

The Price of Mediation

Milan Bradonjic, Gunes Ercal, Adam Meyerson, Alan Roytman

► **To cite this version:**

Milan Bradonjic, Gunes Ercal, Adam Meyerson, Alan Roytman. The Price of Mediation. Discrete Mathematics and Theoretical Computer Science, DMTCS, 2014, Vol. 16 no. 1 (in progress) (1), pp.31–60. <hal-01179210>

HAL Id: hal-01179210

<https://hal.inria.fr/hal-01179210>

Submitted on 22 Jul 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The Price of Mediation

Milan Bradonjić^{1*} Gunes Ercal^{2†} Adam Meyerson^{3‡} Alan Roytman^{3‡}

¹Mathematics of Networks and Communications, Bell Labs, Alcatel-Lucent, Murray Hill, NJ, USA

²Southern Illinois University, Edwardsville, Computer Science Department, USA.

³University of California, Los Angeles, Computer Science Department, USA.

received 12th June 2012, revised 18th Apr. 2013, 21st Oct. 2013, 17th Dec. 2013, accepted 31st Jan. 2014.

We study the relationship between correlated equilibria and 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 social cost (or utility) of the worst correlated equilibrium to the social cost (or utility) of the worst Nash equilibrium. 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 two player, two strategy games we give tight bounds in the non-negative cost model and the non-negative utility model. For larger games (either more players, or more pure strategies per player, or both) we show that the mediator can be arbitrarily harmful. We also have many results on symmetric singleton 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 the 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.

Keywords: Nash equilibria, Correlated equilibria.

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 (see Aumann (1974, 1987)) instead of a Nash equilibrium (see Nash (1950)).

Correlated equilibria are interesting for a number of reasons. In many cases, their computation is tractable (e.g., Papadimitriou (2005); Papadimitriou and Roughgarden (2005)), whereas computing Nash equilibria is PPAD-complete even for 2-player matrix games (see Chen and Deng (2005); Daskalakis et al. (2006)). The multi-agent machine learning community has observed that for *no-regret* algorithms, play converges to a correlated equilibrium as shown in Foster (1995), whereas only under additional restrictions can convergence to a Nash equilibrium be guaranteed, as in Jafari et al. (2001).

*Email: milan@research.bell-labs.com. The work was done at the University of California, Los Angeles and Los Alamos National Laboratory.

†Email: gercal@siue.edu

‡Emails: {awm|alanr}@cs.ucla.edu. Research partially supported by NSF CIF Grant CCF-1016540.

This paper measures the potential social cost or social unhappiness imposed by inept or malicious mediation. We observe that in many cases it is non-trivial to produce an optimum mediator, 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** (or PoM for short). When studying games with non-negative costs, the Price of Mediation is defined as the ratio of the social cost of the worst (i.e., largest cost) correlated equilibrium solution to the social cost of the worst Nash equilibrium solution. When studying games with non-negative utilities, the Price of Mediation is defined as the ratio of the social utility of the worst (i.e., smallest utility) correlated equilibrium solution to the social utility of the worst Nash equilibrium solution. We show that even for two-player matrix games with only two strategies per player and non-negative costs, the PoM can be as large as a factor of two (this is tight). For two-player matrix games with only two strategies per player and non-negative utilities, the PoM can be as small as a factor of $\frac{2}{3}$ (this is tight as well). For larger matrix games, we show that the PoM can be arbitrarily large in the context of the cost model and arbitrarily close to 0 in the context of the utility model. Thus the social harm imposed by inept or malicious mediation can in fact be quite severe.

We also give many results which bound the PoM for symmetric singleton congestion games (also known as load balancing games). These games have been well-studied in previous work from the standpoint of Price of Anarchy (see Koutsoupias and Papadimitriou (1999)), and are a natural place to start in measuring the PoM. We show that for symmetric singleton 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. (2008, 2005) measures the social benefit of optimum mediation in some simple classes of games. The paper considers games with non-negative utilities, whereas we consider games with non-negative costs and games with non-negative utilities. Major results in Ashlagi et al. (2008) 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 with non-negative costs is tightly bounded by 2, and in the non-negative utility model we show that it is tightly bounded by $\frac{2}{3}$. Similar examples to those in Ashlagi et al. (2008) exist to show that the Price of Mediation is unbounded in 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.

A recent result in Roughgarden (2009) compares the worst correlated equilibrium to the social optimum. It is shown 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 equilibrium 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 (in the context of non-negative costs). Our results complement Roughgarden (2009) by bounding the ratio of the worst correlated equilibrium and Nash equilibrium 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 n players, a strategy set S_i for each player, and a cost (or utility) function c_i (or u_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 (utility) function maps each outcome to a real number. In this paper we will deal with games having non-negative costs (or equivalently non-positive utilities) and games having non-negative utilities.

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

$$\sum_s c_i(s_{-i}, s_i^*)p(s).$$

When player i is suggested to play strategy s_i^* , then his expected cost when following the suggestion is:

$$\sum_{s:s_i=s_i^*} c_i(s) \frac{p(s)}{p_i(s_i^*)}.$$

Here, $p_i(s_i^*)$ is the probability that player i is suggested to play strategy s_i^* . When player i is suggested to play strategy s_i^* but defects to strategy s'_i , then his expected cost is:

$$\sum_{s:s_i=s_i^*} c_i(s_{-i}, s'_i) \frac{p(s)}{p_i(s_i^*)}.$$

We can similarly define the expected utility for a player by replacing c_i with u_i . 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 (or does not increase the expected utility). Namely, comparing the two equations above and cancelling the conditioning, a correlated equilibrium is defined by the following inequality:

$$\sum_{s:s_i=s_i^*} c_i(s)p(s) \leq \sum_{s:s_i=s_i^*} c_i(s_{-i}, s'_i)p(s), \forall s_i^*, s'_i$$

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 (or maximize the expected utility), 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 $p(s_1, s_2, \dots, s_n) = p_1(s_1)p_2(s_2)\dots p_n(s_n)$ for all outcomes.

The social cost of a distribution is $C(p) = \sum_s p(s) \sum_i c_i(s)$, and the social utility of a distribution is defined to be $U(p) = \sum_s p(s) \sum_i u_i(s)$. We will consider the **Price of Mediation (PoM)**. For games with non-negative costs, this is defined to be 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. In particular, for the cost

model, if CE denotes the set of correlated equilibria and NE denotes the set of Nash equilibria, we define $\text{PoM} = \frac{\max_{p \in CE} C(p)}{\max_{p \in NE} C(p)}$. For the cost model, if the largest correlated equilibrium is strictly positive and the largest Nash equilibrium is 0, we say that the PoM is infinite. If both are 0, we say the PoM is 1.

Similarly, for games with non-negative utilities, this is defined to be the ratio of the minimum $U(p)$ where p is a correlated equilibrium divided by the minimum $U(p)$ where p is a Nash equilibrium. More formally, for the utilities model we have $\text{PoM} = \frac{\min_{p \in CE} U(p)}{\min_{p \in NE} U(p)}$. If the smallest Nash equilibrium is 0 (and hence the smallest correlated equilibrium is also 0), we say the PoM is 1. In contrast to the cost model, a larger PoM in the utility model is better in the sense that the mediator cannot harm the players as much. The PoM is a property of the particular game, and since every Nash equilibrium is also a correlated equilibrium we will always have that the PoM is at least one for games with non-negative costs and the PoM is at most one for games with non-negative utilities. 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 (or minimum value in the utility model).

It is possible to consider other ways to measure how much harm a mediator can cause to society. For instance, we could also consider the ratio of the worst correlated equilibrium to the best Nash equilibrium. Such a ratio would always be larger than the Price of Mediation in the cost model and smaller in the utility model. However, this ratio is perhaps less interesting to study than the Price of Mediation. For instance, for matrix games with two players and two strategies, it is already easy to construct examples in which the ratio is infinite in the cost model and arbitrarily close to zero in the utility model.

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 in the context of non-negative costs, and as small as $\frac{2}{3}$ in the context of non-negative utilities. For larger games, we show that the PoM can be infinite in the non-negative cost model and arbitrarily close to 0 in the non-negative utility model.

We first briefly discuss why the bounds we obtain within the cost model are not related to the bounds we obtain for the utility model. Typically, in the literature, cost functions and utility functions map outcomes to the set of all real numbers. Moreover, a direct relationship can be made between costs and utilities, namely the cost (or utility) for a player i in outcome s is $c_i(s) = -u_i(s)$. Since we only deal with games which have non-negative costs and non-negative utilities, we no longer have this relationship and hence should not expect identical results in the two different models.

Note that adding a scalar to all costs or to all utilities 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 equilibrium has negative cost and the worst correlated equilibrium has positive cost. For this reason we consider exclusively games with non-negative costs and games with non-negative utilities in this paper (allowing both signs to be present in a single game matrix immediately yields infinite Price of Mediation). Similar issues arise when considering the utility model.

We now give some intuition on how our lower bounds in the case of the cost model and upper bounds in the case of the utility model are constructed. As a general comment, we avoid games which have dominated strategies, as too many constraints are placed on which types of correlated equilibria are con-

structible in such games. For games where we wish to show the Price of Mediation is arbitrarily close to some constant, we first start with games in which the worst Nash equilibrium is 1 and then create outcomes with larger costs (or smaller utilities) for which it is possible to construct correlated equilibria that place positive probability mass on such outcomes. For games in which we wish to show the Price of Mediation is infinite in the cost model, we either construct games with one Nash equilibrium with a social cost of 0 and set other matrix entries such that there exist correlated equilibria that place positive probability mass on outcomes with positive cost, or construct a family of games in which the cost of an outcome is much larger than other outcomes and yet any Nash equilibrium must play such an outcome with much smaller probability than some correlated equilibria. For games in which we wish to show the Price of Mediation is arbitrarily close to zero in the utility model, we construct a family of games in which the worst Nash equilibrium plays outcomes that have arbitrarily large utility while setting other matrix entries such that a correlated equilibrium can be constructed that places low probability mass on such outcomes.

Before we begin, let us quickly review how Nash equilibria are determined. To find all pure Nash equilibria, we consider all combinations of pure strategies for each player and check whether a player wants to defect or not. To determine all mixed Nash equilibria, we check that the game is non-degenerate, and then write a system of equations for all supports of equal size for each player that constrain the players to be indifferent among the strategies being tested for in the support.

3.1 Two-by-Two Games

Consider games with two players, each of whom has two possible strategies. We will first 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, and moreover this is tight.

First, we show that there is no such game where the Price of Mediation actually reaches two.

Lemma 3.1 *For any two-by-two game Γ with non-negative costs, we have $PoM(\Gamma) < 2$.*

Proof: Let Γ be a two-by-two game with $PoM(\Gamma) \geq P > 1$. First, similarly to pages 584 and 585 of Ashlagi et al. (2008), Γ has exactly two pure Nash equilibria which are on a diagonal of Γ . In particular, as in Ashlagi et al. (2008), if Γ has less than two pure Nash equilibria, then every correlated equilibrium is induced by a mixed Nash equilibrium, and hence the PoM is 1. If there are 4 pure Nash equilibria, then clearly the PoM must be 1 as well. If there are 3 pure Nash equilibria, then every correlated equilibrium must place zero probability mass on the strategy profile not in equilibrium, and hence again the PoM is 1 (since the social cost of each correlated equilibrium is a convex combination of the social costs of the pure Nash equilibria). Finally, if there are two pure Nash equilibria which occur on the same row or column, then any correlated equilibrium must place zero probability mass on the other two strategy profiles not in equilibrium, and we again obtain the PoM is 1.

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 equilibrium). In fact, we can conclude that they

must all be strictly negative. Otherwise, every extreme point of the set of correlated equilibria is induced by a mixed Nash equilibrium, which implies the PoM is 1. This holds since the correlated equilibrium of largest social cost must lie on an extreme point, as it corresponds to an optimal solution of a linear program with a linear objective function of maximizing the social cost.

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

Let $\alpha = \frac{\delta_1}{\delta_3}, \beta = \frac{\delta_2}{\delta_4}$. Similar to Ashlagi et al. (2008); Peeters and Potters (1999), the following are the five extremal points of the set of correlated equilibria of Γ , and therefore the only relevant equilibria to consider for our purposes of bounding PoM (again, since the correlated equilibrium of largest social cost must lie on an extreme point).

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

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 social cost of V_i .

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 equilibrium (a, b) is worse than the Nash equilibrium (c, d) (i.e., $a+b \geq c+d$). It can be seen that $c(V_3) \geq c(V_5)$, leaving V_4 as the only candidate to consider for our purposes. So, we have that $\text{PoM}(\Gamma) = \min\{\frac{c(V_4)}{c(V_1)}, \frac{c(V_4)}{c(V_3)}\} \geq P$. In particular, $\frac{c(V_4)}{c(V_3)} \geq P$ for any lower bound P on the PoM (note that $c(V_3) \neq 0$).

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

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

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

and moreover

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

Note the following as well: in Γ if player A has no cost of zero value and N corresponds to the worst Nash equilibrium's expected total costs while C corresponds to the worst non-Nash equilibrium's expected

total costs, then $\hat{\Gamma}$ in which all of A 's costs are decreased by some ϵ (which is at most the smallest cost to player A) has identical equilibria (probability distributions) as Γ . Moreover, $\hat{\Gamma}$ has a strictly worse PoM, since $\text{PoM}(\hat{\Gamma}) = \frac{\hat{C}}{N} = \frac{C-\epsilon}{N-\epsilon} > \frac{C}{N} = \text{PoM}(\Gamma)$ (since we assume that $\text{PoM} \geq P > 1$, we have $C > N$). Note that this operation of decreasing player A 's costs by ϵ does not change any of the δ_i ($1 \leq i \leq 4$), and hence α and β are unchanged (the same holds for subtracting the same value from all of player B 's costs). 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 and due to symmetry. We may also normalize any single remaining non-zero cost to 1 in a consistent manner (i.e., we may choose a single non-zero cost x and divide every cost in the matrix by this value so that the “new” cost of x is 1 after normalization). This does not affect the Price of Mediation.

Therefore, we have three cases to consider.

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

In particular, if $b > d$, then we fall into case 2 (since d is the smallest cost to player B and $b \neq 0$). If $b \leq d$, then we fall into case 1 if $a \neq 0$ (since b is the smallest cost to player B) and case 3 if $a = 0$ (note that $m + n \neq 0$ and $d = 0$ after decreasing the costs to players A and B since $c + d \leq a + b$). As a reminder, we have $j, k, m, n > 0$ since the δ_i ($1 \leq i \leq 4$) are strictly negative.

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

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

This is a contradiction since all values are positive. Therefore, we have that $\text{PoM}(\Gamma) < 2$ for the first case.

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

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

We have $c(V_1) = 1 + a$ and $c(V_4) \geq \text{PoM}(\Gamma)(1 + a)$. Hence, if $\text{PoM}(\Gamma) > 2$, we obtain by substitutions for j, k and noting that $n < 1$ (since $m + 2n < 1$):

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

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 $\text{PoM}(\Gamma) \geq 2$. Then, substituting into inequality (1) we have

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

However, note that $\alpha = \frac{m}{j} \leq \frac{1}{j}$ and $\beta = \frac{k}{n} > k$, since $m, n > 0$ and $m + n = 1$ yields $m, n < 1$ (recalling further that $a = d = 0$ and $\delta_i < 0$ for $i = 1, \dots, 4$). Therefore, we clearly have a contradiction. \square

In fact, this bound is tight: for such games the Price of Mediation asymptotically approaches two:

Lemma 3.2 *For any sufficiently small $\delta > 0$, there exists a two-by-two game Γ with non-negative costs such that $\text{PoM}(\Gamma) > 2 - \delta$.*

Proof: Consider the following game matrix Γ_ϵ for any $0 < \epsilon < \frac{1}{2}$:

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

There are three Nash equilibria for this game (two pure and one mixed). Observe that, in the mixed Nash equilibrium, player one (the *row* player) plays strategy one with probability $\frac{\epsilon}{1-\epsilon}$ and player two (the *column* player) plays strategy one with probability $\frac{1}{1+\epsilon^2-2\epsilon^3}$. The social cost of this mixed Nash equilibrium is given by:

$$\frac{\epsilon + \epsilon(1-2\epsilon)(1+2\epsilon^2-2\epsilon^3) + (1-\epsilon)(1-2\epsilon)}{(1-\epsilon)(1+\epsilon^2-2\epsilon^3)} = \frac{1-\epsilon+2\epsilon^3-6\epsilon^4+4\epsilon^5}{(1-\epsilon)(1+\epsilon^2-2\epsilon^3)} = \frac{1+2\epsilon^3-4\epsilon^4}{1+\epsilon^2-2\epsilon^3} \leq 1.$$

Hence, the worst social cost belongs to the pure Nash equilibrium which plays strategy one for each player and has cost 1.

To lower bound the Price of Mediation, it is sufficient to lower bound any correlated equilibrium (not necessarily the worst one). Hence, consider the correlated equilibrium given by V_C :

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

We can verify that V_C is a correlated equilibrium. Suppose player one is told to play strategy two. Then player one knows that, with probability 1, player two is playing strategy two, and hence player one is getting a cost of 0 as opposed to $\frac{1}{\epsilon^2} > 0$. Now suppose player one is told to play strategy one. In this case, player one is currently getting an expected cost of $\frac{1-2\epsilon}{1+\epsilon^2-2\epsilon^3}$, but if he defects then the expected cost would be identical. Now consider what happens if player two is told to play strategy one. Then player two knows that player one is playing strategy one with probability 1, and hence player two is currently getting a cost of 1 but defecting would increase the cost to $2-2\epsilon > 1$. Finally, if player two is told to play strategy two, then his expected cost is 2ϵ no matter which strategy is being played.

We observe that $1+\epsilon-3\epsilon^2+2\epsilon^3 \leq 1+\epsilon$ and lower bound the social cost of V_C by:

$$c(V_C) = \frac{1+\epsilon^2(1-2\epsilon)(\frac{1}{\epsilon^2}+2-2\epsilon)}{1+\epsilon-3\epsilon^2+2\epsilon^3} \geq \frac{1}{1+\epsilon} \left(1 + \frac{\epsilon^2(1-2\epsilon)}{\epsilon^2} \right) = 2 \left(\frac{1-\epsilon}{1+\epsilon} \right).$$

It is clear that by setting ϵ to be sufficiently small we can get this to be arbitrarily close to two, completing the proof. \square

We now consider two player, two strategy games with non-negative utilities and argue that in this case, it is possible for the Price of Mediation to asymptotically approach $\frac{2}{3}$.

Lemma 3.3 *For any $\epsilon > 0$, there exists a two-by-two game Γ with non-negative utilities such that $PoM(\Gamma) \leq \frac{2}{3}(1+\epsilon)$.*

Proof: Consider the following game matrix with non-negative utilities Γ_ϵ :

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

There are three Nash equilibria for this game (two pure and one mixed). Both pure Nash equilibria have a social utility of $1+\epsilon$. The mixed Nash equilibrium occurs when player 1 plays each of the two strategies

with probability $\frac{1}{2}$ and when player 2 does the same. Notice that the cost of this mixed Nash equilibrium is the worst one, giving a social utility of 1.

A correlated equilibrium for this game is given by V_C :

$$V_C = \begin{pmatrix} \frac{1}{3} & \frac{1}{3} \\ 0 & \frac{1}{3} \end{pmatrix}.$$

We can verify that this is in fact a correlated equilibrium as follows. Suppose player 1 is told to play strategy 1. Then his expected utility is $\frac{1}{2}$, but if he switches he will also get an expected utility of $\frac{1}{2}$, so there is no incentive to defect. Now suppose player 1 is told to play strategy 2. In this case, his expected utility is ϵ , but if he defects he will get 0. Hence, there is no incentive for player 1 to ever defect. The argument for player 2 is symmetric.

Thus, V_C is a correlated equilibrium with a social utility of $\frac{1}{3}(1 + \epsilon + 1 + \epsilon) = \frac{2}{3}(1 + \epsilon)$. This implies that $PoM(\Gamma_\epsilon) \leq \frac{2}{3}(1 + \epsilon)$, giving the lemma. \square

In fact we prove that this bound is also tight, since there is no game where the Price of Mediation actually reaches $\frac{2}{3}$.

Lemma 3.4 *For any two-by-two game Γ with non-negative utilities, we have $PoM(\Gamma) > \frac{2}{3}$.*

Proof: Let Γ be a two-by-two game with non-negative utilities:

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

First, similarly to Ashlagi et al. (2008) and Lemma 3.1, Γ has exactly two pure Nash equilibria which are on a diagonal. We can define $\delta_1, \dots, \delta_4$ as in Lemma 3.1, and note that in the utility model we can assume that the quantities are strictly positive.

We know that a, b, c, d cannot be 0, and we can assume without loss of generality (as in Lemma 3.1), both players have a utility of 0. In particular, either $j = k = 0$ or $j = n = 0$ (the other cases of $m = n = 0$ and $m = k = 0$ are symmetric). Let us first prove the result in the case that $j = k = 0$. Since $a \neq 0$, we can normalize all costs by a so that $a = 1$ without increasing the Price of Mediation.

We let $\alpha = \frac{\delta_1}{\delta_3} = \frac{1-m}{c}$, $\beta = \frac{\delta_2}{\delta_4} = \frac{b}{d-n}$ (notice α and β are positive). We have the same set of extremal points of the set of correlated equilibria as in Lemma 3.1, which we denote by V_1, \dots, V_5 . Moreover, we let $U(V_i)$ denote the expected social utility of V_i . Note that we have the following equations for the social utilities:

$$\begin{aligned} U(V_1) &= 1 + b, \\ U(V_2) &= c + d, \\ U(V_3) &= \frac{1 + b + \beta m + \beta n + \alpha\beta(c + d)}{1 + \alpha + \beta + \alpha\beta} = \frac{1 + \beta d + \beta + \alpha\beta d}{1 + \alpha + \beta + \alpha\beta}, \\ U(V_4) &= \frac{1 + b + \alpha\beta(c + d)}{1 + \alpha + \alpha\beta}, \\ U(V_5) &= \frac{1 + b + \beta(m + n) + \alpha\beta(c + d)}{1 + \beta + \alpha\beta} = \frac{1 + \beta d + \beta + \alpha\beta d}{1 + \beta + \alpha\beta}. \end{aligned}$$

Observe that V_5 cannot be the worst equilibrium, since otherwise we would have $\text{PoM}(\Gamma) \geq \frac{U(V_5)}{U(V_3)} > 1$, which is a contradiction (since we know that the Price of Mediation must be at most 1). Hence, we are only interested in the following three ratios:

$$\begin{aligned} R_1 &= \frac{U(V_4)}{U(V_1)} = \frac{1+b+\alpha\beta(c+d)}{(1+\alpha+\alpha\beta)(1+b)}, \\ R_2 &= \frac{U(V_4)}{U(V_2)} = \frac{1+b+\alpha\beta(c+d)}{(1+\alpha+\alpha\beta)(c+d)}, \\ R_3 &= \frac{U(V_4)}{U(V_3)} = \frac{1+b+\alpha\beta(c+d)}{(1+\alpha+\alpha\beta)} \cdot \frac{1+\alpha+\beta+\alpha\beta}{1+\beta d+\beta+\alpha\beta d}. \end{aligned}$$

We are interesting in minimizing the quantity $\max(R_1, R_2, R_3)$. We will argue that, without loss of generality, these three ratios should all be equal, because if not we can tweak the variables so that this becomes true without increasing the quantity. If we assume that α and β are independent of the variables (which only gives us more flexibility in minimizing the max quantity), we observe that R_1 is a decreasing function of b and an increasing function of c , R_2 is an increasing function of b and a decreasing function of c , and R_3 is an increasing function of b and an increasing function of c .

We can use these facts and modify the variables so that all three ratios become equal without increasing the maximum quantity. Hence, we have the following constraints:

$$\begin{aligned} 1+b=c+d &= \frac{1+\beta d+\beta+\alpha\beta d}{1+\alpha+\beta+\alpha\beta} \implies \\ b &= \frac{\beta d+\alpha\beta d-\alpha-\alpha\beta}{1+\alpha+\beta+\alpha\beta}, \quad c = \frac{1-d+\beta-\alpha d}{1+\alpha+\beta+\alpha\beta}. \end{aligned}$$

If $d \geq 1$, then since c is positive we can conclude that $\alpha < \frac{\beta+1}{d} - 1 \leq \beta$. If $d < 1$, then since b is positive we can conclude that $\alpha < \frac{\beta d}{1+\beta-\beta d} < \beta$. Hence, in either case we have that $\alpha < \beta$.

Since all three ratios are equal, we have that

$$\text{PoM}(\Gamma) \geq \frac{1+b+\alpha\beta(1+b)}{(1+\alpha+\alpha\beta)(1+b)} = \frac{1+\alpha\beta}{1+\alpha+\alpha\beta}.$$

We note that $\frac{1+\alpha\beta}{1+\alpha+\alpha\beta}$ is a decreasing function of α , and since $\alpha < \beta$ we have that

$$\text{PoM}(\Gamma) > \frac{1+\beta^2}{1+\beta+\beta^2}.$$

This function is minimized at $\beta = 1$, giving $\text{PoM}(\Gamma) > \frac{2}{3}$. Hence, we have proved the lemma in the case when $j = k = 0$ (and symmetrically when $m = n = 0$).

We now consider the case when $j = n = 0$. We will solve this case by essentially reducing it to our previous cases. The same Nash and correlated equilibria exist, with $\alpha = \frac{a-m}{c}$ and $\beta = \frac{b-k}{d}$. We have the

following equations for the social utilities.

$$\begin{aligned}
U(V_1) &= a + b, \\
U(V_2) &= c + d, \\
U(V_3) &= \frac{a + b + \alpha k + \beta m + \alpha\beta(c + d)}{1 + \alpha + \beta + \alpha\beta} = \frac{a + b + \alpha b + \beta}{1 + \alpha + \beta + \alpha\beta}, \\
U(V_4) &= \frac{a + b + \alpha k + \alpha\beta(c + d)}{1 + \alpha + \alpha\beta} = \frac{a + b + \alpha b + \alpha\beta c}{1 + \alpha + \alpha\beta}, \\
U(V_5) &= \frac{a + b + \beta m + \alpha\beta(c + d)}{1 + \beta + \alpha\beta} = \frac{a + b + \beta + \alpha\beta d}{1 + \beta + \alpha\beta}.
\end{aligned}$$

We assume that k and m are both strictly positive (otherwise we fall into the previous cases and we are done). Suppose the correlated equilibrium V_4 is the worst correlated equilibrium. Observe that we can increase d by a small $\epsilon > 0$ and adjust k appropriately (by decreasing it) so that β stays fixed. Notice that this has no effect on α and V_4 remains the worst correlated equilibrium (in fact, this operation can only increase the worst Nash equilibrium). Hence, we have just obtained a Price of Mediation which is smaller. For this reason, we can keep repeating this process until $k = 0$ and fall into the previous case when $j = k = 0$, giving the lemma.

If the correlated equilibrium V_5 is the worst correlated equilibrium, then we can increase c by a small $\epsilon > 0$ and adjust m appropriately (by decreasing it) so that α stays fixed (this has no effect on β). This operation guarantees that V_5 remains the worst correlated equilibrium and can only increase the worst Nash equilibrium. Hence, we again have obtained a Price of Mediation which is smaller. We can keep repeating this process until $m = 0$ and fall into the previous case when $m = n = 0$. This proves the lemma. \square

3.2 Two Player, Three Strategy Games

We now prove that a mediator can be arbitrarily bad for games with two players, each having at least three strategies. For games with non-negative costs, we argue that the Price of Mediation is infinite, while for games with non-negative utilities, we argue that the Price of Mediation can be arbitrarily close to 0. Note that a Price of Mediation which is infinite in the cost model is qualitatively the same as a Price of Mediation which is arbitrarily close to 0 in the utility model. Let us first prove the result in the cost model.

Theorem 3.5 *The Price of Mediation for the class of two player games with non-negative costs where each player has at least three strategies is infinite.*

Proof: 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, since the game is non-degenerate. In particular, to see why the game is non-degenerate, we must check for both players that no mixed strategy with a support size of s has more than s pure best responses. For supports of size 1, this holds because each row of B and each column of A has a unique minimum entry. Now consider the case when A plays a mixed strategy of support size 2 with probabilities q and $1 - q$. If A mixes strategies 1 and 2, then B would always prefer strategy 1 over strategy 3 since $q + 2(1 - q) < q + 4(1 - q)$. If A mixes strategies 1 and 3, then B would always prefer strategy 2 over strategy 1 since $2(1 - q) < q + 5(1 - q)$. If A mixes strategies 2 and 3, then player B must be made indifferent between all three strategies to have three pure best responses. Hence, we would have $2q + 5(1 - q) = 3q + 2(1 - q) = 4q$, which is impossible.

Now consider the case when B plays a mixed strategy of support size 2 with probabilities q and $1 - q$. If B mixes strategies 1 and 2, then A would always prefer strategy 3 over strategy 2 since $2(1 - q) < q + 5(1 - q)$. If B mixes strategies 1 and 3, then A must be made indifferent between all three strategies to have three pure best responses. Hence, we would have $10q = 1 = 10(1 - q)$, which is impossible. Finally, if B mixes strategies 2 and 3, then A would always prefer strategy 1 (for a cost of 0) over any other strategy. This implies that the game is non-degenerate, so we can consider supports of equal size.

We first consider the case when the supports are of size 1 (i.e., pure Nash equilibria). Clearly, there is only one pure Nash equilibrium 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:

$$\begin{aligned} 10q_1 &= q_1 + 5q_2, \\ q_1 + q_2 &= 1, \end{aligned}$$

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 $q_2 = -\frac{1}{4}, q_3 = \frac{5}{4}$, which is impossible since $q_2, q_3 > 0$.

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 $p_3 = -\frac{1}{2}$, which is impossible since $p_1, p_3 > 0$. 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 $q_3 = -\frac{1}{4}$, which is impossible since $q_2, q_3 > 0$.

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 $q_1 + 5q_2 = 2q_2 \Rightarrow q_1 = \frac{3}{2}, q_2 = -\frac{1}{2}$, which is impossible since $q_1, q_2 > 0$. 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 costs 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 costs 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:

$$\begin{aligned} p_1 + 2p_2 + 5p_3 &= 3p_2 + 2p_3 = p_1 + 4p_2, \\ p_1 + p_2 + p_3 &= 1. \end{aligned}$$

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 since all probabilities must be strictly positive. 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 $A \cdot V_C^\top$:

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

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. In particular, each entry on the diagonal is the cost to the player when following the mediator's suggestion (i.e., each entry (i, i) gives the cost when player one is suggested and follows strategy i). Moreover, within a column every entry other than the diagonal entry corresponds to the cost of defecting to the corresponding strategy. Similarly, one can verify that player two does not want to defect either by considering the matrix product $V_C^\top \cdot B$:

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

We can use the same reasoning as above, considering now rows instead of columns. So V_C is a correlated equilibrium, which has social cost $3 \cdot \frac{3}{5} + 5 \cdot \frac{1}{15} + 5 \cdot \frac{1}{5} + 4 \cdot \frac{1}{30} + 10 \cdot \frac{1}{30} = \frac{18}{5}$. Hence, since the worst correlated equilibrium has a positive social cost while the worst Nash equilibrium has a social cost of zero, the Price of Mediation is infinite.

Moreover, this proves that the Price of Mediation of two player games with non-negative costs and at least three strategies is also infinite. 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 . \square

We now prove our result in the non-negative utilities model.

Theorem 3.6 *The Price of Mediation for the class of two player games with non-negative utilities where each player has at least three strategies is arbitrarily close to 0.*

Proof: Define game Γ :

$$\begin{pmatrix} 15, 5 & 0, 1 & 0, 10 \\ 10, 5 & \frac{1}{\epsilon}, 10 & 20, 0 \\ 0, 10 & 0, 0 & 50, 0 \end{pmatrix}.$$

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

$$\frac{1}{1 + \epsilon} \begin{pmatrix} \frac{1}{2} & 0 & \frac{1}{8} \\ \frac{1}{8} & \epsilon & 0 \\ \frac{3}{16} & 0 & \frac{1}{16} \end{pmatrix}.$$

We first prove that the social utility of the worst Nash equilibrium is $10 + \frac{1}{\epsilon}$. It is enough to consider supports of equal size for the two players, since the game is non-degenerate. In particular, to see why the game is non-degenerate, we must check for both players that no mixed strategy with a support size of s has more than s pure best responses. If player A plays a mixed strategy of support size 1, then B always has a unique pure best response. If A plays strategy 1, player B 's best unique response is strategy 3 for a utility of 10. If A plays strategy 2, B 's unique best response is strategy 2 for a utility of 10. If A plays strategy 3, then B 's unique best response is strategy 1 for a utility of 10. Now consider the case when A plays a mixed strategy of support size 2 with probabilities q and $1 - q$. If A mixes strategies 1 and 2, then B must be made indifferent between all three strategies to have three pure best responses. Hence, we would have $5 = q + 10(1 - q) = 10q$, which is impossible. If A mixes strategies 1 and 3, then B would always prefer strategy 3 over strategy 2 since $10q > q$. If A mixes strategies 2 and 3, then B would never play strategy 3 which would give zero utility.

Similarly, if B plays a mixed strategy of support size 1, then A always has a unique pure best response. If B plays strategy 1, then A 's best response is strategy 1. If B plays strategy 2, then A 's best response is strategy 2. If B plays strategy 3, then A 's best response is strategy 3. Now consider the case when B plays a mixed strategy of support size 2 with probabilities q and $1 - q$. If B mixes strategies 1 and 2, then A would never play strategy 3, which gives zero utility. If B mixes strategies 1 and 3, then A must be made indifferent between all three strategies to have three pure best responses. Hence, we would have $15q = 10q + 20(1 - q) = 50(1 - q)$, which is impossible. Finally, if B mixes strategies 2 and 3, then A would never play strategy 1, which gives zero utility. This implies that the game is non-degenerate, so we need only consider supports of equal size.

We first consider the case when the supports are of size 1 (i.e., pure Nash equilibria). Clearly, there is only one pure Nash equilibrium which is reached when both players play strategy two, for a social utility of $10 + \frac{1}{\epsilon}$.

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 one must make player two indifferent between strategies one and two:

$$\begin{aligned} p_1 + 10p_2 &= 5, \\ p_1 + p_2 &= 1. \end{aligned}$$

which solves to $p_1 = \frac{5}{9}$ and $p_2 = \frac{4}{9}$. However, player two's strategy would not be a best response since player two would rather defect to 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 must satisfy the following: $10p_1 = 5$. Assuming the probabilities sum to 1, this gives $p_1 = p_2 = \frac{1}{2}$. However, player two's strategy would not be a best response since player two would rather defect and play strategy two. Now, suppose player one still plays \mathbf{p} while player two plays $\mathbf{q} = (0, q_2, q_3)^\top$. In this case, player one cannot be made indifferent between strategies one and two since $\frac{q_2}{\epsilon} + 20q_3 > 0$.

Suppose now that player one plays the mixed strategy $\mathbf{p} = (p_1, 0, p_3)$ and player two plays $\mathbf{q} = (q_1, q_2, 0)^\top$. In this case, player one cannot be made indifferent between the two strategies since $15q_1 > 0$. If player one plays \mathbf{p} and player two plays $\mathbf{q} = (q_1, 0, q_3)^\top$, then we must have $15q_1 = 50q_3$, which solves to $q_1 = \frac{10}{13}$, $q_3 = \frac{3}{13}$. However, this would not be a best response since player one would rather defect and play strategy two. Assuming player one's strategy is the same and player two plays $\mathbf{q} = (0, q_2, q_3)^\top$, player one cannot be made indifferent since $50q_3 > 0$.

Now consider the case when player one plays the mixed strategy $\mathbf{p} = (0, p_2, p_3)$ while player two plays $\mathbf{q} = (q_1, q_2, 0)^\top$. Player one cannot be made indifferent since $10q_1 + \frac{q_2}{\epsilon} > 0$. If player one still plays \mathbf{p} while player two plays $\mathbf{q} = (q_1, 0, q_3)^\top$, then player two cannot be made indifferent since $5p_2 + 10p_3 > 0$. The last supports of size two are when player one plays \mathbf{p} and player two plays the mixed strategy $\mathbf{q} = (0, q_2, q_3)^\top$. In this case, player two cannot be made indifferent since $10p_2 > 0$.

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 one to be indifferent between all three strategies, we must have:

$$\begin{aligned} 15q_1 = 50q_3 &= 10q_1 + \frac{q_2}{\epsilon} + 20q_3, \\ q_1 + q_2 + q_3 &= 1. \end{aligned}$$

This system of equations requires $q_3 = \frac{-3q_2}{10\epsilon}$, which is impossible since the probabilities cannot be negative. Thus, we can conclude that the worst (i.e., smallest social utility) and only Nash equilibrium happens when both players play strategy two with probability 1. The social utility of this Nash equilibrium is $10 + \frac{1}{\epsilon}$.

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 $A \cdot V_C^\top$:

$$\frac{1}{1 + \epsilon} \begin{pmatrix} 15 & 0 & 0 \\ 10 & \frac{1}{\epsilon} & 20 \\ 0 & 0 & 50 \end{pmatrix} \begin{pmatrix} \frac{1}{2} & \frac{1}{8} & \frac{3}{16} \\ 0 & \epsilon & 0 \\ \frac{1}{8} & 0 & \frac{1}{16} \end{pmatrix} = \frac{1}{16(1 + \epsilon)} \begin{pmatrix} 120 & 30 & 45 \\ 120 & 36 & 50 \\ 100 & 0 & 50 \end{pmatrix}.$$

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 maximum value for that column. Similarly, one can verify that player two does not want to defect either by considering the matrix product $V_C^\top \cdot B$:

$$\frac{1}{1+\epsilon} \begin{pmatrix} \frac{1}{2} & \frac{1}{8} & \frac{3}{16} \\ 0 & \epsilon & 0 \\ \frac{1}{8} & 0 & \frac{1}{16} \end{pmatrix} \begin{pmatrix} 5 & 1 & 10 \\ 5 & 10 & 0 \\ 10 & 0 & 0 \end{pmatrix} = \frac{1}{8(1+\epsilon)} \begin{pmatrix} 40 & 14 & 40 \\ 40\epsilon & 80\epsilon & 0 \\ 10 & 1 & 10 \end{pmatrix}.$$

We can use the same reasoning as above, considering now rows instead of columns. So V_C is a correlated equilibrium, which has a social utility of $\frac{1}{1+\epsilon}(\frac{145}{8} + 10\epsilon + 1)$. Thus, the Price of Mediation for the game Γ is at most $\frac{\epsilon}{\epsilon+1} + \frac{145}{8(1+\epsilon)(10+\frac{1}{\epsilon})}$, which we can make arbitrarily close to 0 by letting ϵ approach 0.

Moreover, this proves that the Price of Mediation of two player games with non-negative utilities in which both players have at least three strategies is arbitrarily close to 0 by a similar argument as before. \square

3.3 Three Player, Two Strategy Games

We now consider three-player games with costs and three-player games with utilities. We first prove the following theorem:

Theorem 3.7 *The Price of Mediation for the class of three player, two strategy games with non-negative costs is infinite.*

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 equilibrium 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 the mixed Nash equilibrium will be $2(1 + \epsilon)$.

There exists the following correlated equilibrium:

$$V_C = \frac{1}{2-\epsilon} \begin{pmatrix} \epsilon & 1-\epsilon \\ 1-\epsilon & 0 \end{pmatrix}.$$

If player one is suggested 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 is suggested strategy one, then his expected cost is the same whether he defects or not. So it is a correlated equilibrium.

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

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 $M\epsilon^2 \gg 2$. Now for any Nash equilibrium, the probability that both the first two players select strategy

one is at most ϵ^2 so the total cost will be at most $M\epsilon^2 + 2(1 + \epsilon)$. But for the correlated equilibrium, we get a cost of $\frac{4}{2-\epsilon} + \frac{M\epsilon}{2-\epsilon}$. Hence, this gives a PoM of at least

$$\frac{4 + M\epsilon}{2(M\epsilon^2 + 2(1 + \epsilon))} > \frac{4 + M\epsilon}{4\epsilon(1 + M\epsilon)} > \frac{1}{4\epsilon},$$

which we can make arbitrarily bad by decreasing ϵ . \square

In fact, our proof can be extended to handle more strategies and more players. We can simply add more columns and rows to our two player game with sufficiently large values such that neither player would ever consider playing the new strategies. Adding more players to the game can be done in a similar fashion as in the proof.

It is natural to ask whether a similar result can be proven for utilities, which is answered by the following theorem.

Theorem 3.8 *The Price of Mediation for the class of three player, two strategy games with non-negative utilities can be arbitrarily close to 0.*

Proof: Consider the following game Γ which has three players A , B , and C , each of which has two strategies (here, M is a very large number)

$$\begin{array}{cc} C \text{ picks strategy 1} & C \text{ picks strategy 2} \\ \left(\begin{array}{cc} 0, 0, 0 & 1, 2, 1 \\ 2, 4, 0 & 2, 2, M \end{array} \right), & \left(\begin{array}{cc} 0, 0, 0 & 4, 2, 0 \\ 2, 1, 1 & 2, 2, M \end{array} \right). \end{array}$$

For each matrix entry, the first number represents the utility to player A , the second to player B , and the third to player C . This game can be seen as a variant of the classic chicken game between the row player A and the column player B , where strategy 1 is to Go and strategy 2 is to Stop. Notice that player C gets immense pleasure from watching both A and B play Stop.

The only Nash equilibrium occurs when all three players are mixing both of their strategies. To see this, suppose there is a Nash equilibrium where player A only plays strategy 1. Then player B would prefer to play strategy 2, no matter what player C does. Hence, player C would prefer to play strategy 1, in which case player A 's best response is to play strategy 2, a contradiction. Suppose there is a Nash equilibrium where player A only plays strategy 2. Let r_1, r_2 denote the probabilities player C plays strategy 1 and strategy 2, and let q_1, q_2 denote the probabilities player B plays strategy 1 and strategy 2. There are two cases: either $q_1 = 0$ or $q_1 > 0$. If $q_1 = 0$, then to prevent A from defecting to always playing strategy 1, r_1 and r_2 must satisfy $2r_1 + 2r_2 \geq r_1 + 4r_2$, which implies $r_2 \leq \frac{1}{3}$. However, now player B can improve his utility by always playing strategy 1, which gives an expected utility of $4r_1 + r_2 > 2$, a contradiction. On the other hand, if $q_1 > 0$, then player C 's best response is to always play strategy 2. However, now player B can improve his utility by always playing strategy 2 for a utility of $2 > 1$, a contradiction.

The analysis for the reason player B must mix both strategies is symmetric. To see why player C must mix as well, suppose there is a Nash equilibrium where player C always plays strategy 1. Then player A 's best response is to always play strategy 2, in which case player B 's best response is to always play strategy 1. However, now player C can improve his utility from 0 to 1 by always playing strategy 2. The

reason C cannot always play strategy 2 is symmetric. Hence, each player must play each of their two strategies with strictly positive probability in any Nash equilibrium.

Now, suppose player A plays strategy 1 with probability p and player B plays strategy 1 with probability q . Note that p must be strictly less than 1, and q must be strictly less than 1. If $p > q$, then C is not made indifferent between his two strategies. A similar argument can be made if $p < q$. Thus, we can conclude that in the Nash equilibrium we have $p = q$. If player A plays strategy 2, then player A 's utility will always be 2. If player A plays strategy 1, then player A 's utility will always be at most $4(1 - q)$. Hence, from this we conclude that $p = q \leq \frac{1}{2}$. This implies that the probability players A and B choose strategy 2 is at least $\frac{1}{4}$. These facts combined imply that the social utility in the Nash equilibrium is at least $\frac{M}{4}$.

We now construct a correlated equilibrium, which places a probability of $\frac{1}{4}$ on each of the 2 outcomes where players A and B play different strategies and player C chooses strategy 1, and a probability of $\frac{1}{4}$ on each of the 2 outcomes where players A and B play different strategies and player C chooses strategy 2. It can be easily verified that for each player, defecting will only cause them to be unhappier (for instance, C is indifferent as to which strategy he is told to play in this correlated equilibrium). Notice this correlated equilibrium has a social utility of 5. Hence, the Price of Mediation for this game is at most $5 \cdot \frac{4}{M}$, which we can make arbitrarily close to 0 by making M larger, which proves the theorem. \square

We can extend the result for non-negative utilities to handle larger classes of games (either more players or more strategies or both) in a similar way as before.

4 Symmetric Singleton 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 singleton congestion games or load balancing games). Each of n players selects one of m strategies. If there are n_j players selecting the same strategy j , then each such player i who chooses strategy j incurs a cost of $c_i = f_j(n_j)$ (if the players are mixing, we just have expected costs instead). Note that each function f_j is a nondecreasing function of the number of players selecting the same strategy. The behavior of each function f_j is only relevant on the interval $[0, n]$. In this context, symmetric means that all players have the same set of strategies and the costs of playing a particular strategy depend only on which strategies are being played, not on who plays them (this definition can be found in Ashlagi et al. (2006)). The model we study is a special case of the model in Christodoulou and Koutsoupias (2005), where each player's collection of pure strategies is of size m and each pure strategy set is of size 1 (this corresponds to each player being allowed to choose any of the m links while being restricted to choose only one link per strategy). The social cost function we study is $\sum_{i=1}^n c_i$ (this sum social cost function is also studied in Chien and Sinclair (2009)). These games have been well-studied from the viewpoint of Price of Anarchy (such as in Koutsoupias and Papadimitriou (1999), where the social cost function is the maximum cost over all players rather than the sum of costs). Moreover, these games are a version of the classic model of Roughgarden and Tardos (2000) in which there are a finite number of players.

4.1 Small, Symmetric Singleton 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 Theorem 3.7 for more general games.

Theorem 4.1 *For any symmetric singleton 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 i th strategy be $f_i(x)$. Let the strategies be such that $f_1(n) \leq f_2(n) \leq \dots \leq f_m(n)$. First, any correlated equilibrium has cost at most $nf_1(n)$, since if any player experiences a cost of more than $f_1(n)$, he would choose strategy 1. The game is symmetric with respect to players so there exists a symmetric Nash equilibrium, where the players have the same probability distribution (see Cheng et al. (2004)). Let p_i be the probability of playing the strategy i in the Nash equilibrium. Then the cost of the Nash equilibrium is at least $\sum_i p_i^n n f_i(n)$ (this essentially assumes that $f_i(x) = 0$ for $x < n$; otherwise the cost of the Nash equilibrium is only greater). Using the fact that $f_1(n) \leq f_i(n)$, the Price of Mediation must satisfy:

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

which is obtained by minimizing the sum of p_i^n subject to a sum of 1.

For the other direction, we consider a congestion game where the cost functions for each i are $f_i(n) = 1$ and zero otherwise. Since this is convex, the worst Nash equilibrium will be uniformly at random with expected cost $\frac{n}{m^{n-1}}$ (see Gairing et al. (2004)). 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 randomly chosen strategy and the rest to another random strategy (without replacement) with probability $\frac{mn}{mn+1}$. In this equilibrium, each player pays a cost of $\frac{1}{mn+1}$. However, defecting to a strategy j leads to paying a cost of $\frac{mn}{(mn+1)} \cdot \frac{1}{(m-1)n}$, which is strictly larger. Defecting hurts the player when the mediator randomly chooses the defecting player and sends them to one random strategy while sending the rest of the $n-1$ players to another random strategy (which happens with probability $\frac{mn}{(mn+1)n}$). Then, with probability $\frac{1}{m-1}$, the mediator selects strategy j , so that n players are playing the same strategy. Hence, this is a correlated equilibrium. \square

These results are not quite tight, as there is a gap of nm between the upper and lower bound. For the special case of totally symmetric singleton congestion games in which all strategies have identical cost functions $f(x)$, we have obtained tight bounds for two and three players.

Theorem 4.2 *The PoM for two player totally symmetric games is 1.*

Proof: Consider any correlated equilibrium for a given two player, totally symmetric game $\Gamma_{2,m,f}$. Let p be the probability that the players are placed in the same strategy by the correlated equilibrium, so that $1-p$ is the probability of being alone. Moreover, let p_{ik} be the probability of the outcome in which player 1 is assigned to strategy i and player 2 is assigned to strategy k (note that $\sum_{j=1}^m p_{jj} = p$). We argue that the social cost of the correlated equilibrium is no worse than the social cost of the uniformly at random mixed Nash equilibrium.

Assume towards a contradiction that $p > \frac{1}{m}$. By definition of a correlated equilibrium, since it cannot help player 1 to defect from i to k for any strategies $i \neq k$, we must have:

$$p_{ii}f(2) + f(1) \sum_{j \neq i} p_{ij} \leq p_{ii}f(1) + p_{ik}f(2) + f(1) \sum_{j \neq i,k} p_{ij} \Rightarrow p_{ii}(f(2) - f(1)) \leq p_{ik}(f(2) - f(1)).$$

We may assume $f(2) > f(1)$ (otherwise, the PoM is 1) so that for any strategies i and k , it must be the case that $p_{ii} \leq p_{ik}$. Moreover, since we have a probability distribution, then we must have $1 = \sum_{i=1}^m \sum_{j=1}^m p_{ij} \geq \sum_{i=1}^m m p_{ii} = m p > 1$, a contradiction.

Hence, the social cost of any correlated equilibrium must be at most $2(pf(2) + (1-p)f(1)) = 2(p(f(2) - f(1)) + f(1)) \leq 2(\frac{1}{m}f(2) + \frac{m-1}{m}f(1))$. The result is proved by noting that this is precisely the social cost of the uniformly at random mixed Nash equilibrium. \square

Theorem 4.3 *The PoM for three-player, totally symmetric singleton congestion 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 for each player in the correlated equilibrium is:

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

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 for each player in the correlated equilibrium would be given by:

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

Hence, this gives us the constraint:

$$f(2) \left(\left(\frac{2m-4}{3m-3} \right) p_B - \left(\frac{2}{m-1} \right) p_A \right) + p_C - \left(\frac{1}{3m-3} \right) p_B \leq 0.$$

This implies:

$$p_C \leq \left(\frac{1}{3m-3} \right) p_B - f(2) \left(\frac{2m-4}{3m-3} \right) p_B + f(2) \left(\frac{2}{m-1} \right) p_A.$$

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

$$p_B \left(f(2) \left(\frac{2m-4}{3m-3} \right) - \frac{1}{3m-3} \right) \leq f(2) \left(\frac{2}{m-1} \right) p_A - p_C.$$

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 for each player in the correlated equilibrium is at most $\frac{8}{m-2}$, whereas the cost for each player in the

uniformly random Nash equilibrium is at least $\frac{1}{m^2}$, giving a Price of Mediation 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 for each player in the correlated equilibrium 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 social cost of the worst Nash equilibrium is exactly $\frac{3}{m^2}$. We construct a correlated equilibrium 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 equilibrium has a social cost of $\frac{3}{3m-2}$, for a Price of Mediation of $\Omega(m)$. \square

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 symmetric singleton congestion games.

4.2 Linear, Symmetric Singleton Congestion Games

In this section, we consider symmetric singleton 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 4.4 *The PoM for linear, symmetric singleton congestion games is 1.*

Proof: Let $f_j(x) = a_j x + b_j$ be the cost function for strategy j . We first construct a Nash equilibrium by determining weights \hat{x}_j for each strategy j such that $\sum_j \hat{x}_j = n - 1$. If there are constant functions (i.e., $a_j = 0$), consider the smallest such constant function, and let C denote its value. We first determine values x_j for each strategy j with non-constant cost functions ($a_j > 0$) according to the constraints $f_j(x_j + 1) = C$. There are two cases to consider. In the first case when $\sum_{j:x_j \geq 0} x_j \leq n - 1$, we set $\hat{x}_j = \max(0, x_j)$ (for strategies with $a_j > 0$). For strategies j with $a_j = 0$, we choose a strategy j with $f_j(x) = C$ and assign the remaining weight $\hat{x}_j = n - 1 - \sum_{j:x_j \geq 0} x_j$ to this strategy (and 0 weight to the rest of the strategies with constant cost functions).

In the second case that $\sum_{j:x_j \geq 0} x_j > n - 1$, then there is a unique $0 \leq \mu < C$ such that $f_j(x_j + 1) = \mu$ for each j with $a_j > 0$ and $\sum_{j:x_j \geq 0} x_j = n - 1$. To see this, note that when $\mu = 0$ we have $\sum_{j:x_j \geq 0} x_j = 0$ and when $\mu = C$ we have $\sum_{j:x_j \geq 0} x_j > n - 1$. For strategies with $a_j > 0$, as μ increases, each x_j must increase as well (similarly, as μ decreases, so does each such x_j). For the actual weights, we set $\hat{x}_j = \max(0, x_j)$ for each strategy j with $a_j > 0$ and $\hat{x}_j = 0$ for j with $a_j = 0$.

Let each player be assigned to a strategy j with probability $\hat{x}_j / (n - 1)$. We claim this is a Nash equilibrium. We fix a player i and let λ_j be the number of other players selecting strategy j . In the first case, the expected cost of each strategy in the support of each player (when choosing the strategy) is the same, namely $\mu = C$. There is potentially one constant cost function in the support with cost C . Moreover, all other strategies with constant cost functions (which are not in the support) have cost at least C . For strategies in the support with $a_j > 0$, each player i choosing strategy j has an expected cost $\mathbb{E}[f_j(\lambda_j + 1)] = f_j(\mathbb{E}[\lambda_j] + 1) = C$. For each strategy with $a_j > 0$ and $\hat{x}_j = 0$ (and hence it is not in the support), it must be the case that $f_j(x_j + 1) = C$ implies $x_j \leq 0$, and hence $C \leq a_j + b_j$. Note that $a_j + b_j$ is the cost of having one player on strategy j , so it cannot help player i to defect since this is already larger than C .

The second case is very similar. There are no strategies with constant cost functions in the support for each player, but note that the expected cost of every strategy j in the support of each player (when choosing the strategy) is $f_j(\mathbb{E}[\lambda_j] + 1) = \mu < C$, so a player would never defect to a strategy j with $a_j = 0$. Moreover, every strategy with $a_j > 0$ and $\hat{x}_j = 0$ must again satisfy $\mu \leq a_j + b_j$, and hence it can only hurt the player to defect.

Hence, in both cases we have a Nash equilibrium where every player in expectation pays μ (in the first case, $\mu = C$, and in the second case, $0 \leq \mu < C$). Thus, the social cost of the Nash equilibrium is μn .

We now consider correlated equilibria. 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 $\sum_j \delta_j = n - 1$ so by linearity of expectations it follows that $\sum_j \mathbb{E}[\delta_j] = n - 1$. Since $\sum_j \hat{x}_j = \sum_j \mathbb{E}[\delta_j]$, there is a j such that $\mathbb{E}[\delta_j] \leq \hat{x}_j$ and $\hat{x}_j > 0$. Then by the linearity and monotonicity of the functions it follows that:

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

Of course, player i pays $\mathbb{E}[f_{j'}(\delta_{j'} + 1)]$ when following the mediator's suggestion, 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 $\mathbb{E}[f_{j'}(\delta_{j'} + 1)] \leq \mathbb{E}[f_j(\delta_j + 1)] \leq \mu$. 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 $\text{PoM} \geq 1$, now it follows that $\text{PoM} = 1$. \square

4.3 Concave, Symmetric Singleton 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 4.5 *There exists a concave, symmetric singleton 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 an arbitrary 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, i.e., $\mathbb{E}[N_{ij}] > 1$ for some j in i 's support then there must exist a strategy with $\mathbb{E}[N_{ij}] \leq 1$ (since $\sum_{j=1}^m \mathbb{E}[N_{ij}] = m$) and defecting to this strategy will give player i an expected cost of at most two (thus it is not a Nash equilibrium). 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 $\mathbb{E}[N_{ij}] = 1$ and $\mathbb{E}[f(N_{ij} + 1)] = 2$. From these equations we can conclude that $\mathbb{P}[N_{ij} \geq 3] = 0$. To see this, we observe that $\mathbb{E}[f(N_{ij} + 1)] = 2$ implies $\mathbb{P}[N_{ij} = 0] + 2\mathbb{P}[N_{ij} = 1] + 3\mathbb{P}[N_{ij} \geq 2] = 2$, so that $\mathbb{P}[N_{ij} = 1] + 2\mathbb{P}[N_{ij} \geq 2] = 1$. Hence, we have $\mathbb{E}[N_{ij}] = \sum_{k=0}^m k\mathbb{P}[N_{ij} = k] = 1 + \sum_{k=3}^m (k -$

2) $\mathbb{P}[N_{ij} = k]$. However, since $\mathbb{E}[N_{ij}] = 1$, it must be that $\mathbb{P}[N_{ij} = k] = 0$ for each $k \geq 3$. So for every strategy j , there are exactly *two* players who select this strategy with non-zero probability (since assignments are independent in a Nash equilibrium, if there were more than two players then for some j we would have $\mathbb{P}[N_{ij} \geq 3] > 0$). However, for each of these players that fact that $\mathbb{E}[N_{ij}] = 1$ implies that the other player selects this strategy with probability one. This implies that for i other than these two players, $\mathbb{E}[N_{ij}] = 2$ which is a contradiction. \square

Theorem 4.6 *The PoM of symmetric singleton 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, in a manner similar to Theorem 4.4, we can find weights \hat{x}_j and μ such that $\mu = f_j(\hat{x}_j + 1)$ for all j where $\hat{x}_j > 0$, $\mu \leq f_j(\hat{x}_j + 1)$ for all j where $\hat{x}_j = 0$, and $\sum_j \hat{x}_j = n - 1$. In particular, we compute x_j satisfying $\mu = f_j(x_j + 1)$ for each j . We first handle strategies for which $\mu = f_j(x_j + 1)$ has at most one solution. If $\mu = f_j(x_j + 1)$ has a unique solution where $x_j > 0$, then we simply set $\hat{x}_j = x_j$. If this has no solution for a strategy j , say due to discontinuity, then we just set $\hat{x}_j = 0$. If $\mu = f_j(x_j + 1)$ implies $x_j \leq 0$, then we also set $\hat{x}_j = 0$. We start this process with $\mu = 0$ and keep raising μ until either $\sum_j \hat{x}_j = n - 1$ or we reach a strategy j with multiple solutions to the equation $\mu = f_j(x_j + 1)$, whichever happens first (\hat{x}_j can only increase as we raise μ for each strategy j). Note that multiple solutions means we have hit a flat portion of a strategy j (once a concave, nondecreasing function becomes constant, it must remain constant). If $\sum_j \hat{x}_j = n - 1$ happens first, the process is done. Otherwise, if we find a strategy j' with multiple solutions for the particular value of μ first, so that $\sum_j \hat{x}_j < n - 1$, we simply end the process by increasing $\hat{x}_{j'}$ until $\sum_j \hat{x}_j = n - 1$.

To see why this works, clearly we must have $\sum_j \hat{x}_j = n - 1$. Now, consider any strategy j with positive weight, so that $\hat{x}_j > 0$. Clearly, we have $\hat{x}_j > 0$ if $\mu = f_j(x_j + 1)$ has a unique solution where $x_j > 0$, in which case we set $\hat{x}_j = x_j$ so that $f_j(\hat{x}_j + 1) = \mu$. The only other way we can have $\hat{x}_j > 0$ is if we hit a strategy j with multiple solutions (for instance, a constant function, or a function which becomes constant). In this case, we also have $f_j(\hat{x}_j + 1) = \mu$ for any j with $\hat{x}_j > 0$. On the other hand, consider any strategy j with $\hat{x}_j = 0$. We wish to argue that $f_j(\hat{x}_j + 1) = f_j(1) \geq \mu$. We set $\hat{x}_j = 0$ if $\mu = f_j(x_j + 1)$ has no solution, either because f_j lies entirely above μ (in which case $f_j(1) \geq \mu$) or it is discontinuous, or if $\mu = f_j(x_j + 1)$ has the unique solution where $x_j \leq 0$ (in which case we must also have $f(1) \geq \mu$ since f_j is nondecreasing). The only place a concave, nondecreasing function can be discontinuous is at the first point in its domain, and hence again we must have $f_j(x) \geq \mu$ for any $x > 0$. It is possible that we hit a constant function for the particular value of μ , but again this constant function must be at least μ .

Now, 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 $\mathbb{E}[f_{j'}(\delta_{j'} + 1)] \leq \mathbb{E}[f_j(\delta_j + 1)]$ 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 $\mathbb{E}[\delta_j] \leq \hat{x}_j$ for $\hat{x}_j > 0$ by linearity of expectations. For this j , the expected cost is $\mathbb{E}[f_j(\delta_j + 1)] \leq f_j(\mathbb{E}[\delta_j] + 1) \leq f_j(\hat{x}_j + 1) = \mu$ 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 nondecreasing. 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 equilibrium 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 at some j 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 $f_j(y_j) \geq f_j(\frac{1}{2}(\hat{x}_j + 1)) \geq \frac{1}{2}f_j(\hat{x}_j + 1) \geq \frac{1}{2}\mu$.

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 equilibrium. Thus the Price of Mediation is at most the ratio of the worst correlated equilibrium's social cost (at most $n\mu$) to the worst Nash equilibrium (at least $\frac{1}{4}n\mu$), completing the proof. \square

We prove that for concave, totally symmetric singleton congestion games (i.e., all strategies have the same concave increasing cost function $f(x)$) that the Price of Mediation is at most 2 instead of 4. Note that the example where the PoM was larger than one in Theorem 4.5 was in fact a totally symmetric game.

Theorem 4.7 *The PoM for the class of concave, totally symmetric singleton congestion games is bounded by 2.*

Proof: Consider any concave, totally symmetric singleton congestion game. Let N_{ij} denote the number of players other than i who select strategy j . Since $\sum_j N_{ij} = n - 1$, for any i , there must be some strategy j where $\mathbb{E}[N_{ij}] \leq \frac{n-1}{m}$. It follows from concavity and Jensen's inequality that for this strategy, we must have $\mathbb{E}[f(N_{ij} + 1)] \leq f(\mathbb{E}[N_{ij} + 1]) \leq f(\frac{n-1}{m} + 1)$. If player i has larger expected cost than this, he can simply defect (and we would not have a correlated equilibrium) 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 players choosing strategies uniformly at random is a Nash equilibrium): mq players place 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 $r(q + 1)f(q + 1) + (m - r)qf(q)$.

Note that we may assume $f(x) > 0$ for all $x > 0$. Clearly, if $f(0) > 0$, this must hold (since f is nondecreasing). On the other hand, if $f(0) = 0$ and there exists $x > 0$ such that $f(x) = 0$, then f must be the zero function (otherwise, this would contradict the fact that f is nondecreasing and concave). We observe that if f is the zero function, then every correlated equilibrium and every Nash equilibrium has social cost 0, and hence the Price of Mediation is 1 (by definition).

We first mention a useful inequality: if $x \geq y$, then $\frac{f(x)}{f(y)} \leq \frac{x}{y}$. 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 $f(0)$). We first bound the PoM in the case when $q = 0$ (in this case, we have $r > 0$ and $n = r$). For $q = 0$, the PoM is at most $\frac{nf(1+\frac{n-1}{m})}{rf(1)} \leq 1 + \frac{n-1}{m} = 1 + \frac{r-1}{m} \leq 2$, so we are done (by our assumption, $f(1) > 0$). For the case when $q > 0$, we may bound the PoM from above as follows.

$$\begin{aligned} \frac{nf(1+\frac{n-1}{m})}{r(q+1)f(q+1) + (m-r)qf(q)} &\leq \frac{n(n+m-1)}{mr(q+1)^2 + m(m-r)q^2} \\ &= \frac{(mq+r)(mq+r+m-1)}{m^2q^2 + 2mrq + mr} \\ &\leq \frac{m^2q^2 + 2mrq + m^2q + r^2 + mr}{m^2q^2 + 2mrq + mr} \\ &\leq 1 + \frac{m^2q + r^2}{m^2q^2 + 2mrq + mr} \\ &\leq 2. \end{aligned}$$

The first inequality follows from the inequality $\frac{f(x)}{f(y)} \leq \frac{x}{y}$. If $r = 0$, we need only apply the inequality once with $x = 1 + \frac{n-1}{m}$ and $y = q$. If $r > 0$, we apply the inequality twice, once as before and again with $x = 1 + \frac{n-1}{m}$ and $y = q + 1$. Each time we apply the inequality, we have $y > 0$ and $f(y) > 0$. The rest is algebra, along with applying $n = mq + r$ and observing that $r < m$. \square

4.4 Convex, Symmetric Singleton Congestion Games

For arbitrary convex functions, we can have a Price of Mediation as large as in Theorem 4.1 (in fact the functions in that theorem are convex), namely $\Omega(\frac{1}{n}m^{n-2})$. 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 4.8 *For a symmetric singleton congestion game where all cost functions are polynomials with non-negative coefficients and degree at most s , the PoM is $O((2s)^s)$.*

Proof: We will make repeated use of the following property of polynomials with non-negative coefficients and degree at most s . Consider any $x \geq 0$ and any $\alpha \geq 1$. We must have the property that $f(\alpha x) \leq \alpha^s f(x)$. 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 to that strategy, as in Theorem 4.6. At the end of this process, let x_i be the number of players assigned to strategy i . It is clear that $\sum_i x_i = n$. Let $\mu = \max_i f_i(x_i)$. We will claim that every player in this Nash equilibrium pays at least $\mu 2^{-s}$. If not, then there exists some i such that $x_i > 0$ and $f_i(x_i) < \mu 2^{-s}$. Now using the key property of polynomial functions, we have:

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

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 bounding the maximum ratio of the cost of a correlated equilibrium to the cost of this Nash equilibrium will be sufficient for bounding 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). To see this, note that by definition of a symmetric singleton congestion game, the costs depend only on which strategies are being played, not who is playing them. In particular, the correlated equilibrium will select one of many possible outcomes for the game with probability ρ^j ; each outcome j is defined by an assignment of some number of players to each strategy i , denoted by n_i^j . If $t(k, j)$ denotes the strategy suggested to player k in outcome j , then the expected cost to player k in this correlated equilibrium is given by $C_k = \sum_j \rho^j f_{t(k,j)}(n_{t(k,j)}^j)$. Hence, the social cost of this correlated equilibrium is $\sum_{k=1}^n C_k = \sum_j \rho^j \sum_{k=1}^n f_{t(k,j)}(n_{t(k,j)}^j) = \sum_j \rho^j \sum_{i=1}^m n_i^j f_i(n_i^j)$. On the other hand, if we consider the new correlated equilibrium in which players are shuffled uniformly at random, then the new cost C'_k is the same for each player, namely $\sum_j \rho^j \sum_{i=1}^m \frac{n_i^j}{n} f_i(n_i^j)$. Note that $\frac{n_i^j}{n}$ is the probability that a particular player is assigned to strategy i in outcome j . Hence, we have:

$$\sum_{k=1}^n C'_k = \sum_j \rho^j \sum_{i=1}^m n_i^j f_i(n_i^j) = \sum_{k=1}^n C_k.$$

These equations imply that any correlated equilibrium can be turned into another correlated equilibrium where players are shuffled uniformly at random and the social cost is unchanged (hence, we may assume without loss of generality that players are shuffled at random in the worst correlated equilibrium).

Suppose a player defects by switching to a strategy i that minimizes $\sum_j \rho^j f_i(n_i^j + 1)$ (from any outcome). Indeed, if i^* minimizes this quantity, then because we have a correlated equilibrium, the player cannot improve by defecting to i^* : $\sum_j \rho^j \sum_{i=1}^m \frac{n_i^j}{n} f_i(n_i^j) \leq \sum_j \rho^j f_{i^*}(n_{i^*}^j + 1)$. We can bound each player's expected cost C'_k as follows:

$$C'_k \leq \sum_j \rho^j f_{i^*}(n_{i^*}^j + 1) \leq \sum_j \rho^j \sum_i \frac{x_i}{n} f_i(n_i^j + 1),$$

where the x_i are determined by the pure Nash equilibrium (recall $\sum_i x_i = n$).

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

$$f_i(n_i^j + 1) \leq f_i(sx_i) \leq s^s f_i(x_i) \leq s^s \mu.$$

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

$$f_i(n_i^j + 1) \leq \left(\frac{n_i^j + 1}{n_i^j} \right)^s f_i(n_i^j) \leq e f_i(n_i^j).$$

Summing these two equations gives us:

$$\begin{aligned} C'_k &\leq \sum_j \rho^j \sum_i \left(\frac{x_i}{n} s^s \mu + \frac{e}{ns} n_i^j f_i(n_i^j) \right), \\ C'_k &\leq s^s \mu + \frac{e}{s} C'_k. \end{aligned}$$

Solving the resulting inequality gives us, for $s \geq 3$:

$$C'_k \leq \mu s^{s+1} \frac{1}{s-e}.$$

Combining this with the bound on the cost of the pure Nash equilibrium gives the desired result. Note that the proof holds for polynomials of degree at most s . Hence, if the cost functions all actually have degree $d < 3$ (for instance), then since polynomials of degree $d < 3$ have degree at most 3 (say), we can simply take $s = 3$ and the proof goes through. This increases our bound on the Price of Mediation, but since s is small, the Price of Mediation is just a large constant, so the claim follows immediately. \square

Theorem 4.9 *There is a symmetric singleton congestion game where costs are polynomials of degree s , such that the Price of Mediation grows exponentially in s .*

Proof: We consider a game with m strategies, each of which has identical cost functions $f(x) = x^s$, along with $n = cm$ players, where c is some constant (e.g., we may choose $c = 1$). As shown in Gairing et al. (2004), the Nash equilibrium in which each player chooses each strategy uniformly at random is the worst Nash equilibrium. Hence, we first compute the expected social cost of the uniformly at random Nash equilibrium. Under the binomial distribution $X \sim \text{Bin}(n, 1/m)$, the expected social cost is given by $m\mathbb{E}[Xf(X)] = m\mathbb{E}[X^{s+1}]$.

We first show that the factorial moment satisfies $\mathbb{E}\left[\binom{X}{s+1}\right] = \binom{n}{s+1}m^{-(s+1)}$, and then relate it to the moment $\mathbb{E}[X^{s+1}]$. In order to make the following notation lighter let us introduce $M_\ell := \mathbb{E}\left[\binom{X}{\ell}\right]$ for $\ell \geq 0$ and $B_k := \binom{n}{k}1/m^k(1-1/m)^{n-k}$ for $k \geq 0$. Now consider the generating function of a complex variable z

$$\begin{aligned} \sum_{\ell \geq 0} M_\ell z^\ell &= \sum_{\ell \geq 0} \mathbb{E}\left[\binom{X}{\ell}\right] z^\ell = \sum_{\ell \geq 0} \sum_{k \geq 0} \binom{k}{\ell} B_k z^\ell = \sum_{k \geq 0} B_k \sum_{\ell \geq 0} \binom{k}{\ell} z^\ell = \sum_{k=0}^n B_k \sum_{\ell=0}^k \binom{k}{\ell} z^\ell \\ &= \sum_{k=0}^n B_k (1+z)^k = \left(\frac{1+z}{m} + \frac{m-1}{m}\right)^n = \left(1 + \frac{z}{m}\right)^n = \sum_{\ell=0}^n \binom{n}{\ell} \frac{z^\ell}{m^\ell}, \end{aligned}$$

which gives $\mathbb{E}\left[\binom{X}{s+1}\right] = \binom{n}{s+1}m^{-(s+1)}$. From the fact that $X^{s+1} \leq (s+1)^{s+1}\binom{X}{s+1}$ we obtain $\mathbb{E}[X^{s+1}] \leq (s+1)^{s+1}\mathbb{E}\left[\binom{X}{s+1}\right] = \left(\frac{s+1}{m}\right)^{s+1} \binom{n}{s+1}$. Hence, the expected social cost of the worst Nash equilibrium is at most $m\left(\frac{s+1}{m}\right)^{s+1} \binom{n}{s+1}$. To simplify this expression, we use the sharp Stirling's formula $s! = \sqrt{2\pi s} \left(\frac{s}{e}\right)^s e^{\alpha_s}$, where $\frac{1}{12s+1} < \alpha_s < \frac{1}{12s}$, see Robbins (1955). We also have $(1 - \frac{1}{n})(1 - \frac{2}{n}) \cdots (1 - \frac{s}{n}) = 1 - O(s^2/n)$ for $s = o(\sqrt{n})$. Hence, the expected social cost of the worst Nash equilibrium is at most

$$\frac{n(ce)^{s+1} \left(1 - O\left(\frac{s^2}{n}\right)\right)}{c\sqrt{2\pi(s+1)}} \leq \frac{n(ce)^{s+1}}{c\sqrt{2\pi(s+1)}}.$$

Now we will construct a correlated equilibrium. Suppose that for some $a \in \mathbb{N}$, we allocate n/a players to each of a strategies (we assume $n/a \in \mathbb{N}$ without loss of generality). Since these strategies are selected

at random, each strategy looks the same to every player. Hence, in order for this to produce a correlated equilibrium, we need to ensure that for each player, defecting to any other strategy produces an equal or greater expected cost:

$$\left(\frac{n}{a}\right)^s \leq \frac{a-1}{m-1} \left(\frac{n}{a} + 1\right)^s + \frac{m-a}{m-1}.$$

Thus, it is sufficient to have $\left(\frac{n}{a}\right)^s \leq \frac{a-1}{m-1} \left(\frac{n}{a} + 1\right)^s$. In particular, the inequality is satisfied if we choose a such that

$$\frac{1}{a} \leq \frac{1}{2m} \left(1 + \frac{a}{n}\right)^s.$$

If we select a such that $\frac{2m}{a} = 1 + \frac{as}{n}$, which satisfies this constraint since $s \geq 1$, we will have a correlated equilibrium. Solving the quadratic equation here yields

$$\frac{1}{a} = \frac{1}{4m} \left(1 + \sqrt{1 + 8s/c}\right).$$

If we choose $a = \lceil \frac{4m}{1 + \sqrt{1 + 8s/c}} \rceil$, no player will have an incentive to defect and hence we have a correlated equilibrium. This correlated equilibrium gives us a cost of $(n/a)^s$ for each player, which yields a social cost of at least

$$n \left(\left(\frac{1 + \sqrt{1 + 8s/c}}{8m} \right) n \right)^s = n \left(c \left(\frac{1 + \sqrt{1 + 8s/c}}{8} \right) \right)^s.$$

Dividing this by the bound on the Nash equilibrium gives us a Price of Mediation of at least

$$\left(c \left(\frac{1 + \sqrt{1 + 8s/c}}{8} \right) \right)^s \cdot \frac{c\sqrt{2\pi(s+1)}}{(ce)^{s+1}} \geq \frac{\sqrt{2\pi(s+1)}}{e} \left(\frac{s}{8c\sqrt{e}} \right)^{s/2}.$$

□

Acknowledgement

We would gratefully like to thank Philippe Jacquet and Phil Whiting for useful discussions. We also thank anonymous reviewers for their thorough comments, which helped us to improve this work.

References

- I. Ashlagi, D. Monderer, and M. Tennenholtz. On the value of correlation. In *UAI-05: Proceedings of the 21th Annual Conference on Uncertainty in Artificial Intelligence*, pages 34–41. AUAI Press, 2005.
- I. Ashlagi, D. Monderer, and M. Tennenholtz. Resource selection games with unknown number of players. In *Proceedings of the fifth International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 819–825. ACM, 2006.
- I. Ashlagi, D. Monderer, and M. Tennenholtz. On the value of correlation. *J. Artif. Intell. Res. (JAIR)*, 33:575–613, 2008.
- R. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1:67–96, 1974.

- R. 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>.
- X. Chen and X. Deng. Settling the complexity of 2-player nash equilibrium. *Electronic Colloquium on Computational Complexity*, 2005. ISSN 1433-8092.
- 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.
- S. Chien and A. Sinclair. Strong and pareto price of anarchy in congestion games. In *ICALP '09: Proceedings of the thirty-sixth annual International Colloquium on Automata, Languages and Programming*, pages 279–291, 2009.
- G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *STOC '05: Proceedings of the thirty-seventh annual ACM Symposium on Theory of Computing*, pages 67–73, 2005.
- 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. ISBN 1-59593-134-1. doi: <http://doi.acm.org/10.1145/1132516.1132527>.
- D. Foster. Calibrated learning and correlated equilibria. *Games95 Conference*, 1995.
- M. Gairing, T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. Nash equilibria in discrete routing games with convex latency functions. In *ICALP '04: Proceedings of the thirty-first annual International Colloquium on Automata, Languages and Programming*, pages 645–657, 2004.
- 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.
- E. Koutsoupias and C. Papadimitriou. Worst-case equilibria? In *STACS*, 1999.
- J. Nash. Equilibrium points in n -person games. *Proceedings of the National Academy of Sciences USA*, 36:48–49, 1950.
- 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. ISBN 1-58113-960-8. doi: <http://doi.acm.org/10.1145/1060590.1060598>.
- 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. ISBN 0-89871-585-7.
- R. Peeters and J. Potters. On the structure of the set of correlated equilibria in two-by-two bimatrix games. Discussion Paper 45, Tilburg University, Center for Economic Research, 1999. Available at <http://ideas.repec.org/p/dgr/kubcen/199945.html>.
- H. Robbins. A remark on stirling's formula. *The American Mathematical Monthly*, 62(1):26–29, 1955.
- 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.
- T. Roughgarden and E. Tardos. How bad is selfish routing?. In *IEEE Symposium on Foundations of Computer Science*, pages 93–102, 2000.