

NONPARAMETRIC CONTEXTUAL ONLINE BILATERAL TRADE

Emanuele Coccia, Martino Bernasconi* & Andrea Celli

Department of Computing Sciences
Bocconi University
Milan, Italy
emanuele.coccia@studbocconi.it,
{martino.bernasconi, andrea.celli2}@unibocconi.it

ABSTRACT

We study the problem of contextual online bilateral trade. At each round, the learner faces a seller-buyer pair and must propose a trade price without observing their private valuations for the item being sold. The goal of the learner is to post prices to facilitate trades between the two parties. Before posting a price, the learner observes a d -dimensional context vector that influences the agent's valuations. Prior work in the contextual setting has focused on linear models. In this work, we tackle a general nonparametric setting in which the buyer's and seller's valuations behave according to arbitrary Lipschitz functions of the context. We design an algorithm that leverages contextual information through a hierarchical tree construction and guarantees regret $\tilde{O}(T^{(d-1)/d})$. Remarkably, our algorithm operates under two stringent features of the setting: (1) one-bit feedback, where the learner only observes whether a trade occurred or not, and (2) strong budget balance, where the learner cannot subsidize or profit from the market participants. We further provide a matching lower bound in the full-feedback setting, demonstrating the tightness of our regret bound.

1 INTRODUCTION

Bilateral trade is a fundamental economic model which describes the interaction between a seller and a buyer, each with private valuations for a good, who seek to engage in trade with the goal of maximizing their individual utilities (Vickrey, 1961; Myerson & Satterthwaite, 1983). This scenario arises in many applications such as ridesharing platforms or energy and financial exchanges.

Recently, Cesa-Bianchi et al. (2024) introduced the online bilateral trade problem, in which at each round t , a new seller and buyer arrive with private valuations s_t and b_t , respectively. The valuation of the seller s_t is the lowest price they are willing to sell the item. Analogously, the valuation of the buyer, b_t , is the highest price the buyer is willing to pay for the item. At each round t , the learner must propose a price to both the buyer and the seller without observing their private valuations. After posting price p_t , a trade happens if $s_t \leq p_t \leq b_t$. The *gain from trade* at time t is

$$GFT(p_t | s_t, b_t) := \mathbb{I}(s_t \leq p_t \leq b_t)(b_t - s_t), \quad (1)$$

and represents the total increase in social welfare generated by the trade. The goal of the learner is to maximize the cumulative gain from trade over time, which is equivalent to minimizing the regret with respect to the best policy in hindsight.

In many practical scenarios, the learner has access to side information at each round t about the seller, buyer, or the item being traded. For instance, this is common in online marketplaces, where items are highly differentiated and the learner can observe certain features of the item prior to setting a price. Such contextual information can be leveraged to estimate the relevance of different features based on historical data. Gaucher et al. (2025) introduced the *feature-based* online bilateral trade problem, in which the learner observes a d -dimensional feature vector x_t before posting a price at

*Corresponding author.

round t . They study a linear model where private valuations are linear functions of the form $x_t^\top \theta$ for some unknown parameter vector θ . In this work, we study whether it is possible to move beyond the linear valuation model. A similar question was examined in the context of pricing in Mao et al. (2018).

We provide a comprehensive characterization of the problem when valuations are *arbitrary L -Lipschitz* functions of the context vectors. A key feature of our analysis is that it operates under the *one-bit feedback* model: at each round, the learner observes only the signal $\mathbb{I}(s_t \leq p_t \leq b_t)$, indicating whether a trade occurred or not. This is the most information-constrained setting in online bilateral trade, and existing contextual no-regret algorithms can handle this feedback model only by relaxing the budget balance constraint (Gaucher et al., 2025).

1.1 TECHNICAL CHALLENGES AND TECHNIQUES

Compared to classic online learning problems, the bilateral-trade setting presents additional technical difficulties, making traditional approaches for nonparametric contextual online learning largely inapplicable. The most notable is that the reward function, which is the gain from trade (Equation (1)), is non-continuous both as a function of p or as a function of the valuations s and b . The second challenge, and possibly a more fundamental one, is the lack of observability of the reward we are trying to maximize; indeed, valuations s and b remain hidden, whether the trade happens or not, and so does the GFT. The only feedback received by the algorithm is the indicator of $\mathbb{I}(s \leq p \leq b)$, which only signals whether the trade was accepted or not.

Although our problem has features that separate it from classic online learning, we do not need to consider a generic Lipschitz loss function, and we will heavily exploit the structure of the reward $p \mapsto \text{GFT}(p|s, b)$ and the Lipschitzness of the functions generating s, b from the contexts. The specific shape of the GFT function (and the absence of noise) is what ultimately allows us break the lower bound of $\Omega(T^{(d+2)/(d+3)})$ (from Lu et al., 2009; Slivkins, 2014, with $p = 1$) that we would obtain if our reward function were a generic Lipschitz function (see Section 1.2 for more details).

Similar to the *adaptive zooming* technique employed by Slivkins (2014) or the *chaining* approach of Cesa-Bianchi et al. (2017), we maintain a discretized representation of the context space, which is geometrically refined in regions in which we observe a lot of contexts. Formally, we use a tree with branching factor 2^d , in which each layer ℓ maintains an $O(2^{-\ell})$ discretization of the function mapping contexts to optimal prices. Moreover, at level ℓ , we “mark” a node if we find a price that was accepted when observing a context in the region described by that node. By Lipschitzness, we know that there is a region of size $O(L2^{-\ell})$ that contains a price for all nodes of the sub-tree.

The main challenge then is to effectively utilize this information, as the discontinuity of the reward functions prevents the use of standard techniques. The key trick here is to heavily rely on randomization when posting prices. By playing uniformly at randomly on an ϵ discretization of the identified set of good prices (that has length $O(2^{-\ell})$) we have a probability of at least $\Omega(2^\ell \epsilon)$ of winning any trade with GFT at least ϵ , and thus we will waste only $O(2^{-\ell}/\epsilon)$ turns on expectation before finding a price that is accepted, marking the node with that price and continuing to the children node. On the other hand, losing all the trades with GFT smaller than ϵ only yields a regret of order $O(\epsilon T)$. By tuning the parameter ϵ , and summing over all node of the tree this leads to an overall regret of $\tilde{O}(T^{(d-2)/d})$, which, although being sublinear, is suboptimal.

In order to obtain an optimal rate, the second trick is associated with reducing the regret suffered when guessing the right price and exploits the geometric structure of the problem. Before starting to guess a correct price in the associated interval, we perform a sort of “hide-and-seek game” against the adversary generating the contexts. We probe iteratively two prices. If we win the trade, we are happy since we have zero regret. Otherwise, if we miss a trade for both prices, by geometric reasoning, we can limit the space in which the context can place the valuation, and we can show it has to be very close (order of $O(2^{-\ell})$) to the diagonal. This limits the possible GFT (and thus the regret) by $O(2^{-\ell})$ in all subsequent nodes. This modifies the maximum regret experienced in the guessing phase to be $O(2^{-2\ell}/\epsilon)$ at nodes at layer ℓ , and by appropriately tuning ϵ we can obtain an optimal regret of $O(T^{(d-1)/d})$.

Crucially, this algorithm relies on explicit knowledge of the Lipschitz constant L , which may be unknown in a practical scenario. Building on the main structure of the vanilla algorithm, in Section 4,

we design an algorithm that works by testing geometrically increasing scales of the Lipschitz constant L and only slightly worsens the guarantees, without requiring knowing the Lipschitz constant.

For the matching (up to logarithmic terms) lower bound of $\Omega(T^{(d-1)/d})$, we use a combination of Yao’s weak minimax principle (Yao, 1977), which relates the performance of randomized algorithms in the worst instance to the performance of deterministic algorithms against a distribution over instances, with a simple version of the McShane’s extension theorem (McShane, 1934) that shows the existence of Lipschitz extensions on the continuum of Lipschitz functions on discrete sets.

1.2 FURTHER RELATED WORKS

Bilateral trade In the offline setting, Myerson & Satterthwaite (1983) demonstrated that no mechanism can simultaneously achieve full efficiency, incentive compatibility, individual rationality, and budget balance in general. This impossibility result motivated a long line of research aimed at designing mechanisms that achieve approximate efficiency in Bayesian settings (McAfee, 2008; Blumrosen & Mizrahi, 2016; Brustle et al., 2017; Deng et al., 2022; Kang et al., 2022).

Online bilateral trade In the online setting, Cesa-Bianchi et al. (2024) study a setting where buyer and seller valuations are i.i.d. from an unknown distribution, and the learner must enforce strong budget balance (*i.e.*, at each round they post the same price to both the seller and the buyer). They provide no-regret algorithms for the full-feedback setting, and under partial feedback when the valuations are drawn independently for the seller and the buyer, and the underlying distribution is smooth. In the case in which valuations are generated by an adversary, Azar et al. (2022) provides an algorithm achieving a tight sublinear 2-regret under weak budget balance (*i.e.*, $p_t \leq q_t$ at each round t). Moreover, Bernasconi et al. (2024) shows that sublinear regret is still achievable if the learner is allowed to enforce global budget balance, meaning the constraint must hold over the entire horizon rather than at each round (see also Chen et al. (2025); Lunghi et al. (2025)). The most closely related work in online bilateral trade is by Gaucher et al. (2025), who consider a contextual variant of the problem where valuations at time t take the form $x_t^\top \theta$, possibly with additive noise. They propose an algorithm achieving $\tilde{O}(T^{2/3})$ regret under two-bit feedback and strong budget balance, and show that sublinear regret is also attainable under one-bit feedback if global budget balance is allowed. This linear contextual valuation model builds on a substantial body of work in contextual pricing with linear valuations (Kleinberg & Leighton, 2003; Amin et al., 2014; Cohen et al., 2020; Lobel et al., 2018; Liu et al., 2021).

Other related models A related but structurally different model is that of online brokerage (Bolić et al., 2024), which differs from standard bilateral trade in that traders may act as buyers or sellers depending on market conditions, and their valuations are drawn i.i.d. from a shared unknown distribution. Bachoc et al. (2024) consider a contextual brokerage model, in which valuations are modeled as zero-mean perturbations of a market price that is linear in the agent’s feature vector.

Online nonparametric contextual online learning The study of online learning with nonparametric contexts was initiated by Hazan & Megiddo (2007), in which they proved an upper bound of $\tilde{O}(T^{(d+1)/(d+2)})$ and context dimension d , while providing a $\Omega(T^{(d-1)/d})$ lower bound. The gap was closed in Rakhlin & Sridharan (2015) for full information by providing an upper bound of $\tilde{O}(T^{(d-1)/d})$ and later refined in Cesa-Bianchi et al. (2017) by providing an explicit algorithm. Both of these works also consider an arbitrary Lipschitz reward function. The model was extended to arbitrary metric action spaces of dimension p in Lu et al. (2009) and Slivkins (2014), that show an almost matching rate of $\Theta(T^{(d+p+1)/(d+p+2)})$. In economic settings, the problem of nonparametric contextual information has been studied in Cesa-Bianchi et al. (2017), which examines second-price auctions, and Mao et al. (2018); Chen & Gallego (2021); Tullii et al. (2024) for pricing.

Independent and concurrent work Independently, Cosson et al. (2026) considers the contextual bilateral trade model with linear contexts of Gaucher et al. (2025), focusing on the noiseless setting. They study gain from trade and profit maximization (unlike Gaucher et al. (2025) which only study gain from trade) in various feedback settings and budget balance constraints. The linear noiseless contextual setting allows for substantially better rates with respect to what is achievable in our work.

In particular, the rates they obtain are at most logarithmic in the time horizon and exponential in the dimension, while ours are polynomial with exponent depending on the dimension.

2 PRELIMINARIES

At the beginning of each round t , the learner observes a public context $x_t \in [0, 1]^d$ generated by an oblivious adversary. At the same time, a new pair of buyer and seller arrives, characterized by private valuations $(b_t, s_t) \in [0, 1]^2$. The valuations are generated from the context by two unknown functions $s_t = f_s(x_t)$ and $b_t = f_b(x_t)$. We assume, as it is standard in the nonparametric contextual bandits literature (see, *e.g.*, (Slivkins, 2014; Cesa-Bianchi et al., 2017)), that $f_s, f_b : [0, 1]^d \rightarrow [0, 1]$ are L -Lipschitz functions, *i.e.*, $|f_s(x) - f_s(x')| \leq L\|x - x'\|_\infty$ for all $x, x' \in [0, 1]^d$ (and similarly for f_b). We denote the class of all L -Lipschitz functions as \mathcal{F}_L .

Objective After observing the contextual information x_t , the learning picks one price $p_t \in [0, 1]$, and the trade is accepted if both the seller and the buyer accept it, *i.e.*, $s_t \leq p_t \leq b_t$. When the trade is accepted, it generates a social welfare of $b_t - s_t$ (that is, the increase in total happiness between the two market participants); in particular, the *gain from trade* is defined as in Equation (1). The objective is to minimize the regret, which given a sequence of contexts $\mathbf{x} = \{x_1, \dots, x_T\} \in ([0, 1]^d)^T$ and two L -Lipschitz functions f_s, f_b , is defined as

$$R_T(f_s, f_b, \mathbf{x}) = \sum_{t=1}^T [f_b(x_t) - f_s(x_t)]^+ - \mathbb{E}[\text{GFT}(p_t | f_s(x_t), f_b(x_t))],$$

where the expectation is taken over the randomization of the algorithm. The regret of an algorithm is then defined as $R_T = \sup_{f_s, f_b \in \mathcal{F}_L, \mathbf{x} \in ([0, 1]^d)^T} R_T(f_s, f_b, \mathbf{x})$, which is the worst-case regret attainable across all functions $f_s, f_b \in \mathcal{F}_L$ and all context sequences \mathbf{x} . Since we are interested in regret guarantees that are asymptotic in T , we will work under the assumption that $T \gg 2^d$. In particular, if $T \leq 2^d$, we can only guarantee an upper bound on the regret of 2^d .

Feedback Our algorithm operates under partial feedback, specifically under *one-bit feedback*. In this feedback model, the learner observes only a single bit of information indicating whether a trade occurred, *i.e.*, $\mathbb{I}(s_t \leq p_t \leq b_t)$. In contrast, the lower bound is established under *full feedback*, where the learner observes the valuation pair (s_t, b_t) after posting the price p_t , irrespective of whether the trade was executed. An intermediate setting is the *two-bit feedback* model, in which the learner observes both $\mathbb{I}(s_t \leq p_t)$ and $\mathbb{I}(p_t \leq b_t)$.

Notation We denote by H_d the d -dimensional Boolean hypercube $\{0, 1\}^d$, and by $\mathbf{0}_d$ a d -dimensional vector of zeros. We denote the projection of $x \in \mathbb{R}^n$ onto $\mathcal{C} \subset \mathbb{R}^n$ as $\Pi_{\mathcal{C}}(x)$. Moreover, for $a \leq b$, let $[a, b]_\epsilon$ denote the uniform grid over the interval $[a, b]$ with step size ϵ , where any grid point outside the interval $[0, 1]$ is projected onto the nearest endpoint, *i.e.*, capped to 0 if below 0 and to 1 if above 1. Formally, $[a, b]_\epsilon := \Pi_{[0, 1]}(\{a, a + \epsilon, \dots, a + k\epsilon, b\})$, where $k = \lfloor \frac{b-a}{\epsilon} \rfloor$.

3 REGRET GUARANTEES FOR KNOWN LIPSCHITZ CONSTANT

In this section, we present an algorithm that guarantees regret of order $\tilde{O}(LT^{(d-1)/d})$ in the case in which the Lipschitz parameter L is known. In Section 4 we extend this approach to include the case in which L is not known.

3.1 HIERARCHICAL TREE STRUCTURE

We start by defining a hierarchical partitioning of the d -dimensional hypercube. We define a tree \mathcal{T} in which each node represents a specific subset of the hypercube. A node $N_{\ell, z}$ in the tree is defined by a level ℓ and a reference point z . The subset of the hypercube defined by a node $N_{\ell, z}$ is

$$\mathbf{Region}(N_{\ell, z}) := [z_1, z_1 + 2^{-\ell}] \times \dots \times [z_d, z_d + 2^{-\ell}].$$

The root of the tree corresponds to $\ell = 0$ and the reference point $z = \mathbf{0}_d$. Each node $N_{\ell, z}$ has 2^d children $N_{\ell+1, z_i}$, where $z_i = z + 2^{-(\ell+1)}h_i$ and h_i ranges over all 2^d binary vectors. We denote by

Algorithm 1

Require: Horizon T , Lipschitzness L , context dimension d .

```

1: for  $t \in [T]$  do
2:   Observe context  $x_t$ 
3:   Initialize  $\ell \leftarrow 0, z \leftarrow \mathbf{0}_d$ 
4:   while  $p_{\ell,z} \neq \emptyset$  and  $N_{\ell,z}$  not leaf do      # Traverse tree until active or leaf node is reached
5:     Find  $z'$  such that  $x_t \in \mathbf{Region}(N_{\ell+1,z'})$ , where  $z' \in \mathcal{C}(\ell, z)$ 
6:      $\ell \leftarrow \ell + 1, z \leftarrow z'$ 
7:   if  $\mathbf{REDUCE}(N_{\ell,z})$  has not terminated then
8:     Execute one iteration of the main loop of  $\mathbf{REDUCE}(N_{\ell,z})$ 
9:   else if  $\mathbf{GUESS}(N_{\ell,z})$  has not terminated then
10:    Execute one iteration of the main loop of  $\mathbf{GUESS}(N_{\ell,z})$ 
11:    if  $\mathbf{GUESS}(N_{\ell,z})$  terminates returning price  $p$  then
12:      Set node  $N_{\ell,z}$  as marked by setting  $p_{\ell,z} \leftarrow p$ 
13:    else Play  $p_{\ell,z}$                                 # If node is a marked leaf play corresponding price

```

Algorithm 2 $\mathbf{REDUCE}(N_{\ell,z})$

```

1: Set  $\bar{p} \leftarrow p_{\ell-1, \mathbf{pred}(\ell, z)}$  (i.e., the price associated with the parent node)      # Initialization
2: Set  $p_L \leftarrow \Pi_{[0,1]}(\bar{p} - L2^{-(\ell-1)}), p_U \leftarrow \Pi_{[0,1]}(\bar{p} + L2^{-(\ell-1)})$ 
3: Initialize  $p \leftarrow p_L$ 
4: repeat                                              # Main Loop
5:   Post price  $p$  and observe  $\mathbb{I}(s_t \leq p \leq b_t)$ 
6:   if  $\mathbb{I}(s_t \leq p \leq b_t) = 0$  then
7:     if  $p = p_L$  then set  $p \leftarrow p_U$ ; else terminate  $p \leftarrow \emptyset$ 
8: until  $p = \emptyset$ 

```

$\mathcal{C}(\ell, z) \subset [0, 1]^d$ the set of all the reference points corresponding to the children of $N_{\ell,z}$. Moreover, for any $z' \in \mathcal{C}(\ell, z)$, we use $\mathbf{pred}(\ell + 1, z')$ to denote the reference point of the parent node of $N_{\ell+1,z'}$ (*i.e.*, point z). Under this notation, the parent node of $N_{\ell,z}$ is $N_{\ell-1, \mathbf{pred}(\ell, z)}$.

We denote the set of reference points in nodes at level ℓ with Z_ℓ , which is defined recursively as $Z_0 = \mathbf{0}_d$ and $Z_\ell = \cup_{z \in Z_{\ell-1}} \mathcal{C}(\ell, z)$. Clearly, $\mathbf{Region}(N_{\ell,z}) = \cup_{z' \in \mathcal{C}(\ell, z)} \mathbf{Region}(N_{\ell+1,z'})$. Then, each level ℓ of the tree induces a partition of the d -dimensional hypercube, since by construction, $\cup_{z \in Z_\ell} \mathbf{Region}(N_{\ell,z}) = [0, 1]^d$. We limit the height of the tree to $H = \lfloor \log_2(T) \rfloor = O(\log(T)d^{-1})$.

3.2 ALGORITHM

The algorithm works as outlined in Algorithm 1. In particular, each node $N_{\ell,z}$ of the tree may be “marked with” a price $p_{\ell,z} \in [0, 1]$ (Line 12).

Definition 1. A node $N_{\ell,z}$ is said to be marked if it has been assigned a marking price p such that, for some context $x \in \mathbf{Region}(N_{\ell,z})$, the trade is accepted at p , that is, $\mathbb{I}(f_s(x) \leq p \leq f_b(x)) = 1$.

A node $N_{\ell,z}$ can be marked only if every node along the path from the root to $N_{\ell,z}$ has already been marked. The set of marked nodes, therefore, forms a rooted subtree. The free (*i.e.*, not marked) leaves of this subtree are referred to as *active nodes*. Intuitively, when we reach an active node $N_{\ell,z}$, our goal is to find a price p that results in a successful trade, that is, one satisfying $\mathbb{I}(f_s(x_t) \leq p \leq f_b(x_t))$. To this end, we leverage the information already encoded in the tree: every node $N_{\ell',z'}$ along the path from the root to $N_{\ell,z}$ is marked with a price $p_{\ell',z'}$ at which a trade occurred. These previously observed prices allow us to restrict the interval of candidate prices and form an estimate for a price p that will be accepted by the valuations $f_s(x_t), f_b(x_t)$.

Algorithm 3 GUESS($N_{\ell,z}$)

```

1: Let  $\bar{p} = p_{\ell-1, \text{pred}(\ell, z)}$  be the price associated to the parent node  $N_{\ell-1, \text{pred}(\ell, z)}$  # Initialization
2: Set  $\mathcal{P}_{\ell,z} = [\bar{p} - L2^{-(\ell-1)}, \bar{p} + L2^{-(\ell-1)}]_{\epsilon}$  with  $\epsilon = LT^{-1/\ell}$ 
3: repeat # Main Loop
4:   Post price  $p \sim \text{Unif}(\mathcal{P}_{\ell,z})$  and observe  $\mathbb{I}(s_t \leq p \leq b_t)$ 
5: until  $\mathbb{I}(s_t \leq p \leq b_t) = 1$ 
6: Return  $p$ 

```

Node-wise regret decomposition In the analysis of the algorithm, we decompose the total regret into the contributions to the regret of each node. Let $A \subseteq [T]$ be a subset of rounds. Then, we define

$$\mathcal{R}(A) := \sum_{t \in A} ([f_b(x_t) - f_s(x_t)]^+ - \text{GFT}_t(p_t | f_s(x_t), f_b(x_t))).$$

For any node $N_{\ell,z} \in \mathcal{T}$, we define $T_{\ell,z} \subseteq [T]$ as the set of rounds in which the traversal of the tree stops on $N_{\ell,z}$ (i.e., the number of rounds in which $N_{\ell,z}$ is the node selected at the end of the loop of Line 4). Then, for any node $N_{\ell,z}$, we define the regret ‘‘accumulated on $N_{\ell,z}$ ’’ as the regret accumulated in rounds $T_{\ell,z}$: $R_{\ell,z} = \mathbb{E}[\mathcal{R}(T_{\ell,z})]$. Then, the overall regret can be rewritten by separating contributions incurred at each level of the tree and at the nodes within that level:

$$R_T = \sum_{\ell=0}^H \sum_{z \in Z_{\ell}} R_{\ell,z}. \quad (2)$$

Any node of the tree that is never visited by Algorithm 1 has $T_{\ell,z} = \emptyset$ and thus does not contribute to the overall regret.

For each observed context x_t , we traverse the tree \mathcal{T} by repeatedly selecting the child node whose associated region contains x_t , until reaching an active node or a leaf node (see Line 4 of Algorithm 1). Then, the algorithm proceeds in one of two ways: (i) if the node is not marked, we invoke one of the two routines REDUCE and GUESS described below (Lines 8 and 10); (ii) if the node is a marked leaf, we post the price previously assigned to it (Line 13).

3.3 REDUCE ROUTINE

Consider an active node $N_{\ell,z}$ and its parent $N_{\ell-1, \text{pred}(\ell, z)}$, and focus on the rounds $t \in [T]$ such that $x_t \in \text{Region}(N_{\ell,z})$. Whenever Algorithm 1 selects such node after the tree traversal, the algorithm executes one iteration of the REDUCE routine on $N_{\ell,z}$ until its termination. The number of rounds in which Algorithm 1 calls $\text{REDUCE}(N_{\ell,z})$ is denoted by $T_{\ell,z}^{\text{REDUCE}}$.

The REDUCE routine is described in Algorithm 2. Starting from the price $\bar{p} = p_{\ell-1, \text{pred}(\ell, z)}$ associated to the parent node of $N_{\ell,z}$, the algorithm computes two prices centered around \bar{p} , each with distance $L2^{-(\ell-1)}$. We denote such prices by $p_L = \Pi_{[0,1]}(\bar{p} - L2^{-(\ell-1)})$ and $p_U = \Pi_{[0,1]}(\bar{p} + L2^{-(\ell-1)})$, where the projection onto $[0, 1]$ ensures that prices are well-defined. Then, REDUCE repeatedly posts the price p_L until a trade is rejected (i.e., $\mathbb{I}(s_t \leq p_L \leq b_t) = 0$). It then switches to posting p_U until that price is also rejected by one of the agents. Once both rejections occur, the routine for node $N_{\ell,z}$ terminates. From that point on, whenever the traversal ends at $N_{\ell,z}$, Algorithm 1 invokes the GUESS subroutine (see Section 3.4).

This procedure may appear counterintuitive, as it intentionally seeks rejections of some specific prices. However, the primary objective of REDUCE is to control the regret that may accumulate once the routine concludes. Specifically, upon termination, both prices p_L and p_U must have been rejected once. Then, due to the Lipschitz continuity of the valuation functions, this restricts the range of future contexts for which the node $N_{\ell,z}$ can be selected, effectively ensuring that $f_b(x) - f_s(x) \leq O(L2^{-\ell})$ for any context x in the region corresponding to $N_{\ell,z}$. Formally

Lemma 1. *Consider a node $N_{\ell,z}$ for which the REDUCE procedure terminated. Then, for any $x \in \text{Region}(N_{\ell,z})$, it holds that $f_b(x) - f_s(x) \leq 6L2^{-\ell}$.*

Although simple, this result is rather surprising, as it reveals a fundamental property of the problem. By viewing the context-generation process as adversarial, the previous result can be interpreted as

follows: if the adversary ‘‘hides’’ the valuations both when p_L and p_U are played, then the Lipschitz constraint forces the adversary to place them close to the diagonal. In particular, the REDUCE procedure may not end for some nodes (indeed, for most nodes it is not expected to do so). However, we show that if the adversary forces termination at a specific node, then across all similar contexts (*i.e.*, nodes in the subtree rooted at that node) the adversary’s gain remains limited. This situation is illustrated in Figure 1, which also serves as intuition about the proof.

Now, we show that the regret incurred while running this routine is suitably upper-bounded. Formally, let $R_{\ell,z}^{\text{REDUCE}} := \mathbb{E}[\mathcal{R}(T_{\ell,z}^{\text{REDUCE}})]$ be the regret accumulated by the REDUCE routine executed on node $N_{\ell,z}$. We can upper bound this term as follows.

Lemma 2 (Regret due to REDUCE). *For any node $N_{\ell,z}$ with $\ell \geq 1$ it holds $R_{\ell,z}^{\text{REDUCE}} \leq 24L2^{-\ell}$.*

3.4 GUESS ROUTINE

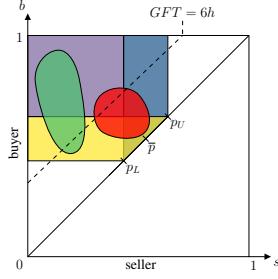


Figure 1: The x -axis (resp., y -axis) represents the seller’s (resp., buyer’s) valuations. The blue (resp., yellow) region denotes valuations for which p_L (resp., p_U) is rejected. Lemma 1 states that if there are valuations both in the yellow and blue one, then there cannot be any with $\text{GFT} \geq 6h = 6L2^{-\ell}$. Let the green and red regions be the image of $\text{Region}(N_{\ell,z})$ under (f_s, f_b) , then Lemma 1 forbids the red region, but allows the green one.

Lemma 3 (Regret due to GUESS). *Let $x \in \text{Region}(N_{\ell,z})$. For any $\epsilon > 0$, consider posting prices uniformly at random from $\mathcal{P}_{\ell,z} = [\bar{p} - L2^{-(\ell-1)}, \bar{p} + L2^{-(\ell-1)}]_{\epsilon}$, where $\bar{p} = p_{\ell-1, \text{pred}(\ell,z)}$ is the marking prices of the parent node. Then the regret incurred by the GUESS procedure is at most*

$$R_{z,\ell}^{\text{GUESS}} \leq \epsilon \mathbb{E}[|T_{\ell,z}|] + 24 \frac{L^2 2^{-2\ell}}{\epsilon}.$$

The intuition about this result is simple: if the trades have small (*i.e.*, less than ϵ) GFT then the total regret is at most ϵ per round. On the other hand, if the GFT is large (*i.e.*, greater than ϵ) then an ϵ -grid over the possible prices is enough to win a trade with probability $\Omega(\epsilon^2/L)$, which happens on average after $O(L2^{-\ell}/\epsilon)$ turns. In each of these turns we are only loosing $O(L2^{-\ell})$ because they take place after the REDUCE procedure of that node.

3.5 REGRET ANALYSIS

By using the results of the previous sections, and for an appropriate choice of ϵ , we can now upper bound the overall regret incurred by the algorithm.

Theorem 1. *For any $T > 2^d$ and any pair of L -Lipschitz valuation functions, Algorithm 1 with $\epsilon = LT^{-1/d}$ guarantees $R_T = O(L \log_2(T) T^{(d-1)/d})$.*

The proof, provided in the appendix, begins by partitioning the rounds $[T]$ across nodes according to Equation (2). For each node, the rounds are then further subdivided based on the phase of the

algorithm. Intuitively, for each round, we look at the node “from which we post a price”. Then, such rounds are partitioned based on whether the algorithm invoked REDUCE, GUESS, or if the node is a marked leaf, as specified in Line 13 of Algorithm 1. Each resulting term can be suitably upper-bounded by applying the previous lemmas.

4 UNKNOWN LIPSCHITZ CONSTANT: A MULTI-SCALE APPROACH

The algorithm presented in Section 3.2, heavily relies on the knowledge of the Lipschitz constant L . The exact knowledge of this constant is crucial in proving the correctness of the algorithm. More precisely, we exploited several Lipschitzness arguments to bound valuations. However, in many real-world scenarios it may be unrealistic to assume that L is known. In this section, we present an algorithm that extends the one introduced in Section 3.2, removing the need for a priori knowledge of L . This comes at the cost of only a slight deterioration in guarantees, yielding a regret of $\tilde{O}(L^2 T^{(d-1)/d})$ instead of $\tilde{O}(LT^{(d-1)/d})$.

We rely on geometric (*i.e.*, multi-scale) extensions of the REDUCE and GUESS subroutines, which we call GEOMETRICREDUCE and GEOMETRICGUESS, respectively. In particular, for each node $N_{\ell,z}$, we define a sequence of geometrically increasing Lipschitz constants $\tilde{L}_{\ell,z}^j = O(2^j)$. In GEOMETRICREDUCE we simply perform the REDUCE routines for each scale j . Crucially, terminating this procedure bounds the GFT in all subsequent nodes by $O(L2^{-\ell})$ (similarly to Lemma 1), but without requiring the knowledge of L .

The construction of GEOMETRICREDUCE is more involved, and The reason is clear if we think about the analysis of GUESS. The analysis relied on a simple fact: for all trades with GFT greater than ϵ , we have a probability of at least $\Omega(\epsilon 2^\ell / L)$ of marking the node, and thus we expect to finish in $O(L2^{-\ell}/\epsilon)$ rounds. However, it is also possible that the process never terminates, since there may be many trades happening with GFT less than ϵ . Therefore, without knowledge of the constant L , the algorithm may spend many rounds working at an excessively small scale, potentially leading to the loss of trades with large GFT. Moreover, it becomes impossible to distinguish whether we are losing these trades because they have small GFT, or simply because the chosen scale is too small for the real Lipschitz constant. The solution is to carefully increase the scale of our Lipschitz constant (more precisely, logarithmically), and to draw a price from the corresponding grid uniformly at random. This process is guaranteed to finish in expectation after $O(L2^{-\ell}/\epsilon)$ rounds against valuations with GFT at least ϵ , and does not require knowing L beforehand. The overall algorithm is detailed in Section C and provides the following regret guarantees, which are only a $\tilde{O}(\log_2(T)L)$ factor greater than the guarantees of Theorem 1.

Theorem 2. *Consider the nonparametric contextual bilateral trade problem, with one-bit feedback and strong budget balance with L -Lipschitz valuation functions. Then there is an algorithm that, without taking as input L , guarantees a regret of $R_T = O(\log_2(T)^2 L^2 T^{(d-1)/d})$.*

5 LOWER BOUND

In this section, we proceed to prove a matching lower bound that proves the tightness of our algorithm. In particular, we will prove the following theorem.

Theorem 3. *Any randomized algorithm under strong-budget balance, d -dimensional context vectors, and L -Lipschitz value functions, suffers a regret of $\Omega(LT^{\frac{d-1}{d}})$, provided that $T > (4L)^d$. This also holds under full feedback.*

From a technical perspective, the proof of Theorem 3 relies on Yao’s weak minimax principle (Yao, 1977), which lower bounds the regret with the expected performance of a deterministic algorithm against a distribution of instances. Formally:

Lemma 4 (Yao’s weak minimax principle). *For any deterministic algorithm generating $\{p_t\}_{t \in [T]}$, and any distribution \mathcal{D} , we have that $R_T \geq \mathbb{E}_{(f_s, f_b) \sim \mathcal{D}} \left[\sum_{t \in [T]} [f_b(x_t) - f_s(x_t)]^+ - \text{GFT}_t(p_t) \right]$.*

To construct a hard distribution over instances, and consequently over valuation functions f_s and f_b , we start by considering all 2^T possible assignments of a discrete set of T context vectors to two

disjoint valuations (*i.e.*, such that no price results in both trades being accepted). We then extend these discrete assignments to Lipschitz-continuous functions over the entire context space. The distribution is then taken as the uniform one over all 2^T possible valuation functions. In particular, we build a distribution \mathcal{D} over instances such that, for any deterministic algorithm, it holds

$$\mathbb{E}_{(f_s, f_b) \sim \mathcal{D}} \left[\sum_{t \in [T]} [f_b(x_t) - f_s(x_t)]^+ - \text{GFT}_t(p_t) \right] \geq \tilde{\Omega}(LT^{\frac{d-1}{d}}).$$

We define the standard δ -grid \mathcal{G} on the $[0, 1]^d$ dimensional hypercube, *i.e.*, $\mathcal{G} = [0, \delta, \dots, 1]^d$. The set \mathcal{G} is the set of possible contexts and, by fixing $\delta = T^{-1/d}$, we have that $|\mathcal{G}| = T$. Fix any ordering of contexts x_1, \dots, x_T over \mathcal{G} . We can define two valuations $(s_0, b_0) = (\frac{1}{2} - \gamma, \frac{1}{2} - \epsilon)$ and $(s_1, b_1) = (\frac{1}{2} + \epsilon, \frac{1}{2} + \gamma)$, where $\gamma > \epsilon > 0$ are two parameters to be specified later. Notice that there is no price p such that p wins both trades, in particular when the valuations are (s_0, b_0) the only prices that get the trade accepted are in the set $[\frac{1}{2} - \gamma, \frac{1}{2} - \epsilon]$, while under valuations (s_1, b_1) prices that make the trade happen are in $[\frac{1}{2} + \epsilon, \frac{1}{2} + \gamma]$. For the same reason, we know that we can focus on deterministic algorithms that post prices only within one of these two intervals, since they yield GFT greater than that of any other price.

Next, for any $L > 0$, we build a family of L -Lipschitz functions from $[0, 1]^d$ to $\{(s_0, b_0), (s_1, b_1)\}$. In order to build our set of instances, consider a binary vector $h \in \mathcal{H}_T$ of length T . Given a context x_i , the i -th bit of h determines the valuations for such a context. In particular, h selects (s_0, b_0) if $h_i = 0$, and (s_1, b_1) if $h_i = 1$. For each of these binary vectors $h \in \mathcal{H}_T$, let $f^h = (f_s^h, f_b^h) : \mathcal{G} \rightarrow \mathbb{R}^2$ be the functions satisfying the corresponding assignment, *i.e.*, $f_s^h(x_i) = s_{h_i}$ and $f_b^h(x_i) = b_{h_i}$. Note that for all $x \neq x' \in \mathcal{G}$ and any $h \in \mathcal{H}$, we have $|f_s^h(x) - f_s^h(x')| \leq \gamma - \epsilon$, and the same holds for f_b^h . By taking $\epsilon = \frac{\gamma}{2}$ and $\gamma = 2LT^{-1/d}$, and by noting that $\inf_{x \neq x', x \in \mathcal{G}} \|x - x'\|_\infty \geq \delta = T^{-1/d}$ we have that f_s^h and f_b^h are L -Lipschitz in \mathcal{G} .¹ Now, we need to extend the functions over the entire context domain $[0, 1]^d$, which can be done in a black box way by the McShane extension theorem that we recall here

Theorem 4 (McShane's extension (McShane, 1934)). *Let $X \subset \mathbb{R}^d$ be a closed set, and f be a \mathbb{R} -valued, L -Lipschitz function on X . Then there exists an extension $\tilde{f} : \mathbb{R}^d \rightarrow \mathbb{R}$ that agrees with f on X and is L -Lipschitz on \mathbb{R}^d .*²

We apply this on each f^h and take \mathcal{D} as the uniform distribution over $\mathcal{F} = \{\tilde{f}^h\}_{h \in \mathcal{H}^T}$. This distribution has the intuitive property that renders its history unhelpful in predicting future valuations.

Lemma 5. *For any $t \in [T]$, the distribution of (s_t, b_t) under \mathcal{D} , conditioned on the filtration up to time $t - 1$, is uniform over $\{(s_0, b_0), (s_1, b_1)\}$.*

We start by observing that the GFT that can be extracted from both (s_0, b_0) and (s_1, b_1) is $\gamma - \epsilon = \gamma/2 > 0$. Therefore, the benchmark will gain $\mathbb{E}_{(f_s, f_b) \sim \mathcal{D}} \left[\sum_{t \in [T]} [f_b(x) - f_s(x)]^+ \right] = \frac{\gamma}{2}T$. On the other hand, by Lemma 5, for any $t \in [T]$ the best expected GFT achievable by any deterministic algorithm is bounded from above by $\frac{\gamma - \epsilon}{2} = \frac{\gamma}{4}$. Since we set $\gamma = 2LT^{-1/d}$, we obtain that

$$R_T \geq \frac{T}{2}(\gamma - \epsilon) = \frac{L}{2}T^{\frac{d-1}{d}},$$

concluding the proof of Theorem 3.

Remark This construction is reminiscent of that of Lu et al. (2009), who establish a rate of $\Omega(T^{(d+1)/(d+2)})$ (instantiated, as an example, in the finite MAB setting), which is higher than both our lower bound of $\Omega(T^{(d-1)/d})$ and our upper bound of $\tilde{O}(T^{(d-1)/d})$, despite our setting appearing more challenging. However, this is not unexpected, since the proof of Lu et al. (2009) leverages a significantly sparser grid in the context space ($O(T^{-1/(d+2)})$ instead of $O(T^{-1/d})$). They are allowed to do this since in their setting contexts can be repeated without incurring zero regret, as noise allows suboptimal actions to remain indistinguishable for longer. In contrast, in our setting, the reward function is identical whenever the same context is observed.

¹Note that we need $T > (4L)^d$ for the valuations to be well defined (*i.e.*, in $[0, 1]^2$).

²In our proof we only care about the existence of such an extension, however, the proof of the theorem is constructive and defines \tilde{f} as $\tilde{f}(x) = \inf_{x' \in X} (f(x) + L\|x - x'\|_\infty)$, (Weaver, 1999, see *e.g.*, Theorem 1.33).

6 CONCLUSIONS AND FUTURE WORK

There are many interesting possible directions for future works. Similar to [Slivkins \(2014\)](#), one can consider a general context space and derive rates that depend on its packing or covering dimension. Alternatively, following [Kuzborskij & Cesa-Bianchi \(2020\)](#); [Cesa-Bianchi et al. \(2017\)](#), one may exploit additional structure in the function class mapping contexts to valuations. From a bilateral trade perspective, other notions of budget balance are also relevant, such as weak and global budget balance ([Cesa-Bianchi et al., 2024](#); [Bernaconi et al., 2024](#)). Under these notions, the learner may post different prices to the buyer and the seller, subject only to the constraint that the market is not subsidized in each round and across all rounds, respectively. Moreover, our tight rates are attainable in part because our model does not allow for noise, unlike some general nonparametric bandit models ([Lu et al., 2009](#); [Slivkins, 2014](#)). The question of what rates can be achieved under the noisy valuations remains open. Finally, [Cossen et al. \(2026\)](#) also studies profit maximization. It would be interesting to study the profit maximization setting with general Lipschitz contexts.

ACKNOWLEDGMENT

The work of MB, and AC was partially funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

This work is supported by an ERC grant (Project 101165466 — PLA-STEER).

REPRODUCIBILITY STATEMENT

Full proofs are included in the Appendix, and assumptions are explicitly stated in the model.

REFERENCES

Kareem Amin, Afshin Rostamizadeh, and Umar Syed. Repeated contextual auctions with strategic buyers. *Advances in Neural Information Processing Systems*, 27, 2014.

Yossi Azar, Amos Fiat, and Federico Fusco. An α -regret analysis of adversarial bilateral trade. In *Advances in Neural Information Processing Systems*, 2022.

François Bachoc, Tommaso Cesari, and Roberto Colombo. A contextual online learning theory of brokerage. *arXiv preprint arXiv:2407.01566*, 2024.

Martino Bernasconi, Matteo Castiglioni, Andrea Celli, and Federico Fusco. No-regret learning in bilateral trade via global budget balance. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pp. 247–258, 2024.

Liad Blumrosen and Yehonatan Mizrahi. Approximating gains-from-trade in bilateral trading. In *WINE*, volume 10123 of *Lecture Notes in Computer Science*, pp. 400–413. Springer, 2016.

Natasa Bolić, Tommaso Cesari, and Roberto Colombo. An online learning theory of brokerage. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*, pp. 216–224, 2024.

Johannes Brustle, Yang Cai, Fa Wu, and Mingfei Zhao. Approximating gains from trade in two-sided markets via simple mechanisms. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pp. 589–590, 2017.

Nicolò Cesa-Bianchi, Pierre Gaillard, Claudio Gentile, and Sébastien Gerchinovitz. Algorithmic chaining and the role of partial feedback in online nonparametric learning. In *Conference on Learning Theory*, pp. 465–481. PMLR, 2017.

Nicolò Cesa-Bianchi, Tommaso Cesari, Roberto Colombo, Federico Fusco, and Stefano Leonardi. Bilateral trade: A regret minimization perspective. *Mathematics of Operations Research*, 49(1): 171–203, 2024.

Houshuang Chen, Yaonan Jin, Pinyan Lu, and Chihao Zhang. Tight regret bounds for fixed-price bilateral trade. *arXiv preprint arXiv:2504.04349*, 2025.

Ningyuan Chen and Guillermo Gallego. Nonparametric pricing analytics with customer covariates. *Operations Research*, 69(3):974–984, 2021.

Maxime C Cohen, Ilan Lobel, and Renato Paes Leme. Feature-based dynamic pricing. *Management Science*, 66(11):4921–4943, 2020.

Romain Cosson, Federico Fusco, Anupam Gupta, Stefano Leonardi, Renato Paes Leme, and Matteo Russo. Contextual online bilateral trade. 2026. Private communication, to appear on the arXiv.

Yuan Deng, Jieming Mao, Balasubramanian Sivan, and Kangning Wang. Approximately efficient bilateral trade. In *STOC*, pp. 718–721. ACM, 2022.

Solenne Gaucher, Martino Bernasconi, Matteo Castiglioni, Andrea Celli, and Vianney Perchet. Feature-based online bilateral trade. In *The Thirteenth International Conference on Learning Representations*, 2025.

Elad Hazan and Nimrod Megiddo. Online learning with prior knowledge. In *International Conference on Computational Learning Theory*, pp. 499–513. Springer, 2007.

Zi Yang Kang, Francisco Pernice, and Jan Vondrák. Fixed-price approximations in bilateral trade. In *SODA*, pp. 2964–2985. SIAM, 2022.

R. Kleinberg and T. Leighton. The value of knowing a demand curve: bounds on regret for online posted-price auctions. In *44th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 594–605, 2003.

Ilja Kuzborskij and Nicolo Cesa-Bianchi. Locally-adaptive nonparametric online learning. *Advances in Neural Information Processing Systems*, 33:1679–1689, 2020.

Allen Liu, Renato Paes Leme, and Jon Schneider. Optimal contextual pricing and extensions. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 1059–1078. SIAM, 2021.

Ilan Lobel, Renato Paes Leme, and Adrian Vladu. Multidimensional binary search for contextual decision-making. *Operations Research*, 66(5):1346–1361, 2018.

Tyler Lu, Dávid Pál, and Martin Pál. Showing relevant ads via context multi-armed bandits. In *Proceedings of AISTATS*, 2009.

Anna Lunghi, Matteo Castiglioni, and Alberto Marchesi. Better regret rates in bilateral trade via sublinear budget violation. *arXiv preprint arXiv:2507.11419*, 2025.

Jieming Mao, Renato Paes Leme, and Jon Schneider. Contextual pricing for lipschitz buyers. *Advances in Neural Information Processing Systems*, 31, 2018.

R Preston McAfee. The gains from trade under fixed price mechanisms. *Applied economics research bulletin*, 1(1):1–10, 2008.

EJ McShane. Extension of range of functions. *Bulletin of the American Mathematical Society*, 40(12):837–842, 1934.

Roger B Myerson and Mark A Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of economic theory*, 29(2):265–281, 1983.

Alexander Rakhlin and Karthik Sridharan. Online nonparametric regression with general loss functions. *arXiv e-prints*, pp. arXiv–1501, 2015.

Aleksandrs Slivkins. Contextual bandits with similarity information. *Journal of Machine Learning Research*, 15:2533–2568, 2014.

Matilde Tullii, Solenne Gaucher, Nadav Merlis, and Vianney Perchet. Improved algorithms for contextual dynamic pricing. *Advances in Neural Information Processing Systems*, 37:126088–126117, 2024.

William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1):8–37, 1961.

Nik Weaver. *Lipschitz Algebras*. WORLD SCIENTIFIC, 1999. doi: 10.1142/4100. URL <https://www.worldscientific.com/doi/abs/10.1142/4100>.

Andrew Chi-Chin Yao. Probabilistic computations: Toward a unified measure of complexity. In *18th Annual Symposium on Foundations of Computer Science (sfcs 1977)*, pp. 222–227. IEEE Computer Society, 1977.

A OMITTED PROOFS FROM SECTION 3 (UPPER BOUND)

Lemma 1. Consider a node $N_{\ell,z}$ for which the REDUCE procedure terminated. Then, for any $x \in \mathbf{Region}(N_{\ell,z})$, it holds that $f_b(x) - f_s(x) \leq 6L2^{-\ell}$.

Proof. Let $p_L = \Pi_{[0,1]}(\bar{p} - L2^{-(\ell-1)})$ and $p_U = \Pi_{[0,1]}(\bar{p} + L2^{-(\ell-1)})$, where $\bar{p} = p_{\ell-1,\mathbf{pred}(\ell,z)}$ is the price with which we marked the parent node.

If REDUCE terminated for node $N_{\ell,z}$, then Line 7 of Algorithm 2 has been executed exactly twice. We denote by $t_L \in [T]$ the first time in which Line 7 is executed (i.e., the round in which we observe a trade rejected for p_L), and analogously we denote by $t_U \in [T]$ the second time in which Line 7 is executed (i.e., the round in which we observe a trade rejected at price p_U). Let $x_L = x_{t_L}$ and $x_U = x_{t_U}$.

At the end of REDUCE we have two pairs (p_L, x_L) and (p_U, x_U) such that

$$f_s(x_L) > p_L \text{ or } f_b(x_L) < p_L,$$

and

$$f_s(x_U) > p_U \text{ or } f_b(x_U) < p_U.$$

Moreover, by the Lipschitz condition and the fact that the parent node $N_{\ell-1,\mathbf{pred}(\ell,z)}$ is marked (see Definition 1), it follows that for every $x \in \mathbf{Region}(N_{\ell,z}) \subseteq \mathbf{Region}(N_{\ell-1,\mathbf{pred}(\ell,z)})$ we have

$$f_s(x) \leq \min\{1, \bar{p} + L2^{-(\ell-1)}\} = p_U \quad \text{and} \quad f_b(x) \geq \max\{0, \bar{p} - L2^{-(\ell-1)}\}2 = p_L.$$

Then, by rewriting the above for $x_U \in \mathbf{Region}(N_{\ell,z})$, we get that the buyer refused the trade at price p_U , since the seller would have accepted it. Hence, $f_b(x_U) < p_U$. Similarly, we know that it was the seller who rejected the trade at x_L , implying $f_s(x_L) > p_L$. Now consider any $x \in \mathbf{Region}(N_{\ell,z})$. By Lipschitzness we have that

$$|f_s(x) - f_s(x_L)| \leq L2^{-\ell} \quad \text{and} \quad |f_b(x) - f_b(x_U)| \leq L2^{-\ell},$$

from which we can conclude that

$$f_b(x) - f_s(x) \leq f_b(x_U) - f_s(x_L) + 2L2^{-\ell} \leq p_U - p_L + 2L2^{-\ell} = 6L2^{-\ell}.$$

This concludes the proof. \square

Lemma 2 (Regret due to REDUCE). *For any node $N_{\ell,z}$ with $\ell \geq 1$ it holds $R_{\ell,z}^{\text{REDUCE}} \leq 24L2^{-\ell}$.*

Proof. First, observe that in rounds belonging to $T_{\ell,z}^{\text{REDUCE}}$ we accumulate regret only twice, once for each rejection of prices p_L and p_U , since accepted prices yield zero regret. Moreover, by construction we have that $N_{\ell,z}$ is active, so we know that the parent node $N_{\ell-1,\mathbf{pred}(\ell,z)}$ has been marked and that REDUCE has been completed for $N_{\ell-1,\mathbf{pred}(\ell,z)}$. For any $x \in \mathbf{Region}(N_{\ell,z}) \subseteq \mathbf{Region}(N_{\ell-1,\mathbf{pred}(\ell,z)})$ we can apply Lemma 1, and we get that $f_b(x) - f_s(x) \leq 6L2^{-(\ell-1)} = 12L2^{-\ell}$. Therefore, we obtain $R_{\ell,z}^{\text{REDUCE}} \leq 24L2^{-\ell}$, which concludes the proof. \square

Lemma 3 (Regret due to GUESS). *Let $x \in \mathbf{Region}(N_{\ell,z})$. For any $\epsilon > 0$, consider posting prices uniformly at random from $\mathcal{P}_{\ell,z} = [\bar{p} - L2^{-(\ell-1)}, \bar{p} + L2^{-(\ell-1)}]_{\epsilon}$, where $\bar{p} = p_{\ell-1,\mathbf{pred}(\ell,z)}$ is the marking prices of the parent node. Then the regret incurred by the GUESS procedure is at most*

$$R_{z,\ell}^{\text{GUESS}} \leq \epsilon \mathbb{E}[|T_{\ell,z}|] + 24 \frac{L^2 2^{-2\ell}}{\epsilon}.$$

Proof. For each $x \in \mathbf{Region}(N_{\ell,z})$ we have that

$$f_s(x) \leq f_s(\bar{x}) + L2^{-(\ell-1)} \leq \bar{p} + L2^{-(\ell-1)},$$

where $\bar{x} \in \mathbf{Region}(N_{\ell-1,\mathbf{pred}(\ell,z)})$ is the context in the parent node's area under which \bar{p} was accepted. Similarly, we have that

$$f_b(x) \geq f_b(\bar{x}) - L2^{-(\ell-1)} \geq \bar{p} - L2^{-(\ell-1)}.$$

Thus, if $f_b(x) \geq f_s(x)$ then the interval $[f_s(x), f_b(x)]$ is contained in the interval $[\bar{p} - L2^{-(\ell-1)}, \bar{p} + L2^{-(\ell-1)}]$.

For any $t \in T_{\ell,z}^{\text{GUESS}}$ we are going to consider two cases: (i) $f_b(x_t) - f_s(x_t) \leq \epsilon$, and (ii) $f_b(x_t) - f_s(x_t) > \epsilon$. We let

$$T_{\ell,z}^{\text{GUESS},1} := \{t \in T_{\ell,z}^{\text{GUESS}} : f_b(x_t) - f_s(x_t) \leq \epsilon\},$$

and $T_{\ell,z}^{\text{GUESS},2} := T_{\ell,z}^{\text{GUESS}} \setminus T_{\ell,z}^{\text{GUESS},1}$. Then, by linearity of expectation, the regret $R_{z,\ell}^{\text{GUESS}}$ can be decomposed as

$$R_{z,\ell}^{\text{GUESS}} = \mathbb{E} [\mathcal{R}(T_{\ell,z}^{\text{GUESS},1})] + \mathbb{E} [\mathcal{R}(T_{\ell,z}^{\text{GUESS},2})].$$

First, we observe that if $t \in T_{\ell,z}^{\text{GUESS},1}$, then the first term in the decomposition can be bounded above by $\epsilon \mathbb{E}[|T_{\ell,z}|]$. Indeed, in this case we have $[f_b(x_t) - f_s(x_t)]^+ \leq \epsilon$, $\text{GFT}_t(\cdot \mid f_s(x_t), f_b(x_t)) \geq 0$, and moreover $T_{\ell,z}^{\text{GUESS},1} \subseteq T_{\ell,z}$.

Similarly, the second term can be upper bounded by $\mathbb{E}[|T_{\ell,z}^{\text{GUESS},2}|] \cdot G$, where G is the maximum GFT attainable. Since the REDUCE procedure must have already terminated for $N_{\ell,z}$ before GUESS is executed, we can apply Lemma 1, which yields $G \leq 6L2^{-\ell}$. Therefore,

$$R_{z,\ell}^{\text{GUESS}} \leq \epsilon \mathbb{E}[|T_{\ell,z}|] + 6L2^{-\ell} \cdot \mathbb{E}[|T_{\ell,z}^{\text{GUESS},2}|].$$

In order to conclude the proof we have to upper bound $\mathbb{E}[|T_{\ell,z}^{\text{GUESS},2}|]$. Notice that for every $t \in T_{\ell,z}^{\text{GUESS},2}$ we have $f_b(x_t) - f_s(x_t) > \epsilon$. Hence, there exists at least one price p in the discretized grid $\mathcal{P}_{\ell,z}$ such that $f_s(x_t) \leq p \leq f_b(x_t)$. This implies that the probability of terminating GUESS at each round $t \in T_{\ell,z}^{\text{GUESS}}$ is at least $|\mathcal{P}_{\ell,z}|^{-1}$, and therefore

$$\mathbb{E}[|T_{\ell,z}^{\text{GUESS},2}|] \leq |\mathcal{P}_{\ell,z}| \leq \frac{4L2^{-\ell}}{\epsilon}.$$

This concludes the proof. \square

Theorem 1. *For any $T > 2^d$ and any pair of L -Lipschitz valuation functions, Algorithm 1 with $\epsilon = LT^{-1/d}$ guarantees $R_T = O(L \log_2(T) T^{(d-1)/d})$.*

Proof. For a node $N_{\ell,z}$, let $T_{\ell,z}^{\text{LEAF}} \subseteq [T]$ be the set of rounds in which Line 13 of Algorithm 1 is executed for that node (i.e., the number of times a leaf node is played after being marked). Moreover, let

$$R_{\ell,z}^{\text{LEAF}} := \mathbb{E} [\mathcal{R}(T_{\ell,z}^{\text{LEAF}})].$$

First, following Equation (2), we decompose the overall regret into the contributions incurred at each level and at the nodes within that level:

$$R_T = \sum_{\ell=0}^H \sum_{z \in Z_\ell} R_{\ell,z} = \sum_{\ell=0}^H \sum_{z \in Z_\ell} (R_{\ell,z}^{\text{REDUCE}} + R_{\ell,z}^{\text{GUESS}} + R_{\ell,z}^{\text{LEAF}})$$

where the second equality come from the fact that $T_{\ell,z}^{\text{REDUCE}}, T_{\ell,z}^{\text{GUESS}}, T_{\ell,z}^{\text{LEAF}}$ form a partition of $T_{\ell,z}$ for each level $\ell \in \{0, 1, \dots, H\}$ and node $z \in Z_\ell$.

REDUCE Now we are going to analyze the first term. By Lemma 2 we have that $R_{\ell,z}^{\text{REDUCE}} \leq 24 \cdot L2^{-\ell}$. Moreover, we recall that $H = \lfloor \log_2 T \rfloor$ and $|Z_\ell| = 2^{d\ell}$. Thus

$$\begin{aligned} \sum_{\ell=0}^H \sum_{z \in Z_\ell} R_{\ell,z}^{\text{REDUCE}} &\leq 24 \cdot L \sum_{\ell=0}^{\lfloor \log_2(T) \rfloor} 2^{d\ell-\ell} \\ &\leq \frac{24L \log T}{d} 2^{dH} \cdot 2^{-H} \\ &\leq \frac{48L \log T}{d} T^{1-\frac{1}{d}}, \end{aligned}$$

where in the last inequality we used that that $\log_{2^d}(T) = \frac{1}{d} \log(T)$ and that $\ell \mapsto 2^{d\ell-\ell}$ is monotone so that $\sum_{\ell=0}^H 2^{d\ell-\ell} \leq H2^{dH-H}$.

GUESS Similarly, for the second term, we are going to proceed as follows:

$$\begin{aligned}
\sum_{\ell=0}^H \sum_{z \in Z_\ell} R_{\ell,z}^{\text{GUESS}} &\leq \sum_{\ell=0}^H \sum_{z \in Z_\ell} \left(\epsilon \mathbb{E}[|T_{\ell,z}|] + 24 \frac{L^2 2^{-2\ell}}{\epsilon} \right) && \text{(Lemma 3)} \\
&= \epsilon T + \frac{24L^2}{\epsilon} \sum_{\ell=0}^H 2^{d\ell-2\ell} && (|Z_\ell| = 2^{d\ell}) \\
&\leq \epsilon T + \frac{24L^2 \log T}{\epsilon d} 2^{d \log_2(T)} \cdot 2^{-2 \log_2(T)} \\
&\leq \epsilon T + \frac{24L^2 \log T}{\epsilon d} T^{1-\frac{2}{d}}.
\end{aligned}$$

Finally, by choosing $\epsilon = LT^{-\frac{1}{d}}$ we obtain

$$\sum_{\ell=0}^H \sum_{z \in Z_\ell} R_{\ell,z}^{\text{GUESS}} \leq LT^{1-\frac{1}{d}} + \frac{24L}{d} \log(T) T^{1-\frac{1}{d}} \leq \frac{d+24}{d} L \log(T) T^{1-\frac{1}{d}}.$$

LEAF To analyze the last term, we recall that $T_{\ell,z}^{\text{LEAF}}$ is non-empty only if $\ell = H$ (i.e., it is a leaf node) and the node was already marked. Thus we can use Lemma 1 which shows that

$$R_{\ell,z}^{\text{LEAF}} \leq 6L 2^{-\log_2(T)} \mathbb{E}[T_{\ell,z}^{\text{LEAF}}] \leq 6LT^{-1/d} \mathbb{E}[T_{\ell,z}].$$

Therefore, by restricting our attention to the H -th layer we get

$$\begin{aligned}
\sum_{\ell=0}^H \sum_{z \in Z_\ell} R_{\ell,z}^{\text{LEAF}} &= \sum_{z \in Z_H} R_{H,z}^{\text{LEAF}} \\
&\leq \sum_{z \in Z_H} 6LT^{-1/d} \mathbb{E}[T_{H,z}] \\
&= 6LT^{\frac{d-1}{d}}.
\end{aligned}$$

Combining all parts Thus, the total regret can be bounded by

$$R_T \leq \frac{72+7d}{d} LT^{1-\frac{1}{d}} \log(T).$$

concluding the proof. \square

B OMITTED PROOFS FROM SECTION 5 (LOWER BOUND)

Lemma 5. *For any $t \in [T]$, the distribution of (s_t, b_t) under \mathcal{D} , conditioned on the filtration up to time $t-1$, is uniform over $\{(s_0, b_0), (s_1, b_1)\}$.*

Proof. Since we have full feedback, we know the sequence of valuations $\{(s_\tau, b_\tau)\}_{\tau \in [t-1]}$. The set of functions \tilde{f}^h that agree on these valuations for the sequence of observed contexts is 2^{t-1} , we call this set $\tilde{\mathcal{F}}_t$. Out of these functions, half will have $h_t = 0$ and half will have $h_t = 1$, and thus the conditioned probability on (s_t, b_t) is uniform over $\{(s_0, b_0), (s_1, b_1)\}$. \square

C DESCRIPTION OF THE ALGORITHM FROM SECTION 4 AND OMITTED PROOFS

The algorithm's main loop is as the vanilla one introduced in Section 3. In particular, we traverse the tree until we reach an active node (Line 4 of Algorithm 4). Then we perform either the **GEO-METRICREDUCE** of Algorithm 5 or **GEOMETRICGUESS** of Algorithm 6. In particular, for a specific

Algorithm 4 Without knowing the Lipschitz constant**Require:** Horizon T , context dimension d .

```

1: for  $t \in [T]$  do
2:   Observe context  $x_t$ 
3:   Initialize  $\ell \leftarrow 0, z \leftarrow \mathbf{0}_d$ 
4:   while  $p_{\ell,z} \neq \emptyset$  and  $N_{\ell,z}$  not leaf do  $\#$  Traverse tree until an active or leaf node is reached
5:     Find  $z'$  such that  $x_t \in \mathbf{Region}(N_{\ell+1,z'})$ , where  $z' \in \mathcal{C}(\ell, z)$ 
6:      $\ell \leftarrow \ell + 1, z \leftarrow z'$ 
7:   if GEOMETRICREDUCE( $N_{\ell,z}$ ) has not terminated then
8:     Execute one iteration of the main loop of GEOMETRICREDUCE( $N_{\ell,z}$ )
9:   else if GEOMETRICGUESS( $N_{\ell,z}$ ) has not terminated then
10:    Execute one iteration of the main loop of GEOMETRICGUESS( $N_{\ell,z}$ )
11:    if GEOMETRICGUESS( $N_{\ell,z}$ ) terminates returning price  $p$  then
12:      Set node  $N_{\ell,z}$  as marked by setting  $p_{\ell,z} \leftarrow p$ 
13:    else Play  $p_{\ell,z}$   $\#$  If node is a marked leaf play corresponding price

```

Algorithm 5 GEOMETRICREDUCE($N_{\ell,z}$)

```

1: Set  $\bar{p} \leftarrow p_{\ell-1,\mathbf{pred}(\ell,z)}$  (i.e., the price associated to the parent node)  $\#$  Initialization
2: Set  $\tilde{L}^j = L_0 2^j$  for  $j \in \{0, \dots, \bar{j}\}$  with  $\bar{j} = \lceil \ell + \log_2(2L_0) \rceil$  and  $L_0 = 1/2$ 
3: for  $j \in \{0, \dots, \bar{j}\}$  do  $\#$  Main Loop
4:   Set  $p_L^j \leftarrow \Pi_{[0,1]}(\bar{p} - \tilde{L}^j 2^{-(\ell-1)})$ ,  $p_U^j \leftarrow \Pi_{[0,1]}(\bar{p} + \tilde{L}^j 2^{-(\ell-1)})$ 
5:   Initialize  $p \leftarrow p_L^j$ 
6:   repeat
7:     Post price  $p$  and observe  $\mathbb{I}(s_t \leq p \leq b_t)$ 
8:     if  $\mathbb{I}(s_t \leq p \leq b_t) = 0$  then
9:       if  $p = p_L^j$  then set  $p \leftarrow p_U^j$ ; else terminate  $p \leftarrow \emptyset$ 
10:  until  $p = \emptyset$ 

```

node, we continue to perform a step of the main loop of Algorithm 5 until it meets its termination condition, or one step of Algorithm 6, until it terminates. Algorithm 5 simply tries sequentially all scales of \tilde{L}^j as in the plain version. Algorithm 6, on the other hand, at iteration τ , uses the same largest scale permissible at that iteration Line 6 of Algorithm 6, and then draws uniform prices from the associated grid (Line 7).

Lemma 6. *For every $\epsilon > 0$ and for every node $N_{\ell,z}$, there exists an algorithm that does not require knowledge of L such that*

$$R_{\ell,z} = O\left(\ell L 2^{-\ell} + \epsilon \mathbb{E}[T_{\ell,z}] + \frac{L^2 2^{-2\ell}}{\epsilon}\right).$$

Proof. Define $\bar{j} = \lceil \ell + \log_2(2L_0) \rceil$, and for each $j \in \{0, \dots, \bar{j}\}$ define $\tilde{L}^j = L_0 2^j$. Note that $2^{-(\ell-1)} \tilde{L}^j > 1$. For each j we now start GEOMETRICREDUCE on prices $p_L^j = \bar{p} - 2^{-(\ell-1)} \tilde{L}^j$ and $p_U^j = \bar{p} + 2^{-(\ell-1)} \tilde{L}^j$, when we get a rejection both for p_L^j and p_U^j we increase j .

Claim 1. *Once the GEOMETRICREDUCE procedure is concluded, then we know that for all $x \in \mathbf{Region}(N_{\ell,z})$ we have $f_b(x) - f_s(x) \leq 10L2^{-\ell}$.*

Proof of Claim 1. If the GEOMETRICREDUCE has concluded we know that we collected tuples $(x_L^j, p_L^j)_{j=0, \dots, \bar{j}}$ and $(x_U^j, p_U^j)_{j=0, \dots, \bar{j}}$ that were refused.

Now take any $x \in \mathbf{Region}(N_{\ell,z})$, since price (\bar{x}, \bar{p}) was accepted at $N_{\ell-1,\mathbf{pred}(\ell,z)}$ we have that

$$f_s(x) \leq f_s(\bar{x}) + L 2^{-(\ell-1)} \leq \bar{p} + L 2^{-(\ell-1)} = \bar{p} + 2L 2^{-\ell},$$

and

$$f_b(x) \geq f_b(\bar{x}) - L 2^{-(\ell-1)} \geq \bar{p} - L 2^{-(\ell-1)} = \bar{p} - 2L 2^{-\ell}.$$

Algorithm 6 GEOMETRICGUESS($N_{\ell,z}$)

```

1: Let  $\bar{p} = p_{\ell-1, \text{pred}(\ell, z)}$  be the price associated to the parent node  $N_{\ell-1, \text{pred}(\ell, z)}$  # Initialization
2: Set  $\tilde{L}^j = L_0 2^j$  for  $j \in \{0, \dots, \bar{j}\}$  with  $\bar{j} = \lceil \ell + \log_2(2L_0) \rceil$  and  $L_0 = 1/2$ 
3: Set  $\mathcal{P}_{\ell,z}^j = [\bar{p} - \tilde{L}^j 2^{-(\ell-1)}, \bar{p} + \tilde{L}^j 2^{-(\ell-1)}]_\epsilon$  with  $\epsilon = T^{-1/d}$ 
4: Set  $\tau = 0$ 
5: repeat # Main Loop
6:    $j^{(\tau)} := \arg \max \{j \in \{0, \dots, \bar{j}\} : |\mathcal{P}_{\ell,z}^j| \leq \tau\}$ 
7:   Post price  $p \sim \text{Unif}(\mathcal{P}_{\ell,z}^{j^{(\tau)}})$  and observe  $\mathbb{I}(s_t \leq p \leq b_t)$ 
8:    $\tau \leftarrow \tau + 1$ 
9: until  $\mathbb{I}(s_t \leq p \leq b_t) = 1$ 
10: Return  $p$ 

```

We take j' such that $\tilde{L}_{\ell-1,z}^{j'} \geq L$. Then, by plugging in the above $x_L^{j'}$ and $x_U^{j'}$, similarly to the proof of Lemma 1, we know that

$$f_b(x) - f_s(x) \leq f_b(x_U^{j'}) - f_s(x_L^{j'}) + 2L2^{-\ell} \leq p_U^{j'} - p_L^{j'} + 2L2^{-\ell} \leq 4\tilde{L}_{\ell-1,z}2^{-\ell} + 2L2^{-\ell},$$

where the first inequality is by Lipschitzness of f_s, f_b . Moreover, note that $2L \geq \tilde{L}^{j'}$ and thus $f_b(x) - f_s(x) \leq 10L2^{-\ell}$ for all $z \in \text{Region}(N_{\ell,z})$. This concludes the proof of Claim 1. \square

By performing GEOMETRICREDUCE at node $N_{\ell,z}$, we only suffer a regret of at most $2\bar{j} \leq 2(\ell + \log_2(2L_0) + 1) \leq 4\ell$ if $\ell \geq 1$ and $L_0 = 1/2$ (the number of rejected trades), which multiplies the maximum regret suffered, which by Claim 1, is at most $20L2^{-\ell}$, and thus a total regret of $80\ell L2^{-\ell}$.

After this phase, we start using GEOMETRICGUESS in the following multi-scale fashion. Let τ be the “internal time” of the routine. At each execution τ of GEOMETRICGUESS, we consider the largest integer $j^{(\tau)}$ so that $|\mathcal{P}_{\ell,z}^{j^{(\tau)}}| \leq \tau$, where we defined $\mathcal{P}_{\ell,z}^j = [\bar{p} - \tilde{L}^j 2^{-(\ell-1)}, \bar{p} + \tilde{L}^j 2^{-(\ell-1)}]_\epsilon$ and thus $|\mathcal{P}_{\ell,z}^j| = 4L_0 2^{j-1} / \epsilon$. Now, at each τ we draw uniformly at random a price from the corresponding uniform grid $\mathcal{P}_{\ell,z}^{j^{(\tau)}}$. Define τ^* as the smallest integer so that $L_0 2^{j^{(\tau^*)}} \geq L$, and note that $\tau^* \leq \frac{8L2^{-\ell}}{\epsilon}$.

After τ^* turns of GEOMETRICGUESS, we will have that, for all valuations with GFT at least ϵ , there is at least one price in $\mathcal{P}_{\ell,z}^{j^{(\tau^*)}}$ that wins it. We then consider blocks \mathcal{B}_i defined as the intervals $\mathcal{B}_i = [2^i \tau^*, 2^{i+1} \tau^*)$. The probability of winning a trade with GFT larger than ϵ for $\tau \in \mathcal{B}_i$ is $(2^i \tau)^{-1}$, since for $\tau \in \mathcal{B}_i$ the size of $\mathcal{P}_{\ell,z}^{j^{(\tau)}}$ remains unchanged for all $\tau \in \mathcal{B}_i$ and smaller than $2^i \tau^*$. We can define an appropriate Markov chain to compute the expected number of times until we stop.

Claim 2. Consider a Markov chain with states $\{0\} \cup \{N, N+1, \dots\}$ and with transition probabilities $p_s = \mathbb{P}(s+1|s) = 1/(2^k N)$ and $\mathbb{P}(0|s) = 1 - p_s$ for all $k \in \mathbb{N}$ and $s \in [2^k N, \dots, 2^{k+1} N - 1]$. Moreover, let $\mathbb{P}(0|0) = 1$. Then, the expected hitting time of the state 0 starting from state N is at most $4N$.

Proof of Claim 2. The probability of reaching $s = 2^{k+1}N$ from $s = 2^k N$ is $(1 - 1/(2^k N))^{2^k N} \leq e^{-1}$ and the probability of reaching $s = 2^k N$ from $s = N$ is thus at most e^{-k} . The expected number of transitions before hitting 0 is the sum of the expected number of transitions in block \mathcal{B}_i , conditioned on reaching $2^i N$. We call X_i the number of transitions inside of block i and T the hitting time of

$s = 0$. Thus:

$$\begin{aligned}\mathbb{E}[T] &= \sum_{i=0}^{\infty} \mathbb{E}[X_i] = \sum_{i=0}^{\infty} \mathbb{E}[X_i | \text{reach } 2^i N] \mathbb{P}(\text{reach } 2^i N) & (\mathbb{E}[X_i | \overline{\text{reach } 2^i N}] = 0) \\ &\leq \sum_{i=0}^{\infty} 2^i N e^{-i} & (\mathbb{P}(\text{reach } 2^i N) \leq e^{-i} \text{ and } X_i \leq 2^i N) \\ &= N \sum_{i=0}^{\infty} \left(\frac{2}{e}\right)^i\end{aligned}$$

which is equal to $N \frac{e}{e-2} \leq 4N$ concluding the proof of Claim 2. \square

We can then use Claim 2 to conclude that the expected number of turns of GEOMETRICGUESS before observing a success is at most $4\tau^*$, and each turn we experience a regret of only $20L2^{-\ell}$ thanks to Claim 1.

Thus, for each $\tau \leq \tau^*$ we only have a regret of $10L2^{-\ell}$ per turn, and after (for $\tau \geq \tau^*$), we have a regret of at most

$$\epsilon \mathbb{E}[T_{\ell,z}] + 80\tau^* L 2^{-\ell} \leq \epsilon \mathbb{E}[T_{\ell,z}] + \frac{640L^2 2^{-2\ell}}{\epsilon},$$

as in the proof of Lemma 3.

The total regret for marking a node is

$$R_{\ell,z} \leq \underbrace{80\ell L 2^{-\ell}}_{\text{GEOMETRICREDUCE}} + \underbrace{10L2^{-\ell}}_{\tau \leq \tau^*} + \underbrace{\epsilon \mathbb{E}[T_{\ell,z}] + \frac{640L^2 2^{-2\ell}}{\epsilon}}_{\tau > \tau^*} \quad (3)$$

$$= O\left(\ell L 2^{-\ell} + \epsilon \mathbb{E}[T_{\ell,z}] + \frac{L^2 2^{-2\ell}}{\epsilon}\right). \quad (4)$$

\square

Lemma 7. *For every marked leaf node $N_{\ell,z}$ the regret $R_{\ell,z}$ is at most $O(L2^{-\ell} \mathbb{E}[T_{\ell,z}])$.*

Proof. Since we reached a leaf node, we must have concluded GEOMETRICREDUCE in the parent node and thus $R_{\ell,z}^{\text{LEAF}} \leq 20L2^{-\ell} \mathbb{E}[T_{\ell,z}]$ by Claim 1. \square

Theorem 2. *Consider the nonparametric contextual bilateral trade problem, with one-bit feedback and strong budget balance with L -Lipschitz valuation functions. Then there is an algorithm that, without taking as input L , guarantees a regret of $R_T = O(\log_2(T)^2 L^2 T^{(d-1)/d})$.*

Proof. The regret of every non-marked node is upper bounded by Lemma 6, while the leaves once marked have a regret upper bounded by Lemma 7. Thus

$$\begin{aligned}R_T &\leq \sum_{\ell=0}^H \sum_{z \in Z_{\ell}} O\left(\ell L 2^{-\ell} + \epsilon \mathbb{E}[T_{\ell,z}] + \frac{L^2 2^{-2\ell}}{\epsilon}\right) + \sum_{z \in Z_H} O(L 2^{-H} \mathbb{E}[T_{z,H}]) \\ &\leq L \sum_{\ell=0}^H O(\ell 2^{-\ell} 2^{d\ell}) + O(\epsilon T) + \frac{L^2}{\epsilon} \sum_{\ell=0}^H O(2^{d\ell} 2^{-2\ell}) + O(L T 2^{-H}) \\ &\leq O(L H^2 T^{(d-1)/d}) + O(\epsilon T) + O\left(\frac{L^2}{\epsilon} H T^{(d-2)/d}\right) + O(L T^{(d-1)/d}).\end{aligned}$$

then, by choosing $\epsilon = T^{-1/d}$, and noting that $H \leq \log_2(T)/d$ we obtain that

$$R_T \leq O(\log_2(T)^2 L^2 T^{(d-1)/d}),$$

concluding the proof. \square

D EXPERIMENTAL EVALUATION

For completeness, we provide an empirical validation on synthetic data to confirm our theoretical findings.

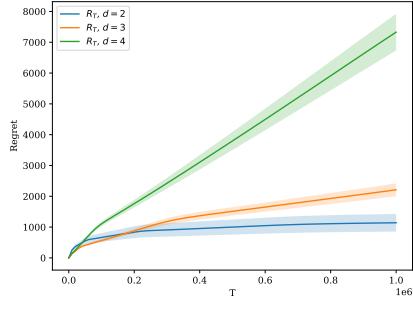


Figure 2: Average regret (together with 95% confidence intervals) in the quadratic setting.

harder as the context dimension d increases, and the behavior of average regrets is consistent with the rates obtained in Theorem 1.

We consider a setting in which valuations are quadratic functions of the context. We run our algorithm on 30 independent instances built as follows. For each run we start by uniformly sampling two matrices $A, B \in [-1, 1]^{d \times d}$. Then, for each t we sample uniformly at random in $[0, 1]^d$ a d -dimensional context x_t . We compute valuations as $x_t^\top Ax_t$ and $x_t^\top Bx_t$ up to normalization to have final valuations belong to the unit square. In particular, the valuations are computed as

$$f_s(x_t) = \frac{x_t^\top Ax_t + d^2}{d^2} \quad f_b(x_t) = \frac{x_t^\top Bx_t + d^2}{d^2}.$$

We test our algorithms on $d = \{2, 3, 4\}$, $T = 10^6$, and run 30 repetitions for each choice of d . The results are reported in Figure 2, and essentially confirm our theoretical findings. As expected, the learning problem becomes

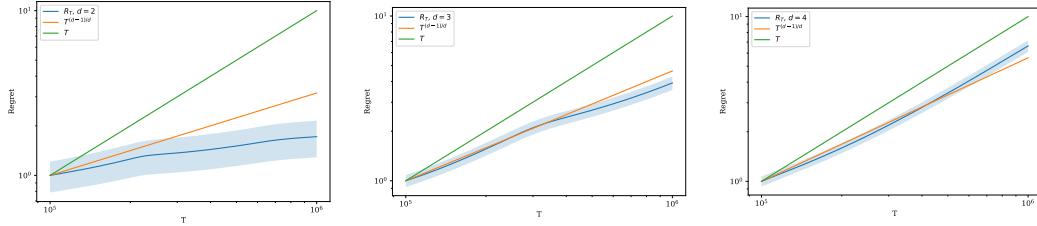


Figure 3: Comparison between average regret and theoretical bounds.

In Figure 3 we report the logarithmic plot of the regret. In particular, for the last 10^4 steps, we plot $\log(t)$ vs $\log(R_t)$ for $d \in \{2, 3, 4\}$, along with the slopes associated with the theoretical bounds (i.e., $t^{(d-1)/d}$) and the linear slope t (we normalize the lines so that they all start from 1 so that we can compare more easily their slopes). We can see that the regret slope matches the theoretical guarantees. In particular, the regret grows at most as fast as $O(t^{(d-1)/d})$ for all d (and for $d = 2$ it appears to be growing even at a better rate).