

Authors are encouraged to submit new papers to INFORMS journals by means of a style file template, which includes the journal title. However, use of a template does not certify that the paper has been accepted for publication in the named journal. INFORMS journal templates are for the exclusive purpose of submitting to an INFORMS journal and should not be used to distribute the papers in print or online or to submit the papers to another publication.

In Congestion Games, Taxes Achieve Optimal Approximation

Dario Paccagnan

Department of Computing, Imperial College London, London, UK, d.paccagnan@imperial.ac.uk

Martin Gairing

Department of Computer Science, University of Liverpool, Liverpool, UK, m.gairing@liverpool.ac.uk

In this work, we address the problem of minimising social cost in atomic congestion games. For this problem, we present lower bounds on the approximation ratio achievable in polynomial time and demonstrate that efficiently computable taxes result in polynomial time algorithms matching such bounds.

Perhaps surprisingly, these results show that indirect interventions, in the form of efficiently computed taxation mechanisms, yield the *same* performance achievable by the best polynomial time algorithm, even when the latter has full control over the agents' actions. It follows that no other tractable approach geared at incentivizing desirable system behavior can improve upon this result, regardless of whether it is based on taxations, coordination mechanisms, information provision, or any other principle. In short: Judiciously chosen taxes achieve optimal approximation. Three technical contributions underpin this conclusion. First, we show that computing the minimum social cost is NP-hard to approximate within a given factor depending solely on the admissible cost functions. Second, we design a polynomially computable taxation mechanism whose efficiency (price of anarchy) matches this hardness factor, and thus is optimal amongst all tractable mechanisms. As these results extend to coarse correlated equilibria, any no-regret algorithm inherits the same performances, allowing us to devise polynomial time algorithms with optimal approximation.

Key words: congestion games, minimum social cost, hardness of approximation, optimal tolls, approximation algorithms, price of anarchy.

1. Introduction

Decision-making in the presence of congestion effects is a central topic in the operations research literature. Within this arena, the celebrated *congestion game* model proposed by Rosenthal (1973) offers a fundamental framework to study resource allocation problems prone to congestion, finding fitting applications in transportation, telecommunications, scheduling, and many other disciplines.

In an atomic congestion game, we are given a finite set of resources and a finite number of agents. Every agent is endowed with a set of feasible allocations, each corresponding to a subset of the set of resources. These subsets may, for example, correspond to machines where an agent wishes to schedule their jobs, or to paths connecting a given origin-destination in a graph. When selecting an allocation, each agent incurs an additive cost over the chosen resources, where the cost of a particular resource depends solely on the number of agents using it, i.e., on the congestion on that resource. At the system level, the quality of a joint allocation is measured through the *social cost* obtained by summing all the agents' costs.

Perhaps the most commonly encountered class of congestion games is that of *network* congestion games, where each feasible allocation corresponds to a path connecting a given pair of origin-destination nodes over a shared graph. This class of congestion games arises in the study of classical optimization problems in road-traffic, wireless-network, and minimum-power routing – see (Rosenthal 1973, Tekin et al. 2012, Andrews et al. 2012) respectively. However, congestion games find interesting applications also beyond network-like settings for example in machine scheduling (Suri et al. 2007), distributed control (Marden and Wierman 2013), sensor allocation (Paccagnan et al. 2020), and factory production (Rosenthal 1973).

Complementary to much of the existing literature, in this work we are concerned with a *system optimal* perspective on congestion games. Specifically, we study the problem of minimising the social cost and consider polynomial running time as a primary concern. We do so in two settings, depending on whether agents are modeled as strategic decision-makers or not.

In the *non-strategic* setting, agents allocations are directly controlled by a planner who aims at minimising social cost. Such problems arise commonly in load balancing (Awerbuch et al. 1995), power routing (Andrews et al. 2012), mobility on-demand (Zardini et al. 2021) and other applications. Interestingly, despite the significant scientific interest and the connections to related questions in operations research, little is known regarding the problem of minimising the social cost in this setting. Most notably, it remains unclear what solution quality polynomial time algorithms can attain¹, and what are the computational barriers to do so (if any). *Our manuscript settles both questions.*

In the *strategic* setting agents are modeled as self-interested decision-makers aiming to minimize their individual cost functions. This setting finds its root in the seminal works of Von Neumann and Morgenstern (1944) and Nash (1950) which laid the foundation for the nascent field of game theory through the definition of Nash equilibria. Crucially, as individual objectives need not be aligned to that of minimising social cost, equilibrium allocations often display a significant performance degradation. A prototypical example of this issue is provided by road-traffic routing: When drivers choose routes that minimize their own travel time, the aggregate congestion could be much higher compared to that of a centrally-imposed routing. While improved performances could be attained by dictating the agents’ choices, imposing such control is often impossible, with traffic routing providing but one illustration. Hence, different approaches including coordination mechanisms, Stackelberg strategies, taxation mechanisms, information provision, and alternative methods for sharing costs have been proposed as *indirect interventions* to influence the equilibrium efficiency – see (Christodoulou et al. 2009, Bonifaci et al. 2010, Caragiannis et al. 2010, Bhaskar et al. 2016, Gkatzelis et al. 2016) respectively. Amongst these, taxation mechanisms, where agents pay an additional cost for utilizing a resource, have attracted significant attention, as witnessed by the growing literature on the topic, e.g., (Harks 2019) and references therein. In spite of that, the problem of designing polynomially computable taxation mechanisms to optimize the equilibrium efficiency remains open even for classical settings such as that of atomic congestion games. Additionally, it remains unclear, if, and to what extent, the use of indirect interventions reduces the performance achievable by centrally imposed decision-making. *Our manuscript resolves also these questions.*

Preview of our contribution

Our manuscript contains three main contributions that settle the questions previously introduced.

¹ As classical in the study of approximation algorithms, we use solution quality to refer to the *worst-case* quality over all instances on which the algorithm will run. Similarly for the other metrics we will introduce. In other words, we are concerned with devising polynomial time algorithms that provide the best solution quality in face of the worst instance. As such, instances encountered in the real-world could exhibit significantly improved performance compared to the worst-case one.

- First, we show that, in congestion games, minimizing the social cost is NP-hard to approximate within an explicitly given factor depending solely on the class of resource costs. We note that this result provides an inherent computational barrier that applies to *both* the strategic and the non-strategic settings. As we will see, the computational lower bounds we provide here are *tight*, i.e., they are attained by a polynomial time algorithm. Prior to our work, no tight computational lower bound was known, even for the special case of linear resource costs.
- Second, we derive polynomially computable taxation mechanisms ensuring that any resulting equilibrium outcome has an efficiency (price of anarchy) matching the hardness factor previously derived. While this result holds for the commonly employed notion of pure Nash equilibrium, it also extends to more general notions of equilibrium that can be computed efficiently (e.g., correlated and coarse correlated equilibria). Moreover, it extends to no-regret algorithms where agents update their allocations and achieve low regret, in the same spirit of the “total price of anarchy” pioneered by Blum et al. (2008).
- Third and final, we derive polynomial time algorithms achieving an approximation that matches the hardness factor previously derived. Jointly with the first contribution, this certifies that the proposed algorithms are *optimal* in the sense that no polynomial time algorithm can yield a better approximation over all instances. While different polynomial-time algorithms might provide a better performance on some instances (not all), we will see that the approximation we derive is near-optimal for many commonly encountered real-world settings, including the case where resource costs increase not too steeply (e.g., polynomials of low degree), or where the uncongested and congested costs are of comparable magnitude (e.g., BPR functions).

When taken together our contributions show that no performance degradation arises when moving from centrally imposed decision-making to the use of indirect interventions, such as taxation mechanisms. All in all, of our work can be summarized as follows:

In congestion games, judiciously designed taxes achieve optimal approximation, and no other tractable intervention, whether based on coordination mechanisms, information provision, or any other principle can improve upon this result.

Significance and comparison with existing results

Our work connects with a number of existing results, in addition to closing different open questions, as we briefly highlight next. We refer to Section 1.3 for a more detailed literature review, and to Section 1.4 for a discussion on the continuous flow counterpart of congestion games.

The problem of determining computational lower bounds for minimizing the social cost in atomic congestion games has been initially studied by Meyers and Schulz (2012). Since then, a number of works have pursued a similar line of research, though no tight bounds were known. Our work provides the best possible inapproximability results, completely settling the hardness question for congestion games with general underlying resource costs.

The study of taxation mechanisms has recently received growing attention, especially within the congestion games literature. Nevertheless, prior to this work, a methodology to design optimal taxation mechanisms (i.e., mechanisms minimizing the price of anarchy) was not known. For congestion games with polynomial costs, Caragiannis et al. (2010) (for linear) and Bilò and Vinci (2019) (for polynomial) propose taxes whose efficiency can be quantified a priori. Both papers conjecture that their design is optimal. Our work resolves the problem of determining optimal taxes for the broader class of congestion games with non-decreasing and semi-convex resource costs, proving these conjectures as a special case.

Roughgarden (2014) studies how lower bounds on the price of anarchy can be derived from computational lower bounds. For congestion games with optimal taxes, we show that such an approach does provide *tight* bounds.

The best known approximation algorithm for minimizing the social cost in congestion games is due to Makarychev and Sviridenko (2018), and leverages a natural linear programming relaxation jointly with a randomized rounding scheme. While their result applies to the more general class of optimization problems with a “diseconomy of scale”, the algorithms we propose here enjoy an equal or strictly better approximation ratio, and can not be further improved, owing to the matching hardness result provided.

Finally, we remark that our work focuses on the original congestion game model proposed by Rosenthal (1973) and not on its continuous flow counterpart. Although connected, the challenges associated with designing optimal taxation mechanisms are significantly different in these two settings, as we carefully discuss in Section 1.4.

1.1. Congestion games and taxation mechanisms

In a congestion game we are given a set of agents $\{1, \dots, N\}$, and a set of resources \mathcal{R} . Each agent can choose a subset of the set of resources which she intends to use. We list all feasible choices for agent i in the set $\mathcal{A}_i \subseteq 2^{\mathcal{R}}$. The cost for using each resource $r \in \mathcal{R}$ depends only on the total number of agents concurrently selecting that resource, and is denoted with $\ell_r : \mathbb{N} \rightarrow \mathbb{R}_{>0}$. Once all agents have made a choice $a_i \in \mathcal{A}_i$, each agent incurs a cost obtained by summing the costs of all resources she selected. Finally, the social cost represents the sum of the resource costs incurred by all agents

$$\text{SC}(a) = \sum_{i=1}^N \sum_{r \in a_i} \ell_r(|a|_r), \quad (1)$$

where $|a|_r$ denotes the number of agents selecting resource r in allocation $a = (a_1, \dots, a_n)$. Given an instance G of a congestion game, we denote with MinSC the problem of globally minimizing the social cost in (1).

Taxation mechanisms. As self-interested decision making often deteriorates the system performance, taxation mechanisms have been proposed to ameliorate this issue. Formally, a taxation mechanism $T : (G, r) \mapsto \tau_r$ associates an instance G , and a resource $r \in \mathcal{R}$ to a taxation function $\tau_r : \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$. Note that each taxation function τ_r is congestion-dependent, that is, it associates the number of agents in resource r to the corresponding tax.² As a consequence, each agent i experiences a cost factoring both the cost associated to the chosen resources, and the tax, i.e.,

$$C_i(a) = \sum_{r \in a_i} [\ell_r(|a|_r) + \tau_r(|a|_r)]. \quad (2)$$

We remark that, in doing so, a taxation mechanism influences the agents’ perceived cost but does not impact the expression of the social cost, which is still of the form in (1).³ As typical in the literature, we measure the performance of a given taxation mechanism T using the ratio between the social cost incurred at the worst-performing outcome, and the minimum social cost. Given the self-interested nature of the agents, an outcome is commonly described by any of the following classical equilibrium notions: pure or mixed Nash equilibria, correlated or coarse correlated equilibria.⁴ When considering pure Nash equilibria, for example, the performance of a taxation mechanism

² We note that congestion-dependent taxation mechanisms are commonly studied in the literature and have been tested in the real world too, see, e.g., (Axhausen et al. 2021) and references therein. In the context of road traffic routing, increased connectivity and the advent of smart vehicles will further ease their implementation.

³ That is, taxation mechanisms are a simple monetary transfer from the agents to the central coordinator.

⁴ It is worth observing that each class of equilibria appearing in this list is a superset of the previous (Roughgarden 2015). Therefore, since pure Nash equilibria are guaranteed to exist even in congestion games where taxes are used (they are, in fact, potential games), all mentioned equilibrium’s sets are non-empty and thus the notion of *price of anarchy* introduced in (3) well defined.

T over the class of congestion games \mathcal{G} is gauged using the notion of *price of anarchy* originally introduced by Koutsoupias and Papadimitriou (1999), i.e.,

$$\text{PoA}(T) = \sup_{G \in \mathcal{G}} \frac{\text{NeCost}(G, T)}{\text{MinCost}(G)}, \quad (3)$$

where $\text{MinCost}(G) = \min_{a \in \mathcal{A}} SC(a)$ is the minimum social cost for instance G , and $\text{NeCost}(G, T)$ denotes the highest social cost at a Nash equilibrium obtained when employing the mechanism T on the game G . By definition, $\text{PoA}(T) \geq 1$ and the lower the price of anarchy, the better performance T guarantees. While it is possible to define the notion of price of anarchy for each and every equilibrium class mentioned, we do not pursue this direction, as we will show that all these metrics coincide within our setting. Thus, we will simply use $\text{PoA}(T)$ to refer to the efficiency values of *any* and *all* these equilibrium classes.

All our results hold for the widely studied case of non-decreasing semi-convex resource costs, corresponding to all resource costs satisfying the following assumption.

ASSUMPTION 1. *The function $\ell : \mathbb{N} \rightarrow \mathbb{R}_{>0}$ is non-decreasing and semi-convex.*⁵

Notation. In our work \mathbb{N} , \mathbb{N}_0 , \mathbb{R} , $\mathbb{R}_{\geq 0}$, $\mathbb{R}_{>0}$ denote the sets of natural numbers, natural numbers including zero, real numbers, non-negative real numbers, and positive real numbers. Further, $\text{Poi}(x)$ denotes a Poisson distribution with parameter $x \in \mathbb{R}_{>0}$. In the remainder of the manuscript, we extend the domain of definition of all resource costs from \mathbb{N} to \mathbb{N}_0 by setting their value to zero. This is without loss of generality. Indeed, the value of $\ell_r(0)$, does not play any role with respect to all quantities we have introduced thus far, e.g., the social cost in (1), since $\ell_r(|a|_r)$ is always evaluated for $|a|_r \geq 1$. However, doing so will allow us to simplify notation.⁶

1.2. Our contribution

The resounding message of this work can be summarized as follows: In congestion games, optimally designed taxation mechanisms can be tractably computed, and achieve the same performance of the best centralized polynomial time algorithm. Three technical contributions substantiate this claim:

- i) We prove a tight NP-hardness result for minimizing the social cost.*
- ii) We design a tractable taxation mechanism whose price of anarchy matches the hardness factor.*
- iii) We obtain a polynomial time algorithm with the best possible approximation ratio for the social cost. We do so combining the previous results with existing algorithms (e.g., no-regret dynamics).*

Results *ii)* and *iii)* extend to network congestion games, as we discuss in the conclusions.

We now turn our attention to discussing each of these contributions in finer detail.

i) Inapproximability of minimum social cost. Our first contribution is concerned with determining tight inapproximability results for the problem of minimizing social cost in congestion games. Our hardness result applies already in the setup where all resources feature the same cost.

THEOREM 1. *Let $\ell : \mathbb{N} \rightarrow \mathbb{R}_{>0}$ satisfying Assumption 1 be given. In congestion games where all resources feature the same cost ℓ , MinSC is NP-hard to approximate within any factor smaller than*

$$\rho_\ell = \sup_{x \in \mathbb{N}} \frac{\mathbb{E}_{P \sim \text{Poi}(x)}[P\ell(P)]}{x\ell(x)}, \quad (4)$$

where we define $\ell(0) = 0$. If $\rho_\ell = \infty$, then MinSC is NP-hard to approximate within any finite factor.

⁵ $\ell : \mathbb{N} \rightarrow \mathbb{R}_{>0}$ is semi-convex if $x\ell(x)$ is convex, i.e., $(x+1)\ell(x+1) - x\ell(x) \geq x\ell(x) - (x-1)\ell(x-1) \forall x \in \mathbb{N}, x \geq 2$.

⁶ For example, it will allow us to write the numerator of (4) compactly as $\mathbb{E}_{P \sim \text{Poi}(x)}[P\ell(P)] = \sum_{k=0}^{\infty} k\ell(k)x^k e^{-x}/k!$, as opposed to $\sum_{k=1}^{\infty} k\ell(k)x^k e^{-x}/k!$ when $\ell(0)$ is undefined. Note here the different extremes of summation.




Figure 1 Inapproximability factors for congestion games with BPR resource costs $\ell_r(x) = \alpha_r x^4 + \beta_r$ (left) and affine resource costs $\ell_r(x) = \alpha_r x + \beta_r$ (right), parametrized by $\max_r \alpha_r / \beta_r$. Commonly encountered instances, e.g., those taken from (Transportation Networks for Research 2016) feature $\alpha_r \ll \beta_r$ yielding an inapproximability factor close to one, see markers on the left panel.

Naturally, Theorem 1 applies directly to richer classes of congestion games, whereby resource costs can differ. For example, if resource costs belong to a set of functions \mathcal{L} , then MinSC is NP-hard to approximate within any factor smaller than that produced by the worst function in \mathcal{L} , i.e., $\sup_{\ell \in \mathcal{L}} \rho_\ell$. As a special case, we obtain hardness results for the thoroughly studied class of polynomial congestion games with maximum degree d . In this case, the highest degree monomial x^d determines the worst factor (see Section 2.4), which reduces to the $(d+1)$ 'st Bell number, as summarized in the following statement.

COROLLARY 1. *In congestion games with resource costs obtained by non-negative combinations of $1, \dots, x^d$, MinSC is NP-hard to approximate within a factor smaller than the $(d+1)$ 'st Bell number.*

We note that Bell numbers grow as function of the degree d , so that the corresponding inapproximability factor also grows with d . However, for many problems of interest, resource costs feature further properties that can be exploited to reduce such a factor. This is interesting, since the algorithms we propose will provide approximations matching this factor. For example, in road-traffic routing, the commonly employed BPR functions take the form $\ell_r(x) = \alpha_r x^4 + \beta_r$ where $\alpha_r \ll \beta_r$ for all resources, see for example the commonly employed Sioux Fall network and other instances from (Transportation Networks for Research 2016). In this context, a direct application of Theorem 1 provides inapproximability factors significantly closer to one. We elucidate this point in the following corollary where we look at congestion games with affine resource costs, as this allows to present a concise analytical result. However, a similar result holds for BPR functions, as demonstrated in Figure 1 (left).

COROLLARY 2. *In congestion games with affine resource costs of the form $\ell_r(x) = \alpha_r x + \beta_r$, where $\max_r \alpha_r / (\alpha_r + \beta_r) \leq q$, MinSC is NP-hard to approximate within a factor smaller than $1 + q$.*

The latter result shows that the inapproximability factor depends only on the ratios between the coefficients α_r and β_r . To see this, denote $\max_r \alpha_r / \beta_r = \kappa$, and observe that $\max_r \alpha_r / (\alpha_r + \beta_r) = \kappa / (\kappa + 1)$ when $\beta_r \neq 0$. Consequently, the inapproximability factor $1 + q$ can be rewritten more conveniently as $1 + \kappa / (\kappa + 1)$, see Fig. 1 (right) for a plot. In the regime where $\alpha_r \ll \beta_r$ for all resources, we have $\kappa \ll 1$, so that the inapproximability factor is close to one. When this ratio grows, the corresponding factor also grows up until it reaches its maximum when $\beta_r \rightarrow 0$ for some resource. In other words, the inapproximability factor depends on the ratio between the speed at which congestion translates into cost (here α_r) and the uncongested cost (here β_r). Interestingly, in many settings the ratio $\alpha_r / \beta_r \ll 1$, thus ensuring that the forthcoming algorithms will provide near-optimal approximation.

Regarding our techniques, the computational lower bounds are shown by reducing the problem of minimizing the social cost in congestion games from the *gap label cover* problem implicitly defined in Feige (1998). A central tool we leverage in our reduction is a gadget called *partitioning system*, generalizing that in Feige (1998). The challenge in defining this object stems from the fact that, in congestion games, the cost experienced over each resource depends on the actual number of agents selecting that resource, contrary to Feige (1998) where the welfare accrued depends solely on whether that resource is selected or not. The expression of the hardness factor (4) falls from the very definition of such an object, which must be duplicated multiple times and carefully arranged in our construction to ensure hardness of the underlying instance. Conceptually, partitioning systems are constructed so that, in our reduction, agents have two types of allocations: either they coordinate among themselves on an allocation with low social cost, or they opt for an allocation with high social cost. The inapproximability factor in (4) is precisely the ratio between these “high” and “low” social costs, see Section 2.1.2. Partitioning systems have been recently applied also by Dudycz et al. (2020) and Barman et al. (2021) in the context of approval voting, and for a generalization of the maximum coverage problem. To the best of our knowledge, our work is the first that employs such gadgets to prove hardness of cost minimization problems.

ii) Taxes achieve optimal approximation. Our second contribution provides a technique to efficiently design taxation mechanisms whose efficiency (price of anarchy) matches the hardness result. The following statement is a succinct variant of that in the forthcoming Theorem 3 where we also include an expression for the taxation mechanism.

THEOREM 2. *Consider the class of congestion games where each resource cost belongs to a set of functions \mathcal{L} satisfying Assumption 1. Then, for any given $\varepsilon > 0$, it is possible to efficiently compute a taxation mechanism whose price of anarchy is upper bounded by $\sup_{\ell \in \mathcal{L}} \rho_\ell + \varepsilon$. The result holds for pure/mixed Nash and correlated/coarse correlated equilibria.*

The extension to coarse correlated equilibria is significant as it gives performance bounds that apply not only to Nash equilibria, but also whenever agents revise their action and achieve low regret. This imposes much weaker assumptions on both the game and its participants’ behavior (Roughgarden 2015). Whilst the theorem applies to a broad class of resource costs, when specialized to polynomial congestion games of degree d , it allows to efficiently design taxation mechanisms whose price of anarchy equals the $(d + 1)$ ’st Bell number plus ε , for any arbitrarily small choice of $\varepsilon > 0$.

We derive taxes whose efficiency matches the hardness factor leveraging two chief ingredients: a parametrized class of taxation mechanisms satisfying a key recursion, and a suitably defined convex optimization program. The convex optimization problem we consider corresponds to a modification of the original MinSC, whereby we relax the integrality constraints and replace the cost $x\ell_r(x)$ produced by each resource with $\mathbb{E}_{P \sim \text{Poi}(x)}[P\ell_r(P)]$, see (12). The solution vector of this program, which is provably convex, is used as set of parameters for the class of mechanisms previously defined. The performance bound on the price of anarchy is finally shown through a smoothness-like approach, where we leverage both the expression of the mechanisms and the solution vector of the convex program.

iii) Optimal polynomial time algorithms. Since the result in Theorem 2 holds also for correlated/coarse correlated equilibria, we can leverage existing polynomial time algorithms to compute such equilibria, and inherit an approximation ratio matching the corresponding price of anarchy. The following statement summarizes this result, while the ensuing discussion provides two possible approaches to do so. We remark that Corollary 3 is a direct consequence of Theorem 2 and polynomial computability of correlated equilibria (Jiang and Leyton-Brown 2015). The fact that the resulting approximation factor we obtain always matches or strictly improves upon that of Makarychev and Sviridenko (2018) is shown in Section 3.1.

COROLLARY 3. *Consider congestion games where each resource cost belongs to a set of functions \mathcal{L} satisfying Assumption 1. Then, for any $\varepsilon > 0$, there exists a polynomial time algorithm computing an allocation with cost lower than $(\sup_{\ell \in \mathcal{L}} \rho_\ell + \varepsilon) \cdot \min_{a \in \mathcal{A}} \text{SC}(a)$.*

The approximation ratio presented in Corollary 3 can be achieved, for example, as follows. Given a desired tolerance $\varepsilon > 0$, we design a taxation mechanism ensuring a price of anarchy of $\sup_{\ell \in \mathcal{L}} \rho_\ell + \varepsilon/2$, which can be done in polynomial time thanks to Theorem 2. We use such taxation mechanisms, and compute an exact correlated equilibrium in polynomial time leveraging the result of Jiang and Leyton-Brown (2015), who propose a variation of the Ellipsoid Against Hope algorithm from Papadimitriou and Roughgarden (2008). Remarkably, Jiang and Leyton-Brown (2015) guarantees that the resulting correlated equilibrium has polynomial-size support, i.e., it places non-zero probability only over a polynomial number of pure strategy profiles. Hence, we compute a correlated equilibrium and enumerate all pure strategy profiles in its support, identifying with a^* that with lowest cost. Since the price of anarchy bounds of Theorem 2 hold for correlated equilibria, the pure strategy profile a^* inherits a matching (or better) approximation ratio.

Alternatively, one can employ the same taxation mechanism as in the above, and let agents simultaneously revise their action for t rounds employing a no-regret algorithm. Well-known families of such algorithms include *Multiplicative-Weights*, and *Follow the Perturbed/Regularized Leader* whereby the average regret decays to an arbitrary ε in a polynomial number of rounds, see Cesa-Bianchi and Lugosi (2006), Kalai and Vempala (2005). We keep track of the pure strategy profile a_t^* with the lowest social cost encountered during the t rounds of any such algorithm. Owing to (Roughgarden 2015, Thm 3.3), its approximation ratio is upper bounded by the corresponding price of anarchy plus an error term that goes to zero with the same rate of the average regret. Waiting for polynomially many rounds suffices to reduce the error as desired.

1.3. Related work

As our work provides tight computational lower bounds, optimal taxation mechanisms, and polynomial time algorithms with the best possible approximation, we review the relevant literature connected with these three areas in the ensuing paragraphs. We conclude comparing our results with those available for the continuous flow counterpart of congestion games (nonatomic model).

Computational lower bounds. The study of computational lower bounds for minimizing the social cost in congestion games has been pioneered by Meyers and Schulz (2012), though remarkable precursors include Chakrabarty et al. (2005), as well as Blumrosen and Dobzinski (2006) who considered a notably different model whereby agent-specific cost functions are utilized. Relative to the classical model of congestion games with convex non-decreasing resource costs, Meyers and Schulz show that minimizing the social cost is strongly NP-hard, and it is hard to approximate within any finite factor, unless $\text{P}=\text{NP}$. Identical results are shown for network congestion games. Whilst this might feel as a contradiction of our results, it is worth noting that their analysis allows for resource costs to be adversarially selected amongst any convex non-decreasing function. On the contrary, our result can be thought of as a refined version of theirs, whereby the computational lower bound we derive is parametrized by the class of admissible resource costs. Naturally, we recover their inapproximability result when resource costs can be arbitrarily selected amongst convex non-decreasing functions.⁷

Motivated by the possibility to translate computational lower bounds to lower bounds on the price of anarchy, Roughgarden (2014) also studied this problem. Relative to polynomial resource costs of maximum degree d and non-negative coefficients, he showed that minimizing the social cost is inapproximable within any factor smaller than $(\beta d)^{d/2}$, for some constant $\beta > 0$. In this setting,

⁷For example MinSC is NP-hard to approximate within any finite factor when $b(x) = e^x$. This is an immediate consequence of Theorem 1 and the fact that $\rho_b = \infty$ for exponentially increasing resource costs.

even without taxes, equilibria with much better performance are guaranteed to exist. In particular, the *price of stability* (measuring the quality of the best-performing equilibrium) is known to grow only linearly with the degree d (Christodoulou and Gairing 2016). While coordinating the agents to one such *good* equilibrium is highly desirable, our hardness result implies that this cannot be achieved in polynomial time.

Spurred by the inapproximability results of Meyers and Schulz (2012), and Roughgarden (2014), a number of works have focused on restricting the allowable class of problems: Pia et al. (2017) study totally unimodular congestion games and show NP-hardness for the asymmetric case; Castiglioni et al. (2021) show NP-hardness even in singleton congestion games with affine resource costs. Similar questions have been explored for online versions of the problem by Klimm et al. (2019).

Taxation mechanisms. Different approaches, such as coordination mechanisms (Christodoulou et al. 2009), Stackelberg strategies (Fotakis 2010), taxation mechanisms (Caragiannis et al. 2010), signaling (Bhaskar et al. 2016), cost-sharing strategies (Gairing et al. 2020), and many more, have been proposed to cope with the performance degradation associated to selfish decision making. Amongst them, taxation mechanisms have attracted significant attention thanks to their ability to indirectly influence the resulting system performance. While the study of taxation mechanisms in road-traffic networks was initiated by Pigou (1920), who utilized a continuous flow model (nonatomic congestion games), the design of taxation mechanisms in (atomic) congestion games was pioneered much more recently by Caragiannis et al. (2010). In this respect, Caragiannis et al. (2010) and many of the subsequent works, build on the solid theoretical ground developed in the years subsequent to the definition of the price of anarchy (Koutsoupias and Papadimitriou 1999), including exact knowledge of the price of anarchy in congestion games with linear (Christodoulou and Koutsoupias 2005, Awerbuch et al. 2013) and polynomial (Aland et al. 2011) resource costs, the advent of the smoothness framework (Roughgarden 2015), as well as primal dual approaches (Bilò 2018, Chandan et al. 2019).

While these results provide us with a strong theory to quantify the price of anarchy, prior to this work, the design of optimal taxation mechanisms (i.e., taxation mechanisms minimizing the price of anarchy) has been an open question even when restricting attention to linear resource costs. Most notably, Caragiannis et al. (2010) considers linear congestion games and designs taxation mechanisms achieving a price of anarchy of 2 for mixed Nash equilibria. More recently, Bilò and Vinci (2019) extend their results to polynomial congestion games of degree d achieving a price of anarchy (for coarse correlated equilibria) equal to the $(d+1)$ 'st Bell number. Our work resolves the problem of designing optimal taxation mechanisms for more general congestion games with semi-convex non-decreasing resource costs, and, as a special case, shows that the mechanisms proposed for linear (Caragiannis et al. 2010) and polynomial (Bilò and Vinci 2019) resource costs are optimal, as conjectured by the authors.

Perhaps closest in spirit to our work, is the recent result by Paccagnan et al. (2021), whereby the authors leverage a tractable linear programming formulation to design optimal taxation mechanisms that utilize solely *local* information. Naturally, the corresponding values of the optimal price of anarchy they achieve are inferior to ours (here, we design optimal taxation mechanisms without any restrictions on what type of information we use), though the efficiency values derived therein are remarkably close to the optimal values obtained here. For example, for affine (resp. quadratic) congestion games, they achieve an optimal price of anarchy of 2.012 (resp. 5.101), to be compared with a value of 2 (resp. 5). This suggests that restricting the attention to taxation mechanisms utilizing solely local information is sufficient to match almost exactly the performance of the best polynomial time algorithm.

We conclude observing that similar questions have been considered for variants of the classical setup studied here. For example, Fotakis and Spirakis (2008) focuses on symmetric network congestion games, Harks et al. (2015), Hoefer et al. (2008), Jelinek et al. (2014) focus on taxing a subset of the resources.

Approximation algorithms. A number of polynomial time algorithms have been proposed for approximating the minimum social cost in congestion games and their network counterpart as discussed in Makarychev and Sviridenko (2018), Andrews et al. (2012), Harks et al. (2016) and references therein. The best known approximation is due to Makarychev and Sviridenko (2018) who use randomization to round the solution of a natural linear programming relaxation. They provide a general expression for the resulting approximation factor as a function of the allowable resource costs. Related works have also considered modifications of the classical setup: Harks et al. (2016) provide approximation algorithms for polymatroid congestion games, whereas Kumar et al. (2009) considers scheduling problems on unrelated machines. While the result of Makarychev and Sviridenko (2018) holds for the more general class of optimization problems with “diseconomy of scale”, the approximation ratios we obtain here always match or strictly improve upon theirs.

1.4. Comparison with nonatomic congestion games.

The continuous flow counterpart of congestion games, commonly referred to as the *nonatomic* model, and its corresponding Wardrop equilibrium were proposed by Wardrop (1952) for his studies on road-traffic routing. The two models are identical in spirit, except that the nonatomic variant postulates the presence of infinitely many agents, each controlling an infinitesimal portion of the traffic flow. Albeit apparently minor, this modification has significant impact both on the applicability of the model and on the corresponding analysis.

Regarding the model’s applicability, nonatomic congestion games require each agent’s decision to have negligible impact on the congestion levels. While this simplifying assumption is acceptable in many settings, there are many others for which this is simply not the case, e.g., scheduling, power routing, sensor allocation and many more.⁸

Regarding the analysis, nonatomic congestion games provide a platform that is, in general, easier to investigate. Specifically, the problem of minimizing the social cost can be formulated as a convex optimization program (under standard semi-convexity assumptions on the resource costs), and thus efficiently solved to optimality even for very large networks. A commonly employed approach is based on the Frank-Wolfe algorithm which leverages efficient algorithms for shortest-path calculations (Geisberger et al. 2012). Contrary to that, in the atomic case, minimizing the social cost is NP-hard and remains NP-hard to approximate within a factor that we give in (4). A similar conclusion holds when turning attention to the design of optimal taxation mechanisms, owing to the uniqueness of the equilibrium flows in the nonatomic model. For this setting, there are well-known taxation mechanisms, referred to as marginal cost taxes, that can be efficiently-computed and incentivize optimal behavior (Pigou 1920). The question is significantly more challenging in the atomic case, where a mechanism is confronted with a multiplicity of equilibria. Here, the same marginal cost taxes do not improve - and instead can deteriorate - the resulting system efficiency (Paccagnan et al. 2021).⁹ Nonetheless, our results show that efficiently-computable taxation mechanisms can still provide an approximation equal to that obtained by the best centralized polynomial time algorithm. Further, they show that such an approximation is near optimal for interesting classes of problems where the nonatomic model is typically utilized (see Figure 1).

Beyond the above, it is important to highlight the connection between non-atomic congestion games and the limiting behavior of atomic congestion games with large number of agents. Specifically, Haurie and Marcotte (1985) first, Paccagnan et al. (2018) and Cominetti et al. (2023) more

⁸ Already within the road-traffic realm, this assumption needs to be carefully considered. For example, in settings where few vehicles can impact the congestion levels significantly (e.g., in a city centre).

⁹ While this statement might feel as a contradiction of the results on the convergence of Nash equilibria to Wardrop equilibria for large number of agents (see following paragraph), we stress that, for such convergence results to hold, each agent needs to have infinitesimal impact on the congestion in the limit. However, there are problems settings in which this assumption is not satisfied. In these cases, i.e., when the number of agents grow but their impact on the congestion levels does not vanish, (Paccagnan et al. 2021) shows that employing the marginal cost mechanisms yields a worse price of anarchy than that obtained without levying *any* toll.

recently, show that the set of Nash equilibria of an atomic congestion game converges to the set of Wardrop equilibria of its nonatomic counterpart if the number of agents grow large and each of them has diminishing impact on the congestion. It is important to highlight, however, that i) the assumption of diminishing impact is not always satisfied (see above and Footnote 8), ii) estimates on the distance between Nash and Wardrop equilibria are available only under further restrictive assumptions on the resource costs, and such estimates depend heavily on the number of agents, the parameters and the structure of the game. As a result, convergence may require a very large number of agents, so that the nonatomic model should be carefully used as a surrogate for its atomic counterpart.¹⁰ However, when this is the case, marginal cost tolls provide an effective way for minimizing social cost.

Finally, it is also worth observing that the price of anarchy in both the non-atomic and atomic settings has been shown to converge to one when the network is fixed but the number of players is increased arbitrarily, see (Colini-Baldeschi et al. 2020) and (Wu et al. 2022) respectively. Stated differently, these results show that when the congestion is extremely high, equilibrium allocations are optimal, and no room is left for improving the social cost regardless of which mechanism is used.

1.5. Roadmap

The remainder of the manuscript is organized as follows. In Section 2 we prove Theorem 1 (hardness of approximation), and Corollaries 1 and 2 (hardness factor for polynomial resource costs). In Section 3 we state and prove Theorem 3, an enriched version of Theorem 2 whereby the expression of the taxation mechanism is also given. A discussion on future research directions concludes the manuscript.

2. NP-hardness of approximation

In this section we prove Theorem 1, i.e., we show that approximating the minimum social cost below the factor ρ_ℓ defined in (4) is NP-hard already for congestion games where all resource costs are identical to ℓ . The proof is based on a reduction from *gap label cover*, where we make use of a generalization of Feige’s partitioning system (Feige 1998). We proceed as follows: In Section 2.1 we introduce *gap label cover* and, independently, the partitioning system. In Section 2.2 we present the reduction and in Section 2.3 prove Theorem 1. Section 2.4 shows that $\sup_{\ell \in \mathcal{L}} \rho_\ell$ reduces to the $(d+1)$ ’st Bell number when resource costs are polynomials of maximum degree d , as claimed in Corollary 1. Throughout, we use $[m]$ to denote the set $\{1, \dots, m\}$ and $\text{Bin}(h, p)$ to denote a Binomial distribution with parameters h, p .

2.1. Background tools

2.1.1. Label Cover We start by introducing `GapLabelCover`, a commonly utilized NP-hard problem to obtain tight inapproximability results. We employ the weak-value formulation of the problem, implicitly used in Feige (1998) and also defined in Dudycz et al. (2020). Informally, a `GapLabelCover` problem, we are given a bipartite graph and a set of colors which we must use to color the left vertices of the graph. Additionally, we are given a function per each edge mapping the color chosen for the left incident vertex to a color for the right incident vertex. The following paragraph formalizes the setting, while Fig. 2 provides an example. As common in the literature, we refer to the set of colors as to the alphabet, and to a color as to a label in the alphabet.

DEFINITION 1. A `LabelCover` instance is described by a tuple $(L, R, E, h, [\alpha], [\beta], \{\pi_e\}_{e \in E})$, where

¹⁰ For example, (Cominetti et al. 2023, Thm 6.2) shows that strictly increasing resource costs are necessary to provide such bounds. Under this (and other milder) assumptions they show that, the distance between Nash and Wardrop equilibria decreases as $\theta \cdot n^{-1/2}$, when agents influence decreases as $1/n$. Here, θ crucially depends on the Lipschitz constant of the resource costs, the total traffic flow, and the size of the strategy space.




Figure 2 Left panel: instance of **LabelCover** described by a bi-regular bipartite graph with $|L| = 3$, $|R| = 1$, $h = 3$. Here, the left and right alphabets have equal size $\alpha = \beta = 2$ and are represented by the colors {red, blue}. The constraint functions π_e are visualized over the corresponding edges, and map any possible choice of color for the left incident vertex into a color for the right incident vertex. Central panel: instance of **LabelCover** and coloring choice that strongly satisfies the right vertex. This is because the choice of left colors maps, through the edge constraints, to the same color on the right vertex. Right panel: instance of **LabelCover** and coloring choice that weakly (but not strongly) satisfies the right vertex. This is because there exists two left vertices (but not all) whose colors match, once mapped to the corresponding right vertex.

- L and R are sets of left and right vertices of a bi-regular bipartite graph with edge set E and right degree h (i.e., the degree of all vertices in R equals h),
- $[\alpha]$ and $[\beta]$ represent left and right alphabets, and
- for every edge $e \in E$, a constraint function $\pi_e : [\alpha] \rightarrow [\beta]$ maps left labels to right labels.

Given a left labeling $\mathcal{L} : L \rightarrow [\alpha]$, i.e., a map that associates every left vertex to a label, we say that a right vertex $u \in R$ is

- *strongly satisfied* if for every pair of neighbors $v, v' \in L$ of u it is $\pi_{(v,u)}(\mathcal{L}(v)) = \pi_{(v',u)}(\mathcal{L}(v'))$;
- *weakly satisfied* if there exist two distinct neighbors $v, v' \in L$ of u , s.t. $\pi_{(v,u)}(\mathcal{L}(v)) = \pi_{(v',u)}(\mathcal{L}(v'))$.

Stated differently, given an assignment of colors to the left vertices (i.e., a labeling), a right vertex $u \in R$ is weakly satisfied if there exists two edges incident to that vertex u whose corresponding right color matches. If this property holds for all edges incident to u , then the right vertex is strongly satisfied. Depending on whether we can find a coloring of the left vertices satisfying all or a fraction of the right vertices, we then distinguish between a **YES** and a **NO** instance, as formalized next.

DEFINITION 2. For any $\delta > 0$, $h \in \mathbb{N}$ let $\text{GapLabelCover}(\delta, h)$ denote the following problem: Given a **LabelCover** instance $(L, R, E, h, [\alpha], [\beta], \{\pi_e\}_{e \in E})$, distinguish between

- **YES**: there exists a labeling that strongly satisfies all right vertices;
- **NO**: no labeling weakly satisfies more than a fraction δ of the right vertices.

Distinguishing between a **YES** and a **NO** instance is known to be difficult (formally **NP-hard**) when the size of the right alphabet, i.e., the number of possible right colors, is sufficiently large as recalled next. This result constitutes the basis for our hardness of approximation.

PROPOSITION 1 (Feige (1998), Dudycz et al. (2020)). $\forall \delta > 0$, $h \in \mathbb{N}$, $h \geq 2$, and β sufficiently large (depending on δ, h), $\text{GapLabelCover}(\delta, h)$ is **NP-hard**.

2.1.2. Partitioning System Similarly to Barman et al. (2021) and Dudycz et al. (2020), we generalize a combinatorial object introduced by Feige (1998), called *partitioning system*, which we also equip with a cost function. This object will be used to construct the agents allocation sets $\{\mathcal{A}_i\}_{i=1}^N$ in our ensuing reduction. In particular, it will ensure that, in our reduction, agents have allocations that are either optimal or incur a very high social cost, and the ratio between these two costs approaches the inapproximability factor we wish to achieve. A general illustration is provided in Fig. 3 while a concrete example is given in Fig. 4.

DEFINITION 3. Given a ground set of elements $[n]$, integers β, h, k such that $kn/h \in \mathbb{N}$, $\beta \geq h \geq k$, a cost function $c: \mathbb{N} \rightarrow \mathbb{R}_{>0}$, and $\eta > 0$, a partitioning system with parameters (n, β, h, k, η) is a collection of partitions $\mathcal{P}_1, \dots, \mathcal{P}_\beta$ of $[n]$ such that:

- P1) Every partition \mathcal{P}_j is a collection of subsets $P_{j,1}, \dots, P_{j,h} \subseteq [n]$ each with kn/h elements and such that each element from $[n]$ is selected by k sets in $P_{j,1}, \dots, P_{j,h}$. Observe that, for any \mathcal{P}_j we have $|\mathcal{P}_j| = h$, and the above implies

$$\sum_{r \in [n]} c(|\mathcal{P}_j|_r) = c(k)n. \quad (5)$$

Here, $|\mathcal{P}_j|_r$ denotes the number of sets in the collection \mathcal{P}_j to which element r belongs, where we extended the definition of the function c to include $c(0) = 0$. This allows to simplify notation in (5) by summing over all resources $r \in [n]$, as opposed to summing over $r \in [n]$ s.t. $|\mathcal{P}_j|_r \geq 1$.

- P2) for any $B \subseteq [\beta]$ with $|B| = h$ and for any function $i: [\beta] \rightarrow [h]$, let $Q = \{P_{j,i(j)}, j \in B\}$. It is

$$\sum_{r \in [n]} c(|Q|_r) \geq (\mathbb{E}_{X \sim \text{Bin}(h, k/h)}[c(X)] - \eta) n. \quad (6)$$

To gain some intuition on properties P1 and P2, we provide a graphical representation of a partitioning system in Fig. 3, whereby each box contains kn/h elements from the ground set $[n]$. Property 1 asserts that every time we select an entire row, we are guaranteed to cover every element in $[n]$ precisely k times. Property 2 asserts that every time we select one and only one set from each row, for a total of h rows, many sets cover the same elements so that the resulting cost is high. A concrete example of partitioning system is provided in Fig. 4.

At this stage, we recall that the probability mass function of the binomial distribution $\text{Bin}(h, k/h)$ converges pointwise for fixed $k \geq h$ to the probability mass function of the Poisson distribution $\text{Poi}(k)$ as h grows large (Durrett 2010). Hence, we informally observe that, when $\rho_\ell < \infty$,

$$\frac{\sum_{r \in Q} c(|Q|_r)}{\sum_{r \in \mathcal{P}_j} c(|\mathcal{P}_j|_r)} = \frac{\mathbb{E}_{X \sim \text{Bin}(h, k/h)}[c(X)] - \eta}{c(k)} \xrightarrow{h \rightarrow \infty} \frac{\mathbb{E}_{X \sim \text{Poi}(k)}[c(X)] - \eta}{c(k)},$$

see Lemma EC.1 in Section EC.1.1 for a proof. If we choose $c(x) = x\ell(x)$, let $h \rightarrow \infty$ and consider the worst case over k , this ratio precisely matches the inapproximability result we aim to derive




Figure 3 A partitioning system with parameters (n, β, h, k, η) . Each box contains kn/h elements from the ground set $[n]$. Property P1 ensures that selecting an entire row results in low cost (left, in green). Property P2 ensures that selecting one and only one box per each row, for a total of h rows, results in a high cost (right, in red).




Figure 4 Examples of partitioning systems obtained with the choice of $c(x) = x^2$, $k = 1$, for increasing values of $\beta = h = 2, 3, 4$ representing the number of rows and columns. Here, each labeled circle represents a resource from the ground set. To verify that (a) is indeed a partitioning system it suffices to observe that, when selecting any entire row, each resource is covered only once (since $k = 1$), so that we incur a cost of $4 \cdot c(1) = 4$ (Property 1), while when selecting one box from each row we incur a cost of $c(1) + 2c(2) + c(1) = 6 = \mathbb{E}_{X \sim \text{Bin}(2, 1/2)}[X^2]$ (Property 2). Analogously for cases (b) and (c). Note how, in these examples, we have been able to construct partitioning systems where the error parameter η entering (6) is identically zero. Finally, we observe that partitioning systems obtained with $c(x) = x^2$, $k = 1$ correspond, in the ensuing reduction, to the special case of congestion games with linear resource cost $\ell(x) = x$.

(cfr. the previous expression and ρ_ℓ in (4)). We are thus left to piece these elements together in the ensuing section. Before doing so, we remark that partitioning systems do exist for every choice of $\eta > 0$ as long as n is taken sufficiently large as stated in the next proposition. Its proof follows the same approach of that in (Barman et al. 2021), and is included in Section EC.1.2 for completeness. We remark that, when used in the upcoming reduction, we will be able to compute a partitioning systems in time that is independent on the size of the instance we reduce from.

PROPOSITION 2. *Let $c : \mathbb{N} \rightarrow \mathbb{R}_{>0}$ non-decreasing be given. For every choice of $\beta \geq h \geq k$ integers with $kn/h \in \mathbb{N}$, $\eta \in (0, 1)$, and $n \geq \frac{c(h)^2}{2\eta^2} [\log(10) + \beta \log(h+1)]$ a partitioning system with parameters (n, β, h, k, η) and cost function c exists. It can be found in time depending solely on h , n and β .*

2.2. Reduction

We first provide the reduction and prove the hardness result in the case when $\rho_\ell < \infty$. We consider the case $\rho_\ell = \infty$ separately at the end of Section 2.3.

Starting from the resource cost function ℓ and a fixed $\varepsilon > 0$, we first construct a partitioning system with appropriately defined parameters. We then reduce an instance of **LabelCover** $\mathcal{C} = (L, R, E, h, [\alpha], [\beta], \{\pi_e\}_{e \in E})$ to an instance of congestion game $G = (N, \mathcal{R}, \{\mathcal{A}_i\}_{i=1}^N, \{\ell_r\}_{r \in \mathcal{R}})$ with identical resource cost. The idea is to define G by creating a copy of the partitioning system for every right vertex and use that to define the players' allocation sets \mathcal{A}_i .

Construction of the partitioning system. Given ℓ and $\varepsilon > 0$, we construct a partitioning system setting its parameters as follows:

- Let k be the maximizer of (4);¹¹

¹¹ For ease of exposition, we show the result when the supremum is attained at some value $k \in \mathbb{N}$. If this is not the case, then the supremum must be achieved at $k \rightarrow \infty$. One then fixes $k \in \mathbb{N}$ and proceeds with the reasoning as is. This will give rise to an additional error term $\nu(k)$ with $\lim_{k \rightarrow \infty} \nu(k) = 0$ in the ensuing Equation (7), where the right hand side will be replaced by $(\rho_\ell + \nu(k))c(k)$. Nevertheless one can select k sufficiently large and control such error to a desired accuracy.

- Let c be the cost function defined from ℓ as in $c(x) = x\ell(x)$ for all $x \in \mathbb{N}$ and $c(0) = 0$;
- Let $\eta < \min\{\varepsilon c(k)/4, 1\}$;
- Let $h \geq k$ integer be such that $|\mathbb{E}_{X \sim \text{Bin}(h, k/h)}[X\ell(X)] - \mathbb{E}_{X \sim \text{Poi}(k)}[X\ell(X)]| \leq \varepsilon c(k)/4$, which exists thanks to the convergence result in Lemma EC.1 in Section EC.1.1;
- Let $\delta \leq \varepsilon/(2\rho_\ell)$ and fix β large enough to ensure hardness of $\text{GapLabelCover}(\delta, h)$ in Proposition 1;
- Let n be so that $kn/h \in \mathbb{N}$ and n is large enough to have existence, from Proposition 2, of the partitioning system with parameters (n, β, h, k, η) and cost function c .

Owing to Proposition 1, since β, h, k, η are now fixed, a partitioning system can be computed in time *independent* of the size of the `LabelCover` instance we wish to reduce from.

Reduction of LabelCover instance to congestion game instance. Now we take an instance of `LabelCover` $\mathcal{C} = (L, R, E, h, [\alpha], [\beta], \{\pi_e\}_{e \in E})$ where h, β are defined above and thus Proposition 1 (NP-hardness) holds. For each right vertex $u \in R$ we use the local partitioning system with parameters (n, β, h, k, η) . We refer to the resources in the partitioning system corresponding to the right vertex $u \in R$ with $\{1^u, \dots, n^u\}$. Similarly we use $\mathcal{P}_j^u = \{P_{j,1}^u, \dots, P_{j,h}^u\}$ for the local partitions. The congestion game $G = (N, \mathcal{R}, \{\mathcal{A}_i\}_{i=1}^N, \{\ell_r\}_{r \in \mathcal{R}})$ is defined as follows

- each left vertex corresponds to an agent, so that the number of agents is $N = |L|$;
- the ground set of resources is the union of the resources introduced by each local partitioning system on every right vertex, i.e., $\mathcal{R} = \cup_{u \in R} \{1^u, \dots, n^u\}$;
- each resource cost is equal to ℓ , i.e., $\ell_r(x) = \ell(x)$ for all $r \in \mathcal{R}, x \in \mathbb{N}$;
- as each left vertex $v \in L$ corresponds to one and only one agent $i \in [N]$, we refer to a left vertex as $i \in [N]$ instead of as $v \in L$ to ease the notation. For agent $i \in [N]$ we construct each pure strategy $a_i \in \mathcal{A}_i$ as follows. We let the left vertex i select a label $l \in [\alpha]$, and correspondingly take the union over all right vertices $u \in \mathcal{N}(i)$ neighboring with i , of the resources belonging to the block $P_{j,i}^u$, where $j = \pi_{(i,u)}(l)$. Repeating over all left labels we obtain the strategy set \mathcal{A}_i . Formally

$$\mathcal{A}_i = \left\{ \cup_{u \in \mathcal{N}(i)} P_{j,i}^u, \quad \text{where } j = \pi_{(i,u)}(l), \quad \forall l \in [\alpha] \right\}.$$

Figure 5 exemplifies the construction. We conclude remarking that the above procedure implicitly defines a map associating a profile of left labels (l_1, \dots, l_N) (one per each left vertex) to an allocation $(a_1, \dots, a_N) \in \mathcal{A}$, and that spanning through all possible choices of (l_1, \dots, l_N) produces all possible allocations in \mathcal{A} . This observation will be useful in proving the hardness result.

2.3. Proof of the result

As anticipated, we first prove the result for the case of $\rho_\ell < \infty$. At the end of this section, we turn the attention to $\rho_\ell = \infty$. For any given instance of `LabelCover` $\mathcal{C} = (L, R, E, h, [\alpha], [\beta], \{\pi_e\}_{e \in E})$, resource cost ℓ , and $\varepsilon > 0$, we consider an instance of congestion game $G = (N, \mathcal{R}, \{\mathcal{A}_i\}_{i=1}^N, \{\ell_r\}_{r \in \mathcal{R}})$ constructed as in the previous section. We will now show that

- *completeness* (Section 2.3.1): If the instance \mathcal{C} is a YES, then $\min_{a \in \mathcal{A}} SC(a) \leq n|R|c(k)$,
- *soundness* (Section 2.3.2): If the instance \mathcal{C} is a NO, then $\min_{a \in \mathcal{A}} SC(a) > (\rho_\ell - \varepsilon)n|R|c(k)$.

An algorithm solving `MinSC` with an approximation ratio smaller than $\rho_\ell - \varepsilon$ will be able to distinguish between YES/NO instances of an NP-hard promise problem. This will conclude the proof.

2.3.1. Completeness We intend to show that if \mathcal{C} is a YES instance, then $\min_{a \in \mathcal{A}} SC(a) \leq n|R|c(k)$. This follows readily. In fact, if \mathcal{C} is a YES instance, there exists a labeling that strongly satisfies all right vertices. This implies that there exists an allocation $a^* \in \mathcal{A}$ whereby, for any given right vertex, all neighboring left vertices (agents) have selected blocks belonging to an entire row of the corresponding partitioning system. Thanks to property P1 the cost of the allocation a^* on every local partitioning system is equal to $nc(k)$. Since the total cost is additive over the local partitioning systems, we obtain the result

$$SC(a^*) = n|R|c(k) \implies \min_{a \in \mathcal{A}} SC(a) \leq n|R|c(k).$$




Figure 5 Given a label cover instance $\mathcal{C} = (L, R, E, h, [\alpha], [\beta], \{\pi_e\}_{e \in E})$, our reduction associates every left vertex in L to an agent in the game G . Here we exemplify how the action set \mathcal{A}_i is generated for agent $i \in L$. To ease the presentation, we consider a left alphabet of size 2 and use **blue** and **orange** to identify the left labels. Since i has two right neighbors, u and u' , we construct two partitioning systems with ground set of resources $\{1^u, \dots, n^u\}$ and $\{1^{u'}, \dots, n^{u'}\}$. Constraints $\pi_{(i,u)}(\text{blue}) = 1$, $\pi_{(i,u)}(\text{orange}) = 2$ and $\pi_{(i,u')}(\text{blue}) = \beta$, $\pi_{(i,u')}(\text{orange}) = 1$ are given, and we represent them graphically with the fact that on the left partitioning system, the label **blue** (resp. **orange**) is associated to the block in row 1 = $\pi_{(i,u)}(\text{blue})$ (resp. row 2 = $\pi_{(i,u)}(\text{orange})$). Similarly for the right partitioning system using $\pi_{(i,u')}$ to determine the row. The set \mathcal{A}_i is readily constructed as $\mathcal{A}_i = \{P_{1,2}^u \cup P_{\beta,1}^{u'}, P_{2,2}^u \cup P_{1,1}^{u'}\}$, where the first (resp. second) allocation corresponds to a **blue** (resp. **orange**) left label.

2.3.2. Soundness We intend to show that if \mathcal{C} is a NO instance, then $\min_{a \in \mathcal{A}} \text{SC}(a) \geq (\rho_\ell - \varepsilon)n|R|c(k)$ which is equivalent to showing $\text{SC}(a) \geq (\rho_\ell - \varepsilon)n|R|c(k)$ for all $a \in \mathcal{A}$. Towards this goal, we build upon the last observation presented in Section 2.2, i.e., the fact that our construction associates each profile of left labels to an allocation, and that spanning through all possible choices of (l_1, \dots, l_N) produces all possible allocations in \mathcal{A} . Hence, it suffices to prove the desired property by considering all possible combinations of profiles (l_1, \dots, l_N) and the corresponding induced cost, instead of considering all $a \in \mathcal{A}$. Since \mathcal{C} is a NO instance, for any possible choice of (l_1, \dots, l_N) , no more than δ fraction of the right vertices are weakly satisfied. Owing to property P2, each partitioning system corresponding to a non weakly satisfied right vertex has a cost larger than $(\mathbb{E}_{X \sim \text{Bin}(h, k/h)}[c(X)] - \eta)n$. Thus, since at least $(1 - \delta)|R|$ right vertices are not weakly satisfied, it is

$$\text{SC}(a) \geq (\mathbb{E}_{X \sim \text{Bin}(h, k/h)}[c(X)] - \eta)(1 - \delta)n|R|, \quad \forall a \in \mathcal{A}.$$

We conclude with some cosmetic manipulation. In particular, we recall that the binomial distribution converges to the Poisson distribution when the number of trials grows large and the success probability of each trial goes to zero (Durrett 2010). In our settings, this corresponds to the fact that the probability mass function of $\text{Bin}(h, k/h)$ converges pointwise for fixed k to that of $\text{Poi}(k)$ as $h \rightarrow \infty$.

Since $\rho_\ell < \infty$, we observe that

$$\lim_{h \rightarrow \infty} \mathbb{E}_{X \sim \text{Bin}(h, k/h)}[c(X)] = \mathbb{E}_{X \sim \text{Poi}(k)}[c(X)] = \rho_\ell c(k), \quad (7)$$

where equality holds thanks to Lemma EC.1 in Section EC.1.1. This implies the existence of a function $\theta(h)$ with $\theta(h) \rightarrow 0$ as $h \rightarrow \infty$ allowing to control the error, and for which $\mathbb{E}_{X \sim \text{Bin}(h, k/h)}[c(X)] \geq$

$\rho_\ell c(k) - \theta(h)$. In other words the LHS can be made arbitrarily close to $\rho_\ell c(k)$ by selecting h sufficiently large. Hence,

$$\begin{aligned}
 \text{SC}(a) &\geq (\rho_\ell c(k) - \theta(h) - \eta)(1 - \delta)n|R| \\
 &= \left[\rho_\ell - \frac{\theta(h)}{c(k)} - \frac{\eta}{c(k)} - \left(\rho_\ell - \frac{\theta(h)}{c(k)} - \frac{\eta}{c(k)} \right) \delta \right] n|R|c(k) \\
 &\geq \left[\rho_\ell - \frac{\theta(h)}{c(k)} - \frac{\eta}{c(k)} - \rho_\ell \delta \right] n|R|c(k) \\
 &> \left[\rho_\ell - \frac{\varepsilon}{4} - \frac{\varepsilon}{4} - \frac{\varepsilon}{2} \right] n|R|c(k) = (\rho_\ell - \varepsilon)n|R|c_k.
 \end{aligned} \tag{8}$$

The last inequality holds by the choice of parameters, ensuring that $\frac{\theta(h)}{c(k)} \leq \frac{\varepsilon}{4}$, $\frac{\eta}{c(k)} < \frac{\varepsilon}{4}$, $\rho_\ell \delta \leq \frac{\varepsilon}{2}$.

Case of $\rho_\ell = \infty$. As anticipated, we treat the case of unbounded ρ_ℓ separately. Towards this goal, we follow the same reduction of Section 2.2, with minor modification on the choice of parameters. We replace ρ_ℓ with a fixed (and conceptually large) constant M . Since $\rho_\ell = \infty$, we note that $\mathbb{E}_{X \sim \text{Poi}(k)}[c(X)]/c(k)$ is unbounded at some k (possibly infinity). Since the probability mass functions of $\text{Bin}(h, k/h)$ and $\text{Poi}(k)$ converge, we can choose the pair h and k so that $\mathbb{E}_{X \sim \text{Bin}(h, k/h)}[c(X)] \geq M c(k)$. Finally, we set $\delta \leq \varepsilon/(2M)$. One then follows the same proof as in the case of bounded ρ_ℓ , whereby (8) is replaced with $\text{SC}(a) \geq (M - \eta/c(k))(1 - \delta)n|R|c(k) \geq (M - \frac{\eta}{c(k)} - M\delta)n|R|c(k) \geq (M - \varepsilon/4 - \varepsilon/2)n|R|c(k) > (M - \varepsilon)n|R|c(k)$. Since M can be taken to be arbitrarily large, the problem is NP-hard to approximate within any finite ratio. Q.E.D.

2.4. Hardness factor for polynomial resource cost

Corollary 1 claims that, when resource costs are obtained by non-negative combinations of monomials of maximum degree $d > 0$, MinSC is hard to approximate within any factor smaller than the $(d+1)$ 'st Bell number. Note that this is a direct consequence of Theorem 1 and its ensuing discussion, which applies since for $d \geq 0$, each monomial x^d is positive, non-decreasing, semi-convex for $x \in \mathbb{N}$. In this setting, MinSC is NP-hard to approximate within any factor smaller than $\sup_{\ell \in \mathcal{L}} \rho_\ell$, where \mathcal{L} contains all polynomials of maximum degree d with non-negative coefficients. Interestingly, Lemma EC.2 in Section EC.1.3 shows that the factor ρ_ℓ is maximized amongst all polynomials in \mathcal{L} by the monomial $m(x) = x^d$. Hence, for polynomial congestion games, it is

$$\sup_{\ell \in \mathcal{L}} \rho_\ell = \rho_m = \sup_{x \in \mathbb{N}} \frac{\mathbb{E}_{P \sim \text{Poi}(x)}[P^{d+1}]}{x^{d+1}} = \sup_{x \in \mathbb{N}} \sum_{i=0}^{d+1} x^{i-(d+1)} \begin{Bmatrix} d+1 \\ i \end{Bmatrix} = \sum_{i=0}^{d+1} \begin{Bmatrix} d+1 \\ i \end{Bmatrix} = \mathcal{B}(d+1).$$

The first equality follows from the fact that $m(x) = x^d$ maximizes the ratio ρ_ℓ , while the second from the definition of ρ_m . The third equality follows from the fact that the $(d+1)$ 'st moment of the Poisson distribution $\text{Poi}(x)$ equals $\sum_{i=0}^{d+1} x^i \begin{Bmatrix} d+1 \\ i \end{Bmatrix}$, where $\begin{Bmatrix} d+1 \\ i \end{Bmatrix}$ is a Stirling number of the second kind (Mansour and Schork 2015, p. 63). The fourth equality holds because each function $x^{i-(d+1)}$ is non-increasing, owing to $i \leq (d+1)$, and thus the supremum is attained at $x = 1$. The last one is due to the definition of the $(d+1)$ 'st Bell number, which we denote with $\mathcal{B}(d+1)$, see (Mansour and Schork 2015, Eq. 1.2).

One can repeat a similar reasoning also when $d \geq 0$ is not integer and show that the expression inside the supremum is non-increasing in $x \geq 1$, e.g., by computing its derivatives. It follows that the supremum is attained at $x = 1$, and the definition of ρ_ℓ gives

$$\rho_\ell = \frac{1}{e} \sum_{i=0}^{\infty} \frac{i^{d+1}}{i!}. \tag{9}$$

The latter expression is sometimes referred to as the fractional Bell number. Note that using Dobiński's formula (Mansour and Schork 2015, Eq. 1.25), one recovers the Bell number $\mathcal{B}(d+1)$ when $d \in \mathbb{N}_0$.

Corollary 2 provides a more refined analysis for affine congestion games where $\ell_r(x) = \alpha_r x + \beta_r$, when an upper bound q on $\max_r \alpha_r / (\alpha_r + \beta_r)$ is given. The proof of Corollary 2 follows readily from the application of Theorem 1 and its ensuing discussion. Indeed, here one can take the set \mathcal{L} to contain all linear cost functions of the form $\ell(x) = \alpha x + \beta$ where $\alpha / (\alpha + \beta) \leq q$. One can then easily compute ρ_ℓ explicitly from (4)

$$\rho_\ell = \sup_{x \in \mathbb{N}} \frac{\mathbb{E}_{P \sim \text{Poi}(x)}[P(\alpha P + \beta)]}{x(\alpha x + \beta)} = \sup_{x \in \mathbb{N}} \frac{\alpha(x^2 + x) + \beta x}{\alpha x^2 + \beta x} = 1 + \frac{\alpha}{\alpha + \beta},$$

and thus obtain the desired hardness factor, i.e., $\sup_{\ell \in \mathcal{L}} \rho_\ell = 1 + q$.

3. Taxes achieve optimal approximation

In this section we show how to compute a taxation mechanism whose price of anarchy matches the hardness factor. Since taxation mechanisms can be utilized to derive polynomial time algorithms with an approximation factor matching the price of anarchy (Section 1.2), in the ensuing Section 3.1 we compare the optimal price of anarchy with the best known polynomial approximation of Makarychev and Sviridenko (2018).

We start by introducing a parameterized family of taxation mechanisms for which we will later provide (efficiently computable) parameters that achieve the desired result. Our taxation mechanisms take as input a congestion game G , with resource costs belonging to a common set \mathcal{L} . Toward this goal, for each $\ell_r : \mathbb{N} \rightarrow \mathbb{R}_{>0}$ given, we extend its definition to $\ell_r : \mathbb{N}_0 \rightarrow \mathbb{R}_{\geq 0}$ by setting $\ell_r(0) = 0$. This is without loss of generality and merely needed to ease the notation (see Page 5 and Footnote 6). For similar reasons, we define $c_r : \mathbb{N}_0 \rightarrow \mathbb{R}_{\geq 0}$ as $c_r(x) = x\ell_r(x)$ for $x \in \mathbb{N}$ and $c_r(0) = 0$. Finally, we associate each resource cost $\ell_r(x)$ with the function $p_r : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ defined as

$$p_r(v) = \mathbb{E}_{P \sim \text{Poi}(v)}[P\ell_r(P)] = \left(\sum_{i=0}^{\infty} i\ell_r(i) \frac{v^i}{i!} \right) e^{-v}, \quad (10)$$

which we use to introduce the following set of parametrized mechanisms.

DEFINITION 4 (PARAMETERIZED TAXATION MECHANISMS). Given a parameter vector $(v_r)_{r \in \mathcal{R}}$ with $v_r \in \mathbb{R}_{\geq 0}$ and a set of resources \mathcal{R} with costs $\ell_r(x)$, define a parameterized taxation function $\tau_r : \mathbb{N}_0 \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$, with $\tau_r(x, v_r) = f_r(x, v_r) - \ell_r(x)$ where $f_r(x, 0) = \ell_r(x)$, $f_r(0, v) = 0$, and

$$f_r(x, v) = \frac{(x-1)!}{v^x} \sum_{i=0}^{x-1} \frac{p_r(v) - i\ell_r(i)}{i!} v^i, \quad x \in \mathbb{N}_0, v \in \mathbb{R}_{>0}. \quad (11)$$

Observe that the previous taxation mechanism effectively substitutes the original resource costs $\ell_r(x)$ with the new resource costs $f_r(x, v_r)$ given in (11). Indeed, when taxes are factored in, the resource cost perceived by each agent on resource r is $\ell_r(x) + \tau_r(x, v_r) = f_r(x, v_r)$. The following lemma provides three important properties of the above class of taxation mechanisms. In particular, Lemma 1 ensures that taxes are non-negative, that modified resource costs introduced in (11) are non-decreasing and that they satisfy a recursion which will be crucial later on. The proof can be found in Section EC.2.2.

LEMMA 1. *For all $v \in \mathbb{R}_{\geq 0}$, $x \in \mathbb{N}_0$, and $r \in \mathcal{R}$ the taxation mechanism introduced in Definition 4 satisfies:*

- (a) $\tau_r(x, v) \geq 0$,
- (b) $f_r(x+1, v) \geq f_r(x, v)$,
- (c) $x\ell_r(x) - xf_r(x, v) + vf_r(x+1, v) = p_r(v)$.

As we will see in the remainder of the section, the mechanism that optimizes the price of anarchy makes use of the parametrized taxation mechanism introduced in Definition 4, with parameters obtained by solving the following convex program in the variables y_i, v_r

$$\begin{aligned} \min \quad & \sum_{r \in \mathcal{R}} p_r(v_r) \\ \text{subject to} \quad & v_r = \sum_{i=1}^N \sum_{k \in [s_i]: r \in a_{i,k}} y_{i,k} \quad \text{for all } r \in \mathcal{R}, \\ & y_i \in \Delta(s_i) \quad \text{for all } i \in [N], \end{aligned} \tag{12}$$

where we let $s_i = |\mathcal{A}_i|$, $a_{i,k}$, denote the k -th action available to player i , and $\Delta(s)$ represent the s -th dimensional simplex. This program can be easily interpreted as the continuous relaxation of the original MinSC problem, where the resource costs ℓ_r have been replaced by the modified resource costs p_r previously defined.

We are now ready to state our main result of this section, which is an extension of Theorem 2. We state the result when $\sup_{\ell \in \mathcal{L}} \rho_\ell < \infty$, else the problem is inapproximable as seen in Theorem 1.

THEOREM 3. *Consider a congestion game where each resource cost belongs to a common set of functions \mathcal{L} satisfying Assumption 1. Let $\sup_{\ell \in \mathcal{L}} \rho_\ell < \infty$ and denote with $(\bar{y}_i)_{i \in [N]}, (\bar{v}_r)_{r \in \mathcal{R}}$ a solution of the convex program (12). Then:*

- *The taxation mechanism introduced in Definition 4 with parameter vector $(\bar{v}_r)_{r \in \mathcal{R}}$ has a price of anarchy no-higher than $\sup_{\ell \in \mathcal{L}} \rho_\ell$;*
- *Moreover, for any choice of $\varepsilon > 0$ one can design in polynomial time, through the approximate solution of (12), a taxation mechanism whose price of anarchy is no-higher than $\sup_{\ell \in \mathcal{L}} \rho_\ell + \varepsilon$.*

Proof. Given a congestion game G , we consider the corresponding program (12). Let us verify that (12) is indeed convex. Since the constraints are linear and the objective function is a sum of univariate functions p_r it suffices to show that each $p_r(v)$ is convex in v . This holds true as p_r is defined in (10) as the expectation of a convex function over a Poisson distributed random variable. For completeness we provide a proof of this fact in Lemma EC.3 in Section EC.2.1.

Let $(\bar{y}_i)_{i \in [N]}, (\bar{v}_r)_{r \in \mathcal{R}}$ be an optimal solution of the convex program (12) and consider the taxation mechanism from Definition 4 with parameter vector $(\bar{v}_r)_{r \in \mathcal{R}}$. Denote $\rho = \sup_{\ell \in \mathcal{L}} \rho_\ell$.

To complete the proof we will use a smoothness approach, with a crucial modification: Instead of comparing an action profile a (e.g., an equilibrium allocation) with another action profile a' (e.g., an optimal allocation), we will compare an action profile a against a mixed profile $y = (y_1, \dots, y_N) \in \Delta(s_1) \times \dots \times \Delta(s_N)$. Specifically, we will choose the mixed profile \bar{y} solving (12), and show that

$$\sum_{i=1}^N \sum_{k=1}^{s_i} \bar{y}_{i,k} [\bar{C}_i(a) - \bar{C}_i(a'_{i,k}, a_{-i})] \geq \text{SC}(a) - \rho \text{SC}(a^{\text{opt}}), \quad \forall a \in \mathcal{A}. \tag{13}$$

where $\bar{C}_i(a)$ denotes the modified cost function perceived by the agents $\bar{C}_i(a) = \sum_{r \in a_i} f_r(|a|_r, \bar{v}_r)$. Once (13) is shown, the desired bound on the price of anarchy follows readily for pure Nash equilibria and more generally extends all the way to coarse correlated equilibria (Roughgarden 2015). In the former case, substituting the profile a with any pure Nash equilibrium a^{ne} , and summing the equilibrium conditions $0 \geq \bar{C}_i(a^{\text{ne}}) - \bar{C}_i(a'_{i,k}, a^{\text{ne}}_{-i})$, one obtains $0 \geq \sum_{i=1}^N \sum_{k=1}^{s_i} \bar{y}_{i,k} [\bar{C}_i(a^{\text{ne}}) - \bar{C}_i(a'_{i,k}, a^{\text{ne}}_{-i})]$, so that

$$0 \geq \sum_{i=1}^N \sum_{k=1}^{s_i} \bar{y}_{i,k} [\bar{C}_i(a^{\text{ne}}) - \bar{C}_i(a'_{i,k}, a^{\text{ne}}_{-i})] \geq \text{SC}(a^{\text{ne}}) - \rho \text{SC}(a^{\text{opt}}),$$

from which one concludes. Since (13) holds for all $a \in \mathcal{A}$, the same bound on the price of anarchy holds for the much broader class of coarse correlated equilibria. To see this, let σ be any coarse correlated equilibrium over $\mathcal{A}_1 \times \cdots \times \mathcal{A}_N$, and consider the expected value of (13). Due to linearity of the expectation and the definition of coarse correlated equilibria, we have $0 \geq \mathbb{E}_{a \sim \sigma} \left[\sum_{i=1}^N \sum_{k=1}^{s_i} \bar{y}_{i,k} [\bar{C}_i(a) - \bar{C}_i(a'_{i,k}, a_{-i})] \right]$, from which one concludes.

We are thus left to prove the smoothness condition (13). Given an optimal allocation $a^{\text{opt}} \in \arg \min_{a \in \mathcal{A}} \text{SC}(a)$, let $v_r^{\text{opt}} = |a^{\text{opt}}|_r$. Inequality (13) follows from

$$\begin{aligned}
& \sum_{i=1}^N \sum_{k=1}^{s_i} \bar{y}_{i,k} [\bar{C}_i(a) - \bar{C}_i(a'_{i,k}, a_{-i})] = \sum_{i=1}^N \bar{C}_i(a) - \sum_{i=1}^N \sum_{k=1}^{s_i} \bar{y}_{i,k} \bar{C}_i(a'_{i,k}, a_{-i}) \\
(\bar{\ell}_r \text{ is non-decreasing by Lemma 1-b}) & \geq \sum_{i=1}^N \bar{C}_i(a) - \sum_{i=1}^N \sum_{k=1}^{s_i} \bar{y}_{i,k} \sum_{r \in a'_{i,k}} f_r(|a|_r + 1, \bar{v}_r) \\
(\text{changing order of summation}) & = \sum_{r \in a} [|a|_r f_r(|a|_r, \bar{v}_r) - \bar{v}_r f_r(|a|_r + 1, \bar{v}_r)] \\
(\text{recursion in Lemma 1-c}) & = \sum_{r \in a} [c_r(|a|_r) - p_r(\bar{v}_r)] \\
(\bar{v}_r \text{ optimal solution of (12)}) & \geq \sum_{r \in a} [c_r(|a|_r) - p_r(v_r^{\text{opt}})] \\
(\text{by def. of } \rho, \text{ it is } p_r(v_r^{\text{opt}}) \leq \rho c_r(v_r^{\text{opt}})) & \geq \sum_{r \in a} [c_r(|a|_r) - \rho c_r(v_r^{\text{opt}})] \\
& = \text{SC}(a) - \rho \text{SC}(a^{\text{opt}}),
\end{aligned} \tag{14}$$

thus completing the proof of the first claim.

We now turn attention to the second claim. Since (12) is a convex program with polynomially many decision variables and constraints, it can be solved to arbitrary precision in polynomial time. In this case, the argument in Eq. (14) will go through with a minor change on the fifth line, where one pays a multiplicative factor $1 + \xi$ with $\xi > 0$. Correspondingly, we obtain a price of anarchy of $(1 + \xi)\rho$ in place of ρ . Selecting ξ sufficiently small, one obtains a price of anarchy of $(1 + \xi)\rho \leq \rho + \varepsilon$ for any choice of $\varepsilon > 0$. Q.E.D.

We conclude observing that the expression of $p_r(v)$ in (10) can be computed analytically for many commonly studied classes of resource costs. Nevertheless, if this is not the case, one can approximate $p_r(v)$ to arbitrary precision by truncating the sum at a finite value. One can then verify the same properties shown in Lemma 1 and apply a reasoning identical to that in (14). An additional multiplicative error will arise in (14), although this can be made as small as desired.

3.1. Comparison with existing approximations

A strength of the approach followed thus far is that optimally designed taxes can be used to derive polynomial time algorithms matching the hardness factor. This can be done relying on existing algorithms, e.g., no-regret dynamics, as discussed in Section 1.2. For this reason, we now turn the attention to comparing the optimal price of anarchy of Theorem 3 with the best known approximation ratio of Makarychev and Sviridenko (2018). Specifically, when all resource costs are identical to $\ell : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$, Makarychev and Sviridenko (2018) give a randomized algorithm with an approximation ratio of $\mu_\ell = \sup_{x \in \mathbb{R}_{> 0}} \mathbb{E}_{P \sim \text{Poi}(1)}[(xP)\ell(xP)] / (x\ell(x))$. While their result applies to the broader class of optimization problems with a “diseconomy of scale”, the approximation ratio in Corollary 3 always matches or strictly improves upon theirs. This follows from

$$\rho_\ell = \sup_{x \in \mathbb{N}} \frac{\mathbb{E}_{P \sim \text{Poi}(x)}[P\ell(P)]}{x\ell(x)} \leq \sup_{x \in \mathbb{N}} \frac{\mathbb{E}_{P \sim \text{Poi}(1)}[(xP)\ell(xP)]}{x\ell(x)} \leq \mu_\ell.$$

The last inequality follows trivially by replacing \mathbb{N} with $\mathbb{R}_{>0}$ and using the definition of μ_ℓ , while the first inequality can be shown using the notion of convex ordering between distributions.¹² An example where the approximation ratios coincide is given by $\ell(x) = x^d$, in which case $\rho_\ell = \mu_\ell$ equal the $(d + 1)$ 'st Bell number, while an instance where the inequality is strict is provided by $\ell(x) = x + 1$, in which case $\rho_\ell = \sup_{x \in \mathbb{N}} (x + 2)/(x + 1) = 3/2 < 2 = \sup_{x \in \mathbb{R}_{>0}} (2x + 1)/(x + 1) = \mu_\ell$.

4. Discussion and Conclusions

Congestion games provide a fundamental framework to study decision-making in the presence of congestion effects. Prior to our work, little was known regarding the classical problem of minimizing the social cost in congestion games, both in the case where we allow for centralized decision-making and for the case where agents can be coordinated only indirectly through the use of interventions.

In these settings, we provide conclusive answers on what approximations can be achieved by efficient (i.e., polynomial time) algorithms, and what are the inherent computational limitations. Interestingly, our technical results show that *no* performance degradation arises when moving from centralized decision-making to the use of interventions. On the contrary, judiciously designed taxation mechanisms can be efficiently computed and achieve the same performance of the best centralized polynomial time algorithm. We achieve this result by providing a tight computational lower bound for the problem of minimizing the social cost, and by designing suitable taxation mechanisms with matching performance. We thus obtain polynomial time algorithms based on taxes matching the hardness factor.

There remain many opportunities for further work on interventions in congestion games. One important research direction is to shed further light on the interplay between (more general) interventions and the achievable performances. It is interesting to understand whether other approaches based on, e.g., information provisioning, are equally powerful. In other words, does the positive result obtained here for taxation mechanisms hold for other classes of interventions?

There are also a number of open questions arising from our work and possible refinements thereof. A first direction is that of considering the *weighted* version of congestion games, where agents increase the congestion of a resource depending on their weight. While the hardness result we obtained in Theorem 1 holds for this more general class of problems, the corresponding hardness factor can be strengthened (i.e., made larger) by considering a partitioning system where all agents share a common weight w , not necessarily equal to one. On the other hand, the question of designing taxation mechanisms with optimal performance is significantly more challenging, and we leave this open as a direction for future work. A second direction is that of considering the variant of *network congestion games*, whereby each strategy set \mathcal{A}_i is implicitly given as the set of paths connecting an origin/destination node in an underlying graph. On its own, this more succinct representation of the strategy space would only increase the computational complexity, but on the other hand the graph also imposes more structure. The results in Theorem 3 (design of optimal taxes) extend to network congestion games by replacing the constraint set in the convex program in (12) with the set of feasible flows on the graph. Similarly, Corollary 3 (polynomial time algorithms) also extends provided that one utilizes no-regret algorithms, such as Follow the Perturbed Leader, that do not require explicit description of all possible paths (which might be exponential in the size of the graph). On the contrary, our reduction in the proof of Theorem 1 (inapproximability of minimum cost) induces a general congestion game. The existence of a reduction to network congestion games remains open.

¹² To see this, we leverage the result in (Shaked and Shanthikumar 2007, Thm 3.A.36) with $X_i, Y \sim x\text{Poi}(1)$ independent, $a_i = 1/x$, and $i = 1, \dots, x \in \mathbb{N}$. Theorem 3.A.36 in (Shaked and Shanthikumar 2007) applies ensuring that $\sum_{i=1}^x a_i X_i \leq_{\text{cx}} Y$, where \leq_{cx} denotes the convex ordering between distributions. Thanks to the choice of X_i, Y, a_i , and to the fact that the sum of x independent Poisson random variables with distribution $\text{Poi}(1)$ is itself a Poisson random variable with distribution $\text{Poi}(x)$, the statement $\sum_{i=1}^x a_i X_i \leq_{\text{cx}} Y$ reads as $\text{Poi}(x) \leq_{\text{cx}} x\text{Poi}(1)$ which implies $\mathbb{E}_{P \sim \text{Poi}(x)}[P\ell(P)] \leq \mathbb{E}_{P \sim \text{Poi}(1)}[(xP)\ell(xP)]$, as $x\ell(x)$ is convex.

Acknowledgments

The authors would like to thank the referees, the associate editor, and the area editor Ilan Lobel, for insightful feedback and suggestions on improving our work. The authors would also like to thank Rahul Chandan, Yiannis Giannakopoulos, Jason Marden, Tim Roughgarden, Rahul Savani, Maximilian Schiffer and Adam Wierman for valuable comments and suggestions. A preliminary version of this manuscript appeared at the 22nd ACM Conference on Economics & Computation (EC'21) (Paccagnan and Gairing 2021).

Dario Paccagnan is an Assistant Professor (UK Lecturer) in the Department of Computing at Imperial College London. Dario's interests are at the interface of game theory, optimization, and control theory, with particular emphasis on the design of incentives and coordination mechanisms in multiagent settings.

Martin Gairing is a Professor in the Department of Computer Science at the University of Liverpool and head of the Economics and Computation research group. His interests are in Theoretical Computer Science and especially in Algorithmic Game Theory. His work shaped the theory of congestion games, for which he is a recognized leading authority.

References

- Aland S, Dumrauf D, Gairing M, Monien B, Schoppmann F (2011) Exact price of anarchy for polynomial congestion games. *SIAM Journal on Computing* 40(5):1211–1233, URL <http://dx.doi.org/10.1137/090748986>.
- Andrews M, Anta AF, Zhang L, Zhao W (2012) Routing for power minimization in the speed scaling model. *IEEE/ACM Transactions on Networking* 20(1):285–294, URL <http://dx.doi.org/10.1109/TNET.2011.2159864>.
- Awerbuch B, Azar Y, Epstein A (2013) The price of routing unsplittable flow. *SIAM Journal on Computing* 42(1):160–177, URL <http://dx.doi.org/10.1137/070702370>.
- Awerbuch B, Azar Y, Grove EF, Kao MY, Krishnan P, Vitter JS (1995) Load balancing in the L_p norm. *Proceedings of IEEE 36th Annual Foundations of Computer Science*, 383–391 (IEEE), URL <http://dx.doi.org/10.1109/SFCS.1995.492494>.
- Axhausen KW, Molloy J, Tchervenkov C, Becker F, Hintermann B, Schoeman B, Götschi T, Castro Fernández A, Tomic U (2021) Empirical analysis of mobility behavior in the presence of pigovian transport pricing. Technical report, ETH Zurich, URL <http://dx.doi.org/10.3929/ethz-b-000500100>.
- Barman S, Fawzi O, Fermé P (2021) Tight approximation guarantees for concave coverage problems. Bläser M, Monmege B, eds., *38th International Symposium on Theoretical Aspects of Computer Science, STACS 2021, March 16-19, 2021, Saarbrücken, Germany (Virtual Conference)*, volume 187 of *LIPICs*, 9:1–9:17, URL <http://dx.doi.org/10.4230/LIPICs.STACS.2021.9>.
- Bhaskar U, Cheng Y, Ko YK, Swamy C (2016) Hardness results for signaling in bayesian zero-sum and network routing games. Conitzer V, Bergemann D, Chen Y, eds., *Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16, Maastricht, The Netherlands, July 24-28, 2016*, 479–496 (ACM), URL <http://dx.doi.org/10.1145/2940716.2940753>.
- Bilò V (2018) A unifying tool for bounding the quality of non-cooperative solutions in weighted congestion games. *Theory of Computing Systems* 62(5):1288–1317, URL <http://dx.doi.org/10.1007/s00224-017-9826-1>.
- Bilò V, Vinci C (2019) Dynamic taxes for polynomial congestion games. *ACM Transactions on Economics and Computation* 7(3):15:1–15:36, URL <http://dx.doi.org/10.1145/3355946>.
- Blum A, Hajiaghayi MT, Ligett K, Roth A (2008) Regret minimization and the price of total anarchy. Dwork C, ed., *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008*, 373–382 (ACM), URL <http://dx.doi.org/10.1145/1374376.1374430>.

- Blumrosen L, Dobzinski S (2006) Welfare maximization in congestion games. Feigenbaum J, Chuang JC, Pennock DM, eds., *Proceedings 7th ACM Conference on Electronic Commerce (EC-2006)*, Ann Arbor, Michigan, USA, June 11–15, 2006, 52–61 (ACM), URL <http://dx.doi.org/10.1145/1134707.1134714>.
- Bonifaci V, Harks T, Schäfer G (2010) Stackelberg routing in arbitrary networks. *Mathematics of Operations Research* 35(2):330–346, URL <http://dx.doi.org/10.1287/moor.1100.0442>.
- Caragiannis I, Kaklamanis C, Kanellopoulos P (2010) Taxes for linear atomic congestion games. *ACM Transactions on Algorithms* 7(1):13:1–13:31, URL <http://dx.doi.org/10.1145/1868237.1868251>.
- Castiglioni M, Celli A, Marchesi A, Gatti N (2021) Signaling in bayesian network congestion games: the subtle power of symmetry. *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 5252–5259, URL <http://dx.doi.org/10.1609/aaai.v35i6.16663>.
- Cesa-Bianchi N, Lugosi G (2006) *Prediction, learning, and games* (Cambridge University Press), URL <http://dx.doi.org/10.1017/CB09780511546921>.
- Chakrabarty D, Mehta A, Nagarajan V (2005) Fairness and optimality in congestion games. Riedl J, Kearns MJ, Reiter MK, eds., *Proceedings 6th ACM Conference on Electronic Commerce (EC-2005)*, Vancouver, BC, Canada, June 5–8, 2005, 52–57 (ACM), URL <http://dx.doi.org/10.1145/1064009.1064015>.
- Chandan R, Paccagnan D, Marden JR (2019) When smoothness is not enough: Toward exact quantification and optimization of the price of anarchy. *arXiv preprint arXiv:1911.07823* URL <http://dx.doi.org/10.48550/arXiv.1911.07823>.
- Christodoulou G, Gairing M (2016) Price of stability in polynomial congestion games. *ACM Transactions on Economics and Computation* 4(2):10:1–10:17, URL <http://dx.doi.org/10.1145/2841229>.
- Christodoulou G, Koutsoupias E (2005) The price of anarchy of finite congestion games. Gabow HN, Fagin R, eds., *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22–24, 2005*, 67–73 (ACM), URL <http://dx.doi.org/10.1145/1060590.1060600>.
- Christodoulou G, Koutsoupias E, Nanavati A (2009) Coordination mechanisms. *Theoretical Computer Science* 410(36):3327–3336, URL <http://dx.doi.org/10.1016/j.tcs.2009.01.005>.
- Colini-Baldeschi R, Cominetti R, Mertikopoulos P, Scarsini M (2020) When is selfish routing bad? the price of anarchy in light and heavy traffic. *Operations Research* 68(2):411–434, URL <http://dx.doi.org/10.1287/opre.2019.1894>.
- Cominetti R, Scarsini M, Schröder M, Stier-Moses N (2023) Approximation and convergence of large atomic congestion games. *Mathematics of Operations Research* 48(2):784–811, URL <http://dx.doi.org/10.1287/moor.2022.1281>.
- Dubhashi DP, Ranjan D (1998) Balls and bins: A study in negative dependence. *Random Structures and Algorithms* 13(2):99–124, URL [http://dx.doi.org/10.1002/\(SICI\)1098-2418\(199809\)13:2%3C99::AID-RSA1%3E3.O.CO;2-M](http://dx.doi.org/10.1002/(SICI)1098-2418(199809)13:2%3C99::AID-RSA1%3E3.O.CO;2-M).
- Dudycz S, Manurangsi P, Marcinkowski J, Sornat K (2020) Tight approximation for proportional approval voting. Bessiere C, ed., *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, 276–282 (ijcai.org), URL <http://dx.doi.org/10.24963/ijcai.2020/39>.
- Durrett R (2010) *Probability: Theory and Examples, 4th Edition* (Cambridge University Press), URL <http://dx.doi.org/10.1017/9781108591034>.
- Feige U (1998) A threshold of $\ln n$ for approximating set cover. *Journal of the ACM* 45(4):634–652, URL <http://dx.doi.org/10.1145/285055.285059>.
- Fotakis D (2010) Stackelberg strategies for atomic congestion games. *Theory of Computing Systems* 47(1):218–249, URL <http://dx.doi.org/10.1007/s00224-008-9152-8>.
- Fotakis D, Spirakis PG (2008) Cost-balancing tolls for atomic network congestion games. *Internet Mathematics* 5(4):343–363, URL <http://dx.doi.org/10.1080/15427951.2008.10129175>.
- Gairing M, Kollias K, Kotsialou G (2020) Existence and efficiency of equilibria for cost-sharing in generalized weighted congestion games. *ACM Transactions on Economics and Computation* 8(2):11:1–11:28, URL <http://dx.doi.org/10.1145/3391434>.

- Geisberger R, Sanders P, Schultes D, Vetter C (2012) Exact routing in large road networks using contraction hierarchies. *Transportation Science* 46(3):388–404, URL <http://dx.doi.org/10.1287/trsc.1110.0401>.
- Gkatzelis V, Kollias K, Roughgarden T (2016) Optimal cost-sharing in general resource selection games. *Operations Research* 64(6):1230–1238, URL <http://dx.doi.org/10.1287/opre.2016.1512>.
- Harks T (2019) Pricing in resource allocation games based on lagrangean duality and convexification. *arXiv preprint arXiv:1907.01976* URL <http://dx.doi.org/10.48550/arXiv.1907.01976>.
- Harks T, Kleinert I, Klimm M, Möhring RH (2015) Computing network tolls with support constraints. *Networks* 65(3):262–285, URL <http://dx.doi.org/10.1002/net.21604>.
- Harks T, Oosterwijk T, Vredeveld T (2016) A logarithmic approximation for polymatroid congestion games. *Operations Research Letters* 44(6):712–717, URL <http://dx.doi.org/10.1016/j.orl.2016.09.001>.
- Haurie A, Marcotte P (1985) On the relationship between Nash-Cournot and Wardrop equilibria. *Networks* 15(3):295–308, URL <http://dx.doi.org/10.1002/net.3230150303>.
- Hoefler M, Olbrich L, Skopalik A (2008) Taxing subnetworks. Papadimitriou CH, Zhang S, eds., *Internet and Network Economics, 4th International Workshop, WINE 2008, Shanghai, China, December 17–20, 2008. Proceedings*, volume 5385 of *Lecture Notes in Computer Science*, 286–294 (Springer), URL http://dx.doi.org/10.1007/978-3-540-92185-1_35.
- Jelinek T, Klaas M, Schäfer G (2014) Computing optimal tolls with arc restrictions and heterogeneous players. Mayr EW, Portier N, eds., *31st International Symposium on Theoretical Aspects of Computer Science (STACS 2014), STACS 2014, March 5–8, 2014, Lyon, France*, volume 25 of *LIPICs*, 433–444 (Schloss Dagstuhl - Leibniz-Zentrum für Informatik), URL <http://dx.doi.org/10.4230/LIPICs.STACS.2014.433>.
- Jiang AX, Leyton-Brown K (2015) Polynomial-time computation of exact correlated equilibrium in compact games. *Games and Economic Behavior* 91:347–359, URL <http://dx.doi.org/10.1016/j.geb.2013.02.002>.
- Joag-Dev K, Proschan F (1983) Negative association of random variables with applications. *The Annals of Statistics* 286–295, URL <http://dx.doi.org/10.1214/aos/1176346079>.
- Kalai AT, Vempala SS (2005) Efficient algorithms for online decision problems. *Journal of Computer and System Sciences* 71(3):291–307, URL <http://dx.doi.org/10.1016/j.jcss.2004.10.016>.
- Klenke A, Mattner L (2010) Stochastic ordering of classical discrete distributions. *Advances in Applied Probability* 42(2):392–410, URL <http://dx.doi.org/10.1239/aap/1275055235>.
- Klimm M, Schmand D, Tönnis A (2019) The online best reply algorithm for resource allocation problems. Fotakis D, Markakis E, eds., *Algorithmic Game Theory - 12th International Symposium, SAGT 2019, Athens, Greece, September 30 - October 3, 2019, Proceedings*, volume 11801 of *Lecture Notes in Computer Science*, 200–215 (Springer), URL http://dx.doi.org/10.1007/978-3-030-30473-7_14.
- Koutsoupias E, Papadimitriou C (1999) Worst-case equilibria. *Annual Symposium on Theoretical Aspects of Computer Science*, 404–413 (Springer), URL http://dx.doi.org/10.1007/3-540-49116-3_38.
- Kumar VSA, Marathe MV, Parthasarathy S, Srinivasan A (2009) A unified approach to scheduling on unrelated parallel machines. *Journal of the ACM* 56(5):28:1–28:31, URL <http://dx.doi.org/10.1145/1552285.1552289>.
- Makarychev K, Sviridenko M (2018) Solving optimization problems with diseconomies of scale via decoupling. *Journal of the ACM* 65(6):42:1–42:27, URL <http://dx.doi.org/10.1145/3266140>.
- Mansour T, Schork M (2015) *Commutation relations, normal ordering, and Stirling numbers* (CRC Press), URL <http://dx.doi.org/10.1201/b18869>.
- Marden JR, Wierman A (2013) Distributed welfare games. *Operations Research* 61(1):155–168, URL <http://dx.doi.org/10.1287/opre.1120.1137>.
- Meyers CA, Schulz AS (2012) The complexity of welfare maximization in congestion games. *Networks* 59(2):252–260, URL <http://dx.doi.org/10.1002/net.20439>.

- Nash JF (1950) Equilibrium points in n-person games. *Proceedings of the national academy of sciences* 36(1):48–49, URL <http://dx.doi.org/10.1073/pnas.36.1.48>.
- Paccagnan D, Chandan R, Ferguson BL, Marden JR (2021) Optimal taxes in atomic congestion games. *ACM Transactions on Economics and Computation* 9(3):33, URL <http://dx.doi.org/10.1145/3457168>.
- Paccagnan D, Chandan R, Marden JR (2020) Utility design for distributed resource allocation—Part I: Characterizing and optimizing the exact price of anarchy. *IEEE Transactions on Automatic Control* 65(11):4616–4631, URL <http://dx.doi.org/10.1109/TAC.2019.2961995>.
- Paccagnan D, Gairing M (2021) In congestion games, taxes achieve optimal approximation. *Proceedings of the 22nd ACM Conference on Economics and Computation*, 743–744, URL <http://dx.doi.org/10.1145/3465456.3467601>.
- Paccagnan D, Gentile B, Parise F, Kamgarpour M, Lygeros J (2018) Nash and wardrop equilibria in aggregative games with coupling constraints. *IEEE Transactions on Automatic Control* 64(4):1373–1388, URL <http://dx.doi.org/10.1109/TAC.2018.2849946>.
- Papadimitriou CH, Roughgarden T (2008) Computing correlated equilibria in multi-player games. *Journal of the ACM* 55(3):14:1–14:29, URL <http://dx.doi.org/10.1145/1379759.1379762>.
- Pia AD, Ferris M, Michini C (2017) Totally unimodular congestion games. Klein PN, ed., *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, 577–588 (SIAM), URL <http://dx.doi.org/10.1137/1.9781611974782.37>.
- Pigou AC (1920) *The Economics of Welfare* (Macmillan), URL <http://dx.doi.org/10.4324/9781351304368>.
- Rosenthal RW (1973) A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory* 2(1):65–67, URL <http://dx.doi.org/10.1007/BF01737559>.
- Roughgarden T (2014) Barriers to near-optimal equilibria. *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, 71–80 (IEEE Computer Society), URL <http://dx.doi.org/10.1109/FOCS.2014.16>.
- Roughgarden T (2015) Intrinsic robustness of the price of anarchy. *Journal of the ACM* 62(5):32:1–32:42, URL <http://dx.doi.org/10.1145/2806883>.
- Shaked M, Shanthikumar JG (2007) *Stochastic orders* (Springer Science & Business Media), URL <http://dx.doi.org/10.1007/978-0-387-34675-5>.
- Suri S, Tóth CD, Zhou Y (2007) Selfish load balancing and atomic congestion games. *Algorithmica* 47(1):79–96, URL <http://dx.doi.org/10.1007/s00453-006-1211-4>.
- Tekin C, Liu M, Southwell R, Huang J, Ahmad SHA (2012) Atomic congestion games on graphs and their applications in networking. *IEEE/ACM Transactions on Networking* 20(5):1541–1552, URL <http://dx.doi.org/10.1109/TNET.2012.2182779>.
- Transportation Networks for Research (2016) <https://github.com/bstabler/TransportationNetworks>, accessed April, 13, 2022.
- Von Neumann J, Morgenstern O (1944) *Theory of games and economic behavior* (Princeton university press), URL <http://dx.doi.org/10.1515/9781400829460>.
- Wardrop JG (1952) Some theoretical aspects of road traffic research. *Proceedings of the institution of civil engineers* 1(3):325–362, URL <http://dx.doi.org/10.1680/ipeds.1952.11362>.
- Wu Z, Moehring RH, Ren C, Xu D (2022) A convergence analysis of the price of anarchy in atomic congestion games. *Mathematical Programming* 1–57, URL <http://dx.doi.org/10.1007/s10107-022-01853-0>.
- Zardini G, Lanzetti N, Pavone M, Frazzoli E (2021) Analysis and control of autonomous mobility-on-demand systems. *Annual Review of Control, Robotics, and Autonomous Systems* 5, URL <http://dx.doi.org/10.1146/annurev-control-042920-012811>.

Additional Proofs

As in the main body of the manuscript, throughout the Electronic Companion we extend the domain of definition of all resource costs from \mathbb{N} to \mathbb{N}_0 by setting their value to zero. This is without loss of generality, and only used to ease the notation, as already clarified on page Page 5 and Footnote 6.

EC.1. Additional material for Section 2

EC.1.1. Binomial expected value converges to Poisson expected value

LEMMA EC.1. *Let $\ell : \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$ satisfy Assumption 1, and $\rho_\ell < \infty$. Then $\forall k \in \mathbb{N}$*

$$\lim_{h \rightarrow \infty} \mathbb{E}_{X \sim \text{Bin}(h, k/h)}[X\ell(X)] = \mathbb{E}_{X \sim \text{Poi}(k)}[X\ell(X)].$$

Proof. We begin noting that the limiting operation is delicate since we do not want to assume boundedness of $x\ell(x)$ for $x \rightarrow \infty$, as this would disqualify interesting cases such as that of polynomials. Since $\rho_\ell < \infty$ by assumption, then $\mathbb{E}_{X \sim \text{Poi}(k)}[X\ell(X)] < \infty$ for any $k \in \mathbb{N}$. Fix k , and let $f_h(i)$ and $f(i)$ denote the probability mass function corresponding to $\text{Bin}(h, k/h)$ and $\text{Poi}(k)$, and recall that $f_h(i) \rightarrow f(i)$ for all i . Define $c(x) = x\ell(x)$ for $x \in \mathbb{N}$ and $c(0) = 0$. The result follows from

$$\mathbb{E}_{X \sim \text{Poi}(k)}[c(X)] = \sum_{i=0}^{\infty} \lim_{h \rightarrow \infty} f_h(i)c(i) \leq \lim_{h \rightarrow \infty} \sum_{i=0}^{\infty} f_h(i)c(i) = \lim_{h \rightarrow \infty} \mathbb{E}_{X \sim \text{Bin}(h, k/h)}[c(X)] \leq \mathbb{E}_{X \sim \text{Poi}(k)}[c(X)],$$

where the first equality holds by definition of expected value and by replacing $f(i) = \lim_{h \rightarrow \infty} f_h(i)$. The following inequality holds by Fatou's lemma and existence of the limit (which hold as this is a series with non-negative terms). As a result, we can interchange the limiting operation with the infinite sum. The next equality is due to the definition of expected value. The last inequality is a consequence of the fact that c is a convex function and the Poisson distribution $\text{Poi}(k)$ dominates the binomial distribution $\text{Bin}(h, k/h)$ in the sense of the convex ordering (Shaked and Shanthikumar 2007), ensuring that $\mathbb{E}_{X \sim \text{Bin}(h, k/h)}[c(X)] \leq \mathbb{E}_{X \sim \text{Poi}(k)}[c(X)]$ for all h . One way to see this is to utilize the fact that the ratio between the probability mass functions $f_h(i)/f(i)$ is unimodal as shown in (Klenke and Mattner 2010, Sec 2.7), and that f_h and f are not ordered by the usual stochastic order, thus concluding thanks to (Shaked and Shanthikumar 2007, Thm. 3.A.53). Alternatively, one can compare the two expectations directly. Q.E.D.

EC.1.2. Proof of Proposition 2

Proof. Existence of a partitioning system is proved through a probabilistic approach similarly to that in Feige (1998). The idea is to construct each \mathcal{P}_i independently from a uniform distribution. Formally, each element in $[n]$ is assigned to k of the $P_{j,i}$ uniformly at random. This ensures that, by construction, each element in $[n]$ appears in exactly k different sets of \mathcal{P}_j . Thus the first property of the partitioning system holds trivially. Define $c(x) = x\ell(x)$ for $x \in \mathbb{N}$ and $c(0) = 0$. In order to prove the second property, we fix $B \subseteq [\beta]$ with $|B| = h$ and correspondingly consider $Q = \{P_{j,i(j)}, j \in B\}$. We intend to show that with high probability $\sum_{r \in Q} c(|Q|_r) \geq (\mathbb{E}_{X \sim \text{Bin}(h, k/h)}[c(X)] - \eta)n$ holds. This would imply that, among all possible ways of constructing $\{\mathcal{P}_i\}_{i=1}^r$ there exists at least one that satisfies the property. To prove that $\sum_{r \in Q} c(|Q|_r) \geq (\mathbb{E}_{X \sim \text{Bin}(h, k/h)}[c(X)] - \eta)n$ holds with high probability, we compute the expected cost that arises from the probabilistic choice outlined above for $\{\mathcal{P}_i\}_{i=1}^r$. In particular

$$\mathbb{E}[c(|Q|_r)] = \mathbb{E}_{X \sim \text{Bin}(h, k/h)}[c(X_r)],$$

since the number of sets in which each resource appears is given by the random variable $X_r = \sum_{j \in B} \mathbb{1}_{r \in P_{j,i(j)}} \sim \text{Bin}(h, k/h)$, owing to the fact that $\mathbb{1}_{r \in P_{j,i(j)}} \sim \text{Ber}(k/h)$ are independent (here

$\mathbb{1}$ denotes the indicator function). We then use Chernoff-Hoeffding bound on the total cost $\sum_{r \in [n]} c(|Q|_r)$, where each term is bounded by $0 \leq c(|Q|_r) \leq c(h)$ owing to c being non-decreasing and non-negative.¹³ Thus, with probability smaller or equal to $2e^{-2n\eta^2/(c(h))^2}$, it is $|\sum_{r \in [n]} c(|Q|_r) - \mathbb{E}_{X \sim \text{Bin}(h, k/h)}[c(X)]| \geq \eta n$. Since there are $\binom{\beta}{h} \cdot h^h \leq (1+h)^\beta$ possible choices for B and Q , a union bound guarantees that with probability higher than $1 - 2(1+h)^\beta \cdot e^{-2n\eta^2/(c(h))^2}$, the cost of all B, Q satisfies $|\sum_{r \in [n]} c(|Q|_r) - \mathbb{E}_{X \sim \text{Bin}(h, k/h)}[c(X)]| < \eta n$. With the specific choice of n as in the statement, we are guaranteed this property with a probability of at least $4/5$. This shows that a partitioning system always exists. One such object can be computed by simple enumeration over all possible choices, which are only a function of h, n and β . Q.E.D.

EC.1.3. Hardness factor for polynomial resource cost

LEMMA EC.2. *Let $d \in \mathbb{N}_0$, and let $\ell : \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$ be $\ell(x) = \sum_{j=0}^d \alpha_j x^j$ where $\alpha_j \geq 0$ for all j . Let $m : \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$ be $m(x) = x^d$. Define ρ_ℓ, ρ_m as in (4). Then $\rho_\ell \leq \rho_m$.*

Proof. Using the definition of ρ_ℓ, ρ_m , we need to show that

$$\rho_\ell = \sup_{x \in \mathbb{N}} \frac{\mathbb{E}_{P \sim \text{Poi}(x)}[\sum_{j=0}^d \alpha_j P^{j+1}]}{\sum_{j=0}^d \alpha_j x^{j+1}} \leq \sup_{x \in \mathbb{N}} \frac{\mathbb{E}_{P \sim \text{Poi}(x)}[P^{d+1}]}{x^{d+1}} = \rho_m.$$

Towards this goal, we will prove that for any $x \in \mathbb{N}$, it is $\sum_{j=0}^d \alpha_j \mathbb{E}_{P \sim \text{Poi}(x)}[P^{j+1}] x^{d+1} \leq \sum_{j=0}^d \alpha_j x^{j+1} \mathbb{E}_{P \sim \text{Poi}(x)}[P^{d+1}]$, from which one concludes by rearranging and taking the supremum on both sides. We will do so term by term, showing that for all $j \in \{0, \dots, d\}$ it is

$$\mathbb{E}_{P \sim \text{Poi}(x)}[P^{j+1}] x^{d-j} \leq \mathbb{E}_{P \sim \text{Poi}(x)}[P^{d+1}].$$

Towards this goal, we recall that the $(j+1)$ 'st moment of the Poisson distribution $\text{Poi}(x)$ equals $\sum_{i=0}^{j+1} x^i \left\{ \begin{smallmatrix} j+1 \\ i \end{smallmatrix} \right\}$, where $\left\{ \begin{smallmatrix} j+1 \\ i \end{smallmatrix} \right\}$ is a Stirling number of the second kind (Mansour and Schork 2015, p. 63). Hence,

$$\mathbb{E}_{P \sim \text{Poi}(x)}[P^{j+1}] x^{d-j} = \sum_{i=0}^{j+1} x^{i+d-j} \left\{ \begin{smallmatrix} j+1 \\ i \end{smallmatrix} \right\} = \sum_{i=d-j}^{d+1} x^i \left\{ \begin{smallmatrix} j+1 \\ i-(d-j) \end{smallmatrix} \right\} \leq \sum_{i=0}^{d+1} x^i \left\{ \begin{smallmatrix} d+1 \\ i \end{smallmatrix} \right\} = \mathbb{E}_{P \sim \text{Poi}(x)}[P^{d+1}],$$

where the key inequality $\left\{ \begin{smallmatrix} j+1 \\ i-(d-j) \end{smallmatrix} \right\} \leq \left\{ \begin{smallmatrix} d+1 \\ i \end{smallmatrix} \right\}$ holds by the recurrence relation defining the Stirling numbers, and the fact that $j \leq d$ (Mansour and Schork 2015, Thm 1.17). Q.E.D.

EC.2. Additional material for Section 3

EC.2.1. Convexity of Poisson expected value

LEMMA EC.3. *Let $\ell : \mathbb{N} \rightarrow \mathbb{R}_{>0}$ satisfy Assumption 1 and ρ_ℓ in (4) be finite. Define $p : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ as*

$$p(v) = \mathbb{E}_{P \sim \text{Poi}(v)}[P\ell(P)] = \left(\sum_{i=0}^{\infty} i b(i) \frac{v^i}{i!} \right) e^{-v}.$$

Then, p is convex and differentiable infinitely many times in $\mathbb{R}_{\geq 0}$.

¹³ Observe that the random variables $\{c(|Q|_r)\}_{r \in [n]} = \{c(X_r)\}_{r \in [n]}$ are negatively associated, which is enough to conclude thanks to Dubhashi and Ranjan (1998) and $\eta \in (0, 1)$. Since $c(X_r) = c(\sum_{j \in B} X_{r,j})$, where $X_{r,j} = \mathbb{1}_{r \in P_{j,i(j)}}$, is non-decreasing in $\{X_{r,j}\}_{j \in [\beta]}$, negative association of $\{c(X_r)\}_{r \in [n]}$ can be shown by proving negative association of $\{X_{r,j}\}_{r \in [n], j \in [\beta]}$ (Joag-Dev and Proschan 1983, Property 6). For fixed $j \in [\beta]$, the variables $\{X_{r,j}\}_{r \in [n]}$ are negatively associated as they are a permutation distribution of $(0, \dots, 0, 1, \dots, 1)$ with $n - kn/h$ zeros and kn/h ones (Joag-Dev and Proschan 1983, Def. 2.10 and Thm. 2.11). Owing to this, and thanks to (Joag-Dev and Proschan 1983, Property 7) negative association of $\{X_{r,j}\}_{r \in [n], j \in [\beta]}$ follows from the above, and from the fact that $\{X_{r,j}\}_{r \in [n]}$ are mutually independent.

Proof. We start by showing that if $\rho_\ell < \infty$, then $p(v)$ is well defined, i.e., the series converges to a finite value for any choice of $v \in \mathbb{R}_{\geq 0}$ (as standard, \mathbb{R} does not include infinity). For $v = 0$, this is immediate. For $v > 0$, define $c(v) = v\ell(v)$ for $v \in \mathbb{N}$ with $c(0) = 0$, and observe that $\rho_\ell < \infty$ implies $p(v)/c(v) < \infty$ for all fixed $v \in \mathbb{N}$ thanks to the definition of ρ_ℓ , so that $p(v) < \infty$ since $0 < c(v) < \infty$ for finite $v \in \mathbb{N}$. Therefore also $\sum_{i=0}^{\infty} c(i) \frac{v^i}{i!} < \infty$ for all $v \in \mathbb{N}$. Since $\sum_{i=0}^{\infty} c(i) \frac{v^i}{i!}$ is increasing in $v \geq 0$, boundedness over the naturals, immediately implies boundedness of the same expression over the non-negative reals. Exploiting the fact that also e^{-v} is bounded, we have shown that $p(v)$ is well defined, and converges to a finite value for any choice of $v \in \mathbb{R}_{\geq 0}$. As a result p is differentiable infinitely many times in its domain, since it is the product of a convergent power series, and of e^{-v} . We can therefore prove convexity of $p(v)$ by computing its second order derivative and verifying that it is non-negative. The first derivative reads as

$$\begin{aligned} p'(v) &= -e^{-v} \left(\sum_{i=0}^{\infty} c(i) \frac{v^i}{i!} \right) + e^{-v} \left(\sum_{i=0}^{\infty} c(i) \frac{iv^{i-1}}{i!} \right) \\ &= -e^{-v} \left(\sum_{i=0}^{\infty} c(i) \frac{v^i}{i!} \right) + e^{-v} \left(\sum_{i=0}^{\infty} c(i+1) \frac{v^i}{i!} \right) \\ &= e^{-v} \sum_{i=0}^{\infty} \frac{v^i}{i!} \Delta c(i), \end{aligned}$$

where we defined $\Delta c(i) = c(i+1) - c(i)$. Following an identical approach the second derivative is

$$p''(v) = e^{-v} \sum_{i=0}^{\infty} \frac{v^i}{i!} [\Delta c(i+1) - \Delta c(i)].$$

The summand corresponding to $i = 0$ reads as $c(2) - 2c(1) = 2(\ell(2) - \ell(1)) > 0$ since ℓ is non-decreasing. The summands corresponding to $i \geq 1$ are non-negative as semi-convexity of ℓ implies convexity of c . Hence, $p''(v) \geq 0$ for all $v \in \mathbb{R}_{\geq 0}$ as desired. Q.E.D.

EC.2.2. Proof of Lemma 1

Lemma 1 consists of three parts. We first prove part (a), then (b), and finally (c).

▷ *Proof of Part (a).* We will show that $f_r(x, v) \geq \ell_r(x)$ for any $x \in \mathbb{N}, v \in \mathbb{R}_{\geq 0}, r \in \mathcal{R}$, as this suffices to conclude since $\tau_r(x, v_r) = f_r(x, v) - \ell_r(x)$. We do so separately for each resource, and thus drop the index r in the following. The case of $v = 0$ follows readily, since $f(x, 0) = \ell(x)$. Similarly, for $x = 0$, it is $f(0, v) = 0 \geq \ell(0) = 0$. In the remaining cases, we are left to prove

$$f(x, v) = \frac{(x-1)!}{v^x} \sum_{i=0}^{x-1} \frac{p(v) - i\ell(i)}{i!} v^i \geq \ell(x) \quad x \in \mathbb{N}, v \in \mathbb{R}_{>0}.$$

Define $c(x) = xb(x)$ for $x \in \mathbb{N}$ with $c(0) = 0$. The latter inequality holds if $p(v) \geq (\sum_{i=0}^x c(i) \frac{v^i}{i!}) / (\sum_{j=0}^{x-1} \frac{v^j}{j!})$ for all $x \in \mathbb{N}, v \in \mathbb{R}_{>0}$. We complete the proof showing that the right hand side in the previous expression is non-decreasing in x , so that the desired property holds for any $x \in \mathbb{N}$ if it holds for x arbitrarily large, i.e., if

$$p(v) \geq \lim_{x \rightarrow \infty} \frac{\sum_{i=0}^x c(i) \frac{v^i}{i!}}{\sum_{j=0}^{x-1} \frac{v^j}{j!}}, \quad \forall v \in \mathbb{R}_{>0},$$

which is indeed satisfied (with equality), as it reduces to $p(v) \geq (p(v)e^v)/e^v$, owing to the definition of $p(v)$ and $\sum_{j=0}^{\infty} \frac{v^j}{j!} = e^v$. To conclude, we thus need to prove that for all $x \in \mathbb{N}$, $v \in \mathbb{R}_{>0}$, it is

$$\begin{aligned}
\frac{\sum_{i=0}^{x+1} c(i) \frac{v^i}{i!}}{\sum_{j=0}^x \frac{v^j}{j!}} &\geq \frac{\sum_{i=0}^x c(i) \frac{v^i}{i!}}{\sum_{j=0}^{x-1} \frac{v^j}{j!}} \iff \sum_{i=0}^{x+1} \sum_{j=0}^{x-1} c(i) \frac{v^{i+j}}{i!j!} \geq \sum_{i=0}^x \sum_{j=0}^x c(i) \frac{v^{i+j}}{i!j!} \\
&\iff \sum_{j=0}^{x-1} c(x+1) \frac{v^{x+1+j}}{(x+1)!j!} \geq \sum_{i=0}^x c(i) \frac{v^{i+x}}{i!x!} \\
&\iff \sum_{j=1}^x c(x+1) \frac{v^{x+j}}{(x+1)!(j-1)!} \geq \sum_{i=1}^x c(i) \frac{v^{i+x}}{i!x!} \\
&\iff \sum_{j=1}^x \left[\frac{c(x+1)}{(x+1)!(j-1)!} - \frac{c(j)}{x!j!} \right] v^{j+x} \geq 0 \\
&\iff \sum_{j=1}^x [\ell(x+1) - \ell(j)] v^{j+x} \geq 0
\end{aligned}$$

which follows from the above chain of implications, and the fact that $\ell(x)$ is non-decreasing. Q.E.D.

▷ *Proof of Part (a)*. To ease the notation, we drop the subscript r , i.e., we show that $f(x+1, v) \geq f(x, v)$ for all $x \in \mathbb{N}_0$ and $v \in \mathbb{R}_{\geq 0}$. When $v = 0$, then $f(x, 0) = \ell(x)$, which is non-decreasing as $\ell(x)$ is so. Thus, in the following we restrict to the case of $v > 0$. When, in addition, $x = 0$ the inequality reduces to $f(1, v) \geq f(0, v) \iff p(v)/v \geq 0$, which holds as $p(v) > 0$ for $v > 0$. We are thus left to consider the case of $x \in \mathbb{N}$ and $v \in \mathbb{R}_{>0}$. Substituting the expression for $f(x+1, v)$ and $f(x, v)$ in the latter inequality and isolating $p(v)$ results in

$$p(v) \left(x \sum_{i=0}^x \frac{v^i}{i!} - v \sum_{i=0}^{x-1} \frac{v^i}{i!} \right) \geq \left(x \sum_{i=0}^x \frac{c(i)}{i!} v^i - v \sum_{i=0}^{x-1} \frac{c(i)}{i!} v^i \right), \quad (\text{EC.1})$$

where we defined $c(i) = i\ell(i)$ for $i \in \mathbb{N}$ and $c(0) = 0$. The term in brackets on the left hand side can be equivalently written as

$$\sum_{i=0}^x x \frac{v^i}{i!} - \sum_{i=0}^{x-1} \frac{v^{i+1}}{i!} = \sum_{i=0}^x \frac{x-i}{i!} v^i.$$

Similarly, the term in the right hand side brackets, is equivalent to

$$\sum_{i=1}^x x \frac{c(i)}{i!} v^i - \sum_{i=0}^{x-1} \frac{c(i)}{i!} v^{i+1} = \sum_{i=1}^x \frac{xc(i) - ic(i-1)}{i!} v^i.$$

Thus, inequality (EC.1) reduces to

$$p(v) \geq \frac{\sum_{i=1}^x \frac{xc(i) - ic(i-1)}{i!} v^i}{\sum_{i=0}^x \frac{x-i}{i!} v^i}, \quad (\text{EC.2})$$

whereby we used the fact that the denominator is positive since $x \geq 1$ and $x \geq i$. Thus, we require (EC.2) to hold for all $x \in \mathbb{N}$ and $v \in \mathbb{R}_{>0}$. Lemma EC.4, ensures that the right hand side of the previous inequality is non-decreasing in $x \in \mathbb{N}$, for each fixed $v \in \mathbb{R}_{>0}$. Therefore, (EC.2) holds, if it holds when x is arbitrarily large, that is if

$$p(v) \geq \lim_{x \rightarrow \infty} \frac{\sum_{i=0}^x \frac{c(i)}{i!} v^i}{\sum_{i=0}^x \frac{v^i}{i!}}. \quad (\text{EC.3})$$

Notice, though, that the right hand side in the last expression is precisely equal to $\frac{p(v)e^v}{e^v} = p(v)$, thanks to the definition of $p(v)$ in (10) and to the fact that $\sum_{i=0}^{\infty} \frac{v^i}{i!} = e^v$. Therefore (EC.3) holds (with equality), which completes the proof of part (a). Q.E.D.

Proof of Part (c). Using the definition of f_r in (11), observe that

$$\begin{aligned} v f_r(x+1, v) - x f_r(x, v) &= v \frac{x!}{v^{x+1}} \sum_{i=0}^x \frac{p_r(v) - i \ell_r(i)}{i!} v^i - x \frac{(x-1)!}{v^x} \sum_{i=0}^{x-1} \frac{p_r(v) - i \ell_r(i)}{i!} v^i \\ &= \frac{x!}{v^x} \frac{p_r(v) - x \ell_r(x)}{x!} v^x \\ &= p_r(v) - x \ell_r(x), \end{aligned}$$

or equivalently $x \ell_r(x) - x f_r(x, v) + v f_r(x+1, v) = p_r(v)$ as needed. Q.E.D.

EC.2.3. Technical Lemma used to prove Lemma 1

LEMMA EC.4. *Let $c: \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$ be convex. Then, the function $g: \mathbb{N} \times \mathbb{R}_{>0} \rightarrow \mathbb{R}_{\geq 0}$,*

$$g(x, v) = \frac{\sum_{i=1}^x \frac{xc(i) - ic(i-1)}{i!} v^i}{\sum_{i=0}^x \frac{v^i}{i!}},$$

where we set $c(0) = 0$, is non-decreasing for all $x \in \mathbb{N}$, for any fixed $v \in \mathbb{R}_{>0}$.

Proof. Proving the claim amounts so showing that for all $v \in \mathbb{R}_{>0}, x \in \mathbb{N}$ it is

$$\frac{\sum_{i=1}^{x+1} a(x+1, i) \frac{v^i}{i!}}{\sum_{j=0}^{x+1} b(x+1, j) \frac{v^j}{j!}} \geq \frac{\sum_{i=1}^x a(x, i) \frac{v^i}{i!}}{\sum_{j=0}^x b(x, j) \frac{v^j}{j!}}, \quad (\text{EC.4})$$

where we let $a(x, i) = xc(i) - ic(i-1)$ and $b(x, j) = x - j$ to ease the notation. Since the denominators on the left and right hand side are positive, and since $b(x+1, x+1) = 0$, (EC.4) is equivalent to

$$\sum_{i=1}^x \sum_{j=0}^x a(x+1, i) b(x, j) \frac{v^{i+j}}{i! j!} + \sum_{j=0}^x b(x, j) a(x+1, x+1) \frac{v^{j+x+1}}{j! (x+1)!} - \sum_{i=1}^x \sum_{j=0}^x a(x, i) b(x+1, j) \frac{v^{i+j}}{i! j!} \geq 0,$$

which we rewrite

$$\underbrace{\sum_{i=1}^x \sum_{j=0}^x (a(x+1, i) b(x, j) - a(x, i) b(x+1, j)) \frac{v^{i+j}}{i! j!}}_{\textcircled{1}} + \underbrace{\sum_{j=0}^x b(x, j) a(x+1, x+1) \frac{v^{j+x+1}}{j! (x+1)!}}_{\textcircled{2}} \geq 0. \quad (\text{EC.5})$$

We now turn our attention to each of the two terms appearing in the previous inequality. In particular, we will show that collecting all the contributions corresponding to the same power of v significantly simplifies the expressions, and allows us to conclude.

We begin with the second term, and substitute the definitions of a, b in $b(x, j) a(x+1, x+1)$ so that

$$\textcircled{2} = \sum_{j=0}^x \frac{x-j}{j! x!} (c(x+1) - c(x)) v^{j+x+1}. \quad (\text{EC.6})$$

We now focus on the first term, and observe that

$$\begin{aligned} a(x+1, i) b(x, j) - a(x, i) b(x+1, j) &= ((x+1)c(i) - ic(i-1))(x-j) - (xc(i) - ic(i-1))(x+1-j) \\ &= -jc(i) + ic(i-1), \end{aligned}$$

where we made use of the definitions of a and b . Hence, we can utilize indices i and $q = i + j$ in place of i, j to rewrite the first term appearing in (EC.5) as

$$\begin{aligned}
\textcircled{1} &= \sum_{i=1}^x \sum_{j=0}^x (ic(i-1) - jc(i)) \frac{v^{i+j}}{i!j!} \\
&= \sum_{q=1}^{2x} \sum_{\substack{i \in [x] \\ \text{s.t. } q-x \leq i \leq q}} \frac{ic(i-1) - (q-i)c(i)}{i!(q-i)!} v^q \\
&= \sum_{q=x+1}^{2x} \sum_{i=q-x}^x \frac{ic(i-1) - (q-i)c(i)}{i!(q-i)!} v^q + \sum_{q=1}^x \sum_{i=1}^q \frac{ic(i-1) - (q-i)c(i)}{i!(q-i)!} v^q \\
&= \sum_{q=x+1}^{2x} \sum_{i=q-x}^x \frac{ic(i-1) - (q-i)c(i)}{i!(q-i)!} v^q \\
&= \sum_{q=x+1}^{2x} \frac{c(q-x-1) - c(x)}{x!(q-x-1)!} v^q \\
&= \sum_{j=0}^{x-1} \frac{1}{j!x!} (c(j) - c(x)) v^{j+x+1}
\end{aligned} \tag{EC.7}$$

where the second line is obtained using the fact that q runs from 1 to $2x$, $i \in [x]$ and $j = q - i$ belongs to $0 \leq q - i \leq x$. The third line follows upon distinguishing the case of $x + 1 \leq q \leq 2x$ and $q \in [x]$. The fourth line is due to the fact that the second summand in the third line vanishes, since

$$\sum_{i=1}^q \frac{ic(i-1) - (q-i)c(i)}{i!(q-i)!} v^q = \sum_{i=2}^q \frac{c(i-1)}{(i-1)!(q-i)!} - \sum_{i=1}^{q-1} \frac{c(i)}{i!(q-i-1)!} = 0.$$

The fifth line follows from

$$\begin{aligned}
\sum_{i=q-x}^x \frac{ic(i-1) - (q-i)c(i)}{i!(q-i)!} &= \sum_{i=q-x}^x \frac{c(i-1)}{(i-1)!(q-i)!} - \sum_{i=q-x}^x \frac{c(i)}{i!(q-i-1)!} \\
&= \frac{c(q-x-1)}{(q-x-1)!x!} + \sum_{i=q-x+1}^x \frac{c(i-1)}{(i-1)!(q-i)!} - \sum_{i=q-x}^{x-1} \frac{c(i)}{i!(q-i-1)!} - \frac{c(x)}{x!(q-x-i)!} \\
&= \frac{c(q-x-1) - c(x)}{x!(q-x-1)!}
\end{aligned}$$

The final line is derived reverting to the original indices i and j .

Thus, in light of (EC.7) and (EC.6), the inequality (EC.4) we need to show reduces to

$$\sum_{j=0}^{x-1} [c(j) - c(x) + (x-j)(c(x+1) - c(x))] \frac{v^{j+x+1}}{j!x!} \geq 0 \quad \forall v \in \mathbb{R}_{>0}, x \in \mathbb{N}.$$

The summand corresponding to $j = 0$ reads as $xc(x+1) - (x+1)c(x) = x(x+1)(b(x+1) - b(x))$, and it is non-negative since b is non-decreasing. All other summands are non-negative thanks to the convexity of c , that guarantees $c(x+1) - c(x) \geq (c(x) - c(j))/(x-j)$ as $1 \leq j < x$, so that $c(j) - c(x) + (x-j)(c(x+1) - c(x)) \geq 0$. One concludes using these observations and $v > 0$. **Q.E.D.**