Lecture 15—Price of Anarchy

Iosif Sakos

November 13, 2025

Abstract

These notes introduce the price of anarchy (PoA) for routing games as a metric of inefficiency. We first treat the nonatomic model, defining flows, latency functions, and Wardrop equilibria, and the PoA for nonatomic routing games. Pigou's example gives a lower bound of 4/3 on the PoA for affine latency functions, and we show this is tight via a reduction to Pigou-like instances. We then introduce atomic routing games with unit-demand players, define Nash equilibria (NE) and the PoA for atomic routing games, and present an affine example yielding a lower bound of 5/2. Finally, we prove that this bound is tight for the entire class of affine latency functions.

Contents

1 Nonatomic routing games

Recall that a Nash equilibrium (NE) is *a stable state* of a game in which no player can improve their payoff by unilaterally deviating. However, a game may admit multiple NE, raising the question of how to assess *the quality* of these equilibria. The price of anarchy (PoA) addresses this by quantifying the inefficiency of an NE relative to a suitably defined *optimal* outcome.

To make these ideas concrete, we begin our discussion in the setting of a nonatomic routing game, in which the analogue of an NE is the Wardrop equilibrium. Under standard continuity and monotonicity assumptions on the cost functions, Wardrop equilibria exist; moreover, all Wardrop equilibria induce the same aggregate cost, and thus *the* equilibrium quality is unambiguous. Finally, in a nonatomic routing game, no-regret learning dynamics drive play, on average, toward the set of Wardrop equilibria [1], which further motivates using the PoA to measure the inefficiency of outcomes arising from a learning process.

1.1 From normal-form games to nonatomic routing games

Consider an n-player normal-form game $\mathscr{G} = \big(\llbracket n \rrbracket, \{\alpha_1, \alpha_2\}^n, \{u_i\}_{i\in\llbracket n \rrbracket}\big)$ in which each player $i \in \llbracket n \rrbracket$ has two pure strategies, α_1 and α_2, and the payoff functions are:

$$u_i(s) = \begin{cases} -\frac{1}{n}\sum_{j=1}^{n} \mathbf{1}_{s_i=s_j}, & \text{if } s_i = \alpha_1, \\ -1, & \text{if } s_i = \alpha_2, \end{cases} \tag{1}$$

for all $i \in \llbracket n \rrbracket$ and all strategy profiles $s \in \{\alpha_1, \alpha_2\}^n$. In words, if player i chooses α_1, their payoff is the negative of the fraction of players (including player i) who also choose α_1. If player i chooses α_2, their payoff is simply -1.

Let $x \in [0, 1]$ be the fraction of players who choose α_1 in a strategy profile s. Then the payoff of any player choosing α_1 is $-x$, while the payoff of any player choosing α_2 is -1. Since payoffs depend only on the fraction x and not on the identities of the players, this description continues to make sense as n grows. In the limit $n \to \infty$, we obtain a nonatomic routing game with a continuum of players, i.e., a population, each choosing between two routes, α_1 and α_2 with *latencies* (costs) x and 1, respectively. You may think of this limit as a game in which each individual player has *negligible influence* on the overall outcome. More generally, a nonatomic routing game is given by a tuple $\big(\mathcal{G}, \{c_e\}_{e\in\mathcal{E}}\big)$, where $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is a directed graph with a designated source–sink pair $(s, t) \in \mathcal{V} \times \mathcal{V}$, and each arc $e \in \mathcal{E}$ has an associated latency function $c_e \colon \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ that specifies the travel time on e as a function of the fraction of the population (the flow) on e. Throughout, we assume that all latency functions are *continuous*, *nonnegative*, and *nondecreasing*.

As an example, consider the simple routing game depicted in Figure 1, known as Pigou's example.

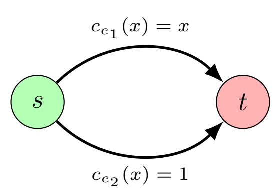

Figure 1: Pigou's example of a nonatomic routing game.

There are two parallel arcs from the source s to the sink t. The latency function on the upper arc e_1 is $c_{e_1}(x) = x$, while the latency function on the lower arc e_2 is $c_{e_2}(x) = 1$. This game is strategically equivalent to the limit of the normal-form game in (1) as $n \to \infty$, in which choosing arc e_1 corresponds to strategy α_1, and choosing arc e_2 corresponds to strategy α_2.

Flows. The analogue of a pure strategy in a nonatomic routing game is a flow. Let $\mathcal{P} \in 2^{\mathcal{E}}$ be the set of all *simple* paths from the source s to the sink t in the graph $\mathcal{G}$. A flow is a $|\mathcal{P}|$-dimensional vector $(f_p)_{p\in\mathcal{P}}$ with $f_p \geq 0$ for all $p \in \mathcal{P}$ and $\sum_{p\in\mathcal{P}} f_p = 1$; it assigns to each path p a value f_p, representing the fraction of the population that chooses path p. We denote the set of all feasible flows by $\mathcal{F}$.

Given a flow $f \in \mathcal{F}$, we may induce a flow on the arcs of the graph. That is, for each arc $e \in \mathcal{E}$, we may compute the fraction of the population that uses arc e as part of their chosen path. Formally, the induced flow is given by

$$f_e = \sum_{p \ni e} f_p, \quad \forall e \in \mathcal{E}. \tag{2}$$

As an example, consider the nonatomic routing game in Figure 2, known as Braess's paradox.

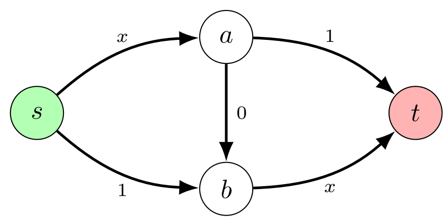

Figure 2: Braess's paradox.

In this game, there are three simple paths from s to t: the upper path $s \to a \to t$, the lower path $s \to b \to t$, and the zigzag path $s \to a \to b \to t$. Suppose that the flow f is such that 50% of the population chooses the upper path $s \to a \to t$, and 50% chooses the lower path $s \to b \to t$. Then the induced flow on each arc is $f_{(s,a)} = f_{(a,t)} = f_{(s,b)} = f_{(b,t)} = 0.5$, while $f_{(a,b)} = 0$. We will return to this example later to see why it is paradoxical.

Path latency and social cost. The latency experienced by players choosing path $p \in \mathcal{P}$ under flow $f \in \mathcal{F}$ is the accumulated latency over all arcs e on path p, i.e.,

$$c_p(f) = \sum_{e \in p} c_e(f_e), \quad \forall p \in \mathcal{P}. \tag{3}$$

In the example above, the latency experienced by players choosing the upper path $s \to a \to t$ is $c_{s\to a\to t}(f) = c_{(s,a)}(0.5) + c_{(a,t)}(0.5) = 0.5 + 1 = 1.5$, while the latency experienced by players choosing the lower path $s \to b \to t$ is $c_{s\to b\to t}(f) = c_{(s,b)}(0.5) + c_{(b,t)}(0.5) = 1 + 0.5 = 1.5$.

Since each infinitesimal player experiences only the latency of their chosen path, the average latency is a natural measure of the overall performance of a flow. We define the social cost $C(f)$ of a flow f to be this average latency. Formally,

$$C(f) = \sum_{p \in \mathcal{P}} f_p \cdot c_p(f) \quad \forall f \in \mathcal{F}. \tag{4}$$

In the example above, the social cost of flow f is $C(f) = 0.5 \cdot 1.5 + 0.5 \cdot 1.5 = 1.5$.

Simplifications. To keep the notation compact, the definition of a nonatomic routing game above assumes a single source–sink pair (s, t) and a total demand (population) of 1 unit of flow to be routed from s to t. The normalization to 1 unit of flow is without loss of generality, since we can always rescale the latency functions accordingly. If there are multiple source–sink pairs, we can view of the game as consisting of multiple populations, each with its own source–sink pair and demand. All results presented in these notes extend in a straightforward way to this more general setting.

1.2 Wardrop equilibria.

A Wardrop equilibrium captures the idea that no player can reduce their latency by *unilaterally* changing their route from s to t. This is the continuum analogue of an NE, in which no infinitesimal player can reduce their latency by unilaterally deviating. Formally, a flow $f^\star \in \mathcal{F}$ is a Wardrop equilibrium if, for all paths $p \in \mathcal{P}$ with positive flow $f^\star_p > 0$, the latency $c_p(f^\star)$ is minimal; i.e.,

$$c_p(f^\star) \le c_{p'}(f^\star), \quad \forall p' \in \mathcal{P}. \tag{5}$$

In other words, all the paths that are used in a Wardrop equilibrium have the *minimal* possible latency. We let $\mathcal{F}^\star \subseteq \mathcal{F}$ denote the set of all Wardrop equilibria.

By definition, it follows that if a flow $f^\star \in \mathcal{F}^\star$ is a Wardrop equilibrium, then

$$c_p(f^\star) = c_{p'}(f^\star), \quad \forall p, p' \in \mathcal{P} \text{ with } f^\star_p, f^\star_{p'} > 0. \tag{6}$$

Indeed, if there were two paths $p, p' \in \mathcal{P}$ with positive flow such that $c_p(f^\star) < c_{p'}(f^\star)$, then players using path p' could reduce their latency by switching to path p, contradicting (5). Consequently, in a Wardrop equilibrium, all players experience the same latency regardless of the path they choose. This implies that the social cost at a Wardrop equilibrium is simply the *common* latency experienced by all players. Let $c(f^\star)$ denote the common latency; the social cost at $f^\star$ is

$$C(f^\star) = \sum_{p\in\mathcal{P}} f^\star_p \cdot c_p(f^\star) = c(f^\star) \cdot \sum_{p\in\mathcal{P}} f^\star_p = c(f^\star) \cdot 1 = c(f^\star). \tag{7}$$

In the next section, we will compare the social cost of Wardrop equilibria to that of an optimal flow, leading to the notion of the PoA in nonatomic routing games. First, there are a few nontrivial facts about the Wardrop equilibria for nonatomic routing games with *continuous* and *nondecreasing* latency functions that are worth mentioning before we define the PoA in nonatomic routing games. In particular, the Wardrop equilibria are the solutions to *a convex optimization problem over a compact convex set* [2]. Consequently, a Wardrop equilibrium *always exists* in such games. Furthermore, all Wardrop equilibria have the same social cost [2], and therefore the quality of a Wardrop equilibrium is *unambiguous*: it is the social cost of *any* a Wardrop equilibrium.

1.3 The price of anarchy of nonatomic routing games

Let us take another look at Braess's paradox in Figure 2. Previously, we considered the flow f in which 50% of the population chooses the upper path $s \to a \to t$ and 50% chooses the lower path $s \to b \to t$. We found that the induced flow on each arc is $f_{(s,a)} = f_{(a,t)} = f_{(s,b)} = f_{(b,t)} = 0.5$, while $f_{(a,b)} = 0$, and that the social cost of this flow is $C(f) = 1.5$. Is this flow a Wardrop equilibrium? To check, we compute the latencies of each path under flow f:

$$c_{s\to a\to t}(f) = c_{(s,a)}(0.5) + c_{(a,t)}(0.5) = 0.5 + 1 = 1.5; \tag{8a}$$
$$c_{s\to b\to t}(f) = c_{(s,b)}(0.5) + c_{(b,t)}(0.5) = 1 + 0.5 = 1.5; \tag{8b}$$
$$c_{s\to a\to b\to t}(f) = c_{(s,a)}(0.5) + c_{(a,b)}(0) + c_{(b,t)}(0.5) = 0.5 + 0 + 0.5 = 1. \tag{8c}$$

We see that players using the zigzag path $s \to a \to b \to t$ experience latency of 1, which is less than the latency of 1.5 experienced by players using the other two paths. Thus, players using either the upper path $s \to a \to t$ or the lower path $s \to b \to t$ could reduce their latency by switching to the zigzag path $s \to a \to b \to t$. Consequently, flow f is not a Wardrop equilibrium.

What is the Wardrop equilibrium of this game? Let us consider the flow $f^\star \in \mathcal{F}$ in which all players choose the zigzag path $s \to a \to b \to t$. In this case, the induced flow on each arc is $f^\star_{(s,a)} = f^\star_{(a,b)} = f^\star_{(b,t)} = 1$, while $f^\star_{(a,t)} = f^\star_{(s,b)} = 0$. The latencies of the paths under flow $f^\star$ are:

$$c_{s\to a\to t}(f^\star) = c_{(s,a)}(1) + c_{(a,t)}(0) = 1 + 1 = 2; \tag{9a}$$
$$c_{s\to b\to t}(f^\star) = c_{(s,b)}(0) + c_{(b,t)}(1) = 1 + 1 = 2; \tag{9b}$$
$$c_{s\to a\to b\to t}(f^\star) = c_{(s,a)}(1) + c_{(a,b)}(1) + c_{(b,t)}(1) = 1 + 0 + 1 = 2. \tag{9c}$$

We see that all players experience latency of 2 regardless of the path they choose. Thus, no player can reduce their latency by unilaterally switching paths. Consequently, flow $f^\star \in \mathcal{F}^\star$ is a Wardrop equilibrium. The social cost of $f^\star$ is $C(f^\star) = c_{s\to a\to b\to t}(f^\star) = 2$.

Braess's paradox arises from the observation that the social cost of the Wardrop equilibrium $f^\star$ is $C(f^\star) = 2$, which is greater than the social cost of the previously considered flow f with $C(f) = 1.5$. In other words, adding the zero-latency arc (a, b) to the network has increased the social cost of the equilibrium from 1.5 to 2!

The price of anarchy. The PoA of a nonatomic routing game quantifies the inefficiency of the Wardrop equilibria by comparing their social cost to that of an optimal flow $f^{\text{opt}} \in \mathcal{F}$, i.e., a flow that minimizes the game's social cost. Since in nonatomic routing games with continuous and nondecreasing latency functions all Wardrop equilibria have the same social cost, the PoA can be expressed in terms of this *common* value $C^\star \stackrel{\text{def}}{=} C(f^\star)$ for *any* $f^\star \in \mathcal{F}^\star$. Formally, the PoA is defined as

$$\text{PoA} = \frac{C^\star}{\min_{f\in\mathcal{F}} C(f)} = \frac{C^\star}{C(f^{\text{opt}})}. \tag{10}$$

In the example of Braess's paradox above, we have $C^\star = C(f^\star) = 2$, and we found a flow f with social cost $C(f) = 1.5$; therefore, the PoA is $\text{PoA} \geq 2/1.5 = 4/3$, indicating that the social cost at equilibrium is at least 33% higher than the optimal social cost. In fact, the PoA for Braess's paradox is exactly 4/3 since, *by symmetry*, the optimal flow must assign equal flow to the upper and lower paths, leading to the social cost $C(f^{\text{opt}}) = 1.5$.

How about the PoA in Pigou's example in Figure 1? Observe that in Pigou's example, the latency on arc e_1 is always less than or equal to the latency on arc e_2, i.e., $c_{e_1}(x) = x \leq 1 = c_{e_2}(x)$ for all $x \in [0, 1]$. Consequently, in a Wardrop equilibrium, all players choose arc e_1, leading to flow $f^\star \in \mathcal{F}^\star$ with induced arc flow $f^\star_{e_1} = 1$ and $f^\star_{e_2} = 0$. The social cost of this Wardrop equilibrium is $C(f^\star) = c_{e_1}(1) = 1$. To find the optimal flow, we consider a flow $f \in \mathcal{F}$ where a fraction $x \in [0, 1]$ of the population chooses arc e_1, and the remaining fraction $1 - x$ chooses arc e_2. The induced arc flow is $f_{e_1} = x$ and $f_{e_2} = 1 - x$, while the social cost of this flow is

$$C(f) = x \cdot c_{e_1}(x) + (1-x) \cdot c_{e_2}(1-x) = x \cdot x + (1-x) \cdot 1 = x^2 - x + 1, \quad \forall x \in [0,1]. \tag{11}$$

It follows that the social cost $C(f)$ is minimized at $x^{\text{opt}} = 0.5$, leading to the optimal social cost $C(f^{\text{opt}}) = (0.5)^2 - 0.5 + 1 = 0.75$. Therefore, the PoA for Pigou's example is $\text{PoA} = 1/0.75 = 4/3$. Is this a coincidence, or is there something special about the value 4/3?

1.4 The price of anarchy bound of 4/3

It turns out that the value 4/3 is not a coincidence. In particular, Pigou's example in Figure 1 attains the *worst-case* PoA among all nonatomic routing games with *linear* latency functions. This statement can, in fact, be generalized further. For every subclass of continuous and nondecreasing latency functions, the worst-case PoA over all nonatomic routing games whose latency functions belong to this subclass is attained by *a Pigou-like routing game*, which we formalize below.

Definition 1 (Pigou-like routing game)**.** Given a continuous and nondecreasing function $c\colon \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$, the Pigou-like routing game with respect to c is the nonatomic routing game:

$$\mathscr{P}_c = \big(\mathcal{G}, \{c_{e_1}, c_{e_2}\}\big) \tag{12}$$

whose graph $\mathcal{G}$ consists of two parallel arcs e_1 and e_2 from the source s to the sink t, as in Figure 1, with corresponding latency functions $c_{e_1} = c$ and $c_{e_2} \equiv c(1)$.

For example, Pigou's example in Figure 1 is a Pigou-like routing game with respect to the latency function $c(x) = x$, and the nonatomic routing game in the following figure is a Pigou-like routing game with respect to the latency function $c(x) = x^p$ for some fixed degree $p > 1$.

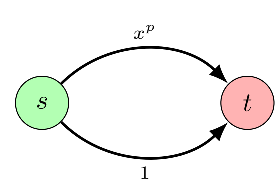

Figure 3: A Pigou-like routing game with respect to the function $c(x) = x^p$ for some fixed degree $p > 1$.

Before proving the statement, we first compute the PoA of the Pigou-like routing game $\mathscr{P}_c$ for a general continuous and nondecreasing function $c\colon \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$. Consider a flow $f \in \mathcal{F}$ in $\mathscr{P}_c$ where a fraction $x \in [0,\, 1]$ of the population chooses arc e_1 and the remaining fraction $1-x$ chooses arc e_2. The induced arc flow is $f_{e_1} = x$ and $f_{e_2} = 1-x$, with the social cost of this flow given by

$$C(f) = x \cdot c_{e_1}(x) + (1-x) \cdot c_{e_2}(1-x) = x \cdot c(x) + (1-x) \cdot c(1), \quad \forall x \in [0,\, 1]. \tag{13}$$

Since the latency function c is continuous and nondecreasing, the Wardrop equilibrium $f^\star \in \mathcal{F}^\star$ has all players choosing arc e_1, leading to induced arc flow $f^\star_{e_1} = 1$ and $f^\star_{e_2} = 0$. The social cost of this Wardrop equilibrium is $C(f^\star) = c_{e_1}(1) = c(1)$. Thus, the PoA of the Pigou-like routing game $\mathscr{P}_c$ is

$$\mathrm{PoA}(\mathscr{P}_c) = \frac{C(f^\star)}{\min_{f \in \mathcal{F}} C(f)} = \max_{x \in [0,\, 1]} \frac{c(1)}{x \cdot c(x) + (1-x) \cdot c(1)}. \tag{14}$$

Denote the worst-case PoA over all Pigou-like routing games whose latency functions belong to a subclass $\mathcal{C}$ of continuous and nondecreasing functions by

$$\alpha(\mathcal{C}) = \sup_{c \in \mathcal{C}} \mathrm{PoA}(\mathscr{P}_c) = \sup_{c \in \mathcal{C}} \max_{x \in [0,\, 1]} \frac{c(1)}{x \cdot c(x) + (1-x) \cdot c(1)}. \tag{15}$$

For many natural subclasses $\mathcal{C}$ of continuous and nondecreasing functions, $\alpha(\mathcal{C})$ can be computed in closed form. For example, for the subclass of linear functions $\mathcal{C}_{\mathrm{lin}} = \left\{c\colon \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0} \mid c(x) = ax + b,\ a, b \geq 0\right\}$, it can be shown that $\alpha(\mathcal{C}_{\mathrm{lin}}) = 4/3$. More generally, for the subclass of polynomial functions of degree at most p with nonnegative coefficients, $\mathcal{C}_{\mathrm{poly},d} = \left\{c\colon \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0} \mid c(x) = \sum_{i=0}^{p} a_i x^i,\ a_i \geq 0,\, i \in \{0, \ldots, p\}\right\}$, it can be shown [3] that

$$\alpha(\mathcal{C}_{\mathrm{poly},d}) = \frac{(d+1)\sqrt[d]{d+1}}{(d+1)\sqrt[d]{d+1} - d} \approx \frac{d}{\ln d}. \tag{16}$$

It remains to show that these bounds are indeed the worst-case PoA over all nonatomic routing games whose latency functions belong to the respective subclasses. This is captured by the following theorem due to Roughgarden [3].

Theorem 2 ([3])**.** *For every subclass $\mathcal{C}$ of continuous and nondecreasing functions and every nonatomic routing game whose latency functions belong to $\mathcal{C}$, the PoA is bounded from above by $\alpha(\mathcal{C})$.*

2 Atomic routing games

In the previous section, we considered nonatomic routing games, in which the flow of traffic between a designated source–sink pair (s, t) in a traffic network is controlled by a continuum of *infinitesimal players*. Each player controls an arbitrarily small fraction of the total flow, so their individual routing decision does not affect the latency on the arcs of the network. Next, we turn our attention to atomic routing games, in which the players control discrete *chunks* of flow, or units of traffic.

Our starting point is, once again, an n-player normal-form game $\mathscr{G} = \left(\llbracket n \rrbracket, \{\alpha_1, \alpha_2\}^n, \{u_i\}_{i \in \llbracket n \rrbracket}\right)$, in which each player $i \in \llbracket n \rrbracket$ has two pure strategies α_1 and α_2. This time, however, the payoff functions are not rescaled by a factor of $1/n$; instead, they are given by

$$u_i(s) = \begin{cases} -\sum_{j=1}^{n} \mathbf{1}_{s_i = s_j}, & \text{if } s_i = \alpha_1, \\ -n, & \text{if } s_i = \alpha_2, \end{cases} \tag{17}$$

for all $i \in \llbracket n \rrbracket$ and all strategy profiles $s \in \{\alpha_1, \alpha_2\}^n$. We may interpret this game as a routing game on a network with two parallel arcs e_1 and e_2 connecting a source node s to a sink node t, with respective latency functions $c_{e_1}(x) = x$ and $c_{e_2}(x) = n$, where x is the number of players (the *load*) choosing arc e_1. As before, choosing arc e_1 corresponds to strategy α_1, and choosing arc e_2 corresponds to strategy α_2. In contrast to the nonatomic case, each player controls *a unit of traffic*, and therefore their individual routing decision affects the load x on arc e_1. For that reason, the routing game is called *atomic*.

Formally, an atomic routing game is a tuple $\big(\mathcal{G}, \{c_e\}_{e\in\mathcal{E}}\big)$, where $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is a directed graph with a designated source–sink pair $(s,t) \in \mathcal{V}\times\mathcal{V}$, and each arc $e \in \mathcal{E}$ has an associated latency function $c_e \colon \mathbb{N} \to \mathbb{R}_{\geq 0}$ that specifies the travel time on e as a function of the load on it. The players' pure strategies are the set $\mathcal{P}$ of simple source–sink paths in $\mathcal{G}$, and the pure-strategy profiles $\mathcal{S} = \mathcal{P}^n$ is the ensembling thereof across all players. The load $x_e(p)$ on an arc $e \in \mathcal{E}$ under a strategy profile $p \in \mathcal{S}$ is the number of players whose chosen path includes e; i.e.,

$$x_e(p) = \sum_{i=1}^{n} \mathbf{1}_{e\in p_i}, \tag{18}$$

where p_i is the path chosen by player $i \in [\![n]\!]$ under strategy profile p, and $\mathbf{1}_{e\in p_i}$ is the indicator that equals 1 if $e \in p_i$ and 0 otherwise. Then the latency experienced by a player $i \in [\![n]\!]$ under a strategy profile p is the accumulated latency over all arcs on their chosen path p_i; i.e.,

$$\ell_i(p) = \sum_{e\in p_i} c_e\big(x_e(p)\big), \tag{19}$$

and, as in nonatomic routing games, the social cost $C(p)$ of a strategy profile $p \in \mathcal{S}$ is the total latency experienced by all players; i.e.,

$$C(p) = \sum_{i=1}^{n} \ell_i(p). \tag{20}$$

Nash equilibria in atomic routing games. A pure-strategy profile $p^\star \in \mathcal{S}$ is an NE if no player can unilaterally deviate to another path and reduce their latency, i.e., for all players $i \in [\![n]\!]$ and all paths $p_i \in \mathcal{P}$,

$$\ell_i(p^\star) \leq \ell_i(p_i, p^\star_{-i}), \tag{21}$$

where $(p_i, p^\star_{-i})$ is the strategy profile obtained by replacing player i's path in $p^\star$ with p_i. We use $\mathcal{S}^\star \subseteq \mathcal{S}$ to denote the set of all pure-strategy NE in the game.

As atomic routing games are special cases of normal-form games, it follows from Nash's theorem [4] that at least one *mixed-strategy* NE exists in such games. Moreover, atomic routing games are strategically equivalent to potential games [5]. Thus, as potential games admit pure-strategy NE [6], it follows that atomic routing games admit at least one *pure-strategy* NE. Therefore, we restrict our attention to the set of pure-strategy profiles $\mathcal{S}$, and we define the PoA of atomic routing games with respect to strategy profiles in $\mathcal{S}$.

As an example, consider the simple routing game depicted in Figure 4, which is the atomic counterpart of Pigou's example in Figure 1.

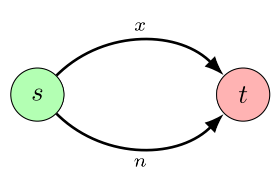

Figure 4: An atomic Pigou-like routing game.

One of the NE in this game is the strategy profile $p^\star \in \mathcal{S}^\star$ in which all n players choose arc e_1, resulting in social cost of $C(p^\star) = n^2$. When n is even, one optimal strategy profiles is p^{opt}, in which half of the players choose one arc, and the remaining players choose the other, resulting in social cost of $C(p^{\text{opt}}) = \frac{3}{4} \cdot n^2$.

2.1 The price of anarchy of atomic routing games

In contrast to the NE in nonatomic routing games, which all share a common social cost, the NE in atomic routing games may have *different social costs*. In this case, the PoA quantifies the inefficiency of the *worst-*

case pure-strategy NE relative to the optimal pure-strategy profile p^{opt}, i.e., $C(p^{\text{opt}}) = \min_{p\in\mathcal{S}} C(p)$. Formally, the PoA in atomic routing games is defined as

$$\text{PoA} = \frac{\max_{p^\star\in\mathcal{S}^\star} C(p^\star)}{\min_{p\in\mathcal{S}} C(p)} = \frac{\max_{p^\star\in\mathcal{S}^\star} C(p^\star)}{C(p^{\text{opt}})}. \tag{22}$$

For example, in the atomic Pigou-like routing game depicted in Figure 4, the PoA (for n even) is *at least*

$$\text{PoA} = \frac{C(p^\star)}{C(p^{\text{opt}})} = \frac{n^2}{\frac{3}{4}\cdot n^2} = \frac{4}{3}. \tag{23}$$

Unlike in the nonatomic case, $p^\star$ is not the only NE in this game. It turns out that there are *multiple NE* $p^{\star,j}$, for $j \in [\![n]\!]$, in which a single player j chooses arc e_2 instead of arc e_1, while the remaining players choose arc e_1. Since the social cost at each of these NE is $C(p^{\star,j}) = (n-1)^2 + n = n^2 - 2n + 1 + n = n^2 - n + 1 \leq n^2 = C(p^\star)$, it follows that the *worst-case* NE is indeed $p^\star$; therefore the PoA is *exactly* 4/3 when n is even.

Thus, the atomic counterpart of Pigou's example has PoA equal to 4/3 (for even n), the same PoA as Pigou's nonatomic example. Can we conclude that the PoA of atomic routing games with affine latency functions is also 4/3? The answer is *no*. We can construct atomic routing games with affine latency functions that exhibit PoA of 5/2 [7]. Furthermore, this is the worst-case PoA attainable by atomic routing games with affine latency functions. Formally, we have the following result due to Christodoulou et al. [7].

Theorem 3 ([7])**.** *In every atomic routing game with affine latency functions, the PoA is at most* $\frac{5}{2}$*.*

Proof. Let $p^\star \in \mathcal{S}^\star$ be the *worst-case* pure-strategy NE, and let $p^{\text{opt}} \in \mathcal{S}$ be an optimal pure-strategy profile minimizing the social cost, i.e., $C(p^{\text{opt}}) = \min_{p\in\mathcal{S}} C(p)$. Suppose that each arc $e \in \mathcal{E}$ is *affine*, i.e., there exist $a_e, b_e \in \mathbb{R}_{\geq 0}$ such that $c_e(x) = a_e x + b_e$ for all $x \in \mathbb{N}$. We now bound the social cost $C(p^\star)$ at the NE $p^\star$ in terms of the social cost $C(p^{\text{opt}})$ at the optimal strategy profile p^{opt}.

For convenience, for each $e \in \mathcal{E}$, let $x_e^\star = x_e(p^\star)$ denote the load on arc $e \in \mathcal{E}$ under the pure-strategy NE $p^\star$, and let $x_e^{\text{opt}} = x_e(p^{\text{opt}})$ denote the load on arc e under the optimal pure-strategy profile p^{opt}.

Bounding the social cost at the NE. Fix an arbitrary player $i \in [\![n]\!]$. By the definition of NE in Equation (21), we have $\ell_i(p^\star) \leq \ell_i(p_i, p_{-i}^\star)$ for all paths $p_i \in \mathcal{P}$. In particular, let p_i^{opt} denote the path chosen by player i in the optimal strategy profile p^{opt}; then

$$\ell_i(p^\star) \leq \ell_i(p_i^{\text{opt}}, p_{-i}^\star) \tag{24a}$$

$$= \sum_{e\in p_i^{\text{opt}}} c_e\big(x_e(p_i^{\text{opt}}, p_{-i}^\star)\big) \tag{24b}$$

$$= \sum_{e\in p_i^{\text{opt}}\cap p_i^\star} c_e\big(x_e(p_i^{\text{opt}}, p_{-i}^\star)\big) + \sum_{e\in p_i^{\text{opt}}\setminus p_i^\star} c_e\big(x_e(p_i^{\text{opt}}, p_{-i}^\star)\big) \tag{24c}$$

$$= \sum_{e\in p_i^{\text{opt}}\cap p_i^\star} c_e(x_e^\star) + \sum_{e\in p_i^{\text{opt}}\setminus p_i^\star} c_e(x_e^\star + 1) \tag{24d}$$

$$\leq \sum_{e\in p_i^{\text{opt}}} c_e(x_e^\star + 1). \tag{24e}$$

Thus summing over all players $i \in [\![n]\!]$, we obtain the following bound on the social cost at the NE $p^\star$:

$$
\begin{aligned}
C(p^\star) &\le \sum_{i=1}^{n} \sum_{e \in p_i^{\mathrm{opt}}} c_e(x_e^\star + 1) && \text{(25a)}\\
&= \sum_{e\in\mathcal{E}} \left(\big(c_e(x_e^\star + 1)\big) \cdot \sum_{i=1}^{n} \mathbf{1}_{e \in p_i^{\mathrm{opt}}} \right) && \text{(25b)}\\
&= \sum_{e\in\mathcal{E}} x_e^{\mathrm{opt}} \cdot \big(c_e(x_e^\star + 1)\big) && \text{(25c)}\\
&= \sum_{e\in\mathcal{E}} a_e \cdot x_e^{\mathrm{opt}}(x_e^\star + 1) + \sum_{e\in\mathcal{E}} b_e \cdot x_e^{\mathrm{opt}}. && \text{(25d)}
\end{aligned}
$$

Disentangling the terms $x_e^{\mathbf{opt}}(x_e^\star + 1)$. It can be shown that for all α, $\beta \in \mathbb{N}$,

$$
\alpha(\beta + 1) \le \frac{5}{3}\alpha^2 + \frac{1}{3}\beta^2. \tag{26}
$$

Thus by applying this inequality to each term $x_e^{\mathrm{opt}}(x_e^\star + 1)$ in the previous bound, we obtain

$$
\begin{aligned}
\sum_{e\in\mathcal{E}} a_e \cdot x_e^{\mathrm{opt}}(x_e^\star + 1) + \sum_{e\in\mathcal{E}} b_e \cdot x_e^{\mathrm{opt}} &\le \frac{5}{3}\sum_{e\in\mathcal{E}} a_e \cdot (x_e^{\mathrm{opt}})^2 + \frac{1}{3}\sum_{e\in\mathcal{E}} a_e \cdot (x_e^\star)^2 + \sum_{e\in\mathcal{E}} b_e \cdot x_e^{\mathrm{opt}} && \text{(27a)}\\
&\le \frac{5}{3}\sum_{e\in\mathcal{E}} x_e^{\mathrm{opt}}(a_e \cdot x_e^{\mathrm{opt}} + b_e) + \frac{1}{3}\sum_{e\in\mathcal{E}} x_e^\star \cdot (a_e \cdot x_e^\star) && \text{(27b)}\\
&\le \frac{5}{3}\sum_{e\in\mathcal{E}} x_e^{\mathrm{opt}}(a_e \cdot x_e^{\mathrm{opt}} + b_e) + \frac{1}{3}\sum_{e\in\mathcal{E}} x_e^\star \cdot (a_e \cdot x_e^\star + b_e) && \text{(27c)}\\
&= \frac{5}{3}\sum_{e\in\mathcal{E}} x_e^{\mathrm{opt}} c_e(x_e^{\mathrm{opt}}) + \frac{1}{3}\sum_{e\in\mathcal{E}} x_e^\star \cdot c_e(x_e^\star), && \text{(27d)}
\end{aligned}
$$

where in second inequality we used the fact that $\sum_{e\in\mathcal{E}} b_e \cdot x_e^{\mathrm{opt}} \ge 0$, and in the third inequality we used the fact that $b_e \ge 0$ for all $e \in \mathcal{E}$.

Bounding the PoA. Observe that for all pure-strategy profiles $p \in \mathcal{S}$, we have

$$
C(p) = \sum_{i=1}^{n} \ell_i(p) = \sum_{i=1}^{n}\sum_{e\in p_i} c_e\big(x_e(p)\big) = \sum_{e\in\mathcal{E}} \left(\Big(c_e\big(x_e(p)\big)\Big) \cdot \sum_{i=1}^{n} \mathbf{1}_{e\in p_i} \right) = \sum_{e\in\mathcal{E}} x_e(p) \cdot c_e(x_e(p)). \tag{28}
$$

Thus by applying this observation to the previous bound, we obtain

$$
\frac{5}{3}\sum_{e\in\mathcal{E}} x_e^{\mathrm{opt}}\big(a_e \cdot x_e^{\mathrm{opt}} + b_e\big) + \frac{1}{3}\sum_{e\in\mathcal{E}} x_e^\star \cdot (a_e \cdot x_e^\star + b_e) = \frac{5}{3}C(p^{\mathrm{opt}}) + \frac{1}{3}C(p^\star). \tag{29}
$$

Putting everything together, we have shown that

$$
C(p^\star) \le \frac{5}{3}C(p^{\mathrm{opt}}) + \frac{1}{3}C(p^\star) \implies C(p^\star) \le \frac{5}{2} \cdot C(p^{\mathrm{opt}}). \tag{30}
$$

Thus, by the definition of PoA, we have

$$
\mathrm{PoA} = \frac{C(p^\star)}{C(p^{\mathrm{opt}})} \le \frac{\frac{5}{2}C(p^{\mathrm{opt}})}{C(p^{\mathrm{opt}})} \le \frac{5}{2}. \tag{31}
$$

□

Together with the example of an atomic routing game with affine latency functions that exhibits PoA 5/2, we conclude that the bound in Theorem 3 is tight. Thus, whereas in the nonatomic case the worst-case PoA is 4/3, in the atomic case the worst-case PoA with affine latency functions is 5/2.

A List of abbreviations

NE **N**ash **E**quilibrium 1, 3, 4, 8–10

PoA **P**rice **o**f **A**narchy 1, 3, 5–10

B Index

References

[1] Avrim Louis Blum, Eyal Even-Dar, and Katrina Ligett. "Routing Without Regret. On Convergence to Nash Equilibria of Regret-Minimizing Algorithms in Routing Games." In: *Theory of Computing* 6.8 (2010), pp. 179–199. ISSN: 1557-2862. DOI: 10.4086/toc.2010.v006a008.

[2] Martin Beckmann, Charles Bartlett McGuire, and Christopher B. Winsten, eds. *Studies in the Economics of Transportation*. New Haven, Connecticut, United States: Yale University Press, 1956.

[3] Tim Roughgarden. "The price of anarchy is independent of the network topology." In: *Journal of Computer and System Sciences* 67.2 (2003), pp. 341–364. ISSN: 1090-2724. DOI: 10.1016/S0022-0000(03)00044-8.

[4] John Forbes Nash Jr. "Non-Cooperative Games." In: *Annals of Mathematics* 54.2 (Sept. 1951), pp. 286–295. ISSN: 0003-486X. DOI: 10.2307/1969529.

[5] Robert Wernick Rosenthal. "A Class of Games Possessing Pure-Strategy Nash Equilibria." In: *International Journal of Game Theory* 2 (Dec. 1973), pp. 65–67. ISSN: 1432-1270. DOI: 10.1007/BF01737559.

[6] Dov Monderer and Lloyd Stowell Shapley. "Potential Games." In: *Games and Economic Behavior* 14.1 (1996), pp. 124–143. ISSN: 0899-8256. DOI: 10.1006/game.1996.0044.

[7] George Christodoulou and Elias Koutsoupias. "The price of anarchy of finite congestion games." In: *STOC 2021. Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. Symposium on Theory of Computing (June 4–6, 2005). New York, New York, United States: Association for Computing Machinery, 2005, pp. 67–73. ISBN: 978-1-4503-8053-9. DOI: 10.1145/1060590.1060600.