

4.2 Linear Symmetric Congestion Games

In this section, we consider symmetric congestion games where all cost functions are linear. We prove that mediation cannot hurt us in this case, regardless of the number of players and strategies.

THEOREM 4.2. *The PoM for linear symmetric congestion games is* 1.

PROOF. Let $f_j(x) = a_j x + b_j$ be the cost function for strategy j. Then there are $\hat{x}_j$ and μ such that $\mu = f(\hat{x}_j + 1)$ for all j, and $\sum \hat{x}_j = n-1$. (Given the linear functions f_j, if $a_j > 0$ then these $\hat{x}_j$ and μ are unique). Let each player be assigned to a strategy j with probability $\hat{x}_j/(n-1)$. Then, every player in expectation pays $E[f(\lambda_j + 1)]$ where λ_j is the number of other players selecting strategy j. Since the function f is linear, this is equal to $f(E[\lambda_j]+1)$ by linearity of expectations, and this is equal to μ. This solution (assigning each player to a strategy j with probability $\hat{x}_j/(n-1)$) is a Nash Equilibrium, since every strategy has the same expected cost for any player i (μ). The total expected cost is μn.

We consider correlated equilibrium. Suppose player i is told to play strategy j', and let δ_j be the number of players other than player i told to play strategy j. Clearly $\Sigma_j \delta_j = n-1$ so by linearity of expectations it follows that $\sum_j E[\delta_j] = n-1$. Since $\sum_j \hat{x}_j = \sum_j E[\delta_j]$, there is j such that $E[\delta_j] \leq \hat{x}_j$. Then by the linearity of the functions it follows:

$$E[f_j(\delta_j + 1)] = f_j(E[\delta_j] + 1) \leq f_j(\hat{x}_j + 1) = \mu.$$

Of course, player i pays $E[f_{j'}(\delta_{j'} + 1)]$, but since this is supposed to be a correlated equilibrium it cannot help player i to defect (disobey the mediator) and select instead strategy j. Thus it must be that $E[f_{j'}(\delta_{j'} + 1)] \leq E[f_j(\delta_j + 1)] \leq \mu$. We conclude that the cost of the correlated equilibrium is at most $n\mu$. From the previous conclusion that the cost of a Nash equilibrium is $n\mu$, and the fact that $PoM \geq 1$, now it follows that $PoM = 1$. □

4.3 Concave Symmetric Congestion Games

In the following, we demonstrate that although a mediator can strictly hurt the social welfare when parallel links have concave rather than linear latency functions, the damage is still bounded by a small constant, 4.

THEOREM 4.3. *There exists a concave congestion game where the PoM is strictly greater than* 1.

PROOF. We consider the function f defined by $f(x) = x$ for $x \leq 3$ and $f(x) = 3$ for $x > 3$. This is clearly a concave function. We will consider a game with m strategies and $n = m + 1$ players, where m is a large odd number.

One correlated equilibrium randomly selects $\frac{m+1}{2}$ strategies and then assigns two players to each of those strategies. This ensures that all players pay a cost of $f(2) = 2$, for a total social cost of $2n$. This is a correlated equilibrium, since a defecting player can do no better than to select a random strategy other than his own, and will pay $f(3) = 3$ with probability $\frac{1}{2}$ and $f(1) = 1$ with probability $\frac{1}{2}$, giving no incentive to defect.

Now consider any Nash equilibrium for this game. For any player i and strategy j, let N_{ij} be the number of players other than i who select strategy j. If player i has an expected cost of more than two, then there must exist a strategy with $E[N_{ij}] \leq 1$ (by pigeonhole principle) and defecting to this strategy will give player i an expected cost of at most two (thus not a Nash). Since every player pays at most 2, in order for the Price of Mediation to be one it must be that all players pay exactly two. From this it follows that for every i and j we must have $E[N_{ij}] = 1$ and $E[f(N_{ij}+1)] = 2$. From these equations we can conclude that $Pr[N_{ij} \geq 3] = 0$, since otherwise the expectation of $f(N_{ij}+1)$ will be too small. So for every strategy j, there are exactly *two* players who select this strategy with non-zero probability (since assignments are independent in Nash, if there were more than two players then for some j we would have $Pr[N_{ij} \geq 3] > 0$). However, for each of these players that fact that $E[N_{ij}] = 1$ implies that the other player selects this strategy with probability one. This implies that for i other than these two players, $E[N_{ij}] = 2$ which is a contradiction. □

THEOREM 4.4. *The PoM of symmetric congestion games with concave latency functions on parallel links is at most* 4.

PROOF. Let $f_j(x)$ be the concave cost function for strategy j. Then there are unique $\hat{x}_j$ and μ such that $\mu = f_j(\hat{x}_j + 1)$ for all j, and $\sum \hat{x}_j = n - 1$. Consider any correlated equilibrium. If player i is told to play strategy j', then let δ_j be the number of players other than i told to play j. It must be the case that $E[f_{j'}(\delta_{j'} + 1)] \leq E[(f_j(\delta_j + 1)]$ for all j, since otherwise player i would defect to strategy j. Since the sum of δ_j is equal to $n-1$, there is some j where $E[\delta_j] \leq \hat{x_j}$ by linearity of expectations. For this j, the expected cost is $E[f_j(\delta_j+1)] \leq f_j(E[\delta_j]+1) \leq f_j(\hat{x_j}+1) = \mu$ by concavity of f. So the expected payment of player i told to play j' is at most μ. Thus, in any correlated equilibrium, the cost to each player must be at most μ.

Now, consider the Nash equilibrium obtained by the following sequential process: number the players 1 through n. Player 1 at time 1 chooses the strategy j that has minimum $f_j(1)$. Let y_{ij} denote the number of players in strategy j right before time i. For time $i > 1$, player i at time i chooses the strategy that minimizes $f_j(y_{ij} + 1)$. Let y_j denote the final strategy profile y_{nj}. This is a pure strategy profile. Consider any pair of strategies j, j' with $y_j > 0$. Let player i be the last player assigned to strategy j. Then $y_j = y_{ij} + 1$ and $y_{j'} \geq y_{ij'}$ since the number of players assigned to j' can only increase as players subsequent to i arrive. Since we assigned i to j, we must have had $f_j(y_{ij} + 1) \leq f_{j'}(y_{ij'} + 1)$ from which we conclude that $f_j(y_j) \leq f_{j'}(y_{j'} + 1)$ since $f_{j'}$

is non-decreasing. Thus no player will want to defect from j to j', implying that this is a Nash equilibrium.

Since $\sum_j y_j = n > n-1 = \sum_j \hat{x}_j$, there exists j such that $y_j > \hat{x}_j$. Since a Nash permits no defections:

$$\forall j' : f_j(y_j) < f_{j'}(y_{j'}+1)$$

If $y_{j'} \geq 1$ (the only relevant case) then $y_{j'}+1 \leq 2y_{j'}$ and $f_j(y_j) < 2f_{j'}(y_{j'})$ by concavity, implying that

$$f_{j'}(y_{j'}) > \frac{1}{2} f_j(y_j)$$

.

So every player must pay at least $\frac{1}{2}f_j(y_j)$. We observe that $2y_j \geq \hat{x}_j + 1$; this is clear because $y_j > \hat{x}_j \geq 0$ and because y_j is integer we must have $y_j \geq 1$. Thus $f_j(y_j) \geq f_j(\frac{1}{2}(\hat{x}_j+1)) \geq \frac{1}{2}f_j(\hat{x}_j+1) = \frac{1}{2}\mu$.

So for all $y_{j'} \geq 1$ we have $f_{j'}(y_{j'}) \geq \frac{1}{4}\mu$ and every player pays at least $\frac{1}{4}\mu$ in this Nash. Thus the Price of Mediation is at most the ratio of the worst correlated equilibrium's social cost (at most μ) to the worst Nash (at least $\frac{1}{4}\mu$), completing the proof. □

We can prove that for totally symmetric games (all cost functions identical) the PoM is bounded by 2 instead of 4; this proof is in appendix D.

4.4 Convex Symmetric Congestion Games

For arbitrary convex functions, we can have Price of Mediation as large as in theorem 4.1 (in fact the functions in that theorem are convex). For polynomial functions, we will show that the Price of Mediation is bounded in terms of the degree of the polynomial, but can be exponential in that degree.

THEOREM 4.5. *For a symmetric congestion game where all cost functions are polynomials with non-negative coefficients and degree at most p, the PoM is at most $O((2p)^p)$.*

PROOF. We will make repeated use of the following property of polynomials with non-negative coefficients and degree at most p. Consider any $x \geq 0$ and any $\alpha \geq 1$. We must have the property that $f(\alpha x) \leq \alpha^p f(x)$. This is easy to see by writing out a polynomial and substituting the appropriate values.

Suppose our congestion game has n players and m strategies. Let the cost function for strategy i be f_i. There exists a pure Nash equilibrium for this game. We can construct one by starting with no players assigned to strategies, then repeatedly assigning players to the strategy with minimum $f_i(x_i+1)$ where x_i is the current number of players assigned that strategy, much as in theorem 4.4. At the end of this process, let x_i be the number of players assigned to strategy i. It is clear that $\Sigma_i x_i = n$. Let $\mu = \max_i f_i(x_i)$. We will claim that every player in this Nash equilibrium pays at least $\mu 2^{-p}$. If not, then there exists some i such that $x_i > 0$ and $f_i(x_i) < \mu 2^{-p}$. Now using the key property of the polynomial functions, we have:

$$f_i(x_i+1) \leq f_i(2x_i) \leq 2^p f_i(x_i) < \mu$$

Since there is some player who pays μ, this player can defect to strategy i and reduce his cost, violating the definition of a Nash equilibrium. Note that the pure Nash thus defined is clearly not the "worst Nash" for this game; however bounding the maximum ratio of the cost of a correlated equilibrium to the cost of this Nash will be sufficient to bound the Price of Mediation.

Now we will consider correlated equilibria. We can assume that the correlated equilibrium is symmetric with respect to players (otherwise we could shuffle the player identities randomly to obtain another correlated equilibrium of equal social cost). The correlated equilibrium will select one of many possible outcomes for the game; each outcome is defined by an assignment of some number of players to each strategy (we can assume that *which* players these actually are is determined uniformly at random). For any possible outcome j, let ρ^j be the probability of selecting that outcome and n_i^j be the number of players assigned to strategy i assuming outcome j. Then the expected cost to any player in this correlated equilibrium is C, where:

$$C = \Sigma_j \rho^j (\Sigma_i \frac{n_i^j}{n} f_i(n_i^j))$$

One possible manner in which players could defect is to assign themselves randomly to strategy i with probability x_i/n (where x_i is from the pure Nash defined earlier). Suppose that a single player defects in this manner. We can bound his new expected cost C' as follows:

$$C' \leq \Sigma_j \rho^j (\Sigma_i \frac{x_i}{n} f_i(n_i^j+1))$$

We now observe that if $n_i^j < px_i$ then since everything is integer we have $n_i^j + 1 \leq px_i$ and we can apply the property of polynomial functions to get:

$$f_i(n_i^j+1) \leq f_i(px_i) \leq p^p f_i(x_i) \leq p^p \mu$$

On the other hand, if $n_i^j \geq px_i$ then if $x_i > 0$ we have:

$$f_i(n_i^j+1) \leq (\frac{n_i^j+1}{n_i^j})^p f_i(n_i^j) \leq e f_i(n_i^j)$$

Summing these two equations gives us:

$$C' \leq \Sigma_j \rho^j \Sigma_i (\frac{x_i}{n} p^p \mu + \frac{e}{np} n_i^j f_i(n_i^j))$$

$$C' \leq p^p \mu + \frac{e}{p} C$$

Of course, for any correlated equilibrium defecting should not help the expected cost, so we must have that $C \leq C'$. Solving the resulting inequality gives us, for $p \geq 3$:

$$C \leq \mu p^{p+1} \frac{1}{p-e}$$

Combining this with the bound on the cost of the pure Nash gives the desired result. Note that for $1 \leq p \leq 3$ this is just a large constant, so the claim follows immediately. □

THEOREM 4.6. *There is a congestion game where costs are polynomials of degree at most p, such that the PoM is exponential in p.*

PROOF. We consider $f(x) = x^p$, for some $p > 0$. We will compute the cost of the worst Nash equilibrium. We will show that the Price of Mediation can be exponential in p. Under the binomial distribution $Bin(n, 1/m)$ we have that the expected cost of a player is

$$\mathrm{E}[f(X)] = \sum_{k=0}^{n} k^p \binom{n}{k} (1/m)^k (1-1/m)^{n-k}. \quad (2)$$

Let us call the expected cost $S_{n,m,p} := \mathrm{E}[f(X)]$. The recurrence equation on $S_{n,m,p}$ can be derived, for $p \geq 1$, to be

$$S_{n,m,p} = \frac{n}{m} \sum_{i=0}^{p-1} \binom{p-1}{i} S_{n-1,m,i},$$

while $S_{n,m,0} = 1$, for $p = 0$.

In general, the p-th moment of the binomial distribution, is a form of a hypergeometric sum, and cannot be represented in a closed form. Instead we will pursue an approximation.

Let us approximate binomial distribution $Bin(n, 1/m)$ by the normal distribution $N(\mu, \sigma^2)$, with expectation $\mu = n/m$ and variance $\sigma^2 = (n/m)(1-1/m)$, and integrate from 0 to n; actually we will integrate up to ∞.

$$\begin{aligned} \mathrm{E}[X^p] &\approx \int_0^\infty x^p \frac{1}{\sigma\sqrt{2\pi}} e^{-(x-\mu)^2/(2\sigma^2)} dx \\ &= \frac{1}{\sigma\sqrt{2\pi}} \int_0^\infty e^{m\phi(x)} dx \end{aligned}$$

where

$$\phi(x) = \frac{p}{m} \ln x - (x-\mu)^2/(2m\sigma^2). \quad (3)$$

Let us discuss the error made in the tail

$$\frac{1}{\sigma\sqrt{2\pi}} \int_n^\infty x^p e^{-(x-\mu)^2/(2\sigma^2)} dx. \quad (4)$$

The last integral can easily be calculated by recursion. Then, using the Stirling approximation for $p!$ we get that the error term, i.e., Eq. (4) tends to 0, as n tends to ∞.

We consider the case of n, m tending to ∞, so we can assume $\sigma^2 \approx \frac{n}{m}$, and consider ϕ to be

$$\phi(x) = \frac{p}{m} \ln x - \frac{1}{2n}(x - \frac{n}{m})^2. \quad (5)$$

The first and second derivatives of $\phi(x)$ are given by

$$\phi(x)' = \frac{p}{m}\frac{1}{x} - \frac{1}{n}(x - \frac{n}{m}) \quad (6)$$

$$\phi(x)'' = -\frac{p}{m}\frac{1}{x^2} - \frac{1}{n}. \quad (7)$$

We want to approximate the integral $\int_0^\infty e^{m\phi(x)}$, when m tends to ∞ and we will use the saddle point method. Let us consider the function $\phi(x)$ on $[0, \infty)$. We call the ratio $c = n/m$. The function has the unique maximum on $[0, \infty)$, obtained by $\phi'(x_0) = 0$

$$x_0 = \frac{c + \sqrt{c^2 + 4pc}}{2} = \frac{c}{2}(1 + \sqrt{1 + 4p/c}). \quad (8)$$

The saddle point method says that

$$\int_0^\infty e^{m\phi(x)} dx \approx \sqrt{\frac{2\pi}{p|\phi''(x_0)|}} e^{m\phi(x_0)}, \quad (9)$$

as $m \to \infty$. We evaluate the second derivative in x_0, that is

$$\begin{aligned} |\phi''(x)| &= 1/x_0^2 + 1/(p\sigma^2) \\ &= \frac{1}{n}\Big(\frac{n}{m}\frac{p}{x_0^2} + 1\Big) \\ &= \frac{1}{n}\Big(\frac{4p/c}{(1+\sqrt{1+4p/c})^2} + 1\Big). \end{aligned}$$

Finally, let us evaluate Eq.(9), which will give us the approximation of $\mathrm{E}[X^p]$.

$$\begin{aligned} \mathrm{E}[X^P] &\approx \frac{1}{\sigma\sqrt{p|\phi''(x_0)|}} e^{m\phi(x_0)} \\ &= \frac{1}{\sqrt{cp|\phi''(x_0)|}} x_0^p e^{-\frac{1}{2c}(x-c)^2} \\ &= \Big\{ \frac{n/(cp)}{\frac{4p/c}{(1+\sqrt{1+4p/c})^2} + 1} \Big\}^{1/2} \\ &\times \Big(\frac{c}{2}(1 + \sqrt{1+4p/c}\Big)^p e^{-\frac{c}{8}(\sqrt{1+4p/c}-1)^2}. \end{aligned}$$

Now we will construct a correlated equilibrium. Suppose that for some $x \leq 1$, we allocate xn players to each of $\frac{1}{x}$ strategies. These strategies are selected at random, so a defecting player can do no better than to select a random strategy other than his own. In order for this to produce a correlated equilibrium, we need:

$$(xn)^p \leq \frac{(1/x) - 1}{m-1}(xn+1)^p. \quad (10)$$

We consider that $(1/x)$ and m are both much larger than one (recall that we're interested in the case where m is large in any case) so this inequality essentially states that:

$$x \leq \frac{1}{m}(1 + \frac{1}{xn})^p. \quad (11)$$

It will be sufficient to select x such that $mx = 1 + \frac{p}{xn}$, which must satisfy this constraint since $p \geq 1$. Solving the quadratic equation here yields:

$$x = \frac{1}{2mn}(n + \sqrt{n^2 + 4mnp}) = \frac{1}{2m}(1 + \sqrt{1 + 4(p/c)}). \quad (12)$$

This correlated equilibrium gives us a cost of $(xn)^p$ for each player, which is exactly:

$$n(\frac{c}{2}(1 + \sqrt{1 + 4(p/c)}))^p. \quad (13)$$

Dividing this by the bound on the Nash gives us a Price of Mediation:

$$e^{\frac{c}{8}(\sqrt{1+4p/c}-1)^2} \sqrt{\frac{\frac{4p/c}{(1+\sqrt{1+4p/c})^2} + 1}{n/(cp)}}. \quad (14)$$

This grows exponentially with p. □

5. REFERENCES

[1] I. Ashlagi, D. Monderer, and M. Tenneholtz. On the value of correlation. In *UAI-05: Proceedings of the 21th Annual Conference on Uncertainty in Artificial Intelligence*, 2005.

[2] R. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1:67–96, 1974.

[3] R. J. Aumann. Correlated equilibrium as an expression of bayesian rationality. *Econometrica*, 55(1):1–18, January 1987. available at http://ideas.repec.org/a/ecm/emetrp/v55y1987i1p1-18.html.

[4] X. Chen and X. Deng. Settling the complexity of 2-player nash equilibrium. *Electronic Colloquium on Computational Complexity*.

[5] S.-G. Cheng, D. M. Reeves, Y. Vorobeychik, and M. P. Wellman. Notes on the equilibria in symmetric games. In *International Joint Conference on Autonomous Agents & Multi Agent Systems, 6th Workshop On Game Theoretic And Decision Theoretic Agents*, August 2004.

[6] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. The complexity of computing a nash equilibrium. In *STOC '06: Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 71–78, New York, NY, USA, 2006. ACM.

[7] D. Foster. Calibrated learning and correlated equilibria. *Games95 Conference*, 1995.

[8] A. Jafari, A. R. Greenwald, D. Gondek, and G. Ercal. On no-regret learning, fictitious play, and nash equilibrium. In *ICML*, pages 226–233, 2001.

[9] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria? In *STACS*, 1999.

[10] J. Nash. Equilibrium points in n-person games. *Proceedings of the National Academny of Sciences USA*, 36:48–49, 1950.

[11] C. H. Papadimitriou. Computing correlated equilibria in multi-player games. In *STOC '05: Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 49–56, New York, NY, USA, 2005. ACM.

[12] C. H. Papadimitriou and T. Roughgarden. Computing equilibria in multi-player games. In *SODA '05: Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 82–91, Philadelphia, PA, USA, 2005. Society for Industrial and Applied Mathematics.

[13] R. Peeters and J. Potters. On the structure of the set of correlated equilibria in two-by-two bimatrix games. Technical report, 1999.

[14] T. Roughgarden. Intrinsic robustness of the price of anarchy. In *STOC '09: Proceedings of the forty-first annual ACM Symposium on Theory of Computing*, New York, NY, USA, 2009. ACM.

[15] T. Roughgarden and E. Tardos. How bad is selfish routing? In *IEEE Symposium on Foundations of Computer Science*, pages 93–102, 2000.

APPENDIX

A. PROOF FOR TWO-BY-TWO GAMES

Proof of Lemma 3.2. Let Γ be a two-by-two game with $PoM(\Gamma) \geq P > 1$. First, similarly to [1] Γ has exactly two pure Nash equilibria which are on a diagonal of Γ and Γ has no dominated strategies, because in every other case we obtain that there is no social utility for correlated equilibria lying outside the convex hull of Nash equilibria. Without loss of generality, we assume that the two pure Nash equilibria involve both players selecting strategy one, or both players selecting strategy two. We write the costs in matrix form as follows:

$$\begin{pmatrix} a,b & j,k \\ m,n & c,d \end{pmatrix}$$

By our assumptions about the pure Nash equilibria, we can conclude that the following quantities are negative (else a player would defect from a pure Nash).

$$\delta_1 = a-m, \delta_2 = b-k, \delta_3 = c-j, \delta_4 = d-n$$

Let $\alpha = \frac{\delta_1}{\delta_3}, \beta = \frac{\delta_2}{\delta_4}$, similar to [1, 13] and the following are the correlated equilibria of Γ:

$$\begin{aligned} V_1 &= \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \\ V_2 &= \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} \\ V_3 &= \frac{1}{1+\alpha+\beta+\alpha\beta} \begin{pmatrix} 1 & \alpha \\ \beta & \alpha\beta \end{pmatrix} \\ V_4 &= \frac{1}{1+\alpha+\alpha\beta} \begin{pmatrix} 1 & \alpha \\ 0 & \alpha\beta \end{pmatrix} \\ V_5 &= \frac{1}{1+\beta+\alpha\beta} \begin{pmatrix} 1 & 0 \\ \beta & \alpha\beta \end{pmatrix}. \end{aligned}$$

V_1, V_2, V_3 are Nash equilibria while V_4, V_5 are non-Nash correlated equilibria. Let $c(V_i)$ denote the expected cost of V_i on Γ, in terms of the absolute values of the payoffs.

We may further say *wlog* that (j,k) is the worst cost entry (i.e. $j+k = max\{j+k, a+b, c+d, m+n\}$) and that the Nash (a,b) is worse than the Nash (c,d) (i.e. $a+b \geq c+d$). Since $j+k$ is taken to be the largest social cost, it follows that $c(V_4) \geq c(V_5)$. In fact, we may say that $j+k$ is the *strictly* worst entry as well for the same reason that otherwise $c(V_3)$ would be the worst equilibrium. So, we have that $PoM(\Gamma) = \min\{\frac{c(V_4)}{c(V_1)}, \frac{c(V_4)}{c(V_3)}\} \geq P$. In particular, $\frac{c(V_4)}{c(V_3)} \geq P$.

$$\frac{c(V_4)}{c(V_3)} = (1+\frac{\beta}{1+\alpha+\alpha\beta})(\frac{1}{1+\beta\frac{m+n}{a+b+\alpha(j+k)+\alpha\beta(c+d)}}) \geq P \tag{15}$$

Therefore, if $PoM(\Gamma) \geq 2$ then

$$(1+\frac{\beta}{1+\alpha+\alpha\beta})(\frac{1}{1+\beta\frac{m+n}{a+b+\alpha(j+k)+\alpha\beta(c+d)}}) \geq 2 \Rightarrow$$

$$\frac{1}{1+\alpha+\alpha\beta} \geq \frac{1}{\beta} + \frac{2(m+n)}{a+b+\alpha(j+k)+\alpha\beta(c+d)}$$

Note the following as well: In Γ if player A has no cost of zero value and N corresponds to the worst Nash expected total payoffs while C corresponds to the worst non-Nash expected total payoffs, then $\hat{\Gamma}$ in which all of A's costs are decreased by some ϵ (which is less than the smallest cost to player A) has identical equilibria (probability distributions) to Γ and a strictly worse PoM (i.e. $PoM(\hat{\Gamma}) = \frac{\hat{C}}{\hat{N}} = \frac{|C|-\epsilon}{|N|-\epsilon} > \frac{|C|}{|N|} = PoM(\Gamma)$). Therefore, it is *wlog* that both players have a cost of value 0. In fact, it is *wlog* that $c = 0$ based on previous remarks. We may also normalize any one remaining non-zero cost to 1 in a consistent manner.

Therefore, the cases to consider are:

1. $c = b = 0$ and $a = 1$
2. $c = d = 0$ and $b = 1$
3. $c = d = a = b = 0$ and $m + n = 1$

We first show that $PoM(\Gamma) < 2$ in the first case: Now noting that in this case, $j = \frac{m-1}{\alpha}$ and $k = \beta(n - d)$, we substitute and obtain

$$\begin{aligned}
\frac{1}{1+\alpha+\alpha\beta} &\geq \frac{1}{\beta} + \frac{2(m+n)}{1+\alpha(\frac{m-1}{\alpha}+\beta(n-d))+\alpha\beta(d)} \\
\frac{1}{1+\alpha+\alpha\beta} &\geq \frac{m+n}{m+\alpha\beta(n)} \\
m+\alpha\beta(n) &> (m+n)(1+\alpha+\alpha\beta) \\
0 &> m\alpha + m\alpha\beta + n + n\alpha
\end{aligned}$$

This is clearly a contradiction since everything is positive. Therefore, we have that $PoM(\Gamma) < 2$ for the first case.

For the second case, we also show that $PoM(\Gamma) < 2$: Note here that $j = \frac{m-a}{\alpha}$ and $k = \beta n + 1$. Assuming $PoM(\Gamma) \geq 2$ and substituting into inequality 16

$$\begin{aligned}
\frac{1}{1+\alpha+\alpha\beta} &\geq \frac{1}{\beta} + \frac{2(m+n)}{1+a+\alpha(\frac{m-a}{\alpha}+\beta n+1)} \\
\frac{1}{1+\alpha+\alpha\beta} &\geq \frac{1}{\beta} + \frac{2(m+n)}{1+m+\alpha\beta n+\alpha} \\
1 &\geq \frac{2(m+n)(1+\alpha+\alpha\beta)}{1+m+n\alpha\beta+\alpha} + \frac{1+\alpha}{\beta} + \alpha \\
1 &> m(1+3\alpha+2\alpha\beta) + n(2+3\alpha+\alpha\beta) \\
m+2n &< 1
\end{aligned}$$

Since $|c(V_1)| = 1$, $|c(V_4)| \geq PoM(\Gamma)(1+a)$, if $PoM(\Gamma) > 2$, we obtain by substitutions for j, k and noting that $n < 1$

$$\begin{aligned}
\frac{1+a+\alpha(j+k)}{1+\alpha(1+\beta)} &\geq 2+2a \\
1+a+\alpha(j+k) &\geq 2+2a+(2+2a)\alpha(1+\beta) \\
\alpha(j+k) &\geq 1+a+2\alpha(1+\beta) \\
j+k &\geq \frac{1+a}{\alpha} + 2(1+\beta) \\
\frac{m-a}{\alpha} + n\beta + 1 &\geq \frac{1+a}{\alpha} + 2 + 2\beta \\
\frac{m-a}{\alpha} &\geq \frac{1+a}{\alpha} + 1 + (2-n)\beta \\
\frac{m-a}{\alpha} &> \frac{1+a}{\alpha} \\
m &> 1+2a
\end{aligned}$$

This gives us inequalities indicating that $m < 1$ and also $m > 1$, which is a contradiction, establishing that in the second case we must have a price of mediation bounded by two.

We complete the proof with the final case: Assume, for the sake of contradiction, that in the third case $PoM(\Gamma) \geq 2$. Then, substituting into Eq. 15 we have

$$\begin{aligned}
(1+\frac{\beta}{1+\alpha+\alpha\beta})(\frac{1}{1+\beta\frac{1}{\alpha(j+k)}}) &\geq 2 \\
\frac{\beta}{1+\alpha(1+\beta)}\frac{1}{1+\frac{\beta}{\alpha(j+k)}} &> 1 \\
\frac{\beta}{1+\frac{\beta}{\alpha(j+k)}+\alpha(1+\beta)+(\beta+1)\frac{\beta}{j+k}} &> 1 \\
\frac{\beta}{j+k}(\frac{1}{\alpha}+\beta+1) &< \beta \\
\frac{1}{\alpha}+\beta+1 &< j+k
\end{aligned}$$

However, noting that $\alpha = \frac{m}{j} < \frac{1}{j}$ and $\beta = \frac{k}{n} > k$ (because $m + n = 1$), we clearly have a contradiction. □

B. THREE-BY-THREE GAMES

PROOF OF THEOREM 3.3. Define game Γ:

$$\begin{pmatrix} 10,1 & 0,0 & 0,1 \\ 1,2 & 5,3 & 1,4 \\ 0,5 & 2,2 & 10,0 \end{pmatrix}$$

Let A be the matrix representing the costs to player 1, and B the matrix representing the costs to player 2. Finally, define the following probability matrix V_C:

$$\begin{pmatrix} 0 & \frac{1}{15} & 0 \\ \frac{3}{5} & 0 & \frac{1}{15} \\ \frac{1}{5} & \frac{1}{30} & \frac{1}{30} \end{pmatrix}$$

We first prove that the social cost of the worst Nash equilibrium is 0. It is enough to consider supports of equal size for the two players. We first consider the case when the supports are of size 1 (i.e. pure Nash). Clearly, there is only one pure Nash which is reached when player one plays strategy one and player two plays strategy two, for a social cost of 0. We now consider the case when the supports are of size 2. If player one plays the mixed strategy $\mathbf{p} = (p_1, p_2, 0)$ and player two plays the mixed strategy $\mathbf{q} = (q_1, q_2, 0)^\top$, then to be at equilibrium player two must make player one indifferent between strategies one and two:

$$10q_1 = q_1 + 5q_2$$

$$q_1 + q_2 = 1$$

which solves to $q_1 = \frac{5}{14}$ and $q_2 = \frac{9}{14}$. However, player one's strategy would not be a best response since $A\mathbf{q} = (\frac{25}{7}, \frac{25}{7}, \frac{9}{7})^\top$. Thus, player one should defect and play strategy three. Consider the case when player one plays $\mathbf{p}$ and player two plays $\mathbf{q} = (q_1, 0, q_3)^\top$. For player two to be indifferent, p_1 and p_2 must satisfy the following: $p_1 + 2p_2 = p_1 + 4p_2$. Assuming the probabilities sum to 1, this gives $p_1 = 1$, $p_2 = p_3 = 0$, which cannot be since we assumed p_2 to be in the support of player one's mixed strategy. Now, suppose player one still plays $\mathbf{p}$ while player two plays $\mathbf{q} = (0, q_2, q_3)^\top$. For player one to be indifferent, we must have $5q_2 + q_3 = 0$, which gives $q_2 = -\frac{1}{4}$, $q_3 = \frac{5}{4}$, a contradiction. Suppose now that player one plays the mixed strategy

$\mathbf{p} = (p_1, 0, p_3)$ and player two plays $\mathbf{q} = (q_1, q_2, 0)^\top$. For player two to be indifferent, we must have $p_1 + 5p_3 = 2p_3$, which gives $p_1 = \frac{3}{2}$ and $p_3 = -\frac{1}{2}$, a contradiction. If player one plays $\mathbf{p}$ and player two plays $\mathbf{q} = (q_1, 0, q_3)^\top$, then we must have $p_1 + 5p_3 = p_1$, which solves to $p_1 = 1$, $p_3 = 0$, a contradiction since we are considering supports of size 2. Assuming player one's strategy is the same and player two plays $\mathbf{q} = (0, q_2, q_3)^\top$, for player one to be indifferent we have $0 = 2q_2 + 10q_3$, which has the unique solution $q_2 = \frac{5}{4}$ and $q_3 = -\frac{1}{4}$, a contradiction. Now consider the case when player one plays the mixed strategy $\mathbf{p} = (0, p_2, p_3)$ while player two plays $\mathbf{q} = (q_1, q_2, 0)^\top$. For player one to be indifferent, we have $q_1 + 5q_2 = 2q_2 \Rightarrow q_1 = \frac{3}{2}, q_2 = -\frac{1}{2}$, a contradiction. If player one still plays $\mathbf{p}$ while player two plays $\mathbf{q} = (q_1, 0, q_3)^\top$, then in order for player two to be indifferent we have $2p_2 + 5p_3 = 4p_2$, which has the unique solution $p_2 = \frac{5}{7}, p_3 = \frac{2}{7}$. With these probabilities, player two's expected payoffs are $\mathbf{p}B = (\frac{20}{7}, \frac{19}{7}, \frac{20}{7})$. Clearly, player two would rather defect and play strategy two, so this is not a Nash equilibrium. The last supports of size two are when player one plays $\mathbf{p}$ and player two plays the mixed strategy $\mathbf{q} = (0, q_2, q_3)^\top$. For player one to be indifferent, we must have $5q_2 + q_3 = 2q_2 + 10q_3 \Rightarrow q_2 = \frac{3}{4}, q_3 = \frac{1}{4}$. If player two plays this mixed strategy, then player one's payoffs are $A\mathbf{q} = (0, 4, 4)^\top$. Once again, player one prefers to defect to strategy one implying that this is not a Nash equilibrium. Finally, the last case to consider is when both supports are of size 3, with player one playing $\mathbf{p} = (p_1, p_2, p_3)$ and player two playing $\mathbf{q} = (q_1, q_2, q_3)^\top$. For player two to be indifferent between all three strategies, we must have:

$$p_1 + 2p_2 + 5p_3 = 3p_2 + 2p_3 = p_1 + 4p_2$$

$$p_1 + p_2 + p_3 = 1$$

This system of equations has the unique solution $p_1 = -\frac{1}{6}, p_2 = \frac{5}{6}, p_3 = \frac{1}{3}$, which is a contradiction. Thus, we can conclude that the worst and only Nash equilibrium happens when player one plays strategy one with probability 1, and player two plays strategy two with probability 1. The social cost of this Nash equilibrium is 0.

We now verify that the probability matrix V_C is actually a correlated equilibrium for the game Γ. To see this, consider the following matrix product $A \cdot V_C^\top$:

$$\begin{pmatrix} 10 & 0 & 0 \\ 1 & 5 & 1 \\ 0 & 2 & 10 \end{pmatrix} \begin{pmatrix} 0 & \frac{3}{5} & \frac{1}{5} \\ \frac{1}{15} & 0 & \frac{1}{30} \\ 0 & \frac{1}{15} & \frac{1}{30} \end{pmatrix} = \begin{pmatrix} 0 & 6 & 2 \\ \frac{1}{3} & \frac{2}{3} & \frac{2}{5} \\ \frac{2}{15} & \frac{2}{3} & \frac{2}{5} \end{pmatrix}$$

It is easy to see that if the mediator tells player one to play strategy i, then he does not want to defect. This is because the i^{th} entry in column i takes on the minimum value for that column. Similarly, one can verify that player two does not want to defect either by considering the matrix product $V_C^\top \cdot B$:

$$\begin{pmatrix} 0 & \frac{3}{5} & \frac{1}{5} \\ \frac{1}{15} & 0 & \frac{1}{30} \\ 0 & \frac{1}{15} & \frac{1}{30} \end{pmatrix} \begin{pmatrix} 1 & 0 & 1 \\ 2 & 3 & 4 \\ 5 & 2 & 0 \end{pmatrix} = \begin{pmatrix} \frac{11}{5} & \frac{11}{5} & \frac{12}{5} \\ \frac{7}{30} & \frac{1}{15} & \frac{1}{15} \\ \frac{3}{10} & \frac{4}{15} & \frac{4}{15} \end{pmatrix}$$

Once again, if the mediator tells player two to play strategy i, then he does not wish to defect since the i^{th} entry in row i takes on the minimum value for that row. So V_C is a correlated equilibrium, which has social cost $3 \cdot \frac{3}{5} + 5 \cdot \frac{1}{15} + 5 \cdot \frac{1}{5} + 4 \cdot \frac{1}{30} + 10 \cdot \frac{1}{30} = \frac{18}{5}$. We can increase all the costs in Γ by ϵ such that the correlated equilibrium is only slightly perturbed while the social cost of the worst and only Nash will look like 2ϵ, resulting in a Price of Mediation of at least $\frac{9}{5\epsilon}$, which we can make arbitrarily bad.

□

C. TOTALLY SYMMETRIC GAMES

We investigate bounds on the Price of Mediation for two and three player *totally symmetric* congestion games. In these games all strategies have cost $f(x)$ (i.e. independent of the strategy).

Theorem C.1. *The PoM for two player totally symmetric games is* 1.

Proof. Given two player totally symmetric game $\Gamma_{2,m,f}$, let CE_2 be the worst social cost correlated equilibrium for $\Gamma_{2,m,f}$. Say that p is the probability that the players are placed in the same strategy by CE_2, so that $1-p$ is the probability of being alone. Assume WLOG that CE_2 is canonical, so the only possible reasonable deviation that either player might consider is to pick a different strategy uniformly at random. Note that due to the cost maximization objective of the worst correlated equilibrium, is it also WLOG that the no-regret inequality becomes a tight equality in CE_2. Namely, where the left side is the current cost, and the right side the cost incurred if the player were to deviate, we obtain:

$$pf(2) + (1-p)f(1) = \frac{1-p}{m-1}f(2) + (1 - \frac{1-p}{m-1})f(1) \quad (16)$$

which implies that $p = \frac{1}{m}$, yielding the same social cost as the uniformly at random mixed Nash Equilibrium. □

Theorem C.2. *The PoM for three-player congestion totally symmetric games is at worst $O(m)$ where m is the number of strategies. Further, there exist games where the PoM is $\Omega(m)$.*

Proof. A correlated equilibrium can place all three players separately with probability p_A, or place two together and one separately with probability p_B, or place all three together with probability p_C. We assume without loss of generality that the players are shuffled at random. So the cost of the correlated is:

$$f(1)(p_A + \frac{1}{3}p_B) + f(2)(\frac{2}{3}p_B) + f(3)p_C$$

In a congestion game, we can assume that $f(1) = 0$ since otherwise we can reduce all costs to obtain a worse Price of Mediation. We can also assume that $f(3) = 1$ by scaling and $0 \leq f(2) \leq 1$.

If we defect, the cost of the correlated would be given by:

$$f(2)(\frac{2}{m-1}p_A + \frac{2}{3}\frac{1}{m-1}p_B) + \frac{1}{3}\frac{1}{m-1}p_B$$

For the worst correlated, these expressions must be equal, which gives us:

$$f(2)(\frac{2m-4}{3m-3}p_B - \frac{2}{m-1}p_A) + p_C - \frac{1}{3m-3}p_B = 0$$

This gives us:

$$p_C = \frac{1}{3m-3}p_B - f(2)\frac{2m-4}{3m-3}p_B + f(2)\frac{2}{m-1}p_A$$

We conclude that $p_C \leq \frac{2}{m-1}$. If we solve for p_B, we get:

$$p_B(f(2)\frac{2m-4}{3m-3} - \frac{1}{3m-3}) = f(2)\frac{2}{m-1}p_A - p_C$$

Multiplying through, we get:

$$p_B(f(2)(2m-4) - 1) \leq 6f(2)$$

Thus if $f(2) \geq \frac{1}{m-2}$ we have $p_B(f(2)(2m-4)) \leq 12f(2)$ and $p_B \leq \frac{6}{m-2}$. From this it follows that the cost of the correlated is at most $\frac{8}{m-2}$ whereas the uniformly random Nash has cost at least $\frac{1}{m^2}$, giving a PoM of $O(m)$. If $f(2) \leq \frac{1}{m-2}$ then we have $p_B f(2) \leq \frac{1}{m-2}$ and thus the cost of the correlated is at most $\frac{3}{m-2}$ for a similar result.

We can construct such a game by setting $f(3) = 1$ and otherwise zero, and then observing that the cost of the worst Nash is exactly $\frac{3}{m^2}$. We construct a correlated which places all three players together with probability $p_C = \frac{1}{3m-2}$ and otherwise places two together with probability $p_B = \frac{3m-3}{3m-2}$. This correlated has cost $\frac{3}{3m-2}$, for a price of mediation of $\Omega(m)$. □

D. TOTALLY SYMMETRIC CONCAVE

We prove that for totally symmetric concave congestion games (i.e. all strategies have the same concave increasing cost function $f(x)$) that the Price of Mediation is at most 2. Note that the example where the PoM was larger than one in theorem 4.3 was in fact a totally symmetric game.

THEOREM D.1. *The PoM for the class of totally symmetric concave congestion games is bounded by* 2.

PROOF. Consider any totally symmetric concave congestion game. Let N_{ij} denote the number of players other than i who select strategy j. Since $\Sigma_j N_{ij} = n-1$, for any i, there must be some strategy j where $E[N_{ij}] \leq \frac{n-1}{m}$. It follow from concavity that for this strategy, we must have $E[f(N_{ij}+1)] \leq f(E[N_{ij}+1]) \leq f(\frac{n-1}{m}+1)$. If player i has larger expected cost than this, he can simply defect (and we would not have a correlated equilibium) so this expression also bounds the expected cost for player i. We conclude that for any correlated equilibrium, the expected cost is at most $nf(\frac{n-1}{m}+1)$.

We write $n = mq + r$ with $0 \leq r < m$. Note that the following is a Nash Equilibrium (for the same reason that u.a.r. selection is a Nash equilibrium): mq players places themselves into the m strategies via deterministic groupings of size q each, while the remaining r players are placed deterministically one per strategy. The social cost of this Nash equilibrium is exactly $r(q+1)f(q+1) + (m-r)(q)f(q)$.

Pick a strategy uniformly at random. The social cost of this Nash equilibrium is at least $nf(q)$. We may bound the PoM as follows:

$$\begin{aligned}
\frac{SC(CE_{worst})}{NR_{worst}} &\leq \frac{nf(1+\frac{n-1}{m})}{(r(q+1)f(q+1)+(m-r)(q)f(q)} \\
&\leq \frac{n(n+m-1)}{mr(q+1)^2+m(m-r)q^2} \\
&\leq \frac{(mq+r)(mq+r+m-1)}{m^2q^2+2mrq+mr} \\
&\leq \frac{m^2q^2+2mrq+m^2q+r^2+mr}{m^2q^2+2mrq+mr} \\
&\leq 1+\frac{m^2q+r^2}{m^2q^2+2mrq+mr} \\
&\leq 2
\end{aligned}$$

The second inequality follows from the concavity of f; in particular if $x > y$ we will have $\frac{f(x)}{f(y)} \leq \frac{x}{y}$. Note that this relies upon function f being non-negative and concave over the domain of all non-negative integers (in particular this includes $f(0) \geq 0$ even though no player will ever actually pay the cost $f(0)$). The rest is algebra, along with applying $n = mq + r$ and observing that $r < m$. □