

The Limits of Smoothness: A Primal-Dual Framework for Price of Anarchy Bounds

Uri Nadav* and Tim Roughgarden**

Abstract. We show a formal duality between certain equilibrium concepts, including the correlated and coarse correlated equilibrium, and analysis frameworks for proving bounds on the price of anarchy for such concepts. Our first application of this duality is a characterization of the set of distributions over game outcomes to which “smoothness bounds” always apply. This set is a natural and strict generalization of the coarse correlated equilibria of the game. Second, we derive a refined definition of smoothness that is specifically tailored for coarse correlated equilibria and can be used to give improved POA bounds for such equilibria.

1 Introduction

A rigorous way to argue that a system with self-interested participants has good performance is to prove that every “plausible outcome” of the system has objective function value close to that of an optimal outcome. For example, one could model a system as a one-shot game, identify “plausible outcomes” with the pure-strategy Nash equilibria (PNE) — outcomes in which each player deterministically picks one strategy so that it has no incentive to unilaterally deviate from it — and prove a relative approximation bound for the PNE of the game.

Such *price of anarchy (POA)* bounds become increasingly robust and compelling as one increases the set of “plausible outcomes”. For example, a POA bound that applies only to the pure-strategy Nash equilibria of a game presumes that the system reaches such a state. This can be a bold assumption, for example in contexts where it is computationally difficult to compute a PNE (see e.g. [7]). A POA bound that applies more generally to “easily learned” outcomes, such as the correlated equilibria [1] or coarse correlated equilibria [8] of a game, presumes far less about the game’s participants [2, 3]. Of course, worst-case approximation bounds typically degrade as the assumptions about play are weakened — for example, the expected performance of the worst coarse correlated equilibrium of a game is typically worse than that of the worst pure-strategy Nash equilibrium.

* Department of Computer Science, Stanford University, 470 Gates Building, 353 Serra Mall, Stanford, CA 94305. Supported by the ONR Young Investigator and PECASE Awards of the second author. Email urinadav@cs.stanford.edu.

** Department of Computer Science, Stanford University, 462 Gates Building, 353 Serra Mall, Stanford, CA 94305. Supported in part by NSF CAREER Award CCF-0448664, an ONR Young Investigator Award, an ONR PECASE Award, and an AFOSR MURI grant. Email: tim@cs.stanford.edu.

This paper shows a precise duality between certain equilibrium concepts, including correlated and coarse correlated equilibria, and analysis frameworks for proving POA bounds for such concepts. This duality makes formal the intuitive trade-off between the plausibility of the rationality assumptions imposed on the game participants and the quality of the corresponding worst-case approximation bound. We offer two applications.

1. Roughgarden [11] showed that every POA bound proved using a “smoothness argument” (see Definition 1) — the most frequently employed method for establishing POA bounds (e.g. [5, 6, 9, 10, 12]) — applies automatically to (at least) all CCE of the game. A basic problem is to characterize the distributions over outcomes to which smoothness bounds always apply. We solve this problem (Theorem 1) and show that the answer is a generalization of CCE in which the *average* regret of players is non-positive, as opposed to the CCE condition that *every* player has non-positive regret (see Definition 2).
2. Applying the duality result in the opposite direction yields analysis frameworks that are guaranteed to be tight for the corresponding equilibrium concepts. We illustrate this idea with the set of CCE, where the corresponding multi-parameter analysis framework refines the simpler two-parameter smoothness paradigm in [11]. This more flexible analysis framework is, by definition, specifically tailored for CCE and can be used to give improved POA bounds for such equilibria.

2 The Primal-Dual Framework

Section 2.1 reviews standard definitions of cost-minimization games, equilibrium concepts, and the price of anarchy. Section 2.2 presents our first contribution and shows that, for every equilibrium concept that can be expressed as the probability distributions over outcomes that are solutions to a set of homogeneous inequalities, there is a corresponding analysis framework that is guaranteed to prove tight bounds on the price of anarchy for that concept. Our second contribution, described in Section 2.3, is an application of this framework: POA bounds proved using the “smoothness paradigm” introduced in [11] apply precisely to a generalization of coarse correlated equilibria that we call “average coarse correlated equilibria”. Section 2.4 demonstrates how a sharper analysis method tailored specifically for coarse correlated equilibria, which follows directly from our primal-dual framework, can be used to prove bounds superior to those that follow from the standard smoothness paradigm.

2.1 Preliminaries

Cost-minimization games. We denote a cost-minimization game by a tuple $\Gamma = (N, \{S_i\}_{i \in N}, \{C_i\}_{i \in N})$, where $N = \{1, \dots, n\}$ is the set of n players, S_i is the set of actions of player i , and $c_i : S \mapsto \mathbb{R}^{++}$ is player’s i positive cost function, where $S = S_1 \times S_2 \times \dots \times S_n$ is the joint action set.¹ We use $\Delta(S)$ to denote the

¹ Our results can be reworked without difficulty for payoff-maximization games.

set of probability distributions over S and \mathbf{s}_{-i} to denote the strategies played in \mathbf{s} by the players other than i .

Equilibrium concepts and the price of anarchy. In this paper, we consider equilibrium concepts that can be described as subsets of $\Delta(S)$. In particular, recall that a *correlated equilibrium (CE)* is a joint probability distribution σ over outcomes of Γ with the property that $\mathbf{E}_{\mathbf{s} \sim \sigma}[C_i(\mathbf{s})|\mathbf{s}_i] \leq \mathbf{E}_{\mathbf{s} \sim \sigma}[C_i(\mathbf{s}'_i, \mathbf{s}_{-i})|\mathbf{s}_i]$ for every i and $\mathbf{s}_i, \mathbf{s}'_i \in S_i$. Thus a distribution σ over outcomes is a CE if the following holds for a random sample $\mathbf{s} \sim \sigma$: for each player i and “recommended strategy” s_i , the player minimizes its expected cost, conditioned on the recommendation s_i and assuming that other players play according to \mathbf{s}_{-i} , by playing s_i . CE are also the limits of sequences of repeated play in which each player has vanishing per-step *swap* or *internal regret* (see [4]). The mixed Nash equilibria of a game are precisely the CE that are also product distributions.

A *coarse correlated equilibrium (CCE)* is a joint probability distribution σ over outcomes of Γ with the property that $\mathbf{E}_{\mathbf{s} \sim \sigma}[C_i(\mathbf{s})] \leq \mathbf{E}_{\mathbf{s} \sim \sigma}[C_i(\mathbf{s}'_i, \mathbf{s}_{-i})]$ for every i and $\mathbf{s}'_i \in S_i$. These equilibrium constraints consider only player deviations that are independent of the recommendation s_i , so every CE is also a CCE (and, generally, the converse fails). CCE are also the limits of sequences of repeated play in which each player has vanishing per-step *external regret* (see [4]).

We assume that the objective function is to minimize the total cost $C(\mathbf{s}) = \sum_{i \in N} C_i(\mathbf{s})$, and use \mathbf{s}^* to denote an optimal outcome. The *price of anarchy (POA)* of a game for an equilibrium concept $\text{EQ} \subseteq \Delta(S)$ is the ratio between the expected total cost of the worst (i.e., highest-cost) equilibrium $\sigma \in \text{EQ}$ and the social cost of \mathbf{s}^* .

2.2 A Primal-Dual Framework for POA Bounds

This section describes our primal-dual framework, which formalizes a duality between equilibrium concepts that can be represented as solutions of homogeneous inequalities and analysis methods that are necessary and sufficient to prove tight bounds on the POA for such concepts.

Fix a game Γ , and an equilibrium concept EQ that can be written as $\text{EQ} = \{\sigma \in \Delta(S) : A\sigma \leq 0\}$, where $A \in \mathbb{R}^{|S| \times m}$ is a matrix that can depend on players’ cost functions in Γ . For example, the equilibrium concepts CE and CCE can be described in this way:

Example 1 (Correlated Equilibria). We can express the CE of a cost-minimization game as the probability distributions over outcomes that satisfy

$$\text{CE} = \left\{ \sigma : \sum_{\mathbf{s} : \mathbf{s}_i = s_i} \sigma_{\mathbf{s}} (C(s'_i, \mathbf{s}_{-i}) - C(\mathbf{s})) \leq 0, \text{ for every } i \in N, \text{ and } s_i, s'_i \in S_i, \sigma_{\mathbf{s}} \geq 0 \right\}.$$

Example 2 (Coarse Correlated Equilibrium). We can express the CCE of a cost-minimization game as the probability distributions that satisfy

$$\text{CCE} = \left\{ \sigma : \sum_{\mathbf{s}} \sigma_{\mathbf{s}} (C(s'_i, \mathbf{s}_{-i}) - C(\mathbf{s})) \leq 0, \text{ for every } i \in N, \text{ and } s'_i \in S_i, \sigma_{\mathbf{s}} \geq 0 \right\}.$$

A third example will arise naturally in Section 2.3.

We now develop our simple primal-dual framework. We can formally write the POA of a game Γ and an equilibrium concept EQ as

$$\text{POA}_{\text{EQ}}(\Gamma) = \sup_{\sigma \in \text{EQ}} \left\{ \frac{\mathbf{E}_{\mathbf{s} \sim \sigma} [C(\mathbf{s})]}{C(\mathbf{s}^*)} \right\}.$$

After scaling by $C(\mathbf{s}^*)$, this maximization problem can be expressed as the solution of the following linear program:

$$\begin{aligned} \text{PRIMAL-EQ : Maximize } & \sum_{\mathbf{s} \in S} \sigma_{\mathbf{s}} C(\mathbf{s}) \\ \text{subject to } & \sum_{\mathbf{s} \in S} \sigma_{\mathbf{s}} C(\mathbf{s}^*) = 1 \\ & A\sigma \leq 0, \sigma_{\mathbf{s}} \geq 0 \end{aligned}$$

The dual problem of PRIMAL-EQ is

$$\begin{aligned} \text{DUAL-EQ : Minimize } & p \\ \text{subject to } & C(\mathbf{s}^*)p \cdot \mathbf{1}^n + zA^T \geq \mathbf{0}, \\ & z \geq 0, p \geq 0, \quad p \in \mathbb{R}, z \in \mathbb{R}^m \end{aligned}$$

where $\mathbf{1}^n$ is the n dimensional vector with all entries 1, and m is the number of inequalities in A .

We say that a game is p -bounded for the equilibrium concept EQ if there exists a vector $z \in \mathbb{R}^m$ such that the pair (p, z) is feasible for DUAL-EQ, or simply p -bounded when the equilibrium concept is clear. We refer to z as a *dual certificate* for Γ and EQ.

Strong linear programming duality immediately implies the following.

Proposition 1. *For every cost-minimization game Γ and equilibrium concept EQ representable as the solution of homogeneous inequalities, $\text{POA}_{\text{EQ}}(\Gamma) \leq p$ if and only if Γ is p -bounded for EQ.*

The following example instantiates Proposition 1 for correlated equilibria. The next two sections provide further examples.

Example 3 (Primal-Dual Framework for Correlated Equilibria). For a cost-minimization game Γ , the quantity $\text{POA}_{\text{CE}}(\Gamma)$ is, by definition, the optimal solution to the problem PRIMAL-CE:

$$\begin{aligned} \text{PRIMAL-CE : Maximize } & \sum_{\mathbf{s} \in S} \sigma_{\mathbf{s}} C(\mathbf{s}) / C(\mathbf{s}^*) \\ \text{subject to } & \sum_{\mathbf{s} \in S} \sigma_{\mathbf{s}} = 1 \\ & \sum_{\mathbf{s}: \mathbf{s}_i = a} \sigma_{\mathbf{s}} (C(b, \mathbf{s}_{-i}) - C(\mathbf{s})) \leq 0, \text{ for every } i \in N, \\ & \quad \text{and } a, b \in S_i \\ & \sigma_{\mathbf{s}} \geq 0. \end{aligned}$$

The corresponding DUAL-CE problem is then

$$\begin{aligned} \text{DUAL-CE : Minimize } & p \\ \text{subject to } & pC(\mathbf{s}^*) + \sum_i \sum_{b \in S_i} z_{\mathbf{s}_i, b}^i (C_i(\mathbf{s}) - C_i(b, \mathbf{s}_{-i})) \geq C(\mathbf{s}), \\ & \quad \text{for all } \mathbf{s} \in S \\ & z \geq 0, p \geq 0. \end{aligned}$$

Hence, to prove an upper bound of p on the POA for correlated equilibrium, it suffices to show that the game is p -bounded for CE — that is, to find a dual certificate $z = \{z_{a,b}^i\}_{i \in N, a, b \in S_i}$ so that (p, z) is feasible for DUAL-CE.

2.3 The Limits of (λ, μ) -Smoothness

Roughgarden [11] defined a smooth game as follows.

Definition 1 (Smooth Games). *A cost-minimization game with minimum-cost outcome \mathbf{s}^* is (λ, μ) -smooth if*

$$\sum_{i=1}^k C_i(\mathbf{s}_i^*, \mathbf{s}_{-i}) \leq \lambda \cdot C(\mathbf{s}^*) + \mu \cdot C(\mathbf{s}) \quad (1)$$

for every outcome \mathbf{s} .

One of the main results in [11] is that $\text{POA}_{\text{CCE}}(\Gamma) \leq \lambda/(1 - \mu)$ whenever Γ is (λ, μ) -smooth.² In addition, many known POA bounds — often stated only for pure or mixed Nash equilibria — are or can be recast as smoothness bounds (see [11]), and thus these bounds “extend automatically” to the more general concept of CCE.

This section addresses the basic question of characterizing the distributions over outcomes to which a (λ, μ) -smoothness bound applies. The answer, which we derive via the primal-dual framework in the previous section, turns out to be a strict generalization of CCE that we call an *average coarse correlated equilibrium, with respect to \mathbf{s}^** (ACCE*).

Definition 2 (ACCE*). *For a fixed game and an outcome $r \in S$*

$$\text{ACCE}^r = \{\sigma \in \Delta(S) : E_{\mathbf{s} \sim \sigma}[C(\mathbf{s})] \leq E_{\mathbf{s} \sim \sigma}[\sum_i C_i(r_i, \mathbf{s}_{-i})]\}.$$

When r is the minimum-cost outcome \mathbf{s}^* , we abbreviate $\text{ACCE}^{\mathbf{s}^*}$ by ACCE^* .

Conceptually, there are two differences between a CCE and an ACCE^* . In a CCE, the expected cost incurred by a player is at most that of unconditionally deviating to an any fixed action — i.e., every player has non-positive “regret”. ACCE^* is a more permissive equilibrium concept. First, we measure the regret of a player i by comparing its expected cost only to that incurred under a deviation to s_i^* , rather than to an arbitrary (or best) strategy. Second, in an ACCE^* , some players i can have negative regret with respect to s_i^* as long as the *average* (over players) such regret is non-positive. Unsurprisingly, many games have ACCE^* that are not CCE; the proof of Proposition 2 provides one concrete example.

² In [11] the definition of (λ, μ) -smoothness requires that inequality (1) holds for every pair \mathbf{s}, \mathbf{s}^* outcomes. The weaker requirement stated here still translates, via the same proofs, to an upper bound on the POA for CCE.

The next theorem shows that every (λ, μ) -smoothness argument bounds the worst-case expected cost of *precisely* the set of $ACCE^*$. This characterization has both positive and negative implications. First, even the $ACCE^*$ distributions of a (λ, μ) -smooth game have good expected cost (and not only the CCE, as proved in [11]). Second, conversely, the worst-case $ACCE^*$ constrains the best-possible upper bound that can be proved via a (λ, μ) -smoothness argument.

Theorem 1 (Duality Between (λ, μ) -Smoothness and $ACCE^*$). *For every cost-minimization game Γ , the best smoothness upper bound for Γ equals its POA for the equilibrium concept $ACCE^*$:*

$$\inf \left\{ \frac{\lambda}{1-\mu} : (\lambda, \mu) \text{ s.t. the game } \Gamma \text{ is } (\lambda, \mu)\text{-smooth} \right\} = \text{POA}_{ACCE^*}(\Gamma).$$

Proof. We prove that the (λ, μ) -smoothness requirements are equivalent to the constraints of the DUAL problem for the equilibrium concept $ACCE^*$. We consider the linear fractional problem for obtaining the best (i.e., least) upper bound using (λ, μ) -smoothness:

$$\begin{aligned} \text{(LFP)} : & \text{Minimize } \frac{\lambda}{1-\mu} \\ & \text{subject to } \sum_{i \in N} C_i(\mathbf{s}_i^*, \mathbf{s}_{-i}) \leq \lambda C(\mathbf{s}^*) + \mu C(\mathbf{s}), \text{ for all } \mathbf{s} \in S \\ & \mu < 1. \end{aligned}$$

By rearranging terms in the first inequality of problem (LFP) and dividing through by $1 - \mu > 0$ we obtain

$$\begin{aligned} \text{(LFP2)} : & \text{Minimize } \frac{\lambda}{1-\mu} \\ & \text{subject to } \frac{\lambda}{1-\mu} C(\mathbf{s}^*) + \frac{1}{1-\mu} (C(\mathbf{s}) - \sum_{i \in N} C_i(\mathbf{s}_i^*, \mathbf{s}_{-i})) \geq C(\mathbf{s}) \\ & \text{for all } \mathbf{s} \in S \\ & \mu < 1. \end{aligned}$$

Now, re-writing (LFP2) with a change of variables $p = \frac{\lambda}{1-\mu}$, and $z = \frac{1}{1-\mu}$ gives the following linear program (LP):

$$\begin{aligned} \text{(LP)} : & \text{Minimize } p \\ & \text{subject to } pC(\mathbf{s}^*) + z(C(\mathbf{s}) - \sum_i C_i(\mathbf{s}_i^*, \mathbf{s}_{-i})) \geq C(\mathbf{s}), \text{ for all } \mathbf{s} \in S \\ & z > 0. \end{aligned}$$

The dual problem of (LP) is:

$$\begin{aligned} \text{(D)} : & \text{Maximize } \sum_{\mathbf{s} \in S} \sigma_{\mathbf{s}} C(\mathbf{s}) \\ & \text{subject to } \sum_{\mathbf{s} \in S} \sigma_{\mathbf{s}} C(\mathbf{s}^*) \leq 1 \\ & \sum_{\mathbf{s}} \sigma(\mathbf{s}) (\sum_i C_i(\mathbf{s}_i^*, \mathbf{s}_{-i}) - C(\mathbf{s})) \geq 0 \\ & \sigma_{\mathbf{s}} \geq 0, \text{ for all } \mathbf{s} \in S. \end{aligned}$$

We can replace the first inequality in (D) with an equality since the social cost function is positive by assumption. Then, after scaling by $C(\mathbf{s}^*)$ we get an equiv-

alent linear program that corresponds to the POA for ACCE*:

$$\begin{aligned}
(\text{PRIMAL-ACCE}^*) : & \text{ Maximize } \sum_{\mathbf{s} \in S} \sigma_{\mathbf{s}} \frac{C(\mathbf{s})}{C(\mathbf{s}^*)} \\
& \text{ subject to } \sum_{\mathbf{s} \in S} \sigma_{\mathbf{s}} = 1 \\
& \sum_{\mathbf{s}} \sigma(\mathbf{s}) (\sum_i C(\mathbf{s}_i^*, \mathbf{s}_{-i}) - C(\mathbf{s})) \geq 0 \\
& \sigma_{\mathbf{s}} \geq 0, \text{ for all } \mathbf{s} \in S.
\end{aligned}$$

2.4 Better Dual Certificates Give Better POA Upper Bounds

Theorem 1 shows that the smoothness analysis framework in [11] corresponds precisely to worst-case upper bounds on the set of ACCE*. In this section we assume that the goal is to prove upper bounds on the quantity $\text{POA}_{\text{CCE}}(\Gamma)$, and view the fact that (λ, μ) -smoothness bounds the expected cost of a strictly larger set of outcome distributions as an “accident”. Motivated by this perspective, this section uses the primal-dual framework of Section 2.2 to derive a condition tailored for CCE that is sharper than (λ, μ) -smoothness and that can be used to prove better upper bounds on $\text{POA}_{\text{CCE}}(\Gamma)$.

Let CCE* denote the equilibrium concept where each player’s expected cost is at most that of deviating to its action in \mathbf{s}^* , i.e., $\sigma \in \text{CCE}^*$, if and only if

$$\mathbf{E}_{\mathbf{s} \sim \sigma} [C_i(\mathbf{s})] \leq \mathbf{E}_{\mathbf{s} \sim \sigma} [C_i(\mathbf{s}_i^*, \mathbf{s}_{-i})]$$

for every $i \in N$. Obviously, $\text{CCE} \subseteq \text{CCE}^* \subseteq \text{ACCE}^*$.

Proposition 1 shows that every equilibrium concept that is the solution to homogeneous inequalities, such as CCE*, has a corresponding tight analysis framework. To bound the POA for CCE*, we only need to find a suitable dual certificate. The DUAL-CCE* problem is

$$\begin{aligned}
(\text{DUAL-CCE}^*) : & \text{ Minimize } p \\
& \text{ subject to } pC(\mathbf{s}^*) + \sum_{i \in N} z_i (C_i(\mathbf{s}) - C_i(\mathbf{s}_i^*, \mathbf{s}_{-i})) \geq C(\mathbf{s}), \\
& \hspace{15em} \text{for all } \mathbf{s} \in S \\
& z_i \geq 0, \text{ for all } i \in N.
\end{aligned}$$

Thus, a dual certificate for CCE* is an n -dimensional vector z such that $pC(\mathbf{s}^*) + \sum_{i \in N} z_i (C_i(\mathbf{s}) - C_i(\mathbf{s}_i^*, \mathbf{s}_{-i})) \geq C(\mathbf{s})$, for all $\mathbf{s} \in S$. This is evidently more flexible than the single-parameter dual certificate one is forced to use for ACCE*. Can this flexibility lead to better worst-case upper bounds? The following proposition gives an affirmative answer.

Proposition 2. *There is a game Γ such that $\text{POA}_{\text{CCE}}(\Gamma) < \text{POA}_{\text{ACCE}^*}(\Gamma)$.*

Proof (sketch): Consider a load balancing game with two jobs J_1, J_2 , with weights 2 and 1 respectively, and two machines M_1, M_2 with latency functions

$$\begin{aligned}
\ell_1(1) &= 1, & \ell_1(2) &= 2, & \ell_1(3) &= 3; \\
\ell_2(1) &= 1 + \epsilon, & \ell_2(2) &= 2, & \ell_2(3) &= 4.
\end{aligned}$$

The optimal outcome \mathbf{s}^* assigns J_1 to M_1 and J_2 to M_2 , and has a social cost 3. For small enough ϵ , the best ACCE* dual certificate³ is $z \approx 7/9$ which corresponds to a POA of $16/9 \approx 1.77$. For CCE* a dual certificate $(z_1, z_2) \approx (23/24, 5/12)$ exists, for a better POA bound of $49/30 \approx 1.63$. ■

Remark 1. There are games with an arbitrary gap between $\text{POA}_{\text{CCE}^*}$ and $\text{POA}_{\text{ACCE}^*}$, e.g., by changing the latency function ℓ_2 in the proof of Proposition 2 to $\ell_2(3) = H$, for a large enough H .

Remark 2. In contrast to Proposition 2, in symmetric games — where all players have the same strategy set and each player’s cost depends only on its own strategy and the number of players that choose each strategy — the POA for ACCE* is equal to the POA for CCE*. We omit the easy argument.

References

1. R. J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1):67–96, 1974.
2. A. Blum, E. Even-Dar, and K. Ligett. Routing without regret: On convergence to Nash equilibria of regret-minimizing algorithms in routing games. In *Proceedings of the 25th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 45–52, 2006.
3. A. Blum, M. Hajiaghayi, K. Ligett, and A. Roth. Regret minimization and the price of total anarchy. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC)*, pages 373–382, 2008.
4. A. Blum and Y. Mansour. Learning, regret minimization, and equilibria. In N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 4, pages 79–101. Cambridge University Press, 2007.
5. G. Christodoulou and E. Koutsoupias. On the price of anarchy and stability of correlated equilibria of linear congestion games. In *Proceedings of the 13th Annual European Symposium on Algorithms (ESA)*, volume 3669 of *Lecture Notes in Computer Science*, pages 59–70, 2005.
6. J. R. Correa, A. S. Schulz, and N. E. Stier Moses. On the inefficiency of equilibria in congestion games. *Games and Economic Behavior*, 64(2):457–469, 2008.
7. A. Fabrikant, C. H. Papadimitriou, and K. Talwar. The complexity of pure Nash equilibria. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC)*, pages 604–612, 2004.
8. J. Hannan. Approximation to Bayes risk in repeated plays. In M. Dresher, A. Tucker, and P. Wolfe, editors, *Contributions to the Theory of Games*, volume 3, pages 97–139. Princeton University Press, 1957.
9. T. Harks. Stackelberg strategies and collusion in network games with splittable flow. In *Proceedings of the 6th International Workshop on Approximation and Online Algorithms (WAOA)*, volume 5426 of *Lecture Notes in Computer Science*, pages 133–146, 2008.
10. G. Perakis. The price of anarchy when costs are non-separable and asymmetric. *Mathematics of Operations Research*, 32(3):614–628, 2007.

³ The equivalent, and optimal (λ, μ) pair for this upper bound is $(16/7, -2/7)$.

11. T. Roughgarden. Intrinsic robustness of the price of anarchy. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC)*, pages 513–522, 2009.
12. A. Vetta. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In *Proceedings of the 43rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 416–425, 2002.