Learning-Augmented Coordination Mechanisms

George Christodoulou
Aristotle University of Thessaloniki
Archimedes, Athena Research Center
gichristo@csd.auth.gr

Vasilis Christoforidis
Aristotle University of Thessaloniki
Archimedes, Athena Research Center
vgchrist@csd.auth.gr

Alkmini Sgouritsa
Athens University of Economics and Business
Archimedes, Athena Research Center
alkmini@aueb.gr

Ioannis Vlachos
Athens University of Economics and Business
Archimedes, Athena Research Center
ioa.vlahos@aueb.gr

Abstract

The inefficiency of decentralized resource allocation, a core challenge in Algorithmic Game Theory, is classically measured by the *Price of Anarchy (PoA)* and the *Price of Stability (PoS)*. The most widely studied and foundational class of these models are *atomic and non-atomic congestion games*, for which a robust theory quantifying the inefficiency of equilibria is well-established. To mitigate the effects of this selfish behavior, we employ *Coordination Mechanisms*, which modify resource costs to incentivize socially improved equilibrium outcomes. In their standard form, Coordination Mechanisms have been limited by a pessimistic, worst-case view; it is assumed that the demand is *unknown* and *adversarially* chosen. Consequently, positive results remain sparse, applying mainly to specialized network topologies like parallel links. We address this limitation by studying the design and analysis of *learning-augmented coordination mechanisms*, where the mechanism is endowed with a potentially inaccurate *prediction of the demand* $\bar{r}$. We study general *atomic and non-atomic congestion games* with polynomial latency functions of degree d. Our main positive result is a learning-augmented coordination mechanism tuned by a confidence parameter β which achieves a "best of both worlds" type of result: it achieves approximately optimal performance for an accurate prediction without sacrificing the worst-case guarantee of a good equilibrium (PoS). We further establish bounds on a stronger robustness measure, the PoA-Robustness (worst equilibrium performance under cost modification). The mechanism's construction relies on transforming approximate equilibria of the original game into exact equilibria of the modified game using a carefully defined *approximate potential function*. We complement our positive results by providing *impossibility results*. We first show that a stronger notion of consistency, requiring that *all* equilibria have optimal social cost for an accurate prediction, fails to guarantee bounded robustness. We also show a strong *error tolerance* limit: any mechanism achieving optimal consistency in non-atomic games suffers a sharp, discontinuous performance that collapses to the original *PoA* value under even infinitesimally small prediction errors, highlighting the limits of efficiency guarantees in non-atomic congestion games.

1 Introduction

For the past three decades, a core research area in Algorithmic Game Theory has focused on investigating and quantifying the *inefficiency of equilibria*. Two standard concepts are used to

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

measure this inefficiency. On the one hand, the *Price of Anarchy* (PoA) [KP99] serves as a robust measure for this assessment, comparing the worst-case equilibrium with the system optimal solution. On the other hand, the *Price of Stability* (PoS) [ADK+08] offers an optimistic measure that uses the best-case equilibrium for this comparison. The former is a very robust measure, while the latter is perhaps a more appropriate concept from an engineering perspective, analyzing the ideal solution that we would like our protocols to produce.

The richest and most well-studied classes of games among these models are arguably *atomic and non-atomic congestion games* (see, e.g., [NRTV07] for a detailed discussion), where a set of players compete for a set of resources. The resources are sensitive to the amount of traffic they receive, which is typically modeled via a latency function. This highly descriptive class of games captures a wide variety of application domains, most prominently *routing games*.

The PoA and PoS of congestion games are thoroughly understood. Roughgarden and Tardos's seminal work [RT02, RT04] delivered tight bounds for the PoA of *non-atomic congestion games*, where individual players exert negligible influence on traffic. For *atomic congestion games*, the PoA was progressively resolved: initially for affine latency functions [AAE05, CK05b], then for polynomials [ADG+11], and finally for more general monotone cost functions [Rou15]. These developments underscore congestion games' role as a central paradigm within the *Smoothness Framework* [Rou15], an elegant and unified methodology for bounding the PoA. The PoS[1] of congestion games retains much smaller values compared to the respective PoA values: approximately 1.577 for affine functions [CK05a, CFK+11], as opposed to the PoA bound of 2.5; and $\Theta(d)$ [CG16] for polynomials of degree d as opposed to the much higher PoA bound of $(\Theta(d/\ln d))^{d+1}$. All aforementioned PoA and PoS results are *tight*, and collectively they provide a solid theory for estimating the inefficiency of equilibria in congestion games.

Coordination Mechanisms. While the PoA and PoS offer valuable insights into the performance of equilibria, they are inherently *post-factum* measures; they assess the quality of an equilibrium that has already been reached. This observation motivated the introduction of *coordination mechanisms* [CKN09, CMP14] for congestion games. A coordination mechanism aims to improve the equilibrium performance of a given congestion game by modifying (e.g., increasing) the latency functions of the resources in a way that encourages equilibrium behavior with good social cost. The challenge is that these modifications must occur *without prior knowledge of the traffic demand*.[2] [CMP14] took the worst-case approach where demand is chosen adversarially after latency modifications are applied. This robust, yet overly pessimistic, approach has yielded positive results only in the special case of parallel links with affine latencies [CMP14]. However, in many practical settings, it is often realistic to assume that traffic can be accurately predicted from historical data [LDK+15].

In light of this, we adopt the *learning-augmented framework* (see, e.g., [MV22] for an introduction), where algorithms are equipped with an additional, potentially unreliable, prediction about the input. This paradigm has proven effective in the design of *online algorithms*, where the algorithm is augmented with a prediction of future inputs. In this work, we study the design of *learning-augmented coordination mechanisms*, where the mechanism is augmented with a *prediction of the demand rate*.[3]

Due to space limitations, we defer a more detailed discussion of related work to Appendix A.

1.1 Contributions

We study coordination mechanisms through the lens of the learning-augmented framework. A learning-augmented coordination mechanism takes as input an instance of a congestion game and a *prediction of the demand* $\bar{r}$ (see Sec. 2 for formal definitions), but not the actual demand r which is unknown. Our results apply to a broad class of games: non-atomic (Section 3) and atomic (Appendix C) congestion games with polynomial latency functions with nonnegative coefficients of maximum degree d. We present the non-atomic case in the main text and defer the treatment of atomic games to Appendix due to space limitations.

[1] We note that for non-atomic congestion games PoA and PoS coincide.

[2] In a routing game, the demand refers to the amount of traffic that needs to move from each source to each destination of a given network. If the demand is known in advance, it is possible to modify the latencies to enforce the optimal flow [CMP14].

[3] The work of [CCSV26] is closely related to our work and we compare our results with the results of [CCSV26] in Section A.

1.1.1 Positive Results

Consistency - Robustness trade-off. We propose a learning-augmented coordination mechanism, called POTENTIAL, which takes as input a predicted demand $\bar{r}$ and can be tuned by a *confidence parameter* $\beta \in [1, d+1]$. POTENTIAL modifies (increases) the latency functions of the game to ensure that the modified game satisfies the following:

- when the prediction is accurate (i.e., $r = \bar{r}$), there exists an equilibrium with approximately optimal social cost as a function of β **(Consistency)**;
- even when the predicted demand is *arbitrarily inaccurate*, there exists an equilibrium whose social cost is at most a factor of $\beta(d+1)$ away from the optimal social cost **(Robustness)**.

We remark that in both cases, we compare the equilibrium cost with the *increased* latencies to the *original* optimal cost, prior to the modification. The exact statements can be found in Theorem 1 for the case of non-atomic congestion games and in Theorem 3 for the case of atomic congestion games, respectively.

It is instructive to compare our results with the performance of the best Nash Equilibrium in the original game which is measured by the PoS. For atomic congestion games, we are able to improve the performance of the best equilibrium by a factor of β when the prediction is accurate, while for arbitrarily inaccurate predictions, we recover the worst-case guarantee for the best equilibrium up to a factor of β. For non-atomic games we obtain analogous guarantees, by losing at most a factor depending on β and a small factor of $O(\ln(d))$. In Table 1 we provide the asymptotic PoS upper bounds in the original games as $d \to \infty$, alongside upper bounds on consistency and robustness in the modified games. For the non-atomic consistency entry, the table reports a simplified upper bound; the exact expression is given in Theorem 1.

Model	PoS ($d \to \infty$)	Consistency	Robustness
Non-atomic	$\Theta\left(\frac{d}{\ln d}\right)$ [RT04]	$\frac{d+1}{\beta}$ (Thm. 1)	$\beta(d+1)$ (Thm. 1)
Atomic	$\Theta(d)$ [CG16]	$\frac{d+1}{\beta}$ (Thm. 3)	$\beta(d+1)$ (Thm. 3)

Table 1: Summary of the main guarantees for atomic and non-atomic congestion games with polynomial latency functions of degree d. The first column reports the classical PoS benchmark without predictions, while the last two columns give the consistency and robustness guarantees of our learning-augmented coordination mechanism with confidence parameter β. For non-atomic games, as $d \to \infty$, the consistency bound can be refined to $O(d/(\beta + \ln d))$.

Our results demonstrate that incorporating predictions of the input may dramatically improve the performance when the prediction is accurate, without significantly compromising robustness for moderate values of β. We further emphasize that our results hold for *general* congestion games with polynomial latency functions, as opposed to previous work on coordination mechanisms which considered parallel links [CMP14, WYC16, CCSV26] and affine latency functions [CMP14].

PoA-Robustness. We also provide a bound on the performance guarantee of *any* equilibrium under arbitrarily inaccurate predicted demand (**PoA-Robustness**). As before, we compare the worst-case equilibrium cost measured using the *modified* latencies, with the *original* optimal cost, prior to the modification. In Table 2 we provide the asymptotic PoA bounds for both the unmodified and modified games when $d \to \infty$.

Note that, for atomic games, the performance degrades in the worst-case only from $(\Theta(d/\ln d))^{d+1}$ to $(O(d))^{d+1}$.For non-atomic games, the PoA-Robustness can be substantially worse than in the unmodified game. This degradation occurs because our intervention introduces discontinuities, effectively transforming the game into an atomic one. Note, however, that such discontinuities are necessary: without them, no improvement is possible, as established in [CMP14]. Nonetheless, we emphasize that all our bounds are essentially constant, as they do not depend on the demand (i.e., the number of agents or the input rate) but only on the maximum degree d of the polynomial.

	PoA	PoA-Robustness
Non-atomic	$\Theta\left(\frac{d}{\ln d}\right)$ [RT04]	$\max\{d+1, \beta^{d+1}\}$ (Thm. 2)
Atomic	$\left(\Theta\left(\frac{d}{\ln d}\right)\right)^{d+1}$ [ADG+11]	$O\left(\left(\frac{d}{W(d/\beta)}\right)^{d+1}\right)$ (Thm. 4)

Table 2: Asymptotic behavior, as $d \to \infty$, of the classical PoA bounds for atomic and non-atomic congestion games with polynomial latency functions of degree d, together with the corresponding PoA-Robustness guarantees of our learning-augmented coordination mechanism. For non-atomic games, the displayed bound is the simplified uniform upper bound implied by Theorem 2: the exact first branch is $\Theta(d/\ln d)$ at $\beta = 1$ and increases up to $d+1$ as β approaches $(d+1)^{1/(d+1)}$, while the second branch is β^{d+1}. In the atomic PoA-Robustness bound, W denotes the Lambert W function.

1.1.2 Impossibility Results

We complement our results by providing impossibility results that hold for *all* learning-augmented coordination mechanisms.

PoA - Consistency. At this point, a natural question may arise: "*Is it possible to obtain good consistency - robustness trade-off if one insists that* all *equilibria of the modified game have optimal social cost when $r = \bar{r}$?*" We present strong negative results for both atomic (Theorem 6) and non-atomic games (Theorem 5), showing that any coordination mechanism that demands this stronger notion of consistency fails to guarantee bounded robustness for any arbitrarily inaccurate predicted demand.

Error tolerance. Lastly, inspired by the notion of *smoothness* for learning-augmented algorithms, we study how approximation guarantees degrade under small prediction errors in non-atomic congestion games (see Section D.2). In Theorem 8, we establish a strong impossibility result, showing that there is *no* mechanism that can achieve optimal performance under perfectly accurate predictions while simultaneously maintaining approximately optimal performance when the prediction is slightly inaccurate. More precisely, for any 1-consistent mechanism, we show that even an infinitesimal prediction error causes the performance to revert back to the original PoA obtained without predictions. Furthermore, we show that this negative result applies to any class of games beyond single-commodity network congestion games (see Section D.2). For instance, in multi-commodity networks or symmetric congestion games, no mechanism can be 1-consistent while also maintaining an approximation ratio that approaches 1 as the prediction error tends to zero. Nonetheless, in the special case of single-commodity network congestion games, we exploit the network structure to obtain a positive result (Theorem 11).

1.1.3 Techniques

Design of coordination mechanisms. At a high level, our coordination mechanism transforms *approximate* equilibria with approximately optimal social cost in the original game into *exact* equilibria of the modified game that retain the same social cost. A central building block in the definition of our mechanism is an *approximate potential function* w.r.t. a twisted set of latency functions h, tuned by the confidence parameter β. The minimizers of this approximate potential are approximate equilibria with bounded performance. The associated latencies h, together with the potential minimizer $\bar{\mathbf{s}}$ (for $r = \bar{r}$), are carefully used via a threshold mechanism to define the POTENTIAL coordination mechanism. The construction ensures that, when $r = \bar{r}$, the potential minimizer is also a potential minimizer of the modified game, and hence an exact equilibrium for that game, while retaining its original social cost. Remarkably, this improvement comes at little cost: even under inaccurate predictions, any potential minimizer of the modified game continues to provide strong performance guarantees.

Analysis. When predictions are inaccurate, we propose a framework to provide robustness bounds for learning-augmented coordination mechanisms. To achieve this, we transfer standard techniques for deriving bounds on the inefficiency of equilibria, most notably the smoothness framework for the PoA and the potential function method for the PoS, to relate the equilibria of the modified game to those of the original game.

We remark that the games that we consider in this work are guaranteed to possess pure Nash equilibria. For simplicity, our upper bound proofs are presented for pure equilibria. However, all our proofs can be directly extended to mixed, correlated, and coarse-correlated equilibria via standard extension arguments.

Our approach is quite general and provides a unified framework for designing and analyzing coordination mechanisms in both non-atomic and atomic congestion games. In most cases, we tailor the analysis to the specific parameters of an approximate potential function, yielding tighter bounds than those obtained through general approximate equilibrium guarantees.

2 Model and Notation

We study both atomic and non-atomic congestion games. We introduce the relevant notation and the definitions of each setting separately.

Atomic congestion games. An (atomic) *congestion game* is a tuple $(N, E, (\mathcal{S}_i)_{i\in N}, (\ell_e)_{e\in E})$, where $N = \{1, \ldots, n\}$ is a set of n players and E is a set of resources. Each player $i \in N$ chooses a pure *strategy* $s_i \subseteq E$, from a set of available strategies $\mathcal{S}_i \subseteq 2^E$. Each resource $e \in E$ is associated with a nonnegative *latency function* $\ell_e : \mathbb{N} \to \mathbb{R}^+$. We consider polynomial latency functions with maximum degree d and nonnegative coefficients; that is every latency function is of the form $\ell_e(x) = \sum_{k=0}^{d} a_{e,k} \cdot x^k$ with $a_{e,k} \geq 0$ for all k. We refer to the pair $r = (N, (\mathcal{S}_i)_{i\in N})$ as the *demand* of the game.

A *pure strategy profile* is a choice of strategies $\mathbf{s} = (s_1, s_2, \ldots, s_n) \in \mathcal{S} = \mathcal{S}_1 \times \cdots \times \mathcal{S}_n$ one for each player. We adopt the standard notation $\mathbf{s}_{-i} = (s_1, \ldots, s_{i-1}, s_{i+1}, \ldots s_n)$, $\mathcal{S}_{-i} = \mathcal{S}_1 \times \cdots \times \mathcal{S}_{i-1} \times \mathcal{S}_{i+1} \times \cdots \times \mathcal{S}_n$, and $\mathbf{s} = (s_i, \mathbf{s}_{-i})$. For a pure strategy profile $\mathbf{s}$ define the *load* $x_e(\mathbf{s}) = |i \in N : e \in s_i|$ as the number of players that use resource e. The *cost* for player i is defined by $C_i(\mathbf{s}) = \sum_{e\in s_i} \ell_e(x_e(\mathbf{s}))$.

A pure strategy profile $\mathbf{s}$ is a pure *Nash equilibrium* if and only if for every player $i \in N$ and for all $s_i' \in \mathcal{S}_i$, we have $C_i(\mathbf{s}) \leq C_i(s_i', \mathbf{s}_{-i})$. The social cost of a strategy profile $\mathbf{s}$ is the sum of the players' costs

$$C(\mathbf{s}) = \sum_{i\in N} C_i(\mathbf{s}) = \sum_{e\in E} x_e(\mathbf{s})\ell_e(x_e(\mathbf{s})).$$

We further denote by C_{opt} the minimum social cost, i.e. $C_{opt} = \min_s C(s)$. The *Price of Anarchy* of a congestion game is the social cost of the worst-case Nash equilibrium over the optimum social cost

$$\mathsf{PoA} = \max_{\mathbf{s}\in \mathrm{Eq}(r)} \frac{C(\mathbf{s})}{C_{opt}},$$

where $\mathrm{Eq}(r)$ denotes the set of all equilibria under demand r. The PoA for a class of games is the largest PoA among all games in the class.

Congestion games admit a *potential* function $\Phi(\mathbf{s}) = \sum_{e\in E} \Phi_e(\mathbf{s}) = \sum_{e\in E} \sum_{j=1}^{x_e(\mathbf{s})} \ell_e(j)$, which was introduced by Rosenthal [Ros73]. The set of pure Nash equilibria corresponds to local optima of the potential function [Ros73]. This mapping suggests a special pure Nash equilibrium, the *potential minimizer*, which is the profile $\mathbf{s}^\diamond$ that minimizes the potential, i.e., $\Phi(\mathbf{s}^\diamond) \leq \Phi(\mathbf{s}')$ for all $\mathbf{s}' \in \mathcal{S}$. We sometimes use the notation Φ^ℓ, Φ_e^ℓ to emphasize the dependency of the latency function in the computation. Similarly, non-atomic congestion games admit a potential function $\Phi(f) = \sum_{e\in E} \int_0^{f_e} \ell_e(x)dx$ and equilibria correspond to global potential minimizers of the potential function [BMW56]. The inefficiency of potential minimizers has been well-analyzed as it is used as a means to show upper bounds for the Price of Stability[4] [ADK+08], see e.g. [CK05a, CKS11, CFK+11, CG16, Bil18, CGGS19].

Non-atomic congestion games. Rather than formally defining the class of non-atomic congestion games, we focus on the selfish routing problem, as it serves as a clear and more illustrative example. The key difference is that, in non-atomic congestion games, players' strategies are not necessarily paths within a network (as in the selfish routing problem), but rather arbitrary subsets of resources.

[4]The price of stability of a game is the social cost of the best-case Nash equilibrium over the optimum social cost.

Let $G = (V, E)$ be a directed graph, where V is the set of vertices and E is the set of edges. We consider a set $\mathcal{I}$ of k commodities: origin-destination pairs $\{s_i, t_i\}_{i\in\mathcal{I}}$ that define the origins and the destinations in the graph. We denote by $\mathcal{P}_i$ the set of all $s_i - t_i$ simple paths of G, and by $\mathcal{P} = \bigcup_i \mathcal{P}_i$ the union of all feasible paths. A flow f, is a mapping from the set of paths to the set of nonnegative reals $f : \mathcal{P} \to \mathbb{R}_{\geq 0}$. For a given flow f, the flow on an edge is defined as the sum of the flows of all the paths that use this edge $f_e = \sum_{P\in\mathcal{P}, e\in P} f_P$. We relate with every commodity $\{s_i, t_i\}$ a demand rate r_i, as the total traffic that needs to move from s_i to t_i. A flow f is feasible, if for every commodity $\{s_i, t_i\}$, the traffic rate equals the total flow routed on paths in $\mathcal{P}_i$, i.e., $r_i = \sum_{P\in\mathcal{P}_i} f_P$. We refer to $r = (r_i)_{i\in\mathcal{I}}$ as the *demand* of the game.

Every edge introduces a delay in the network which depends on the load of the edge and is determined by a nonnegative latency function, $\ell_e : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$. In our case this is a polynomial of maximum degree d with nonnegative coefficients. An instance of a routing game is denoted by the triple (G, r, ℓ). The latency of a path P, for a given flow f, is defined as the sum of all the latencies of the edges that belong to P, i.e., $\ell_P(f) = \sum_{e\in P} \ell_e(f_e)$. The social cost w.r.t. a given flow f is the total delay due to f

$$C(f) = \sum_{P\in\mathcal{P}} \ell_P(f) f_P = \sum_{e\in E} \ell_e(f_e) f_e.$$

We denote by $C_{opt}(r)$ the minimum cost of a flow for demand r. A feasible flow f that routes r_i units of flow from s_i to t_i for each commodity $i \in \mathcal{I}$, is at Nash (or Wardrop) equilibrium if and only if it is routed on shortest $s_i - t_i$ paths, i.e., if for any two paths $p_1, p_2 \in \mathcal{P}_i$ with $f_{p_1} > 0$, $\ell_{p_1}(f) \leq \ell_{p_2}(f)$. We will often use discontinuous latency functions, under which standard Wardrop equilibria might not exist. Hence, we resort to the notion of *user equilibria* [BS94]. A feasible flow f that routes r_i units of flow for each commodity $i \in \mathcal{I}$ is a user equilibrium, if and only if for all paths $p_1, p_2 \in \mathcal{P}_i$ with $f_{p_1} > 0$,

$$\ell_{p_1}(f) \leq \liminf_{\epsilon\downarrow 0} \ell_{p_2}(f + \epsilon\mathbf{1}_{p_2} - \epsilon\mathbf{1}_{p_1}),$$

where $\mathbf{1}_p$ denotes a unit of flow passing through p. Throughout the paper, we refer to user equilibrium simply as equilibrium.

Learning-augmented coordination mechanisms. A *learning-augmented coordination mechanism* takes as input a partial instance of a game, where the actual demand r is missing. This instance is augmented with a *prediction*[5] of the demand $\bar{r}$. For the case of *atomic games*, this input is a tuple $(\bar{N}, E, (\bar{\mathcal{S}}_i)_{i\in\bar{N}}, (\ell_e)_{e\in E})$, where the predicted demand $\bar{r}$ consists of a prediction of the set of players $\bar{N}$, and a prediction of the strategy space[6] $\bar{\mathcal{S}}_i$ for each player $i \in \bar{N}$. For the case of *non-atomic games*, the input is a tuple $(G, \bar{r}, \ell)$ where $\bar{r} = (\bar{r}_i)_{i\in\bar{\mathcal{I}}}$ refers to the prediction on the set of commodities $\bar{\mathcal{I}}$ along with a prediction of the demand $\bar{r}_i$ for each commodity $i \in \bar{\mathcal{I}}$. For simplicity, we will often refer only to the predicted demand $\bar{r}$ as the input of the coordination mechanism, emphasizing the predicted element of the input, as the other parameters are usually inferred from the context.

Given a predicted demand $\bar{r}$, a learning-augmented mechanism defines a set of modified latencies $\hat{\ell} = (\hat{\ell}_e)_{e\in E}$ such that $\hat{\ell}_e(x) \geq \ell_e(x)$ for all $x \geq 0$. The goal of this modification is to create equilibria $\mathbf{s}$ with good social cost denoted by $\hat{C}_{\bar{r}}(\mathbf{s})$, in the modified game with the *increased* latency functions, compared to the optimal social cost in the original game using the *original* latencies, when the unknown actual demand r arrives. Adopting the standard practice of the learning augmented framework, we evaluate the performance of a mechanism in two different regimes of r: when the demand is accurate, $r = \bar{r}$, and when r is arbitrarily inaccurate.

We define the *Consistency* of the mechanism as the social cost of the prediction-derived potential minimizer $\bar{\mathbf{s}}$ used to define the modified latencies, compared to the cost of the optimal solution (w.r.t. the original latencies) when the prediction is accurate ($r = \bar{r}$), and we take the worst-case $\bar{r}$. We define the *Robustness* of a mechanism as the worst-case guarantee of a potential minimizer $\mathbf{s}^{\diamond}$[7] when

[5] This is a key difference with the standard definition of a coordination mechanism [CKN09, CMP14], where r is an unknown parameter chosen adversarially.

[6] In a network congestion game, the strategy space can be implicitly defined by the simple paths of the network that connect a set of origin-destination pairs $(\bar{s}_i, \bar{t}_i)$, one for each $i \in \bar{N}$.

[7] Technically the potential minimizer may not be unique, in which case we consider the worst-case one.

the predicted demand may be *arbitrarily inaccurate*, taking the worst-case over r and $\bar{r}$. Formally,

$$\textbf{(Consistency)}\ \max_{\bar{r}} \frac{\hat{C}_{\bar{r}}(\bar{\mathbf{s}})}{C_{\mathrm{opt}}(\bar{r})}\,, \qquad \textbf{(Robustness)}\ \max_{r,\bar{r}} \max_{\mathbf{s}^\diamond \in \mathrm{Pot}(r)} \frac{\hat{C}_{\bar{r}}(\mathbf{s}^\diamond)}{C_{opt}(r)}.$$

where $\mathrm{Pot}(r)$ denotes the set of all potential minimizers of $\hat{\Phi} := \Phi^{\hat{\ell}}$ w.r.t. the modified latencies for demand equal to r. We also define a stronger notion of robustness, which we refer to as *PoA*-Robustness, as the worst-case guarantee for *any* equilibrium when the predicted demand may be arbitrarily inaccurate[8]

$$\textbf{(PoA-Robustness)}\ \max_{r,\bar{r}} \max_{\mathbf{s}\in \mathrm{Eq}(r)} \frac{\hat{C}_{\bar{r}}(\mathbf{s})}{C_{opt}(r)}\,,$$

where $\mathrm{Eq}(r)$ denotes the set of all equilibria under demand r. We also define the approximate versions of these concepts. A learning-augmented mechanism is α-consistent if its consistency is at most α, and ρ-robust if its robustness is at most ρ.

3 Non-atomic Congestion Games

In this section, we study learning-augmented coordination mechanisms for the class of non-atomic congestion games with polynomial latencies of maximum degree d. We are given an instance augmented with a prediction of the demand $(G, \bar{r}, \ell)$, where $\bar{r}$ denotes the predicted demand $\bar{r} = (\bar{r}_i)_{i\in\bar{\mathcal{I}}}$, and $\bar{\mathcal{I}}$ is a prediction of the set of $\bar{k}$ commodities $(\bar{s}_1, \bar{t}_1), \ldots, (\bar{s}_{\bar{k}}, \bar{t}_{\bar{k}})$.

We propose a learning-augmented mechanism, which we call POTENTIAL (Definition 2). Given $\bar{r}$, it appropriately modifies the latency functions, achieving a trade-off between consistency and robustness which is tuned by a confidence parameter $\beta \in [1, d+1]$. The following theorem, which is the main result of this section, summarizes the consistency–robustness trade-off as a function of β.

Theorem 1. *Given the predicted input rates $\bar{r}$ and a parameter $\beta \in [1, d+1]$, the* POTENTIAL(β,$\bar{r}$) *mechanism has*

$$\textit{Consistency:}\ \left(\beta\left(1 - \frac{d}{d+1}\left(\frac{\beta}{d+1}\right)^{1/d}\right)\right)^{-1}; \qquad \textit{Robustness:}\ \beta(d+1)\,.$$

The consistency guarantee is evaluated on the prediction-derived potential minimizer used to define the modified latencies, while the robustness guarantee holds for every potential minimizer of the modified game. Increasing β improves the consistency guarantee but worsens the robustness guarantee linearly in β. At $\beta = 1$, the mechanism attains its strongest robustness guarantee, namely $d+1$, which is asymptotically a factor of $\Theta(\ln d)$ away from the PoS/PoA bound for non-atomic congestion games with polynomial latencies. At $\beta = d+1$, the mechanism is 1-consistent, while its robustness guarantee becomes $(d+1)^2$. The above theorem gives the following corollary for the interesting and well-studied case of affine latency functions, where $d = 1$.

Corollary 1. *Given the predicted input rates $\bar{r}$ and a parameter $\beta \in [1, 2]$, the* POTENTIAL(β,$\bar{r}$) *mechanism for affine latency functions is $\frac{4}{(4-\beta)\beta}$-consistent and 2β-robust.*

We additionally provide a stronger PoA-Robustness guarantee that applies to the worst-case social cost across *all* equilibria, not just to potential minimizers.

Theorem 2. *Given the predicted input rates $\bar{r}$ and a parameter $\beta \in [1, d+1]$, the* POTENTIAL(β,$\bar{r}$) *mechanism has*

$$\textit{PoA-Robustness:}\ \begin{cases} \left(1 - \beta\frac{d}{d+1}\left(\frac{\beta}{d+1}\right)^{1/d}\right)^{-1}, & \textit{for } \beta \in [1, (d+1)^{\frac{1}{d+1}}) \\ \beta^{d+1}, & \textit{for } \beta \in [(d+1)^{\frac{1}{d+1}}, d+1] \end{cases}$$

For the special case of affine latency functions, we obtain the following result, which serves as a warm-up for the proof of Theorem 2.

[8] The PoA-Robustness guarantee holds for *any* actual demand and coincides with the *engineered Price of Anarchy* of coordination mechanisms [CMP14].

Corollary 2. *Given the predicted input rates $\bar{r}$ and a parameter $\beta \in [1, 2]$, the* POTENTIAL$(\beta,\bar{r})$ *mechanism for affine latency functions has PoA-Robustness $\frac{4}{4-\beta^2}$ for $\beta \leq \sqrt{2}$, and β^2 otherwise.*

Two cases of the parameter β are of special interest; when $\beta = d + 1$, the consistency bound becomes 1 and the PoA-Robustness bound becomes $(d+1)^{d+1}$, while for $\beta = 1$, the consistency and PoA-Robustness bounds are both equal to $\left(1 - \frac{d}{d+1}\left(\frac{1}{d+1}\right)^{1/d}\right)^{-1}$, which matches the standard PoA bound of the original game *without predictions*.

The structure of this section is as follows. In Section 3.1, we introduce the general toolbox to design coordination mechanisms based on predicted rates, including the POTENTIAL mechanism and the main technical lemmas used in its analysis; the proofs of these lemmas are deferred to Appendix B.1. We then give proof overviews for the consistency–robustness guarantee of Theorem 1 and for the PoA-Robustness guarantee of Theorem 2. The full proofs, including the affine special cases of Corollaries 1 and 2, are deferred to Appendices B.2, B.3, and B.4. We also include a threshold-sensitive refinement of the robustness bound in Appendix B.3.1.

3.1 The POTENTIAL Mechanism

We now present the main machinery for designing learning-augmented mechanisms for non-atomic congestion games. We first introduce a generic coordination mechanism, called THRESHOLD, which takes as input a threshold flow f and a set of latency functions h_e, one for each resource $e \in E$. We then state two useful properties of THRESHOLD, summarized in Lemmas 1 and 2. Finally, we define the main mechanism of this section, called POTENTIAL. This learning-augmented mechanism takes as input an instance augmented with a demand prediction $(G, \bar{r}, \ell)$, together with a confidence parameter β. Based on the predicted demand $\bar{r}$, it modifies the latencies to $\hat{\ell}$, and the resulting equilibria are evaluated on the actual instance $(G, r, \hat{\ell})$.

We begin with the formal definition of the THRESHOLD mechanism. Intuitively, THRESHOLD is triggered once the actual flow *exceeds some threshold* f_e, and modifies the latency of this edge according to h_e. If the actual flow does not exceed the threshold, it uses the original latency function. Since coordination mechanisms with continuous modifications have been shown to yield no improvement [CMP14], we employ discontinuous latencies instead.

Definition 1 (THRESHOLD(h, f))**.** *Given a set of latency functions $h = (h_e)_{e\in E}$ and a "threshold" flow $f = (f_e)_{e\in E}$, such that $h_e(x) \geq \ell_e(x)$ for all $x \geq 0$, we define the coordination mechanism* THRESHOLD(h, f) *with the following modified latency functions:*

$$\hat{\ell}_e(x) = \begin{cases} h_e(x), & x > f_e, \\ \ell_e(x), & \textit{otherwise}. \end{cases}$$

An important property of the THRESHOLD mechanism is that, when the threshold flow is chosen as a minimizer of the potential Φ^h, this same flow is also a potential minimizer with respect to the modified latency functions $\hat{\ell}$. Moreover, this particular flow preserves its original social cost, since it never triggers the increased latencies. This property is established in the following lemma.

Lemma 1. *Given a predicted input rate $\bar{r}$ and a set of non-decreasing latency functions $(h_e)_{e\in E}$, with $\ell_e(x) \leq h_e(x)$ for all $x \geq 0$ and $e \in E$, let $\bar{f}$ be a minimizer of Φ^h for $r = \bar{r}$, and let $\hat{\ell}$ be the modified latencies produced by* THRESHOLD$(h, \bar{f})$*. Then $\bar{f}$ is a minimizer of $\Phi^{\hat{\ell}}$ for demand $\bar{r}$ and preserves the original social cost, i.e., $\hat{C}(\bar{f}) = C(\bar{f})$.*

The second important property of THRESHOLD helps to provide robustness upper bounds. The following lemma provides a variational inequality that relates the social cost of any equilibrium w.r.t. the *modified* latencies to a benchmark flow under the auxiliary latencies h.

Lemma 2 (Lifting Lemma)**.** *Given a set of continuous and non-decreasing latency functions $h = (h_e)_{e\in E}$, such that $\ell_e(x) \leq h_e(x)$, for all $x \geq 0$ and $e \in E$, and a "threshold" flow $f^h = (f^h_e)_{e\in E}$, let $\hat{\ell}$ be the modified latencies of* THRESHOLD(h, f^h)*. For any demand r, let f be an equilibrium flow under $\hat{\ell}$ and f^* be any other feasible flow. Then,*

$$\sum_{e\in E} \hat{\ell}_e(f_e) f_e \leq \sum_{e\in E} h_e(f_e) f^*_e \,.$$

We are now ready to define the POTENTIAL mechanism, that takes as input a predicted rate $\bar{r}$ as well as a confidence parameter β. It is then defined as a specific instantiation of the THRESHOLD mechanism, for a carefully selected latency function $h(\beta)$ parametrized by β. For the case of polynomials, $\beta \in [1, d+1]$, where d is the maximum degree over all latency functions. We define the latency functions $h(\beta)$ according to the potential function of the best β-approximate equilibrium provided in [CKS11], and with threshold flow equal to the prediction-derived potential minimizer.

Definition 2 (POTENTIAL(β,$\bar{r}$))**.** *Given a parameter $\beta \in [1, d+1]$, we define the polynomial latency function $h(\beta) = (h_e)_{e \in E}$ as,*

$$h_e(x) = \sum_{k=0}^{\lfloor \beta \rfloor - 1} (k+1) a_{e,k} x^k + \sum_{k=\lfloor \beta \rfloor}^{d} \beta a_{e,k} x^k \,.$$

Let $\bar{f}$ be a potential minimizer of Φ^h for the predicted demand $\bar{r}$. Then, the POTENTIAL$(\beta, \bar{r})$ *mechanism is the learning-augmented mechanism that uses the modified latencies produced by* THRESHOLD$(h(\beta), \bar{f})$.

We analyze the POTENTIAL mechanism through the proof overviews below; the full proofs for affine and polynomial latencies are deferred to Appendix B.

3.2 Proof Overviews

Consistency–Robustness. We give an overview of the proof of Theorem 1

Consistency. We first establish the consistency upper bound. We obtain it by relating the predicted potential minimizer to the best β-approximate equilibrium. Suppose that the prediction is accurate, i.e., $\bar{r} = r$. The POTENTIAL(β,$\bar{r}$) mechanism modifies the latencies according to THRESHOLD$(h, \bar{f})$, for

$$h_e(x) = \sum_{k=0}^{\lfloor \beta \rfloor - 1} (k+1) a_{e,k} x^k + \sum_{k=\lfloor \beta \rfloor}^{d} \beta a_{e,k} x^k \,,$$

$\forall e \in E$, and $\bar{f}$ being a potential minimizer of Φ^h when the input rate is $\bar{r} = r$. Since $\bar{f}$ never triggers the increased latencies, we have by Lemma 1 that $\hat{C}(\bar{f}) = C(\bar{f})$. In [CKS11] it is shown that any flow f that minimizes the approximate potential

$$\Phi^\beta(f) = \sum_{e \in E} \int_0^{f_e} \frac{h_e(x)}{\beta} dx$$

has social cost w.r.t. the original latencies ℓ at most a factor of $\alpha = \left(\beta \left(1 - \frac{d}{d+1} \left(\frac{\beta}{d+1} \right)^{1/d} \right) \right)^{-1}$ away from the optimal cost. Note that $\bar{f}$ also minimizes $\Phi^\beta = \Phi^h / \beta$, and since $\hat{C}(\bar{f}) = C(\bar{f})$, the mechanism is α-consistent.

Robustness. Let f be any feasible flow and let E_h be the set of edges for which $f_e > \bar{f}_e$, i.e. $\hat{\ell}_e(f_e) = h_e(f_e)$, and let $E_\ell = E \setminus E_h$ be the rest of the edges, where the original latencies are used. Then $\hat{\Phi}(f)$ can be written as

$$\hat{\Phi}(f) = \sum_{e \in E} \int_0^{f_e} \hat{\ell}_e(t) dt = \sum_{e \in E_\ell} \int_0^{f_e} \ell_e(t) dt + \sum_{e \in E_h} \left(\int_0^{\bar{f}_e} \ell_e(t) dt + \int_{\bar{f}_e}^{f_e} h_e(t) dt \right) .$$

We split the rest of the proof into two technical inequalities, whose proofs are deferred to Appendix B:

$$\hat{\Phi}(f) \leq C(f) \qquad \text{and} \qquad \hat{\Phi}(f) \geq \frac{1}{\beta(d+1)} \hat{C}(f).$$

We finally combine them to get the robustness bound. Let $f^\diamond$ be any potential minimizer of $\hat{\Phi}$, and let f^* be an optimal flow w.r.t. the original latency functions, i.e., f^* minimizes $C(f)$. Then, by utilizing the above inequalities that hold for any f, and therefore for $f^\diamond$ and f^*, it follows that $\frac{1}{\beta(d+1)} \hat{C}(f^\diamond) \leq \hat{\Phi}(f^\diamond) \leq \hat{\Phi}(f^*) \leq C(f^*)$, where the middle inequality follows from the fact that $f^\diamond$ minimizes $\hat{\Phi}$. This establishes the desired robustness bound for every potential minimizer, and so completes the proof of the theorem.

PoA-Robustness. We give an overview of the proof of Theorem 2. The proof follows a different route and applies to every equilibrium of the modified game. Let f be any equilibrium flow under the modified latencies and let f^* be an optimal flow under the original latencies. The Lifting Lemma 2 gives the key comparison $\hat{C}(f) = \sum_{e\in E} \hat{\ell}_e(f_e) f_e \leq \sum_{e\in E} h_e(f_e) f_e^*$. We then bound the right-hand side by a refined smoothness inequality tailored to the auxiliary latencies h. In particular, for every $\lambda \geq 1/\beta$,

$$\hat{C}(f) \leq \beta \max\left\{ \frac{1}{4\lambda}, \frac{d}{(d+1)\sqrt[d]{\lambda(d+1)}} \right\} \hat{C}(f) + \lambda\beta C(f^*).$$

Rearranging and optimizing over λ yields $\lambda = 1/\beta$ for $\beta \in [1, (d+1)^{1/(d+1)})$ and $\lambda = \beta^d/(d+1)$ for $\beta \in [(d+1)^{1/(d+1)}, d+1]$, which gives

$$\frac{\hat{C}(f)}{C(f^*)} \leq \begin{cases} \left(1 - \beta \frac{d}{d+1} \left(\frac{\beta}{d+1}\right)^{1/d}\right)^{-1}, & \beta \in [1, (d+1)^{1/(d+1)}), \\ \beta^{d+1}, & \beta \in [(d+1)^{1/(d+1)}, d+1]. \end{cases}$$

Acknowledgments and Disclosure of Funding

This work has been partially supported by project MIS 5154714 of the National Recovery and Resilience Plan Greece 2.0 funded by the European Union under the NextGenerationEU Program.

The authors declare no competing interests.

References

[AAE05] Baruch Awerbuch, Yossi Azar, and Amir Epstein. The price of routing unsplittable flow. In *STOC*, pages 57–66. ACM, 2005.

[ABG+24] Priyank Agrawal, Eric Balkanski, Vasilis Gkatzelis, Tingting Ou, and Xizhi Tan. Learning-augmented mechanism design: Leveraging predictions for facility location. *Math. Oper. Res.*, 49(4):2626–2651, 2024.

[ACH14] Fidaa Abed, José R. Correa, and Chien-Chung Huang. Optimal coordination mechanisms for multi-job scheduling games. In *ESA*, volume 8737 of *Lecture Notes in Computer Science*, pages 13–24. Springer, 2014.

[ADG+11] Sebastian Aland, Dominic Dumrauf, Martin Gairing, Burkhard Monien, and Florian Schoppmann. Exact price of anarchy for polynomial congestion games. *SIAM J. Comput.*, 40(5):1211–1233, 2011.

[ADK+08] Elliot Anshelevich, Anirban Dasgupta, Jon M. Kleinberg, Éva Tardos, Tom Wexler, and Tim Roughgarden. The price of stability for network design with fair cost allocation. *SIAM J. Comput.*, 38(4):1602–1623, 2008.

[AFJ+15] Yossi Azar, Lisa Fleischer, Kamal Jain, Vahab S. Mirrokni, and Zoya Svitkina. Optimal coordination mechanisms for unrelated machine scheduling. *Oper. Res.*, 63(3):489–500, 2015.

[AH12] Fidaa Abed and Chien-Chung Huang. Preemptive coordination mechanisms for unrelated machines. In *ESA*, volume 7501 of *Lecture Notes in Computer Science*, pages 12–23. Springer, 2012.

[BCKS16] Umang Bhaskar, Yu Cheng, Young Kun-Ko, and Chaitanya Swamy. Hardness results for signaling in bayesian zero-sum and network routing games. In *EC*, pages 479–496. ACM, 2016.

[BGR14] Kshipra Bhawalkar, Martin Gairing, and Tim Roughgarden. Weighted congestion games: The price of anarchy, universal worst-case examples, and tightness. *ACM Trans. Economics and Comput.*, 2(4):14:1–14:23, 2014.

[BGT23a] Eric Balkanski, Vasilis Gkatzelis, and Xizhi Tan. Mechanism design with predictions: An annotated reading list. *SIGecom Exch.*, 21(1):54–57, 2023.

[BGT23b] Eric Balkanski, Vasilis Gkatzelis, and Xizhi Tan. Strategyproof scheduling with predictions. In *ITCS*, volume 251 of *LIPIcs*, pages 11:1–11:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023.

[BHLR08] Avrim Blum, MohammadTaghi Hajiaghayi, Katrina Ligett, and Aaron Roth. Regret minimization and the price of total anarchy. In *STOC*, pages 373–382. ACM, 2008.

[BHS10] Vincenzo Bonifaci, Tobias Harks, and Guido Schäfer. Stackelberg routing in arbitrary networks. *Math. Oper. Res.*, 35(2):330–346, 2010.

[BIKM14] Sayan Bhattacharya, Sungjin Im, Janardhan Kulkarni, and Kamesh Munagala. Coordination mechanisms from (almost) all scheduling policies. In *ITCS*, pages 121–134. ACM, 2014.

[Bil18] Vittorio Bilò. A unifying tool for bounding the quality of non-cooperative solutions in weighted congestion games. *Theory Comput. Syst.*, 62(5):1288–1317, 2018.

[BKM14] Sayan Bhattacharya, Janardhan Kulkarni, and Vahab S. Mirrokni. Coordination mechanisms for selfish routing over time on a tree. In *ICALP (1)*, volume 8572 of *Lecture Notes in Computer Science*, pages 186–197. Springer, 2014.

[BMW56] M.J. Beckmann, C.B. McGuire, and C.B. Winsten. *Studies in the Economics of Transportation*. Cowles Commission for Research in Economics, 1956.

[BP25] Ziyad Benomar and Vianney Perchet. On tradeoffs in learning-augmented algorithms. In *AISTATS*, volume 258 of *Proceedings of Machine Learning Research*, pages 802–810. PMLR, 2025.

[BS94] David Bernstein and Tony E. Smith. Equilibria for networks with lower semicontinuous costs: With an application to congestion pricing. *Transp. Sci.*, 28(3):221–235, 1994.

[BV19] Vittorio Bilò and Cosimo Vinci. Dynamic taxes for polynomial congestion games. *ACM Trans. Economics and Comput.*, 7(3):15:1–15:36, 2019.

[Car13] Ioannis Caragiannis. Efficient coordination mechanisms for unrelated machine scheduling. *Algorithmica*, 66(3):512–540, 2013.

[CCG+15] Richard Cole, José R. Correa, Vasilis Gkatzelis, Vahab S. Mirrokni, and Neil Olver. Decentralized utilitarian mechanisms for scheduling games. *Games Econ. Behav.*, 92:306–326, 2015.

[CCSV26] George Christodoulou, Vasilis Christoforidis, Alkmini Sgouritsa, and Ioannis Vlachos. Improving the price of anarchy via predictions in parallel-link networks. In *WWW*, pages 309–317. ACM, 2026.

[CDN11] Johanne Cohen, Christoph Dürr, and Kim Thang Nguyen. Non-clairvoyant scheduling games. *Theory Comput. Syst.*, 49(1):3–23, 2011.

[CDR03] Richard Cole, Yevgeniy Dodis, and Tim Roughgarden. Pricing network edges for heterogeneous selfish users. In *STOC*, pages 521–530. ACM, 2003.

[CDR06] Richard Cole, Yevgeniy Dodis, and Tim Roughgarden. How much can taxes help selfish routing? *J. Comput. Syst. Sci.*, 72(3):444–467, 2006.

[CF19] Ioannis Caragiannis and Angelo Fanelli. An almost ideal coordination mechanism for unrelated machine scheduling. *Theory Comput. Syst.*, 63(1):114–127, 2019.

[CF21] Ioannis Caragiannis and Angelo Fanelli. On approximate pure nash equilibria in weighted congestion games with polynomial latencies. *J. Comput. Syst. Sci.*, 117:40–48, 2021.

[CFK+11] Ioannis Caragiannis, Michele Flammini, Christos Kaklamanis, Panagiotis Kanellopoulos, and Luca Moscardelli. Tight bounds for selfish and greedy load balancing. *Algorithmica*, 61(3):606–637, 2011.

[CG16] George Christodoulou and Martin Gairing. Price of stability in polynomial congestion games. *ACM Trans. Economics and Comput.*, 4(2):10:1–10:17, 2016.

[CGGS19] George Christodoulou, Martin Gairing, Yiannis Giannakopoulos, and Paul G. Spirakis. The price of stability of weighted congestion games. *SIAM J. Comput.*, 48(5):1544–1582, 2019.

[CGS24] George Christodoulou, Vasilis Gkatzelis, and Alkmini Sgouritsa. Resource-aware cost-sharing methods for scheduling games. *Oper. Res.*, 72(1):167–184, 2024.

[CGV17] Ioannis Caragiannis, Vasilis Gkatzelis, and Cosimo Vinci. Coordination mechanisms, cost-sharing, and approximation algorithms for scheduling. In *WINE*, volume 10660 of *Lecture Notes in Computer Science*, pages 74–87. Springer, 2017.

[CK05a] George Christodoulou and Elias Koutsoupias. On the price of anarchy and stability of correlated equilibria of linear congestion games. In *ESA*, volume 3669 of *Lecture Notes in Computer Science*, pages 59–70. Springer, 2005.

[CK05b] George Christodoulou and Elias Koutsoupias. The price of anarchy of finite congestion games. In *STOC*, pages 67–73. ACM, 2005.

[CKK10] Ioannis Caragiannis, Christos Kaklamanis, and Panagiotis Kanellopoulos. Taxes for linear atomic congestion games. *ACM Trans. Algorithms*, 7(1):13:1–13:31, 2010.

[CKN09] George Christodoulou, Elias Koutsoupias, and Akash Nanavati. Coordination mechanisms. *Theor. Comput. Sci.*, 410(36):3327–3336, 2009.

[CKS08] George Christodoulou, Elias Koutsoupias, and Paul G. Spirakis. On the performance of approximate equilibria in congestion games. *CoRR*, abs/0804.3160, 2008.

[CKS11] George Christodoulou, Elias Koutsoupias, and Paul G. Spirakis. On the performance of approximate equilibria in congestion games. *Algorithmica*, 61(1):116–140, 2011.

[CKS18] Riccardo Colini-Baldeschi, Max Klimm, and Marco Scarsini. Demand-independent optimal tolls. In *ICALP*, volume 107 of *LIPIcs*, pages 151:1–151:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.

[CKST24] Riccardo Colini-Baldeschi, Sophie Klumper, Guido Schäfer, and Artem Tsikiridis. To trust or not to trust: Assignment mechanisms with predictions in the private graph model. In *EC*, pages 1134–1154. ACM, 2024.

[CMP14] Giorgos Christodoulou, Kurt Mehlhorn, and Evangelia Pyrga. Improving the price of anarchy for selfish routing via coordination mechanisms. *Algorithmica*, 69(3):619–640, 2014.

[CSM04] José R. Correa, Andreas S. Schulz, and Nicolás E. Stier Moses. Selfish routing in capacitated networks. *Math. Oper. Res.*, 29(4):961–976, 2004.

[CSV24] George Christodoulou, Alkmini Sgouritsa, and Ioannis Vlachos. Mechanism design augmented with output advice. In *NeurIPS*, 2024.

[EADL24] Alex Elenter, Spyros Angelopoulos, Christoph Dürr, and Yanni Lefki. Overcoming brittleness in pareto-optimal learning augmented algorithms. In *NeurIPS*, 2024.

[FJM04] Lisa Fleischer, Kamal Jain, and Mohammad Mahdian. Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games. In *FOCS*, pages 277–285. IEEE Computer Society, 2004.

[FKS05] Dimitris Fotakis, Spyros C. Kontogiannis, and Paul G. Spirakis. Selfish unsplittable flows. *Theor. Comput. Sci.*, 348(2-3):226–239, 2005.

[Fot10] Dimitris Fotakis. Stackelberg strategies for atomic congestion games. *Theory Comput. Syst.*, 47(1):218–249, 2010.

[FS08] Dimitris Fotakis and Paul G. Spirakis. Cost-balancing tolls for atomic network congestion games. *Internet Math.*, 5(4):343–363, 2008.

[GKR16] Vasilis Gkatzelis, Konstantinos Kollias, and Tim Roughgarden. Optimal cost-sharing in general resource selection games. *Oper. Res.*, 64(6):1230–1238, 2016.

[GKST22] Vasilis Gkatzelis, Kostas Kollias, Alkmini Sgouritsa, and Xizhi Tan. Improved price of anarchy via predictions. In *EC*, pages 529–557. ACM, 2022.

[GNS22] Yiannis Giannakopoulos, Georgy Noarov, and Andreas S. Schulz. Computing approximate equilibria in weighted congestion games via best-responses. *Math. Oper. Res.*, 47(1):643–664, 2022.

[GP23] Yiannis Giannakopoulos and Diogo Poças. A unifying approximate potential for weighted congestion games. *Theory Comput. Syst.*, 67(4):855–876, 2023.

[GST25] Vasilis Gkatzelis, Daniel Schoepflin, and Xizhi Tan. Clock auctions augmented with unreliable advice. In *SODA*, pages 2629–2655. SIAM, 2025.

[HKS14] Christoph Hansknecht, Max Klimm, and Alexander Skopalik. Approximate pure nash equilibria in weighted congestion games. In *APPROX-RANDOM*, volume 28 of *LIPIcs*, pages 242–257. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2014.

[ILMS09] Nicole Immorlica, Li (Erran) Li, Vahab S. Mirrokni, and Andreas S. Schulz. Coordination mechanisms for selfish scheduling. *Theor. Comput. Sci.*, 410(17):1589–1598, 2009.

[Kas26] Danish Kashaev. Selfish, local and online scheduling via vector fitting. In *Proceedings of the 2026 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3841–3871, 2026.

[KK04] George Karakostas and Stavros G. Kolliopoulos. Edge pricing of multicommodity networks for heterogeneous selfish users. In *FOCS*, pages 268–276. IEEE Computer Society, 2004.

[Kol13] Konstantinos Kollias. Nonpreemptive coordination mechanisms for identical machines. *Theory Comput. Syst.*, 53(3):424–440, 2013.

[KP99] Elias Koutsoupias and Christos H. Papadimitriou. Worst-case equilibria. In *STACS*, volume 1563 of *Lecture Notes in Computer Science*, pages 404–413. Springer, 1999.

[LDK+15] Yisheng Lv, Yanjie Duan, Wenwen Kang, Zhengxi Li, and Fei-Yue Wang. Traffic flow prediction with big data: A deep learning approach. *IEEE Trans. Intell. Transp. Syst.*, 16(2):865–873, 2015.

[LM23] Alexander Lindermayr and Nicole Megow. Repository of works on algorithms with predictions. https://algorithms-with-predictions.github.io/, 2023.

[LWZ24] Pinyan Lu, Zongqi Wan, and Jialin Zhang. Competitive auctions with imperfect predictions. In *EC*, pages 1155–1183. ACM, 2024.

[MS96] Dov Monderer and Lloyd S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, 1996.

[MV22] Michael Mitzenmacher and Sergei Vassilvitskii. Algorithms with predictions. *Commun. ACM*, 65(7):33–35, 2022.

[NRTV07] Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.

[PG24] Dario Paccagnan and Martin Gairing. In congestion games, taxes achieve optimal approximation. *Oper. Res.*, 72(3):966–982, 2024.

[Ros73] Robert W Rosenthal. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory*, 2(1):65–67, December 1973.

[Rou02] Tim Roughgarden. *Selfish Routing*. PhD thesis, Cornell University, USA, 2002.

[Rou04] Tim Roughgarden. Stackelberg scheduling strategies. *SIAM J. Comput.*, 33(2):332–350, 2004.

[Rou15] Tim Roughgarden. Intrinsic robustness of the price of anarchy. *J. ACM*, 62(5):32:1–32:42, 2015.

[RT02] Tim Roughgarden and Éva Tardos. How bad is selfish routing? *J. ACM*, 49(2):236–259, 2002.

[RT04] Tim Roughgarden and Éva Tardos. Bounding the inefficiency of equilibria in nonatomic congestion games. *Games and Economic Behavior*, 47(2):389–403, may 2004.

[WYC16] Te-Li Wang, Chih-Kuan Yeh, and Ho-Lin Chen. An improved tax scheme for selfish routing. In *ISAAC*, volume 64 of *LIPIcs*, pages 61:1–61:12. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016.

A Further Related Work

Inefficiency of equilibria. Of particular relevance to our work is the study of the performance degradation resulting from selfish behavior and methods to mitigate it. Extensive work has characterized the PoA and PoS of congestion games; apart from the works already discussed in the introduction, bounds were shown for mixed and (coarse) correlated equilibria [CK05a, BHLR08, Rou15], as well as for approximate equilibria [CKS08] and weighted congestion games [BGR14, CGGS19, GP23].

(Approximate) Potentials. Since the work of [Ros73], that introduced the class of congestion games, a long line of work has been devoted to identifying classes of games that admit pure Nash equilibria. A key technical tool of Rosenthal's work was the use of a *potential* function which captures the gains in each player's deviating moves in an elegant way. Congestion games are *exact potential* games [MS96]. Beyond proving the existence of equilibria, potential functions can also provide important structural properties on the performance of potential minimizers.

This has motivated the development of *approximate* potential functions, which extend the power of potential-based arguments to approximate equilibria (see e.g. [CKS11, HKS14, CGGS19, CF21, GNS22, GP23]). Moreover, the well-known fact that *weighted* congestion games do not always have equilibria in pure strategies [FKS05] further underlined the need for approximate potentials. Minimizers of approximate potential functions are guaranteed to be approximate equilibria, in contrast to Rosenthal's case.

Coordination mechanisms. Coordination mechanisms were first introduced in [CKN09] as a method to mitigate the inefficiency of equilibria. They have since been extensively studied in the design of local policies across various scheduling settings and congestion games [ILMS09, CDN11, AH12, Car13, Kol13, ACH14, BIKM14, BKM14, CCG+15, AFJ+15, CF19, Kas26]. Specifically, [CMP14, WYC16] considered coordination mechanisms in the context of *selfish routing* over parallel-link networks. [CGV17] established intricate connections between cost-sharing protocols and coordination mechanisms, while [CGS24] applied the main principle of coordination mechanisms to cost-sharing games, aiming to improve social welfare by overcharging.

Coping with Selfishness. Besides coordination mechanisms, an important approach explored in the literature is the use of *taxation* and *tolls*, where an agent may be charged a tax that influences their perceived latency. The idea of introducing taxes to improve system efficiency has been extensively studied in transportation science (cf. [CDR03] and references therein), and gained significant traction in the algorithmic game theory literature; both in non-atomic games (see e.g., [FJM04, KK04, CDR06]) and atomic ones (e.g., [CKK10, FS08, BV19, PG24]). [CKS18] deviate from the marginal edge pricing approach and seek for optimal tolls without the exact knowledge of the demand. They show conditions under which demand-independent optimal tolls exist for any input rate in the network.

Another approach suggests that the game designer can select a fraction of players to be centrally routed using a *Stackelberg strategy* [Rou04, BHS10, Fot10]. *Signaling* strategies have also been used to leverage the designer's informational advantage and steer the traffic equilibrium to a more efficient outcome [BCKS16]. *Cost-sharing protocols* determine how a common resource is allocated among players [GKR16].

Learning-augmented framework. The learning-augmented framework has gained significant attention in the design of algorithms beyond the worst-case. Rather than surveying the growing literature, we highlight a few representative works on canonical problems such as auctions, strategic facility location, and strategic scheduling [ABG+24, BGT23b, CKST24, CSV24, LWZ24, GST25]. In the direction of improving the PoA, [GKST22] incorporated demand predictions into the design of cost-sharing mechanisms for scheduling games. For a broader overview, we refer the interested reader to [MV22, BGT23a] and the references therein, and to [LM23] for a frequently updated online repository of works in this area.

The recent paper [CCSV26] is closely related to our work. They also explore the use of demand predictions in the design of coordination mechanisms for games with polynomial latencies, but their results are incomparable with ours. The focus of [CCSV26] is on *non-atomic* congestion games and on *parallel-link* networks, while our results apply to a much more general setting; to *general* congestion games, for both atomic and non-atomic variants. However, they are able to obtain stronger results for the stronger definitions of consistency and robustness. They are able to show an improved PoA-Robustness upper bound of $d+1$ for 1-PoA-Consistent mechanisms and they also show that this result is tight. As we show in Theorems 6 and 5 (for atomic and non-atomic congestion games, respectively), it is impossible to achieve this in general congestion games.

B Missing Proofs of Section 3

B.1 The POTENTIAL Mechanism

Proof of Lemma 1

Proof. Let $\hat{\Phi}$, Φ^h, and Φ^ℓ be the potentials induced by the latencies $\hat{\ell}$, h, and ℓ, respectively. Let f be any feasible flow for demand $\bar{r}$. Define E_h to be the set of edges for which the increased latencies h are triggered, i.e., $f_e > \bar{f}_e$, and let $E_\ell = E \setminus E_h$ be the rest of the edges, for which the original latencies are used. The potential of f is bounded as follows:

$$
\begin{aligned}
\hat{\Phi}(f) &\geq \hat{\Phi}(f) + \Phi^h(\bar{f}) - \Phi^h(f) \\
&= \sum_{e\in E_h} \left(\hat{\Phi}_e(f) + \Phi^h_e(\bar{f}) - \Phi^h_e(f)\right) + \sum_{e\in E_\ell} \left(\hat{\Phi}_e(f) + \Phi^h_e(\bar{f}) - \Phi^h_e(f)\right) \\
&= \sum_{e\in E_h} \left(\left(\int_0^{\bar{f}_e} \ell_e(t)dt + \int_{\bar{f}_e}^{f_e} h_e(t)dt\right) - \int_{\bar{f}_e}^{f_e} h_e(t)dt\right) + \sum_{e\in E_\ell} \left(\Phi^\ell_e(f) + \Phi^h_e(\bar{f}) - \Phi^h_e(f)\right) \\
&= \sum_{e\in E_h} \Phi^\ell_e(\bar{f}) + \sum_{e\in E_\ell} \left(\Phi^\ell_e(f) + \int_{f_e}^{\bar{f}_e} h_e(t)dt\right) \\
&= \sum_{e\in E_h} \Phi^\ell_e(\bar{f}) + \sum_{e\in E_\ell} \left(\int_0^{f_e} \ell_e(t)dt + \int_{f_e}^{\bar{f}_e} h_e(t)dt\right) \\
&\geq \sum_{e\in E_h} \Phi^\ell_e(\bar{f}) + \sum_{e\in E_\ell} \Phi^\ell_e(\bar{f}) \\
&= \hat{\Phi}(\bar{f}),
\end{aligned}
$$

where the first inequality is due to the fact that $\bar{f}$ minimizes Φ^h, and the second due to $h_e(x) \geq \ell_e(x)$ for all $e \in E$ and $x \geq 0$. The last equality is due to the fact that the profile $\bar{f}$ never triggers the increased latencies h, hence $\hat{\Phi}(\bar{f}) = \Phi^\ell(\bar{f})$.

The inequality above shows that $\hat{\Phi}(f) \geq \hat{\Phi}(\bar{f})$ for every feasible flow f for demand $\bar{r}$. Therefore, $\bar{f}$ is a minimizer of $\hat{\Phi}$ for demand $\bar{r}$. Moreover, since $\bar{f}$ never triggers the increased latencies, we also have $\hat{C}(\bar{f}) = C(\bar{f})$. □

Proof of Lemma 2

Proof. We prove the equivalent path-based inequality

$$\sum_{p \in \mathcal{P}} \hat{\ell}_p(f) f_p \leq \sum_{p \in \mathcal{P}} h_p(f) f^*_p \,.$$

For each commodity i, let $p_i \in \arg\min_{p \in \mathcal{P}_i} h_p(f)$. Since f is an equilibrium with respect to $\hat{\ell}$, then for any path $p \in \mathcal{P}_i$ with $f_p > 0$, it holds that $\hat{\ell}_p(f) \leq \liminf_{\epsilon \downarrow 0} \hat{\ell}_{p_i}(f + \epsilon \mathbf{1}_{p_i} - \epsilon \mathbf{1}_p) \leq h_{p_i}(f)$. Then,

$$\begin{aligned}
\sum_{p \in \mathcal{P}} \hat{\ell}_p(f) f_p &\leq \sum_{i \in \mathcal{I}} \sum_{p \in \mathcal{P}_i} h_{p_i}(f) f_p \\
&= \sum_{i \in \mathcal{I}} h_{p_i}(f) \sum_{p \in \mathcal{P}_i} f_p \\
&= \sum_{i \in \mathcal{I}} h_{p_i}(f) \sum_{p \in \mathcal{P}_i} f^*_p \\
&\leq \sum_{i \in \mathcal{I}} \sum_{p \in \mathcal{P}_i} h_p(f) f^*_p \\
&= \sum_{p \in \mathcal{P}} h_p(f) f^*_p \,,
\end{aligned}$$

where the second equality is due to the fact that $\sum_{p \in \mathcal{P}_i} f_p = \sum_{p \in \mathcal{P}_i} f^*_p = r_i$, for all $i \in \mathcal{I}$, and the last inequality is derived by the definition of p_i. □

B.2 Warmup - Affine Latencies

In this section, we provide consistency and robustness analysis of the POTENTIAL mechanism for the case of affine latency functions of the form $\ell_e(x) = a_e x + b_e$, where a_e and b_e are non-negative coefficients, specific to each resource e. Let β be the confidence parameter and $\bar{r}$ be the predicted demand, then the threshold function $h(\beta)$ used by the POTENTIAL$(\beta, \bar{r})$ boils down to $h_e(x) = \beta a_e x + b_e$, for all $e \in E$, and the respective potential to

$$\Phi^\beta(f) = \sum_{e \in E} \int_{x=0}^{f_e} \frac{h_e(x)}{\beta} dx = \sum_{e \in E} \left(\frac{1}{2} a_e f_e^2 + \frac{1}{\beta} b_e f_e \right).$$

Let $\bar{f}$ be a potential minimizer of Φ^h for demand equal to $\bar{r}$. Then the modified latencies of POTENTIAL$(\beta, \bar{r})$ are defined as follows

$$\hat{\ell}_e(x) = \begin{cases} \beta a_e x + b_e, & x > \bar{f}_e \\ a_e x + b_e, & \text{otherwise.} \end{cases}$$

Corollary 1 (Restated)**.** *Given the predicted input rates $\bar{r}$ and a parameter $\beta \in [1, 2]$, the* POTENTIAL$(\beta,\bar{r})$ *mechanism for affine latency functions is $\frac{4}{(4-\beta)\beta}$-consistent and 2β-robust.*

Proof. **Consistency.** We first show the consistency bound. Suppose that the prediction is accurate, i.e., $\bar{r} = r$. Let f^* be the optimal flow w.r.t. the original latencies ℓ. By Lemma 1, the potential minimizer $\bar{f}$ is also a minimizer of $\hat{\Phi}$ and preserves its original social cost, i.e., $\hat{C}(\bar{f}) = C(\bar{f})$. Moreover, $\bar{f}$ minimizes Φ^h. Equivalently, it minimizes Φ^h/β, which is the approximate potential used in [CKS11]. Therefore, by the bound of [CKS11] for β-approximate Wardrop equilibria,[9]

[9] This approach was used in [CKS11] to bound the PoS of β-approximate Wardrop equilibria.

$C(\bar{f}) \leq \frac{4}{(4-\beta)\beta} C(f^*)$. Combining this with $\hat{C}(\bar{f}) = C(\bar{f})$ yields $\hat{C}(\bar{f}) \leq \frac{4}{(4-\beta)\beta} C(f^*)$, which gives the required consistency bound.

Robustness. Let f be any feasible flow and let E_h be the set of edges for which $f_e > \bar{f}_e$, i.e., $\hat{\ell}_e(f_e) = h_e(f_e)$, and let $E_\ell \subseteq E \setminus E_h$ be the rest of the edges, for which the original latencies are used. Then, $\hat{\Phi}(f)$ can be written as

$$\begin{aligned}
\hat{\Phi}(f) &= \sum_{e \in E} \int_0^{f_e} \hat{\ell}_e(t) dt \\
&= \sum_{e \in E_\ell} \int_0^{f_e} (a_e t + b_e) dt + \sum_{e \in E_h} \left(\int_0^{\bar{f}_e} (a_e t + b_e) dt + \int_{\bar{f}_e}^{f_e} (\beta a_e t + b_e) dt \right) \\
&= \sum_{e \in E_\ell} \left(\frac{1}{2} a_e f_e^2 + b_e f_e \right) + \sum_{e \in E_h} \left(\frac{\beta}{2} a_e f_e^2 + b_e f_e - \frac{\beta - 1}{2} a_e \bar{f}_e^2 \right) \\
&\leq \sum_{e \in E} a_e f_e^2 + b_e f_e = C(f) \,,
\end{aligned}$$

where the inequality uses $\beta \leq 2$ and $\bar{f}_e < f_e$ for every $e \in E_h$.

Moreover,

$$\begin{aligned}
\hat{\Phi}(f) &= \sum_{e \in E_\ell} \left(\frac{1}{2} a_e f_e^2 + b_e f_e \right) + \sum_{e \in E_h} \left(\frac{\beta}{2} a_e f_e^2 + b_e f_e - \frac{\beta - 1}{2} a_e \bar{f}_e^2 \right) \\
&\geq \sum_{e \in E_\ell} \left(\frac{1}{2} a_e f_e^2 + b_e f_e \right) + \sum_{e \in E_h} \left(\frac{1}{2} a_e f_e^2 + b_e f_e \right) \\
&\geq \frac{1}{2\beta} \left(\sum_{e \in E_\ell} (a_e f_e^2 + b_e f_e) + \sum_{e \in E_h} (\beta a_e f_e^2 + b_e f_e) \right) \\
&= \frac{1}{2\beta} \hat{C}(f) \,.
\end{aligned}$$

Combining the above inequalities yields $\frac{1}{2\beta} \hat{C}(f^\diamond) \leq \hat{\Phi}(f^\diamond) \leq \hat{\Phi}(f^*) \leq C(f^*)$, since they hold for any f, and therefore, for $f^\diamond$ and f^* too, while the middle inequality follows from the fact that $f^\diamond$ is a potential minimizer. This establishes the desired robustness bound and concludes the proof. □

B.3 Bounded-degree Polynomials

We now show how to extend our analysis to the general case where latency functions are polynomials of maximum degree d. Let $\beta \in [1, d+1]$. The consistency argument remains unchanged, hence, it is equal to the PoS of β-approximate equilibria in non-atomic congestion games for polynomial latency functions. Regarding robustness, we bound the inefficiency of potential minimizers to get a guarantee for the PoS for the modified game.

Proof of Theorem 1

Proof. **Consistency.** We first establish the consistency upper bound. Similarly to the proof for the affine latency functions, we obtain that the consistency is equal to the PoS of the best β-approximate equilibrium. Suppose that the prediction is accurate, i.e., $\bar{r} = r$. The POTENTIAL(β,$\bar{r}$) mechanism modifies the latencies according to THRESHOLD($h, \bar{f}$), for

$$h_e(x) = \sum_{k=0}^{\lfloor \beta \rfloor - 1} (k+1) a_{e,k} x^k + \sum_{k=\lfloor \beta \rfloor}^{d} \beta a_{e,k} x^k \,, \qquad \forall e \in E,$$

and $\bar{f}$ is a potential minimizer of Φ^h for input rate $\bar{r} = r$. By Lemma 1, $\bar{f}$ is a minimizer of $\hat{\Phi} = \Phi^{\hat{\ell}}$ and preserves its original social cost, i.e., $\hat{C}(\bar{f}) = C(\bar{f})$. Moreover, $\bar{f}$ minimizes Φ^h. In [CKS11] it

is shown that any flow f that minimizes the approximate potential

$$\Phi^{\beta}(f) = \sum_{e \in E} \int_0^{f_e} \frac{h_e(x)}{\beta} dx$$

has social cost w.r.t. the original latencies ℓ at most a factor of $\alpha = \left(\beta\left(1 - \frac{d}{d+1}\left(\frac{\beta}{d+1}\right)^{1/d}\right)\right)^{-1}$ away from the optimal cost. Note that $\bar{f}$ also minimizes $\Phi^{\beta} = \Phi^{h}/\beta$, and since $\hat{C}(\bar{f}) = C(\bar{f})$, the mechanism is α-consistent.

Robustness. Let f be any feasible flow and let E_h be the set of edges for which $f_e > \bar{f}_e$, i.e. $\hat{\ell}_e(f_e) = h_e(f_e)$ and let $E_\ell = E \setminus E_h$ be the rest of the edges, where the original latencies are used. Then $\hat{\Phi}(f)$ can be written as

$$\hat{\Phi}(f) = \sum_{e \in E} \int_0^{f_e} \hat{\ell}_e(t)dt = \sum_{e \in E_\ell} \int_0^{f_e} \ell_e(t)dt + \sum_{e \in E_h} \left(\int_0^{\bar{f}_e} \ell_e(t)dt + \int_{\bar{f}_e}^{f_e} h_e(t)dt \right).$$

We split the rest of the proof into two claims. We first show that $\hat{\Phi}(f) \leq C(f)$ and then that $\hat{\Phi}(f) \geq \frac{1}{\beta(d+1)}\hat{C}(f)$. We finally combine them to get the robustness bound.

Claim 1. $\hat{\Phi}(f) \leq C(f)$.

Proof. The part of the potential containing edges in E_h can be written as follows

$$\begin{aligned}
\sum_{e \in E_h} \left(\int_0^{\bar{f}_e} \ell_e(t)dt + \int_{\bar{f}_e}^{f_e} h_e(t)dt \right) &\leq \sum_{e \in E_h} \int_0^{f_e} h_e(t)dt \\
&= \sum_{e \in E_h} \sum_{k=0}^{\lfloor \beta \rfloor - 1} (k+1) \int_0^{f_e} a_{e,k} t^k dt + \sum_{e \in E_h} \sum_{k=\lfloor \beta \rfloor}^{d} \beta \int_0^{f_e} a_{e,k} t^k dt \\
&= \sum_{e \in E_h} \sum_{k=0}^{\lfloor \beta \rfloor - 1} a_{e,k} f_e^{k+1} + \sum_{e \in E_h} \sum_{k=\lfloor \beta \rfloor}^{d} \frac{\beta}{k+1} a_{e,k} f_e^{k+1} \\
&\leq \sum_{e \in E_h} \sum_{k=0}^{d} a_{e,k} f_e^{k+1}
\end{aligned}$$

Note that for every index $k \geq \lfloor \beta \rfloor$, we have $k+1 > \beta$, hence $\frac{\beta}{k+1} \leq 1$, which justifies the inequality above. For the edges in E_ℓ, it holds that:

$$\sum_{e \in E_\ell} \int_0^{f_e} \ell_e(t)dt = \sum_{e \in E_\ell} \sum_{k=0}^{d} \frac{a_{e,k}}{k+1} f_e^{k+1} \leq \sum_{e \in E_\ell} \sum_{k=0}^{d} a_{e,k} f_e^{k+1}$$

Overall, the two inequalities above give

$$\hat{\Phi}(f) \leq \sum_{e \in E} \sum_{k=0}^{d} a_{e,k} f_e^{k+1} = C(f).$$

□

Claim 2. $\hat{\Phi}(f) \geq \frac{1}{\beta(d+1)}\hat{C}(f)$.

Proof. We have that

$$\begin{aligned}
\hat{\Phi}(f) &\geq \sum_{e\in E_\ell}\sum_{k=0}^{d}\frac{1}{\beta(d+1)}a_{e,k}f_e^{k+1} \\
&\quad+\sum_{e\in E_h}\left(\sum_{k=0}^{\lfloor\beta\rfloor-1}\frac{1}{\beta(d+1)}(k+1)a_{e,k}f_e^{k+1}+\sum_{k=\lfloor\beta\rfloor}^{d}\frac{1}{\beta(d+1)}\beta a_{e,k}f_e^{k+1}\right) \\
&= \frac{1}{\beta(d+1)}\left(\sum_{e\in E_\ell}\sum_{k=0}^{d}a_{e,k}f_e^{k+1}+\sum_{e\in E_h}\left(\sum_{k=0}^{\lfloor\beta\rfloor-1}(k+1)a_{e,k}f_e^{k+1}+\sum_{k=\lfloor\beta\rfloor}^{d}\beta a_{e,k}f_e^{k+1}\right)\right) \\
&= \frac{1}{\beta(d+1)}\hat{C}(f)\,.
\end{aligned}$$

The inequality uses $\bar{f}_e < f_e$ for every $e \in E_h$, $\beta \geq 1$, and $k+1 \leq d+1$ for every $k \in \{0,\ldots,d\}$. □

Let $f^\diamond$ be any potential minimizer of $\hat{\Phi}$, and f^* be an optimal profile w.r.t. the original latency functions, i.e., f^* minimizes $C(f)$. Then, by utilizing the above inequalities that hold for any f, and therefore for $f^\diamond$ and f^*, it follows that $\frac{1}{\beta(d+1)}\hat{C}(f^\diamond) \leq \hat{\Phi}(f^\diamond) \leq \hat{\Phi}(f^*) \leq C(f^*)$, which concludes the proof for robustness, and so completes the proof of the theorem. □

B.3.1 A threshold-sensitive robustness refinement

We present a refined threshold-sensitive analysis of the POTENTIAL mechanism, depending on how close the triggered edges are to their predicted thresholds. Let f be any feasible flow and let $E_h(f) = \{e \in E : f_e > \bar{f}_e\}$ be the set of triggered edges. For every $e \in E_h(f)$, define $\eta_e = \bar{f}_e/f_e \in [0,1)$. If $E_h(f) \neq \emptyset$, let $\eta(f,\bar{f}) = \max_{e\in E_h(f)}\eta_e$, and otherwise set $\eta(f,\bar{f}) = 0$. Recall that $\hat{\Phi}_e(f) = \int_0^{f_e}\hat{\ell}_e(t)\,dt$, $\hat{\Phi}(f) = \sum_{e\in E}\hat{\Phi}_e(f)$, and $\hat{C}_e(f) = f_e\hat{\ell}_e(f_e)$.

We will use two elementary facts. First, $h_e(x) \leq \beta\ell_e(x)$ for every edge e and every $x \geq 0$. Indeed, by the definition of h_e, each monomial coefficient is either multiplied by β, or by $k+1$ with $k+1 \leq \beta$; all coefficients are nonnegative. Second, since h_e is a polynomial of degree at most d with nonnegative coefficients, writing $h_e(x) = \sum_{k=0}^{d} b_{e,k}x^k$ gives $\Phi_e^h(x) = \int_0^x h_e(t)\,dt = \sum_{k=0}^{d} b_{e,k}x^{k+1}/(k+1) \geq xh_e(x)/(d+1)$.

Lemma B.1 (Threshold-sensitive potential lower bound)**.** *For every feasible flow f,*

$$\hat{\Phi}(f) \geq \frac{1-\left(1-\frac{1}{\beta}\right)\eta(f,\bar{f})}{d+1}\,\hat{C}(f).$$

Proof. We prove the inequality edge by edge. First consider an untriggered edge $e \notin E_h(f)$, so $f_e \leq \bar{f}_e$. Then $\hat{\ell}_e(f_e) = \ell_e(f_e)$, and hence $\hat{\Phi}_e(f) = \int_0^{f_e}\ell_e(t)\,dt \geq f_e\ell_e(f_e)/(d+1) = \hat{C}_e(f)/(d+1)$, where the inequality uses that ℓ_e is a polynomial of degree at most d with nonnegative coefficients. Since $1-(1-1/\beta)\eta(f,\bar{f}) \leq 1$, this already implies the desired edge-wise bound for every untriggered edge.

Now consider a triggered edge $e \in E_h(f)$, so $f_e > \bar{f}_e$. Then

$$\begin{aligned}
\hat{\Phi}_e(f) &= \int_0^{\bar{f}_e}\ell_e(t)\,dt+\int_{\bar{f}_e}^{f_e}h_e(t)\,dt \\
&= \Phi_e^\ell(\bar{f}_e)+\Phi_e^h(f_e)-\Phi_e^h(\bar{f}_e) \\
&\geq \Phi_e^h(f_e)-\left(1-\frac{1}{\beta}\right)\Phi_e^h(\bar{f}_e) \\
&\geq \Phi_e^h(f_e)-\left(1-\frac{1}{\beta}\right)\eta_e\Phi_e^h(f_e)
\end{aligned}$$

$$
\begin{aligned}
&= \left(1 - \left(1 - \frac{1}{\beta}\right)\eta_e\right)\Phi_e^h(f_e) \\
&\geq \frac{1 - \left(1 - \frac{1}{\beta}\right)\eta_e}{d+1} f_e h_e(f_e) \\
&= \frac{1 - \left(1 - \frac{1}{\beta}\right)\eta_e}{d+1} \hat{C}_e(f).
\end{aligned}
$$

The first inequality uses $h_e(t) \leq \beta \ell_e(t)$, which implies $\Phi_e^\ell(\bar{f}_e) \geq \Phi_e^h(\bar{f}_e)/\beta$. The second inequality uses $\Phi_e^h(\bar{f}_e) \leq \eta_e \Phi_e^h(f_e)$: indeed, Φ_e^h has nonnegative coefficients and zero constant term, and $(\bar{f}_e/f_e)^q \leq \bar{f}_e/f_e = \eta_e$ for every $q \geq 1$. The last inequality uses $\Phi_e^h(f_e) \geq f_e h_e(f_e)/(d+1)$. Since $\eta_e \leq \eta(f, \bar{f})$ and $1 - (1 - 1/\beta)\eta$ is decreasing in η, this gives

$$
\hat{\Phi}_e(f) \geq \frac{1 - \left(1 - \frac{1}{\beta}\right)\eta(f, \bar{f})}{d+1} \hat{C}_e(f)
$$

for every triggered edge. Summing over all edges completes the proof. □

Corollary B.1 (Threshold-sensitive robustness)**.** *Let $f^\diamond$ be any potential minimizer of the modified game, and let f^* be an optimal flow under the original latencies. Then*

$$
\hat{C}(f^\diamond) \leq \frac{d+1}{1 - \left(1 - \frac{1}{\beta}\right)\eta(f^\diamond, \bar{f})} C(f^*).
$$

Proof. By Lemma B.1,

$$
\frac{1 - \left(1 - \frac{1}{\beta}\right)\eta(f^\diamond, \bar{f})}{d+1} \hat{C}(f^\diamond) \leq \hat{\Phi}(f^\diamond).
$$

Since $f^\diamond$ minimizes $\hat{\Phi}$, and since the proof of Theorem 1 shows that $\hat{\Phi}(f) \leq C(f)$ for every feasible flow f, we have $\hat{\Phi}(f^\diamond) \leq \hat{\Phi}(f^*) \leq C(f^*)$. Combining the two inequalities and rearranging gives the claim. □

This refinement interpolates between the strongest robustness bound $d + 1$, obtained when no triggered edge is close to its threshold, and the worst-case guarantee $\beta(d + 1)$ of Theorem 1 as $\eta(f^\diamond, \bar{f})$ approaches 1. Thus, the robustness loss depends not only on the confidence parameter β, but also on how close the realized potential minimizer is to the prediction-derived threshold flow on the edges where the mechanism is triggered.

B.4 PoA-Robustness

In this section, we analyze the PoA-Robustness of POTENTIAL, first for the case of affine latencies (Corollary 2) and subsequently for the case of polynomial latencies (Theorem 2).

Corollary 2 (Restated)**.** *Given the predicted input rates $\bar{r}$ and a parameter $\beta \in [1, 2]$, the* POTENTIAL(β,$\bar{r}$) *mechanism for affine latency functions has PoA-Robustness $\frac{4}{4-\beta^2}$ for $\beta \leq \sqrt{2}$, and β^2 otherwise.*

Proof of Corollary 2

Proof. Given any input rate r, let f be an equilibrium flow under the modified latencies $\hat{\ell}$ and let f^* be the optimal flow under the original latencies ℓ. By using the Lifting Lemma (Lemma 2) and substituting for $h_e(x) = \beta a_e x + b_e$, we obtain

$$
\hat{C}(f) = \sum_{e \in E} \hat{\ell}_e(f_e) f_e \leq \sum_{e \in E} h_e(f_e) f_e^* = \sum_{e \in E} (\beta a_e f_e f_e^* + b_e f_e^*). \tag{1}
$$

A quick —but not tight— upper bound of PoA-Robustness can be obtained as follows. The last term of (1) can be bounded above by $\beta \sum_{e \in E} \ell_e(f_e) f_e^*$, which gives $\hat{C}(f) \leq \beta \sum_{e \in E} \ell_e(f_e) f_e^*$.

By observing that the above inequality is also satisfied by β-approximate equilibria w.r.t. the original latencies, one can follow a similar analysis with that of [CKS11] to upper bound the PoA of β-approximate equilibria. This gives a PoA-Robustness bound of at most $\frac{4\beta}{4-\beta}$, for $\beta \in [1, 2]$.

However, we are able to obtain an improved bound for all $\beta \in (1, 2)$ (see Figure 1.(a) for a comparison) by a more careful analysis. For this, we need to establish the following claim.

Claim 3. *For any parameter $\lambda \geq 1/\beta$ it holds that*

$$\hat{C}(f) \leq \frac{\beta}{4\lambda}\hat{C}(f) + \lambda\beta C(f^*)\,.$$

Proof. By the first inequality of (1) and the fact that for any real numbers x, y, λ it holds that $xy \leq \frac{1}{4\lambda}x^2 + \lambda y^2$ (Lemma 2 in [CKS08]), we get

$$\begin{aligned}
\hat{C}(f) &\leq \sum_{e\in E}(\beta a_e f_e f_e^* + b_e f_e^*) \\
&\leq \sum_{e\in E}\left(\beta a_e\left(\frac{1}{4\lambda}(f_e)^2 + \lambda(f_e^*)^2\right) + b_e f_e^*\right) \\
&\leq \frac{\beta}{4\lambda}\sum_{e\in E}(a_e(f_e)^2 + b_e f_e) + \lambda\beta\sum_{e\in E}(a_e(f_e^*)^2 + b_e f_e^*) \qquad \text{(since } \lambda\beta \geq 1\text{)} \\
&\leq \frac{\beta}{4\lambda}\hat{C}(f) + \lambda\beta C(f^*)\,.
\end{aligned}$$

□

Rearranging yields the following bound on the robustness:

$$\frac{\hat{C}(f)}{C(f^*)} \leq \frac{\lambda\beta}{1-\frac{\beta}{4\lambda}} = \frac{4\lambda^2\beta}{4\lambda-\beta}\,.$$

We appropriately choose $\lambda \geq 1/\beta$ to minimize the above expression, which is $\lambda = 1/\beta$ for $\beta \leq \sqrt{2}$, and $\lambda = \beta/2$ otherwise. This overall results in robustness of $\frac{4}{4-\beta^2}$ for $\beta \leq \sqrt{2}$, and robustness of β^2 otherwise. We remark that this bound works for *any* equilibrium flow under $\hat{\ell}$. □

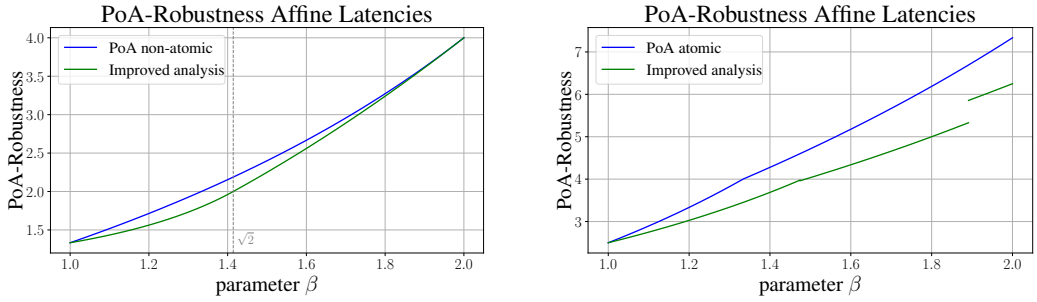

(a) Non-atomic (b) Atomic

Figure 1: PoA-Robustness guarantees for affine latency functions. The left panel corresponds to the non-atomic setting and the right one to the atomic setting. In green, we highlight the improved bounds obtained by our analysis, compared to the PoA for β-approximate equilibria.

Proof of Theorem 2

Proof. Given any input rate r, suppose that f is the equilibrium flow under the modified latencies $\hat{\ell}$, and f^* is the optimal flow under the original latencies ℓ. By the Lifting Lemma 2 we get:

$$\hat{C}(f) = \sum_{e \in E} \hat{\ell}_e(f_e) f_e \leq \sum_{e \in E} h_e(f_e) f_e^* \,.$$

Similarly to the affine latency case, the last term can be bounded above by $\beta \sum_{e \in E} \ell_e(f_e) f_e^*$, which results in $\hat{C}(f) \leq \beta \sum_{e \in E} \ell_e(f_e) f_e^*$. One can thus follow the analysis to upper bound the PoA of β-approximate equilibria provided by [CKS11] and obtain a PoA-Robustness bound of at most $\left(\frac{1}{\beta} - \frac{d}{(d+1)^{1+\frac{1}{d}}}\right)^{-1}$ for $\beta \leq \sqrt[d]{d+1}$, and β^{d+1} for $\sqrt[d]{d+1} < \beta \leq d+1$. However, we are able to provide a tighter analysis for our setting that improves this bound for all $\beta \in (1, \sqrt[d]{d+1})$. For this, we show the following claim:

Claim 4. *For any parameter $\lambda \geq 1/\beta$ it holds that*

$$\hat{C}(f) \leq \max\left\{\frac{1}{4\lambda}, \frac{\beta d}{(d+1)\sqrt[d]{\lambda(d+1)}}\right\} \hat{C}(f) + \lambda\beta C(f^*) \,.$$

Proof. First, consider the inequality stated in Lemma 5 (see Appendix E). Combining this result with Lemma 2, and noting that $\lambda\beta \geq 1$, we obtain

$$\begin{aligned}
\hat{C}(f) &\leq \sum_{e \in E} h_e(f_e) f_e^* \\
&= \sum_{e \in E} \left(\sum_{k=0}^{\lfloor \beta \rfloor - 1} (k+1) a_{e,k} f_e^k f_e^* + \sum_{k=\lfloor \beta \rfloor}^{d} \beta a_{e,k} f_e^k f_e^* \right) \\
&\leq \sum_{e \in E} \left(a_{e,0} f_e^* + \sum_{k=1}^{\lfloor \beta \rfloor - 1} \beta a_{e,k} f_e^k f_e^* + \sum_{k=\lfloor \beta \rfloor}^{d} \beta a_{e,k} f_e^k f_e^* \right) \\
&\leq \sum_{e \in E} \left(\sum_{k=0}^{d} \beta a_{e,k} \max\left\{\frac{1}{4\lambda}, \frac{d}{(d+1)\sqrt[d]{\lambda(d+1)}}\right\} f_e^{k+1} + \sum_{k=0}^{d} \beta a_{e,k} \lambda (f_e^*)^{k+1} \right) \\
&\leq \beta \max\left\{\frac{1}{4\lambda}, \frac{d}{(d+1)\sqrt[d]{\lambda(d+1)}}\right\} \hat{C}(f) + \lambda\beta C(f^*) \,.
\end{aligned}$$

□

Rearranging yields the following bound on the PoA-Robustness:

$$\frac{\hat{C}(f)}{C(f^*)} \leq \frac{\lambda\beta}{1 - \beta \max\left\{\frac{1}{4\lambda}, \frac{d}{(d+1)\sqrt[d]{\lambda(d+1)}}\right\}} \,.$$

For $d = 1$, it holds that $\frac{1}{4\lambda} = \frac{d}{(d+1)\sqrt[d]{\lambda(d+1)}}$, and we derive the bounds of Corollary 2.

Suppose now that $d \geq 2$. Define $\lambda_0 = (\frac{1}{4d})^{\frac{d}{d-1}} (d+1)^{\frac{d+1}{d-1}}$ to be the value of λ such that $\frac{1}{4\lambda} = \frac{d}{(d+1)\sqrt[d]{\lambda(d+1)}}$. Note that for $\lambda \leq \lambda_0$ it is $\frac{1}{4\lambda} \geq \frac{\beta d}{(d+1)\sqrt[d]{\lambda(d+1)}}$, and vice-versa for $\lambda > \lambda_0$. We appropriately choose the value of $\lambda \geq 1/\beta$ to minimize the above expression. Observe that in case $\beta < \frac{1}{\lambda_0}$, it holds that $\lambda \geq \frac{1}{\beta} > \lambda_0$, and the above expression becomes $\frac{\lambda\beta}{1 - \frac{\beta d}{(d+1)\sqrt[d]{\lambda(d+1)}}}$. This is

minimized for $\lambda = 1/\beta$ when $\beta \in [1, (d+1)^{\frac{1}{d+1}})$ and $\lambda = \frac{\beta^d}{d+1}$ when $\beta \in [(d+1)^{\frac{1}{d+1}}, \frac{1}{\lambda_0})$. For $\beta \in [\frac{1}{\lambda_0}, d+1]$, it holds that $\beta \geq 1/\lambda_0 \geq 4\lambda_0$ and therefore if $\lambda \leq \lambda_0$ the PoA-Robustness is negative. Thus, again the expression is optimized for $\lambda = \frac{\beta^d}{d+1} \geq \lambda_0$. Overall, this gives:

$$\lambda = \begin{cases} \frac{1}{\beta}, & \text{for } \beta \in [1, (d+1)^{\frac{1}{d+1}}) \\ \frac{\beta^d}{d+1}, & \text{for } \beta \in [(d+1)^{\frac{1}{d+1}}, d+1] \end{cases}$$

In total, PoA-Robustness is bounded by:

$$\frac{\hat{C}(f)}{C(f^*)} \leq \begin{cases} \left(1 - \beta \frac{d}{d+1} \left(\frac{\beta}{d+1}\right)^{1/d}\right)^{-1}, & \text{for } \beta \in [1, (d+1)^{\frac{1}{d+1}}) \\ \beta^{d+1}, & \text{for } \beta \in [(d+1)^{\frac{1}{d+1}}, d+1) \end{cases}$$

□

C Atomic Congestion Games

In this section, we switch to the class of atomic congestion games with polynomial latency functions. As in the non-atomic case, we propose the POTENTIAL mechanism and analyze its performance. Given an instance of a congestion game augmented with a prediction of the demand $(\bar{N}, E, (\bar{\mathcal{S}}_i)_{i \in \bar{N}}, (\ell_e)_{e \in E})$, where $\bar{r} = (\bar{N}, (\bar{\mathcal{S}}_i)_{i \in \bar{N}})$ denotes the predicted demand, the POTENTIAL mechanism modifies the latencies according to a THRESHOLD coordination mechanism.

The main result of this section offers a trade-off between consistency and robustness, tuned by a confidence parameter β, and is summarized in the next theorem.

Theorem 3. *Given a predicted input demand $\bar{r}$ and a parameter $\beta \in \left[\frac{2(d+1)}{d+3}, d+1\right]$, the* POTENTIAL($\beta$,$\bar{r}$) *mechanism in atomic congestion games has*

$$\textit{Consistency: } \frac{d+1}{\beta}\,; \qquad \textit{Robustness: } \beta(d+1)\,.$$

For the case of affine latency functions, we are able to get an improved analysis for the consistency bound.

Corollary 3. *Given the predicted input rates $\bar{r}$ and a parameter $\beta \in [1, 2]$, the* POTENTIAL(β,$\bar{r}$) *mechanism is $\frac{\sqrt{3}+1}{\sqrt{3}+\beta-1}$-consistent, and 2β-robust for affine latency functions.*

As in the case of non-atomic congestion games, we additionally provide a stronger PoA-Robustness guarantee that applies to the worst-case social cost across all equilibria, not only to potential minimizers.

Theorem 4. *Given a predicted demand $\bar{r}$ and a parameter $\beta \in [2(d+1)/(d+3), d+1]$, the PoA-Robustness of* POTENTIAL(β,$\bar{r}$) *for atomic congestion games is at most $\phi_{d,\beta}^{d+1}$, where $\phi_{d,\beta}$ is the solution of $z^{d+1} = \beta(z+1)^d$.*

For the special case of affine latency functions, we obtain a refined analysis.

Corollary 4. *Given the predicted input rates $\bar{r}$ and a parameter $\beta \in [1, 2]$, the* POTENTIAL(β,$\bar{r}$) *mechanism has PoA-Robustness $\frac{\beta z^2 + (\beta+2)z + \frac{\beta+3}{4}}{2z+1-\beta}$, for affine latency functions, where z is the maximum integer that satisfies $z \leq \frac{\beta - 1 + \sqrt{\beta^4 + 4\beta^2 - 1}}{2}$.*

For the case of affine latencies, in the extreme value of $\beta = 2$, Corollary 3 gives an (optimal) consistency bound of 1, and Corollary 4 gives a PoA-Robustness of 6.25, whereas for the other extreme value of $\beta = 1$, the consistency becomes approximately 1.577 (which coincides with the PoS of exact equilibria), while PoA-Robustness is $5/2$ (which coincides with the PoA of exact equilibria).

Remark 1. *We remark that the bounds of Corollaries 3 and 4 provide better consistency and PoA-Robustness guarantees compared to the results of Theorems 3 and 4 for $d = 1$. Applying $d = 1$ to Theorems 3 and 4 gives a higher consistency bound of 2 (instead of $\simeq 1.577$) and a higher PoA-Robustness bound of approximately 7.33 (instead of 6.25), for $\beta = 1$ and $\beta = 2$, respectively.*

The remainder of this section is organized as follows. Section C.1 introduces the definitions and lemmas used to design learning-augmented coordination mechanisms. In Section C.2, we analyze the performance of the POTENTIAL mechanism for the case of affine latency functions (Corollary 3). Section C.3 extends the analysis to the more general setting of bounded-degree polynomials and proves Theorem 3. Finally, in Section C.4, we analyze PoA-Robustness for both affine (Corollary 4) and polynomial latency functions (Theorem 4). We also include a threshold-sensitive robustness refinement in Section C.3.1.

C.1 The POTENTIAL Mechanism

In this section we introduce our main learning-augmented mechanism for atomic congestion games, namely the POTENTIAL mechanism, along with the technical lemmas used to bound its performance. Analogous to the non-atomic setting, we begin by presenting a base coordination mechanism, denoted THRESHOLD, which modifies each latency function based on a specified "threshold" congestion level $x_e(s)$ and a corresponding set of target latencies h_e. In particular, for each resource e, the mechanism increases the latency $\ell_e(x)$ to $h_e(x)$ once the congestion level exceeds the threshold. Lemmas 3 and 4 establish the key structural properties of this mechanism.

We then proceed to the description of the POTENTIAL mechanism. The mechanism takes as input the predicted demand $\bar{r}$ and a parameter $\beta \in [2(d+1)/(d+3), d+1]$, and modifies the latency functions according to the following THRESHOLD mechanism.

Definition 3 (THRESHOLD$(h, \mathbf{x})$)**.** *Given a set of latency functions $h = (h_e)_{e\in E}$ and a threshold load $\mathbf{x} = (x_e)_{e\in E}$, such that $\ell_e(x) \le h_e(x)$ for all $x \ge 1$, we define the following modified latencies:*

$$\hat{\ell}_e(x) = \begin{cases} h_e(x), & x > x_e \\ \ell_e(x), & \textit{otherwise} \end{cases}$$

An important property of the THRESHOLD mechanism is that, when the threshold profile is chosen as a minimizer of the potential Φ^h, this same profile is also a potential minimizer w.r.t. the modified latencies $\hat{\ell}$. Moreover, this particular profile preserves its original social cost, since it never triggers the increased latencies. We show this property of THRESHOLD in the following lemma.

Lemma 3. *Given a predicted demand $\bar{r}$ and a set of non-decreasing latency functions $(h_e)_{e\in E}$, with $\ell_e(x) \le h_e(x)$ for all $x \ge 1$ and $e \in E$, let $\bar{\mathbf{s}}$ be a minimizer of Φ^h for demand $\bar{r}$, and let $\hat{\ell}$ be the modified latencies produced by* THRESHOLD$(h, \bar{\mathbf{s}})$*. Then $\bar{\mathbf{s}}$ is a minimizer of $\Phi^{\hat{\ell}}$ for demand $\bar{r}$ and preserves the original social cost, i.e., $\hat{C}(\bar{\mathbf{s}}) = C(\bar{\mathbf{s}})$.*

Proof. Let $\hat{\Phi} = \Phi^{\hat{\ell}}$, Φ^h, and Φ^ℓ be the potential under the latencies $\hat{\ell}$, h, and ℓ respectively. Let $\mathbf{s}$ be any strategy profile and let E_h be the set of edges for which the increased latencies h are triggered, i.e., $x_e(\mathbf{s}) > x_e(\bar{\mathbf{s}})$, and let $E_\ell = E \setminus E_h$ be the rest of the edges, where the original latencies are used. The potential of $\mathbf{s}$ is bounded as follows

$$\begin{aligned}
\hat{\Phi}(\mathbf{s}) &\ge \hat{\Phi}(\mathbf{s}) + \Phi^h(\bar{\mathbf{s}}) - \Phi^h(\mathbf{s}) \\
&= \sum_{e\in E_h} \left(\hat{\Phi}_e(\mathbf{s}) + \Phi^h_e(\bar{\mathbf{s}}) - \Phi^h_e(\mathbf{s})\right) + \sum_{e\in E_\ell} \left(\hat{\Phi}_e(\mathbf{s}) + \Phi^h_e(\bar{\mathbf{s}}) - \Phi^h_e(\mathbf{s})\right) \\
&= \sum_{e\in E_h} \left(\left(\sum_{i=1}^{x_e(\bar{\mathbf{s}})} \ell_e(i) + \sum_{i=x_e(\bar{\mathbf{s}})+1}^{x_e(\mathbf{s})} h_e(i) \right) - \sum_{i=x_e(\bar{\mathbf{s}})+1}^{x_e(\mathbf{s})} h_e(i) \right) + \sum_{e\in E_\ell} \left(\Phi^\ell_e(\mathbf{s}) + \Phi^h_e(\bar{\mathbf{s}}) - \Phi^h_e(\mathbf{s})\right) \\
&= \sum_{e\in E_h} \Phi^\ell_e(\bar{\mathbf{s}}) + \sum_{e\in E_\ell} \left(\Phi^\ell_e(\mathbf{s}) - \Phi^h_e(\mathbf{s}) + \Phi^h_e(\bar{\mathbf{s}})\right)
\end{aligned}$$

$$
\begin{aligned}
&= \sum_{e \in E_h} \Phi_e^{\ell}(\bar{\mathbf{s}}) + \sum_{e \in E_\ell} \left(\Phi_e^{\ell}(\mathbf{s}) + \sum_{i=x_e(\mathbf{s})+1}^{x_e(\bar{\mathbf{s}})} h_e(i) \right) \\
&= \sum_{e \in E_h} \Phi_e^{\ell}(\bar{\mathbf{s}}) + \sum_{e \in E_\ell} \left(\sum_{i=1}^{x_e(\mathbf{s})} \ell_e(i) + \sum_{i=x_e(\mathbf{s})+1}^{x_e(\bar{\mathbf{s}})} h_e(i) \right) \\
&\geq \sum_{e \in E_h} \Phi_e^{\ell}(\bar{\mathbf{s}}) + \sum_{e \in E_\ell} \Phi_e^{\ell}(\bar{\mathbf{s}}) \\
&= \hat{\Phi}(\bar{\mathbf{s}}),
\end{aligned}
$$

where the first inequality is due to the fact that $\bar{\mathbf{s}}$ minimizes Φ^h, and the second due to $h_e(x) \geq \ell_e(x)$ for all $e \in E$ and $x \geq 0$. The last equality is due to the fact that the profile $\bar{\mathbf{s}}$ never triggers the increased latencies h, hence $\hat{\Phi}(\bar{\mathbf{s}}) = \Phi^{\ell}(\bar{\mathbf{s}})$.

The inequality above shows that $\hat{\Phi}(\mathbf{s}) \geq \hat{\Phi}(\bar{\mathbf{s}})$ for every strategy profile $\mathbf{s}$ feasible for demand $\bar{r}$. Therefore, $\bar{\mathbf{s}}$ is a minimizer of $\hat{\Phi} = \Phi^{\hat{\ell}}$ for demand $\bar{r}$. Moreover, since the profile $\bar{\mathbf{s}}$ never triggers the increased latencies, we also have $\hat{C}(\bar{\mathbf{s}}) = C(\bar{\mathbf{s}})$. □

The second key property of THRESHOLD is the one that enables us to derive upper bounds on PoA-Robustness using smoothness arguments that are commonly used in PoA proofs. The following lemma relates the latencies of the *modified network* to the original latencies, serving a role analogous to the Lifting Lemma for non-atomic congestion games (Lemma 2). Its proof is straightforward.

Lemma 4 (Lifting Lemma)**.** *Given a set of non-decreasing latency functions $h = (h_e)_{e \in E}$, such that $\ell_e(x) \leq h_e(x)$, for all $x \geq 0$ and $e \in E$, and a threshold load $\mathbf{x} = (x_e(s))_{e \in E}$, let $\hat{\ell}$ be the modified latencies of* THRESHOLD$(h, \mathbf{x})$*. For any demand r, let $\mathbf{s}$ be an equilibrium under $\hat{\ell}$ and s^* any other feasible strategy profile. Then,*

$$
\sum_{e \in E} \hat{\ell}_e(x_e(\mathbf{s})) x_e(\mathbf{s}) \leq \sum_{e \in E} h_e(x_e(\mathbf{s}) + 1) x_e(\mathbf{s}^*)
$$

Proof. Since $\mathbf{s}$ is a Nash equilibrium w.r.t. $\hat{\ell}$,

$$
\sum_{e \in s_i} \hat{\ell}_e(x_e(\mathbf{s})) \leq \sum_{e \in s_i^*} \hat{\ell}_e(x_e(\mathbf{s}) + 1) \leq \sum_{e \in s_i^*} h_e(x_e(\mathbf{s}) + 1).
$$

By summing up over all players the lemma follows. □

We are now ready to define the POTENTIAL mechanism that receives as input a predicted demand $\bar{r}$ and a parameter $\beta \in [2(d+1)/(d+3), d+1]$, and modifies the latencies by using a specific instantiation of THRESHOLD depending on $\bar{r}$ and β.

For simplicity, we assume, without loss of generality, that each resource is decomposed into $d+1$ resources, where the latency of each such resource is one of the monomials of the latency of the original resource. In the general case, where latencies are polynomials of degree d, all our expressions can be derived by the summation of the related expressions of the monomials. For each resource e in the equivalent setting of monomials, we denote by d_e the degree of the monomial representing the latency on e.

Definition 4 (POTENTIAL$(\beta, \bar{r})$)**.** *Given a parameter $\beta \in \left[\frac{2(d+1)}{d+3}, d+1\right]$, we define the polynomial latency function $h(\beta) = (h_e)_{e \in E}$ as follows. For $d_e \geq 1$,*

$$
\begin{aligned}
h_e(x) = &\frac{a_e \beta (d+1)}{\beta(d+1) + (d_e+1)(d+1-\beta)} (x^{d_e+1} - (x-1)^{d_e+1}) + \\
&\frac{a_e (d_e+1)(d+1-\beta)}{\beta(d+1) + (d_e+1)(d+1-\beta)} (x^{d_e} - (x-1)^{d_e}) \,.
\end{aligned}
$$

For $d_e = 0$, h_e is defined as $h_e(x) = a_e$. Let $\bar{\mathbf{s}}$ be a potential minimizer of Φ^h for the predicted demand $\bar{r}$. Then, we define the POTENTIAL$(\beta, \bar{r})$ *mechanism as the learning-augmented mechanism that uses the modified latencies of* THRESHOLD$(h(\beta), \bar{\mathbf{s}})$.

Remark 2. *We remark that h_e is chosen as being the function satisfying $\Phi^h_e(x) = \sum_{i=1}^{x} h_e(i)$, for the following $\Phi^h_e(x)$, for all $e \in E$:*

$$\Phi^h_e(x) = a_e \frac{\beta(d+1)}{\beta(d+1) + (d_e+1)(d+1-\beta)} x^{d_e+1} + a_e \frac{(d_e+1)(d+1-\beta)}{\beta(d+1) + (d_e+1)(d+1-\beta)} x^{d_e}\,, \tag{2}$$

for $d_e \geq 1$; for $d_e = 0$, it is defined as $\Phi^h_e(x) = a_e x$. This potential function is used in [CGGS19] to bound the PoS of β-approximate equilibria; the strategy profile $\mathbf{s}$ that minimizes $\Phi^h(\mathbf{s}) = \sum_{e \in E} \Phi^h_e(x_e(\mathbf{s}))$ is a β-approximate equilibrium under ℓ. The function h_e is simply $h_e(x) = \Phi^h_e(x) - \Phi^h_e(x-1)$ for $x \geq 1$, and satisfies $\ell_e(x) \leq h_e(x)$, for all $x \geq 1$, as needed [CGGS19], therefore, THRESHOLD$(h(\beta), \mathbf{x})$ *is well defined.*

We analyze the performance of POTENTIAL in the following subsections: Section C.2 addresses the case of affine latency functions, while Section C.3 extends the analysis to general polynomial latencies.

C.2 Warmup - Affine Latencies

In this section, we provide consistency-robustness guarantees for the POTENTIAL mechanism for affine latencies of the form $\ell_e(x) = a_e x + b_e$, for each resource $e \in E$, where a_e and b_e are nonnegative coefficients.

Corollary 3 (Restated)**.** *Given the predicted input rates $\bar{r}$ and a parameter $\beta \in [1, 2]$, the* POTENTIAL$(\beta,\bar{r})$ *mechanism is $\frac{\sqrt{3}+1}{\sqrt{3}+\beta-1}$-consistent, and 2β-robust for affine latency functions.*

Proof. For a given confidence parameter $\beta \in [1, 2]$, h_e is defined as follows

$$h_e(x) = \frac{2a_e\beta(2x-1)}{2\beta + 2(2-\beta)} + \frac{2a_e(2-\beta)}{2\beta + 2(2-\beta)} + b_e = \beta(a_e x + b_e) - (\beta-1)(a_e + b_e) = \beta a_e x + b_e - (\beta-1)a_e\,.$$

Let $\bar{\mathbf{s}}$ be a potential minimizer of Φ^h when the demand is $\bar{r}$. The modified latency functions $\hat{\ell}$ of POTENTIAL$(\beta, \bar{r})$ are defined by THRESHOLD$(h, \mathbf{x})$, and $\mathbf{x} = (x_e)_{e \in E}$ with $x_e = x_e(\bar{\mathbf{s}})$.

Consistency. We first show the consistency guarantee. For that, suppose that the prediction is accurate, i.e., $r = \bar{r}$. By Lemma 3, $\bar{\mathbf{s}}$ is a minimizer of $\hat{\Phi}$ and preserves its original social cost, i.e., $\hat{C}(\bar{\mathbf{s}}) = C(\bar{\mathbf{s}})$. Moreover, by definition, $\bar{\mathbf{s}}$ minimizes Φ^h, and hence also Φ^h/β. By [CKS11], any minimizer of Φ^h/β[10] has social cost w.r.t. the original latencies ℓ at most a factor of $\frac{\sqrt{3}+1}{\sqrt{3}+\beta-1}$ away from the optimal cost. Since $\hat{C}(\bar{\mathbf{s}}) = C(\bar{\mathbf{s}})$, the consistency bound follows.

Robustness. Let $\mathbf{s}$ be any strategy profile and let E_h be the set of edges for which $x_e(\mathbf{s}) > x_e(\bar{\mathbf{s}})$, i.e., $\hat{\ell}_e(x_e(\mathbf{s})) = h_e(x_e(\mathbf{s}))$, and let $E_\ell \subseteq E \setminus E_h$ be the rest of the edges that are used in $\mathbf{s}$ by at least one agent; note that for each $e \in E_\ell$, it is $\hat{\ell}_e(x_e(\mathbf{s})) = \ell_e(x_e(\mathbf{s}))$. Then $\hat{\Phi}(\mathbf{s})$ can be written as

$$\begin{aligned}
\hat{\Phi}(\mathbf{s}) &= \sum_{e \in E_\ell} \left(a_e \frac{x_e(\mathbf{s})(x_e(\mathbf{s})+1)}{2} + b_e x_e(\mathbf{s}) \right) \\
&\quad + \sum_{e \in E_h} \left(\beta a_e \frac{x_e(\mathbf{s})(x_e(\mathbf{s})+1)}{2} + (b_e - (\beta-1)a_e) x_e(\mathbf{s}) - (\beta-1)a_e \frac{x_e(\bar{\mathbf{s}})(x_e(\bar{\mathbf{s}})-1)}{2} \right) \\
&\leq \sum_{e \in E_\ell} \left(a_e \frac{x_e(\mathbf{s})(x_e(\mathbf{s})+1)}{2} + b_e x_e(\mathbf{s}) \right) \\
&\quad + \sum_{e \in E_h} \left(\beta a_e \frac{x_e(\mathbf{s})(x_e(\mathbf{s})+1)}{2} + (b_e - (\beta-1)a_e) x_e(\mathbf{s}) \right)
\end{aligned}$$

[10] In fact, Φ^h/β is used in [CKS11] to upper bound the PoS of β-approximate equilibria.

$$\leq \sum_{e\in E_\ell} x_e(\mathbf{s})\ell_e(x_e(\mathbf{s})) + \sum_{e\in E_h} x_e(\mathbf{s})\ell_e(x_e(\mathbf{s}))$$
$$\leq C(\mathbf{s})$$

where the second inequality uses $x_e(\mathbf{s}) + 1 \leq 2x_e(\mathbf{s})$ for $e \in E_\ell$, and for $e \in E_h$. Note that for any $e \in E_h$, we have $a_e(x_e(\mathbf{s}))^2 + b_e x_e(\mathbf{s}) + (\beta/2-1)a_e(x_e(\mathbf{s}))^2 - (\beta/2-1)a_e x_e(\mathbf{s}) \leq x_e(\mathbf{s})\ell_e(x_e(\mathbf{s}))$, where for the inequality we use that $(\beta/2-1)a_e(x_e(\mathbf{s}))^2 \leq (\beta/2-1)a_e x_e(\mathbf{s})$, since $x_e(\mathbf{s}) \geq 1$ and $\beta \leq 2$.

Since

$$\sum_{i=1}^{x_e(\bar{\mathbf{s}})} \ell_e(i) + \sum_{i=x_e(\bar{\mathbf{s}})+1}^{x_e(\mathbf{s})} h_e(i) \geq \frac{1}{\beta}\sum_{i=1}^{x_e(\mathbf{s})} h_e(i) \geq \frac{x_e(\mathbf{s})}{2\beta} h_e(x_e(\mathbf{s})),$$

where the first inequality follows from $\ell_e(i) \geq h_e(i)/\beta$ for the prefix $i \leq x_e(\bar{\mathbf{s}})$, and from $h_e(i) \geq h_e(i)/\beta$ for the suffix $i > x_e(\bar{\mathbf{s}})$, using $\beta \geq 1$ and $h_e(i) \geq 0$. The second inequality follows from $\sum_{i=1}^{x_e(\mathbf{s})} h_e(i) \geq \frac{x_e(\mathbf{s})}{2} h_e(x_e(\mathbf{s}))$, and since also $\sum_{i=1}^{x_e(\mathbf{s})} \ell_e(i) \geq \frac{x_e(\mathbf{s})}{2\beta}\ell_e(x_e(\mathbf{s}))$ for every $e \in E_\ell$, we get

$$\hat{\Phi}(\mathbf{s}) \geq \sum_{e\in E_\ell} \frac{x_e(\mathbf{s})}{2\beta}\ell_e(x_e(\mathbf{s})) + \sum_{e\in E_h} \frac{x_e(\mathbf{s})}{2\beta} h_e(x_e(\mathbf{s})) = \frac{1}{2\beta}\hat{C}(\mathbf{s})\,.$$

We combine the above inequalities as follows: Let $\mathbf{s}^\diamond$ be any potential minimizer of $\hat{\Phi}$, and $\mathbf{s}^*$ be an optimal profile w.r.t. the original latency functions, i.e., $\mathbf{s}^*$ minimizes $C(\mathbf{s})$. Then, by utilizing the above inequalities that hold for any $\mathbf{s}$, and therefore for $\mathbf{s}^\diamond$ and $\mathbf{s}^*$, it follows that $\frac{1}{2\beta}\hat{C}(\mathbf{s}^\diamond) \leq \hat{\Phi}(\mathbf{s}^\diamond) \leq \hat{\Phi}(\mathbf{s}^*) \leq C(\mathbf{s}^*)$, which concludes the proof for robustness, and so completes the proof of the theorem. □

C.3 Bounded-degree Polynomials

In this section, we extend our analysis to polynomial latency functions and present the proof of the main theorem. Similarly to the non-atomic case, we relate the proposed equilibrium to a suitable β-approximate equilibrium to derive the consistency bound. Robustness is established via a potential inequality argument.

Theorem 3 (Restated). *Given a predicted input demand $\bar{r}$ and a parameter $\beta \in \left[\frac{2(d+1)}{d+3}, d+1\right]$, the* POTENTIAL($\beta$,$\bar{r}$) *mechanism in atomic congestion games has*

$$\textit{Consistency: } \frac{d+1}{\beta}\,; \qquad \textit{Robustness: } \beta(d+1)\,.$$

Proof. **Consistency.** We first bound consistency. We obtain that the consistency is equal to the PoS of the β-approximate equilibria. To that end, suppose that $r = \bar{r}$, i.e., the prediction is correct. By Definition 4, the POTENTIAL($\beta, \bar{r}$) mechanism modifies the latencies according to THRESHOLD($h, \bar{\mathbf{s}}$), where

$$h_e(x) = \frac{a_e\beta(d+1)}{\beta(d+1)+(d_e+1)(d+1-\beta)}(x^{d_e+1} - (x-1)^{d_e+1}) + \frac{a_e(d_e+1)(d+1-\beta)}{\beta(d+1)+(d_e+1)(d+1-\beta)}(x^{d_e} - (x-1)^{d_e})$$

$\forall e \in E$ with $d_e \geq 1$, and $h_e(x) = a_e$, $\forall e \in E$, with $d_e = 0$, and $\bar{\mathbf{s}}$ is a minimizer of $\Phi^h = \sum_{e\in E}\sum_{i=1}^{x_e(\mathbf{s})} h_e(i)$; recall that $\Phi^h(\mathbf{s}) = \sum_{e\in E}\Phi_e^h(x_e(\mathbf{s}))$,, where the functions Φ_e^h are as defined in (2). By Lemma 3, $\bar{\mathbf{s}}$ is a minimizer of $\hat{\Phi} = \Phi^{\hat{\ell}}$ and preserves its original social cost, i.e., $\hat{C}(\bar{\mathbf{s}}) = C(\bar{\mathbf{s}})$. Moreover, by definition, $\bar{\mathbf{s}}$ minimizes Φ^h, and hence also Φ^h/β. Then, due to [CGGS19] any minimizer of Φ^h/β[11] has social cost w.r.t. the original latencies ℓ at most a factor of $\frac{d+1}{\beta}$ away

[11] Any minimizer of Φ^h/β corresponds to a β-approximate equilibrium.

from the optimal cost, for $\beta \in \left[\frac{2(d+1)}{d+3}, d+1\right]$. Since $\hat{C}(\bar{\mathbf{s}}) = C(\bar{\mathbf{s}})$, we obtain $\hat{C}(\bar{\mathbf{s}}) \leq \frac{d+1}{\beta} C_{\text{opt}}(\bar{r})$. Therefore, for every $\beta \in \left[\frac{2(d+1)}{d+3}, d+1\right]$, the mechanism's consistency is at most $\frac{d+1}{\beta}$.

Robustness. Let $\mathbf{s}$ be any strategy profile and let E_h be the set of edges for which $x_e(\mathbf{s}) > x_e(\bar{\mathbf{s}})$, i.e. $\hat{\ell}_e(x_e(\mathbf{s})) = h_e(x_e(\mathbf{s}))$ and let $E_\ell = E \setminus E_h$ be the rest of the edges, where the original latencies are used by at least one agent. Then $\hat{\Phi}(\mathbf{s})$ can be written as:

$$\hat{\Phi}(\mathbf{s}) = \sum_{e \in E} \sum_{i=1}^{x_e(\mathbf{s})} \hat{\ell}_e(i) = \sum_{e \in E_\ell} \sum_{i=1}^{x_e(\mathbf{s})} \ell_e(i) + \sum_{e \in E_h} \left(\sum_{i=1}^{x_e(\bar{\mathbf{s}})} \ell_e(i) + \sum_{i=x_e(\bar{\mathbf{s}})+1}^{x_e(\mathbf{s})} h_e(i) \right).$$

We will use the following properties of the auxiliary latencies h: for every edge e and every $x \geq 1$, $\ell_e(x) \leq h_e(x) \leq \beta \ell_e(x)$.[12] We first show that $\hat{\Phi}(\mathbf{s}) \leq C(\mathbf{s})$ and then that $\hat{\Phi}(\mathbf{s}) \geq \frac{1}{\beta(d+1)} \hat{C}(\mathbf{s})$. We will use Lemma 6, whose statement and proof are deferred to Appendix E.

Claim 5. $\hat{\Phi}(\mathbf{s}) \leq C(\mathbf{s})$.

Proof. For any $e \in E_h$, with $d_e \geq 1$, by using Lemma 6 (1), it holds that

$$\begin{aligned}
&\sum_{i=1}^{x_e(\bar{\mathbf{s}})} \ell_e(i) + \sum_{i=x_e(\bar{\mathbf{s}})+1}^{x_e(\mathbf{s})} h_e(i) \leq \sum_{i=1}^{x_e(\mathbf{s})} h_e(i) \\
&= \frac{a_e \beta(d+1)}{\beta(d+1) + (d_e+1)(d+1-\beta)} \sum_{i=1}^{x_e(\mathbf{s})} (i^{d_e+1} - (i-1)^{d_e+1}) + \\
&\quad \frac{a_e (d_e+1)(d+1-\beta)}{\beta(d+1) + (d_e+1)(d+1-\beta)} \sum_{i=1}^{x_e(\mathbf{s})} (i^{d_e} - (i-1)^{d_e}) \\
&= \frac{a_e \beta(d+1)}{\beta(d+1) + (d_e+1)(d+1-\beta)} (x_e(\mathbf{s}))^{d_e+1} + \frac{a_e (d_e+1)(d+1-\beta)}{\beta(d+1) + (d_e+1)(d+1-\beta)} (x_e(\mathbf{s}))^{d_e} \\
&\leq \frac{a_e \beta(d+1)}{\beta(d+1) + (d_e+1)(d+1-\beta)} (x_e(\mathbf{s}))^{d_e+1} + \frac{a_e (d_e+1)(d+1-\beta)}{\beta(d+1) + (d_e+1)(d+1-\beta)} (x_e(\mathbf{s}))^{d_e+1} \\
&= a_e x_e(\mathbf{s})^{d_e+1} = x_e(\mathbf{s}) \ell_e(x_e(\mathbf{s}))
\end{aligned}$$

The same also holds with equality when $d_e = 0$. Then, by using the above inequality for any $e \in E_h$ and since any ℓ_e is a non-decreasing function

$$\begin{aligned}
\hat{\Phi}(\mathbf{s}) &\leq \sum_{e \in E_\ell} \sum_{i=1}^{x_e(\mathbf{s})} \ell_e(x_e(\mathbf{s})) + \sum_{e \in E_h} x_e(\mathbf{s}) \ell_e(x_e(\mathbf{s})) \\
&= \sum_{e \in E_\ell} x_e(\mathbf{s}) \ell_e(x_e(\mathbf{s})) + \sum_{e \in E_h} x_e(\mathbf{s}) \ell_e(x_e(\mathbf{s})) = C(\mathbf{s}).
\end{aligned}$$

□

Claim 6. $\hat{\Phi}(\mathbf{s}) \geq \frac{1}{\beta(d+1)} \hat{C}(\mathbf{s})$.

Proof. By using Lemma 6 (2) for any $e \in E_h$, with $d_e \geq 1$:

$$\sum_{i=1}^{x_e(\bar{\mathbf{s}})} \ell_e(i) + \sum_{i=x_e(\bar{\mathbf{s}})+1}^{x_e(\mathbf{s})} h_e(i) \geq \frac{1}{\beta} \sum_{i=1}^{x_e(\mathbf{s})} h_e(i)$$

[12] Regarding the first inequality, for $d_e = 0$, we have $h_e(x) = \ell_e(x) = a_e$, so this is immediate. For $d_e \geq 1$, it follows by expanding $h_e(x) - \ell_e(x)$ and using $x^{d_e} - (x-1)^{d_e} \geq d_e (x-1)^{d_e - 1}$; the remaining condition is $\beta(d+1)d_e \geq (d_e+1)(d+1-\beta)$, equivalently $\beta \geq \frac{(d_e+1)(d+1)}{d_e(d+1)+d_e+1}$,which holds for all $\beta \geq \frac{2(d+1)}{d+3}$ for every $1 \leq d_e \leq d$.

$$
\begin{aligned}
&= \frac{1}{\beta}\left(\frac{a_e\beta(d+1)}{\beta(d+1)+(d_e+1)(d+1-\beta)}(x_e(\mathbf{s}))^{d_e+1}\right.\\
&\quad\left.+\frac{a_e(d_e+1)(d+1-\beta)}{\beta(d+1)+(d_e+1)(d+1-\beta)}(x_e(\mathbf{s}))^{d_e}\right)\\
&\geq \frac{1}{\beta}\left(\frac{a_e\beta(d+1)}{\beta(d+1)+(d_e+1)(d+1-\beta)}x_e(\mathbf{s})\frac{(x_e(\mathbf{s}))^{d_e+1}-(x_e(\mathbf{s})-1)^{d_e+1}}{d_e+1}\right.\\
&\quad\left.+\frac{a_e(d_e+1)(d+1-\beta)}{\beta(d+1)+(d_e+1)(d+1-\beta)}x_e(\mathbf{s})\frac{(x_e(\mathbf{s}))^{d_e}-(x_e(\mathbf{s})-1)^{d_e}}{d_e}\right)\\
&\geq \frac{x_e(\mathbf{s})h_e(x_e(\mathbf{s}))}{\beta(d+1)}.
\end{aligned}
$$

The first inequality comes from the fact that for the prefix $i \leq x_e(\bar{\mathbf{s}})$, we use $\ell_e(i) \geq h_e(i)/\beta$. For the suffix $i > x_e(\bar{\mathbf{s}})$, we use $h_e(i) \geq h_e(i)/\beta$, which follows from $\beta \geq 1$ and $h_e(i) \geq 0$.

The same inequality also holds when $d_e = 0$. Moreover, using Lemma 6 (3) for any $e \in E_\ell$:

$$
\sum_{i=1}^{x_e(\mathbf{s})}\ell_e(i) = \sum_{i=1}^{x_e(\mathbf{s})}a_e i^{d_e} \geq \frac{a_e(x_e(\mathbf{s}))^{d_e+1}}{d_e+1} \geq \frac{x_e(\mathbf{s})\ell_e(x_e(\mathbf{s}))}{d+1}\,.
$$

Overall,

$$
\begin{aligned}
\hat{\Phi}(\mathbf{s}) &= \sum_{e\in E_\ell}\sum_{i=1}^{x_e(\mathbf{s})}\ell_e(i) + \sum_{e\in E_h}\left(\sum_{i=1}^{x_e(\bar{\mathbf{s}})}\ell_e(i) + \sum_{i=x_e(\bar{\mathbf{s}})+1}^{x_e(\mathbf{s})}h_e(i)\right)\\
&\geq \sum_{e\in E_\ell}\frac{x_e(\mathbf{s})\ell_e(x_e(\mathbf{s}))}{\beta(d+1)} + \sum_{e\in E_h}\frac{x_e(\mathbf{s})h_e(x_e(\mathbf{s}))}{\beta(d+1)} = \frac{1}{\beta(d+1)}\hat{C}(\mathbf{s})\,.
\end{aligned}
$$

□

We complete the proof for the robustness upper bound by using the above two claims. Let $\mathbf{s}^\diamond$ be any potential minimizer of $\hat{\Phi}$, and $\mathbf{s}^*$ be an optimal profile w.r.t. the original latency functions, i.e., $\mathbf{s}^*$ minimizes $C(\mathbf{s})$. Then, by utilizing the above two claims that hold for any $\mathbf{s}$, and therefore for $\mathbf{s}^\diamond$ and $\mathbf{s}^*$, it follows that $\frac{1}{\beta(d+1)}\hat{C}(\mathbf{s}^\diamond) \leq \hat{\Phi}(\mathbf{s}^\diamond) \leq \hat{\Phi}(\mathbf{s}^*) \leq C(\mathbf{s}^*)$, which concludes the proof for robustness, and so completes the proof of the theorem. □

C.3.1 A threshold-sensitive robustness refinement

Similarly to the non-atomic case, we present a refined threshold-sensitive analysis of the POTENTIAL mechanism, depending on how close the triggered resources are to their predicted thresholds. Let $\mathbf{s}$ be any strategy profile and let $E_h(\mathbf{s}) = \{e \in E : x_e(\mathbf{s}) > x_e(\bar{\mathbf{s}})\}$ be the set of triggered resources. For every $e \in E_h(\mathbf{s})$, define $\eta_e = x_e(\bar{\mathbf{s}})/x_e(\mathbf{s}) \in [0,1)$. If $E_h(\mathbf{s}) \neq \emptyset$, let $\eta(\mathbf{s},\bar{\mathbf{s}}) = \max_{e\in E_h(\mathbf{s})}\eta_e$, and otherwise set $\eta(\mathbf{s},\bar{\mathbf{s}}) = 0$. Recall that $\hat{\Phi}_e(\mathbf{s}) = \sum_{i=1}^{x_e(\mathbf{s})}\hat{\ell}_e(i)$, $\hat{\Phi}(\mathbf{s}) = \sum_{e\in E}\hat{\Phi}_e(\mathbf{s})$, and $\hat{C}_e(\mathbf{s}) = x_e(\mathbf{s})\hat{\ell}_e(x_e(\mathbf{s}))$.

We will use two elementary facts. First, $h_e(x) \leq \beta\ell_e(x)$ for every resource e and every integer $x \geq 1$, as established in the proof of Theorem 3. Second, if $\Phi_e^h(x) = \sum_{i=1}^{x}h_e(i)$, then $\Phi_e^h(x) \geq xh_e(x)/(d+1)$ for every integer $x \geq 1$. This follows from the definition of the auxiliary potential Φ_e^h and Lemma 6.

Lemma C.1 (Threshold-sensitive potential lower bound)**.** *For every strategy profile* $\mathbf{s}$*,*

$$
\hat{\Phi}(\mathbf{s}) \geq \frac{1-\left(1-\frac{1}{\beta}\right)\eta(\mathbf{s},\bar{\mathbf{s}})}{d+1}\,\hat{C}(\mathbf{s}).
$$

Proof. We prove the inequality resource by resource. First consider an untriggered resource $e \notin E_h(\mathbf{s})$, so $x_e(\mathbf{s}) \leq x_e(\bar{\mathbf{s}})$. Then $\hat{\ell}_e(i) = \ell_e(i)$ for every $1 \leq i \leq x_e(\mathbf{s})$, and hence $\hat{\Phi}_e(\mathbf{s}) =$

$\sum_{i=1}^{x_e(\mathbf{s})} \ell_e(i) \geq x_e(\mathbf{s})\ell_e(x_e(\mathbf{s}))/(d+1) = \hat{C}_e(\mathbf{s})/(d+1)$, where the inequality follows from Lemma 6 and the nonnegativity of the coefficients of ℓ_e. Since $1-(1-1/\beta)\eta(\mathbf{s},\bar{\mathbf{s}}) \leq 1$, this already implies the desired resource-wise bound for every untriggered resource.

Now consider a triggered resource $e \in E_h(\mathbf{s})$, so $x_e(\mathbf{s}) > x_e(\bar{\mathbf{s}})$. Then

$$\begin{aligned}
\hat{\Phi}_e(\mathbf{s}) &= \sum_{i=1}^{x_e(\bar{\mathbf{s}})} \ell_e(i) + \sum_{i=x_e(\bar{\mathbf{s}})+1}^{x_e(\mathbf{s})} h_e(i) \\
&= \Phi_e^{\ell}(x_e(\bar{\mathbf{s}})) + \Phi_e^{h}(x_e(\mathbf{s})) - \Phi_e^{h}(x_e(\bar{\mathbf{s}})) \\
&\geq \Phi_e^{h}(x_e(\mathbf{s})) - \left(1-\frac{1}{\beta}\right)\Phi_e^{h}(x_e(\bar{\mathbf{s}})) \\
&\geq \Phi_e^{h}(x_e(\mathbf{s})) - \left(1-\frac{1}{\beta}\right)\eta_e \Phi_e^{h}(x_e(\mathbf{s})) \\
&= \left(1-\left(1-\frac{1}{\beta}\right)\eta_e\right)\Phi_e^{h}(x_e(\mathbf{s})) \\
&\geq \frac{1-\left(1-\frac{1}{\beta}\right)\eta_e}{d+1} x_e(\mathbf{s}) h_e(x_e(\mathbf{s})) \\
&= \frac{1-\left(1-\frac{1}{\beta}\right)\eta_e}{d+1} \hat{C}_e(\mathbf{s}).
\end{aligned}$$

The first inequality uses $h_e(i) \leq \beta \ell_e(i)$, which implies $\Phi_e^{\ell}(x_e(\bar{\mathbf{s}})) \geq \Phi_e^{h}(x_e(\bar{\mathbf{s}}))/\beta$. The second inequality uses $\Phi_e^{h}(x_e(\bar{\mathbf{s}})) \leq \eta_e \Phi_e^{h}(x_e(\mathbf{s}))$: indeed, Φ_e^{h} has nonnegative coefficients and zero constant term as a polynomial in the load, and $(x_e(\bar{\mathbf{s}})/x_e(\mathbf{s}))^q \leq x_e(\bar{\mathbf{s}})/x_e(\mathbf{s}) = \eta_e$ for every $q \geq 1$. The last inequality uses $\Phi_e^{h}(x_e(\mathbf{s})) \geq x_e(\mathbf{s}) h_e(x_e(\mathbf{s}))/(d+1)$. Since $\eta_e \leq \eta(\mathbf{s},\bar{\mathbf{s}})$ and $1-(1-1/\beta)\eta$ is decreasing in η, this gives

$$\hat{\Phi}_e(\mathbf{s}) \geq \frac{1-\left(1-\frac{1}{\beta}\right)\eta(\mathbf{s},\bar{\mathbf{s}})}{d+1} \hat{C}_e(\mathbf{s})$$

for every triggered resource. Summing over all resources completes the proof. □

Corollary C.1 (Threshold-sensitive robustness)**.** *Let $\mathbf{s}^{\diamond}$ be any potential minimizer of the modified game, and let $\mathbf{s}^*$ be an optimal strategy profile under the original latencies. Then*

$$\hat{C}(\mathbf{s}^{\diamond}) \leq \frac{d+1}{1-\left(1-\frac{1}{\beta}\right)\eta(\mathbf{s}^{\diamond},\bar{\mathbf{s}})} C(\mathbf{s}^*).$$

Proof. By Lemma C.1,

$$\frac{1-\left(1-\frac{1}{\beta}\right)\eta(\mathbf{s}^{\diamond},\bar{\mathbf{s}})}{d+1} \hat{C}(\mathbf{s}^{\diamond}) \leq \hat{\Phi}(\mathbf{s}^{\diamond}).$$

Since $\mathbf{s}^{\diamond}$ minimizes $\hat{\Phi}$, and since the robustness proof of Theorem 3 shows that $\hat{\Phi}(\mathbf{s}) \leq C(\mathbf{s})$ for every strategy profile $\mathbf{s}$, we have $\hat{\Phi}(\mathbf{s}^{\diamond}) \leq \hat{\Phi}(\mathbf{s}^*) \leq C(\mathbf{s}^*)$. Combining the two inequalities and rearranging gives the claim. □

As in the non-atomic case, this bound refines the worst-case robustness guarantee of Theorem 3 by tracking how close the triggered resources are to their predicted threshold loads. When the realized potential minimizer exceeds the threshold loads by a large margin, the factor is close to $d+1$; in the worst case, as $\eta(s^{\diamond},\bar{s})$ approaches 1, the bound recovers the general guarantee $\beta(d+1)$.

C.4 PoA-Robustness

In this section, we analyze the PoA-Robustness of POTENTIAL for the case of atomic games with polynomial latencies with nonnegative coefficients of degree d. The main result shows an upper bound of PoA-Robustness as a function of the confidence parameter β (Theorem 4).

The next result presents an upper bound of PoA-Robustness for the special case of affine latency functions. We stress that this is an improved bound compared to the bound obtained by Theorem 4 for $d = 1$.

Corollary 4 (Restated)**.** *Given the predicted input rates $\bar{r}$ and a parameter $\beta \in [1, 2]$, the* POTENTIAL(β,$\bar{r}$) *mechanism has PoA-Robustness $\frac{\beta z^2+(\beta+2)z+\frac{\beta+3}{4}}{2z+1-\beta}$, for affine latency functions, where z is the maximum integer that satisfies $z \leq \frac{\beta-1+\sqrt{\beta^4+4\beta^2-1}}{2}$.*

Proof. Given an arbitrary input rate r, let $\mathbf{s}$ be any equilibrium profile under the modified latencies $\hat{\ell}$, and $\mathbf{s}^*$ be the optimal profile under the original latencies ℓ.

By using the Lifting Lemma for atomic games (Lemma 4) and substituting for $h_e(x) = \beta(a_e x + b_e) - (\beta - 1)(a_e + b_e)$, we obtain

$$\hat{C}(\mathbf{s}) = \sum_{e \in E} \hat{\ell}_e(x_e(\mathbf{s}))x_e(\mathbf{s}) \leq \sum_{e \in E} (\beta a_e x_e(\mathbf{s}) + a_e + b_e)x_e(\mathbf{s}^*). \tag{3}$$

A quick —but not tight— upper bound of PoA-Robustness can be obtained as follows. The last term of (3) can be bounded above by $\beta \cdot \sum_{e \in E} \ell_e(x_e(\mathbf{s}) + 1)x_e(\mathbf{s}^*)$, which gives

$$\hat{C}(\mathbf{s}) \leq \beta \sum_{e \in E} \ell_e(x_e(\mathbf{s}) + 1)x_e(\mathbf{s}^*).$$

By observing that the above inequality is also satisfied by β-approximate equilibria w.r.t. the original latencies, one can follow a similar analysis with that of [CKS11] to upper bound the PoA of β-approximate equilibria. This gives a PoA-Robustness bound of at most $\beta\frac{z^2+3z+1}{2z-\beta+1}$, where $z = \left\lfloor \frac{\beta+\sqrt{5+6(\beta-1)+(\beta-1)^2}}{2} \right\rfloor$.

We are able to obtain an improved bound by a more careful analysis (see Figure 1 (b)). In particular, by applying Lemma 7 (we defer the proof to Appendix E) to the last term of (3) we establish the following claim.

Claim 7. *For any parameter $\beta \in [1, 2]$ it holds that*

$$\hat{C}(\mathbf{s}) \leq \frac{\beta}{2z+1}\hat{C}(\mathbf{s}) + \beta\frac{z^2 + (1 + \frac{2}{\beta})z + \frac{\beta+3}{4\beta}}{2z+1}C(\mathbf{s}^*)\,.$$

Proof.

$$\begin{aligned}
\hat{C}(\mathbf{s}) &\leq \sum_{e \in E} (\beta a_e x_e(\mathbf{s}) + a_e + b_e)x_e(\mathbf{s}^*) \\
&\leq \sum_{e \in E} \beta \left(\frac{1}{2z+1}\right) a_e(x_e(\mathbf{s}))^2 + \beta\frac{z^2 + \left(1 + \frac{2}{\beta}\right)z + \frac{\beta+3}{4\beta}}{2z+1} a_e(x_e(\mathbf{s}^*))^2 + b_e x_e(\mathbf{s}^*) \\
&\leq \frac{\beta}{2z+1}\hat{C}(\mathbf{s}) + \frac{\beta z^2 + (\beta+2)z + \frac{\beta+3}{4}}{2z+1}C(\mathbf{s}^*)\,,
\end{aligned}$$

where the last inequality follows from the fact that $\beta z^2 + (\beta+2)z + \frac{\beta+3}{4} \geq 2z+1$ for all non-negative integers z and $\beta \geq 1$. □

Rearranging yields the following bound:

$$\frac{\hat{C}(\mathbf{s})}{C(\mathbf{s}^*)} \leq \frac{\beta z^2 + (\beta+2)z + \frac{\beta+3}{4}}{2z+1-\beta}\,.$$

Taking z to be the integer specified in the statement gives the claimed upper bound.

□

Now we are ready to restate the main result of this section.

Theorem 4 (Restated)**.** *Given a predicted demand $\bar{r}$ and a parameter $\beta \in [2(d+1)/(d+3), d+1]$, the PoA-Robustness of* POTENTIAL(β,$\bar{r}$) *for atomic congestion games is at most $\phi_{d,\beta}^{d+1}$, where $\phi_{d,\beta}$ is the solution of $z^{d+1} = \beta(z+1)^d$.*

Proof. In fact we will show that

$$\text{PoA-Robustness: } \frac{\beta((z+1)^{2d+1} - z^{d+1}(z+2)^d)}{(z+1)^{d+1} - z^{d+1} - \beta((z+2)^d - (z+1)^d)},$$

where z is the maximum integer that satisfies $\frac{z^{d+1}}{(z+1)^d} \leq \beta$. Then, similarly to [ADG+11] it holds that the above expression is upper bounded by $z^{d+1} = \phi_{d,\beta}^{d+1}$, for z that makes $\frac{z^{d+1}}{(z+1)^d} = \beta$.

Given any input rate r, let $\mathbf{s}$ be an equilibrium profile under the modified latencies $\hat{\ell}$ and $\mathbf{s}^*$ the optimal profile under the original latencies ℓ. Using the Lifting Lemma for atomic games (Lemma 4), we obtain that:

$$\hat{C}(\mathbf{s}) = \sum_{e \in E} \hat{\ell}_e(x_e(s))x_e(s) \leq \sum_{e \in E} h_e(x_e(s)+1)x_e(s^*) \tag{4}$$

Now, the right part can be upper bounded by $\beta \sum_{e \in E} \ell_e(x_e(s)+1)x_e(s^*)$. The resulting inequality enables us to apply the analysis of the atomic PoA for β-approximate equilibria in [CKS11], by observing that $\mathbf{s}$ corresponds to a β-approximate equilibrium under the original latencies ℓ. Thus, we can upper bound the robustness via the PoA of β-approximate equilibria provided by [CKS11]. Hence, we obtain a robustness bound of at most $\frac{\beta((z+1)^{2d+1} - z^{d+1}(z+2)^d)}{(z+1)^{d+1} - z^{d+1} - \beta((z+2)^d - (z+1)^d)}$, where z is the maximum integer that satisfies $\frac{z^{d+1}}{(z+1)^d} \leq \beta$. □

D Impossibility Results

In this section, we present impossibility results concerning the design of learning-augmented coordination mechanisms that satisfy further desirable properties. In Section D.1 we show that when a stronger notion of consistency (PoA-Consistency) is required, any mechanism that is 1-consistent with respect to this stronger criterion must suffer unbounded robustness. Furthermore, in Section D.2 we establish impossibility results for mechanisms whose approximation preserves a continuity property for arbitrarily small prediction errors.

D.1 PoA-Consistency

The following two theorems demonstrate that designing mechanisms that are 1-consistent with respect to the PoA is excessively demanding, thereby, motivating our choice of the potential minimizers for the consistency and robustness guarantees. In particular, we show that mechanisms that are 1-consistent w.r.t. PoA exhibit unbounded robustness (Theorem 5 and Theorem 6 for non-atomic and atomic congestion games, respectively). Furthermore, for the atomic games we generalize our impossibility result for mechanisms with PoA-Consistency equal to $\alpha \geq 1$ (see Appendix D.1.1).

Theorem 5. *There is no learning-augmented coordination mechanism for non-atomic congestion games that has simultaneously PoA-Consistency equal to* 1 *and finite robustness.*

Proof. Consider a non-atomic congestion game defined as follows. There are 4 commodities, denoted by $0, 1, 2, 3$, and a set of resources $E = \{e_0, e_1, e_2, e_3\}$. The predicted demand is $\bar{r}_0 = 2$, $\bar{r}_1 = \bar{r}_2 = 1$, $\bar{r}_3 = 0$. The latencies of the resources are: $\ell_0(x) = 2$, $\ell_1(x) = \ell_2(x) = x$, and $\ell_3(x) = \epsilon$. The set of strategies is $\Sigma_0 = \{\{e_0\}, \{e_1\}, \{e_2\}\}$ for commodity 0, $\Sigma_1 = \{\{e_1\}, \{e_2\}\}$ for commodity 1, $\Sigma_2 = \{\{e_1, e_3\}, \{e_2\}\}$ for commodity 2, and $\Sigma_3 = \{e_3\}$.

The optimal solution assigns $\bar{r}_0$ to e_0, $\bar{r}_1$ to e_1, and $\bar{r}_2$ to e_2 for the predicted demand $\bar{r}$ resulting in a social cost $C_{opt}(\bar{r}) = 2\ell_0(2) + \ell_1(1) + \ell_2(1) = 6$. There exists an equilibrium profile f in which

commodity 0 routes half of its demand through e_1 and the rest through e_2, commodity 1 routes all its demand through e_1, while commodity 2 routes all its demand through $\{e_2\}$, yielding a social cost $C(f) = 8$.

Now consider any mechanism that is required to be 1-consistent w.r.t. PoA. So, when $r = \bar{r}$, the mechanism must enforce commodity 0 to route all its demand through e_0, and for this it must modify the latencies of e_1 and e_2, so that $\hat{\ell}_1(x) > 2$ for $x > 1$ and $\hat{\ell}_2(x) > 2$ for $x > 1$. Then,

the profile where commodity 0 routes all its demand through e_0, commodity 1 routes all its demand through e_2 and commodity 2 routes all its demand through $\{e_1, e_3\}$ must not be an equilibrium for $r = \bar{r}$. Eliminating such an equilibrium without harming the predicted optimal solution requires $\hat{\ell}_3(x) > 1$ for $x > 0$.

Now consider an actual instance in which only commodity 3's demand is realized, i.e., $r_0 = r_1 = r_2 = 0$ and $r_3 = 1$. In this case, the social cost of the unique outcome is equal to the modified latency of e_3 for the realized demand. Hence, the mechanism's robustness is at least $1/\epsilon$, which tends to ∞ as $\epsilon \to 0$. □

Theorem 6. *There is no learning-augmented coordination mechanism for atomic congestion games that has simultaneously PoA-Consistency equal to* 1 *and finite robustness.*

Proof. Consider an atomic congestion game defined as follows. The set of players is $N = \{1, 2, 3\}$ and the set of resources is given by $E = \{e_1, e_2, e_3\}$. The strategy sets are $\Sigma_1 = \{\{e_1\}, \{e_2\}\}$, $\Sigma_2 = \{\{e_1, e_3\}, \{e_2\}\}$, and $\Sigma_3 = \{e_3\}$, for players $1, 2$, and 3 respectively. The latency functions are defined as $\ell_1(x) = \ell_2(x) = x$ and $\ell_3(x) = \epsilon$, for $\epsilon > 0$ arbitrarily small.

Let the predicted demand vector be $\bar{r} = \{1, 2\}$, i.e., indicating that only players 1 and 2 will be active. For this predicted instance, the socially optimal allocation assigns agent 1 to e_1 and agent 2 to e_2, resulting in a social cost of $C_{opt}(\bar{r}) = \ell_1(1) + \ell_2(1) = 2$. However, for the same predicted active set $\{1, 2\}$, the strategy profile $\mathbf{s}$ in which player 1 chooses $\{e_2\}$ and player 2 chooses $\{e_1, e_3\}$ is a Nash equilibrium with a resulting social cost of $C(s) = \ell_1(1) + \ell_2(1) + \ell_3(1) = 2 + \epsilon$, so the PoA for the predicted instance is equal to $1 + \epsilon/2 > 1$.

Now consider any mechanism that is required to be 1-consistent w.r.t. the PoA: the mechanism must modify the latencies such that any equilibrium has optimal social cost. Concretely, the mechanism must eliminate the bad equilibrium $\mathbf{s}$. Observe that the resources e_1 and e_2 are also used in the optimal solution for the predicted instance, so increasing their latencies for load 1 would also yield an increased social cost of the optimal allocation. Therefore, the only way to eliminate the bad equilibrium without harming the predicted optimum (i.e., maintaining 1-consistency) is to enforce player 2 to deviate to $\{e_2\}$ by satisfying $\hat{\ell}_3(1) \geq c$, for some constant $c > 1$.

Now consider an actual realization, that differs from the predicted demand, in which only player 3 is active. In this case, the social cost of the unique outcome is the latency of e_3, under the mechanism's modified latencies that is $\hat{\ell}_3(1) > c$. Hence, robustness is equal to c/ϵ, which is arbitrarily large after choosing ϵ arbitrarily small. □

These impossibilities show that the PoA is a poorly suited consistency target for mechanisms that are required to preserve optimality when the prediction is correct.

D.1.1 PoA-Consistency lower bounds

We further show lower bounds relaxing the 1-PoA-consistency requirement. We show that any α-consistent mechanism must suffer the following PoA-Robustness.

Theorem 7. *Any α-consistent mechanism (with respect to PoA) must have robustness*

$$\rho \geq \begin{cases} (3 - 2\alpha)/(2\alpha - 2), & 1 \leq \alpha \leq (1 + \sqrt{2})/2 \\ 2/(2\alpha - 1), & (1 + \sqrt{2})/2 \leq \alpha < (1 + \sqrt{17})/4 \end{cases}$$

Proof. Consider the same construction as in Theorem 6, but let the ground set of players be $N = \{1, 2, 3, 4\}$, the resources be $M = \{e_1, e_2, e_3\}$, and the strategy sets be as before for players $1, 2, 3$,

with player 4 having the singleton strategy set $\Sigma_4 = \{\{e_1\}\}$. The predicted active set is $\bar{N} = \{1, 2\}$. The latency functions are $\ell_1(x) = \ell_2(x) = x$ and $\ell_3(x) = 2\alpha - 2 + \epsilon$, where $\alpha \geq 1$ denotes the consistency guarantee and $\epsilon > 0$ is arbitrarily small and chosen such that $2\alpha - 2 + \epsilon < 1$.

For the predicted instance, the socially optimal allocation assigns player 1 to e_1 and player 2 to e_2, and has social cost 2. As in Theorem 6, the profile s in which player 1 selects e_2 and player 2 selects $\{e_1, e_3\}$ is an equilibrium under the original latencies. Its social cost is $2 + (2\alpha - 2 + \epsilon) = 2\alpha + \epsilon$, and therefore its approximation ratio is $\alpha + \epsilon/2 > \alpha$. Thus, any α-consistent mechanism with respect to PoA must prevent s from being an equilibrium under the modified latencies.

In the profile s, player 1's cost is $\hat{\ell}_2(1)$, while deviating to e_1 would create load 2 on e_1, and hence would cost at least 2, since the mechanism only increases latencies. Similarly, player 2's cost is $\hat{\ell}_1(1) + \hat{\ell}_3(1)$, while deviating to e_2 would create load 2 on e_2, and hence would cost at least 2. Therefore, in order to eliminate s, the mechanism must satisfy

$$\hat{\ell}_2(1) \geq 2 \qquad \text{or} \qquad \hat{\ell}_1(1) + \hat{\ell}_3(1) \geq 2.$$

We first argue that the case $\hat{\ell}_2(1) \geq 2$ is incompatible with α-PoA-consistency in the range of the theorem. Indeed, since $\alpha < (1 + \sqrt{17})/4 < 3/2$, we have $2\alpha < 3$. If $\hat{\ell}_2(1) \geq 2$, then the predicted optimal profile has modified social cost at least $\hat{\ell}_1(1) + \hat{\ell}_2(1) \geq 1 + 2 = 3 > 2\alpha$, because $\hat{\ell}_1(1) \geq \ell_1(1) = 1$. The non-optimal profile s has modified social cost at least $\hat{\ell}_1(1) + \hat{\ell}_2(1) + \hat{\ell}_3(1) \geq 1 + 2 + (2\alpha - 2 + \epsilon) > 2\alpha$. The remaining two profiles place both active players on the same load-dependent resource, and therefore have modified social cost at least $4 > 2\alpha$. Since every finite congestion game has a pure equilibrium, some profile must be an equilibrium, and all possible profiles have modified social cost greater than 2α, contradicting α-PoA-consistency. Hence, for an α-consistent mechanism, we may assume $\hat{\ell}_1(1) + \hat{\ell}_3(1) \geq 2$. Under this condition, the non-optimal profile s has modified social cost at least $\hat{\ell}_1(1) + \hat{\ell}_2(1) + \hat{\ell}_3(1) \geq 1 + 2 = 3 > 2\alpha$. The two profiles in which both active players create load 2 on either e_1 or e_2 also have modified social cost at least $4 > 2\alpha$. Thus, in any α-consistent modified game, the only possible equilibrium for the predicted active set is the predicted optimal profile. Since this equilibrium must have modified social cost at most 2α, we get $\hat{\ell}_1(1) + \hat{\ell}_2(1) \leq 2\alpha$. Using $\hat{\ell}_2(1) \geq 1$, this implies $\hat{\ell}_1(1) \leq 2\alpha - 1$. Combining this with $\hat{\ell}_1(1) + \hat{\ell}_3(1) \geq 2$, we obtain $\hat{\ell}_3(1) \geq 2 - \hat{\ell}_1(1)$.

We now derive the robustness lower bound. Consider an actual realization in which only player 4 is active. Since player 4 is forced to use e_1, the unique equilibrium has modified social cost $\hat{\ell}_1(1)$, while the optimal social cost under the original latencies is 1. Hence, $\rho \geq \hat{\ell}_1(1)$. Next consider an actual realization in which only player 3 is active. Since player 3 is forced to use e_3, the unique equilibrium has modified social cost $\hat{\ell}_3(1)$, while the optimal social cost under the original latencies is $2\alpha - 2 + \epsilon$. Hence, $\rho \geq \frac{\hat{\ell}_3(1)}{2\alpha - 2 + \epsilon}$. Therefore,

$$\rho \geq \max\left\{\hat{\ell}_1(1), \frac{\hat{\ell}_3(1)}{2\alpha - 2 + \epsilon}\right\} \geq \max\left\{\hat{\ell}_1(1), \frac{2 - \hat{\ell}_1(1)}{2\alpha - 2 + \epsilon}\right\},$$

with $\hat{\ell}_1(1) \leq 2\alpha - 1$. Letting $\epsilon \to 0$, the right-hand side is minimized by balancing the two terms, namely when $\hat{\ell}_1(1) = \frac{2 - \hat{\ell}_1(1)}{2\alpha - 2}$, which gives $\hat{\ell}_1(1) = \frac{2}{2\alpha - 1}$. This balancing point is feasible precisely when $\frac{2}{2\alpha - 1} \leq 2\alpha - 1$, equivalently, when $\alpha \geq (1 + \sqrt{2})/2$. Therefore, if $(1 + \sqrt{2})/2 \leq \alpha$, then $\rho \geq \frac{2}{2\alpha - 1}$. If $1 \leq \alpha \leq (1 + \sqrt{2})/2$, then the balancing point is infeasible, and the best lower bound is obtained at the endpoint $\hat{\ell}_1(1) = 2\alpha - 1$, yielding

$$\rho \geq \frac{2 - (2\alpha - 1)}{2\alpha - 2} = \frac{3 - 2\alpha}{2\alpha - 2}.$$

For $\alpha = 1$, this expression is interpreted as an unbounded lower bound by taking $\epsilon \to 0$.

Finally, the second lower bound is relevant as long as it is at least the consistency parameter itself, i.e., $\frac{2}{2\alpha - 1} \geq \alpha$. Solving this inequality gives $\alpha \leq (1 + \sqrt{17})/4$. Hence, for $(1 + \sqrt{2})/2 \leq \alpha < (1 + \sqrt{17})/4$, we obtain

$$\rho \geq \frac{2}{2\alpha - 1}.$$

This completes the proof.

□

D.2 Error Tolerance

In this section, we examine whether it is possible to design learning-augmented coordination mechanisms whose performance degrades gracefully as the prediction error increases in non-atomic games. Inspired by the notion of *brittleness* in learning-augmented algorithms [EADL24] (see also [BP25]), we formalize this notion of smooth performance degradation by requiring that the performance measure varies continuously with the prediction error, for sufficiently small errors. For our lower bounds, we use the standard monotonicity assumption, also used in [CMP14], that both the original and modified latencies are nondecreasing.

Our main impossibility (Theorem 8) shows that no mechanism can be 1-consistent while simultaneously achieving a graceful degradation to $4/3$, which coincides with the PoA without predictions. In other words, in order to improve the PoA of the original network, any mechanism achieving optimal consistency suffers a discontinuous performance that collapses to the original PoA value, even for arbitrarily small prediction error. This phenomenon highlights fundamental barriers to designing 1-consistent mechanisms in this setting.

We complement this negative result with a refined analysis that yields both positive and negative results: we show that this impossibility persists (albeit with weaker bounds) even in symmetric games (Theorem 9) and (multi-commodity) network congestion games (Theorem 10). Nonetheless, in the special case of single-commodity network congestion games, we exploit the network structure to obtain a positive result (Theorem 11).

Congestion Games. We consider the most general setting: non-atomic congestion games with no underlying network structure. We establish the following impossibility result.

Theorem 8. *There is no* 1*-consistent learning-augmented coordination mechanism for non-atomic congestion games that guarantees the existence of an equilibrium with approximation ratio strictly better than* $4/3$ *as the prediction error tends to zero.*

Proof. Consider a non-atomic congestion game defined as follows: there are two commodities 1 and 2, and a set of resources $E = \{e_1, e_2, e_3, e_4\}$. The set of latencies is given below: $\ell_1(x) = \ell_4(x) = 1$ and $\ell_2(x) = \ell_3(x) = x$. The set of strategies is $\Sigma_1 = \{\{e_1, e_2\}, \{e_2, e_3\}\} = \{S_1, S_2\}$ for commodity 1, and $\Sigma_2 = \{\{e_2, e_3\}, \{e_3, e_4\}\} = \{S_2, S_3\}$. For each commodity i the predicted demand is $\bar{r}_i = 1/2$. The optimal solution assigns $\bar{r}_1$ to S_1 and $\bar{r}_2$ to S_3, with a social cost $C_{opt} = 1/2(\ell_1(1/2) + \ell_2(1/2)) + 1/2(\ell_3(1/2) + \ell_4(1/2)) = 3/2$. There exists a Nash equilibrium in which both commodities route all their demands through S_2; its social cost is $C_N = 1/2(\ell_2(1) + \ell_3(1)) + 1/2(\ell_2(1) + \ell_3(1)) = 2$. Hence, the PoA is equal to $4/3$.

Consider any 1-consistent coordination mechanism. Suppose that $\hat{\ell}_2(x) < 1$ for some $x > 1/2$; then commodity 2 has a profitable deviation from the optimal allocation: commodity 2 improves its latency by routing $x - 1/2$ via S_2 and the rest via S_3 compared to routing all its demand via S_3. Observe that this also holds due to symmetry for $\hat{\ell}_3(x)$ and commodity 1. Combining the two facts, we obtain that $\hat{\ell}_2(x) \geq 1$ and $\hat{\ell}_3(x) \geq 1$ for all $x > 1/2$.

Now let the actual demands be $r_1 = r_2 = (1 + \epsilon)/2$ for $\epsilon > 0$ infinitesimally small. Note that e_2 belongs to all strategies of commodity 1 and e_3 belongs to all strategies of commodity 2; therefore, no matter the strategy profile, the load of e_2 and e_3 is more than $1/2$, resulting in their latency being at least 1. Hence, the social cost of any equilibrium in the modified game is at least $2(1 + \epsilon)$ which tends to 2 as ϵ goes to 0, while the optimal solution in the original network has social cost that tends to $3/2$ as ϵ goes to 0. We conclude that any equilibrium for this almost accurate demand is at least $4/3$. □

Theorem 9. *There is no* 1*-consistent coordination mechanism in non-atomic symmetric congestion games that guarantees the existence of an equilibrium with approximation ratio strictly better than* $7/6$ *as the prediction error tends to zero.*

Proof. Consider the symmetric congestion game (G, ℓ) defined as follows. The predicted demand is $\bar{r} = 1$ and the set of resources is given by $E = \{e_1, e_2, e_3, e_4\}$. The set of available strategies is $\Sigma = \{S_1, S_2, S_3\}$, where $S_1 = \{e_1, e_3\}$, $S_2 = \{e_2, e_4\}$, and $S_3 = \{e_1, e_2\}$, and the latency functions are defined as $\ell_1(x) = \ell_2(x) = x$ and $\ell_3(x) = \ell_4(x) = 1$.

Notice that this setup closely resembles the Braess network. The optimal solution assigns half of the demand to resources in S_1 and the other half to S_2, resulting in a total latency of $C_{opt}(\bar{r}) = 1/2(\ell_1(1/2) + \ell_3(1/2)) + 1/2(\ell_2(1/2) + \ell_4(1/2)) = 3/2$. However, in the unique equilibrium f, all players choose strategy S_3, yielding a total latency of $C_N(f) = 2$.

Regarding consistency, let $\bar{f}^*$ denote the optimal solution under the predicted input rate $\bar{r}$, and let $\hat{\ell}_e$ denote the modified latency function for each resource $e \in E$. Similarly to the characterization of 1-consistent mechanisms for parallel link networks given by [CCSV26], we establish necessary conditions on the modified latencies for $x > \bar{f}^*_e$.

Claim 8. *For any 1-consistent mechanism with modified latencies $\hat{\ell}$, the following hold for any $e \in E$:*

- $\hat{\ell}_e(\bar{f}^*_e) = \ell_e(\bar{f}^*_e)$ *and*
- $\liminf_{\epsilon \to 0^+} \hat{\ell}_e(\bar{f}^*_e + \epsilon) \geq 1$.

Proof. Let $r = \bar{r} = 1$. We first establish the first condition. Observe that the optimal solution of the congestion game is unique. Thus, in order to guarantee 1-consistency, $\bar{f}^*$ should be an equilibrium with respect to $\hat{\ell}$, and the latencies must remain untouched for any used resource in the optimal solution. Since all resources are used in the optimal solution, it should be $\hat{\ell}_e(\bar{f}^*_e) = \ell_e(\bar{f}^*_e)$, for all $e \in E$.

Regarding the second condition, it should be $\hat{\ell}_3(x) \geq \ell_3(x) = 1$ and $\hat{\ell}_4(x) \geq \ell_4(x) = 1$ for all $x > 0$. Regarding the rest of the resources, assume for the sake of contradiction that $\liminf_{\epsilon \to 0^+} \hat{\ell}_1(\bar{f}^*_1 + \epsilon) < 1$; the proof for e_2 is similar. Then, the optimal solution (that splits half of the rate through S_1 and half through S_2) is not an equilibrium. This is because the cost for strategy S_2 equals $3/2$, whereas deviating to strategy S_3 results in a cost strictly lower than $3/2$. □

Let now $r = 1 + \epsilon$, where $\epsilon > 0$. Then, in any equilibrium f, either $f_1 > 1/2$ or $f_2 > 1/2$ (or both). If w.l.o.g. $f_1 > 1/2$ and $f_2 \leq 1/2$, then $f_{S_2} + f_{S_3} = f_2 \leq 1/2$, leaving $f_{S_1} = r - f_2 > 1/2$ units of flow to experience a cost of at least 2. The remaining f_2 units experience a cost of at least $1 + f_2$. If both $f_1 > 1/2$ and $f_2 > 1/2$, all r units experience a cost of at least 2. Overall, let $x = f_{S_1}$ in the first case and $x = r$ in the second. Then $x > 1/2$ units experience a cost of at least 2, and the remaining $r - x$ units experience a cost of at least $1 + r - x$. Therefore,

$$\hat{C}(f) \geq 2x + (r - x)(1 + r - x) = 2r - \frac{1}{4} + \left(r - x - \frac{1}{2}\right)^2 \geq 2r - \frac{1}{4}.$$

In the original game, splitting the demand equally between S_1 and S_2 gives social cost $r + r^2/2$. Hence $C_{opt}(r) \leq r + r^2/2$. Taking the lower limit as ϵ goes to 0, for any equilibrium f at demand $1 + \epsilon$, yields

$$\liminf_{\epsilon \downarrow 0} \frac{\hat{C}(f)}{C_{opt}} \geq \lim_{\epsilon \downarrow 0} \frac{2(1+\epsilon) - 1/4}{(1+\epsilon) + (1+\epsilon)^2/2} = \frac{7}{6}.$$

□

Network congestion games: Multi-commodity. Having established a strong negative result in the general setting, we now examine whether additional structure can facilitate perfect consistency along with error tolerance guarantees. The natural restriction is to consider multi-commodity network congestion games, where players' strategies correspond to paths in a graph. Despite the added structure, we show that the impossibility persists even in this more constrained setting. Specifically, we establish that 1-consistency cannot be combined with an approximation ratio that approaches 1 as the prediction error tends to zero.

Theorem 10. *There is no* 1*-consistent coordination mechanism in multi-commodity network non-atomic congestion games that guarantees the existence of an equilibrium with approximation ratio strictly better than* $7/6$ *as the prediction error tends to zero.*

Proof. Consider a multi-commodity game defined over the Wheatstone network (also known as the Braess paradox network). The network consists of five edges $E = \{e_1, e_2, e_3, e_4, e_5\}$ with latencies $\ell_1(x) = \ell_4(x) = x$, $\ell_2(x) = \ell_3(x) = 1$, and $\ell_5(x) = 0$, and two commodities, (s_1, t_1) and (s_2, t_2) (as illustrated in Figure 2). Commodity 2 has a unique strategy: all flow from s_2 to t_2 is routed along edge e_5.

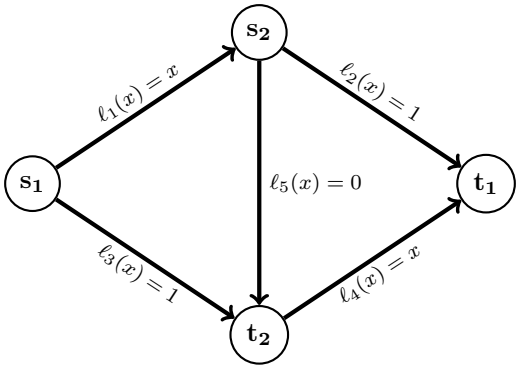

Figure 2: Network multicommodity game.

Let $\bar{r}_1 = \bar{r}_2 = 1$. In the optimal flow $\bar{r}_1$ uses only paths (s_1, s_2, t_1) and (s_1, t_2, t_1), whereas in the equilibrium flow it uses only the path (s_1, s_2, t_2, t_1); $\bar{r}_2$ uses only e_5 in both optimum and equilibrium flow. Any 1-consistent mechanism should preserve $\hat{\ell}_5(1) = 0$.

Then, 1-consistency requires that the predicted optimal flow be an equilibrium under the modified latencies, with $\hat{\ell}_1(1/2) = \hat{\ell}_4(1/2) = 1/2$ and $\hat{\ell}_2(1/2) = \hat{\ell}_3(1/2) = 1$. Moreover, $\liminf_{\epsilon \downarrow 0}[\hat{\ell}_1(1/2 + \epsilon) + \hat{\ell}_5(1 + \epsilon)] \geq 1$. Otherwise, commodity 1 could profitably deviate from (s_1, t_2, t_1) to (s_1, s_2, t_2, t_1): edge e_4 is common to both paths, while replacing e_3 by e_1 and e_5 would incur a cost strictly below 1. By symmetry, the same inequality holds with $\hat{\ell}_4$ in place of $\hat{\ell}_1$.

Let now $r_1 = r_2 = 1 + \epsilon$, where $\epsilon > 0$, and let f be any equilibrium. The flow identities give $f_1 + f_4 = f_5 \geq r_2 > 1$ and $f_2 + f_3 = 2r_1 - f_1 - f_4$. Thus, either $f_1 > 1/2$ or $f_4 > 1/2$. Assume w.l.o.g. that $f_1 > 1/2$. By monotonicity and the preceding inequality, $\hat{\ell}_1(f_1) + \hat{\ell}_5(f_5) \geq 1$. Using also $\hat{\ell}_4(f_4) \geq f_4$, $\hat{\ell}_5(f_5) \geq 0$, and $\hat{\ell}_2(f_2), \hat{\ell}_3(f_3) \geq 1$, we obtain

$$\begin{aligned}
\hat{C}(f) &\geq f_1[\hat{\ell}_1(f_1) + \hat{\ell}_5(f_5)] + f_4[\hat{\ell}_4(f_4) + \hat{\ell}_5(f_5)] + f_2 + f_3 \\
&\geq f_1 + f_4^2 + 2r_1 - f_1 - f_4 \\
&= 2r_1 - \frac{1}{4} + \left(f_4 - \frac{1}{2}\right)^2 \geq 2r_1 - \frac{1}{4}.
\end{aligned}$$

In the original game, splitting the demand of commodity 1 equally between the two outer paths gives social cost $r_1 + r_1^2/2$, while commodity 2 incurs zero cost. Hence, $C_{opt}(r_1, r_2) \leq r_1 + r_1^2/2$. Taking the lower limit as ϵ goes to zero, for any equilibrium at demands $r_1 = r_2 = 1 + \epsilon$, yields

$$\liminf_{\epsilon \downarrow 0} \frac{\hat{C}(f)}{C_{opt}} \geq \lim_{\epsilon \downarrow 0} \frac{2(1+\epsilon) - 1/4}{(1+\epsilon) + (1+\epsilon)^2/2} = \frac{7}{6}.$$

□

Network congestion games: Single-commodity. We now turn to single-commodity network congestion games and show that, unlike the multi-commodity case, the special topology of these games enables a positive result. In particular, given a prediction $\bar{r}$ and an error tolerance parameter $\bar{\eta}$, we construct a 1-consistent mechanism that admits an equilibrium whose approximation ratio varies

continuously with the input rate in $[\max\{0, \bar{r} - \bar{\eta}\}, \bar{r} + \bar{\eta}]$. We assume $\bar{r} > 0$ and $C_{opt}(\bar{r}) > 0$. For polynomial latencies with nonnegative coefficients, the latter assumption implies $C_{opt}(r) > 0$ for every $r > 0$.

Definition 5 (ERRORTOLERANT$(\bar{\eta}, \bar{r})$)**.** *Let $0 \leq \bar{\eta} < \infty$, and let $\bar{f}^*$ be an optimal flow at the predicted rate $\bar{r}$ whose positive edge support is acyclic*[13] *Choose a maximum-cost s–t path in its positive edge support, i.e., $p^* \in \arg\max_{p \in \mathcal{P}:\ \bar{f}^*_e > 0 \text{ for all } e \in p} \ell_p(\bar{f}^*)$. Let $u_e = \bar{f}^*_e + \bar{\eta}\mathbf{1}_{\{e \in p^*\}}$ and choose $W \geq \sum_{e \in p^*} \ell_e(u_e)$. Define*

$$\hat{\ell}_e(x) = \begin{cases} \ell_e(x), & 0 \leq x \leq u_e, \\ \max\{W, \ell_e(x)\}, & x > u_e. \end{cases}$$

In the next theorem we show that the ERRORTOLERANT mechanism achieves the combination of 1-consistency and error tolerance.

Theorem 11. *In single-commodity network non-atomic congestion games, the* ERRORTOLERANT$(\bar{\eta}, \bar{r})$ *mechanism is 1-consistent and error tolerant for any input rate $r \in [\max\{0, \bar{r} - \bar{\eta}\}, \bar{r} + \bar{\eta}]$.*

Proof. Given a predicted input rate $\bar{r}$, let $\bar{f}^*$ be an acyclic optimal flow under $\bar{r}$, and let p^* be defined as above. Consider any input rate r, with $\bar{r} \leq r \leq \bar{r} + \bar{\eta}$, and let f be the flow that routes $\bar{r}$ as $\bar{f}^*$ does, and the extra $r - \bar{r}$ via p^*, i.e., $f_e = \bar{f}^*_e$ for any $e \notin p^*$, and $f_e = \bar{f}^*_e + (r - \bar{r})$ for any $e \in p^*$. Since $f_e \leq u_e$ for every $e \in E$, the modified latencies coincide with the original latencies at f. We first show that f is an equilibrium for every $r \in [\bar{r}, \bar{r} + \bar{\eta}]$. For this, let p_1 be any path used by f, i.e., $f_{p_1} > 0$, and let p_2 be any other path in the network. We establish the following claim:

Claim 9. *For every path p_1 used by f and every alternative path p_2, it holds that*

$$\sum_{e \in p_1 \setminus p_2} \hat{\ell}_e(f_e) \leq \liminf_{\epsilon \downarrow 0} \sum_{e \in p_2 \setminus p_1} \hat{\ell}_e(f_e + \epsilon).$$

Proof. First, observe that p^* remains a maximum-cost path in the positive edge support of $\bar{f}^*$ under f. Indeed, the additional flow is routed exclusively through p^*, the original latencies are nondecreasing, and $\hat{\ell}_e(f_e) = \ell_e(f_e)$ for every $e \in E$. We distinguish two cases based on whether there exists an edge in $p_2 \setminus p_1$ that does not belong to p^*.

If $(p_2 \setminus p_1) \setminus p^* \neq \emptyset$, then there exists an edge $e' \in p_2 \setminus p_1$ such that $e' \notin p^*$. Since $f_{e'} = \bar{f}^*_{e'} = u_{e'}$, by the definition of the mechanism, for any $\epsilon > 0$, $\sum_{e \in p_2 \setminus p_1} \hat{\ell}_e(f_e + \epsilon) \geq \hat{\ell}_{e'}(f_{e'} + \epsilon) \geq W$. Consequently, $\liminf_{\epsilon \downarrow 0} \sum_{e \in p_2 \setminus p_1} \hat{\ell}_e(f_e + \epsilon) \geq W$. On the other hand, since p_1 belongs to the positive edge support of $\bar{f}^*$,

$$\begin{aligned} \sum_{e \in p_1 \setminus p_2} \hat{\ell}_e(f_e) &\leq \hat{\ell}_{p_1}(f) \\ &\leq \hat{\ell}_{p^*}(f) \\ &\leq \hat{\ell}_{p^*}(u) \\ &= \sum_{e \in p^*} \ell_e(\bar{f}^*_e + \bar{\eta}) \leq W. \end{aligned}$$

This establishes the claim in the first case.

If $(p_2 \setminus p_1) \setminus p^* = \emptyset$, then every edge in $p_2 \setminus p_1$ belongs to p^*. We consider the alternative path $p' = \big(p^* \setminus (p_2 \setminus p_1)\big) \cup (p_1 \setminus p_2)$. This is a valid s–t path in the positive edge support of $\bar{f}^*$. Since p^* remains a maximum-cost path in this support under f, we have $\hat{\ell}_{p^*}(f) \geq \hat{\ell}_{p'}(f)$. Furthermore, by the construction of p', $\hat{\ell}_{p_1}(f) + \hat{\ell}_{p^*}(f) = \hat{\ell}_{p_2}(f) + \hat{\ell}_{p'}(f)$. Combining these inequalities and cancelling the common edges of p_1 and p_2 yields $\hat{\ell}_{p_2 \setminus p_1}(f) \geq \hat{\ell}_{p_1 \setminus p_2}(f)$. Consequently,

$$\liminf_{\epsilon \downarrow 0} \sum_{e \in p_2 \setminus p_1} \hat{\ell}_e(f_e + \epsilon) \geq \sum_{e \in p_2 \setminus p_1} \hat{\ell}_e(f_e) \geq \sum_{e \in p_1 \setminus p_2} \hat{\ell}_e(f_e),$$

[13] Such a flow is guaranteed to exist [Rou02].

where the first inequality follows from the nondecreasing modified latencies. □

By Claim 9, we next show that the condition for f bein an equilibrium is satisfied:

$$\begin{aligned}\hat{\ell}_{p_1}(f) &= \sum_{e\in p_1\cap p_2} \hat{\ell}_e(f_e) + \sum_{e\in p_1\setminus p_2} \hat{\ell}_e(f_e) \\ &\le \sum_{e\in p_1\cap p_2} \hat{\ell}_e(f_e) + \liminf_{\epsilon\downarrow 0} \sum_{e\in p_2\setminus p_1} \hat{\ell}_e(f_e+\epsilon) \\ &= \liminf_{\epsilon\downarrow 0} \hat{\ell}_{p_2}(f+\epsilon\mathbf{1}_{p_2}-\epsilon\mathbf{1}_{p_1}).\end{aligned}$$

Thus, f is a user equilibrium of the modified game for every $r \in [\bar{r}, \bar{r}+\bar{\eta}]$. Moreover, since $f_e \le u_e$ for every edge, the modified latencies coincide with the original latencies at f. Consequently, $\hat{C}(f) = C(f)$ varies continuously with r, as the additional flow is routed through p^*. Since the mechanism ensures that the cost on this path does not increase abruptly for small changes in the flow, the error tolerance condition is satisfied. In particular, for $r = \bar{r}$, we recover $\bar{f}^*$ with $\hat{C}(\bar{f}^*) = C_{\mathrm{opt}}(\bar{r})$.

For $r \le \bar{r}$, consider a Beckmann–McGuire–Winsten (BMW) equilibrium f in the network with edge capacities u_e, i.e., a minimizer of the original Beckmann potential subject to $f_e \le u_e$ for every $e \in E$ [CSM04]. Existence and the corresponding variational inequalities follow from [CSM04]. It is easy to see that this also a user equilibrium under the modified latencies. Any deviation that exceeds a capacity incurs a modified latency of at least W on some edge, and therefore, cannot be profitable. Finally, all BMW equilibria at a fixed rate have the same social cost, as follows from [CSM04, Lemma 3.3] and monotonicity. Since the capacitated feasible sets vary continuously with the input rate, compactness and continuity of the Beckmann potential imply that this common social cost is continuous in r. At $r = \bar{r}$, the optimal flow $\bar{f}^*$ is such a minimizer, since it minimizes the modified potential and the two potentials coincide below the thresholds. Consequently, the equilibrium cost converges to $C_{\mathrm{opt}}(\bar{r})$ as $r \uparrow \bar{r}$. □

E Arithmetic Lemmas

Lemma 5. *For any integer $k \in [1, d]$, non-negative real numbers x, y and real number $\lambda > 0$, it holds that*

$$x^k y \le \max\left\{\frac{1}{4\lambda}, \frac{d}{(d+1)\sqrt[d]{\lambda(d+1)}}\right\} x^{k+1} + \lambda y^{k+1}.$$

Proof. Consider first the following arithmetic inequality. For any integer $k \in [1, d]$, non-negative real numbers x, y and real number $\lambda \ge 1$, it holds that:

$$x^k y \le \frac{k}{(k+1)\sqrt[k]{\lambda(k+1)}} x^{k+1} + \lambda y^{k+1}.$$

This is used in [CKS11], but here we parameterize it over a parameter λ which will be determined later in order to optimize the robustness bound. Considering only positive $k > 0$, the coefficient $\frac{k}{(k+1)\sqrt[k]{\lambda(k+1)}}$ is zero for $k = 0$ and strictly increasing when $\lambda \ge 1$. In the case of $\lambda < 1$ it achieves its minimum at some value $k_{min} > 0$ and is strictly decreasing for $k < k_{min}$ and strictly increasing for $k > k_{min}$. Overall, since $k \in [1, d]$, the expression achieves its maximum either for $k = 1$ or for $k = d$, and therefore:

$$x^k y \le \max\left\{\frac{1}{4\lambda}, \frac{d}{(d+1)\sqrt[d]{\lambda(d+1)}}\right\} x^{k+1} + \lambda y^{k+1}.$$

□

Lemma 6. *For any integers $x \geq 1$, the following hold:*

(1) $\sum_{i=1}^{x} \left(i^k - (i-1)^k\right) = x^k$*, for any integer* $k \geq 1$;

(2) $(k+1)x^k \geq x^{k+1} - (x-1)^{k+1}$*, for any integer* $k \geq 0$;

(3) $(k+1)\sum_{i=1}^{x} i^k \geq x^{k+1}$*, for any integer* $k \geq 0$;

Proof. The first equality holds for $k \geq 1$ because

$$\sum_{i=1}^{x} \left(i^k - (i-1)^k\right) = x^k + \sum_{i=1}^{x-1} i^k - \sum_{i=1}^{x} (i-1)^k = x^k + \sum_{i=1}^{x-1} i^k - \sum_{i=0}^{x-1} i^k = x^k \,.$$

For the second inequality we will show that $f(x,k) = (k+1)x^k - x^{k+1} + (x-1)^{k+1} \geq 0$ by induction on k. For $k = 0$, $f(x,0) = 0$ for any $x \geq 1$. Consider any $k \geq 1$, and suppose for the inductive hypothesis that $f(x,k-1) \geq 0$ for all $x \geq 1$. Taking the derivative w.r.t. x, $df(x,k)/dx = (k+1)f(x,k-1) \geq 0$, for all $x \geq 1$, meaning that $f(x,k)$ is non-decreasing w.r.t. x. Since, $f(1,k) = k \geq 0$, $f(x,k) \geq 0$, for any $x \geq 1$ and the statement follows.

The third inequality is derived immediately by the second one as follows:

$$(k+1)\sum_{i=1}^{x} i^k = \sum_{i=1}^{x} (k+1)i^k \geq \sum_{i=1}^{x} \left(i^{k+1} - (i-1)^{k+1}\right) = x^{k+1} \,.$$

□

Lemma 7. *For every $x, y, z \in \mathbb{N}$ and real number $\beta \geq 1$:*

$$y\left(x + \frac{1}{\beta}\right) \leq \frac{1}{2z+1}x^2 + \frac{z^2 + (1+\frac{2}{\beta})z + \frac{\beta+3}{4\beta}}{2z+1}y^2$$

Proof. We subtract the left hand side of the statement from the right hand side, multiply by $2z+1$, and use $f(x,y,\beta)$ to denote the result. Then, we have:

$$\begin{aligned} f(x,y,\beta) &= x^2 + \left(z^2 + \left(1 + \frac{2}{\beta}\right)z + \frac{\beta+3}{4\beta}\right)y^2 - (2z+1)y\left(x + \frac{1}{\beta}\right) \\ &= \left(x - \frac{2z+1}{2}y\right)^2 + \left(\frac{2}{\beta}z + \frac{3}{4\beta}\right)y^2 - \frac{2z+1}{\beta}y \\ &= \left(x - \frac{2z+1}{2}y\right)^2 + \frac{(8z+3)y - 8z - 4}{4}\frac{y}{\beta} \,. \end{aligned}$$

We want to show that $f(x,y,\beta) \geq 0$ for all nonnegative integers x, y, z. For $y = 0$, this is clearly true. For any $y \geq 2$, $f(x,y,\beta)$ is positive since $2(8z+3) > 8z+4$. It remains to check the case $y = 1$; we have that $f(x,1,\beta) = x^2 - (2z+1)x + z^2 + z + \frac{1}{4} - \frac{1}{4\beta} = (x-z)(x-z-1) + \frac{1}{4} - \frac{1}{4\beta} \geq 0$, which is nonnegative for x, z integers and $\beta \geq 1$. So, the lemma follows. □

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The abstract and introduction state the paper's theoretical contributions and scope, including the assumptions under which the results hold.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: The paper discusses the limitations of its theoretical results, including the assumptions under which the theorems hold and the resulting scope of the claims. Since the paper is theoretical, there are no empirical, dataset, or implementation limitations to report.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: The paper is theoretical, and each formal result states or references the assumptions under which it holds; see Section 3. Complete proofs are provided in Appendix B, C and D, with supporting definitions and lemmas cross-referenced as needed.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [N/A]

   Justification: The paper does not include experiments, and its claims are theoretical.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [N/A]

Justification: The paper does not include experiments, datasets, or implementation-based results requiring code. Its contributions are theoretical and are supported by formal statements and proofs.

Guidelines:

- The answer [N/A] means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://neurips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://neurips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

Answer: [N/A]

Justification: The paper does not include experiments, training procedures, test sets, data splits, hyperparameters, or optimizers. The results are theoretical and are specified through formal assumptions, statements, and proofs.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [N/A]

Justification: The paper does not include experiments or empirical measurements.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [N/A]

   Justification: The paper does not include experiments.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

   Answer: [Yes]

   Justification: The authors have reviewed the NeurIPS Code of Ethics, and the research conforms to it. The paper is theoretical and does not involve human subjects, datasets, model training, or released research artifacts that would raise additional ethics considerations.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [N/A]

    Justification: The paper is a theoretical contribution and does not introduce a deployed system, dataset, model, or algorithmic implementation with direct societal impact. While the motivation is related to congestion games, social welfare, and improved understanding

of decentralized resource allocation, including domains such as transportation, assessing concrete positive or negative societal impacts would require additional application-specific modeling beyond the scope of this work.

Guidelines:

- The answer [N/A] means that there is no societal impact of the work performed.
- If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: The contribution is theoretical. The paper does not release data, models, code, or other artifacts that pose a high risk for misuse.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [N/A]

    Justification: The paper does not use existing assets such as code, datasets, models, benchmarks, or software packages. The paper only cites prior scholarly work where relevant.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.

- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification: The paper's contribution is theoretical and is documented through its formal statements and proofs. The paper does not release new assets such as datasets or models.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: The paper does not involve crowdsourcing or data collection from human subjects. The contribution is theoretical.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

Justification: The paper does not involve crowdsourcing or data collection from human subjects. The contribution is theoretical.

Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [N/A]

Justification: LLMs were only used for editing and formatting assistance.

Guidelines:

- The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.