

Chasing Convex Functions with Long-term Constraints

Adam Lechowicz¹ Nicolas Christianson² Bo Sun³ Noman Bashir⁴ Mohammad Hajiesmaili¹
 Adam Wierman² Prashant Shenoy¹

Abstract

We introduce and study a family of online metric problems with long-term constraints. In these problems, an online player makes decisions \mathbf{x}_t in a metric space (X, d) to simultaneously minimize their hitting cost $f_t(\mathbf{x}_t)$ and switching cost as determined by the metric. Over the time horizon T , the player must satisfy a long-term demand constraint $\sum_t c(\mathbf{x}_t) \geq 1$, where $c(\mathbf{x}_t)$ denotes the fraction of demand satisfied at time t . Such problems can find a wide array of applications to online resource allocation in sustainable energy/computing systems. We devise optimal competitive and learning-augmented algorithms for the case of bounded hitting cost gradients and weighted ℓ_1 metrics, and further show that our proposed algorithms perform well in numerical experiments.

1. Introduction

This paper introduces and studies a novel class of online metric problems with *long-term demand constraints* motivated by emerging applications in the design of sustainable systems. In *convex function chasing with a long-term constraint*, an online player aims to satisfy a demand by making decisions in a normed vector space, paying a hitting cost based on time-varying convex cost functions which are revealed online, and switching cost defined by the norm. The player is constrained to ensure that the entire demand is satisfied at or before the time horizon T ends, and their objective is to minimize their total cost. The generality of this problem makes it applicable to a wide variety of online

resource allocation problems; in this paper, we consider one such special case, discussing its connections to other online settings and suggestions towards broad new areas of inquiry in *online optimization with long-term constraints*.

Our motivation to introduce these problems is rooted in an emerging class of *carbon-aware* control problems for sustainable systems. A shared objective involves minimizing carbon emissions by shifting flexible workloads temporally and/or spatially to better leverage low-carbon electricity generation (e.g., renewables such as solar and wind). Examples which have recently seen significant interest include carbon-aware electric vehicle (EV) charging (Cheng et al., 2022) and carbon-aware compute shifting (Wiesner et al., 2021; Bashir et al., 2021; Radovanovic et al., 2022; Acun et al., 2023; Hanafy et al., 2023).

The problems we introduce in this paper build on a long line of related work in online algorithms. Most existing work can be roughly classified into two types: *online metric problems*, where many works consider multidimensional decision spaces and switching costs but do not consider long-term constraints (Borodin et al., 1992; Koutsoupias, 2009; Chen et al., 2018; Bubeck et al., 2019; 2021; Bansal & Coester, 2022; Bubeck et al., 2023), and *online search problems*, which feature long-term demand constraints but do not consider multidimensional decision spaces or switching costs (El-Yaniv et al., 2001; Lorenz et al., 2008; Mohr et al., 2014; Sun et al., 2021b).

We briefly review the direct precursors of our work below. In the online metric literature, the problem we study is an extension of *convex function chasing* (CFC) introduced by Friedman & Linial (1993), where an online player makes online decisions \mathbf{x}_t in a normed vector space $(X, \|\cdot\|)$ over a sequence of time-varying cost functions in order to minimize their total hitting and switching cost. In the online search literature, the problem we study is a generalization of *one-way trading* (OWT) introduced by El-Yaniv et al. (2001), in which an online player must sell an entire asset in fractional shares over a sequence of time-varying prices while maximizing their profit.

Despite extensive existing work in the online metric and online search tracks, few works simultaneously consider long-term demand constraints (as in OWT) and move-

¹Manning College of Information and Computer Sciences, **University of Massachusetts Amherst**, USA. ²Computing & Mathematical Sciences, **California Institute of Technology**, USA.
³Cheriton School of Computer Science, **University of Waterloo**, Ontario, Canada. ⁴Computer Science & Artificial Intelligence Laboratory, **Massachusetts Institute of Technology**, USA.. Correspondence to: Adam Lechowicz <alechowicz@cs.umass.edu>.

Proceedings of the 41st International Conference on Machine Learning, Vienna, Austria. PMLR 235, 2024. Copyright 2024 by the author(s).

ment/switching costs (as in CFC). The existing prior works (Lechowicz et al., 2023; 2024) that consider both components are restricted to unidimensional decision spaces, as is typical in the online search literature. However, generalizing from the unidimensional case is highly non-trivial; e.g., in convex function chasing with a long-term constraint, the problem cannot simply be decomposed over dimensions due to the shared constraint function and multidimensional switching cost. Thus, in this work we tackle the following question: *Is it possible to design algorithms for the studied problems that operate in multidimensional decision spaces while simultaneously considering long-term constraints, hitting costs, and switching costs?*

Although the aforementioned literature focuses on competitive algorithms in adversarial settings, there has recently been significant interest in moving beyond worst-case analysis, which can result in overly pessimistic algorithms. The field of *learning-augmented algorithms* (Lykouris & Vassilvitskii, 2018; Kumar et al., 2018) has emerged as a paradigm for designing and analyzing algorithms that incorporate untrusted machine-learned advice to improve average-case performance without sacrificing worst-case performance bounds. Such algorithms are evaluated through the metrics of *consistency* and *robustness* (see Def. 2.1). Recent studies have proposed learning-augmented algorithms for related problems, including convex function chasing (Christianson et al., 2022), one-way trading (Sun et al., 2021a), metrical task systems (Christianson et al., 2023), and online search (Lee et al., 2024). While the literature in each of these tracks considers a spectrum of different advice models, their results prompt a natural open question: *Can we design algorithms for online metric problems with long-term constraints that effectively utilize untrusted advice (such as machine-learned predictions) to improve performance while preserving worst-case competitive guarantees?*

Contributions. Despite extensive prior literature on adjacent problems, the problems we propose in this paper are the first online settings to combine long-term demand constraints with multidimensional decision spaces and switching costs. We introduce convex function chasing with a long-term constraint, and a special case called *online metric allocation with a long-term constraint*. The general forms of both are independently interesting for further study.

We obtain positive results for both of the questions posed above under problem instantiations that are especially relevant for motivating applications (namely, weighted ℓ_1 norms, and cost functions with bounded gradients). We provide the first competitive results for online problems of this form in Section 3, and show that our proposed algorithm (Algorithm 1) achieves the best possible competitive ratio. In Section 4, we propose a learning-augmented algorithm, CLIP (Algorithm 2), and show it achieves the provably optimal trade-off between consistency and robustness.

To achieve these results, the proposed algorithms must tackle technical challenges distinct from prior work on adjacent problems. Motivated by difficulties in directly applying algorithms for unidimensional decision spaces in the online search literature, we build on a generalization of threshold-based design called *pseudo-cost minimization*. While this framework is well-known in the online search literature, it is a new idea in the context of online metric problems and multidimensional decision spaces (see Section 3).

Our learning-augmented algorithm CLIP uses a novel adaptive optimization-based approach to achieve specific target consistency and robustness bounds. In recent years, there has been interest in understanding fundamental trade-offs between consistency and robustness and designing algorithms that exactly match those trade-offs, which has proven to be non-trivial (Wei & Zhang, 2020). CLIP’s approach directly incorporates the lower bound and exactly matches it, distinguishing it from algorithms that achieve, e.g., an asymptotically optimal trade-off. To achieve this, CLIP introduces a *projected consistency constraint* designed to guarantee consistency against the advice ADV by continuously comparing solutions in terms of cost incurred so far, switching cost trajectories, and the projected worst-case cost required to complete the long-term constraint. We believe CLIP’s high-level approach is applicable to other problems and may improve current results in the broader field of learning-augmented algorithms.

2. Problem Formulation and Preliminaries

This section formalizes convex function chasing and online metric allocation with long-term constraints, motivating them with a sustainability application. We also provide preliminaries used throughout the paper, and give initial results to build algorithmic connections between both problems.

Convex function chasing with a long-term constraint. A player chooses decisions $\mathbf{x}_t \in X \subseteq \mathbb{R}^d$ online from a normed vector space $(X, \|\cdot\|)$ in order to minimize their total cost $\sum_{t=1}^T f_t(\mathbf{x}_t) + \sum_{t=1}^{T+1} \|\mathbf{x}_t - \mathbf{x}_{t-1}\|$, where $f_t(\cdot) : X \rightarrow \mathbb{R}$ is a convex “hitting” cost that is revealed just before the player chooses \mathbf{x}_t , and $\|\mathbf{x}_t - \mathbf{x}_{t-1}\|$ is a switching cost associated with changing decisions between rounds. Additionally, the player must satisfy a long term constraint of the form $\sum_{t=1}^T c(\mathbf{x}_t) = 1$, where $c(\mathbf{x}) : X \rightarrow [0, 1]$ gives the fraction of the constraint satisfied by a decision \mathbf{x} . The offline version of the problem is formalized as follows:

$$\min_{\{\mathbf{x}_t\}_{t \in [T]}} \underbrace{\sum_{t=1}^T f_t(\mathbf{x}_t)}_{\text{Convex hitting cost}} + \underbrace{\sum_{t=1}^{T+1} \|\mathbf{x}_t - \mathbf{x}_{t-1}\|}_{\text{Switching cost}} \quad (1)$$

$$\text{s.t. } \underbrace{\sum_{t=1}^T c(\mathbf{x}_t)}_{\text{Long-term constraint}} \geq 1, \quad (2)$$

$$\mathbf{x}_t^i \in [0, 1] \quad \forall i \in [d], \forall t \in [T]. \quad (3)$$

We denote the *utilization* at time t by $z^{(t)} = \sum_{\tau=1}^t c(\mathbf{x}_\tau)$, which gives the total fraction of the long-term constraint satisfied up to and including time t .

Assumptions. Here, we describe the precise variant of convex function chasing with a long-term constraint for which we design algorithms in the remainder of the paper. Let $\|\mathbf{x} - \mathbf{x}'\| := \|\mathbf{x} - \mathbf{x}'\|_{\ell_1(\mathbf{w})}$, where $\|\cdot\|_{\ell_1(\mathbf{w})}$ denotes the weighted ℓ_1 norm with weight vector $\mathbf{w} \in \mathbb{R}^d$.

We define the long-term constraint such that $c(\mathbf{x}) := \|\mathbf{x}\|_{\ell_1(\mathbf{c})}$, i.e., the weighted ℓ_1 norm with weight vector $\mathbf{c} \in \mathbb{R}^d$. Then let the metric space X be the ℓ_1 ball defined by $X := \{\mathbf{x} \in \mathbb{R}^d : c(\mathbf{x}) \leq 1\}$. Note that \mathbf{x} is restricted to lie in the positive orthant by (3).

For all cost functions $f_t(\cdot) : X \rightarrow \mathbb{R}$, we assume bounded gradients such that $L \leq [\nabla f_t]^i / c^i \leq U \forall i \in [d], t \in [T]$, where i denotes the i th dimension of the corresponding vector, and L, U are known positive constants. This also gives as a corollary that $f_t(\mathbf{x}) \geq 0$ for any valid \mathbf{x} .

Letting $\mathbf{0}$ denote the origin in \mathbb{R}^d (w.l.o.g), we have the property $f_t(\mathbf{0}) = 0$ for all $t \in [T]$, i.e., that “satisfying none of the long-term constraint costs nothing”, since $c(\mathbf{0}) = 0$. We assume the player starts and ends at the origin, i.e., $\mathbf{x}_0 = \mathbf{0}$ and $\mathbf{x}_{T+1} = \mathbf{0}$, to enforce switching “on” and “off”. These assumptions are intuitive and reasonable in practice, e.g., in our example motivating application below.

For analysis, it will be useful to establish a shorthand for the magnitude of the switching cost. Let $\beta := \max(\mathbf{w}^i / c^i)$, which gives the greatest magnitude of the switching cost coefficient when normalized by the constraint function. We assume that β is bounded on the interval $[0, U-L/2]$; if β is “large” (i.e., $> U-L/2$), we can show that the player should prioritize minimizing the switching cost.¹

Recall that the long-term constraint must be satisfied before the sequence ends. If the player has satisfied $z^{(t)}$ of the constraint at time t , we assume a compulsory trade begins at time j as soon as $(T - (j+1)) \cdot \mathbf{c}^i < (1 - z^{(j)}) \forall i \in [d]$ (i.e., when the time steps after j are not sufficient to satisfy the constraint by moving to a point at the boundary defined by (3)). During this compulsory trade, a cost-agnostic algorithm takes over, making maximal decisions to satisfy the constraint. T is unknown in advance, but the player is notified when this compulsory trade should begin.² For the problem to remain technically interesting, we assume that

¹As brief justification for the bounds on β , consider that a feasible solution may have objective value $L + 2\beta$. If $\beta > U-L/2$, $L + 2\beta > U$, and we argue that the incurred switching cost is more important than the cost functions accepted.

²Note that pre-notification of T is necessary to ensure constraint satisfaction – if T is completely unknown, the algorithm has very little flexibility to do anything without risking a violation of the long-term constraint.

this compulsory trade is a small portion of the sequence.³

For brevity, we henceforth use CFL to refer to the variant of convex function chasing with a long-term constraint under the assumptions outlined here. While this setting is challenging theoretically and useful for real-world applications as outlined below, it is worth mentioning that our solution techniques in Sections 3 and 4 do not heavily rely on the idiosyncrasies of, e.g., the ℓ_1 norm. Furthermore, as is common in the literature, the results in the rest of this paper can extend to other metrics by leveraging finite-dimension norm equivalencies (Johnson, 2020).

An example motivating application. CFL can model a variety of applications, including specific applications that motivate this study. Consider a *carbon-aware temporal load shifting* application with heterogeneous servers. Here, each of the d dimensions corresponds to one of d heterogeneous servers. An algorithm makes decisions $\mathbf{x}_t \in \mathbb{R}^d$, where $\mathbf{x}_t^i \in [0, 1]$ denotes the load of the i th server at time t . The long-term constraint $\sum_{t=1}^T c(\mathbf{x}_t) \geq 1$ enforces that an entire workload should be finished before time T , and each coefficient \mathbf{c}^i represents the throughput of the i th server. Each cost function $f_t(\mathbf{x}_t)$ represents the carbon emissions due to the electricity usage of the servers configured according to \mathbf{x}_t , and the switching cost $\|\cdot\|_{\ell_1(\mathbf{w})}$ captures the carbon emissions overhead (e.g., extra latency) of pausing, resuming, scaling, and moving the workload between servers.

Our motivation for the assumptions placed on CFL are deeply rooted in this and other similar applications, where the switching cost and constraint function are both typically best modeled as a linear (i.e., ℓ_1) function, and bounds on the marginal hitting cost (i.e., the bounded gradient assumption) are reasonable to obtain.

Online metric allocation with a long-term constraint. Bansal & Coester (2022) introduced the online metric allocation problem (MAP), which connects several online metric problems. MAP on a star metric is equivalent to CFC when cost functions are separable over dimensions and supported on the unit simplex Δ_n .⁴ Furthermore, the randomized metrical task systems problem (MTS) is a special case of MAP when cost functions are linear and increasing.

We build on this formulation in our setting and introduce *online metric allocation with a long-term constraint*, which

³We assume the first time j' where $(T - (j'+1)) \cdot \mathbf{c}^i < 1 \forall i$ satisfies $j' \gg 1$, implying that T and \mathbf{c} are both appropriate for the constraint. This is reasonable for our motivating applications, since short deadlines (small T) or low throughput (small $\mathbf{c}^i \forall i$) imply that even offline solutions suffer a lack of flexibility in reducing the overall cost.

⁴Given metric space X , consider $\Delta(X)$, which represents the set of probability measures over the points of X . Since X is finite, we have that $|X| = n$ and $\Delta(X)$ is denoted as Δ_n .

captures a particularly interesting special case of CFL. The general version of the problem considers an n -point metric space (X, d) , and a unit resource which can be allocated in arbitrary fractions to the points of X . At each time $t \in [T]$, convex cost functions $f_t^a(\cdot) : [0, 1] \rightarrow \mathbb{R}$ arrive at each point a in the metric space. The online player chooses an allocation x_t^a to each point a in the metric space, such that $\sum_{a=1}^n x_t^a = 1$ for all $t \in [T]$. When changing this allocation between time steps, the player pays a switching cost defined by $d(a, b)$ for any distinct points $a, b \in X$. As in CFL, the long-term constraint enforces that $\sum_{t=1}^T c(\mathbf{x}_t) \geq 1$, where $c(\mathbf{x})$ is a linear and separable function of the form $c(\mathbf{x}) = \sum_{a=1}^n \mathbf{c}^a x^a$. As previously, the player's objective is to minimize the total cost (hitting plus switching costs) incurred while satisfying the long-term constraint.

Assumptions. In the rest of the paper, we consider an instantiation of online metric allocation with a long-term constraint on *weighted star metrics* that is particularly relevant to a wide class of resource allocation problems.

To ensure the long-term constraint is non-trivial, we denote at least one point a' in the metric space as an “OFF state”, where $\mathbf{c}^{a'} = 0$ and $f_t^{a'}(x) = 0 \forall t \in [T], \forall x \in [0, 1]$. We do not require that an OFF state be situated at the center of the star, although this condition creates a useful special case of MAL (see Lemma 2.2 for details). For all other cost functions, we carry forward the assumptions that $L \leq \left(\frac{df_t^a}{dx^a}/\mathbf{c}^a\right) \leq U$, $f_t^a(0) = 0 \forall t \in [T]$. We define $\beta := \max_{a', a} d(a', a)/\mathbf{c}^a$, i.e., the maximum distance between any OFF state and any ON state in the weighted star, normalized by the value of \mathbf{c}^a at a specific ON state. We inherit the same assumption that $\beta \in [0, U-L/2]$. For brevity, we henceforth use MAL to refer to the problem on weighted star metrics with the assumptions described above.

Competitive analysis. Our goal is to design an algorithm that guarantees a small *competitive ratio* (Manasse et al., 1988; Borodin et al., 1992), i.e., performs nearly as well as the offline optimal solution. Formally, let $\mathcal{I} \in \Omega$ denote a valid input sequence, where Ω is the set of all feasible inputs for the problem. Let $\text{OPT}(\mathcal{I})$ denote the cost of an optimal offline solution for instance \mathcal{I} , and let $\text{ALG}(\mathcal{I})$ denote the cost incurred by running an online algorithm ALG over the same instance. The competitive ratio is then defined as $\text{CR}(\text{ALG}) := \sup_{\mathcal{I} \in \Omega} \frac{\text{ALG}(\mathcal{I})}{\text{OPT}(\mathcal{I})} = \eta$, and ALG is said to be η -competitive. Note that $\text{CR}(\text{ALG})$ is always ≥ 1 , and a lower competitive ratio implies that the online algorithm is guaranteed to be *closer* to the offline optimal solution.

Learning-augmented consistency and robustness. In the emerging literature on learning-augmented algorithms, competitive ratio is interpreted via the notions of *consistency* and *robustness*, introduced by (Lykouris & Vassilvitskii,

2018; Kumar et al., 2018).

Definition 2.1. Let LALG denote a learning-augmented online algorithm provided with advice denoted by ADV. Then LALG is said to be **b -consistent** if it is b -competitive with respect to ADV. Conversely, LALG is **r -robust** if it is r -competitive with respect to OPT when given any ADV (i.e., regardless of the performance of ADV).

A connection between CFL and MAL. Below we state two useful results connecting the CFL and MAL settings that influence our algorithm design for each problem.

Lemma 2.2. For any MAL instance on a weighted star metric (X, d) , there is a corresponding CFL instance on $(\mathbb{R}^{n-1}, \|\cdot\|_{\ell_1(\mathbf{w}^*)})$ that preserves $f_t^a(\cdot) \forall t, c(\cdot) \forall a \in X$. Furthermore, for any points $(a, b) \in X$, their distance is upper bounded by a weighted ℓ_1 norm between the corresponding vectors $(\mathbf{a}, \mathbf{b}) \in \mathbb{R}^{n-1}$, i.e., $d(a, b) \leq \|\mathbf{a} - \mathbf{b}\|_{\ell_1(\mathbf{w}^*)}$. If the MAL instance contains an OFF state at the center of the weighted star X , distances are preserved exactly, i.e., $d(a, b) = \|\mathbf{a} - \mathbf{b}\|_{\ell_1(\mathbf{w}^*)}$.

Leveraging Lemma 2.2, the following result explicitly connects the competitive results of the CFL and MAL settings.

Proposition 2.3. Given an algorithm ALG for CFL, any competitive bound for ALG gives an identical competitive bound for MAL with parameters corresponding to the CFL instance constructed in Lemma 2.2.

The proofs of both are deferred to Appendix B.3. At a high-level, Proposition 2.3 shows that if ALG is η -competitive against OPT which pays no switching cost, Lemma 2.2 implies it is also η -competitive on MAL. In the next section, our proposed algorithms will be presented using CFL notation, but these results provide the necessary condition which allows them to solve MAL as well.

3. Designing Competitive Algorithms

In this section, we present our robust algorithm design. We start by discussing some inherent challenges in the problem, highlighting reasons why existing algorithms fail. Next, we introduce a generalization of techniques from online search called pseudo-cost minimization, which underpins our competitive algorithm, ALG1 (Algorithm 1). Finally, we state (and prove in Appendix B) two bounds, which jointly imply that ALG1 achieves the optimal competitive ratio for CFL and MAL.

Challenges. Canonical algorithms for CFC (Chen et al., 2018; Sellke, 2020; Zhang et al., 2021) make decisions that attempt to minimize (or nearly minimize) the hitting cost of cost functions $f_t(\cdot)$ and switching cost across all time steps. As discussed in the introduction, the structure of the problem with a long-term constraint means that such

myopic cost-minimization algorithms will fail in general. To illustrate this, consider the actions of a minimizer-driven algorithm on an arbitrary sequence with length T . For each $t < T$, the algorithm chooses a point at or near $\mathbf{0}$, since $\mathbf{0}$ is the minimizer of each f_t . However, since $c(\mathbf{0}) = 0$, such an algorithm must subsequently satisfy all or almost all of the long-term constraint during the *compulsory trade*, incurring an arbitrarily bad hitting cost.

This challenge motivates an algorithm design that balances between the two extremes of finishing the long-term constraint “immediately” (i.e., at early time steps), and finishing the long-term constraint “when forced to” (i.e., during the compulsory trade). Both extremes result in a poor competitive ratio. Algorithms in the online search literature (e.g., online knapsack, OWT) leverage a *threshold-based design* to address precisely this problem, as in (Zhou et al., 2008; Sun et al., 2021b; Lechowicz et al., 2024). However, such threshold-based algorithms are traditionally derived for unidimensional decision spaces with no switching costs.

Generalizing beyond unidimensional decision spaces proves to be non-trivial – consider the following example. Towards the application of carbon-aware load shifting, it may reasonable to expect that cost functions are separable over dimensions. It is thus reasonable to consider whether an existing unidimensional algorithm (i.e., as shown by Lechowicz et al. (2024)) can solve the problem via *decomposition*, e.g., by running d instances of the unidimensional algorithm. However, such a technique fails from the perspective of competitive analysis, because the necessary coupling between d independently determined unidimensional decisions and the multidimensional long-term constraint is broken. A multidimensional switching cost complicates such a decomposition even further, as decisions in one dimension increase cost in other dimensions, potentially beyond the competitive threshold used to determine a decision in the first place.

In what follows, we describe a *pseudo-cost minimization approach*, which generalizes the threshold-based design to operate in the CFL setting. A key enabling result for this approach is the ability to simultaneously consider all dimensions for both the hitting costs and the long-term constraint, leveraging the definition of the constraint function $c(\mathbf{x})$.

Algorithm description. Recall that $z^{(t)}$ gives the fraction of the long-term constraint satisfied at time t . Building off of the intuition of threshold-based design, we define a function ϕ , which will be used to compute a *pseudo-cost minimization* problem central to our robust algorithm.

Definition 3.1 (Pseudo-cost threshold function ϕ for CFL). *For any utilization $z \in [0, 1]$, ϕ is defined as:*

$$\phi(z) = U - \beta + (U/\alpha - U + 2\beta) \exp(z/\alpha), \quad (5)$$

where α is the competitive ratio and is defined in (6).

Algorithm 1 Pseudo-cost minimization algorithm (ALG1)

input: long-term constraint function $c(\cdot)$, distance metric $\|\cdot\|_{\ell_1(\mathbf{w})}$, pseudo-cost threshold function $\phi(z)$

initialize: $z^{(0)} = 0$;

while cost function $f_t(\cdot)$ is revealed and $z^{(t-1)} < 1$ **do**

 solve pseudo-cost minimization problem:

$$\mathbf{x}_t = \arg \min_{\mathbf{x} \in X: c(\mathbf{x}) \leq 1 - z^{(t-1)}} f_t(\mathbf{x}) + \|\mathbf{x} - \mathbf{x}_{t-1}\|_{\ell_1(\mathbf{w})} - \int_{z^{(t-1)}}^{z^{(t-1)} + c(\mathbf{x}_t)} \phi(u) du \quad (4)$$

 update utilization $z^{(t)} = z^{(t-1)} + c(\mathbf{x}_t)$

end while

Then our algorithm (Algorithm 1, referred to as ALG1) solves the pseudo-cost minimization problem defined in (4) to obtain a decision \mathbf{x}_t at each time step. At a high level, the inclusion of ϕ in this pseudo-cost problem enforces that, upon arrival of a cost function, the algorithm satisfies “just enough” of the long-term constraint. Concretely, the structure of the ϕ function enforces that $\phi(z^{(t)}) - \beta$ corresponds to the “best cost function seen so far”. Then, if a good cost function arrives, the pseudo-cost minimization problem solves for the \mathbf{x}_t which guarantees a competitive ratio of α against the current estimate of OPT.

At a glance, it is not obvious that the minimization problem in (4) is tractable; however, in Appendix B.1, we show that the problem is *convex*, implying that it can be solved efficiently. In Theorem 3.2, we state the competitive result for ALG1. We discuss the significance of the result below, and relegate the full proof to Appendix B.2.

Theorem 3.2. ALG1 is α -competitive for CFL, where α is the solution to $\frac{U-L-2\beta}{U-U/\alpha-2\beta} = \exp(1/\alpha)$, given by

$$\alpha := \left[W \left(\left(\frac{2\beta}{U} + \frac{L}{U} - 1 \right) e^{\frac{2\beta}{U}-1} \right) - \frac{2\beta}{U} + 1 \right]^{-1}, \quad (6)$$

where W is the Lambert W function (Corless et al., 1996).

Intuitively, parameters of CFL (L , U , and β) appear in the competitive bound. While results for OWT and CFC are not directly comparable with CFL results, we discuss a few connections here.

When $\beta \rightarrow 0$, α matches the optimal competitive ratio of $[W((L/U - 1)e^{-1}) + 1]^{-1}$ for the minimization variant of OWT (Lorenz et al., 2008; Sun et al., 2021a). In the intermediate case (i.e., when $\beta \in (0, U-L/2)$), Taylor expanding the expression for α yields a leading order approximation of $O(\sqrt{U/L}) + O(\beta)$ – thus, CFL adds a new *linear* dependence on β compared to OWT. Furthermore, as $\beta \rightarrow U-L/2$, α approaches U/L , which is the competitive ratio achievable by e.g., a myopic cost minimization algorithm. In the Appendix, in Fig. 12, we plot α as a function of these dependencies (i.e., U/L , β).

Since α does not feature a dependence on the dimension d of the vector space, we note a connection with CFC: it is known that “dimension-free” bounds are achievable in CFC with necessary structural assumptions on the hitting cost that are evocative of our bounded gradient assumptions (Argue et al., 2020). In particular, Chen et al. (2018) give an algorithm for CFC that achieves a dimension-independent competitive ratio on ℓ_2 metrics when hitting cost functions increase away from the minimizer at a certain rate. They extend this to a dimension-dependent bound for general metrics using norm equivalency results.

Via Proposition 2.3, we obtain an immediate corollary to Theorem 3.2 which gives the following competitive bound when ALG1 is used to solve MAL. The full proof of Corollary 3.3 can be found in Appendix B.3.

Corollary 3.3. ALG1 is α -competitive for MAL.

On the tightness of competitive ratios. It is important to highlight that the bounds in Theorem 3.2 and Corollary 3.3 are the first competitive bounds for any variant of convex function chasing or online metric allocation imbued with long-term constraints. A natural follow-up question concerns whether any online algorithm for CFL (or MAL) can achieve a better competitive bound. In the following, we answer this question in the negative, showing that ALG1’s competitive ratio is the best that any deterministic online algorithm for CFL and/or MAL can achieve. We state the result here, and defer the full proof to Appendix B.4.

Theorem 3.4. For any L, U , and $\beta \in [0, U-L/2)$, there exists a family of CFL instances such that any deterministic online algorithm for CFL is at least α -competitive, where α is as defined in (6).

Since ALG1 is α -competitive by Theorem 3.2, this implies that ALG1 achieves the optimal competitive ratio for CFL. Furthermore, by leveraging Lemma 2.2, this result gives an immediate corollary result in the MAL setting by constructing a corresponding family of MAL instances, which forces any algorithm to achieve a competitive ratio of α . We state the result here, deferring the full proof to Appendix B.5.

Corollary 3.5. The CFL instances in Theorem 3.4 correspond to instances of MAL such that any deterministic online algorithm for MAL is at least α -competitive.

As previously, since ALG1 is α -competitive by Corollary 3.3, it achieves the optimal competitive ratio for MAL. We note that beyond the settings of CFL and MAL considered in this paper, Theorem 3.4 and Corollary 3.5 are the first lower bound results for convex function chasing and online metric allocation with long-term constraints, and may thus give useful insight into the achievable competitive bounds for different or more general settings of these problems.

4. Learning-augmented Algorithms

In this section, we consider how *untrusted black-box advice* can help improve the average-case performance of a learning-augmented algorithm for CFL and MAL while retaining worst-case guarantees. We first consider a sub-optimal “baseline” algorithm that directly combines advice with a robust algorithm such as ALG1. We then propose a unified algorithm called CLIP, which integrates advice more efficiently and achieves the optimal trade-off between *consistency* and *robustness* (Definition 2.1).

Advice model. For a CFL or MAL instance $\mathcal{I} \in \Omega$, let ADV denote untrusted black-box decision advice, i.e., $\text{ADV} := \{\mathbf{a}_t \in X : t \in [T]\}$. If the advice is correct, it achieves the optimal objective value (i.e., $\text{ADV}(\mathcal{I}) = \text{OPT}(\mathcal{I})$).

A simple baseline. Lechowicz et al. (2024) show that a straightforward “fixed-ratio” learning-augmented approach works well in practice for unidimensional online search with switching costs. Here we show that a similar technique (playing a convex combination of the solutions chosen by the advice and a robust algorithm) achieves bounded but sub-optimal consistency and robustness for CFL.

Let $\text{ROB} := \{\tilde{\mathbf{x}}_t : t \in [T]\}$ denote the actions of a robust algorithm for CFL (e.g., ALG1). For any value $\epsilon \in (0, \alpha - 1]$, the fixed-ratio algorithm (denoted as **Baseline** for brevity) sets a combination factor $\lambda := \frac{\alpha-1-\epsilon}{\alpha-1}$. Then at each time step, **Baseline** makes a decision according to $\mathbf{x}_t = \lambda \mathbf{a}_t + (1 - \lambda) \tilde{\mathbf{x}}_t$. We present consistency and robustness results for **Baseline** below, deferring the full proof to Appendix C.1.

Lemma 4.1. Letting ROB denote the actions of ALG1 and setting a parameter $\epsilon \in (0, \alpha - 1]$, **Baseline** is $(1 + \epsilon)$ -consistent and $\left(\frac{(U+2\beta)/L(\alpha-1-\epsilon)+\alpha\epsilon}{(\alpha-1)}\right)$ -robust for CFL.

Although this fixed-ratio algorithm verifies that an algorithm for CFL can utilize untrusted advice to improve performance, it remains an open question of whether the trade-off between consistency and robustness given in Lemma 4.1 is optimal. Thus, we study whether a learning-augmented algorithm for CFL can be designed which *does* achieve the provably optimal consistency-robustness trade-off. In the next section, we start by considering a more sophisticated method of incorporating advice into an algorithm design.

An optimal learning-augmented algorithm. We present CLIP (Consistency-Limited Pseudo-cost minimization, Algorithm 2) which achieves the optimal trade-off between consistency and robustness for CFL.

To start, for any $\epsilon \in (0, \alpha - 1]$, we define a corresponding target robustness factor γ^ϵ , the unique positive solution to:

$$\gamma^\epsilon = \epsilon + \frac{U}{L} - \frac{\gamma^\epsilon}{L}(U - L) \ln \left(\frac{U - L - 2\beta}{U - U/\gamma^\epsilon - 2\beta} \right). \quad (10)$$

Algorithm 2 Consistency Limited Pseudo-cost minimization (CLIP)

input: consistency parameter ϵ , long-term constraint function $c(\cdot)$, pseudo-cost threshold function $\phi^\epsilon(\cdot)$

initialize: $z^{(0)} = 0$; $p^{(0)} = 0$; $A^{(0)} = 0$; $\text{CLIP}_0 = 0$; $\text{ADV}_0 = 0$

while cost function $f_t(\cdot)$ is revealed, untrusted advice \mathbf{a}_t is revealed, and $z^{(t-1)} < 1$ **do**

 update advice cost $\text{ADV}_t = \text{ADV}_{t-1} + f_t(\mathbf{a}_t) + \|\mathbf{a}_t - \mathbf{a}_{t-1}\|_{\ell_1(\mathbf{w})}$ and advice utilization $A^{(t)} = A^{(t-1)} + c(\mathbf{a}_t)$

 solve *constrained* pseudo-cost minimization problem:

$$\mathbf{x}_t = \arg \min_{\mathbf{x} \in X: c(\mathbf{x}) \leq 1 - z^{(t-1)}} f_t(\mathbf{x}) + \|\mathbf{x} - \mathbf{x}_{t-1}\|_{\ell_1(\mathbf{w})} - \int_{p^{(t-1)}}^{p^{(t-1)} + c(\mathbf{x})} \phi^\epsilon(u) du \quad (7)$$

 such that

$$\begin{aligned} \text{CLIP}_{t-1} + f_t(\mathbf{x}) + \|\mathbf{x} - \mathbf{x}_{t-1}\|_{\ell_1(\mathbf{w})} + \|\mathbf{x} - \mathbf{a}_t\|_{\ell_1(\mathbf{w})} + \|\mathbf{a}_t\|_{\ell_1(\mathbf{w})} + (1 - z^{(t-1)} - c(\mathbf{x}))L + \max((A^{(t)} - z^{(t-1)} - c(\mathbf{x})), 0)(U - L) \\ \leq (1 + \epsilon)[\text{ADV}_t + \|\mathbf{a}_t\|_{\ell_1(\mathbf{w})} + (1 - A^{(t)})L] \end{aligned} \quad (8)$$

 update cost $\text{CLIP}_t = \text{CLIP}_{t-1} + f_t(\mathbf{x}_t) + \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_{\ell_1(\mathbf{w})}$ and utilization $z^{(t)} = z^{(t-1)} + c(\mathbf{x}_t)$

 solve *unconstrained* pseudo-cost minimization problem:

$$\bar{\mathbf{x}}_t = \arg \min_{\mathbf{x} \in X: c(\mathbf{x}) \leq 1 - z^{(t-1)}} f_t(\mathbf{x}) + \|\mathbf{x} - \mathbf{x}_{t-1}\|_{\ell_1(\mathbf{w})} - \int_{p^{(t-1)}}^{p^{(t-1)} + c(\mathbf{x})} \phi^\epsilon(u) du \quad (9)$$

 update pseudo-utilization $p^{(t)} = p^{(t-1)} + \min(c(\bar{\mathbf{x}}_t), c(\mathbf{x}_t))$

end while

Note that $\gamma^{\alpha-1} = \alpha$, and $\gamma^0 = U/L$. We use γ^ϵ to define a pseudo-cost threshold function ϕ^ϵ that will be used in a minimization problem to choose a decision at each step of the CLIP algorithm.

Definition 4.2 (Pseudo-cost threshold function ϕ^ϵ).

Given γ^ϵ from (10), $\phi^\epsilon(p)$ for $p \in [0, 1]$ is defined as:

$$\phi^\epsilon(p) = U - \beta + (U/\gamma^\epsilon - U + 2\beta) \exp(p/\gamma^\epsilon). \quad (11)$$

For each time step $t \in [T]$, we define a *pseudo-utilization* $p^{(t)} \in [0, 1]$, where $p^{(t)} \leq z^{(t)} \forall t$, and $p^{(t)}$ describes the fraction of the long-term constraint which been satisfied “robustly” (as defined by the pseudo-cost) at time t .

Then CLIP (see Algorithm 2) solves a *constrained* pseudo-cost minimization problem (defined in (7)) to obtain a decision \mathbf{x}_t at each time step. The objective of this problem is mostly inherited from ALG1, but the inclusion of a *consistency constraint* allows the framework to accommodate untrusted advice for bounded consistency and robustness.

The high-level intuition behind this consistency constraint (defined in (8)) is to directly compare the solutions of CLIP and ADV *so far*, while “*hedging*” against worst-case scenarios which may cause CLIP to violate the desired $(1 + \epsilon)$ -consistency. We introduce some notation to simplify the expression of the constraint. We let CLIP_t denote the cost of CLIP up to time t , i.e., $\text{CLIP}_t := \sum_{\tau=1}^t f_\tau(\mathbf{x}_\tau) + \|\mathbf{x}_\tau - \mathbf{x}_{\tau-1}\|_{\ell_1(\mathbf{w})}$. Similarly, we let $\text{ADV}_t := \sum_{\tau=1}^t f_\tau(\mathbf{a}_\tau) + \|\mathbf{a}_\tau - \mathbf{a}_{\tau-1}\|_{\ell_1(\mathbf{w})}$ denote the cost of ADV up to time t . Additionally, we let $A^{(t)}$ denote the utilization of ADV at time t , i.e., $A^{(t)} := \sum_{\tau=1}^t c(\mathbf{a}_\tau)$

The constraint defined in (8) considers the cost of both CLIP and ADV so far, and the current hitting and switching cost

$f_t(\mathbf{x}) + \|\mathbf{x} - \mathbf{x}_{t-1}\|_{\ell_1(\mathbf{w})}$, ensuring that $(1 + \epsilon)$ -consistency is preserved. Both sides of the constraint also include terms which consider the cost of potential future situations. First, $\|\mathbf{x} - \mathbf{a}_t\|_{\ell_1(\mathbf{w})} + \|\mathbf{a}_t\|_{\ell_1(\mathbf{w})}$ ensures that if CLIP pays a switching cost to follow ADV and/or pays a switching cost to “switch off” (move to 0) in e.g., the next time step, that cost has been paid for “in advance”. As $\mathbf{x}_{T+1} = 0$, the constraint also charges ADV in advance for the mandatory switching cost at the end of the sequence ($\|\mathbf{a}_t\|_{\ell_1(\mathbf{w})}$); this ensures that there is always a feasible setting of \mathbf{x}_t .

In the term $(1 - A^{(t)})L$, the consistency constraint assumes that ADV can satisfy the rest of the long-term constraint at the best marginal cost L . Respectively, in the term $(1 - z^{(t-1)} - c(\mathbf{x}))L + \max((A^{(t)} - z^{(t-1)} - c(\mathbf{x})), 0)(U - L)$, the constraint assumes CLIP can satisfy *up to* $(1 - A^{(t)})$ of the remaining long-term constraint at the best cost L , but any excess (i.e., $(A^{(t)} - z^{(t)})$) must be satisfied at the worst cost U (e.g., during the compulsory trade). This worst-case assumption ensures that when actual hitting costs replace the above terms, the desired $(1 + \epsilon)$ -consistency holds.

At each step, CLIP also solves an *unconstrained* pseudo-cost minimization problem to obtain $\bar{\mathbf{x}}_t$, which updates the pseudo-utilization $p^{(t)}$. This ensures that when ADV has accepted a cost function which would not be accepted by the unconstrained pseudo-cost minimization, the threshold function ϕ^ϵ can “start from zero” in subsequent time steps.

At a high level, CLIP’s consistency constraint combined with the pseudo-cost minimization generates decisions which are *as robust as possible* while preserving consistency. In Theorem 4.3, we state the consistency and robustness of CLIP; we relegate the full proof to Appendix C.2.

Theorem 4.3. For any $\epsilon \in [0, \alpha - 1]$, CLIP is $(1 + \epsilon)$ -consistent and γ^ϵ -robust for CFL (γ^ϵ as defined in (10)).

The previous result gives an immediate corollary when CLIP is used to solve MAL, which we state below. The full proof of Corollary 4.4 can be found in Appendix C.3.

Corollary 4.4. For any $\epsilon \in [0, \alpha - 1]$, CLIP is $(1 + \epsilon)$ -consistent and γ^ϵ -robust for MAL.

Optimal trade-offs between robustness and consistency.

Although the trade-off given by CLIP implies that achieving 1-consistency requires a large robustness bound of U/L in the worst-case, in the following theorem we show that this is the best we can obtain from any consistent and robust algorithm. We state the result and discuss its significance here, deferring the full proof to Appendix C.4.

Theorem 4.5. Given untrusted advice ADV and $\epsilon \in (0, \alpha - 1]$, any $(1 + \epsilon)$ -consistent learning-augmented algorithm for CFL is at least γ^ϵ -robust, where γ^ϵ is defined in (10).

This result implies that CLIP achieves the *optimal* trade-off between consistency and robustness for CFL. Furthermore, via Lemma 2.2, this result immediately gives Corollary 4.6, which we state here and prove in Appendix C.5.

Corollary 4.6. Any $(1 + \epsilon)$ -consistent learning-augmented algorithm for MAL is at least γ^ϵ -robust (γ^ϵ defined by (10)).

As previously, this implies CLIP achieves the optimal consistency-robustness trade-off for MAL. Beyond the settings of CFL and MAL, these Pareto-optimality results may give useful insight into the achievable consistency-robustness trade-offs for more general settings.

5. Numerical Experiments




Figure 1. CDFs of empirical competitive ratios for various algorithms.




Figure 2. Varying adversarial factor ξ , with $U/L = 250$, $\beta = 50$, $d = 5$, and $\sigma = 50$.

In this section, we conduct numerical experiments on synthetic CFL instances. We evaluate ALG1 and CLIP against the offline optimal solution, three heuristics adapted from related work, and the learning-augmented Baseline.

Setup. We construct a d -dimensional decision space, where d is picked from the set $\{5, 7, \dots, 21\}$. The competitive ratio of our proposed algorithms depends on both

U/L and $\beta = \max_i \mathbf{w}^i$, as the switching cost. Hence, we evaluate their performance over the range of these parameters. We set different cost fluctuation ratios $U/L \in \{50, 150, \dots, 1250\}$ by setting L and U accordingly, and β is picked from the set $\beta \in \{0, 5, \dots, U/2.5\}$. For each experiment, $c(\mathbf{x}) = \|\mathbf{x}\|_1$.

For a given setting of d , U/L , and β , we generate 1,000 random instances as follows. First, each term of the weight vector \mathbf{w} for the weighted ℓ_1 norm is drawn randomly from the uniform distribution on $[0, \beta]$. Next, the time horizon T is generated randomly from a uniform distribution on $[6, 24]$. For each time $t \in [T]$, a cost function is generated as follows: Let $f_t(\mathbf{x}) = \mathbf{f}_t^\top \mathbf{x}$, where \mathbf{f}_t is a d -dimensional *cost vector*. To generate \mathbf{f}_t , we first draw μ_t from the uniform distribution on $[L, U]$, and then draw each term of \mathbf{f}_t from a normal distribution centered at μ_t with standard deviation σ (i.e., $\mathbf{f}_t^i \sim \mathcal{N}(\mu_t, \sigma)$). Any terms which are outside the assumed interval $[L, U]$ (i.e. $\mathbf{f}_t^i < L$ or $\mathbf{f}_t^i > U$) are truncated appropriately. For each instance, we report the empirical competitive ratios as the evaluation metric, comparing the tested algorithms against an offline optimal benchmark.

In the setting with advice, we obtain simulated advice as follows: Let $\xi \in [0, 1]$ denote an *adversarial factor*. When $\xi = 0$, ADV gives the optimal solution, and when $\xi = 1$, ADV is fully adversarial. Formally, letting $\{\mathbf{x}_t^* : t \in [T]\}$ denote the decisions made by an optimal solution, and letting $\{\check{\mathbf{x}}_t : t \in [T]\}$ represent the decisions made by a solution which maximizes the objective (rather than minimizing it), we have that $\text{ADV} = \{(1 - \xi)\mathbf{x}_t^* + \xi\check{\mathbf{x}}_t : t \in [T]\}$. We note that although $\{\check{\mathbf{x}}_t : t \in [T]\}$ is adversarial from the perspective of the objective, it is still a feasible solution for the problem (i.e., it satisfies the long-term constraint).

Comparison algorithms. We use CVXPY (Diamond & Boyd, 2016) to compute the offline optimal solution for each instance using a convex optimization solver with access to all cost functions in advance. Since we are the first to study CFL, there are no directly comparable algorithms with competitive guarantees – thus, we consider three heuristic techniques based on literature for adjacent problems.

The first technique is termed “**agnostic**”, which chooses the minimum dimension of the cost function in the first time step $t = 1$ (i.e., $k = \arg \min_{i \in [d]} \mathbf{c}_1^i$), sets $\mathbf{x}_1^k = 1$, and $\mathbf{x}_t = \mathbf{0} \forall t > 1$. The second technique is termed “**move to minimizer**”, which takes inspiration from algorithms for CFC (Zhang et al., 2021) and satisfies $1/T$ fraction of the long-term constraint at each time step by moving to the minimum dimension of each cost function. Formally, at each time step t , letting $k_t = \arg \min_{i \in [d]} \mathbf{c}_t^i$, “move to minimizer” sets $\mathbf{x}_t^{k_t} = 1/T$. Finally, the third technique is termed “**simple threshold**”, which takes inspiration from algorithms for online search (El-Yaniv et al., 2001). This algorithm




Figure 3. Varying U/L , with $\beta = U/5$, $d = 5$, $\xi = 0$, and $\sigma = U/5$. Figure 4. Varying β , with $U/L = 250$, $d = 5$, $\xi = 0$, and $\sigma = 50$.

Figure 5. Varying d with $\beta = 50$, $U/L = 250$, $\sigma = 50$, and $\xi = 0$. Figure 6. Varying σ , with $\beta = 50$, $U/L = 250$, $d = 5$, and $\xi = 0$.

sets a fixed threshold $\psi = \sqrt{UL}$, and completes the long-term constraint at the first time step and dimension where the hitting cost does not exceed ψ . Formally, at the first time step τ satisfying $\exists k \in [d] : \mathbf{f}_\tau^k \leq \psi$, “simple threshold” sets $\mathbf{x}_\tau^k = 1$. In the setting with advice, we compare our proposed CLIP learning-augmented algorithm against the Baseline learning-augmented algorithm described in Section 4 (e.g., Lemma 4.1).

Experimental results. Fig. 1 summarizes the main results for ALG1, the comparison heuristic algorithms, and one setting of CLIP ($\epsilon = 2$) in a CDF plot of the empirical competitive ratios across several experiments. In these experiments, we fix $U/L = 250$, $\xi = 0$, $\sigma = 50$, while varying β and d . Notably, ALG1 outperforms in both average-case and worst-case performance, improving on the closest “simple threshold” by an average of 18.2%, and outperforming “agnostic” and “move to minimizer” by averages of 56.1% and 71.5%, respectively. With correct advice, CLIP sees significant performance gains across the board.

In Fig. 3-6, we investigate the impact of parameters on the average empirical competitive ratio for each algorithm. In Appendix A.1, we give corresponding plots for the 95th percentile (“worst-case”) results. Fig. 3 plots competitive ratios for different values of U/L . We fix $\beta = U/5$, $d = 5$, $\xi = 0$, $\sigma = U/5$, while varying U/L . Since there is a dependence on U/L in our competitive results, the performance of ALG1 degrades as U/L grows, albeit at a favorable pace compared to the heuristics. Fig. 4 plots competitive ratios for different values of β . We fix $U/L = 250$, $d = 5$, $\xi = 0$, $\sigma = 50$. As β grows, the “agnostic” and “move to minimizer” heuristics improve because the switching cost paid by OPT grows.

In Fig. 6, we plot competitive ratios for different values of σ . We fix $U/L = 250$, $\beta = 50$, $d = 5$, $\xi = 0$, while varying σ . As cost functions become more variable, the performance of all algorithms degrades, with the exception of CLIP. There is a plateau as σ grows, because a large σ implies that more terms in each \mathbf{f}_t must be truncated to the interval $[L, U]$. Finally, Fig. 5 plots competitive ratios for different values of d . We fix $U/L = 250$, $\beta = 50$, $\xi = 0$, $\sigma = 50$, while varying d . As d grows, ALG1 and CLIP’s performance

degrades slower compared to the heuristics, as predicted by their *dimension-free* theoretical bounds.⁵

Fig. 2 plots the effect of prediction error on the learning-augmented algorithms CLIP and Baseline. We test several values of $\xi \in [0, 1/2]$ (recall that $\xi = 0$ recovers correct advice), while fixing $U/L = 250$, $\beta = 50$, $d = 5$, and $\sigma = 50$. We also test Baseline and CLIP for several values of $\epsilon \in \{2, 5, 10\}$ (note that ADV corresponds to Baseline and CLIP with $\epsilon = 0$). Notably, we find that CLIP significantly outperforms the Baseline algorithm as ξ grows, showing an average improvement of 60.8% when $\xi > 0.1$. This result implies that CLIP is more empirically robust to prediction errors than the simple fixed ratio technique of Baseline.

6. Conclusion

We study online metric problems with long-term constraints, motivated by emerging problems in sustainability. These are the first such problems to concurrently incorporate multidimensional decision spaces, switching costs, and long-term demand constraints. Our main results instantiate the CFL and MAL problems towards a motivating application. We design competitive and learning-augmented algorithms, show that their performance bounds are tight, and validate them in numerical experiments. Several interesting open questions are prompted by our work. Specifically, (i) what is achievable in non- ℓ_1 vector spaces e.g., the Euclidean setting, and (ii) can our results for MAL inform algorithm designs for e.g., tree metrics, and by extension, arbitrary metric spaces?

⁵While the plot suggests a correlation between d and ALG1’s performance, this seems to be a side effect of the random cost functions, which indirectly introduce more dimension-wise variability as d increases. There is a large gap between empirical performance and the theoretical upper bounds – the asymptotic performance of ALG1 would be constant for large d .

Impact Statement

This paper presents work whose goal is to advance the field of online optimization and study problems relevant to sustainability applications. There are many potential societal consequences of our work, none of which we feel must be explicitly highlighted here.

Acknowledgements

This research is supported by National Science Foundation grants CAREER-2045641, CNS-2102963, CNS-2106299, CNS-2146814, CNS-1518941, CNS-2106403, CPS-2136197, CPS-2136199, NGSDI-2105494, NGSDI-2105648, 1908298, 2020888, 2021693, 2045641, 2213636, 2211888, and an NSF Graduate Research Fellowship (DGE-2139433).

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Department of Energy Computational Science Graduate Fellowship under Award Number DE-SC0024386.

Disclaimers

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

References

Acun, B., Lee, B., Kazhamiaka, F., Maeng, K., Gupta, U., Chakkaravarthy, M., Brooks, D., and Wu, C.-J. Carbon Explorer: A Holistic Framework for Designing Carbon Aware Datacenters. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2*, ASPLOS 2023, pp. 118–132, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9781450399166. doi: 10.1145/3575693.3575754. URL <https://doi.org/10.1145/3575693.3575754>.

Argue, C. J., Gupta, A., and Guruganesh, G. Dimension-Free Bounds for Chasing Convex Functions. In *Proceedings of Thirty Third Conference on Learning Theory*, pp. 219–241. PMLR, July 2020.

Bansal, N. and Coester, C. Online Metric Allocation and Time-Varying Regularization. In Chechik, S., Navarro, G., Rotenberg, E., and Herman, G. (eds.), *30th Annual European Symposium on Algorithms (ESA 2022)*, volume 244 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pp. 13:1–13:13, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. ISBN 978-3-95977-247-1. doi: 10.4230/LIPIcs.ESA.2022.13. URL <https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.ESA.2022.13>.

Bashir, N., Guo, T., Hajiesmaili, M., Irwin, D., Shenoy, P., Sitaraman, R., Souza, A., and Wierman, A. Enabling Sustainable Clouds: The Case for Virtualizing the Energy System. In *Proceedings of the ACM Symposium on Cloud Computing*, SoCC ’21, pp. 350–358, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450386388. doi: 10.1145/3472883.3487009. URL <https://doi.org/10.1145/3472883.3487009>.

Borodin, A., Linial, N., and Saks, M. E. An Optimal On-Line Algorithm for Metrical Task System. *J. ACM*, 39(4):745–763, Oct 1992. ISSN 0004-5411. doi: 10.1145/146585.146588. URL <https://doi.org/10.1145/146585.146588>.

Bubeck, S., Klartag, B., Lee, Y. T., Li, Y., and Sellke, M. Chasing Nested Convex Bodies Nearly Optimally. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, Proceedings, pp. 1496–1508. Society for Industrial and Applied Mathematics, December 2019. doi: 10.1137/1.9781611975994.91.

Bubeck, S., Cohen, M. B., Lee, J. R., and Lee, Y. T. Metrical Task Systems on Trees via Mirror Descent and Unfair Gluing. *SIAM Journal on Computing*, 50(3):909–923, January 2021. ISSN 0097-5397, 1095-7111. doi: 10.1137/19M1237879.

Bubeck, S., Coester, C., and Rabani, Y. The Randomized \$k\$-\$\\$-\$Server Conjecture Is False! In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing (STOC 2023)*, STOC 2023, pp. 581–594, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9781450399135. doi: 10.1145/3564246.3585132. URL <https://doi.org/10.1145/3564246.3585132>.

Chen, N., Goel, G., and Wierman, A. Smoothed Online Convex Optimization in High Dimensions via Online

Balanced Descent. In *Proceedings of the 31st Conference On Learning Theory*, pp. 1574–1594. PMLR, July 2018.

Cheng, K.-W., Bian, Y., Shi, Y., and Chen, Y. Carbon-Aware EV Charging. In *2022 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm)*, pp. 186–192, 2022. doi: 10.1109/SmartGridComm52983.2022.9960988.

Christianson, N., Handina, T., and Wierman, A. Chasing Convex Bodies and Functions with Black-Box Advice. In *Proceedings of the 35th Conference on Learning Theory*, volume 178, pp. 867–908. PMLR, 02–05 Jul 2022.

Christianson, N., Shen, J., and Wierman, A. Optimal robustness-consistency tradeoffs for learning-augmented metrical task systems. In *International Conference on Artificial Intelligence and Statistics*, 2023.

Corless, R. M., Gonnet, G. H., Hare, D. E., Jeffrey, D. J., and Knuth, D. E. On the Lambert W function. *Advances in Computational mathematics*, 5:329–359, 1996.

Diamond, S. and Boyd, S. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83):1–5, 2016.

El-Yaniv, R., Fiat, A., Karp, R. M., and Turpin, G. Optimal Search and One-Way Trading Online Algorithms. *Algorithmica*, 30(1):101–139, May 2001. doi: 10.1007/s00453-001-0003-0. URL <https://doi.org/10.1007/s00453-001-0003-0>.

Friedman, J. and Linial, N. On convex body chasing. *Discrete & Computational Geometry*, 9(3):293–321, March 1993. doi: 10.1007/bf02189324. URL <https://doi.org/10.1007/bf02189324>.

Hanafy, W. A., Liang, Q., Bashir, N., Irwin, D., and Shenoy, P. CarbonScaler: Leveraging Cloud Workload Elasticity for Optimizing Carbon-Efficiency. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 7(3), Dec 2023.

Johnson, S. G. Notes on the equivalence of norms. <https://math.mit.edu/~stevenj/18.335/norm-equivalence.pdf>, 2020.

Koutsoupias, E. The k-server problem. *Computer Science Review*, 3(2):105–118, May 2009. doi: 10.1016/j.cosrev.2009.04.002. URL <https://doi.org/10.1016/j.cosrev.2009.04.002>.

Kumar, R., Purohit, M., and Svitkina, Z. Improving Online Algorithms via ML Predictions. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.

Lechowicz, A., Christianson, N., Zuo, J., Bashir, N., Hajiesmaili, M., Wierman, A., and Shenoy, P. The Online Pause and Resume Problem: Optimal Algorithms and An Application to Carbon-Aware Load Shifting. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 7(3), Dec 2023.

Lechowicz, A., Christianson, N., Sun, B., Bashir, N., Hajiesmaili, M., Wierman, A., and Shenoy, P. Online Conversion with Switching Costs: Robust and Learning-augmented Algorithms, 2024. URL <https://arxiv.org/abs/2310.20598>.

Lee, R., Sun, B., Hajiesmaili, M., and Lui, J. C. S. Online Search with Predictions: Pareto-optimal Algorithm and its Applications in Energy Markets. In *Proceedings of the 15th ACM International Conference on Future Energy Systems*, e-Energy '24, New York, NY, USA, June 2024. Association for Computing Machinery.

Lorenz, J., Panagiotou, K., and Steger, A. Optimal Algorithms for k-Search with Application in Option Pricing. *Algorithmica*, 55(2):311–328, August 2008. doi: 10.1007/s00453-008-9217-8. URL <https://doi.org/10.1007/s00453-008-9217-8>.

Lykouris, T. and Vassilvitskii, S. Competitive Caching with Machine Learned Advice. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 3296–3305. PMLR, 10–15 Jul 2018. URL <https://proceedings.mlr.press/v80/lykouris18a.html>.

Manasse, M., McGeoch, L., and Sleator, D. Competitive Algorithms for On-Line Problems. In *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing*, STOC '88, pp. 322–333, New York, NY, USA, 1988. Association for Computing Machinery. ISBN 0897912640. doi: 10.1145/62212.62243.

Mitrinovic, D. S., Pečarić, J. E., and Fink, A. M. *Inequalities Involving Functions and Their Integrals and Derivatives*, volume 53. Springer Science & Business Media, 1991.

Mohr, E., Ahmad, I., and Schmidt, G. Online algorithms for conversion problems: A survey. *Surveys in Operations Research and Management Science*, 19(2):87–104, July 2014. doi: 10.1016/j.sorms.2014.08.001. URL <https://doi.org/10.1016/j.sorms.2014.08.001>.

Radovanovic, A., Koningstein, R., Schneider, I., Chen, B., Duarte, A., Roy, B., Xiao, D., Haridasan, M., Hung, P., Care, N., et al. Carbon-Aware Computing for Datacenters. *IEEE Transactions on Power Systems*, 2022.

Sellke, M. Chasing Convex Bodies Optimally. In *Proceedings of the Thirty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '20, pp. 1509–1518, USA, January 2020. Society for Industrial and Applied Mathematics.

Sun, B., Lee, R., Hajiesmaili, M., Wierman, A., and Tsang, D. Pareto-Optimal Learning-Augmented Algorithms for Online Conversion Problems. In Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., and Vaughan, J. W. (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 10339–10350. Curran Associates, Inc., 2021a.

Sun, B., Zeynali, A., Li, T., Hajiesmaili, M., Wierman, A., and Tsang, D. H. Competitive Algorithms for the Online Multiple Knapsack Problem with Application to Electric Vehicle Charging. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 4(3), June 2021b. doi: 10.1145/3428336. URL <https://doi.org/10.1145/3428336>.

Wei, A. and Zhang, F. Optimal robustness-consistency trade-offs for learning-augmented online algorithms. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NeurIPS '20, Red Hook, NY, USA, 2020. Curran Associates Inc. ISBN 9781713829546.

Wiesner, P., Behnke, I., Scheinert, D., Gontarska, K., and Thamsen, L. Let's Wait AWhile: How Temporal Workload Shifting Can Reduce Carbon Emissions in the Cloud. In *Proceedings of the 22nd International Middleware Conference*, pp. 260–272, New York, NY, USA, 2021. Association for Computing Machinery. doi: 10.1145/3464298.3493399.

Zhang, L., Jiang, W., Lu, S., and Yang, T. Revisiting Smoothed Online Learning, 2021. URL <https://arxiv.org/abs/2102.06933>.

Zhou, Y., Chakrabarty, D., and Lukose, R. Budget Constrained Bidding in Keyword Auctions and Online Knapsack Problems. In *Lecture Notes in Computer Science*, pp. 566–576. Springer Berlin Heidelberg, 2008.

Appendix

A. Numerical Experiments (continued)

In this section, we give supplemental results examining the 95th percentile (“worst-case”) empirical competitive ratio results, following the same general structure as in the main body.

A.1. Supplemental Results

To complement the results for the average empirical competitive ratio shown in Fig. 5, in this section we plot the 95th percentile empirical competitive ratios for each tested algorithm, which primarily serve to show that the improved performance of our proposed algorithm holds in both average-case and tail (“worst-case”) scenarios.

In Fig. 8-11, we investigate the impact of different parameters on the performance of each algorithm. In Fig. 8, we plot 95th percentile empirical competitiveness for different values of U/L – in this experiment, we fix $\beta = U/5$, $d = 5$, $\xi = 0$, and $\sigma = U/5$, while varying $U/L \in \{50, \dots, 1250\}$. As observed in the average competitive ratio plot (Fig. 3), the performance of ALG1 degrades as U/L grows, albeit at a favorable pace compared to the comparison algorithms. Fig. 9 plots the 95th percentile empirical competitiveness for different values of β – in this experiment, we fix $U/L = 250$, $d = 5$, $\xi = 0$, and $\sigma = 50$. As previously in the average competitive results (Fig. 4), “agnostic” and “move to minimizer” heuristics perform better when β grows, because the switching cost paid by the optimal solution grows as well.

In Fig. 10, we plot the 95th percentile empirical competitiveness for different values of d – in this experiment, we fix $U/L = 250$, $\beta = 50$, $\xi = 0$, and $\sigma = 50$, while varying d . Mirroring the previous results (Fig. 5), ALG1 and CLIP’s competitive performance degrades slower as d grows compared to the comparison heuristics, as predicted by their *dimension-free* theoretical bounds. Finally, Fig. 11 plots the 95th percentile empirical competitiveness for different values of σ , which is the dimension-wise variability of each cost function. Here we fix $U/L = 250$, $\beta = 50$, $d = 5$, and $\xi = 0$, while varying $\sigma \in \{0, \dots, U/2\}$. Intuitively, as cost functions become more variable, the competitive ratios of all tested algorithms degrade, with the exception of our learning-augmented algorithm CLIP. This degradation plateaus as σ grows, as a large standard deviation forces more of the terms of each cost vector \mathbf{c}_t to be truncated to the interval $[L, U]$.

In Fig. 7, we plot the 95th percentile empirical competitive ratio companion to Fig. 2, which measures the effect of prediction error on the learning-augmented algorithms CLIP and Baseline. We test several values of $\xi \in [0, 1]$, the adversarial factor (recall that $\xi = 0$ implies the advice is correct), while fixing $U/L = 250$, $\beta = 50$, $d = 5$, $\sigma = 50$. We test Baseline and CLIP for several values of $\epsilon \in \{2, 5, 10\}$ (note that ADV corresponds to running either Baseline or CLIP with $\epsilon = 0$). Notably, in these 95th percentile “worst-case” results, we find that CLIP continues to significantly outperforms the Baseline algorithm as ξ grows, further validating that CLIP is more empirically robust to prediction errors than the simple fixed ratio technique of Baseline.




Figure 8. Varying U/L , with $\beta = U/5$, $d = 5$, $\xi = 0$, and $\sigma = U/5$. Figure 9. Varying β , with $U/L = 250$, $d = 5$, $\xi = 0$, and $\sigma = 50$.

Figure 10. Varying d with $\beta = 50$, $U/L = 250$, $\sigma = 50$, and $\xi = 0$. Figure 11. Varying σ , with $\beta = 50$, $U/L = 250$, $d = 5$, and $\xi = 0$.




Figure 7. Varying adversarial factor ξ , with $U/L = 250$, $\beta = 50$, $d = 5$, and $\sigma = 50$.




Figure 12. Plotting α , the competitive upper bound of ALG1 (see Equation 6), as a function of the CFL problem's parameters (i.e., $U/L, \beta$).

B. Proofs for Section 3 (Competitive Algorithms)

B.1. Convexity of the pseudo-cost minimization problem in ALG1

In this section, we show that the pseudo-cost minimization problem central to the design of ALG1 is a convex minimization problem, implying that it can be solved efficiently.

Define $h_t(\mathbf{x}) : t \in [T]$ to represent the pseudo-cost minimization problem for a single arbitrary time step:

$$h_t(\cdot) = f_t(\mathbf{x}) + d(\mathbf{x}, \mathbf{x}_{t-1}) - \int_{z^{(t-1)}}^{z^{(t-1)} + c(\mathbf{x})} \phi(u) du. \quad (12)$$

Theorem B.1. *Under the assumptions of the CFL and MAL problem settings, $h_t(\cdot)$ is always convex.*

Proof. We prove the above statement by contradiction.

By definition, we know that the sum of two convex functions gives a convex function. Since we have that $d(\mathbf{x}, \mathbf{x}')$ is defined as some norm, by definition and by observing that \mathbf{x}' is fixed, $d(\mathbf{x}, \mathbf{x}')$ is convex. We have also assumed as part of the problem setting that each $f_t(\mathbf{x})$ is convex. Thus, $f_t(\mathbf{x}) + d(\mathbf{x}, \mathbf{x}')$ must be convex.

We turn our attention to the term $-\int_{z^{(t-1)}}^{z^{(t-1)} + c(\mathbf{x})} \phi(u) du$. Let $k(c(\mathbf{x})) = \int_{z^{(t-1)}}^{z^{(t-1)} + c(\mathbf{x})} \phi(u) du$. By the fundamental theorem of calculus, $\nabla k(c(\mathbf{x})) = \phi(z^{(t-1)} + c(\mathbf{x})) \nabla c(\mathbf{x})$

Let $g(c(\mathbf{x})) = \phi(z^{(t-1)} + c(\mathbf{x}))$. Then $\nabla^2 k(c(\mathbf{x})) = \nabla^2 c(\mathbf{x}) k(c(\mathbf{x})) + \nabla c(\mathbf{x}) g'(c(\mathbf{x})) \nabla c(\mathbf{x})^\top$. Since $c(\mathbf{x})$ is piecewise linear (CFL and MAL both assume it is linear), we know that $\nabla^2 c(\mathbf{x}) g(c(\mathbf{x})) = 0$. Since ϕ is monotonically decreasing on the interval $[0, 1]$, we know that $g'(c(\mathbf{x})) < 0$, and thus $\nabla c(\mathbf{x}) g'(c(\mathbf{x})) \nabla c(\mathbf{x})^\top$ is negative semidefinite. This implies that $k(c(\mathbf{x}))$ is concave in \mathbf{x} .

Since the negation of a concave function is convex, this causes a contradiction, because the sum of two convex functions gives a convex function.

Thus, $h_t(\cdot) = f_t(\mathbf{x}) + d(\mathbf{x}, \mathbf{x}_{t-1}) - \int_{z^{(t-1)}}^{z^{(t-1)} + c(\mathbf{x})} \phi(u) du$ is always convex under the assumptions of CFL and MAL. \square

By showing that $h_t(\cdot)$ is convex, it follows that the pseudo-cost minimization (4) in ALG1 is a convex minimization problem (i.e., it can be solved efficiently using numerical methods).

B.2. Proof of Theorem 3.2

In this section, we prove Theorem 3.2, which shows that α as given by (6) is an upper bound on the worst-case competitive ratio of ALG1 (given by Algorithm 1) for the CFL problem.

Proof of Theorem 3.2. Let $z^{(j)} = \sum_{t \in [T]} c(\mathbf{x}_t)$ denote the fraction of the long-term constraint satisfied by ALG1 before the compulsory trade on an arbitrary CFL instance $\mathcal{I} \in \Omega$. Also note that $z^{(t)} = \sum_{m \in [t]} c(\mathbf{x}_m)$ is non-decreasing over n .

Lemma B.2. *The offline optimal solution $\text{OPT}(\mathcal{I})$ for any CFL instance $\mathcal{I} \in \Omega$ is lower bounded by $\phi(z^{(j)}) - \beta$.*

Proof of Lemma B.2. We prove this lemma by contradiction. Note that the offline optimum will stay at $\mathbf{0}$ whenever possible, and satisfy the long-term constraint using the cost functions with the minimum gradient (i.e., the best marginal cost). Assume that $\text{OPT}(\mathcal{I}) < \phi(z^{(j)}) - \beta$, and that $z^{(j)} < 1$ (implying that $\text{OPT}(\mathcal{I}) > L$).

Recall that any cost function $f_t(\cdot) : X \rightarrow \mathbb{R}$ is minimized exactly at $\mathbf{0}$, since $f_t(\mathbf{0}) = 0 \forall t \in [T]$. By convexity of the cost functions, this implies that the gradient of some cost function f_t is similarly minimized at the point $\mathbf{0}$, and thus the *best marginal cost* for f_t can be obtained by taking an infinitesimally small step away from $\mathbf{0}$ in at least one direction, which we denote (without loss of generality) as $i \in [d]$. For brevity, we denote this best marginal cost in f_t by $[\nabla f_t]^i$.

The assumption that $\text{OPT}(\mathcal{I}) < \phi(z^{(j)}) - \beta$ implies that instance \mathcal{I} must contain a cost function $f_m(\cdot)$ at some arbitrary time step m ($m \in [T]$) which satisfies $[\nabla f_m]^i < \phi(z^{(j)}) - \beta$ for any dimension $i \in [d]$.

Prior work (Lorenz et al., 2008; Sun et al., 2021b) has shown that the worst-case for online search problems with long-term demand constraints occurs when cost functions arrive online in descending order, so we henceforth adopt this assumption. Recall that at each time step, ALG1 solves the pseudo-cost minimization problem defined in (4). Without loss of generality, assume that $z^{(m-1)} = z^{(j)}$, i.e. the cost function $f_m(\cdot)$ arrives when ALG1 has already reached its final utilization (before the compulsory trade). This implies that $\mathbf{x}_m = \mathbf{0}$, and further that $c(\mathbf{x}_m) = 0$. This implies that $f_m(\mathbf{x}) + \|\mathbf{x} - \mathbf{x}_{m-1}\|_{\ell_1(\mathbf{w})} > \int_{z^{(m-1)}}^{z^{(m-1)} + c(\mathbf{x})} \phi(u) du$, since the pseudo-cost minimization problem should be minimized when ALG1 sets $\mathbf{x}_m = \mathbf{0}$.

The pseudo-cost minimization problem at time step m can be expressed as follows:

$$\mathbf{x}_m = \arg \min_{\mathbf{x} \in \mathbb{R}^d : c(\mathbf{x}) \leq 1 - z^{(m-1)}} f_m(\mathbf{x}) + \|\mathbf{x} - \mathbf{x}_{m-1}\|_{\ell_1(\mathbf{w})} - \int_{z^{(m-1)}}^{z^{(m-1)} + c(\mathbf{x})} \phi(u) du.$$

We note that $\|\mathbf{x} - \mathbf{x}_{m-1}\|_{\ell_1(\mathbf{w})}$ is upper bounded by $\beta(z^{(m-1)} + c(\mathbf{x}))$, since in the worst case, the previous online decision \mathbf{x}_{m-1} built up all of ALG1's utilization ($z^{(m-1)}$) so far, and in the next step it will have to switch dimensions to ramp up to \mathbf{x} .

Since the function ϕ is monotonically decreasing on $z \in [0, 1]$, the \mathbf{x}_m solving the true pseudo-cost minimization problem is lower-bounded by the $\check{\mathbf{x}}_m$ solving the following minimization problem (i.e., $c(\check{\mathbf{x}}_m) \leq c(\mathbf{x}_m)$):

$$\check{\mathbf{x}}_m = \arg \min_{\mathbf{x} \in \mathbb{R}^d : c(\mathbf{x}) \leq 1 - z^{(m-1)}} f_m(\mathbf{x}) + \beta(z^{(m-1)} + c(\mathbf{x})) - \int_{z^{(m-1)}}^{z^{(m-1)} + c(\mathbf{x})} \phi(u) du.$$

This further gives the following:

$$\begin{aligned} & f_m(\mathbf{x}) + \beta(z^{(t)} + c(\mathbf{x})) - \int_{z^{(m-1)}}^{z^{(m-1)} + c(\mathbf{x})} \phi(u) du \\ & f_m(\mathbf{x}) + \beta(z^{(t)} + c(\mathbf{x})) - \int_{z^{(m-1)}}^{z^{(m-1)} + c(\mathbf{x})} \left[U - \beta + \left(\frac{U}{\alpha} - U + 2\beta \right) \exp(u/\alpha) \right] du \\ & f_m(\mathbf{x}) - (U - \beta)c(\mathbf{x}) + \beta(z^{(t)} + c(\mathbf{x})) - [U - U\alpha + 2\beta\alpha] \left(\exp\left(\frac{z^{(m-1)} + c(\mathbf{x})}{\alpha}\right) - \exp\left(\frac{z^{(m-1)}}{\alpha}\right) \right) \end{aligned}$$

By assumption, since $f_m(\cdot)$ is convex and satisfies $[\nabla f_m]^i < \phi(z^{(j)}) - \beta$ at $\mathbf{x} = \mathbf{0}$, there must exist a dimension i in f_m where an incremental step away from $\mathbf{0}$ in direction i satisfies the following inequality: $f_m(\mathbf{x}) \lesssim [\nabla f_m]^i \cdot c(\mathbf{x}) < [\phi(z^{(j)}) - \beta]c(\mathbf{x})$ for some \mathbf{x} where $c(\mathbf{x}) > 0$. Thus, we have the following in the pseudo-cost minimization problem:

$$([\nabla f_m]^i - U + \beta)c(\mathbf{x}) + \beta(z^{(t)} + c(\mathbf{x})) - [U - U\alpha + 2\beta\alpha] \left(\exp\left(\frac{z^{(m-1)} + c(\mathbf{x})}{\alpha}\right) - \exp\left(\frac{z^{(m-1)}}{\alpha}\right) \right)$$

Letting $c(x)$ be some scalar y (which is valid since we assume there is at least one dimension in $f_t(\cdot)$ where the cost function growth rate is at most ∇f_m), the pseudo-cost minimization problem finds the value y which minimizes the following

quantity:

$$([\nabla f_m]^i - U + \beta)y + \beta(z^{(t)} + y) - [U - U\alpha + 2\beta\alpha] \left(\exp\left(\frac{z^{(m-1)} + y}{\alpha}\right) - \exp\left(\frac{z^{(m-1)}}{\alpha}\right) \right)$$

Taking the derivative of the above with respect to y yields the following:

$$\begin{aligned} \frac{d}{dy} \left[([\nabla f_m]^i - U + \beta)y + \beta(z^{(t)} + y) - [U - U\alpha + 2\beta\alpha] \left(\exp\left(\frac{z^{(m-1)} + y}{\alpha}\right) - \exp\left(\frac{z^{(m-1)}}{\alpha}\right) \right) \right] = \\ = [\nabla f_m]^i + 2\beta - U + \frac{(U\alpha - 2\alpha\beta - U) \exp\left(\frac{z^{(m-1)} + y}{\alpha}\right)}{\alpha} \end{aligned}$$

If $y = 0$, we have the following by assumption that $[\nabla f_m]^i < \phi(z^{(j)}) - \beta$ and that $z^{(j)} = z^{(m-1)}$:

$$\begin{aligned} [\nabla f_m]^i + 2\beta - U + (U - 2\beta - U/\alpha) \exp\left(\frac{z^{(m-1)}}{\alpha}\right) &< \phi(z^{(j)}) + \beta - U + (U - 2\beta - U/\alpha) \exp\left(\frac{z^{(m-1)}}{\alpha}\right) \\ &< \phi(z^{(j)}) - (\phi(z^{(m-1)})) = 0 \end{aligned}$$

The above derivation implies that the derivative of the cost minimization problem at $c(\mathbf{x}) = 0$ (which corresponds to the case where $\mathbf{x} = \mathbf{0}$) is strictly less than 0. This further implies that $\check{\mathbf{x}}_m$ must be *non-zero*, since the minimizer must satisfy $c(\check{\mathbf{x}}_m) > 0$. Since $c(\check{\mathbf{x}}_m)$ lower bounds the true $c(\mathbf{x}_m)$, this causes a contradiction, as it was assumed that the utilization after time step m would satisfy $z^{(m)} = z^{(m-1)} = z^{(j)}$, but if $c(\mathbf{x}_m) > 0$, $z^{(m)}$ must satisfy $z^{(m)} > z^{(m-1)}$.

It then follows by contradiction that $\text{OPT}(\mathcal{I}) \geq \phi(z^{(j)}) - \beta$.

Lemma B.3. *The cost of ALG1 on any valid CFL instance $\mathcal{I} \in \Omega$ is upper bounded by*

$$\text{ALG1}(\mathcal{I}) \leq \int_0^{z^{(j)}} \phi(u) du + \beta z^{(j)} + (1 - z^{(j)})U. \quad (13)$$

Proof of Lemma B.3. First, recall that $z^{(t)} = \sum_{\tau \in [t]} c(\mathbf{x}_\tau)$ is non-decreasing over $t \in [T]$.

Observe that whenever $c(\mathbf{x}_t) > 0$, we know that $f_t(\mathbf{x}_t) + \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_{\ell_1(\mathbf{w})} < \int_{z^{(t-1)}}^{z^{(t-1)} + c(\mathbf{x}_t)} \phi(u) du$. Then, if $c(\mathbf{x}_t) = 0$, which corresponds to the case when $\mathbf{x}_t = \mathbf{0}$, we have the following:

$$f_t(\mathbf{x}_t) + \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_{\ell_1(\mathbf{w})} - \int_{z^{(t-1)}}^{z^{(t-1)} + c(\mathbf{x}_t)} \phi(u) du = 0 + \|\mathbf{x}_{t-1}\|_{\ell_1(\mathbf{w})} - 0 \leq \beta c(\mathbf{x}_{t-1})$$

This gives that for any time step where $c(\mathbf{x}_t) = 0$, we have the following inequality:

$$f_t(\mathbf{x}_t) + \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_{\ell_1(\mathbf{w})} \leq \beta c(\mathbf{x}_{t-1}), \forall t \in [T] : c(\mathbf{x}_t) = 0. \quad (14)$$

And thus, since any time step where $c(\mathbf{x}_t) > 0$ implies $f_t(\mathbf{x}_t) + \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_{\ell_1(\mathbf{w})} < \int_{z^{(t-1)}}^{z^{(t-1)} + c(\mathbf{x}_t)} \phi(u) du$, we have the following inequality for all time steps (*i.e.*, *an upper bound on the excess cost not accounted for in the pseudo-cost threshold function or compulsory trade*)

$$f_t(\mathbf{x}_t) + \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_{\ell_1(\mathbf{w})} - \int_{z^{(t-1)}}^{z^{(t-1)} + c(\mathbf{x}_t)} \phi(u) du \leq \beta c(\mathbf{x}_{t-1}), \forall t \in [T]. \quad (15)$$

Thus, we have

$$\beta z^{(j)} = \sum_{t \in [j]} \beta c(\mathbf{x}_{t-1}) \geq \sum_{t \in [j]} \left[f_t(\mathbf{x}_t) + \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_{\ell_1(\mathbf{w})} - \int_{z^{(t-1)}}^{z^{(t-1)} + c(\mathbf{x}_t)} \phi(u) du \right] \quad (16)$$

$$= \sum_{t \in [j]} [f_t(\mathbf{x}_t) + \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_{\ell_1(\mathbf{w})}] - \int_0^{z^{(j)}} \phi(u) du \quad (17)$$

$$= \text{ALG1} - (1 - z^{(j)})U - \int_0^{z^{(j)}} \phi(u) du. \quad (18)$$

Combining Lemma B.2 and Lemma B.3 gives

$$\frac{\text{ALG1}(\mathcal{I})}{\text{OPT}(\mathcal{I})} \leq \frac{\int_0^{z^{(j)}} \phi(u) du + \beta z^{(j)} + (1 - z^{(j)})U}{\phi(z^{(j)}) - \beta} \leq \alpha, \quad (19)$$

where the last inequality holds since for any $z \in [0, 1]$

$$\int_0^z \phi(u) du + \beta z + (1 - z)U = \int_0^z [U - \beta + (U/\alpha - U + 2\beta) \exp(z/\alpha)] + \beta z + (1 - z)U \quad (20)$$

$$= (U - \beta)z + \alpha(U/\alpha - U + 2\beta)[\exp(z/\alpha) - 1] + \beta z + (1 - z)U \quad (21)$$

$$= \alpha(U/\alpha - U + 2\beta)[\exp(z/\alpha) - 1] + U \quad (22)$$

$$= \alpha[U - 2\beta + (U/\alpha - U + 2\beta) \exp(z/\alpha)] \quad (23)$$

$$= \alpha[\phi(z) - \beta]. \quad (24)$$

Thus, we conclude that ALG1 is α -competitive for CFL. \square

B.3. Proof of Corollary 3.3

In this section, we prove Corollary 3.3, which shows that the worst-case competitive ratio of ALG1 for MAL is again upper bounded by α as defined in (6).

Proof of Corollary 3.3. To show this result, we first prove a result stated in the main body, namely Lemma 2.2, which states the following: For any MAL instance on a weighted star metric (X, d) , there is a corresponding CFL instance on $(\mathbb{R}^{n-1}, \|\cdot\|_{\ell_1(\mathbf{w}^*)})$ which preserves $f_t^a(\cdot) \forall t, c(\cdot) \forall a \in X$, and upper bounds $d(a, b) \forall (a, b) \in X$.

Before the proof, we note that Bansal & Coester (2022) showed online metric allocation on a weighted star metric (X, d) is identical to convex function chasing (with separable cost functions) on the normed vector space $(\Delta_n, \|\cdot\|_{\ell_1(\mathbf{w})})$, where Δ_n is the n -point simplex in \mathbb{R}^n and $\|\cdot\|_{\ell_1(\mathbf{w})}$ is the weighted ℓ_1 norm, with weights given by the corresponding edge weight in the underlying star metric as follows:

$$\|\mathbf{x}\|_{\ell_1(\mathbf{w})} = \sum_{a \in X} \mathbf{w}^a |\mathbf{x}^a|.$$

Proof of Lemma 2.2. Recall that by assumption, the MAL instance contains at least one OFF point denoted by $a' \in X$ in the MAL instance, where $\mathbf{c}^{a'} = 0$. Without loss of generality, let the *first dimension* in Δ_n correspond to this OFF point.

We define a linear map $\Phi : \Delta_n \rightarrow \mathbb{R}^{n-1}$, where Φ has $n - 1$ rows and n columns, and is specified as follows:

$$\Phi_{i,j} = \begin{cases} 1 & \text{if } j = i+1 \\ 0 & \text{otherwise} \end{cases}$$

It is straightforward to see that $\Phi \mathbf{x} \in \mathbb{R}^{n-1}$, $\forall \mathbf{x} \in \Delta_n$.

Recall that a CFL decision space is the ℓ_1 ball defined by the long-term constraint function in \mathbb{R}^{n-1} . For any MAL instance with constraint function $c(\mathbf{x}) : \Delta_n \rightarrow \mathbb{R}$, we can define a long-term constraint function $c'(\mathbf{x}') : \mathbb{R}^{n-1} \rightarrow \mathbb{R}$ as follows. The MAL constraint function $c(\mathbf{x})$ is defined as $\|\cdot\|_{\ell_1(\mathbf{c})}$ for some vector $\mathbf{c} \in \mathbb{R}^{n-1}$. Then

$$\begin{aligned}\mathbf{c}' &= \Phi\mathbf{c} \\ c'(\mathbf{x}') &= \|\mathbf{x}'\|_{\ell_1(\mathbf{c}')} \quad \forall \mathbf{x}' \in \mathbb{R}^{n-1}\end{aligned}$$

Furthermore, for any $z \in [0, 1]$, let $\mathbf{x} \in \Delta_n : c(\mathbf{x}) < 1 - z$. Then it follows that $\Phi\mathbf{x}$ is in $\mathbb{R}^{n-1} : c'(\mathbf{x}') < 1 - w$.

Recall that cost functions in the MAL instance are convex and linearly separable as follows:

$$f_t(\mathbf{x}) = \sum_{a \in X} f_t^a(\mathbf{x}^a)$$

Next, again letting $\mathbf{x} \in \Delta_n$, note that the i th term in \mathbf{x} is identical to the $(i-1)$ th term in $\Phi\mathbf{x}$ (excluding the first term in \mathbf{x}). Then we can construct cost functions in the CFL instance as follows:

$$f'_t(\mathbf{x}') = \sum_{i \in [n-1]} f_t^{i+1}(\mathbf{x}^i)$$

Under the mapping Φ , note that it is straightforward to show that $f_t(\mathbf{x}) = f'_t(\Phi\mathbf{x})$ for any $\mathbf{x} \in \Delta_n$.

Finally, consider the distances in the MAL instance's weighted star metric, which can be expressed as a weighted ℓ_1 norm defined by \mathbf{w} , where the terms of \mathbf{w} correspond to the weighted edges of the star metric. Recall that $\beta := \max_{a', a} \|a' - a\|_{\ell_1(\mathbf{w})}$, i.e., the maximum distance between the OFF point and any other point in the weighted star.

Then we define a corresponding distance metric in the CFL instance, which is an ℓ_1 norm weighted by $\mathbf{w}' \in \mathbb{R}^{n-1}$, which is defined as follows:

$$\mathbf{w}'^i = \mathbf{w}^{i+1} + \mathbf{w}^0.$$

Note that \mathbf{w}^0 is the edge weight associated with the OFF point. Then for any $(\mathbf{x}, \mathbf{y}) \in \Delta_n$, it is straightforward to show the following:

$$\|\mathbf{x} - \mathbf{y}\|_{\ell_1(\mathbf{w})} \leq \|\Phi\mathbf{x} - \Phi\mathbf{y}\|_{\ell_1(\mathbf{w}')}$$

This follows since for any $(\mathbf{x}, \mathbf{y}) \in \Delta_n$ where $\mathbf{x}^0 = 0$ and $\mathbf{y}^0 = 0$ (i.e., allocations which do not allocate anything to the OFF point), $\|\Phi\mathbf{x} - \Phi\mathbf{y}\|_{\ell_1(\mathbf{w}')} = \|\mathbf{x} - \mathbf{y}\|_{\ell_1(\mathbf{w})} + \|\mathbf{x} - \mathbf{y}\|_{\ell_1} \cdot \mathbf{w}^0$.

Conversely, if either \mathbf{x} or \mathbf{y} have $\mathbf{x}^0 > 0$ or $\mathbf{y}^0 > 0$, we have $\|\mathbf{x} - \mathbf{y}\|_{\ell_1(\mathbf{w})} \leq \|\Phi\mathbf{x} - \Phi\mathbf{y}\|_{\ell_1(\mathbf{w}')} \leq \|\mathbf{x} - \mathbf{y}\|_{\ell_1(\mathbf{w})} + \|\mathbf{x} - \mathbf{y}\|_{\ell_1} \cdot \mathbf{w}^0$. Finally, supposing that (without loss of generality) \mathbf{x} has $\mathbf{x}^0 = 1$, we have that $\|\mathbf{x} - \mathbf{y}\|_{\ell_1(\mathbf{w})} = \|\Phi\mathbf{x} - \Phi\mathbf{y}\|_{\ell_1(\mathbf{w}')}$. Thus, $\|\Phi\mathbf{x} - \Phi\mathbf{y}\|_{\ell_1(\mathbf{w}')}$ upper bounds $\|\mathbf{x} - \mathbf{y}\|_{\ell_1(\mathbf{w})}$. Furthermore, the constructed distance metric preserves β , i.e. given $(a', a) = \arg \max_{a', a} \|a' - a\|_{\ell_1(\mathbf{w})}$, we have that $\|\Phi a' - \Phi a\|_{\ell_1(\mathbf{w}')} = \beta$.

Next, we show that the transformation Φ is bijective. We define the affine map $\Phi^{-1} : \mathbb{R}^{n-1} \rightarrow \Delta_n$ as follows: Φ^{-1} has n rows and $n-1$ columns, where the first row is all -1 , and the bottom n rows are the $n \times n$ identity matrix. Let $\mathbf{b} \in \mathbb{R}^{n-1}$ denote the vector with $\mathbf{b}^0 = 1$ and all other terms are zero, i.e., $\mathbf{b}_i = 0 \forall i \geq 1$.

For any $\mathbf{x}' \in \mathbb{R}^{n-1} : c'(\mathbf{x}') \leq 1$, it is straightforward to show that $\Phi^{-1}\mathbf{x}' + \mathbf{b}$ is in Δ_n , since by definition we have that $\sum_{i \in [n-1]} (\Phi^{-1}\mathbf{x}' + \mathbf{b})_i = 1$. Furthermore, by definition of $c'(\mathbf{x}')$, we have that $c(\Phi^{-1}\mathbf{x}' + \mathbf{b}) = c'(\mathbf{x}')$, because the i th term (excluding the first term) of $\Phi^{-1}\mathbf{x}' + \mathbf{b}$ is identical to the $(i-1)$ th term of \mathbf{x}' . Similarly, by definition of f'_t , we have that $f_t(\Phi^{-1}\mathbf{x}' + \mathbf{b}) = f'_t(\mathbf{x}')$.

Finally, considering the distance metric, we have that for any $(\mathbf{x}', \mathbf{y}') \in \mathbb{R}^{n-1} : c'(\mathbf{x}') \leq 1$:

$$\|(\Phi^{-1}\mathbf{x}' + \mathbf{b}) - (\Phi^{-1}\mathbf{y}' + \mathbf{b})\|_{\ell_1(\mathbf{w})} \leq \|\mathbf{x}' - \mathbf{y}'\|_{\ell_1(\mathbf{w}')}$$

This follows by considering that for any \mathbf{x}' , $\Phi^{-1}\mathbf{x}' + \mathbf{b}$ adds a dimension (corresponding to the OFF point) and sets $(\Phi^{-1}\mathbf{x}' + \mathbf{b}) = 1 - \|\mathbf{x}'\|_1$. Then the distance between any two points which allocate a non-negative fraction to the OFF

point in Δ_n is \leq the distance in \mathbb{R}^{n-1} by definition of the weight vector \mathbf{w}' , and the distance between e.g., the allocation fully in the OFF point (a') and any other allocation is exactly preserved.

Furthermore, note that if $\mathbf{w}^0 = 0$ (i.e., the weight of the OFF state in the weighted star metric is 0), Φ is a bijective isometry between $(\Delta_n, \|\cdot\|_{\ell_1(\mathbf{w})})$ and $(\mathbb{R}^{n-1}, \|\cdot\|_{\ell_1(\mathbf{w}')})$. \square

The transformation defined by Φ in [Lemma 2.2](#) allows us to put decisions on the CFL instance $(\mathbb{R}^{n-1}, \|\cdot\|_{\ell_1(\mathbf{w}')})$ in one-to-one correspondence with decisions in $(\Delta_n, \|\cdot\|_{\ell_1(\mathbf{w})})$.

Below, we formalize this by proving a result stated in the main body ([Proposition 2.3](#)) which states the following: Given an algorithm ALG for CFL, any performance bound on ALG which assumes OPT does not pay any switching cost will translate to an identical performance bound for MAL whose parameters depend on the corresponding CFL instance constructed according to [Lemma 2.2](#).

Proof of Proposition 2.3. The cost of ALG on the CFL instance is an upper bound on the cost of the ALG's decisions mapped into the MAL instance. This follows since the cost functions are preserved exactly between the two instances, the long-term constraint function is preserved exactly, and the CFL switching cost is by definition an upper bound on the MAL switching cost.

If the CFL performance bound assumes that OPT does not pay any switching cost (e.g., as in [Theorem 3.2](#)), lower bounding the cost of OPT on the CFL instance is equivalent to lower bounding the cost of OPT on the MAL instance, as the cost functions and constraint functions are preserved exactly.

Thus, we have that any such performance bound for ALG on the CFL instance constructed appropriately (as in [Lemma 2.2](#)) immediately gives an identical performance bound for the MAL instance, yielding the result. \square

By [Lemma 2.2](#), we have that since ALG1 is α -competitive for CFL ([Theorem 3.2](#)), ALG1 is α -competitive for any CFL instance constructed based on a MAL instance. Furthermore, by [Proposition 2.3](#), ALG1 is also α -competitive on the underlying MAL instance, where α is given by (6). \square

B.4. Proof of Theorem 3.4

In this section, we prove [Theorem 3.4](#), which shows that α as given by (6) is the best competitive ratio achievable for CFL.

To show this lower bound, we first define a family of special adversaries, and then show that the competitive ratio for any deterministic algorithm is lower bounded under the instances provided by these adversaries.

Prior work has shown that difficult instances for online search problems with a minimization objective occur when inputs arrive at the algorithm in an decreasing order of cost ([El-Yaniv et al., 2001](#); [Lorenz et al., 2008](#); [Sun et al., 2021b](#); [Lechowicz et al., 2023](#)). For CFL, we additionally consider how an adaptive adversary can essentially force an algorithm to incur a large switching cost in the worst-case. We now formalize such a family of adversaries $\{\mathcal{A}_y\}_{y \in [L, U]}$, where \mathcal{A}_y is called a y -adversary.

Definition B.4 (y -adversary for CFL). Let $w, m \in \mathbb{Z}$ be sufficiently large, and $\delta := (U - L)/w$.

Without loss of generality, let $k = \arg \max_{i \in [d]} \mathbf{w}_i$, where \mathbf{w} is the weight vector for $\|\cdot\|_{\ell_1(\mathbf{w})}$, and let $\beta = \max_{i \in [d]} \mathbf{w}_i$. For $y \in [L, U]$, an adaptive adversary \mathcal{A}_y sequentially presents two types of cost functions $f_t(\cdot)$ to both ALG and OPT.

These types of cost functions are $\mathbf{Up}(\mathbf{x}) = U\mathbf{1}\mathbf{x}^\top$, and $\mathbf{Down}^i(\mathbf{x}) = \sum_{j \neq k}^d U\mathbf{x}^j + (U - i\delta)\mathbf{x}^k$.

The adversary sequentially presents cost functions from these two types in an alternating, “continuously decreasing” order. Specifically, they start by presenting cost function $\mathbf{Up}(\mathbf{x})$, up to m times.

Then, they present $\mathbf{Down}^1(\mathbf{x})$, which has linear cost coefficient U in every direction except direction k , which has cost coefficient $(U - 1 \cdot \delta)$. $\mathbf{Down}^1(\mathbf{x})$ is presented up to m times. If ALG ever “accepts” a cost function $\mathbf{Down}^1(\mathbf{x})$ (i.e., if ALG makes a decision \mathbf{x} where $c(\mathbf{x}) > 0$), the adaptive adversary immediately presents $\mathbf{Up}(\mathbf{x})$ starting in the next time step until either ALG moves to the origin (i.e. online decision $\mathbf{x} = \mathbf{0}$) or ALG's utilization $z = 1$.

The adversary continues alternating in this manner, presenting $\mathbf{Down}^2(\mathbf{x})$ up to m times, followed by $\mathbf{Up}(\mathbf{x})$ if ALG accepts anything, followed by $\mathbf{Down}^3(\mathbf{x})$ up to m times, and so on. This continues until the adversary presents $\mathbf{Down}^{w_y}(\mathbf{x})$, where

$y := (U - w_y \delta)$, up to m times. After presenting $\mathbf{Down}^{w_y}(\mathbf{x})$, \mathcal{A}_y will present $\mathbf{Up}(\mathbf{x})$ until either ALG moves to the origin or has utilization $z = 1$. Finally, the adversary presents exactly m cost functions of the form $\sum_{j \neq k}^d U \mathbf{x}^j + (y + \varepsilon) \mathbf{x}^k$, followed by m cost functions $\mathbf{Up}(\mathbf{x})$.

The mechanism of this adaptive adversary is designed to present “good cost functions” (i.e., $\mathbf{Down}^i(\mathbf{x})$) in a worst-case decreasing order, interrupted by blocks of “bad cost functions” $\mathbf{Up}(\mathbf{x})$ which force a large switching cost in the worst case.

\mathcal{A}_U is simply a stream of m cost functions U , and the final cost functions in any y -adversary instance are always $\mathbf{Up}(\mathbf{x})$.

Proof of Theorem 3.4. Let $g(y)$ denote a *conversion function* $[L, U] \rightarrow [0, 1]$, which fully describes the progress towards the long-term constraint (before the compulsory trade) of a deterministic ALG playing against adaptive adversary \mathcal{A}_y . Note that for large w , the adaptive adversary $\mathcal{A}_{y-\delta}$ is equivalent to first playing \mathcal{A}_y (besides the last two batches of cost functions), and then processing batches with cost functions $\mathbf{Down}^{w_y+1}(\mathbf{x})$ and $\mathbf{Up}(\mathbf{x})$. Since ALG is deterministic and the conversion is unidirectional (irrevocable), we must have that $g(y - \delta) \geq g(y)$, i.e. $g(y)$ is non-increasing in $[L, U]$. Intuitively, the entire capacity should be satisfied if the minimum possible price is observed, i.e $g(L) = 1$.

Note that for $\varepsilon \rightarrow 0$, the optimal solution for adversary \mathcal{A}_y is $\text{OPT}(\mathcal{A}_y) = y + 2\beta/m$, and for m sufficiently large, $\text{OPT}(\mathcal{A}_y) \rightarrow y$.

Due to the adaptive nature of each y -adversary, any deterministic ALG incurs a switching cost proportional to $g(y)$, which gives the amount of utilization obtained by ALG before the end of \mathcal{A}_y ’s sequence.

Whenever ALG accepts some cost function with coefficient $U - i\delta$ in direction k , the adversary presents $\mathbf{Up}(\mathbf{x})$ starting in the next time step. Any ALG which does not switch away immediately obtains a competitive ratio strictly worse than an algorithm which does switch away (if an algorithm accepts c fraction of a good price and switches away immediately, the switching cost it will pay is $2\beta c$. An algorithm may continue accepting c fraction of coefficient U in the subsequent time steps, but a sequence exists where this decision will take up too much utilization to recover when better cost functions are presented later. In the extreme case, if an algorithm continues accepting c fraction of these U coefficients, it might fill its utilization and then OPT can accept a cost function which is arbitrarily better).

Since accepting any price by a factor of c incurs a switching cost of $2\beta c$, the switching cost paid by ALG on adversary \mathcal{A}_y is $2\beta g(y)$. We assume that ALG is notified of the compulsory trade, and does not incur a significant switching cost during the final batch.

Then the total cost incurred by an α^* -competitive online algorithm ALG on adversary \mathcal{A}_y is $\text{ALG}(\mathcal{A}_y) = g(U/\alpha^*)U/\alpha^* - \int_{U/\alpha^*}^y udg(u) + 2\beta g(y) + (1 - g(y))U$, where $udg(u)$ is the cost of buying $dg(u)$ utilization at cost coefficient u , the last term is from the compulsory trade, and the second to last term is the switching cost incurred by ALG. Note that any deterministic ALG which makes conversions when the price is larger than U/α^* can be strictly improved by restricting conversions to prices $\leq U/\alpha^*$.

For any α^* -competitive online algorithm, the corresponding conversion function $g(\cdot)$ must satisfy $\text{ALG}(\mathcal{A}_y) \leq \alpha^* \text{OPT}(\mathcal{A}_y) = \alpha^* y, \forall y \in [L, U]$. This gives a necessary condition which the conversion function must satisfy as follows:

$$\text{ALG}(\mathcal{A}_y) = g(U/\alpha^*)U/\alpha^* - \int_{U/\alpha^*}^y udg(u) + 2\beta g(y) + (1 - g(y))U \leq \alpha^* y, \quad \forall y \in [L, U].$$

By integral by parts, the above implies that the conversion function must satisfy $g(y) \geq \frac{U - \alpha^* y}{U - y - 2\beta} - \frac{1}{U - y - 2\beta} \int_{U/\alpha^*}^y \frac{U - \alpha^* u}{U - u - 2\beta} du$. By Grönwall’s Inequality (Mitrinovic et al., 1991, Theorem 1, p. 356), we have that

$$\begin{aligned} g(y) &\geq \frac{U - \alpha^* y}{U - y - 2\beta} - \frac{1}{U - y - 2\beta} \int_{U/\alpha^*}^y \frac{U - \alpha^* u}{U - u - 2\beta} \cdot \exp \left(\int_u^y \frac{1}{U - r - 2\beta} dr \right) du \\ &\geq \frac{U - \alpha^* y}{U - y - 2\beta} - \int_{U/\alpha^*}^y \frac{U - \alpha^* u}{(U - u - 2\beta)^2} du \\ &\geq \frac{U - \alpha^* y}{U - y - 2\beta} - \left[\frac{U\alpha^* - U - 2\beta\alpha^*}{u + 2\beta - U} - \alpha^* \ln(u + 2\beta - U) \right]_{U/\alpha^*}^y \\ &\geq \alpha^* \ln(y + 2\beta - U) - \alpha^* \ln(U/\alpha^* + 2\beta - U), \quad \forall y \in [L, U]. \end{aligned}$$

$g(L) = 1$ by the problem definition – we can combine this with the above constraint to give the following condition for an α^* -competitive online algorithm:

$$\alpha^* \ln(L + 2\beta - U) - \alpha^* \ln(U/\alpha^* + 2\beta - U) \leq g(L) = 1.$$

The optimal α^* is obtained when the above inequality is binding, so solving for the value of α^* which solves $\alpha^* \ln(L + 2\beta - U) - \alpha^* \ln(U/\alpha^* + 2\beta - U) = 1$ yields that the best competitive ratio for any ALG solving CFL is $\alpha^* \geq \left[W\left(\frac{e^{2\beta/U}(L/U+2\beta/U-1)}{e}\right) - \frac{2\beta}{U} + 1 \right]^{-1}$. \square

B.5. Proof of Corollary 3.5

In this section, we prove [Corollary 3.5](#), which shows that α as given by (6) is the best competitive ratio achievable for MAL.

To show this lower bound, we build off of the family of adversaries in [Definition B.4](#), which are designed to force an algorithm to incur a large switching cost while satisfying the long-term constraint. In [Definition B.5](#) we define this family of adversarial instances tailored for MAL.

Definition B.5 (y -adversary for MAL). Let $w, m \in \mathbb{Z}$ be sufficiently large, and $\delta := (U-L)/w$.

Recall that \mathbf{w} denotes the vector of edge weights for each point in the weighted star metric X , and the OFF point is defined (without loss of generality) as the point $a' \in X$ where $\mathbf{c}^{a'} = 0$ and $f_t^{a'}(\mathbf{x}^a) = 0 \forall t \in [T], \forall \mathbf{x}^a \in [0, 1]$. We will assume that $\mathbf{c}^a = 1 \forall a \in X : a \neq a'$.

Then we set $\mathbf{w}^{a'} = 0$, i.e., the OFF point is connected to the interior vertex of the weighted star with an edge of weight 0. Without loss of generality, we let $k = \arg \max_{a \in [n]} \mathbf{w}^a$ denote the largest edge weight of any other (non-OFF) point in the metric. By definition, recall that $\beta = \mathbf{w}^k$.

For $y \in [L, U]$, an adaptive adversary \mathcal{A}_y sequentially presents two different sets of cost functions $f_t^a(\cdot)$ at each point in the metric space.

These sets of cost functions are **Up** = $\{f^a(x) = U\mathbf{x}^a \ \forall a \in X \setminus \{a'\}\}$, and **Down**ⁱ = $\{f^k(\mathbf{x}^k) = (U - i\delta)\mathbf{x}^k\} \cap \{f^a(\mathbf{x}^a) = U\mathbf{x}^a \ \forall a \in X \setminus \{a', k\}\}$. Note that the adversary only ever presents cost functions with a coefficient $< U$ at the point k which corresponds to the largest edge weight.

The adversary sequentially presents either of these two sets of cost functions in an alternating, “continuously decreasing” order. Specifically, they start by presenting **Up**, up to m times.

Then, they present **Down**, which has cost coefficient U in every point except point k , which has cost coefficient $(U - 1 \cdot \delta)$. **Down**¹ is presented up to m times. If ALG ever “accepts” a cost function in **Down**¹ (i.e., if ALG makes a decision x where $c(x) > 0$), the adaptive adversary immediately presents **Up** starting in the next time step until either ALG moves entirely to the OFF point (i.e. online decision $x^{a'} = 1$) or ALG’s utilization $z = 1$.

The adversary continues alternating in this manner, presenting **Down**² up to m times, followed by **Up** if ALG accepts anything, followed by **Down**³ up to m times, and so on. This continues until the adversary presents **Down** ^{w_y} , where $y = (U - w_y\delta)$, up to m times. After presenting **Down** ^{w_y} , \mathcal{A}_y will present **Up**(x) until either ALG moves to the OFF point or has utilization $z = 1$. Finally, the adversary presents the set of cost functions $\{f^k(\mathbf{x}^k) = (y + \varepsilon)\mathbf{x}^k\} \cap \{f^a(\mathbf{x}^a) = U\mathbf{x}^a \ \forall a \in X \setminus \{a', k\}\}$ m times, followed by **Up** m times.

The mechanism of this adaptive adversary is designed to present “good cost functions” (i.e., **Down**ⁱ) in a worst-case decreasing order, interrupted by blocks of “bad cost functions” **Up** which force a large switching cost in the worst case.

As in [Theorem 3.4](#), \mathcal{A}_U is simply a stream of m **Up** sets of cost functions, and the final cost functions in any y -adversary instance are always **Up**.

Proof of Corollary 3.5. As previously, we let $g(y)$ denote a *conversion function* $[L, U] \rightarrow [0, 1]$, which fully describes the progress towards the long-term constraint (before the compulsory trade) of a deterministic ALG playing against adaptive adversary \mathcal{A}_y . Since ALG is deterministic and the conversion is unidirectional (irrevocable), $g(y)$ is non-increasing in $[L, U]$. Intuitively, the entire long-term constraint should be satisfied if the minimum possible price is observed, i.e $g(L) = 1$. For $\varepsilon \rightarrow 0$, the optimal solution for adversary \mathcal{A}_y is $\text{OPT}(\mathcal{A}_y) = y + 2\beta/m$, and for m sufficiently large, $\text{OPT}(\mathcal{A}_y) \rightarrow y$.

As in [Theorem 3.4](#), the adaptive nature of each y -adversary forces any deterministic ALG to incur a switching cost of $2\beta g(y)$ on adversary \mathcal{A}_y , and we assume that ALG does not incur a significant switching cost during the final batch (i.e., during the compulsory trade).

Then the total cost incurred by an α^* -competitive online algorithm ALG on adversary \mathcal{A}_y is $\text{ALG}(\mathcal{A}_y) = g(U/\alpha^*)U/\alpha^* - \int_{U/\alpha^*}^y u dg(u) + 2\beta g(y) + (1 - g(y))U$, where $udg(u)$ is the cost of buying $dg(u)$ utilization at cost coefficient u , the last term is from the compulsory trade, and the second to last term is the switching cost incurred by ALG. Note that this expression for the cost is exactly as defined in [Theorem 3.4](#).

Thus by [Theorem 3.4](#), for any α^* -competitive online algorithm, the conversion function $g(\cdot)$ must satisfy $\text{ALG}(\mathcal{A}_y) \leq \alpha^* \text{OPT}(\mathcal{A}_y) = \alpha^* y, \forall y \in [L, U]$. Via integral by parts and Grönwall's Inequality ([Mitrovic et al., 1991](#), Theorem 1, p. 356), we have the following condition on $g(y)$:

$$g(y) \geq \alpha^* \ln(y + 2\beta - U) - \alpha^* \ln(U/\alpha^* + 2\beta - U), \quad \forall y \in [L, U].$$

$g(L) = 1$ by the problem definition – combining this with the previous condition gives the following condition for an α^* -competitive online algorithm:

$$\alpha^* \ln(L + 2\beta - U) - \alpha^* \ln(U/\alpha^* + 2\beta - U) \leq g(L) = 1.$$

As in [Theorem 3.4](#), the optimal α^* is obtained when the above inequality is binding, yielding that the best competitive ratio for any ALG solving MAL is $\alpha^* \geq \left[W \left(\frac{e^{2\beta/U} (L/U + 2\beta/U - 1)}{e} \right) - \frac{2\beta}{U} + 1 \right]^{-1}$. \square

C. Proofs for Section 4 (Learning-Augmentation)

C.1. Proof of Lemma 4.1

In this section, we prove [Lemma 4.1](#), which shows that the baseline fixed-ratio combination algorithm (**Baseline**) is $(1 + \epsilon)$ -consistent and $\left(\frac{(U+2\beta)/L(\alpha-1-\epsilon)+\alpha\epsilon}{(\alpha-1)} \right)$ -robust for CFL, given any $\epsilon \in [0, \alpha - 1]$ and where α is as defined in [\(6\)](#). Recall that [Lemma 4.1](#) specifies **ALG1** as the “robust algorithm” to use for the following analysis.

Proof of Lemma 4.1. Under the assumption that ADV satisfies the long-term constraint, (i.e., that $\sum_{t=1}^T c(\mathbf{a}_t) \geq 1$), we first observe that the online solution of **Baseline** must also satisfy the long-term constraint.

Under the assumptions of CFL, note that $c(\mathbf{x})$ is linear (i.e., a weighted ℓ_1 norm with weight vector \mathbf{c}). By definition, denoting the decisions of **ALG1** by $\tilde{\mathbf{x}}_t$, we know that $\sum_{t=1}^T c(\tilde{\mathbf{x}}_t) \geq 1$.

Thus, we have the following:

$$\sum_{t=1}^T c(\mathbf{x}_t) = \sum_{t=1}^T c(\lambda \mathbf{a}_t + (1 - \lambda) \tilde{\mathbf{x}}_t) = \lambda \sum_{t=1}^T c(\mathbf{a}_t) + (1 - \lambda) \sum_{t=1}^T c(\tilde{\mathbf{x}}_t) \geq \lambda + (1 - \lambda) = 1.$$

Let $\mathcal{I} \in \Omega$ be an arbitrary valid CFL sequence. We denote the *hitting* and *switching* costs of the robust advice by $\text{ALG1}_{\text{hitting}}$ and $\text{ALG1}_{\text{switch}}$, respectively. Likewise, the hitting and switching cost of the black-box advice ADV is denoted by $\text{ADV}_{\text{hitting}}$ and $\text{ADV}_{\text{switch}}$.

The total cost of **Baseline** is upper bounded by the following:

$$\begin{aligned}
 \text{Baseline}(\mathcal{I}) &= \sum_{t=1}^T f_t(\mathbf{x}_t) + \sum_{t=1}^{T+1} \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_{\ell_1(\mathbf{w})}, \\
 &= \sum_{t=1}^T f_t(\lambda \mathbf{a}_t + (1-\lambda) \tilde{\mathbf{x}}_t) + \sum_{t=1}^{T+1} \|\lambda \mathbf{a}_t + (1-\lambda) \tilde{\mathbf{x}}_t - \lambda \mathbf{a}_{t-1} - (1-\lambda) \tilde{\mathbf{x}}_{t-1}\|_{\ell_1(\mathbf{w})}, \\
 &\leq \lambda \sum_{t=1}^T f_t(\mathbf{a}_t) + (1-\lambda) \sum_{t=1}^T f_t(\tilde{\mathbf{x}}_t) + \sum_{t=1}^{T+1} \|\lambda \mathbf{a}_t - \lambda \mathbf{a}_{t-1}\|_{\ell_1(\mathbf{w})} + \sum_{t=1}^{T+1} \|(1-\lambda) \tilde{\mathbf{x}}_t - (1-\lambda) \tilde{\mathbf{x}}_{t-1}\|_{\ell_1(\mathbf{w})}, \\
 &\leq \lambda \text{ADV}_{\text{hitting}}(\mathcal{I}) + (1-\lambda) \text{ALG1}_{\text{hitting}}(\mathcal{I}) + \lambda \sum_{t=1}^{T+1} \|\mathbf{a}_t - \mathbf{a}_{t-1}\|_{\ell_1(\mathbf{w})} + (1-\lambda) \sum_{t=1}^{T+1} \|\tilde{\mathbf{x}}_t - \tilde{\mathbf{x}}_{t-1}\|_{\ell_1(\mathbf{w})}, \\
 &\leq \lambda \text{ADV}_{\text{hitting}}(\mathcal{I}) + (1-\lambda) \text{ALG1}_{\text{hitting}}(\mathcal{I}) + \lambda \text{ADV}_{\text{switch}}(\mathcal{I}) + (1-\lambda) \text{ALG1}_{\text{switch}}(\mathcal{I}), \\
 &\leq \lambda \text{ADV}(\mathcal{I}) + (1-\lambda) \text{ALG1}(\mathcal{I}).
 \end{aligned}$$

Since $\text{ALG1} \leq \alpha \cdot \text{OPT} \leq \alpha \cdot \text{ADV}$, this gives the following:

$$\text{Baseline}(\mathcal{I}) \leq \lambda \text{ADV}(\mathcal{I}) + (1-\lambda) \alpha \text{ADV}(\mathcal{I}), \quad (25)$$

$$\text{Baseline}(\mathcal{I}) \leq (\lambda + (1-\lambda)\alpha) \cdot \text{ADV}(\mathcal{I}) \quad (26)$$

$$\text{Baseline}(\mathcal{I}) \leq (1+\epsilon) \cdot \text{ADV}(\mathcal{I}). \quad (27)$$

Furthermore, since $\text{ADV} \leq U + 2\beta \leq \frac{\text{OPT}}{L/(U+2\beta)}$, we have:

$$\text{Baseline}(\mathcal{I}) \leq \lambda \frac{\text{OPT}(\mathcal{I})}{L/(U+2\beta)} + (1-\lambda) \alpha \text{OPT}(\mathcal{I}), \quad (28)$$

$$\text{Baseline}(\mathcal{I}) \leq \left[\frac{\lambda(U+2\beta)}{L} + (1-\lambda)\alpha \right] \cdot \text{OPT}(\mathcal{I}), \quad (29)$$

$$\text{Baseline}(\mathcal{I}) \leq \left(\frac{(U+2\beta)/L(\alpha-1-\epsilon)+\alpha\epsilon}{(\alpha-1)} \right) \cdot \text{OPT}(\mathcal{I}). \quad (30)$$

By combining (27) and (30), we conclude that **Baseline** is $(1+\epsilon)$ -consistent with respect to black-box advice **ADV**, and $\left(\frac{(U+2\beta)/L(\alpha-1-\epsilon)+\alpha\epsilon}{(\alpha-1)}\right)$ -robust. \square

C.2. Proof of Theorem 4.3

In this section, we prove [Theorem 4.3](#), which shows that **CLIP** is $(1+\epsilon)$ -consistent and γ^ϵ -robust for **CFL**, where γ^ϵ is defined as the solution to the following (as in (10)):

$$\gamma^\epsilon = \epsilon + \frac{U}{L} - \frac{\gamma^\epsilon}{L}(U-L) \ln \left(\frac{U-L-2\beta}{U-U/\gamma^\epsilon-2\beta} \right).$$

Proof of Theorem 4.3. We show the above result by separately considering consistency (the competitive ratio when advice is correct) and robustness (the competitive ratio when advice is not correct) in turn.

Recall that the black-box advice **ADV** is denoted by a decision \mathbf{a}_t at each time t . Throughout the following proof, we use shorthand notation CLIP_t to denote the cost of **CLIP** up to time t , and ADV_t to denote the cost of **ADV** up to time t . We start with the following lemma to prove consistency.

Lemma C.1. **CLIP** is $(1+\epsilon)$ -consistent.

Proof. First, we note that the constrained optimization enforces that the possible cost so far plus a compulsory term is always within $(1+\epsilon)$ of the advice. Formally, if time step $j \in [T]$ denotes the time step marking the start of the compulsory trade, we have that the constraint given by (8) holds for every time step $t \in [j]$.

Thus, to show $(1 + \epsilon)$ consistency, we must resolve the cost during the *compulsory trade* and show that the final cumulative cost of CLIP is upper bounded by $(1 + \epsilon)\text{ADV}$.

Let $\mathcal{I} \in \Omega$ be an arbitrary valid CFL sequence. If the compulsory trade begins at time step $j < T$, both CLIP and ADV must greedily fill their remaining utilization during the last m time steps $[j, T]$. This is assumed to be feasible, and the switching cost is assumed to be negligible as long as m is sufficiently large.

Let $(1 - z^{(j-1)})$ denote the remaining long-term constraint that must be satisfied by CLIP at the final time step, and let $(1 - A^{(j-1)})$ denote the remaining long-term constraint to be satisfied by ADV.

We consider the following two cases, which correspond to the cases where CLIP has *under-* and *over-* provisioned with respect to ADV, respectively.

Case 1: CLIP(\mathcal{I}) has “underprovisioned” ($(1 - z^{(j-1)}) > (1 - A^{(j-1)})$). In this case, CLIP must satisfy *more* of the long-term constraint during the compulsory trade compared to ADV.

From the previous time step, we know that the following constraint holds: $\text{CLIP}_{j-1} + \|\mathbf{x}_{j-1} - \mathbf{a}_{j-1}\|_{\ell_1(\mathbf{w})} + \|\mathbf{a}_{j-1}\|_{\ell_1(\mathbf{w})} + (1 - A^{(j-1)})L + (A^{(j-1)} - z^{(j-1)})U \leq (1 + \epsilon) [\text{ADV}_{j-1} + \|\mathbf{a}_{j-1}\|_{\ell_1(\mathbf{w})} + (1 - A^{(j-1)})L]$.

Let $\{\mathbf{x}_t\}_{t \in [j, T]}$ and $\{\mathbf{a}_t\}_{t \in [j, T]}$ denote the decisions made by CLIP and ADV during the compulsory trade, respectively. By definition, we have that $\sum_{t=j}^T c(\mathbf{x}_t) = (1 - z^{(j-1)})$ and $\sum_{t=j}^T c(\mathbf{a}_t) = (1 - A^{(j-1)})$.

Considering $\{f_t(\cdot)\}_{t \in [j, T]}$, we know that by definition $\sum_{t=j}^T f_t(\mathbf{a}_t) \geq L \sum_{t=j}^T c(\mathbf{a}_t)$, and by convex assumptions on the cost functions, $\sum_{t=j}^T f_t(\mathbf{x}_t) \leq \sum_{t=j}^T f_t(\mathbf{a}_t) + U(\sum_{t=j}^T c(\mathbf{x}_t) - \sum_{t=j}^T c(\mathbf{a}_t))$.

Note that the worst case for CLIP occurs when $\sum_{t=j}^T f_t(\mathbf{a}_t) = L \sum_{t=j}^T c(\mathbf{a}_t)$, as ADV is able to satisfy the rest of the long-term constraint at the best possible price.

By the constraint in the previous time step, we have the following:

$$\begin{aligned} \text{CLIP}_{j-1} + \|\mathbf{a}_{j-1}\|_{\ell_1(\mathbf{w})} + (1 - A^{(j-1)})L + (A^{(j-1)} - z^{(j-1)})U \\ \leq (1 + \epsilon) [\text{ADV}_{j-1} + \|\mathbf{a}_{j-1}\|_{\ell_1(\mathbf{w})} + (1 - A^{(j-1)})L], \\ \text{CLIP}_{j-1} + \|\mathbf{a}_{j-1}\|_{\ell_1(\mathbf{w})} + L \sum_{t=j}^T c(\mathbf{a}_t) + U \left(\sum_{t=j}^T c(\mathbf{x}_t) - \sum_{t=j}^T c(\mathbf{a}_t) \right) \\ \leq (1 + \epsilon) \left[\text{ADV}_{j-1} + \|\mathbf{a}_{j-1}\|_{\ell_1(\mathbf{w})} + L \sum_{t=j}^T c(\mathbf{a}_t) \right], \\ \text{CLIP}(\mathcal{I}) \leq (1 + \epsilon) [\text{ADV}(\mathcal{I})]. \end{aligned}$$

Case 2: CLIP(\mathcal{I}) has “overprovisioned” ($(1 - z^{(j-1)}) \leq (1 - A^{(j-1)})$). In this case, CLIP must satisfy *less* of the long-term constraint during the compulsory trade compared to ADV.

From the previous time step, we know that the following constraint holds: $\text{CLIP}_{j-1} + \|\mathbf{x}_{j-1} - \mathbf{a}_{j-1}\|_{\ell_1(\mathbf{w})} + \|\mathbf{a}_{j-1}\|_{\ell_1(\mathbf{w})} + (1 - A^{(j-1)})L + (A^{(j-1)} - z^{(j-1)})U \leq (1 + \epsilon) [\text{ADV}_{j-1} + \|\mathbf{a}_{j-1}\|_{\ell_1(\mathbf{w})} + (1 - A^{(j-1)})L]$.

Let $\{\mathbf{x}_t\}_{t \in [j, T]}$ and $\{\mathbf{a}_t\}_{t \in [j, T]}$ denote the decisions made by CLIP and ADV during the compulsory trade, respectively. By definition, we have that $\sum_{t=j}^T c(\mathbf{x}_t) = (1 - z^{(j-1)})$ and $\sum_{t=j}^T c(\mathbf{a}_t) = (1 - A^{(j-1)})$.

Considering $\{f_t(\cdot)\}_{t \in [j, T]}$, we know that by definition, $\sum_{t=j}^T f_t(\mathbf{x}_t) \geq L \sum_{t=j}^T c(\mathbf{x}_t)$, and $\sum_{t=j}^T f_t(\mathbf{a}_t) \geq L \sum_{t=j}^T c(\mathbf{a}_t)$. By convexity, because $\sum_{t=j}^T c(\mathbf{x}_t) \leq \sum_{t=j}^T c(\mathbf{a}_t)$, $\sum_{t=j}^T f_t(\mathbf{x}_t) \leq \sum_{t=j}^T f_t(\mathbf{a}_t)$.

By the constraint in the previous time step, we have:

$$\frac{\text{CLIP}_{j-1} + \|\mathbf{x}_{j-1} - \mathbf{a}_{j-1}\|_{\ell_1(\mathbf{w})} + \|\mathbf{a}_{j-1}\|_{\ell_1(\mathbf{w})} + (1 - z^{(j-1)})L}{\text{ADV}_{j-1} + \|\mathbf{a}_{j-1}\|_{\ell_1(\mathbf{w})} + (1 - A^{(j-1)})L} =$$

$$\frac{\text{CLIP}_{j-1} + \|\mathbf{x}_{j-1} - \mathbf{a}_{j-1}\|_{\ell_1(\mathbf{w})} + \|\mathbf{a}_{j-1}\|_{\ell_1(\mathbf{w})} + L \sum_{t=j}^T c(\mathbf{x}_t)}{\text{ADV}_{j-1} + \|\mathbf{a}_{j-1}\|_{\ell_1(\mathbf{w})} + L \sum_{t=j}^T c(\mathbf{a}_t)} \leq (1 + \epsilon).$$

Let $y = \sum_{t=j}^T f_t(\mathbf{x}_t) - L \sum_{t=j}^T c(\mathbf{x}_t)$, and let $y' = \sum_{t=j}^T f_t(\mathbf{a}_t) - L \sum_{t=j}^T c(\mathbf{a}_t)$. By definition, $y \geq 0$ and $y' \geq 0$.

Note that $\text{CLIP}_{j-1} + \|\mathbf{x}_{j-1} - \mathbf{a}_{j-1}\|_{\ell_1(\mathbf{w})} + \|\mathbf{a}_{j-1}\|_{\ell_1(\mathbf{w})} + (1 - z^{(j-1)})L + y \geq \text{CLIP}(\mathcal{I})$ and $\text{ADV}_{j-1} + \|\mathbf{a}_{j-1}\|_{\ell_1(\mathbf{w})} + L \sum_{t=j}^T c(\mathbf{a}_t) + y' = \text{ADV}(\mathcal{I})$.

Furthermore, by definition and convexity of the cost functions $f_t(\cdot)$, we have that $y \leq y'$.

Combined with the constraint from the previous time step, we have the following bound:

$$\frac{\text{CLIP}(\mathcal{I})}{\text{ADV}(\mathcal{I})} \leq \frac{\text{CLIP}_{j-1} + \|\mathbf{x}_{j-1} - \mathbf{a}_{j-1}\|_{\ell_1(\mathbf{w})} + \|\mathbf{a}_{j-1}\|_{\ell_1(\mathbf{w})} + (1 - z^{(j-1)})L + y}{\text{ADV}_{j-1} + \|\mathbf{a}_{j-1}\|_{\ell_1(\mathbf{w})} + (1 - A^{(j-1)})L + y'}$$

$$\leq \frac{\text{CLIP}_{j-1} + \|\mathbf{a}_{j-1}\|_{\ell_1(\mathbf{w})} + L \sum_{t=j}^T c(\mathbf{x}_t)}{\text{ADV}_{j-1} + \|\mathbf{a}_{j-1}\|_{\ell_1(\mathbf{w})} + L \sum_{t=j}^T c(\mathbf{a}_t)} \leq (1 + \epsilon).$$

Thus, by combining the bounds in each of the above two cases, the result follows, and we conclude that CLIP is $(1 + \epsilon)$ -consistent with accurate advice. \square

Having proved the consistency of CLIP, we proceed to show robustness in the next lemma.

Lemma C.2. CLIP is γ^ϵ -robust, where γ^ϵ is as defined in (10).

Proof. Let $\epsilon \in (0, \alpha - 1]$ be the target consistency (recalling that CLIP is $(1 + \epsilon)$ consistent), and let $\mathcal{I} \in \Omega$ denote an arbitrary valid CFL sequence.

To prove the robustness of CLIP, we consider two “bad cases” for the advice $\text{ADV}(\mathcal{I})$, and show that in the worst-case, CLIP’s competitive ratio is bounded by γ^ϵ .

Case 1: $\text{ADV}(\mathcal{I})$ is “inactive”. Consider the case where ADV accepts nothing during the main sequence and instead satisfies the entire long-term constraint in the final time step. In the worst-case, this gives that $\text{ADV}(\mathcal{I}) = U + 2\beta$.

Based on the consistency constraint (and using the fact that CLIP will always be “overprocuring” w.r.t. ADV throughout the main sequence), we can derive an upper bound on the amount that CLIP is allowed to accept from the robust pseudo-cost minimization. Recall the following constraint:

$$\text{CLIP}_{t-1} + f_t(\mathbf{x}_t) + \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_{\ell_1(\mathbf{w})} + \|\mathbf{x}_t - \mathbf{a}_t\|_{\ell_1(\mathbf{w})} + \|\mathbf{a}_t\|_{\ell_1(\mathbf{w})} + (1 - z^{(t-1)} - c(\mathbf{x}_t))L$$

$$\leq (1 + \epsilon) \left[\text{ADV}_t + \|\mathbf{a}_t\|_{\ell_1(\mathbf{w})} + (1 - A^{(t)})L \right].$$

Proposition C.3. z_{PCM} is an upper bound on the amount that CLIP can accept from the pseudo-cost minimization without violating $(1 + \epsilon)$ consistency, and is defined as:

$$z_{PCM} = \gamma^\epsilon \ln \left[\frac{U - L - 2\beta}{U - U/\gamma^\epsilon - 2\beta} \right]$$

Proof. Consider an arbitrary time step t . When CLIP is *not* allowed to accept anything more from the robust pseudo-cost minimization, we have that $c(\mathbf{x}_t)$ is restricted to be 0 (recall that $\mathbf{a}_t = \mathbf{0}$ for any time steps before T , because the advice is assumed to be inactive).

By definition, since any cost functions accepted in CLIP_{t-1} can be attributed to the robust pseudo-cost minimization, we have the following in the worst-case:

$$\text{CLIP}_{t-1} = \int_0^{z^{(t-1)}} \phi^\epsilon(u) du + \beta z^{(t-1)}.$$

Combining the above with the left-hand side of the consistency constraint, we have the following by observing that $\mathbf{x}_t = \mathbf{0}$ and $\mathbf{a}_t = \mathbf{0}$, and the switching cost to “ramp-up” is absorbed into the pseudo-cost ϕ :

$$\text{CLIP}_{t-1} + (1 - z^{(t-1)})L = \int_0^{z^{(t-1)}} \phi^\epsilon(u) du + \beta z^{(t-1)} + (1 - z^{(t-1)})L.$$

As stated, let $z^{(t-1)} = z_{\text{PCM}}$. Then by properties of the pseudo-cost,

$$\begin{aligned} \text{CLIP}_{t-1} + (1 - z_{\text{PCM}})L &= \int_0^{z_{\text{PCM}}} \phi(u) du + \beta z_{\text{PCM}} + (1 - z_{\text{PCM}})U + (1 - z_{\text{PCM}})L - (1 - z_{\text{PCM}})U, \\ &= \gamma^\epsilon [\phi^\epsilon(z_{\text{PCM}}) - \beta] + (1 - z_{\text{PCM}})L - (1 - z_{\text{PCM}})U, \\ &= \gamma^\epsilon L + (L - U) \left(1 - \gamma^\epsilon \ln \left[\frac{U - L - 2\beta}{U - U/\gamma^\epsilon - 2\beta} \right] \right), \\ &= \gamma^\epsilon L + L - U - (L - U) \gamma^\epsilon \ln \left[\frac{U - L - 2\beta}{U - U/\gamma^\epsilon - 2\beta} \right]. \end{aligned}$$

Substituting for the definition of γ^ϵ , we obtain:

$$\begin{aligned} \text{CLIP}_{t-1} + (1 - z_{\text{PCM}})L &= \gamma^\epsilon L + L - U - (L - U) \gamma^\epsilon \ln \left[\frac{U - L - 2\beta}{U - U/\gamma^\epsilon - 2\beta} \right], \\ &= \left[\epsilon L + U - \gamma^\epsilon (U - L) \ln \left[\frac{U - L - 2\beta}{U - U/\gamma^\epsilon - 2\beta} \right] \right] + L - U + (U - L) \gamma^\epsilon \ln \left[\frac{U - L - 2\beta}{U - U/\gamma^\epsilon - 2\beta} \right], \\ &= \epsilon L + L = (1 + \epsilon)L. \end{aligned}$$

As $(1 + \epsilon)L$ is exactly the right-hand side of the consistency constraint (i.e., $(1 + \epsilon) [\text{ADV}_t + \|\mathbf{a}_t\|_{\ell_1(\mathbf{w})} + (1 - A_t)L] = (1 + \epsilon)L$), this completes the proposition. \square

If CLIP is constrained to use at most z_{PCM} of its utilization to be robust, the remaining $(1 - z_{\text{PCM}})$ utilization must be used for the compulsory trade and/or to follow ADV. Thus, we have the following worst-case competitive ratio for CLIP, specifically for Case 2:

$$\frac{\text{CLIP}(\mathcal{I})}{\text{OPT}(\mathcal{I})} \leq \frac{\int_0^{z_{\text{PCM}}} \phi^\epsilon(u) du + \beta z_{\text{PCM}} + (1 - z_{\text{PCM}})U}{L}$$

By the definition of $\phi^\epsilon(p)$, we have the following:

$$\begin{aligned} \frac{\text{CLIP}(\mathcal{I})}{\text{OPT}(\mathcal{I})} &\leq \frac{\int_0^{z_{\text{PCM}}} \phi^\epsilon(u) du + \beta z_{\text{PCM}} + (1 - z_{\text{PCM}})U}{L} \\ &\leq \frac{\gamma^\epsilon [\phi^\epsilon(z_{\text{PCM}}) - \beta]}{L} \leq \frac{\gamma^\epsilon [L + \beta - \beta]}{L} \leq \gamma^\epsilon. \end{aligned}$$

Case 2: $\text{ADV}(\mathcal{I})$ is “overactive”. We now consider the case where ADV accepts bad cost functions which it “should not” accept (i.e. $\text{ADV}(\mathcal{I}) \gg \text{OPT}(\mathcal{I})$). Let $\text{ADV}(\mathcal{I}) = v \gg \text{OPT}_T$ (i.e. the final total hitting and switching cost of ADV is v for some $v \in [L, U + 2\beta]$, and this is much greater than the optimal solution).

This is without loss of generality, since we can assume that v is the “best cost function” accepted by ADV and the consistency ratio changes strictly in favor of ADV. Based on the consistency constraint, we can derive a lower bound on the amount that CLIP *must* accept from ADV in order to stay $(1 + \epsilon)$ -consistent.

To do this, we consider the following sub-cases:

- **Sub-case 2.1:** Let $v \geq \frac{U+\beta}{1+\epsilon}$.

In this sub-case, CLIP can fully ignore the advice, because the following consistency constraint is never binding (note that $\text{ADV}_t \geq \frac{U+\beta}{1+\epsilon} A^{(t)}$):

$$\begin{aligned} \text{CLIP}_{t-1} + f_t(\mathbf{x}_t) + \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_{\ell_1(\mathbf{w})} + \|\mathbf{x}_t - \mathbf{a}_t\|_{\ell_1(\mathbf{w})} + \|\mathbf{a}_t\|_{\ell_1(\mathbf{w})} + (1 - A^{(t)})L + (A^{(t)} - z^{(t-1)} - c(\mathbf{x}_t))U \\ \leq (1 + \epsilon) \left[\text{ADV}_t + \|\mathbf{a}_t\|_{\ell_1(\mathbf{w})} + (1 - A^{(t)})L \right], \\ (1 - A^{(t)})L + (A^{(t)})U + \|\mathbf{a}_t\|_{\ell_1(\mathbf{w})} \leq (1 + \epsilon) \left[\text{ADV}_t + \|\mathbf{a}_t\|_{\ell_1(\mathbf{w})} + (1 - A^{(t)})L \right], \\ (1 - A^{(t)})L + UA^{(t)} + \beta A^{(t)} \leq (1 + \epsilon) \left[\frac{U + \beta}{1 + \epsilon} A^{(t)} + (1 - A^{(t)})L \right] \end{aligned}$$

- **Sub-case 2.2:** Let $v \in (L, \frac{U+\beta}{1+\epsilon})$.

To remain $(1 + \epsilon)$ consistent, CLIP must accept some of these “bad cost functions” denoted by v in the worst-case. We would like to derive a lower bound z_{ADV} , such that z_{ADV} describes the minimum amount that CLIP must accept from ADV in order to always satisfy the $(1 + \epsilon)$ consistency constraint.

Based on the consistency constraint, we have the following:

$$\begin{aligned} \text{CLIP}_{t-1} + f_t(\mathbf{x}_t) + \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_{\ell_1(\mathbf{w})} + \|\mathbf{x}_t - \mathbf{a}_t\|_{\ell_1(\mathbf{w})} + \|\mathbf{a}_t\|_{\ell_1(\mathbf{w})} + (1 - A^{(t)})L + (A^{(t)} - z^{(t-1)} - c(\mathbf{x}_t))U \\ \leq (1 + \epsilon) \left[\text{ADV}_t + \|\mathbf{a}_t\|_{\ell_1(\mathbf{w})} + (1 - A^{(t)})L \right]. \end{aligned}$$

We let $f_t(\mathbf{x}_t) + \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_{\ell_1(\mathbf{w})} + \|\mathbf{x}_t - \mathbf{a}_t\|_{\ell_1(\mathbf{w})} + \|\mathbf{a}_t\|_{\ell_1(\mathbf{w})} \leq vc(\mathbf{x}_t)$ for any $\mathbf{x}_t : c(\mathbf{x}_t) < c(\mathbf{a}_t)$, which holds by convexity of the cost functions $f_t(\cdot)$ and a prevailing condition that $c(\mathbf{x}_t) \leq c(\mathbf{a}_t)$ for the “bad cost functions” accepted by ADV. Note that $v - U$ is negative (by the condition of Sub-case 2.2):

$$\begin{aligned} \text{CLIP}_{t-1} + vc(\mathbf{x}_t) + L - LA^{(t)} + UA^{(t)} - Uz^{(t-1)} - Uc(\mathbf{x}_t) &\leq (1 + \epsilon) \left[vA^{(t-1)} + vc(\mathbf{a}_t) + L - LA^{(t)} \right], \\ vc(\mathbf{x}_t) - Uc(\mathbf{x}_t) &\leq (1 + \epsilon) \left[vA^{(t-1)} + vc(\mathbf{a}_t) + L - LA^{(t)} \right] - \text{CLIP}_{t-1} - L + LA^{(t)} - UA^{(t)} + Uz^{(t-1)}, \\ vc(\mathbf{x}_t) - Uc(\mathbf{x}_t) &\leq vA^{(t)} - UA^{(t)} - \text{CLIP}_{t-1} + Uz^{(t-1)} + \epsilon \left[vA^{(t-1)} + vc(\mathbf{a}_t) + L - LA^{(t)} \right], \\ c(\mathbf{x}_t) &\geq \frac{vA^{(t)} - UA^{(t)} - \text{CLIP}_{t-1} + Uz^{(t-1)} + \epsilon [vA^{(t)} + L - LA^{(t)}]}{v - U}. \end{aligned}$$

In the event that $A^{(t-1)} = 0$ (i.e. nothing has been accepted so far by either ADV or CLIP), we have the following:

$$\begin{aligned} c(\mathbf{x}_t) &\geq \frac{vc(\mathbf{a}_t) - Uc(\mathbf{a}_t) + \epsilon [vc(\mathbf{a}_t) + L - Lc(\mathbf{a}_t)]}{v - U}, \\ c(\mathbf{x}_t) &\geq c(\mathbf{a}_t) - \frac{\epsilon [vc(\mathbf{a}_t) + L - Lc(\mathbf{a}_t)]}{U - v}. \end{aligned}$$

Through a recursive definition, we can show that for any $A^{(t)}$, given that CLIP has accepted $z^{(t-1)}$ of ADV’s suggested prices so far, it must set \mathbf{x}_t such that:

$$z^{(t)} \geq z^{(t-1)} + c(\mathbf{a}_t) - \frac{\epsilon [vc(\mathbf{a}_t) + L - Lc(\mathbf{a}_t)]}{U - v}.$$

Continuing the assumption that v is constant, if CLIP has accepted $z^{(t-1)}$ thus far, we have the following if we assume that the acceptance up to this point happened in a single previous time step m :

$$\begin{aligned}
 c(\mathbf{x}_t) &\geq A^{(t)} + \frac{Uc(\mathbf{x}_m) - \text{CLIP}_{t-1} + \epsilon [vA^{(t)} + L - LA^{(t)}]}{v - U}, \\
 c(\mathbf{x}_t) &\geq c(\mathbf{a}_t) + c(\mathbf{a}_m) + \frac{Uc(\mathbf{x}_m) - vc(\mathbf{x}_m) + \epsilon [v(c(\mathbf{a}_t) + c(\mathbf{a}_m)) + L - L(c(\mathbf{a}_t) + c(\mathbf{a}_m))]}{v - U}, \\
 c(\mathbf{x}_t) &\geq c(\mathbf{a}_t) + c(\mathbf{a}_m) - \mathbf{x}_m + \frac{\epsilon [v(c(\mathbf{a}_t) + c(\mathbf{a}_m)) + L - L(c(\mathbf{a}_t) + c(\mathbf{a}_m))]}{v - U}, \\
 c(\mathbf{x}_t) + c(\mathbf{x}_m) &\geq c(\mathbf{a}_t) + c(\mathbf{a}_m) + \frac{\epsilon [v(c(\mathbf{a}_t) + c(\mathbf{a}_m)) + L - L(c(\mathbf{a}_t) + c(\mathbf{a}_m))]}{v - U}, \\
 z^{(t)} &\geq A^{(t)} + \frac{\epsilon [vA^{(t)} + L - LA^{(t)}]}{v - U}.
 \end{aligned}$$

This gives intuition into the desired z_{ADV} bound. The above describes and motivates that the *aggregate* acceptance by CLIP at any given time step t must satisfy a lower bound. Consider that the worst case for Sub-case 2.2 occurs when all of the v prices accepted by ADV arrive first, before any prices which would be considered by the pseudo-cost minimization. Then let $A^{(t)} = 1$ for some arbitrary time step t , and we have the following lower bound on z_{ADV} :

$$z_{\text{ADV}} \geq 1 - \frac{v\epsilon}{U - v}.$$

If CLIP is forced to use z_{ADV} of its utilization to be $(1 + \epsilon)$ consistent against ADV, that leaves at most $(1 - z_{\text{ADV}})$ utilization for robustness.

We define $z' = \min(1 - z_{\text{ADV}}, z_{\text{PCM}})$ and consider the following two cases.

- **Sub-case 2.2.1:** if $z' = z_{\text{PCM}}$, the worst-case competitive ratio is bounded by the following. Note that if $z' = z_{\text{PCM}}$, the amount of utilization that CLIP can use to “be robust” is exactly the same as in **Case 1**:

$$\begin{aligned}
 \frac{\text{CLIP}(\mathcal{I})}{\text{OPT}(\mathcal{I})} &\leq \frac{\int_0^{z_{\text{PCM}}} \phi(u)du + \beta z_{\text{PCM}} + (1 - z_{\text{ADV}} - z_{\text{PCM}})U + z_{\text{ADV}}v}{L}, \\
 &\leq \frac{\int_0^{z_{\text{PCM}}} \phi(u)du + \beta z_{\text{PCM}} + (1 - z_{\text{PCM}})U}{L} \leq \gamma^\epsilon.
 \end{aligned}$$

- **Sub-case 2.2.2:** if $z' = 1 - z_{\text{ADV}}$, the worst-case competitive ratio is bounded by the following. Note that CLIP *cannot* use z_{PCM} of its utilization for robustness, so the following bound assumes that the cost functions accepted by CLIP are bounded by the *worst* $(1 - z_{\text{ADV}})$ fraction of the pseudo-cost threshold function ϕ^ϵ (which follows since ϕ^ϵ is non-decreasing on $z \in [0, 1]$):

$$\frac{\text{CLIP}(\mathcal{I})}{\text{OPT}(\mathcal{I})} \leq \frac{\int_0^{1-z_{\text{ADV}}} \phi(u)du + \beta(1 - z_{\text{ADV}}) + z_{\text{ADV}}v}{L}.$$

Note that if $z' = 1 - z_{\text{ADV}}$, we know that $1 - z_{\text{ADV}} < z_{\text{PCM}}$, which further gives the following by definition of z_{ADV} :

$$\begin{aligned}
 1 - z_{\text{PCM}} &< 1 - \frac{v\epsilon}{U - v}, \\
 v\epsilon &< (U - v)z_{\text{PCM}}, \\
 v &< \frac{U}{(1 + \frac{\epsilon}{z_{\text{PCM}}})}.
 \end{aligned}$$

By plugging v back into the definition of z_{ADV} , we have that $z_{\text{ADV}}v \leq \left(\frac{(1-z_{\text{PCM}})U}{1+\frac{\epsilon}{z_{\text{PCM}}}} \right)$, giving the following:

$$\begin{aligned} \frac{\text{CLIP}(\mathcal{I})}{\text{OPT}(\mathcal{I})} &\leq \frac{\int_0^{1-z_{\text{ADV}}} \phi(u)du + \beta(1-z_{\text{ADV}}) + \left(\frac{(1-z_{\text{PCM}})U}{1+\frac{\epsilon}{z_{\text{PCM}}}} \right)}{L}, \\ &\leq \frac{\int_0^{z_{\text{PCM}}} \phi(u)du + \beta z_{\text{PCM}} + (1-z_{\text{PCM}})U}{L} \leq \gamma^\epsilon. \end{aligned}$$

Thus, by combining the bounds in each of the above two cases, the result follows, and we conclude that CLIP is γ^ϵ -robust for any advice ADV. \square

Having proven [Lemma C.1](#) (consistency) and [Lemma C.2](#) (robustness), the statement of [Theorem 4.3](#) follows – CLIP is $(1+\epsilon)$ -consistent and γ^ϵ -robust given any advice for CFL. \square

C.3. Proof of [Corollary 4.4](#)

In this section, we prove [Corollary 4.4](#), which shows that CLIP is $(1+\epsilon)$ -consistent and γ^ϵ -robust for MAL, where γ^ϵ is defined in [\(10\)](#).

Proof of [Corollary 4.4](#). We show the above result by separately considering consistency (the competitive ratio when advice is correct) and robustness (the competitive ratio when advice is not correct), relying on the proof of [Theorem 4.3](#).

Consistency. By definition, MAL on a weighted star metric is identical to an instance of convex function chasing with a long-term constraint on $(\Delta_n, \|\cdot\|_{\ell_1(\mathbf{w}')})$, where Δ_n is the n -point simplex in \mathbb{R}^n and $\|\cdot\|_{\ell_1(\mathbf{w}')}$ is the weighted ℓ_1 norm, with weights \mathbf{w}' given by the corresponding edge weight in the underlying star metric.

Observe that the consistency proof given in [Lemma C.1](#) holds when the consistency constraint at each time step is defined as follows:

$$\begin{aligned} \text{CLIP}_{t-1} + f_t(\mathbf{x}) + \|\mathbf{x} - \mathbf{x}_{t-1}\|_{\ell_1(\mathbf{w}')} + \|\mathbf{x} - \mathbf{a}_t\|_{\ell_1(\mathbf{w}')} + \|\mathbf{a}_t\|_{\ell_1(\mathbf{w}')} + (1 - z^{(t-1)} - c(\mathbf{x}))L + \max((A^{(t)} - z^{(t-1)} - c(\mathbf{x})), 0)(U - L) \\ \leq (1+\epsilon)[\text{ADV}_t + \|\mathbf{a}_t\|_{\ell_1(\mathbf{w}')} + (1 - A^{(t)})L], \end{aligned} \quad (31)$$

where \mathbf{x} and \mathbf{a} denote decisions by CLIP and ADV (respectively) supported on Δ_n . Thus, since the consistency proof in [Lemma C.1](#) exactly holds under the CFL vector space corresponding to MAL, we conclude that CLIP is $(1+\epsilon)$ -consistent for MAL.

Robustness. First, we note that the robustness proof given in [Lemma C.2](#) assumes OPT does not pay any switching cost. This implies that the proof of [Lemma C.2](#) meets the conditions of [Proposition 2.3](#), which states that any performance bound for an arbitrary ALG solving CFL which assumes OPT pays no switching cost translates to an identical bound for MAL, where the problem's parameters can be recovered by constructing a corresponding CFL instance according to [Lemma 2.2](#).

Thus, by [Proposition 2.3](#), we conclude that CLIP is γ^ϵ -robust for MAL, where γ^ϵ is defined in [\(10\)](#).

By combining the two results, the statement of [Corollary 4.4](#) follows – CLIP is $(1+\epsilon)$ -consistent and γ^ϵ -robust given any advice ADV for MAL. \square

C.4. Proof of [Theorem 4.5](#)

In this section, we prove [Theorem 4.5](#), which shows that any $(1+\epsilon)$ -consistent algorithm for CFL is at least γ^ϵ -robust, where γ^ϵ is as defined in [\(10\)](#).

Proof of [Theorem 4.5](#). To show this result, we leverage the same special family of y -adversaries for CFL defined in [Definition B.4](#), where $y \in [L, U]$. Recall that $k = \arg \max_{i \in [d]} \mathbf{w}_i$, where \mathbf{w} is the weight vector for $\|\cdot\|_{\ell_1(\mathbf{w})}$.

As in the proof of [Theorem 3.4](#), we note that for adversary \mathcal{A}_y , the optimal offline solution is $\text{OPT}(\mathcal{A}_y) = y + 2\beta/m$, and that as m grows large, $\text{OPT}(\mathcal{A}_y) \rightarrow y$.

Against these adversaries, we consider two types of advice – the first is *bad* advice, which sets $\mathbf{a}_t = \mathbf{0}$ for all time steps $t < T$ (i.e., before the compulsory trade), incurring a final cost of $U + 2\beta$.

On the other hand, *good* advice sets $\mathbf{a}_t = \mathbf{0}$ for all time steps up to the first time step when y is revealed, at which point it sets $\mathbf{a}_t^k = 1/m$ to achieve final cost $\text{ADV}(\mathcal{A}_y) = \text{OPT}(\mathcal{A}_y) = y + 2\beta/m$.

We let $g(y)$ denote a *robust conversion function* $[L, U] \rightarrow [0, 1]$, which fully quantifies the actions of a learning-augmented algorithm LALG playing against adaptive adversary \mathcal{A}_y , where $g(y)$ gives the progress towards the long-term constraint under the instance \mathcal{A}_y before (either) the compulsory trade or the black-box advice sets $\mathbf{a}_t^k > 0$. Note that for large w , the adaptive adversary $\mathcal{A}_{y-\delta}$ is equivalent to first playing \mathcal{A}_y (besides the last two batches of cost functions), and then processing batches with cost functions $\mathbf{Down}^{w_y+1}(\mathbf{x})$ and $\mathbf{Up}(\mathbf{x})$. Since LALG is deterministic and the conversion is unidirectional (irrevocable), we must have that $g(y - \delta) \geq g(y)$, i.e. $g(y)$ is non-increasing in $[L, U]$.

As in the proof of [Theorem 3.4](#), the adaptive nature of each y -adversary forces any algorithm to incur a switching cost proportional to $g(y)$, specifically denoted by $2\beta g(y)$.

For any γ -robust online algorithm LALG given any arbitrary black-box advice, the following must hold:

$$\text{LALG}(\mathcal{A}_y) \leq \gamma \text{OPT}(\mathcal{A}_y) = \gamma y, \quad \forall y \in [L, U].$$

The cost of LALG with conversion function g on an instance \mathcal{A}_y is $\text{LALG}(\mathcal{A}_y) = g(U/\gamma)U/\gamma - \int_{U/\gamma}^y u dg(u) + 2\beta g(y) + (1 - g(y))U$, where $udg(u)$ is the cost of buying $dg(u)$ utilization at price u , the last term is from the compulsory trade, and the second to last term is the switching cost incurred by LALG.

This implies that $g(y)$ must satisfy the following:

$$g(U/\gamma)U/\gamma - \int_{U/\gamma}^y u dg(u) + 2\beta g(y) + (1 - g(y))U \leq \gamma y, \quad \forall y \in [L, U].$$

By integral by parts, the above implies that the conversion function must satisfy $g(y) \geq \frac{U - \gamma y}{U - y - 2\beta} - \frac{1}{U - y - 2\beta} \int_{U/\gamma}^y g(u) du$. By Grönwall's Inequality ([Mitrinovic et al., 1991](#))[Theorem 1, p. 356], we have that

$$g(y) \geq \frac{U - \gamma y}{U - y - 2\beta} - \frac{1}{U - y - 2\beta} \int_{U/\gamma}^y \frac{U - \gamma u}{U - u - 2\beta} \cdot \exp \left(\int_u^y \frac{1}{U - r - 2\beta} dr \right) du \quad (32)$$

$$\geq \frac{U - \gamma y}{U - y - 2\beta} - \int_{U/\gamma}^y \frac{U - \gamma u}{(U - u - 2\beta)^2} du \quad (33)$$

$$\geq \frac{U - \gamma y}{U - y - 2\beta} - \left[\frac{U\gamma - U - 2\beta\gamma}{u + 2\beta - U} - \gamma \ln(u + 2\beta - U) \right]_{U/\gamma}^y \quad (34)$$

$$\geq \gamma \ln(y + 2\beta - U) - \gamma \ln(U/\gamma + 2\beta - U), \quad \forall y \in [L, U]. \quad (35)$$

In addition, to simultaneously be η -consistent when the advice *is* correct, LALG must satisfy $\text{LALG}(\mathcal{A}_L) \leq \eta \text{OPT}(\mathcal{A}_L) = \eta L$. If the advice is correct (and m is sufficiently large), we assume that LALG pays no switching cost to satisfy the long-term constraint at the best cost functions L . It must still pay for switching incurred by the robust algorithm (recall that OPT pays no switching cost).

$$\int_{U/\gamma}^L g(u) du + 2\beta g(L) \leq \eta L - L. \quad (36)$$

By combining equations (35) and (36), the conversion function $g(y)$ of any γ -robust and η -consistent online algorithm must satisfy the following:

$$\gamma \int_{U/\gamma}^L \ln \left(\frac{u + 2\beta - U}{U/\gamma + 2\beta - U} \right) du + 2\beta \left[\gamma \ln \left(\frac{u + 2\beta - U}{U/\gamma + 2\beta - U} \right) \right] \leq \eta L - L. \quad (37)$$

When all inequalities are binding, this equivalently gives that

$$\eta \geq \gamma + 1 - \frac{U}{L} + \frac{\gamma(U - L)}{L} \ln \left(\frac{U - L - 2\beta}{U - U/\gamma^\epsilon - 2\beta} \right). \quad (38)$$

We define η such that $\eta := (1 + \epsilon)$. By substituting for η into (38), we recover the definition of γ^ϵ as given by (10), which subsequently completes the proof. Thus, we conclude that any $(1 + \epsilon)$ -consistent algorithm for CFL is at least γ^ϵ -robust. \square

C.5. Proof of Corollary 4.6

In this section, we prove Corollary 4.6, which shows that any $(1 + \epsilon)$ -consistent algorithm for MAL is at least γ^ϵ -robust, where γ^ϵ is as defined in (10).

Proof of Corollary 4.6. To show this result, we leverage the same special family of y -adversaries for CFL defined in Definition B.5, where $y \in [L, U]$. Recall that $k = \arg \max_{a \in [n]} \mathbf{w}^a$, denotes the largest edge weight of any (non-0FF) point in the metric space, and $\beta = \mathbf{w}^k$.

As in the proof of Theorem 3.4, we note that for adversary \mathcal{A}_y , the optimal offline solution is $\text{OPT}(\mathcal{A}_y) = y + 2\beta/m$, and that as m grows large, $\text{OPT}(\mathcal{A}_y) \rightarrow y$.

Against these adversaries, we consider two types of advice – the first is *bad* advice, which sets $\mathbf{a}_t^{a'} = 1$ (i.e., ADV stays in the 0FF point) for all time steps $t < T$ (i.e., before the compulsory trade), incurring a final cost of $U + 2\beta$.

On the other hand, *good* advice sets $\mathbf{a}_t^{a'} = 1$ for all time steps up to the first time step when y is revealed, at which point it sets $\mathbf{a}_t^k = 1/m$ to achieve final cost $\text{ADV}(\mathcal{A}_y) = \text{OPT}(\mathcal{A}_y) = y + 2\beta/m$.

As previously, we let $g(y)$ denote a *robust conversion function* $[L, U] \rightarrow [0, 1]$, which fully quantifies the actions of a learning augmented algorithm LALG playing against adaptive adversary \mathcal{A}_y . Since LALG is deterministic and the conversion is unidirectional (irrevocable), $g(y)$ is non-increasing in $[L, U]$. Intuitively, the entire long-term constraint should be satisfied if the minimum possible price is observed, i.e $g(L) = 1$.

As in Theorem 4.5, the adaptive nature of each y -adversary forces any deterministic ALG to incur a switching cost of $2\beta g(y)$ on adversary \mathcal{A}_y , and we assume that ALG does not incur a significant switching cost during the final batch (i.e., during the compulsory trade).

For any γ -robust LALG given any arbitrary black-box advice, the following must hold:

$$\text{LALG}(\mathcal{A}_y) \leq \gamma \text{OPT}(\mathcal{A}_y) = \gamma y, \quad \forall y \in [L, U].$$

The cost of LALG with conversion function g on an instance \mathcal{A}_y is $\text{LALG}(\mathcal{A}_y) = g(U/\gamma)U/\gamma - \int_{U/\gamma}^y u dg(u) + 2\beta g(y) + (1 - g(y))U$, where $udg(u)$ is the cost of buying $dg(u)$ utilization at price u , the last term is from the compulsory trade, and the second to last term is the switching cost incurred by LALG. Note that this expression for the cost is exactly as defined in Theorem 4.5.

Thus by Theorem 4.5, for any learning-augmented algorithm LALG which is simultaneously η -consistent and γ -robust, the conversion function $g(\cdot)$ must satisfy the following inequality (via integral by parts and Grönwall's Inequality (Mitrinovic et al., 1991, Theorem 1, p. 356)):

$$\gamma \int_{U/\gamma}^L \ln \left(\frac{u + 2\beta - U}{U/\gamma + 2\beta - U} \right) du + 2\beta \left[\gamma \ln \left(\frac{u + 2\beta - U}{U/\gamma + 2\beta - U} \right) \right] \leq \eta L - L. \quad (39)$$

When all inequalities are binding, this equivalently gives that the optimal η and γ satisfy:

$$\eta \geq \gamma + 1 - \frac{U}{L} + \frac{\gamma(U - L)}{L} \ln \left(\frac{U - L - 2\beta}{U - U/\gamma^\epsilon - 2\beta} \right). \quad (40)$$

We define η such that $\eta := (1 + \epsilon)$. By substituting for η into (38), we recover the definition of γ^ϵ as given by (10), which subsequently completes the proof. Thus, we conclude that any $(1 + \epsilon)$ -consistent algorithm for CFL is at least γ^ϵ -robust. \square