

SCM class	Upper bounds	Class effect	Lower bounds
General	$\mathcal{O}\left(Kd^{L-1}\sqrt{T\log T}\varpi_{\mathcal{F}}\right)$	$\varpi_{\mathcal{F}} = \dim(\mathcal{F}) \left[1 + \frac{\log N + \log \text{cn}(\mathcal{F})}{\log T}\right]$	–
Linear	$\mathcal{O}\left(Kd^L\sqrt{T\log T}\varpi_{\mathcal{L}}\right)$	$\varpi_{\mathcal{L}} = \log T \left[1 + \frac{\log N}{d\log T}\right]$	$\Omega\left(Kd^{\frac{L}{2}-1}\sqrt{T}\right)$
Polynomial	$\mathcal{O}\left(Kd^{L+1}\sqrt{T\log T}\varpi_{\mathcal{P}}\right)$ for $p = 2$	$\varpi_{\mathcal{P}} = \log T \left[1 + \frac{\log N}{d\log T}\right]$	$\Omega\left(K\sqrt{T}\right)$
Neural network (width $s \geq d$)	$\mathcal{O}\left(\bar{K}d^{L-1}\sqrt{T\log T}\varpi_{\mathcal{N}}\right)$	$\varpi_{\mathcal{N}} = s \log T \left[1 + \frac{\log N}{d\log T}\right]$	$\Omega\left(Kd^{\frac{L}{2}-1}\sqrt{T}\right)$

Table 1: Comparison of upper and lower bounds on regret for general SCMs and special cases.

bandit (GCB) algorithm to design interventions in this general setting. Lower and upper bounds are established for the cumulative regret of the algorithm, that is, the regret accumulated over time with respect to the best intervention in hindsight. The regret results are subsequently specialized and refined for three main classes of SCMs (linear, polynomial, and neural networks).

Key observations. The SCMs are assumed to belong to general classes of functions \mathcal{F} (with mild regularity conditions). We provide an overview below of how the regret bounds are affected by graph topology, SCMs, and intervention models. The regret bounds are summarized in Table 1.

- **SCM class:** The achievable regret depends on two intrinsic measures associated with the SCM class \mathcal{F} , namely the eluder dimension of \mathcal{F} ($\dim(\mathcal{F})$) and the covering number of \mathcal{F} ($\text{cn}(\mathcal{F})$).
- **Graph topology:** The regret lower and upper bounds have diminishing dependence on the graph size N and instead depend on the graph through two key parameters: the maximum in-degree (d) and the length of the longest causal path (L).
- **Intervention granularity:** The regret lower and upper bounds are *independent* of the cardinality of the parameters specifying the interventions.
- **Lower bound:** A general minimax lower bound on regret is given without restrictions on \mathcal{F} . This lower bound is further refined for special classes of functions.
- **Dependence on horizon:** The regret lower and upper bounds both scale with the time horizon T with the rate $\tilde{\mathcal{O}}(\sqrt{T})$.

Relevant literature. The existing literature on causal bandits can be categorized along four dimensions: (i) whether the causal graph is known or unknown, (ii) whether the pre- and post-intervention probability distributions of the causal graph are known or unknown, (iii) the class of SCMs adopted, and

(iv) the nature of interventions (e.g., granularity). Most existing studies focus on the first two dimensions, with the last two aspects highly under-investigated. Specifically, based on the availability of the causal graph and the interventional distributions, there are broadly four distinct categories of studies. The initial study by Lattimore et al. (2016) assumes knowledge of both the causal graph and all the interventional distributions. Similar assumptions are adopted by Bareinboim et al. (2015); Sen et al. (2017); Lu et al. (2020); Nair et al. (2021). There has been progress towards relaxing these assumptions. Specifically, the study by Bilodeau et al. (2022) assumes that the causal graph structure is unknown while the interventional distributions are known. In contrast, settings in which the causal graph is known but the interventional distributions are unknown are studied by Yabe et al. (2018); Maiti et al. (2022); Xiong and Chen (2023); Feng and Chen (2023); Sawarni et al. (2023). Finally, the case of unknown graph structure and unknown interventional distributions is investigated by Lu et al. (2021); De Kroon et al. (2022); Bilodeau et al. (2022); Feng et al. (2023); Konobeev et al. (2023); Malek et al. (2023).

These studies, generally, focus on idealized models that involve (i) ideal *do* interventions, which completely remove parental causal effects and can be restrictive in many applications, and (ii) simple classes of structural causal models. The SCMs investigated include linear (Maiti et al., 2022) or binary generalized linear functions (Xiong and Chen, 2023; Jaber et al., 2020).

Two recent studies by Varici et al. (2023); Sussex et al. (2023) consider *soft* interventions, which, instead of cutting off parental causal effects, only modify the conditional distribution of the nodes. The study by Varici et al. (2023) assumes a *linear* SCM with a *binary* intervention model per node. In contrast, we assume *general* SCMs with an arbitrary *continuum* of interventions. Furthermore, our general model can be specialized to linear SCMs and in this case, our achievable regret bound improves on that of (Varici et al., 2023).

The study by Sussex et al. (2023) investigates the class of *Gaussian processes* within a reproducing kernel Hilbert space (RKHS), assumes controllable variables for specifying interventions, and characterizes achievable regret bounds. In contrast, we focus on general SCMs, and our achievable regret has two significant distinctions. First, the achievable regret of Sussex et al. (2023) scales *linearly* with the graph size N . In contrast, in the most general form, our regret scales with $\sqrt{\log N}$, and more importantly, for most special cases, it scales with $\sqrt{\frac{\log N}{\log T}}$, implying a diminishing dependence on N as the horizon T grows. Furthermore, the regret upper bound in Sussex et al. (2023) relies on the maximum information gain (Srinivas et al., 2010). This renders an achievable regret in which all model parameters grow exponentially with L . As a result, the associated cumulative regret bound loses its sublinear growth guarantee in T in some models, e.g., Gaussian processes under bounded parameter RKHS with polynomial spectral decay (see Appendix D for more details). Furthermore, besides the achievable bound, we also characterize a general minimax lower bound and its refined counterparts for special cases.

Notations. For a positive integer $N \in \mathbb{N}$, we define $[N] \triangleq \{1, \dots, N\}$. The Euclidean norm of vector $X \in \mathbb{R}^d$ is denoted by $\|X\|$. For a function f and a set of its arguments \mathcal{Z} , we denote the empirical quadratic error of f with respect to the set \mathcal{Z} by $\|f\|_{\mathcal{Z}}^2 \triangleq \sum_{z \in \mathcal{Z}} f^2(z)$. The infinity norm of function $f : \mathcal{Z} \rightarrow \mathbb{R}$ is denoted by $\|f\|_{\infty} \triangleq \sup_{z \in \mathcal{Z}} |f(z)|$. For a set \mathcal{A} , we use $|\mathcal{A}|$ to denote the cardinality of the set.

2 GCB MODEL AND OBJECTIVE

2.1 Causal Graph with General SCM

Consider a directed acyclic graph (DAG) $\mathcal{G}(\mathcal{V}, \mathcal{E})$, in which $\mathcal{V} = [N]$ is the set of nodes, and \mathcal{E} denotes the set of directed edges. The ordered tuple $(i, j) \in \mathcal{E}$ indicates a directed edge from node $i \in [N]$ to node $j \in [N]$. We refer to the parents of node $i \in [N]$ by $\text{pa}(i)$. We define the *causal depth* of node i as the length of the longest directed causal path that ends at node $i \in [N]$ and denote it by L_i . We define $d_i \triangleq |\text{pa}(i)|$ as the in-degree of node i and the maximum degree in \mathcal{G} is denoted by $d \triangleq \max_{i \in [N]} d_i$, and the length of the longest directed causal path in \mathcal{G} is denoted by $L = \max_{i \in [N]} L_i$. We define $\mathbf{X} \triangleq (X_1, \dots, X_n)^\top$ as the vector of random variables associated with the nodes. We consider the following general SCM:

$$X_i = f_i(X_{\text{pa}(i)}) + \epsilon_i, \quad \forall i \in [N], \quad (1)$$

where functions $\{f_i : i \in [N]\}$ are **unknown** and f_i belongs to a specified class of functions denoted by

\mathcal{F}_i (e.g., linear, polynomial, or neural network). Furthermore, our analysis is sometimes specialized to *parameterized* (i.e., finite-dimensional) families of distributions. In such cases, we denote the parameters of model f_i by θ_i . The terms $\{\epsilon_i : i \in [N]\}$ account for additive noise terms, according to which we define the noise vector $\epsilon \triangleq (\epsilon_1, \epsilon_2, \dots, \epsilon_N)^\top$. Noise terms are independent across nodes. Our study centers on the causal bandit problem over a *known* DAG.

2.2 Generalised Soft Interventions

The conditional distribution of X_i given its parents, i.e., $\mathbb{P}(X_i | X_{\text{pa}(i)})$, is assumed to be **unknown**. A soft intervention on node i changes this conditional distribution. We consider a *generalized* soft intervention mechanism such that intervening on node i changes $\mathbb{P}(X_i | X_{\text{pa}(i)})$ to a *class* of distinct distributions. Interventions on node i are parameterized by $a_i \in \mathcal{A}_i \subseteq \mathbb{R}$, and the intervention spaces $\{\mathcal{A}_i : i \in [N]\}$ can be discrete or a continuum of models. This generalized model subsumes the widely used atomic intervention models. To formalize the effect of interventions, we incorporate their impact as changes on the functions $\{f_i : i \in [N]\}$. Specifically, under an interventional on node i parameterized by a_i , function $f_i \in \mathcal{F}_i$ changes to another function in the same class \mathcal{F}_i , which we denote by $f_i(\cdot; a_i)$. Hence, the interventional counterpart of the SEMs in (1) under an intervention parameterized by $a_i \in \mathcal{A}_i$ becomes

$$X_i = f_i(X_{\text{pa}(i)}; a_i) + \epsilon_i, \quad \forall i \in [N]. \quad (2)$$

To unify the observational and interventional model in (1) and (2), without loss of generality, we use the convention that setting $a_i = 0$ in (2) generates the observational model. We allow simultaneous intervention at any desired set of nodes. Hence, an intervention is specified by the intervention vector $\mathbf{a} \triangleq (a_1, \dots, a_N)^\top$, generated from the intervention space $\mathcal{A} \triangleq \prod_{i=1}^N \mathcal{A}_i$. Under intervention $\mathbf{a} \in \mathcal{A}$, we denote the *unknown* probability distribution of \mathbf{X} by $\mathbb{P}_{\mathbf{a}}$. Finally, we define function vector $\mathbf{f} \triangleq (f_1, \dots, f_N)^\top$ and function space $\mathcal{F} \triangleq \prod_{i=1}^N \mathcal{F}_i$.

2.3 Problem Statement

The objective in causal bandits is designing a sequence of intervention vectors to optimize a reward measure. Following the convention of the causal bandit literature, we designate N as the reward node. Hence, the expected reward under intervention $\mathbf{a} \in \mathcal{A}$ is given by

$$\mu_{\mathbf{a}} \triangleq \mathbb{E}_{\mathbf{a}} [X_N], \quad (3)$$

where $\mathbb{E}_{\mathbf{a}}$ denotes expectation under measure $\mathbb{P}_{\mathbf{a}}$. The optimal intervention $\mathbf{a}^* \in \mathcal{A}$ is defined as the interven-

tion that maximizes the reward, i.e.,

$$\mathbf{a}^* \triangleq \arg \max_{\mathbf{a} \in \mathcal{A}} \mu_{\mathbf{a}} . \quad (4)$$

To identify \mathbf{a}^* , the learner sequentially interacts with the underlying causal system. The sequence of interventions over time is denoted by $\{\mathbf{a}(t) \in \mathcal{A} : t \in \mathbb{N}\}$. Upon intervention $\mathbf{a}(t)$ at time $t \in \mathbb{N}$, the learner observes the graph data $\mathbf{X}(t) = (X_1(t), \dots, X_N(t))^\top$. We denote the filtration generated by the choices of interventions and associated values by

$$\mathcal{H}_t \triangleq \sigma(\mathbf{a}(1), \mathbf{X}(1), \dots, \mathbf{a}(t), \mathbf{X}(t), \mathbf{a}(t+1)) . \quad (5)$$

The learner’s objective is to minimize the cumulative regret incurred over time with respect to the best intervention \mathbf{a}^* in hindsight. We define $r(t) \triangleq \mu_{\mathbf{a}^*} - \mu_{\mathbf{a}(t)}$ as the average regret incurred at time t and denote the associated average cumulative regret over horizon T by $R(T) = \sum_{t=1}^T r(t)$. We consider the following two standard cumulative regret measures.

Frequentist regret: We assess the expected cumulative regret denoted by

$$\mathbb{E}[R(T)] \triangleq T\mu_{\mathbf{a}^*} - \sum_{t=1}^T \mu_{\mathbf{a}(t)} . \quad (6)$$

Bayesian regret: We evaluate the Bayesian regret as the average cumulative regret over the entire class \mathcal{F} . For this purpose, we denote the cumulative regret associated with $\mathbf{f} \in \mathcal{F}$ by $R_{\mathbf{f}}(T)$ and accordingly define

$$\text{BR}(T) \triangleq \mathbb{E}_{\mathcal{F}} \mathbb{E}[R_{\mathbf{f}}(T)] . \quad (7)$$

We adopt the following regularity assumptions.

Assumption 2.1 (Sub-Gaussian noise). We assume the noise terms are conditionally sub-Gaussian, i.e., for all $\lambda \in \mathbb{R}$ we have

$$\mathbb{E}[\exp(\lambda \epsilon_i(t)) \mid \mathcal{H}_{t-1}] \leq \exp(\lambda^2/2) , \quad \lambda \in \mathbb{R} . \quad (8)$$

Assumption 2.2 (Lipschitz continuity). The functions f_i in class \mathcal{F}_i are Lipschitz-continuous with Lipschitz constant K_i , i.e., for all $a_i \in \mathcal{A}_i$ and $x, y \in \mathbb{R}^{d_i}$, we have

$$|f_i(x; a_i) - f_i(y; a_i)| \leq K_i \|x - y\| . \quad (9)$$

Assumption 2.3 (Function boundedness). For each class \mathcal{F}_i there exists a constant $C_i \in \mathbb{R}$ such that

$$|f_i(X_{\text{pa}(i)}; a_i)| \leq C_i . \quad (10)$$

This assumption is natural for some SCMs, e.g., neural network SCMs with sigmoid activation functions. In some other SCMs, this ensures the system’s boundedness, a common assumption in causal bandits (e.g., Varici et al., 2023; Sussex et al., 2023).

3 PRELIMINARIES

We provide some definitions that are instrumental to the analysis and characterizing the regret.

3.1 Eluder Dimension and Covering Number

Our regret analyses show that the achievable regret is a function of the inherent complexity of the function classes $\{\mathcal{F}_i : i \in [N]\}$. In this section, we specify two measures of functional complexity that are instrumental in specifying achievable regrets. These measures aim to capture an approximate structure of the function class based on target approximation tolerances. The first measure is the *eluder dimension* of a function class (Russo and Van Roy, 2013). It measures the degree of dependence among the function class as the worst-case sample complexity required to predict the values of unobserved interventions using the observed samples. We refer to Huang et al. (2021) and Li et al. (2022) for more discussion on the eluder dimension. The second measure is the *covering number*, which provides insight into the function class’s propensity towards statistical over-fitting. This complexity measure is widely used in analyzing infinite function spaces (Wainwright, 2019).

To formalize these definitions, we use an extended definition of ϵ -dependence. This definition is instrumental to characterizing how distinguishable the functions $\mathbf{f} \in \mathcal{F}$ are, which can be used to specify a measure of complexity for \mathcal{F} . Corresponding to each node i , consider n arbitrary inputs to function f_i as specified in (2). Denote these inputs by $\{Z_i(m) : m \in [n]\}$ where $Z_i(m) \triangleq (X_{\text{pa}(i)}(m), a_i(m))$. Denote the space of these inputs by $\mathcal{Z}_i \triangleq \mathbb{R}^{d_i} \times \mathcal{A}_i$. Based on these, we define ϵ -dependence among the members of \mathcal{Z}_i as follows.

Definition 3.1 (ϵ -dependent). We say $Z_i \in \mathcal{Z}_i$ is ϵ -dependent on $\{Z_i(m) : m \in [n]\}$ with respect to \mathcal{F}_i if any pair of functions $f_i, \tilde{f}_i \in \mathcal{F}_i$ that satisfy

$$\sum_{m=1}^n (f_i(Z_i(m)) - \tilde{f}_i(Z_i(m)))^2 \leq \epsilon^2 , \quad (11)$$

also satisfies $|f(Z_i) - \tilde{f}(Z_i)| \leq \epsilon$.

Based on this similarity measure, the ϵ -eluder dimension of the function class \mathcal{F}_i is defined as follows.

Definition 3.2 (ϵ -eluder dimension). The ϵ -eluder dimension $\text{dim}(\mathcal{F}_i, \epsilon)$ is the largest n such that for some $\epsilon' \geq \epsilon$, there exists an ordered set $\{Z_i(m) : m \in [n]\}$ where each element is ϵ' -independent of its predecessors. Based on the individual eluder dimensions of the

classes, we define the maximum eluder dimension as

$$\dim(\mathcal{F}, \epsilon) \triangleq \max_{i \in [N]} \dim(\mathcal{F}_i, \epsilon). \quad (12)$$

Definition 3.3 (Covering number). For a given function class \mathcal{F}_i and any constant $\alpha > 0$, we say that the set $\mathcal{C} \subseteq \mathcal{F}_i$ is an α -cover of \mathcal{F}_i if for any function $f_i \in \mathcal{F}_i$ we can always find another function $\tilde{f}_i \in \mathcal{C}$ such that $\|f_i - \tilde{f}_i\|_\infty \leq \alpha$. The α -covering number associated with \mathcal{F}_i and α , denoted by $\text{cn}_\alpha(\mathcal{F}_i)$, is the smallest cardinality \mathcal{C} among all possible α -cover sets \mathcal{C} . Accordingly, for the class \mathcal{F} we define

$$\text{cn}_\alpha(\mathcal{F}) \triangleq \max_{i \in [N]} \text{cn}_\alpha(\mathcal{F}_i). \quad (13)$$

4 GCB ALGORITHMS

This section serves a two-fold purpose. First, it provides constructive proofs for regret guarantee results for any arbitrary family of functions \mathcal{F} . Secondly, it provides algorithms for parametric families of functions that achieve the regret guarantees. More specifically, some decision rules involve solving *functional* optimization problems for choosing an optimal choice of $f_i \in \mathcal{F}_i$. Solving such problems has different natures in the parametric (finite-dimensional) and non-parametric (infinite-dimensional) function spaces \mathcal{F} .

In the infinite-dimensional setting, it requires functional optimization. Unless solvable via the calculus of variations, solving such problems is difficult to carry out in practice. Nevertheless, a guaranteed existence of a solution suffices for establishing the regret guarantees for the class. On the other hand, in the parametric class of functions, which are the classes encountered in all practical settings, finding an optimal f_i becomes equivalent to solving a parametric optimization problem, which often can be solved numerically, rendering the parameterized family of distributions amenable to practical implementations. For such cases, in addition to the regret guarantees, we also have algorithms that achieve the regret upper bounds. We present algorithms for the frequentist and Bayesian settings and provide the associated regret analysis in Section 5.

Algorithm overview. The central part of our algorithm for the frequentist setting is designing a procedure for estimating the optimal intervention vector \mathbf{a}^* specified in (4). The estimator is designed such that we can progressively, over time, control the size of the confidence set we form for vector function \mathbf{f} . Given the confidence intervals, we adopt a UCB-based strategy to choose interventions \mathbf{a}^* . We refer to this algorithm as the **General Causal Bandit UCB** (GCB-UCB). For the Bayesian setting, we find the posterior likelihood of the function vector $\mathbf{f} \in \mathcal{F}$ based on the available data

and adopt an approach based on Thompson Sampling (TS) to identify an intervention that maximizes the expected posterior reward. We refer to this algorithm as the **General Causal Bandit TS** (GCB-TS).

Algorithm 1 GCB-UCB

- 1: **Input:** Horizon T , causal graph \mathcal{G} , intervention set \mathcal{A} and function spaces \mathcal{F} .
 - 2: **Initialization:** Set $\beta_0(\mathcal{F}_i, \delta, \alpha)$ as in (18).
 - 3: **for** $t = 1, 2, \dots, T$ **do**
 - 4: Compute the least squares estimates $\tilde{f}_{i,t}$ as in (16) for $i \in [N]$.
 - 5: Construct the confidence set $\mathcal{C}_{i,t}$ as in (17) for $i \in [N]$.
 - 6: Select interventions $\mathbf{a}(t)$ according to (20).
 - 7: Pull $\mathbf{a}(t)$, observe $\mathbf{X}(t) = [X_1(t), \dots, X_N(t)]$.
 - 8: **end for**
-

Algorithm 2 GCB-TS

- 1: **Input:** Horizon T , causal graph \mathcal{G} , intervention set \mathcal{A} , function spaces \mathcal{F} and prior distribution π_0 .
 - 2: **for** $t = 1, 2, \dots, T$ **do**
 - 3: $\tilde{\mathbf{f}} \sim \pi_{t-1}(\cdot | \mathcal{Z}(t))$.
 - 4: Select interventions $\mathbf{a}(t)$ according to (22).
 - 5: Pull $\mathbf{a}(t)$, observe $\mathbf{X}(t) = [X_1(t), \dots, X_N(t)]$.
 - 6: **end for**
-

To formalize the algorithms, by recalling the definition $Z_i(t) \triangleq (X_{\text{pa}(i)}(t), a_i(t))$, we define

$$\mathcal{Z}_i(t) \triangleq \{Z_i(s) : s \in [t-1]\}, \quad (14)$$

which is the set of $t-1$ initial interventions on node i and the data generated by its parents. Similarly, we denote the set of all interventions and data generated in the system up to time $t-1$ by

$$\mathcal{Z}(t) \triangleq \{(\mathbf{X}(s), \mathbf{a}(s)) : s \in [t-1]\}. \quad (15)$$

4.1 GCB-UCB Algorithm

We adopt the ordinary least-squares method to estimate the functions. Based on the choices of interventions and the interventional data up to time t collected in $\mathcal{Z}_i(t)$, we form an estimate for f_i at time t as follows.

$$\tilde{f}_{i,t} \triangleq \arg \min_{f_i \in \mathcal{F}_i} \sum_{s=1}^{t-1} \left[f_i(X_{\text{pa}(i)}(s); a_i(s)) - X_i(s) \right]^2. \quad (16)$$

Based on this estimate, we construct a confidence interval for the estimated function $\tilde{f}_{i,t}$ as follows.

$$\mathcal{C}_{i,t} \triangleq \left\{ g \in \mathcal{F}_i : \left\| g - \tilde{f}_{i,t} \right\|_{\mathcal{Z}_i(t)} \leq \sqrt{\beta_t(\mathcal{F}_i, \delta, \alpha_i)} \right\}, \quad (17)$$

in which $\beta_t(\mathcal{F}_i, \delta, \alpha_i)$ controls the radius of the confidence interval, and it is specified based on the α_i -covering number of the class \mathcal{F}_i , i.e., $\text{cn}_{\alpha_i}(\mathcal{F}_i)$, as follows.

$$\beta_t(\mathcal{F}_i, \delta, \alpha_i) \triangleq 8 \log \left(\frac{\text{cn}_{\alpha_i}(\mathcal{F}_i)}{\delta} \right) + 2\alpha_i t \left(8C_i + \sqrt{8 \ln \frac{4t^2}{\delta}} \right). \quad (18)$$

Besides the covering number, an intrinsic property of the function space, the radius $\beta_t(\mathcal{F}_i, \delta, \alpha_i)$ depends on α_i and δ . The role of α_i is controlling the discretization parameter for covering number, which also influences the second summand that is related to the uncertainty of estimators. Furthermore, the role of δ is controlling the confidence level of the confidence set. Based on this, the confidence radius has two components. One component is a constant (time-invariant) that depends on the covering number $\text{cn}_{\alpha}(\mathcal{F}_i)$ and is decreasing with δ . The second component is increasing in time t and is decreasing with δ . Finally, we define the confidence space for the entire causal system as $\mathcal{C}_t \triangleq \prod_{i=1}^N \mathcal{C}_{i,t}$. Based on these definitions, corresponding to intervention $\mathbf{a} \in \mathcal{A}$ at time t we define the upper confidence bound as follows.

$$\text{UCB}_{\mathbf{a}}(t) \triangleq \max_{\mathbf{f} \in \mathcal{C}_t} \mathbb{E}_{\mathbf{a}}[X_N | \mathbf{f}]. \quad (19)$$

Accordingly, the intervention at time t is specified by

$$\mathbf{a}(t) \triangleq \arg \max_{\mathbf{a} \in \mathcal{A}} \text{UCB}_{\mathbf{a}}(t), \quad (20)$$

and finally, we define

$$\tilde{\mathbf{f}}_t \triangleq \arg \max_{\mathbf{f} \in \mathcal{C}_t} \text{UCB}_{\mathbf{a}(t)}(t). \quad (21)$$

4.2 GCB-TS Algorithm

A TS-based algorithm gradually refines the posterior distributions of the reward for each intervention and selects the interventions by sampling from their posterior distributions. The interventions are selected sequentially in a way that they balance the exploitation and the exploration processes. To this end, a TS algorithm must update all arms' posterior distributions. For the GCB-TS algorithm, we assume a prior π_0 is known in advance, and it is periodically updated over time to construct the posterior π_t after gathering the observed samples $\mathcal{Z}(t)$. The GC-TS algorithm at time t , samples a function \mathbf{f} from the posterior π_{t-1} . Subsequently, an intervention is selected such that it maximizes the expected reward under $\tilde{\mathbf{f}}$, i.e.,

$$\mathbf{a}(t) = \arg \max_{\mathbf{a} \in \mathcal{A}} \mathbb{E}_{\mathbf{a}}[X_N | \tilde{\mathbf{f}}]. \quad (22)$$

5 REGRET GUARANTEE

This section has two parts. In Section 5.1, we establish regret upper and lower bounds for a general class of functions \mathcal{F} . These bounds explicitly capture the impact of the causal graph structure (through d and L) and the function class \mathcal{F} (via its eluder dimension, covering number, and Lipschitz constants). In Section 5.2, we specialize and refine these bounds for three special classes of SCMs: linear, polynomial, and neural network, which cover the prevalent SCMs.

5.1 Regret Upper Bounds

We first provide upper bounds that capture the performance of the GCB-UCB and GCB-TS algorithms, followed by a minimax lower bound that is independent of the function class \mathcal{F} . The key step in these analyses is evaluating the impact of the causal structure on the regret. Specifically, once we form an estimate of the function f_i for any node $i \in [N]$, the estimation error for this node propagates through the causal path from node i to node N , affecting the estimation errors of all the nodes on the path as well as the final reward node. We refer readers to this as the *compounding* effect of the causal graph.

We begin with recalling a concentration lemma provided in Russo and Van Roy (2013). This lemma addresses the confidence set's precision. Specifically, this lemma shows that the confidence set $\mathcal{C}_{i,t}$ contains f_i with a high probability at all times and specifies the precision of the estimates $\{\tilde{f}_{i,t} : t \in \mathbb{N}\}$ formed by (16).

Lemma 5.1. *For any $i \in [N]$ and $\forall \delta > 0$ and $\forall \alpha > 0$*

$$\mathbb{P}(f_i \in \mathcal{C}_{i,t}) \geq 1 - 2\delta, \quad \forall t \in [T]. \quad (23)$$

Next, we specify the discretization parameter α_i based on the function space \mathcal{F}_i . When the function space is finite, we choose α_i to be the minimum distance with respect to the infinity norm between any two functions in \mathcal{F}_i . If the function space is continuous, however, the minimum distance may be zero. In the latter case, we truncate α_i at $\frac{1}{T}$, resulting in the following specification overall:

$$\alpha_i \triangleq \max \left\{ \frac{1}{T}, \inf_{f_1 \neq f_2 \in \mathcal{F}_i} \|f_1 - f_2\|_{\infty} \right\}. \quad (24)$$

The choice of $\frac{1}{T}$ can be understood by considering the two opposing effects of α_i in the two summands of $\beta_t(\mathcal{F}_i, \delta, \alpha_i)$, the confidence radii defined in (18). On the one hand, increasing α_i decreases the covering number $\text{cn}_{\alpha_i}(\mathcal{F}_i)$, which in turn decreases the first summand. On the other hand, increasing α_i increases the second summand in $\beta_t(\mathcal{F}_i, \delta, \alpha_i)$. Hence, selecting $\alpha_i = \frac{1}{T}$ balances the two summands. Throughout the

rest of the paper, we use the shorthand $\dim(\mathcal{F})$ for $\dim(\mathcal{F}, \frac{1}{T})$ and the shorthand $\text{cn}(\mathcal{F})$ for $\text{cn}_{\frac{1}{T}}(\mathcal{F})$.

For characterizing the *compounding* effect of the causal graph along various causal paths, we define the following derived Lipschitz constants. First recall that we have defined L_i as the causal depth of node i , which is the length of the longest causal path that starts from the root nodes of the graph and ends at node i . Accordingly, we define $K^{(\ell)}$ as the maximum Lipschitz constant among the constants of the nodes that have causal depth $\ell \in [L]$, i.e.,

$$K^{(\ell)} = \max_{i \in [N], L_i = \ell} K_i. \quad (25)$$

Finally, we define $K = \prod_{\ell=2}^L K^{(\ell)}$.

We present the following lemma to bound the cumulative error at a node $i \in [N]$ due to the compounding of estimator uncertainty along causal paths. This is a key step in characterizing the regret at the reward node because the latter hinges on the compounded estimation errors at all intermediate nodes along the causal paths to the reward node. In our analysis, we find that the contribution of node i at time t is determined by the difference in the expected value of X_i under the optimistic function $\hat{\mathbf{f}}_t$, which maximizes UCB over the confidence set, and the true function \mathbf{f} . The next lemma bounds these differences.

Lemma 5.2 (Compounding Error). *If $f_i \in \mathcal{C}_{i,t}$ for all $i \in [N]$ and $t \in [T]$, then the following error bound holds*

$$\begin{aligned} & \sum_{t=1}^T \left| \mathbb{E}_{\mathbf{a}(t)}[X_i | \hat{\mathbf{f}}_t] - \mathbb{E}_{\mathbf{a}(t)}[X_i | \mathbf{f}] \right| \\ & \leq \mathcal{B}(\mathcal{F}, \delta) \sum_{\ell=2}^{L_i} d^{\ell-1} \prod_{k=2}^{\ell} K^{(k)}. \end{aligned} \quad (26)$$

The quantity $\mathcal{B}(\mathcal{F}, \delta)$ appearing in Lemma 5.2 is defined as follows, where for each node $i \in [N]$, we define $\mathcal{B}_i(\mathcal{F}_i, \delta)$ to represent the maximum accumulated error when $L_1 = 1$ for $i \in [N]$.

$$\begin{aligned} \mathcal{B}_i(\mathcal{F}_i, \delta) & \triangleq 1 + \min\{\dim(\mathcal{F}_i, \alpha_i), T\} C_i \\ & \quad + 4\sqrt{\dim(\mathcal{F}_i, \alpha_i) \beta_T(\mathcal{F}_i, \alpha_i, \delta) T}, \quad (27) \\ \mathcal{B}(\mathcal{F}, \delta) & \triangleq \max_{i \in [N]} \mathcal{B}_i(\mathcal{F}_i, \delta). \quad (28) \end{aligned}$$

When we have eluder dimension $\dim(\mathcal{F}_i, \alpha_i) = o(\sqrt{T})$, i.e., when $\mathcal{B}_i(\mathcal{F}_i, \delta) = o(T)$, for all nodes $i \in [N]$, $\mathcal{B}(\mathcal{F}, \delta)$ defined in (28) scales as

$$\mathcal{O}\left(\sqrt{T \dim(\mathcal{F}) \beta_T\left(\mathcal{F}, \frac{1}{T}, \delta\right)}\right). \quad (29)$$

Leveraging the result of Lemma 5.2, we characterize the achievable regrets of GCB-UCB and GCB-TS.

Theorem 5.3 (Regret Upper Bound). *Under Assumptions 2.1-2.3, by setting $\delta = \frac{1}{NT}$, for the GCB-UCB algorithm, we have*

$$\begin{aligned} & \mathbb{E}[R(T)] \\ & \leq 4C_N + \mathcal{B}\left(\mathcal{F}, \frac{1}{NT}\right) \sum_{\ell=0}^L d^{\ell-1} \prod_{j=2}^{\ell} K^{(j-1)} \quad (30) \\ & = \mathcal{O}\left(K d^{L-1} \sqrt{T \dim(\mathcal{F}) \log(NT \text{cn}(\mathcal{F}))}\right). \quad (31) \end{aligned}$$

Corollary 5.4 (Bayesian Regret Upper Bound). *Under Assumptions 2.1-2.3, by setting $\delta = \frac{1}{NT}$, for the GCB-TS algorithm, we have*

$$\text{BR}(T) = \mathcal{O}\left(K d^{L-1} \sqrt{T \dim(\mathcal{F}) \log(NT \text{cn}(\mathcal{F}))}\right). \quad (32)$$

Based on the results established in Theorem 5.3, we have the following key observations.

- **Graph size.** One important observation is that the regret bounds depend only logarithmically on graph size N . More importantly, we will show that for all the widely-used special cases, the dependence on N behaves as $\sqrt{\frac{\log N}{\log T}}$, indicating that the graph size has a diminishing effect in T . To highlight the significance of this, note that in the state-of-the-art results (Varici et al., 2023; Sussex et al., 2023), the achievable regret bounds for less general SCMs (linear and Gaussian processes) scale with \sqrt{N} and N , respectively. Furthermore, the regret bounds depend on the graph through its connectivity structure specified by d and L .
- **Impact of function classes.** The upper bound depends linearly on the square root of the eluder dimension and logarithm of the covering number, reflecting the influence of the function space's complexity on the regret's growth. Furthermore, regret bounds depend on the Lipschitz coefficient K , which is expected.
- **Graph parameters.** The regret upper bound has polynomial dependence on maximum degree d and exponential dependence on maximum causal depth L .
- **Intervention space.** Another observation is the lack of dependence on the size of the intervention space. This allows for adopting intervention spaces with any desired level of granularity.

5.2 Refined Regret Bounds for Special SCMs

In this section, we discuss three special classes of SCMs and provide refined upper and corresponding lower bounds on regret for them. These results are summarized in Table 1. The proofs and additional details are presented in Appendices B and C.

For the lower bound, we consider the minimax lower bound on the regret when considering the specific function classes \mathcal{F} . We define \mathcal{G} as the set of all causal graphs specified over N nodes with parameters d and L , and Π as the set of all possible bandit policies. The minimax regret is defined as

$$\inf_{\pi \in \Pi} \sup_{\mathcal{F} \in \mathcal{F}} \sup_{\mathcal{G} \in \mathcal{G}} \mathbb{E}[R(T)]. \quad (33)$$

Linear SCMs. For given vectors $\theta_i, \bar{\theta}_i \in \mathbb{R}^{d_i}$ and scalar $a_i \in \mathcal{A}_i$, set

$$f_i(X_{\text{pa}(i)}; a_i) = \langle \theta_i(1 - a_i) + a_i \bar{\theta}_i, X_{\text{pa}(i)} \rangle, \quad (34)$$

and accordingly, define

$$\mathcal{F}_i = \{f_i : \theta_i, \bar{\theta}_i \in \mathbb{R}^{d_i}, \|\theta_i\| \leq K_i, \|\bar{\theta}_i\| \leq K_i\}, \quad (35)$$

where the norm constraints are enforced to satisfy Assumptions 2.2-2.3. By setting $\mathcal{A}_i = \{0, 1\}$, the generalized soft intervention model reduces to the linear SCMs with binary atomic interventions. We note that Assumption 2.3 implies that the noise terms are bounded under linear SCMs, which is similar to the assumption adopted in Varici et al. (2023).

Theorem 5.5. *For the linear class of SCMs specified in (35) and $\mathcal{A}_i \in \{0, 1\}$ we have*

$$\dim(\mathcal{F}_i) = \tilde{\mathcal{O}}(d \log T), \quad \text{and} \quad \log \text{cn}(\mathcal{F}_i) = \tilde{\mathcal{O}}(d \log T),$$

and the upper bound in Theorem 5.3 becomes

$$\mathcal{O}(Kd^L \sqrt{T \log T \varpi_L}), \quad \varpi_L = \log T \left[1 + \frac{\log N}{d \log T} \right].$$

Theorem 5.5 shows that achievable regret has a diminishing dependence on N (as T increases). This is a substantial improvement over the known regret result, i.e., $\mathcal{O}(Kd^{L-\frac{1}{2}} \sqrt{NT} \log T)$, which grows with \sqrt{N} .

Theorem 5.6 (Minimax Lower Bound for Linear SCMs). *Under linear SCMs defined in (35), given graph parameters d and L and Lipschitz constant K , there exists a causal bandit instance for which the expected regret of any algorithm is at least*

$$\mathbb{E}[R(T)] \geq \Omega\left(Kd^{\frac{L}{2}-1} \sqrt{T}\right). \quad (36)$$

Polynomial SCMs. For a given vector $\theta_i \in \mathbb{R}^{d_i+1}$ and a polynomial maximum degree $p \in \mathbb{N}$, set

$$f_i(X_{\text{pa}(i)}; a_i) = \left(\langle \theta_i, [X_{\text{pa}(i)}, a_i] \rangle \right)^p, \quad (37)$$

in which the intervention parameter $|a_i| \leq 1$ controls the coefficients in the polynomial. Note that the polynomial can be defined in different forms, and the one selected is only a representative. Accordingly, define

$$\mathcal{F}_i = \left\{ f_i : \theta_i \in \mathbb{R}^{d_i+1}, \|\theta_i\| \leq \frac{K_i^{1/p}}{p^{1/p}(dC+1)} \right\}, \quad (38)$$

where $C = \max_{i \in [N]} C_i$ and the norm constraints are enforced to satisfy Assumptions 2.2-2.3. Similarly, Assumption 2.3 implies the noises are bounded in this case.

Theorem 5.7. *For the polynomial class of SCMs specified in (38) for $p = 2$ we have*

$$\dim(\mathcal{F}_i) = \tilde{\mathcal{O}}((d+1)^2 \log T), \quad (39)$$

$$\text{and} \quad \log \text{cn}(\mathcal{F}_i) = \tilde{\mathcal{O}}((d+1)^2 \log T), \quad (40)$$

and the upper bound in Theorem 5.3 becomes

$$\mathcal{O}(Kd^{L+1} \sqrt{T \log T \varpi_P}), \quad \varpi_P = \log T \left[1 + \frac{\log N}{d \log T} \right].$$

Theorem 5.8 (Minimax Lower Bound for Polynomial SCMs). *Under polynomial SCMs defined in (38), given graph parameters d and L and Lipschitz constant K , there exists a causal bandit instance for which the expected regret of any algorithm is at least*

$$\mathbb{E}[R(T)] \geq \Omega\left(K \sqrt{T}\right). \quad (41)$$

Neural Network SCMs. For a given matrix $\Theta_1 \in \mathbb{R}^{s \times (d_i+1)}$ and vector $\Theta_2 \in \mathbb{R}^s$, set f_i as a two-layer neural network with the input vector $[a_i, X_{\text{pa}(i)}]$, i.e.,

$$f_i(X_{\text{pa}(i)}; a_i) = \sigma(\Theta_2 \cdot \sigma(\Theta_1 \cdot [a_i, X_{\text{pa}(i)}])), \quad (42)$$

where σ is an increasing continuous activation function with bounded gradients. Accordingly, we define \mathcal{F}_i as

$$\mathcal{F}_i = \{f_i : \Theta_1 \in \mathbb{R}^{s \times (d_i+1)}, \Theta_2 \in \mathbb{R}^s\}. \quad (43)$$

Theorem 5.9. *For the neural network class of SCMs specified in (43) with width $s \geq d$, the upper bound in Theorem 5.3 becomes*

$$\mathcal{O}\left(\bar{K}d^{L-1} \sqrt{T \log T \varpi_N}\right), \quad \varpi_N = s \log T \left[1 + \frac{\log N}{d \log T} \right],$$

where \bar{K} is an upper bound of Lipschitz constant K .

Theorem 5.10 (Minimax Lower Bound for Neural Network SCMs). *Given graph parameters d and L and Lipschitz constant K , there exists a causal bandit instance with neural network SCM defined in (43) for which the expected regret of any algorithm is at least*

$$\mathbb{E}[R(T)] \geq \Omega\left(Kd^{\frac{L}{2}-1}\sqrt{T}\right). \quad (44)$$

We note that in our definition of the class of neural networks in (42), we have an activation function in the second layer, which is not standard in neural networks. This choice was made to simplify the results presented. Besides, if activation functions are allowed to vary between layers, we can set the activation function in the second layer to be the identity mapping, which will result in a similar result for the widely used definition of neural network.

We note that the difference between lower and upper bounds in linear and neural network SCMs (i.e., $d^{L/2}$ versus d^L) has reasons fundamentally beyond our setting. Specifically, the regret of the linear bandit algorithm called optimism in the face of uncertainty (OFUL) (Abbasi-Yadkori et al., 2011) is $\mathcal{O}(d\sqrt{T})$. In contrast, the lower bound for linear bandits with finite arms is $\mathcal{O}(\sqrt{dT})$, yielding a difference of $d^{1/2}$ versus d in the lower and upper bounds. The linear bandit problem can be considered a special case of our framework with linear SCM and $L = 1$. In essence, we have a similar gap compounded over L layers. Besides, we expect the missing graph parameters in the lower bound of quadratic SCMs are due to the large Lipschitz constant in this case (see Appendix C.2 for details).

6 CONCLUSION

In this paper, we have provided regret analysis for causal bandits with general structural equation models (SCMs) and generalized soft interventions. The achievable regret bounds are shown to depend on inherent complexity measures of the class of functions encompassing the SCMs, namely the eluder dimension and the covering number. Furthermore, the achievable regret depends on the graph topology primarily through its connectivity parameters (maximum degree and the length of the causal paths) and has a diminishing dependence on the size of the graph. Furthermore, we have delineated a universal minimax regret lower bound. Finally, we have customized the results and refined the bounds for three special classes of SCMs: linear, polynomial, and neural network SCMs.

Acknowledgments

This work was supported by the Rensselaer-IBM AI Research Collaboration (AIRC), part of the IBM AI Horizons Network.

References

- Abbasi-Yadkori, Y., Pál, D., and Szepesvári, C. (2011). Improved algorithms for linear stochastic bandits. In *Proc. Advances in Neural Information Processing Systems*, Granada, Spain.
- Agrawal, R. (1995). Sample mean based index policies by $o(\log n)$ regret for the multi-armed bandit problem. *Advances in Applied Probability*, 27(4):1054–1078.
- Ahmed, O., Träuble, F., Goyal, A., Neitz, A., Wuthrich, M., Bengio, Y., Schölkopf, B., and Bauer, S. (2021). Causalworld: A robotic manipulation benchmark for causal structure and transfer learning. In *Proc. International Conference on Learning Representations*, virtual.
- Badsha, M. B. and Fu, A. Q. (2019). Learning causal biological networks with the principle of mendelian randomization. *Frontiers in Genetics*, 10.
- Bareinboim, E., Forney, A., and Pearl, J. (2015). Bandits with unobserved confounders: A causal approach. In *Proc. Advances in Neural Information Processing Systems*, Montréal, Canada.
- Bilodeau, B., Wang, L., and Roy, D. M. (2022). Adaptively exploiting d -separators with causal bandits. In *Proc. Advances in Neural Information Processing Systems*, New Orleans, LA.
- De Kroon, A., Mooij, J., and Belgrave, D. (2022). Causal bandits without prior knowledge using separating sets. In *Proc. Conference on Causal Learning and Reasoning*, Eureka, CA.
- Feng, S. and Chen, W. (2023). Combinatorial causal bandits. In *Proc. the AAAI Conference on Artificial Intelligence*, Washington, D.C.
- Feng, S., Xiong, N., and Chen, W. (2023). Combinatorial causal bandits without graph skeleton. *arXiv:2301.13392*.
- Huang, K., Kakade, S. M., Lee, J. D., and Lei, Q. (2021). A short note on the relationship of information gain and eluder dimension. *arXiv:2107.02377*.
- Jaber, A., Kocaoglu, M., Shanmugam, K., and Bareinboim, E. (2020). Causal discovery from soft interventions with unknown targets: Characterization and learning. In *Proc. Advances in Neural Information Processing Systems*, Virtual.
- Jun, K.-S., Bhargava, A., Nowak, R., and Willett, R. (2017). Scalable generalized linear bandits: Online

- computation and hashing. *Advances in Neural Information Processing Systems*, 30.
- Konobeev, M., Etesami, J., and Kiyavash, N. (2023). Causal bandits without graph learning. *arXiv:2301.11401*.
- Kveton, B., Zaheer, M., Szepesvari, C., Li, L., Ghavamzadeh, M., and Boutilier, C. (2020). Randomized exploration in generalized linear bandits. In *International Conference on Artificial Intelligence and Statistics*, pages 2066–2076.
- Lattimore, F., Lattimore, T., and Reid, M. D. (2016). Causal bandits: Learning good interventions via causal inference. In *Proc. Advances in Neural Information Processing Systems*, Barcelona, Spain.
- Lattimore, T. and Szepesvári, C. (2020). *Bandit Algorithms*. Cambridge University Press, Cambridge, UK.
- Li, G., Kamath, P., Foster, D. J., and Srebro, N. (2022). Understanding the eluder dimension. *Proc. Advances in Neural Information Processing Systems*, 35.
- Liu, S., See, K. C., Ngiam, K. Y., Celi, L. A., Sun, X., and Feng, M. (2020). Reinforcement learning for clinical decision support in critical care: Comprehensive review. *Journal of Medical Internet Research*, 22(7).
- Lu, Y., Meisami, A., and Tewari, A. (2021). Causal bandits with unknown graph structure. In *Proc. Advances in Neural Information Processing Systems*, virtual.
- Lu, Y., Meisami, A., Tewari, A., and Yan, W. (2020). Regret analysis of bandit problems with causal background knowledge. In *Proc. Conference on Uncertainty in Artificial Intelligence*, virtual.
- Maiti, A., Nair, V., and Sinha, G. (2022). A causal bandit approach to learning good atomic interventions in presence of unobserved confounders. In *Proc. Conference on Uncertainty in Artificial Intelligence*, Eindhoven, Netherlands.
- Malek, A., Aglietti, V., and Chiappa, S. (2023). Additive causal bandits with unknown graph. In *Proc. International Conference on Machine Learning*, Honolulu, Hawaii.
- Nair, V., Patil, V., and Sinha, G. (2021). Budgeted and non-budgeted causal bandits. In *Proc. International Conference on Artificial Intelligence and Statistics*, virtual.
- Osband, I. and Van Roy, B. (2014). Model-based reinforcement learning and the eluder dimension. In *Proc. Advances in Neural Information Processing Systems*, Montréal, Canada.
- Russo, D. and Van Roy, B. (2013). Eluder dimension and the sample complexity of optimistic exploration. In *Proc. Advances in Neural Information Processing Systems*, Stateline, NV.
- Sawarni, A., Madhavan, R., Sinha, G., and Barman, S. (2023). Learning good interventions in causal graphs via covering. In *Proc. Conference on Uncertainty in Artificial Intelligence*, Pittsburgh, PA.
- Sen, R., Shanmugam, K., Dimakis, A. G., and Shakkottai, S. (2017). Identifying best interventions through online importance sampling. In *Proc. International Conference on Machine Learning*, Sydney, Australia.
- Srinivas, N., Krause, A., Kakade, S. M., and Seeger, M. (2010). Gaussian process optimization in the bandit setting: No regret and experimental design. In *Proc. International Conference on Machine Learning*, Haifa, Israel.
- Sussex, S., Makarova, A., and Krause, A. (2023). Model-based causal bayesian optimization. In *Proc. International Conference on Learning Representations*, Kigali, Rwanda.
- Varici, B., Shanmugam, K., Sattigeri, P., and Tajer, A. (2023). Causal bandits for linear structural equation models. *The Journal of Machine Learning Research*, 24(297):1–59.
- Wainwright, M. J. (2019). *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, Cambridge, UK.
- Xiong, N. and Chen, W. (2023). Combinatorial pure exploration of causal bandits. In *Proc. International Conference on Learning Representations*, Kigali, Rwanda.
- Yabe, A., Hatano, D., Sumita, H., Ito, S., Kakimura, N., Fukunaga, T., and Kawarabayashi, K. (2018). Causal bandits with propagating inference. In *Proc. International Conference on Machine Learning*, Stockholm, Sweden.
- Yang, Z., Jin, C., Wang, Z., Wang, M., and Jordan, M. I. (2020). On function approximation in reinforcement learning: Optimism in the face of large state spaces. In *Proc. Advances in Neural Information Processing Systems*, Virtual.
- Zhao, Y., Goodman, M., Kanase, S., Xu, S., Kimmel, Y., Payne, B., Khan, S., and Grao, P. (2022). Mitigating targeting bias in content recommendation with causal bandits. In *Proc. ACM Conference on Recommender Systems Workshop on Multi-Objective Recommender Systems*, Seattle, WA.

Checklist

1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes. See Section 2–4 for more details.](#)
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes. We provide the analysis of cumulative regret in Section 5.](#)
 - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Yes. And we only need standard libraries.](#)
2. For any theoretical claim, check if you include:
 - (a) Statements of the full set of assumptions of all theoretical results. [Yes. See Assumption 2.1–2.3.](#)
 - (b) Complete proofs of all theoretical results. [Yes. See Appendix A–C.](#)
 - (c) Clear explanations of any assumptions. [Yes. Assumptions are standard and we have provided discussion on Assumption 2.3.](#)
3. For all figures and tables that present empirical results, check if you include:
 - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Yes.](#)
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes. See Appendix E.](#)
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Not Applicable.](#)
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Not Applicable. Our codes require single cpu.](#)
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
 - (a) Citations of the creator If your work uses existing assets. [Not Applicable.](#)
 - (b) The license information of the assets, if applicable. [Not Applicable.](#)
 - (c) New assets either in the supplemental material or as a URL, if applicable. [Not Applicable.](#)
 - (d) Information about consent from data providers/curators. [Not Applicable.](#)
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable.](#)
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
 - (a) The full text of instructions given to participants and screenshots. [Not Applicable.](#)
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable.](#)
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable.](#)

Causal Bandits with General Causal Models and Interventions

Supplementary Materials

Table of Contents

A	Regret Upper Bound – General SCM	12
A.1	Proof of Lemma 5.2	12
A.2	Proof of Theorem 5.3	16
B	Regret Upper Bounds for Special SCMs	17
B.1	Proof of Theorem 5.5 (Linear SCMs)	17
B.2	Proof of Theorem 5.7 (Polynomial SCMs)	21
B.3	Proof of Theorem 5.9 (Neural Network SCMs)	21
C	Regret Lower Bounds	23
C.1	Proof of Theorem 5.6 (Linear SCMs)	24
C.2	Proof of Theorem 5.8 (Polynomial SCMs)	26
C.3	Proof of Theorem 5.10 (Neural Network SCMs)	29
D	Further Comparisons with (Sussex et al., 2023)	32
D.1	Information Gain	33
D.2	Achievable Regret of the MCBO Algorithm	33
D.3	Achievable Regret of the GCB Algorithm	34
E	Experiments	34
F	Extension to <i>do</i> interventions	36

A Regret Upper Bound – General SCM

A.1 Proof of Lemma 5.2

Throughout this proof, we use β_t as a shorthand to refer to $\beta_t(\mathcal{F}_i, \delta, \alpha_i)$. We define $\bar{\mathbf{f}}_t \triangleq (\bar{f}_{1,t}, \dots, \bar{f}_{N,t})$. Based on these, we first prove that for the nodes with causal depth 0, i.e., $L_i = 0$, we have Consider node $i \in [N]$ with causal depth $L_i = 0$, in this case we have

$$[X_i(t) \mid \bar{\mathbf{f}}_t] = [X_i(t) \mid \mathbf{f}] = \epsilon_i(t) . \tag{45}$$

This holds since the nodes with causal depth 0 have no parent nodes, and their values are entirely controlled by noise. Hence, we obtain

$$\sum_{t=1}^T \mathbb{E}_{\mathbf{a}(t)} |[X_i(t) \mid \bar{\mathbf{f}}_t] - [X_i(t) \mid \mathbf{f}]| = \sum_{t=1}^T \mathbb{E}_{\mathbf{a}(t)} |\epsilon_i(t) - \epsilon_i(t)| = 0 . \tag{46}$$

In the next step, we show that for all nodes $i \in [N]$ with causal depth at least 1, i.e., $L_i \geq 1$, we have the following upper bound on the cumulative average error terms:

$$\sum_{t=1}^T \mathbb{E}_{\mathbf{a}(t)} \left| [X_i(t) \mid \bar{\mathbf{f}}_t] - [X_i(t) \mid \mathbf{f}] \right| \leq \mathcal{B}(\mathcal{F}, \delta) \sum_{\ell=1}^{L_i} d^{\ell-1} \prod_{k=2}^{\ell} K^{(k)}, \quad (47)$$

in which we adopt the convention of $\prod_{k=2}^1 K^{(k)} = 1$. We establish this via induction on the causal depth L_i .

Base Step: $L_i = 1$. For node $i \in [N]$ with causal depth $L_i = 1$, we show that

$$\sum_{t=1}^T \mathbb{E}_{\mathbf{a}(t)} \left| [X_i(t) \mid \bar{\mathbf{f}}_t] - [X_i(t) \mid \mathbf{f}] \right| \leq \mathcal{B}(\mathcal{F}, \delta). \quad (48)$$

When the causal path of a node is $L_i = 1$, according to SCM defined in (2), we have the following expansion:

$$\mathbb{E}_{\mathbf{a}(t)} \left| [X_i(t) \mid \bar{\mathbf{f}}_t] - [X_i(t) \mid \mathbf{f}] \right| = \mathbb{E}_{\mathbf{a}(t)} \left| \bar{f}_{i,t}(X_{\text{pa}(i)}(t); a_i(t)) - f_i(X_{\text{pa}(i)}(t); a_i(t)) \right| \quad (49)$$

$$= \mathbb{E}_{\mathbf{a}(t)} [\Delta_t(X_{\text{pa}(i)}(t); a_i(t))], \quad (50)$$

where we have defined the difference function Δ_t for input $Z_i(t) \in \mathcal{Z}_i$ as

$$\Delta_t(Z_i(t)) \triangleq \left| \bar{f}_{i,t}(Z_i(t)) - f_i(Z_i(t)) \right|. \quad (51)$$

Next, we provide the following lemma for bounding the sum of difference functions over time.

Lemma A.1. *For each node $i \in [N]$, and any given sequence $\{Z_i(t) : t \in [T]\}$, where we had defined $Z_i(t) \triangleq (X_{\text{pa}(i)}(t), a_i(t))$ we have*

$$\sum_{t=1}^T \Delta_t(Z_i(t)) \leq \mathcal{B}(\mathcal{F}, \delta). \quad (52)$$

Proof. We first leverage the following inequality

$$\Delta_t(Z_i(t)) \leq \sup_{f, \bar{f} \in \mathcal{C}_{i,t}} \bar{f}(Z_i(t)) - f(Z_i(t)). \quad (53)$$

Based on this inequality, we show the following inequality, which specifies the number of times that the error exceeds a certain threshold for any realization of noise vectors $\mathcal{E} = \{\epsilon(t), t \in [T]\}$ and actions vectors $\{\mathbf{a}(t), t \in [T]\}$. Specifically, we show that for the confidence sets $\mathcal{C}_{i,t}$ defined in (17), for any inputs sequence $Z_i(t) \in \mathcal{Z}_i$, for any horizon $T \in \mathbb{N}$, and any $\epsilon \in \mathbb{R}^+$, we almost surely have

$$\sum_{t=1}^T \mathbb{1}(\Delta_t(Z_i(t)) > \epsilon) \leq \left(\frac{4\beta_T}{\epsilon^2} + 1 \right) \dim(\mathcal{F}_i, \epsilon). \quad (54)$$

This inequality is a result of (Russo and Van Roy, 2013, Proposition 3) and (53). Next, we reorder the sequence $\{\Delta_t(Z_i(t)) : t \in [T]\}$ in the descending order denoted by $\Delta_{i_1} \geq \Delta_{i_2} \geq \dots \geq \Delta_{i_T}$. Subsequently, for any given α_i we have the following expansion:

$$\sum_{t=1}^T \Delta_t(Z_i(t)) = \sum_{t=1}^T \Delta_{i_t} \mathbb{1}\{\Delta_{i_t} \leq \alpha_i\} + \sum_{t=1}^T \Delta_{i_t} \mathbb{1}\{\Delta_{i_t} > \alpha_i\}. \quad (55)$$

Next, we provide upper bounds for the two summands of (55). For the first one, if $\alpha_i = \frac{1}{T}$, we have

$$\sum_{t=1}^T \Delta_{i_t} \mathbb{1}\{\Delta_{i_t} \leq \alpha_i\} \leq \sum_{t=1}^T \frac{1}{T} = 1. \quad (56)$$

On the other hand, when $\alpha_i > \frac{1}{T}$, it means that the minimum distance of any two functions is larger than $\frac{1}{T}$. This, in turn, indicates that event $\{\Delta_{i_t} \leq \alpha_i\}$ never occurs since the minimum gap between two functions is less than α_i . Hence, for $\alpha_i > \frac{1}{T}$ we have

$$\sum_{t=1}^T \Delta_{i_t} \mathbb{1}\{\Delta_{i_t} \leq \alpha_i\} = 0. \quad (57)$$

As a result, from (56) and (57) we have

$$\sum_{t=1}^T \Delta_{i_t} \mathbb{1}\{\Delta_{i_t} \leq \alpha_i\} \leq 1. \quad (58)$$

Next, we bound the second summand in (55). Based on Assumption 2.3 we know that

$$\Delta_{i_t} \leq C_i. \quad (59)$$

Furthermore, note that by the definition of Δ_{i_t} for any $\gamma > 0$ we have

$$\sum_{t=1}^T \mathbb{1}\{\Delta_{i_t} > \Delta_{i_s} - \gamma\} \geq s, \quad s \in [T]. \quad (60)$$

Combining the observations in (54) and (60) we have

$$s \leq \left(\frac{4\beta_T}{(\Delta_{i_s} - \gamma)^2} + 1 \right) \dim(\mathcal{F}_i, \Delta_{i_s} - \gamma), \quad s \in [T]. \quad (61)$$

Next, note that the eluder dimension is a non-increasing function of the discretization parameter. Thus, when $\Delta_{i_s} - \gamma \geq \alpha_i$ we have $\dim(\mathcal{F}_i, \Delta_{i_s} - \gamma) \leq \dim(\mathcal{F}_i, \alpha_i)$, resulting in

$$s \leq \left(\frac{4\beta_T}{(\Delta_{i_s} - \gamma)^2} + 1 \right) \dim(\mathcal{F}_i, \alpha_i), \quad s \in [T]. \quad (62)$$

By rearranging the terms in the above inequality, when $s > \dim(\mathcal{F}_i, \alpha_i)$, we find that for any $\gamma > 0$ have

$$\Delta_{i_s} - \gamma < \sqrt{\frac{4\beta_T \dim(\mathcal{F}_i, \alpha_i)}{s - \dim(\mathcal{F}_i, \alpha_i)}}, \quad s \in [T]. \quad (63)$$

By taking supremum of the left-hand side over γ in (63), we conclude that when $\Delta_{i_s} > \alpha_i$ and $s > \dim(\mathcal{F}_i, \alpha_i)$ we have

$$\Delta_{i_s} \leq \sqrt{\frac{4\beta_T \dim(\mathcal{F}_i, \alpha_i)}{s - \dim(\mathcal{F}_i, \alpha_i)}}, \quad s \in [T]. \quad (64)$$

Hence, we can expand and upper bound the second summand of (55) as follows.

$$\sum_{t=1}^T \Delta_{i_t} \mathbb{1}\{\Delta_{i_t} > \alpha_i\} = \sum_{t=1}^{\dim(\mathcal{F}_i, \alpha_i)} \Delta_{i_t} \mathbb{1}\{\Delta_{i_t} > \alpha_i\} + \sum_{t=\dim(\mathcal{F}_i, \alpha_i)+1}^T \Delta_{i_t} \mathbb{1}\{\Delta_{i_t} > \alpha_i\} \quad (65)$$

$$\stackrel{(59), (64)}{\leq} \min\{\dim(\mathcal{F}_i, \alpha_i), T\} C_i + \sum_{t=\dim(\mathcal{F}_i, \alpha_i)+1}^T \sqrt{\frac{4\beta_T \dim(\mathcal{F}_i, \alpha_i)}{t - \dim(\mathcal{F}_i, \alpha_i)}} \quad (66)$$

$$\leq \min\{\dim(\mathcal{F}_i, \alpha_i), T\} C_i + 2\sqrt{\beta_T \dim(\mathcal{F}_i, \alpha_i)} \sum_{t=\dim(\mathcal{F}_i, \alpha_i)+1}^T \frac{1}{\sqrt{t}} \quad (67)$$

$$\leq \min\{\dim(\mathcal{F}_i, \alpha_i), T\} C_i + 2\sqrt{\beta_T \dim(\mathcal{F}_i, \alpha_i)} \int_{t=0}^T \frac{1}{\sqrt{t}} dt \quad (68)$$

$$\leq \min\{\dim(\mathcal{F}_i, \alpha_i), T\} C_i + 4\sqrt{\dim(\mathcal{F}_i, \alpha_i) \beta_T T}. \quad (69)$$

Thus, in conclusion, based on (55), (58), and (69) we have

$$\sum_{t=1}^T \Delta_t(Z_i(t)) \leq 1 + \min\{\dim(\mathcal{F}, \alpha_i), T\} C_i + 4\sqrt{\dim(\mathcal{F}, \alpha_i) \beta_T T} \quad (70)$$

$$\leq \mathcal{B}(\mathcal{F}, \delta). \quad (71)$$

□

Subsequently,

$$\sum_{t=1}^T \mathbb{E}_{\mathbf{a}(t)} \left| [X_i(t) \mid \bar{\mathbf{f}}_t] - [X_i(t) \mid \mathbf{f}] \right| = \sum_{t=1}^T \mathbb{E}_{\mathbf{a}(t)} [\Delta_t(Z(t))] \leq \mathcal{B}(\mathcal{F}, \delta), \quad (72)$$

which proves the result for $L_i = 1$.

Induction Step. Assume that the property holds true for causal depths up to $L_i = m$. We show that it will also hold for $L_i = m + 1$. For this purpose, we start with the following expansion and apply the triangular inequality to find an upper bound for it.

$$\sum_{t=1}^T \mathbb{E}_{\mathbf{a}(t)} \left| [X_i(t) \mid \bar{\mathbf{f}}_t] - [X_i(t) \mid \mathbf{f}] \right| \quad (73)$$

$$= \sum_{t=1}^T \mathbb{E}_{\mathbf{a}(t)} \left| \bar{f}_{i,t}([X_{\text{pa}(i)}(t) \mid \bar{\mathbf{f}}_t]; a_i(t)) - f_i([X_{\text{pa}(i)}(t) \mid \mathbf{f}]; a_i(t)) \right| \quad (74)$$

$$\leq \sum_{t=1}^T \mathbb{E}_{\mathbf{a}(t)} \left| \bar{f}_{i,t}([X_{\text{pa}(i)}(t) \mid \bar{\mathbf{f}}_t]; a_i(t)) - \bar{f}_{i,t}([X_{\text{pa}(i)}(t) \mid \mathbf{f}]; a_i(t)) \right| \\ + \sum_{t=1}^T \mathbb{E}_{\mathbf{a}(t)} \left| \bar{f}_{i,t}([X_{\text{pa}(i)}(t) \mid \mathbf{f}]; a_i(t)) - f_i([X_{\text{pa}(i)}(t) \mid \mathbf{f}]; a_i(t)) \right| \quad (75)$$

$$\leq K_i \sum_{t=1}^T \mathbb{E}_{\mathbf{a}(t)} \left\| [X_{\text{pa}(i)}(t) \mid \bar{\mathbf{f}}_t] - [X_{\text{pa}(i)}(t) \mid \mathbf{f}] \right\| + \mathcal{B}(\mathcal{F}, \delta) \quad (76)$$

$$\leq K_i \sum_{j \in \text{pa}(i)} \sum_{t=1}^T \mathbb{E}_{\mathbf{a}(t)} \left| [X_j(t) \mid \bar{\mathbf{f}}_t] - [X_j(t) \mid \mathbf{f}] \right| + \mathcal{B}(\mathcal{F}, \delta). \quad (77)$$

where the transition to (75) holds due to the triangular inequality via adding and subtracting terms $\bar{f}_{i,t}([X_{\text{pa}(i)}(t) \mid \mathbf{f}]; a_i(t))$; (76) holds due to Lipschitz-continuity in conjunction with Lemma A.1; and (77) holds since the triangle inequality of L_2 norm.

Next, we find an upper bound on the first summand in (77). We notice that the summation is taken over all parents of node i , thus, we aim to find an upper bound for the error bound for each parent. Based on the induction assumption, for each node $j \in \text{pa}(i)$, we have

$$\sum_{t=1}^T \mathbb{E}_{\mathbf{a}(t)} \left| [X_j(t) \mid \bar{\mathbf{f}}_t] - [X_j(t) \mid \mathbf{f}] \right| \leq \mathcal{B}(\mathcal{F}, \delta) \sum_{\ell=1}^{L_j} d^{\ell-1} \prod_{k=2}^{\ell} K^{(k)} \quad (78)$$

$$\leq \mathcal{B}(\mathcal{F}, \delta) \sum_{\ell=1}^m d^{\ell-1} \prod_{k=2}^{\ell} K^{(k)}, \quad (79)$$

where (79) holds due to the causal depth of parent nodes is less than that of node i , i.e., $L_j \leq L_i - 1 = m$ for all $j \in \text{pa}(i)$. Subsequently, plugging (79) into (77), we obtain

$$K_i \sum_{j \in \text{pa}(i)} \sum_{t=1}^T \mathbb{E}_{\mathbf{a}(t)} \left| [X_j(t) \mid \bar{\mathbf{f}}_t] - [X_j(t) \mid \mathbf{f}] \right| \leq K_i \sum_{j \in \text{pa}(i)} \mathcal{B}(\mathcal{F}, \delta) \sum_{\ell=1}^m d^{\ell-1} \prod_{k=2}^{\ell} K^{(k)} \quad (80)$$

$$\leq d K_i \mathcal{B}(\mathcal{F}, \delta) \sum_{\ell=1}^m d^{\ell-1} \prod_{k=2}^{\ell} K^{(k)} \quad (81)$$

$$\leq \mathcal{B}(\mathcal{F}, \delta) \sum_{\ell=1}^m d^{\ell} \prod_{k=2}^{\ell+1} K^{(k)}, \quad (82)$$

where the inequality in (82) follows from $K_i \leq K^{(m+1)}$. Combining the results in (77) and (82), we conclude

$$\sum_{t=1}^T \mathbb{E}_{\mathbf{a}(t)} | [X_i(t) | \bar{\mathbf{f}}_t] - [X_i(t) | \mathbf{f}] | \leq \mathcal{B}(\mathcal{F}, \delta) \sum_{\ell=1}^{m+1} d^{\ell-1} \prod_{k=2}^{\ell} K^{(k)}, \quad (83)$$

which establishes the desired property in (47). Finally, by applying Jensen's inequality we find

$$\sum_{t=1}^T |\mathbb{E}_{\mathbf{a}(t)} [X_i(t) | \bar{\mathbf{f}}_t] - \mathbb{E}_{\mathbf{a}(t)} [X_i(t) | \mathbf{f}]| \leq \mathbb{E}_{\mathbf{a}(t)} \left| \sum_{t=1}^T [X_i(t) | \bar{\mathbf{f}}_t] - [X_i(t) | \mathbf{f}] \right| \quad (84)$$

$$\leq \mathcal{B}(\mathcal{F}, \delta) \sum_{\ell=1}^{L_i} d^{\ell-1} \prod_{k=2}^{\ell} K^{(k)}, \quad (85)$$

which completes the proof.

A.2 Proof of Theorem 5.3

We start by setting $\delta = \frac{1}{NT}$ and defining the event in which over T rounds, all confidence sets $\mathcal{C}_{i,t}$ contain the ground truth function f_i . Specifically,

$$\mathcal{E}_i \triangleq \{f_i \in \mathcal{C}_{i,t}, \forall t \in [T]\}. \quad (86)$$

Accordingly, define the event

$$\mathcal{E} \triangleq \bigcap_{i=1}^N \mathcal{E}_i. \quad (87)$$

According to Lemma 5.1, we know $\mathbb{P}(\mathcal{E}_i^c) \leq \frac{2}{NT}$. By invoking the union bound on probability, we have

$$\mathbb{P}(\mathcal{E}^c) \leq \sum_{i=1}^N \mathbb{P}(\mathcal{E}_i^c) \leq \sum_{i=1}^N \frac{2}{NT} = \frac{2}{T}. \quad (88)$$

Next, we decompose the regret defined in (6) under the events \mathcal{E} and \mathcal{E}^c .

$$\mathbb{E}[R(T)] = \sum_{t=1}^T [\mu_{\mathbf{a}^*} - \mu_{\mathbf{a}(t)}] \quad (89)$$

$$= \sum_{t=1}^T \mathbb{E}_{\mathbf{a}^*} [X_N | \mathbf{f}] - \mathbb{E}_{\mathbf{a}(t)} [X_N | \mathbf{f}] \quad (90)$$

$$= \sum_{t=1}^T \mathbb{E} \left[\mathbf{1}\{\mathcal{E}^c\} \underbrace{\left(\mathbb{E}_{\mathbf{a}^*} [X_N | \mathbf{f}] - \mathbb{E}_{\mathbf{a}(t)} [X_N | \mathbf{f}] \right)}_{\leq 2C_N} \right] + \sum_{t=1}^T \mathbb{E} \left[\mathbf{1}\{\mathcal{E}\} \underbrace{\left(\mathbb{E}_{\mathbf{a}^*} [X_N | \mathbf{f}] - \mathbb{E}_{\mathbf{a}(t)} [X_N | \mathbf{f}] \right)}_{\leq \text{UCB}_{\mathbf{a}^*}(t)} \right]$$

$$\leq 2C_N \sum_{t=1}^T \underbrace{\mathbb{P}(\mathcal{E}^c)}_{\leq \frac{2}{T}} + \sum_{t=1}^T \mathbb{E} \left[\text{UCB}_{\mathbf{a}(t)}(t) - \mathbb{E}_{\mathbf{a}(t)} [X_N | \mathbf{f}] \right] \quad (91)$$

$$\leq 4C_N + \sum_{t=1}^T \mathbb{E} \left[\mathbb{E}_{\mathbf{a}(t)} [X_N | \bar{\mathbf{f}}_t] - \mathbb{E}_{\mathbf{a}(t)} [X_N | \mathbf{f}] \right], \quad (92)$$

where we have used the set of inequalities

$$\mathbb{E}_{\mathbf{a}^*} [X_N | \mathbf{f}] \leq \text{UCB}_{\mathbf{a}^*}(t) \leq \text{UCB}_{\mathbf{a}(t)}(t) \leq \mathbb{E}_{\mathbf{a}(t)} [X_N | \bar{\mathbf{f}}_t], \quad (93)$$

in which $\bar{\mathbf{f}}_t$ is a function that achieves the upper bound $\text{UCB}_{\mathbf{a}(t)}(t)$, i.e.,

$$\bar{\mathbf{f}}_t \in \arg \max_{\mathbf{f} \in \mathcal{C}_t} \text{UCB}_{\mathbf{a}(t)}(t). \quad (94)$$

Then using the lemma 5.2, we have

$$\mathbb{E}[R(T)] \leq 4C_N + \mathcal{B}(\mathcal{F}, \delta) \sum_{\ell=1}^L d^{\ell-1} \prod_{k=2}^{\ell} K^{(k)} \quad (95)$$

$$= \mathcal{O}\left(\prod_{\ell=2}^L K^{(\ell)} d^{L-1} \sqrt{T \dim(\mathcal{F}) \log(NT \text{cn}(\mathcal{F}))}\right) \quad (96)$$

$$= \mathcal{O}\left(K d^{L-1} \sqrt{T \dim(\mathcal{F}) \log(NT \text{cn}(\mathcal{F}))}\right), \quad (97)$$

where the (96) is due to the fact that we have $\alpha_i \leq \frac{1}{T}$ and the fact that for $d \geq 2$

$$\sum_{\ell=1}^L d^{\ell-1} = \frac{d^L - 1}{d - 1} \leq 2d^{L-1} = \tilde{\mathcal{O}}(d^{L-1}), \quad (98)$$

and (97) holds due to the definition $K = \prod_{\ell=2}^L K^{(\ell)}$.

B Regret Upper Bounds for Special SCMs

We start by providing an equivalent definition of eluder dimension $\dim(\mathcal{F}_i, \epsilon)$, which follows the definitions of ϵ -dependence and ϵ -eluder dimension in Definition 3.1 and Definition 3.2. For a class of functions \mathcal{F}_i with input space \mathcal{Z}_i , an ordered sequence of m inputs $\mathcal{Z}_i(n) = \{Z_i(m) \in \mathcal{Z}_i : m \in [n]\}$, and $\epsilon \in \mathbb{R}^+$, for some $\epsilon' \geq \epsilon$ we define

$$W_s(\mathcal{F}_i, \mathcal{Z}_i(n)) \triangleq \sup \left\{ \left| f(Z_i(s)) - \tilde{f}(Z_i(s)) \right| \mid f, \tilde{f} \in \mathcal{F}_i, \sqrt{\sum_{m=1}^{s-1} (f(Z_i(m)) - \tilde{f}(Z_i(m)))^2} \leq \epsilon' \right\}. \quad (99)$$

The eluder dimension $\dim(\mathcal{F}_i, \epsilon)$ is the length of the longest ordered sequence $\mathcal{Z}_i(n)$ such that we have $W_s(\mathcal{F}_i, \mathcal{Z}_i(n)) \geq \epsilon'$ for all $t \in [n]$. Throughout the rest of this section, when it is clear from the context, we suppress the dependence of $W_s(\mathcal{F}_i, \mathcal{Z}_i(n))$ on $\mathcal{Z}_i(n)$ and \mathcal{F}_i , and use the shorthand W_s .

B.1 Proof of Theorem 5.5 (Linear SCMs)

Lemma B.1. *Consider the class of linear SCMs $\{\mathcal{F}_i : i \in [N]\}$ specified in (35) with the binary intervention space $\mathcal{A}_i = \{0, 1\}$. The eluder dimension of \mathcal{F}_i satisfies*

$$\dim(\mathcal{F}_i) = \tilde{\mathcal{O}}(d_i \log(T)). \quad (100)$$

Proof. For the selected class, we have

$$f_i(X_{\text{pa}(i)}; a_i) = \langle \theta_i(1 - a_i) + a_i \bar{\theta}_i, X_{\text{pa}(i)} \rangle. \quad (101)$$

Consider two functions $f_i, \tilde{f}_i \in \mathcal{F}_i$ specified by model parameters $(\theta_i, \bar{\theta}_i)$ and $(\psi_i, \bar{\psi}_i)$. Accordingly, define

$$\Delta = \theta_i - \psi_i, \quad \text{and} \quad \bar{\Delta} = \bar{\theta}_i - \bar{\psi}_i. \quad (102)$$

Leveraging the norm bound on the model parameters specified in (35), we have $\|\Delta\| \leq 2K_i$ and $\|\bar{\Delta}\| \leq 2K_i$. We also define the following Gram matrices:

$$\Phi_{s,0} \triangleq \sum_{m=1}^{s-1} \mathbf{1}\{a_i(m) = 0\} X_{\text{pa}(i)}(m) X_{\text{pa}(i)}^\top(m), \quad (103)$$

$$\Phi_{s,1} \triangleq \sum_{m=1}^{s-1} \mathbf{1}\{a_i(m) = 1\} X_{\text{pa}(i)}(m) X_{\text{pa}(i)}^\top(m). \quad (104)$$

By setting

$$\lambda \triangleq \left(\frac{\epsilon'}{2K_i} \right)^2, \quad (105)$$

we also define the $d_i \times d_i$ matrices

$$V_{t,0} \triangleq \Phi_{t,0} + \lambda I, \quad \text{and} \quad V_{t,1} \triangleq \Phi_{t,1} + \lambda I. \quad (106)$$

Based on (103) and (104) we also have

$$\sum_{m=1}^{s-1} \left(f_i(X_{\text{pa}(i)}(m); a_i(m)) - \tilde{f}_i(X_{\text{pa}(N)}(m); a_i(m)) \right)^2 = \Delta^\top \Phi_{s,0} \Delta + \bar{\Delta}^\top \Phi_{s,1} \bar{\Delta}. \quad (107)$$

The proof involves bounding the number of times that the event $W_s > \epsilon'$ occurs for any ordered sequence $\{Z_i(m) \in \mathcal{Z}_i : m \in [n]\}$, and it consists of three steps.

Step 1: We start by showing that if for a given $t \in \mathbb{N}$, the event $W_t \geq \epsilon'$ implies that

$$X_{\text{pa}(i)}^\top(t) V_{t,a_i(t)}^{-1} X_{\text{pa}(i)}(t) \geq \frac{1}{2}. \quad (108)$$

We provide the proof for $a_i(t) = 0$, and that of $a_i(t) = 1$ follows the same steps. We first find an upper bound on W_t as follows.

$$W_t = \sup \left\{ \left| \langle \Delta(1 - a_i(t)) + a_i(t)\bar{\Delta}, X_{\text{pa}(i)}(t) \rangle \right| \mid \|\Delta\| \leq 2K_i, \|\bar{\Delta}\| \leq 2K_i, \Delta^\top \Phi_{t,0} \Delta + \bar{\Delta}^\top \Phi_{t,1} \bar{\Delta} \leq (\epsilon')^2 \right\}. \quad (109)$$

$$= \sup \{ |\langle \Delta, X_{\text{pa}(i)}(t) \rangle| \mid \|\Delta\| \leq 2K_i, \|\bar{\Delta}\| \leq 2K_i, \Delta^\top \Phi_{t,0} \Delta + \bar{\Delta}^\top \Phi_{t,1} \bar{\Delta} \leq (\epsilon')^2 \} \quad (110)$$

$$\leq \sup \{ |\langle \Delta, X_{\text{pa}(i)}(t) \rangle| \mid \Delta^\top \Phi_{t,0} \Delta \leq (\epsilon')^2 \} \quad (111)$$

$$\leq \sup \{ |\langle \Delta, X_{\text{pa}(i)}(t) \rangle| \mid \Delta^\top V_{t,0} \Delta \leq 2(\epsilon')^2 \} \quad (112)$$

$$= \sqrt{2(\epsilon')^2} \|X_{\text{pa}(i)}(t)\|_{V_{t,0}^{-1}}, \quad (113)$$

where we use the definition of w_t in (109), we set $a_i(t) = 1$ in (110), and we relax the conditions in (111) by dropping the norm constraints and noting that $\Phi_{t,1}$ is positive semidefinite. The last equality in (112) holds due to $\Delta^\top \lambda I \Delta \leq (\epsilon')^2$ and we relax the condition. Hence, the event $W_t \geq \epsilon'$ implies that

$$\|X_{\text{pa}(i)}(t)\|_{V_{t,0}^{-1}} \geq \frac{1}{2}. \quad (114)$$

Step 2: In this step, we provide upper and lower bounds for the determinants of matrices $V_{t,0}$ and $V_{t,1}$. For this purpose, let us define the counting numbers $N_0(t)$ and $N_1(t)$ as the numbers that count the number of interventions 0 and 1 chosen up to time, i.e., $t - 1$, respectively.

$$N_0(t) = \sum_{s=1}^{t-1} \mathbb{1}\{a_i(s) = 0\}, \quad \text{and} \quad N_1(t) = \sum_{s=1}^{t-1} \mathbb{1}\{a_i(s) = 1\}. \quad (115)$$

Then, if $W_s \geq \epsilon'$ holds for all $s \in [t]$ we have the following inequalities.

$$\det(V_{t,0}) \geq \lambda^{d_i} \left(\frac{3}{2} \right)^{N_0(t)}, \quad \text{and} \quad \det(V_{t,1}) \geq \lambda^{d_i} \left(\frac{3}{2} \right)^{N_1(t)}. \quad (116)$$

From the definition of $V_{t,0}$ in (106), for all $s \in [t]$ we have

$$V_{s,0} = V_{s-1,0} + \mathbb{1}\{a_i(s-1) = 0\} X_{\text{pa}(i)}^\top(s-1) X_{\text{pa}(i)}(s-1). \quad (117)$$

By using the matrix determinant lemma, we obtain

$$\det(V_{s,0}) = \begin{cases} \det(V_{s-1,0}) \left(1 + \|X_{\text{pa}(i)}(s-1)\|_{V_{s-1,0}^{-1}}\right) \geq \frac{3}{2} \det(V_{s,0}) & \text{if } a_i(s-1) = 0 \\ \det(V_{s-1,0}) & \text{if } a_i(s-1) = 1 \end{cases}, \quad (118)$$

where for the case of $a_i(s-1) = 0$ we use the results in Step 1 stating that $\|X_{\text{pa}(i)}(s)\|_{V_{s,0}^{-1}} \geq \frac{1}{2}$ for all $s \in [t]$. By noting that for the initial value of $s = 0$ we have $\det(V_{0,0}) = \det(\lambda I) = \lambda^{d_i}$, by recursively applying (118) we have

$$\det(V_{t,0}) \geq \lambda^{d_i} \left(\frac{3}{2}\right)^{N_0(t)}. \quad (119)$$

The results for $V_{t,1}$ follow the same line of arguments. Since matrices $V_{t,0}$ and $V_{t,1}$ are symmetric and positive-definite, we have the following arithmetic-geometric inequality and leveraging (106) we have

$$\det(V_{t,0}) \leq \left(\frac{\text{tr}(V_{t,0})}{d_i}\right)^{d_i} \leq \left(\frac{N_0(t)}{d_i}c + \lambda\right)^{d_i}, \quad \text{and} \quad \det(V_{t,1}) \leq \left(\frac{\text{tr}(V_{t,1})}{d_i}\right)^{d_i} \leq \left(\frac{N_1(t)}{d_i}c + \lambda\right)^{d_i}. \quad (120)$$

where we have set $c = \sup_{X_{\text{pa}(i)}} \|X_{\text{pa}(i)}\|^2$. Next, recall that $\det(V_{t,0})$ is the product of the eigenvalues of $V_{t,0}$, whereas the trace $\text{tr}(V_{t,0})$ is the summation of the eigenvalues. As $\det(V_{s,0})$ is positive definite, the determinant is maximized when all eigenvalues are equal

Step 3: Based on the determinant bounds in (119) and (120), we can find the following upper bound on $N_0(t)$:

$$\left(\frac{3}{2}\right)^{\frac{N_0(t)}{d_i}} \leq \frac{N_0(t)}{d_i} \bar{c} + 1, \quad (121)$$

where we have defined $\bar{c} = c/\lambda$. To find an upper bound on $N_0(t)$, we find the largest $N_0(t)$ that satisfies (121). For this purpose, we define the function

$$B(x) \triangleq \max\{y \geq 1 : (1+x)^y \leq \bar{c}y + 1\}, \quad (122)$$

based on which we have $N_0(t) \leq d_i B(\frac{1}{2})$. Next, we find an explicit upper bound on $B(x)$ for $x \leq 1$. By noting that $\log(1+x) \geq \frac{x}{1+x}$ for $x \leq 1$, from the definition of $B(x)$ we have

$$B(x) \cdot \frac{x}{1+x} \leq B(x) \log(1+x) \quad (123)$$

$$\leq \log(\bar{c}B(x) + 1) \quad (124)$$

$$\leq \log(\bar{c}B(x) + \bar{B}(x)) \quad (125)$$

$$= \log B(x) + \log(1 + \bar{c}) \quad (126)$$

$$= \log \left(B(x) \cdot \frac{x}{1+x} \right) + \log \frac{1+x}{x} + \log(1 + \bar{c}) \quad (127)$$

$$\leq B(x) \cdot \frac{x}{1+x} \cdot \frac{1}{e} + \log \frac{1+x}{x} + \log(1 + \bar{c}). \quad (128)$$

By rearranging the terms, we have

$$B(x) \cdot \frac{x}{1+x} \left(1 - \frac{1}{e}\right) \log \frac{1+x}{x} + \log(1 + \bar{c}), \quad (129)$$

or equivalently

$$B(x) \leq \frac{1+x}{x} \cdot \frac{e}{e-1} \left(\log \frac{1+x}{x} + \log(1 + \bar{c}) \right), \quad (130)$$

which specifies an upper bound on $B(x)$ in terms of x . By setting $x = \frac{1}{2}$ we obtain

$$N_0(t) \leq d_i B\left(\frac{1}{2}\right) \leq 4.8 \left(1.1 + \log(1 + \bar{c})\right) d_i. \quad (131)$$

Following the same steps, we can find a similar bound for $N_1(t)$, based on which by recalling the definitions $\bar{c} = c/\lambda$, $\lambda = (\epsilon'/2K_i)^2$, and by setting $\epsilon' = \frac{1}{T}$ we obtain

$$\dim(\mathcal{F}_i, \epsilon) \leq N_0(t) + N_1(t) \leq 9.6 \left(1.1 + \log(1 + \bar{c}) \right) d_i = \tilde{\mathcal{O}}(d_i \log(T)) . \quad (132)$$

□

Lemma B.2 (Covering number.). *Consider the class of linear SCMs $\{\mathcal{F}_i : i \in [N]\}$ specified in (35) with the binary intervention space $\mathcal{A}_i = \{0, 1\}$. The covering number of \mathcal{F}_i satisfies*

$$\log \text{cn}(\mathcal{F}_i) = \tilde{\mathcal{O}}(d \log(T)) . \quad (133)$$

Proof. We adopt the following lemma (described in terms of the parameters in this paper) to find an upper bound on the covering number for linear SCMs.

Lemma B.3. (Russo and Van Roy, 2013) *Suppose the parameters of the Linear SCM class \mathcal{F}_i belong to the hypercube $[-A, A]^{2d_i}$. Then $\text{cn}_\alpha(\mathcal{F}_i) \leq (1 + 2Ac/\alpha)^{2d_i}$.*

By setting $A = K_i$, our hyper-sphere parameter space $\|\theta_i\| \leq K_i$ and $\|\bar{\theta}_i\| \leq K_i$ will be confined within the hypercube $[\theta_i, \bar{\theta}_i] \subset [-K_i, K_i]^{2d_i}$. Subsequently, the associate covering number will be upper-bounded by $(1 + 2cK_i/\alpha)^{2d_i}$. By setting $\alpha = \frac{1}{T}$ we have

$$\ln \text{cn}_\alpha(\mathcal{F}_i) \leq 2d_i \ln(1 + 2cK_i T) = \tilde{\mathcal{O}}(d \log(T)) . \quad (134)$$

□

Based on the results of Lemma B.1 and Lemma B.2, we the regret bound in Theorem 5.3 becomes

$$\mathbb{E}[R(T)] = \mathcal{O} \left(K d^{L-1} \sqrt{T \dim(\mathcal{F}) \log(NT \text{cn}(\mathcal{F}))} \right) \quad (135)$$

$$= \mathcal{O} \left(K d^{L-1} \sqrt{T d \log T \left(\log NT + d \log T \right)} \right) \quad (136)$$

$$= \mathcal{O}(K d^L \sqrt{T \log T \varpi_L}) , \quad (137)$$

where

$$\varpi_L = \log T \left[1 + \frac{\log N}{d \log T} \right] . \quad (138)$$

Note that the term d in (138) can be refined to $d + 1$, but that will not change the scaling behavior of the regret with respect to T .

The proofs presented for Lemma B.1 and Lemma B.2 can be readily extended to discrete and finite spaces. The main change in the proof pertains to repeating the arguments we presented for bounding $N_0(t)$ and $N_1(1)$, for all other interventions as well. This results in bounding the eluder dimension as

$$\dim(\mathcal{F}_i) = \sum_{j=1}^{|\mathcal{A}_i|} N_j(t) \leq |\mathcal{A}_i| 4.8 \left(1.1 + \ln(1 + \bar{c}) \right) d_i = \tilde{\mathcal{O}}(d_i \log(T)) . \quad (139)$$

Since the \mathcal{A}_i is finite, we can find a constant c_a such that $|a_{i,j}| \leq c_a$ covering number is minor by replacing c with $c + c_a$, which yields

$$\ln \text{cn}_\alpha(\mathcal{F}_i) \leq 2d_i \ln(1 + 2(c + c_a)K_i T) = \tilde{\mathcal{O}}(d \log(T)) . \quad (140)$$

B.2 Proof of Theorem 5.7 (Polynomial SCMs)

For $p = 2$ for the functions specified by (37) we have

$$f_i(X_{\text{pa}(i)}; a_i) = (\langle \theta_i, [X_{\text{pa}(i)}, a_i] \rangle)^2, \quad (141)$$

which can be recast as

$$f_i(X_{\text{pa}(i)}; a_i) = [X_{\text{pa}(i)}, a_i]^\top \Theta [X_{\text{pa}(i)}, a_i], \quad (142)$$

where $\Theta \in \mathbb{R}^{(d+1) \times (d+1)}$ is a matrix with element $\Theta_{j,k} = \theta_{i,j} \theta_{i,k}$ with $\theta_{i,j}$ is j -th element of θ_i . Hence, the associated class \mathcal{F}_i specified by (38) is a subspace of the space of quadratic function. Hence, the eluder dimension of \mathcal{F}_i can be upper bounded based on the following lemmas (Osband and Van Roy, 2014).

Lemma B.4. *For the class of polynomial SCMs defined in (38) with $p = 2$ and bounded intervention space such that $|a_i| \leq 1$ for all $a_i \in \mathcal{A}$, the eluder dimension satisfies*

$$\dim(\mathcal{F}_i, \epsilon) = \tilde{\mathcal{O}}((d+1)^2 \log(1/\epsilon)). \quad (143)$$

By noting that the set of parameters captured by Θ has $(d+1)^2$ degrees of freedom, thus we have

Lemma B.5. *For the class of polynomial SCMs defined in (38) with $p = 2$ and bounded intervention space such that $|a_i| \leq 1$ for all $a_i \in \mathcal{A}$, the covering number satisfies*

$$\log \text{cn}_\alpha(\mathcal{F}_i) = \tilde{\mathcal{O}}((d+1)^2 \log(1/\epsilon)). \quad (144)$$

Then with the choice of $\epsilon = \alpha_i = \frac{1}{T}$, we have

$$\dim(\mathcal{F}_i) = \tilde{\mathcal{O}}((d+1)^2 \log T), \quad \text{and } \log \text{cn}(\mathcal{F}_i) = \tilde{\mathcal{O}}((d+1)^2 \log T). \quad (145)$$

Then the regret bound in Theorem 5.3 becomes

$$\mathbb{E}[R(T)] = \mathcal{O} \left(K d^{L-1} \sqrt{T \dim(\mathcal{F}) \log(NT \text{cn}(\mathcal{F}))} \right) \quad (146)$$

$$= \mathcal{O} \left(K d^{L-1} \sqrt{T (d+1)^2 \log T \log(NT \times T^{(d+1)^2})} \right) \quad (147)$$

$$= \mathcal{O}(K d^{L+1} \sqrt{T \log T \varpi_P}), \quad (148)$$

where

$$\varpi_P = \log T \left[1 + \frac{\log N}{d \log T} \right]. \quad (149)$$

B.3 Proof of Theorem 5.9 (Neural Network SCMs)

Corresponding to the class of neural network SCMs specified by (43), we augment the causal graph \mathcal{G} to generate another causal graph $\tilde{\mathcal{G}}$ as follows. As shown in Figure 1, for each node $i \in [N]$, We insert $s \geq d+1$ nodes between the set of parents of node i and i and add directed edges to all parents of i to the s intermediate nodes and from the s intermediate nodes to node i . Hence, graph $\tilde{\mathcal{G}}$ has $N + s|\{i : \text{pa}(i) \neq \emptyset\}|$ nodes, its maximum in-degree is s , and its longest causal path has length $2L$.

In the augmented graph $\tilde{\mathcal{G}}$, we denote the data of the s intermediate nodes between $\text{pa}(i)$ and i by $\bar{\mathbf{X}}_i \in \mathbb{R}^s$ and for the data of node i , we keep the notation X_i . The SCM $\tilde{\mathcal{G}}$ follows a generalized linear model. Specifically consider node i with neural network parameters¹ Θ_1 and Θ_2 . If node i has an even causal depth, its SCM will be specified by $X_i = \sigma(\Theta_2 \bar{\mathbf{X}}_i)$, which is a generalized linear model. On the other hand, if node i has an odd causal depth, its SCM will be specified by $\bar{\mathbf{X}}_i = \sigma(\Theta_1 [a_i, X_{\text{pa}(i)}])$, which is also a generalized linear SCM. Note that the interventions will impact only the nodes with even causal depths. Their impact, subsequently, will be

¹We have different matrices Θ_1 and Θ_2 for different nodes, which we can specify by $\Theta_{i,1}$ and $\Theta_{i,2}$. However, we are suppressing the implicit dependence on i for notational clarity.

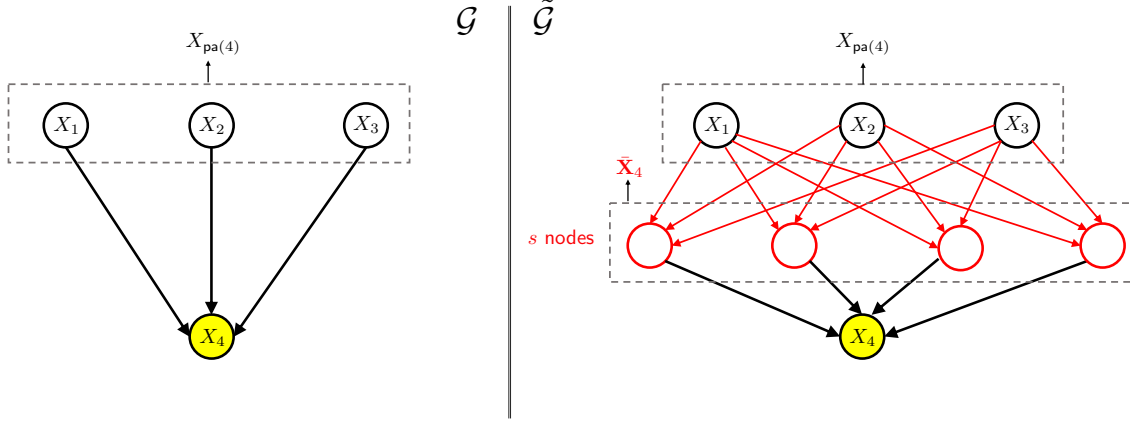


Figure 1: Construction of $\tilde{\mathcal{G}}$ based on \mathcal{G} by embedding s nodes between a node (node 4) and its parents (nodes 1,2,3).

passed by the intermediate nodes to the nodes with even causal depths, which are the nodes in the original graph \mathcal{G} . Finally, for any node i , if it has an odd causal depth, we denote its Lipschitz constants by $K_{i,o}$ and if it has an even causal depth, we denote its Lipschitz constants by $K_{i,e}$. Accordingly, for even i define

$$\bar{K}_i = sK_{i,e} \cdot K_{i-1,o}, \quad \forall i \in \{2\ell : \ell \in [N]\}, \quad (150)$$

and

$$\bar{K}^{(k)} = \max\{\bar{K}_i : L_i = 2k, i \in \{2\ell : \ell \in [N]\}\}, \quad \forall k \in [L], \quad (151)$$

so that $\bar{K}^{(k)} \geq K^{(k)}$ is an upper bound of Lipschitz constant.

We can adjust the GCB algorithms to form estimates and confidence sets for all the parameters in the augmented graph $\tilde{\mathcal{G}}$. We denote the class of generalized SCMs that specify the cause-effect relationships in $\tilde{\mathcal{G}}$ by \mathcal{H} . Based on this, we provide a counterpart of Lemma 5.2 for bounding the cumulative error at each node i due to the compounding estimate uncertainties along the causal paths.

Corollary B.6 (Compounding Error for Neural Networks). *Consider the neural network SCMs defined in (43). If $f_i \in \mathcal{C}_{i,t}$ for all $i \in [N]$ and $t \in [T]$, then the following error bound holds*

$$\sum_{t=1}^T \left| \mathbb{E}_{\mathbf{a}(t)}[X_i | \bar{\mathbf{f}}_t] - \mathbb{E}_{\mathbf{a}(t)}[X_i | \mathbf{f}] \right| \leq \mathcal{B}(\mathcal{H}, \delta) \sum_{\ell=1}^{\frac{L_i}{2}} (d+1)^{\ell-1} \prod_{k=2}^{\ell} \bar{K}^{(k)}. \quad (152)$$

Proof. This proof follows the flow of the one in Lemma 5.2, with the key difference being the addition of an intermediate layer between nodes i and their parents. Consequently, each induction step is now comprised of two stages: firstly, upper bounding the error from the intermediate nodes $\bar{\mathbf{X}}_i$ to X_i , and secondly, upper bounding the error from $X_{\text{pa}(i)}$ to the intermediate nodes. The second stage aligns with the procedure in the proof of Lemma 5.2, contributing to $(d+1)K_{j,o}$ for node $X_j \in \bar{\mathbf{X}}_i$. As for the contribution from $\bar{\mathbf{X}}_i$ to X_i , it amounts to $sK_{i,e}$. This results in the error bound presented in (152), with the Lipschitz constants $\bar{K}^{(k)}$ defined in (151). It is noteworthy that in the special case where $s = d+1$, our results reduce to the one in Lemma 5.2 with a maximum causal depth of $2L$. \square

Then plug (152) into (92) in the proof of Theorem 5.3, the upper bound in Theorem 5.3 becomes

$$\mathbb{E}[R(T)] \leq 4C_N + \mathcal{B}(\mathcal{H}, \delta) \sum_{\ell=1}^L (d+1)^{\ell-1} \prod_{k=2}^{\ell} \bar{K}^{(k)} \quad (153)$$

$$= \mathcal{O}\left(\bar{K} d^{L-1} \sqrt{T \dim(\mathcal{H}) \log(NT \text{cn}(\mathcal{H}))}\right). \quad (154)$$

Then, the eluder dimension and covering number of the class of functions \mathcal{H} upper bounded by the following lemma.

Lemma B.7. (Russo and Van Roy, 2013) For the generalized linear SCM class of functions \mathcal{H} with degree d , we have

$$\dim(\mathcal{H}) = \tilde{\mathcal{O}}(rd \log T), \quad \text{and} \quad \log \text{cn}(\mathcal{H}) = \tilde{\mathcal{O}}(rd \log T), \quad (155)$$

where r is defined as the ratio of maximum gradient over minimum derivative of the activation function²

$$r = \frac{\sup \nabla \sigma(\cdot)}{\inf \nabla \sigma(\cdot)}. \quad (156)$$

Next, since the input dimension is at most $d+1$ for hidden layers and s for the outputs X_i for $i \in [N]$. Applying Lemma B.7 in the result in (154), we conclude the regret is upper bounded by

$$\mathbb{E}[R(T)] = \mathcal{O}\left(r \bar{K} d^{L-1} \sqrt{T \dim(\mathcal{H}) \log(NT \text{cn}(\mathcal{H}))}\right) \quad (157)$$

$$= \mathcal{O}\left(r \bar{K} d^{L-1} \sqrt{T \log T \varpi_N}\right), \quad (158)$$

where

$$\varpi_N = \log T \max d, s \left[1 + \frac{\log N}{d \log T}\right]. \quad (159)$$

C Regret Lower Bounds




Figure 2: Hierarchical graph with degree d and maximum causal path length L .

Proof Sketch: We find lower bounds for the term in (33) by lower bounding the inner supremum terms $\sup_{\mathcal{G} \in \bar{\mathcal{G}}} \mathbb{E}[R(T)]$, which captures the model with the worst-case regret. We can do this by shrinking the graph set $\bar{\mathcal{G}}$ to a set of only two possible graph instances $\{\mathcal{G}_1, \mathcal{G}_2\}$, for which the regret will not exceed that of the worst case regret over all possible choices in $\bar{\mathcal{G}}$. For this purpose, we choose a hierarchical graph shown in Figure 2. These two graphs consist of $L+1$ layers. Each of the first L layers has d nodes and the nodes within a layer are not connected. The nodes between two adjacent layers are fully connected. Finally, the last layer consists of one node (reward node) that is fully connected to the nodes in layer L . We use the bounded noise to ensure the boundness of the causal system. Subsequently, we provide an information-theoretic analysis that captures the hardness of distinguishing the two models, and facilitates finding a universal minimax lower bound for each class for each special family of models. Throughout the analysis, we will be using the following lemma (Lattimore and Szepesvári, 2020, Theorem 14.2).

²This can be generalized to the maximum and minimum derivative within the input space.

Lemma C.1 (Bretagnolle-Huber Inequality). *Let \mathbb{P}_1 and \mathbb{P}_2 be two probability measures on the same measurable space (Ω, \mathcal{F}) and let $A \in \mathcal{F}$ be an arbitrary event. Then,*

$$\mathbb{P}_1(A) + \mathbb{P}_2(A^c) \geq \frac{1}{2} \exp(-D_{\text{KL}}(\mathbb{P}_1 \parallel \mathbb{P}_2)) , \quad (160)$$

where D_{KL} denotes the Kullback–Leibler divergence.

C.1 Proof of Theorem 5.6 (Linear SCMs)

We consider two linear SCM causal bandit instances \mathcal{G}_1 and \mathcal{G}_2 that share the same hierarchical graph specified earlier (Figure 2) with degree d and maximum causal path length L . Subsequently, our minimax lower bound of interest is lower bounded by

$$\inf_{\pi \in \Pi} \sup_{\mathcal{F} \in \bar{\mathcal{F}}} \sup_{\mathcal{G} \in \bar{\mathcal{G}}} \mathbb{E}[R(T)] \geq \inf_{\pi \in \Pi} \sup_{\mathcal{G} \in \{\mathcal{G}_1, \mathcal{G}_2\}} \mathbb{E}[R(T)] . \quad (161)$$

We assume the two bandit instances share same weights $\{\theta_i : i \in [N]\}$ and $\{\bar{\theta}_i : i \in [N]\}$, which are set to

$$\theta_i = \underbrace{\left[\frac{K^{(L_i)}}{\sqrt{d_i}}, \dots, \frac{K^{(L_i)}}{\sqrt{d_i}} \right]}_{d_i \text{ times}} , \quad \text{and} \quad \bar{\theta}_i = \underbrace{\left[\frac{K^{(L_i)}}{\sqrt{d_i}} - \delta, \dots, \frac{K^{(L_i)}}{\sqrt{d_i}} - \delta \right]}_{d_i \text{ times}} , \quad (162)$$

where $\delta > 0$ is a small value to be chosen later. The intervention sets for two bandit instances \mathcal{G}_1 and \mathcal{G}_2 are set to binary, i.e. $\mathcal{A}_i = \{0, 1\}$. The difference between the two bandit instances is the noise distribution in the system. For the noise for the node 1, in first bandit instance \mathcal{G}_1 , it is defined as

$$\epsilon_1 = \begin{cases} \text{Bern}(1/2 + \delta) & \text{if } a_1 = 0 \\ \text{Bern}(1/2) & \text{if } a_1 = 1 \end{cases} , \quad (163)$$

where $\text{Bern}(p)$ denotes the Bernoulli random distribution with probability p . The noise is reversed in the second bandit instance \mathcal{G}_2 , which is

$$\epsilon_1 = \begin{cases} \text{Bern}(1/2) & \text{if } a_1 = 0 \\ \text{Bern}(1/2 + \delta) & \text{if } a_1 = 1 \end{cases} . \quad (164)$$

For the noise terms associated with the rest of the random variables, we assume that they follow the the following Bernoulli model.

$$\epsilon_i = \begin{cases} -1 & \text{with probability } \frac{1}{2} \\ +1 & \text{with probability } \frac{1}{2} \end{cases} , \quad \text{for } i \neq 1 . \quad (165)$$

Since the weights are non-negative and the noise terms are zero-mean for $i \neq 1$, the optimal action is the one that maximizes each weight and the average value of the node 1. As a result, the optimal action at node 1 is the action that maximizes the mean of noise, while the optimal actions at other nodes are the actions that maximize the weights, i.e., $a_i = 0$ for $i \neq 1$. Hence, the optimal intervention for bandit instances \mathcal{G}_1 and \mathcal{G}_2 are

$$\mathbf{a}_{\mathcal{G}_1}^* = (0, 0, \dots, 0) , \quad \text{and} \quad \mathbf{a}_{\mathcal{G}_2}^* = (1, 0, \dots, 0) . \quad (166)$$

Next, consider a fixed bandit policy π that generates the filtration \mathcal{H}_t over time. The decision of π at time t is \mathcal{H}_{t-1} -measurable. Accordingly, define the $\mathbb{P}_{\mathcal{G}_1}$ and $\mathbb{P}_{\mathcal{G}_2}$ as the probability measures induced by \mathcal{H}_T by T rounds of interactions of π with \mathcal{G}_1 and \mathcal{G}_2 , respectively. We show that for any algorithm π , it cannot get small regret in both causal bandit instances simultaneously. Since the difference between two bandit instances is on the first node, we are interested in the event in which $a_1 = 1$ is chosen at least $\frac{T}{2}$ times after T rounds of interactions. We denote this event by

$$\mathcal{E}_\pi \triangleq \left\{ \sum_{i=1}^T \mathbb{1}\{a_1(t) = 1\} \geq \frac{T}{2} \right\} . \quad (167)$$

We observe that algorithm π will incur a regret when choosing intervention with $a_1 = 1$ on bandit instance \mathcal{G}_1 . Furthermore, since the weights in the causal system are positive, we have $\mu_{\mathbf{a}_{\mathcal{G}_2}^*} \geq \mu_{\mathbf{a}}$ for all intervention \mathbf{a} with

$a_1 = 1$. Therefore, for \mathcal{G}_1 , the expected value of instantaneous regret is at least $\mu_{\mathbf{a}_{\mathcal{G}_1}^*} - \mu_{\mathbf{a}_{\mathcal{G}_2}^*}$ if $a_i(t) = 1$. Besides, since ϵ_i has a zero mean for $i \geq 2$, only the paths that start at node 1 and end at the reward node N contribute to the expected regret. Thus, for the gap between $\mu_{\mathbf{a}_{\mathcal{G}_1}^*}$ and $\mu_{\mathbf{a}_{\mathcal{G}_2}^*}$, each path contributes $\prod_{\ell=1}^L K^{(\ell)} d^{-\frac{\ell}{2}} \delta$ and we have d^{L-1} such paths. Therefore, we have

$$\mu_{\mathbf{a}_{\mathcal{G}_1}^*} - \mu_{\mathbf{a}_{\mathcal{G}_2}^*} = d^{\frac{L}{2}-1} \delta \prod_{\ell=1}^L K^{(\ell)}. \quad (168)$$

Consequently, we can lower bound the regret associated with \mathcal{G}_1 by conditioning on the event \mathcal{E}_π being true, i.e.,

$$\mathbb{E}_{\mathcal{G}_1}[R(T)] = \mathbb{E}_{\mathcal{G}_1} \left[\sum_{t=1}^T (\mu^* - \mu_{\mathbf{a}(t)}) \right] \quad (169)$$

$$\geq \mathbb{E}_{\mathcal{G}_1} \left[\sum_{t \in [T]: a_1(t)=1} (\mu_{\mathbf{a}_{\mathcal{G}_1}^*} - \mu_{\mathbf{a}(t)}) \right] \quad (170)$$

$$\geq \mathbb{E}_{\mathcal{G}_1} \left[\sum_{t \in [T]: a_1(t)=1} \mathbb{1}\{\mathcal{E}_\pi\} (\mu_{\mathbf{a}_{\mathcal{G}_1}^*} - \mu_{\mathbf{a}_{\mathcal{G}_2}^*}) \right] + \mathbb{E}_{\mathcal{G}_1} \left[\sum_{t \in [T]: a_1(t)=1} \mathbb{1}\{\mathcal{E}_\pi^c\} (\mu_{\mathbf{a}_{\mathcal{G}_1}^*} - \mu_{\mathbf{a}_{\mathcal{G}_2}^*}) \right] \quad (171)$$

$$\geq \mathbb{P}_{\mathcal{G}_1}(\mathcal{E}_\pi) \frac{T}{2} (\mu_{\mathbf{a}_{\mathcal{G}_1}^*} - \mu_{\mathbf{a}_{\mathcal{G}_2}^*}) \quad (172)$$

$$= \mathbb{P}_{\mathcal{G}_1}(\mathcal{E}_\pi) \frac{T}{2} d^{\frac{L}{2}-1} \delta \prod_{\ell=1}^L K^{(\ell)}, \quad (173)$$

where (170) lower bound the regret by considering the time instance that regret incurred when $a_1(t) = 1$; (171) holds due to $\mu_{\mathbf{a}_{\mathcal{G}_2}^*} \geq \mu_{\mathbf{a}}$ for all intervention \mathbf{a} with $a_1 = 1$ and we split the sum based on whether \mathcal{E}_π holds; and (172) uses the definition of \mathcal{E}_π , lower bounding for the occurrence time of regret by $\frac{T}{2}$ for \mathcal{E}_π , and by 0 for the \mathcal{E}_π^c . The last equality in (173) holds due to equation (168).

Similarly, for bandit instance \mathcal{G}_2 , we have $\mu_{\mathbf{a}_{\mathcal{G}_1}^*} \geq \mu_{\mathbf{a}}$ for all intervention \mathbf{a} that $a_i = 0$. Therefore, for \mathcal{G}_2 , the expected value of instantaneous regret is at least $\mu_{\mathbf{a}_{\mathcal{G}_2}^*} - \mu_{\mathbf{a}_{\mathcal{G}_1}^*}$ if $a_i(t) = 0$. Hence, we can similarly show that

$$\mathbb{E}_{\mathcal{G}_2}[R(T)] \geq \mathbb{P}_{\mathcal{G}_2}(\mathcal{E}_\pi^c) \frac{T}{2} d^{\frac{L}{2}-1} \delta \prod_{\ell=1}^L K^{(\ell)}. \quad (174)$$

Combining (173) and (174), we have

$$\mathbb{E}_{\mathcal{G}_1}[R(T)] + \mathbb{E}_{\mathcal{G}_2}[R(T)] \geq (\mathbb{P}_{\mathcal{G}_1}(\mathcal{E}_\pi) + \mathbb{P}_{\mathcal{G}_2}(\mathcal{E}_\pi^c)) \frac{T}{2} d^{\frac{L}{2}-1} \delta \prod_{\ell=1}^L K^{(\ell)}, \quad (175)$$

Next, we leverage Lemma C.1 and find a lower bound on the term $\mathbb{P}_{\mathcal{G}_1}(\mathcal{E}_\pi) + \mathbb{P}_{\mathcal{G}_2}(\mathcal{E}_\pi^c)$, resulting in

$$\mathbb{E}_{\mathcal{G}_1}[R(T)] + \mathbb{E}_{\mathcal{G}_2}[R(T)] \geq \frac{T}{4} d^{\frac{L}{2}-1} \delta \exp(-D_{\text{KL}}(\mathbb{P}_{\mathcal{G}_1} \parallel \mathbb{P}_{\mathcal{G}_2})) \prod_{\ell=1}^L K^{(\ell)}. \quad (176)$$

It remains to compute the KL-divergence between two bandit instances \mathcal{G}_1 and \mathcal{G}_2 , which is formalized by the following lemma.

Lemma C.2. *For the causal bandit instance \mathcal{G}_1 and \mathcal{G}_2 , the KL-divergence between $\mathbb{P}_{\mathcal{G}_1}$ and $\mathbb{P}_{\mathcal{G}_2}$ is upper bounded by*

$$D_{\text{KL}}(\mathbb{P}_{\mathcal{G}_1} \parallel \mathbb{P}_{\mathcal{G}_2}) \leq T \log \left(\frac{1}{(1+2\delta)(1-2\delta)} \right). \quad (177)$$

Proof. The two bandit instances differ only in the mechanism of node 1, which is a root node. Hence, $D_{\text{KL}}(\mathbb{P}_{\mathcal{G}_1} \parallel \mathbb{P}_{\mathcal{G}_2})$ can be simplified as

$$D_{\text{KL}}(\mathbb{P}_{\mathcal{G}_1} \parallel \mathbb{P}_{\mathcal{G}_2}) = \sum_{i=1}^N D_{\text{KL}}(\mathbb{P}_{\mathcal{G}_1}(X_i \mid X_{\text{pa}(i)}) \parallel \mathbb{P}_{\mathcal{G}_2}(X_i \mid X_{\text{pa}(i)})) = D_{\text{KL}}(\mathbb{P}_{\mathcal{G}_1}(X_1) \parallel \mathbb{P}_{\mathcal{G}_2}(X_1)), \quad (178)$$

where the second equality holds because the noise terms for all nodes $i \geq 2$ are identical in \mathcal{G}_1 and \mathcal{G}_2 . Next, we calculate $D_{\text{KL}}(\mathbb{P}_{\mathcal{G}_1}(X_1) \parallel \mathbb{P}_{\mathcal{G}_2}(X_1))$ under two cases: when $a_1 = 0$ and $a_1 = 1$. By noting that the KL-divergence between two Bernoulli random variables is given by

$$D_{\text{KL}}(\text{Bern}(p) \parallel \text{Ber}(q)) = p \log\left(\frac{p}{q}\right) + (1-p) \log\left(\frac{1-p}{1-q}\right). \quad (179)$$

from the above we obtain

$$\begin{aligned} & D_{\text{KL}}(\mathbb{P}_{\mathcal{G}_1}(X_1) \parallel \mathbb{P}_{\mathcal{G}_2}(X_1)) \\ &= \sum_{t \in [T]: a_i(t)=0} D_{\text{KL}}(\text{Bern}(1/2 + \delta) \parallel \text{Ber}(1/2)) + \sum_{t \in [T]: a_i(t)=1} D_{\text{KL}}(\text{Bern}(1/2) \parallel \text{Ber}(1/2 + \delta)) \end{aligned} \quad (180)$$

$$= \sum_{s=1}^T \mathbb{1}\{a_1(t) = 0\} \left[\left(\frac{1}{2} + \delta\right) \log(1 + 2\delta) + \left(\frac{1}{2} - \delta\right) \log(1 - 2\delta) \right] \quad (181)$$

$$+ \sum_{s=1}^T \mathbb{1}\{a_1(t) = 1\} \frac{1}{2} \log\left(\frac{1}{(1+2\delta)(1-2\delta)}\right) \quad (182)$$

$$< \frac{T}{2} \log\left(\frac{1}{(1+2\delta)(1-2\delta)}\right), \quad (183)$$

where the last inequality holds since $(\frac{1}{2} + \delta) \log(1 + 2\delta) + (\frac{1}{2} - \delta) \log(1 - 2\delta) < \frac{1}{2} \log\left(\frac{1}{(1+2\delta)(1-2\delta)}\right)$ for $0 < \delta < 1/2$. \square

If we choose $\delta = \frac{1}{\sqrt{T}}$ to balance the terms in the lower bound, we obtain

$$\max\{\mathbb{E}_{\mathcal{G}_1}[R(T)], \mathbb{E}_{\mathcal{G}_2}[R(T)]\} \geq \frac{1}{2} (\mathbb{E}_{\mathcal{G}_1}[R(T)] + \mathbb{E}_{\mathcal{G}_2}[R(T)]) \quad (184)$$

$$\stackrel{(176)}{\geq} \frac{T}{8} d^{\frac{L}{2}-1} \delta \exp(-D_{\text{KL}}(\mathbb{P}_{\mathcal{G}_1} \parallel \mathbb{P}_{\mathcal{G}_2})) \prod_{\ell=1}^L K^{(\ell)} \quad (185)$$

$$\stackrel{(177)}{\geq} \frac{T}{8} d^{\frac{L}{2}-1} \delta [(1+2\delta)(1-2\delta)]^T \prod_{\ell=1}^L K^{(\ell)} \quad (186)$$

$$= \frac{1}{8} d^{\frac{L}{2}-1} \sqrt{T} \times \left(1 - \frac{4}{T}\right)^T \prod_{\ell=1}^L K^{(\ell)}. \quad (187)$$

for $T \geq 5$, the term $\left(1 - \frac{4}{T}\right)^T$ is an increasing function of T . Hence, for $T \geq 5$, we have the lower bound $\left(1 - \frac{4}{T}\right)^T \geq 0.2^5$. By setting $c = \frac{1}{8} \times 0.2^5$, we have

$$\max\{\mathbb{E}_{\mathcal{G}_1}[R(T)], \mathbb{E}_{\mathcal{G}_2}[R(T)]\} \geq cK d^{\frac{L}{2}-1} \sqrt{T}. \quad (188)$$

Hence, the policy π incurs a regret $cK d^{\frac{L}{2}-1} \sqrt{T}$ in at least one of the two bandit instances.

C.2 Proof of Theorem 5.8 (Polynomial SCMs)

Similarly to the linear SCM setting, we consider two polynomial SCM causal bandit instances \mathcal{G}_1 and \mathcal{G}_2 that share the same hierarchical graph specified earlier (Figure 2) with degree d and maximum causal path length L . Subsequently, the minimax regret of interest is lower bounded by

$$\inf_{\pi \in \Pi} \sup_{\mathcal{F} \in \bar{\mathcal{F}}} \sup_{\mathcal{G} \in \bar{\mathcal{G}}} \mathbb{E}[R(T)] \geq \inf_{\pi \in \Pi} \sup_{\mathcal{G} \in \{\mathcal{G}_1, \mathcal{G}_2\}} \mathbb{E}[R(T)]. \quad (189)$$

We assume the two bandit instances share same weights $\{\theta_i : i \in [N]\}$, which are set to be the same $\sqrt{\beta}$ and our function in polynomial SCM becomes

$$f_i(X_{\text{pa}(i)}; a_i) = \beta \left(\sum_{j \in \text{pa}(i)} X_j \right)^p, \quad (190)$$

where $\beta > 0$ is a constant defined later that controls the scaling behavior of the system, which we control later. The intervention sets for two bandit instances \mathcal{G}_1 and \mathcal{G}_2 are set to $\mathcal{A}_i = [0, 1]$. The difference between the two bandit instances is the noise distribution in the system. For the noise for the node 1, in first bandit instance \mathcal{G}_1 , it is defined as

$$\epsilon_1 = \begin{cases} \text{Bern}(1/2 + \delta) & \text{if } a_1 \leq \frac{1}{2} \\ \text{Bern}(1/2) & \text{if } a_1 > \frac{1}{2} \end{cases}, \quad (191)$$

where $\delta \in (0, \frac{1}{2})$ is a small value to be chosen later. The noise is reversed in the second bandit instance \mathcal{G}_2 , which is

$$\epsilon_1 = \begin{cases} \text{Bern}(1/2) & \text{if } a_1 \leq \frac{1}{2} \\ \text{Bern}(1/2 + \delta) & \text{if } a_1 > \frac{1}{2} \end{cases}. \quad (192)$$

We assume the rest of the random variables are noiseless, i.e. $\epsilon_i = 0$ for $i \neq 1$. Otherwise, the system will be more difficult to learn and cause larger regret. Since the weights are fixed and non-negative and interventions at other nodes do not affect the reward, the optimal action is the one that maximizes the average value of the node 1. In contrast, other interventions can be chosen freely. Hence, the optimal intervention set for \mathcal{G}_1 and \mathcal{G}_2 are

$$\mathcal{A}_{\mathcal{G}_1}^* = \left\{ \mathbf{a} \in \mathcal{A}, a_i \leq \frac{1}{2} \right\} \quad \text{and} \quad \mathcal{A}_{\mathcal{G}_2}^* = \left\{ \mathbf{a} \in \mathcal{A}, a_i > \frac{1}{2} \right\}. \quad (193)$$

We choose $\mathbf{a}_{\mathcal{G}_1}^* \in \mathcal{A}_{\mathcal{G}_1}^*$ and $\mathbf{a}_{\mathcal{G}_2}^* \in \mathcal{A}_{\mathcal{G}_2}^*$ to represent the optimal interventions. Since all the nodes $i \neq 1$ are noiseless, the variables in the same layer will have the same values. Then, we have the following recursive relationships in the system

$$X_i = \beta \left(\sum_{j \in \text{pa}(i)} X_j \right)^p = \beta d^p X_j, \quad \forall j \in \text{pa}(i). \quad (194)$$

Based on the recursive relationship, when we dissect the structure to the first layer from node $i \in [N]$, we get the distribution of X_i as follows:

$$X_i = \begin{cases} 0 & \text{if } \epsilon_1 = 0 \\ \beta \frac{p^{L_i-1}}{p-1} d^{\frac{p(L_i-1-1)}{p-1}} & \text{if } \epsilon_1 = 1 \end{cases}. \quad (195)$$

Accordingly, the Lipschitz constant, which is the maximum gradient magnitude, becomes

$$K_i = \sup \left\| \nabla_{X_{\text{pa}(i)}} \beta \left(\sum_{j \in \text{pa}(i)} X_j \right)^p \right\| \quad (196)$$

$$= \sqrt{d} \beta p \sup \left\| X_j^{p-1} \right\| \quad (197)$$

$$= \sqrt{d} \beta p \beta^{p^{L_i-1}-1} d^{p^{L_i-1}-p} \quad (198)$$

$$= p \beta^{p^L} d^{p^{L_i-1}-p+\frac{1}{2}}, \quad (199)$$

where $j \in \text{pa}(i)$ is a fixed index, \sqrt{d} term (197) is due to the maximum gradient magnitude, which is maximized when

$$X_j = \beta \frac{p^{L_j-1}}{p-1} d^{\frac{p(L_j-1-1)}{p-1}}, \quad \forall j \in \text{pa}(i), \quad (200)$$

and we plug in the characterization of X_i specified in (195) in (198). Since the interventions on nodes other than node 1 do not influence the reward, we have $\mu_{\mathbf{a}_{\mathcal{G}_2}^*} = \mu_{\mathbf{a}}$ for all intervention \mathbf{a} with $\mathcal{A}_{\mathcal{G}_2}^*$. Therefore, for \mathcal{G}_1 , the expected value of instantaneous regret is $\mu_{\mathbf{a}_{\mathcal{G}_1}^*} - \mu_{\mathbf{a}_{\mathcal{G}_2}^*}$ if $a_1 > \frac{1}{2}$. Furthermore, since the noises $\epsilon_i = 0$ for $i \geq 2$, based on (195), we have

$$\mu_{\mathbf{a}_{\mathcal{G}_1}^*} - \mu_{\mathbf{a}_{\mathcal{G}_2}^*} = \beta \frac{p^L - 1}{p - 1} d^{\frac{p(p^L - 1 - 1)}{p - 1}} \delta. \quad (201)$$

By setting β as

$$\beta = (p^L d^{L(p - \frac{1}{2})})^{-p^L - p}. \quad (202)$$

we have the following relationship

$$\mu_{\mathbf{a}_{\mathcal{G}_1}^*} - \mu_{\mathbf{a}_{\mathcal{G}_2}^*} = K \delta, \quad (203)$$

in which $K = \prod_{\ell=2}^L K^{(\ell)}$. Next, consider a fixed bandit policy π that generates the filtration \mathcal{H}_t over time. The decision of π at time t is \mathcal{H}_{t-1} -measurable. Accordingly, define the $\mathbb{P}_{\mathcal{G}_1}$ and $\mathbb{P}_{\mathcal{G}_2}$ as the probability measures induced by \mathcal{H}_T by T rounds of interaction of π with \mathcal{G}_1 and \mathcal{G}_2 , respectively. We show that any algorithm π cannot simultaneously achieve small regret in both causal bandit instances. Since the difference between the two bandit instances is caused by their differences in the first node, we are interested in the event in which $a_1 > \frac{1}{2}$ is chosen at least $\frac{T}{2}$ times after T rounds of interactions. We denote this event by

$$\mathcal{E}_\pi \triangleq \left\{ \sum_{i=1}^T \mathbb{1}\left\{a_1(t) > \frac{1}{2}\right\} \geq \frac{T}{2} \right\}. \quad (204)$$

Consequently, we can lower bound the regret associated with \mathcal{G}_1 by conditioning on the event \mathcal{E}_π being true, i.e.

$$\mathbb{E}_{\mathcal{G}_1}[R(T)] = \mathbb{E}_{\mathcal{G}_1} \left[\sum_{t=1}^T (\mu^* - \mu_{\mathbf{a}(t)}) \right] \quad (205)$$

$$\geq \mathbb{E}_{\mathcal{G}_1} \left[\sum_{t \in [T]: a_1(t) > \frac{1}{2}} (\mu_{\mathbf{a}_{\mathcal{G}_1}^*} - \mu_{\mathbf{a}(t)}) \right] \quad (206)$$

$$\geq \mathbb{E}_{\mathcal{G}_1} \left[\sum_{t \in [T]: a_1(t) > \frac{1}{2}} \mathbb{1}\{\mathcal{E}_\pi\} (\mu_{\mathbf{a}_{\mathcal{G}_1}^*} - \mu_{\mathbf{a}_{\mathcal{G}_2}^*}) \right] + \mathbb{E}_{\mathcal{G}_1} \left[\sum_{t \in [T]: a_1(t) > \frac{1}{2}} \mathbb{1}\{\mathcal{E}_\pi^c\} (\mu_{\mathbf{a}_{\mathcal{G}_1}^*} - \mu_{\mathbf{a}_{\mathcal{G}_2}^*}) \right] \quad (207)$$

$$= \mathbb{P}_{\mathcal{G}_1}(\mathcal{E}_\pi) \frac{T}{2} (\mu_{\mathbf{a}_{\mathcal{G}_1}^*} - \mu_{\mathbf{a}_{\mathcal{G}_2}^*}) \quad (208)$$

$$= \mathbb{P}_{\mathcal{G}_1}(\mathcal{E}_\pi) \frac{T}{2} K \delta, \quad (209)$$

where (206) is a lower bound on the regret by considering the regret incurred when $a_1(t) > \frac{1}{2}$; (207) holds due to $\mu_{\mathbf{a}_{\mathcal{G}_2}^*} = \mu_{\mathbf{a}}$ for all intervention $\mathbf{a} \in \mathcal{A}_{\mathcal{G}_2}^*$ and we split the sum based on whether \mathcal{E}_π holds; and (208) uses the definition of \mathcal{E}_π , lower bounding for the occurrence time of regret by $\frac{T}{2}$ for \mathcal{E}_π , and by 0 for the \mathcal{E}_π^c . The last equality in (209) holds due to (203). Similarly, for bandit instance \mathcal{G}_2 , we have $\mu_{\mathbf{a}_{\mathcal{G}_1}^*} = \mu_{\mathbf{a}}$ for all intervention $\mathbf{a} \in \mathcal{A}_{\mathcal{G}_1}^*$. Therefore, for \mathcal{G}_2 , the expected value of instantaneous regret is $\mu_{\mathbf{a}_{\mathcal{G}_2}^*} - \mu_{\mathbf{a}_{\mathcal{G}_1}^*}$ if $a_i(t) \leq \frac{1}{2}$. Hence, we can similarly show that

$$\mathbb{E}_{\mathcal{G}_2}[R(T)] \geq \mathbb{P}_{\mathcal{G}_2}(\mathcal{E}_\pi^c) \frac{T}{2} K \delta. \quad (210)$$

Combining (209) and (210), we have

$$\mathbb{E}_{\mathcal{G}_1}[R(T)] + \mathbb{E}_{\mathcal{G}_2}[R(T)] \geq (\mathbb{P}_{\mathcal{G}_1}(\mathcal{E}_\pi) + \mathbb{P}_{\mathcal{G}_2}(\mathcal{E}_\pi^c)) K \delta. \quad (211)$$

Next, we use the Lemma C.1 and find a lower bound on the term $\mathbb{P}_{\mathcal{G}_1}(\mathcal{E}_\pi) + \mathbb{P}_{\mathcal{G}_2}(\mathcal{E}_\pi^c)$, resulting in

$$\mathbb{E}_{\mathcal{G}_1}[R(T)] + \mathbb{E}_{\mathcal{G}_2}[R(T)] \geq \frac{T}{4} K \delta \exp(-\text{D}_{\text{KL}}(\mathbb{P}_{\mathcal{G}_1} \parallel \mathbb{P}_{\mathcal{G}_2})). \quad (212)$$

It remains to compute the KL-divergence between the two bandit instances \mathcal{G}_1 and \mathcal{G}_2 , which is formalized by the following lemma.

Lemma C.3. For the causal bandit instance \mathcal{G}_1 and \mathcal{G}_2 , the KL-divergence between $\mathbb{P}_{\mathcal{G}_1}$ and $\mathbb{P}_{\mathcal{G}_2}$ is upper bounded by

$$D_{\text{KL}}(\mathbb{P}_{\mathcal{G}_1} \parallel \mathbb{P}_{\mathcal{G}_2}) \leq T \log \left(\frac{1}{(1+2\delta)(1-2\delta)} \right). \quad (213)$$

Proof. The proof follows the same steps as that of Lemma C.2 with the difference that we split the summation in (180) by conditioning on the events $\{a_i(t) \leq \frac{1}{2}\}$ and $\{a_i(t) > \frac{1}{2}\}$. \square

If we choose $\delta = \frac{1}{\sqrt{T}}$ to balance the terms in the lower bound, we obtain

$$\max\{\mathbb{E}_{\mathcal{G}_1}[R(T)], \mathbb{E}_{\mathcal{G}_2}[R(T)]\} \geq \frac{1}{2} (\mathbb{E}_{\mathcal{G}_1}[R(T)] + \mathbb{E}_{\mathcal{G}_2}[R(T)]) \quad (214)$$

$$\stackrel{(212)}{\geq} \frac{T}{8} K \delta \exp(-D_{\text{KL}}(\mathbb{P}_{\mathcal{G}_1} \parallel \mathbb{P}_{\mathcal{G}_2})) \quad (215)$$

$$\stackrel{(213)}{\geq} \frac{T}{8} K \delta [(1+2\delta)(1-2\delta)]^T \quad (216)$$

$$= \frac{1}{8} K \sqrt{T} \times \left(1 - \frac{4}{T}\right)^T. \quad (217)$$

for $T \geq 5$, the term $(1 - 4/T)^T$ is an increasing function of T . Hence, for $T \geq 5$, we have the lower bound $(1 - 4/T)^T \geq 0.2^5$. By setting $c = \frac{1}{8} \times 0.2^5$, we have

$$\max\{\mathbb{E}_{\mathcal{G}_1}[R(T)], \mathbb{E}_{\mathcal{G}_2}[R(T)]\} \geq cK\sqrt{T}. \quad (218)$$

Hence, the policy π incurs a regret $cK\sqrt{T}$ in at least one of the two bandit instances.

C.3 Proof of Theorem 5.10 (Neural Network SCMs)

Similarly to the linear and polynomial settings, we consider two neural network SCM causal bandit instances \mathcal{G}_1 and \mathcal{G}_2 that share the same hierarchical graph specified earlier (Figure 2) with degree d and maximum causal path length L . Subsequently, our minimax lower bound of interest is lower bounded by

$$\inf_{\pi \in \Pi} \sup_{\mathcal{F} \in \bar{\mathcal{F}}} \sup_{\mathcal{G} \in \bar{\mathcal{G}}} \mathbb{E}[R(T)] \geq \inf_{\pi \in \Pi} \sup_{\mathcal{G} \in \{\mathcal{G}_1, \mathcal{G}_2\}} \mathbb{E}[R(T)]. \quad (219)$$

Within the class of neural network models, we use the parametric Rectified Linear Unit (ReLU) function as the activation function, which is specified by

$$\sigma(x) = \begin{cases} \beta x & \text{if } x \geq 0 \\ \alpha x & \text{if } x < 0 \end{cases}, \quad (220)$$

where the constant $\beta > \alpha > 0$ control the slope of the function. We assume the two bandit instances share same weights $\{\Theta_{i,1} : i \in [N]\}$ and $\{\Theta_{i,2} : i \in [N]\}$, which are specified as follows.

$$[\Theta_{i,1}]_{j,k} = \begin{cases} 0, & k = 1 \\ \frac{\sqrt{K^{(L_i)}}}{\beta\sqrt{s}}, & k \in [d_i + 1], k \geq 2 \end{cases}, \quad \forall j \in [s]. \quad (221)$$

and

$$[\Theta_{i,2}]_j = \frac{\sqrt{K^{(L_i)}}}{\beta\sqrt{s}}, \quad \forall j \in [s]. \quad (222)$$

The intervention sets for two bandit instances \mathcal{G}_1 and \mathcal{G}_2 are set to $\mathcal{A}_i = [0, 1]$. The difference between the two bandit instances is the noise distribution in the system. For the noise for the node 1, in first bandit instance \mathcal{G}_1 , it is defined as

$$\epsilon_1 = \begin{cases} \text{Bern}(1/2 + \delta) & \text{if } a_1 \leq \frac{1}{2} \\ \text{Bern}(1/2) & \text{if } a_1 > \frac{1}{2} \end{cases}, \quad (223)$$

where $\delta \in (0, \frac{1}{2})$ is a small value to be chosen later. The noise is reversed in the second bandit instance \mathcal{G}_2 , which is

$$\epsilon_1 = \begin{cases} \text{Bern}(1/2) & \text{if } a_1 \leq \frac{1}{2} \\ \text{Bern}(1/2 + \delta) & \text{if } a_1 > \frac{1}{2} \end{cases}. \quad (224)$$

We assume that the noise terms associated with the rest of the random variables follow the Bernoulli model.

$$\epsilon_i = \begin{cases} -1 & \text{with probability } \frac{1}{2} \\ +1 & \text{with probability } \frac{1}{2} \end{cases} \quad \text{for } i \neq 1. \quad (225)$$

Since the weights are fixed and non-negative, and the noise terms are zero-mean when $i \neq 1$, the optimal action is the one that maximizes the average value of the node 1, and the interventions at other nodes do not affect the reward. As a result, the optimal action at node 1 is the action that maximizes the mean of noise, while other interventions can be chosen freely. Hence, the optimal intervention set for \mathcal{G}_1 and \mathcal{G}_2 are

$$\mathcal{A}_{\mathcal{G}_1}^* = \left\{ \mathbf{a} \in \mathcal{A}, a_1 \leq \frac{1}{2} \right\} \quad \text{and} \quad \mathcal{A}_{\mathcal{G}_2}^* = \left\{ \mathbf{a} \in \mathcal{A}, a_1 > \frac{1}{2} \right\}. \quad (226)$$

We pick $\mathbf{a}_{\mathcal{G}_1}^* \in \mathcal{A}_{\mathcal{G}_1}^*$ and $\mathbf{a}_{\mathcal{G}_2}^* \in \mathcal{A}_{\mathcal{G}_2}^*$ as the represent of the optimal intervention. Next, consider a fixed bandit policy π that generates the filtration \mathcal{H}_t over time. The decision of π at time t is \mathcal{H}_{t-1} -measurable. Accordingly, define the $\mathbb{P}_{\mathcal{G}_1}$ and $\mathbb{P}_{\mathcal{G}_2}$ as the probability measures induced by \mathcal{H}_T by T rounds of interaction of π with \mathcal{G}_1 and \mathcal{G}_2 , respectively. We show that any algorithm π cannot get small regret in both causal bandit instances simultaneously. Since the difference between two bandit instances is on the first node, we are interested in the event in which $a_1 > \frac{1}{2}$ is chosen at least $\frac{T}{2}$ times after T rounds of interactions. We denote this event by

$$\mathcal{E}_\pi \triangleq \left\{ \sum_{i=1}^T \mathbb{1}\left\{a_1(t) > \frac{1}{2}\right\} \geq \frac{T}{2} \right\}. \quad (227)$$

Since interventions on nodes other than node 1 do not influence the reward, we have $\mu_{\mathbf{a}_{\mathcal{G}_2}^*} = \mu_{\mathbf{a}}$ for all intervention $\mathbf{a} \in \mathcal{A}_{\mathcal{G}_2}^*$. Therefore, for \mathcal{G}_1 , the expected value of instantaneous regret is $\mu_{\mathbf{a}_{\mathcal{G}_1}^*} - \mu_{\mathbf{a}_{\mathcal{G}_2}^*}$ if $a_1 > \frac{1}{2}$. Besides, since ϵ_i has a zero mean for $i \geq 2$, only the paths that start at node 1 and end at the reward node N contribute to the expected regret. Thus, for the gap between $\mu_{\mathbf{a}_{\mathcal{G}_1}^*}$ and $\mu_{\mathbf{a}_{\mathcal{G}_2}^*}$, each path contribute $\prod_{\ell=1}^L K^{(\ell)} d^{-\frac{L}{2}} s^{-L} \delta$ and we have $d^{L-1} s^L$ such paths. Therefore, we have

$$\mu_{\mathbf{a}_{\mathcal{G}_1}^*} - \mu_{\mathbf{a}_{\mathcal{G}_2}^*} = d^{\frac{L}{2}-1} \delta \prod_{\ell=1}^L K^{(\ell)}. \quad (228)$$

Consequently, we can lower bound the regret associated with \mathcal{G}_1 by conditioning on the event \mathcal{E}_π being true, i.e.,

$$\mathbb{E}_{\mathcal{G}_1}[R(T)] = \mathbb{E}_{\mathcal{G}_1} \left[\sum_{t=1}^T (\mu^* - \mu_{\mathbf{a}(t)}) \right] \quad (229)$$

$$\geq \mathbb{E}_{\mathcal{G}_1} \left[\sum_{t \in [T]: a_1(t) > \frac{1}{2}} (\mu_{\mathbf{a}_{\mathcal{G}_1}^*} - \mu_{\mathbf{a}(t)}) \right] \quad (230)$$

$$\geq \mathbb{E}_{\mathcal{G}_1} \left[\sum_{t \in [T]: a_1(t) > \frac{1}{2}} \mathbb{1}\{\mathcal{E}_\pi\} (\mu_{\mathbf{a}_{\mathcal{G}_1}^*} - \mu_{\mathbf{a}_{\mathcal{G}_2}^*}) \right] + \mathbb{E}_{\mathcal{G}_1} \left[\sum_{t \in [T]: a_1(t) > \frac{1}{2}} \mathbb{1}\{\mathcal{E}_\pi^c\} (\mu_{\mathbf{a}_{\mathcal{G}_1}^*} - \mu_{\mathbf{a}_{\mathcal{G}_2}^*}) \right] \quad (231)$$

$$= \mathbb{P}_{\mathcal{G}_1}(\mathcal{E}_\pi) \frac{T}{2} (\mu_{\mathbf{a}_{\mathcal{G}_1}^*} - \mu_{\mathbf{a}_{\mathcal{G}_2}^*}) \quad (232)$$

$$= \mathbb{P}_{\mathcal{G}_1}(\mathcal{E}_\pi) \frac{T}{2} d^{\frac{L}{2}-1} \delta \prod_{\ell=1}^L K^{(\ell)}, \quad (233)$$

where (230) lower bound the regret by considering the regret incurred when $a_1(t) > \frac{1}{2}$; (231) holds due to $\mu_{\mathbf{a}_{\mathcal{G}_2}^*} = \mu_{\mathbf{a}}$ for all intervention $\mathbf{a} \in \mathcal{A}_{\mathcal{G}_2}^*$ and we split the sum based on whether \mathcal{E}_π holds; and (232) uses the definition of \mathcal{E}_π , lower bounding for the occurrence time of regret by $\frac{T}{2}$ for \mathcal{E}_π , and by 0 for the \mathcal{E}_π^c . The last equality in (233) holds due to (228). Similarly, for bandit instance \mathcal{G}_2 , we have $\mu_{\mathbf{a}_{\mathcal{G}_1}^*} = \mu_{\mathbf{a}}$ for all intervention $\mathbf{a} \in \mathcal{A}_{\mathcal{G}_1}^*$. Therefore, for \mathcal{G}_1 , the expected value of instantaneous regret is $\mu_{\mathbf{a}_{\mathcal{G}_2}^*} - \mu_{\mathbf{a}_{\mathcal{G}_1}^*}$ if $a_i(t) \leq \frac{1}{2}$. Hence, we can similarly show that

$$\mathbb{E}_{\mathcal{G}_2}[R(T)] \geq \mathbb{P}_{\mathcal{G}_2}(\mathcal{E}_\pi^c) \frac{T}{2} d^{\frac{L}{2}-1} \delta \prod_{\ell=1}^L K^{(\ell)}. \quad (234)$$

Combining (233) and (234), we have

$$\mathbb{E}_{\mathcal{G}_1}[R(T)] + \mathbb{E}_{\mathcal{G}_2}[R(T)] \geq (\mathbb{P}_{\mathcal{G}_1}(\mathcal{E}) + \mathbb{P}_{\mathcal{G}_2}(\mathcal{E}^c)) \frac{T}{2} d^{\frac{L}{2}-1} \delta \prod_{\ell=1}^L K^{(\ell)}. \quad (235)$$

Next, we use the Lemma C.1 and find a lower bound on the term $\mathbb{P}_{\mathcal{G}_1}(\mathcal{E}_\pi) + \mathbb{P}_{\mathcal{G}_2}(\mathcal{E}_\pi^c)$, resulting in

$$\mathbb{E}_{\mathcal{G}_1}[R(T)] + \mathbb{E}_{\mathcal{G}_2}[R(T)] \geq \frac{T}{4} d^{\frac{L}{2}-1} \delta \exp(-D_{\text{KL}}(\mathbb{P}_{\mathcal{G}_1} \parallel \mathbb{P}_{\mathcal{G}_2})) \prod_{\ell=1}^L K^{(\ell)}. \quad (236)$$

We provide the following lemma to characterize the KL-divergence between bandit instances \mathcal{G}_1 and \mathcal{G}_2 .

Lemma C.4. *For causal bandit instance \mathcal{G}_1 and \mathcal{G}_2 , the KL-divergence between $\mathbb{P}_{\mathcal{G}_1}$ and $\mathbb{P}_{\mathcal{G}_2}$ is upper bounded by*

$$D_{\text{KL}}(\mathbb{P}_{\mathcal{G}_1} \parallel \mathbb{P}_{\mathcal{G}_2}) \leq T \log \left(\frac{1}{(1+2\delta)(1-2\delta)} \right). \quad (237)$$

Proof. The proof follows the same steps as that of Lemma C.2 with the difference that we split the summation in (180) by event $\{a_i(t) \leq \frac{1}{2}\}$ and $\{a_i(t) > \frac{1}{2}\}$. \square

If we choose $\delta = \frac{1}{\sqrt{T}}$ to balance the terms in the lower bound, we obtain

$$\max\{\mathbb{E}_{\mathcal{G}_1}[R(T)], \mathbb{E}_{\mathcal{G}_2}[R(T)]\} \geq \frac{1}{2} (\mathbb{E}_{\mathcal{G}_1}[R(T)] + \mathbb{E}_{\mathcal{G}_2}[R(T)]) \quad (238)$$

$$\stackrel{(236)}{\geq} \frac{T}{8} d^{\frac{L}{2}-1} \delta \exp(-D_{\text{KL}}(\mathbb{P}_{\mathcal{G}_1} \parallel \mathbb{P}_{\mathcal{G}_2})) \prod_{\ell=1}^L K^{(\ell)} \quad (239)$$

$$\stackrel{(237)}{\geq} \frac{T}{8} d^{\frac{L}{2}-1} \delta [(1+2\delta)(1-2\delta)]^T \prod_{\ell=1}^L K^{(\ell)} \quad (240)$$

$$= \frac{1}{8} d^{\frac{L}{2}-1} \sqrt{T} \times \left(1 - \frac{4}{T}\right)^T \prod_{\ell=1}^L K^{(\ell)}. \quad (241)$$

Similarly to the previous two models, by setting $c = \frac{1}{8} \times 0.2^5$, we have

$$\max\{\mathbb{E}_{\mathcal{G}_1}[R(T)], \mathbb{E}_{\mathcal{G}_2}[R(T)]\} \geq cKd^{\frac{L}{2}-1}\sqrt{T}. \quad (242)$$

Hence, the policy π incurs a regret $cKd^{\frac{L}{2}-1}\sqrt{T}$ in at least one of the two bandit instances.

D Further Comparisons with (Sussex et al., 2023)

In this section, we analytically compare the upper bounds presented in this paper, and those of the model-based causal Bayesian optimization (MCBO) approach presented in (Sussex et al., 2023). We present two comparisons.

- **Scaling with N .** The important distinction pertains to the scaling rate of the regret terms with respect to the graph size. Specifically, we discuss that the achievable regret of MCBO scales *linearly* with graph size N , while as we discussed earlier, the regret of GCB algorithms scales with N according to $\sqrt{\frac{\log N}{\log T}}$, which for finite T is substantially slower than the linear growth rate, and diminishes as T grows.
- **Scaling with T .** We also provide a class of functions \mathcal{F} for which the GCB algorithm is guaranteed to have sublinear regret in terms of T , while the MCBO algorithm is not.

Since the regret analyses of the two methods are carried out via entirely distinct methods, we need to find a way to unify and interpret the results. More specifically, our results are presented in terms of the eluder dimensions and the covering numbers of function classes. On the other hand, the results of MCBO for Gaussian processes are presented in terms of *information gain* associated with the corresponding reproducing kernel Hilbert spaces. To facilitate comparisons, we will provide a unifying way of presenting both results in terms of information gain measures, which is the basis for MCBO's analysis. We compare these results for a specific model, described next, and specialize the regret bounds of the GCB and MCBO algorithms for this model.

RKHS with polynomial spectral decay. For a given vector $\theta \in \mathbb{R}^{d_i}$, scalar $a_i \in \mathcal{A}_i$, and RKHS \mathcal{H}_i with respect to a known kernel function k_i , set

$$f_i(X_{\text{pa}(i)}; a_i) = \langle \theta, k_i(\cdot, [X_{\text{pa}(i)}, a_i]) \rangle_{\mathcal{H}_i}, \quad (243)$$

and accordingly, set

$$\mathcal{F}_i = \{f_i \mid \|\theta\| \leq 1, k_i \text{ satisfies Assumption D.1}\}. \quad (244)$$

Assumption D.1. (Polynomial decay). Denote the space of inputs to f_i by $\mathcal{Z}_i \in \mathbb{R}^{d_i} \times \mathcal{A}_i$. Assume that the kernel function k_i admits the following orthonormal decomposition such that for any $z, z' \in \mathcal{Z}_i$:

$$k_i(z, z') = \sum_{j=1}^{\infty} \lambda_j \phi_j(z) \phi_j(z'), \quad (245)$$

where $\{\lambda_j : j \in \mathbb{N}\}$ are countable and positive eigenvalues with corresponding eigenfunctions $\{\phi_j : j \in \mathbb{N}\}$ that satisfy $\|\phi_j(z)\|_{\infty} \leq C_{\phi}$ for $z \in \mathcal{Z}_i$ and the eigenvalues λ_j have a β -polynomial decay, i.e., $\lambda_j \leq C_p j^{-\beta}$, for some $\beta > 2 + \frac{1}{d}$ and constant C_p . We show that for the RKHS class specified in (244), MCBO is not guaranteed to have a sublinear regret bound in T , whereas GCB maintains a sub-linear regret.

D.1 Information Gain

We provide the definitions of information gain and the associated maximum and critical information gains customized for Gaussian processes (Srinivas et al., 2010).

Definition D.2. (Information Gain). Corresponding to node $i \in [N]$, consider the RKHS \mathcal{H}_i with input space $\mathcal{Z}_i \triangleq \mathbb{R}^{d_i} \times \mathcal{A}_i$ and an associated kernel k_i . For any $\lambda > 0$, the information gain upon observing $\{z_\ell \in \mathcal{Z}_i : \ell \in [T]\}$ is defined as

$$\gamma_i(\lambda; \{z_1, \dots, z_T\}) \triangleq \log \det \left(I + \frac{1}{\lambda} \sum_{i=1}^T z_i z_i^\top \right), \quad (246)$$

where I is the identity matrix.

Based on the notion of information gain, the maximum information gain is defined as follows, which captures the maximum volume of the ellipsoid generated by T points in \mathcal{Z}_i .

Definition D.3. (Maximum Information Gain). For any $\lambda > 0$ and $T \in \mathbb{N}$, the maximum information gain upon T observations is defined as

$$\gamma_{i,T}(\lambda) \triangleq \max_{\{z_\ell: \ell \in [T]\} \in \mathcal{Z}_i^T} \gamma_i(\lambda; \{z_1, \dots, z_T\}). \quad (247)$$

Accordingly, we define

$$\gamma_T \triangleq \max_{i \in [N]} \gamma_{i,T}(\lambda). \quad (248)$$

Finally, the following critical information gain is the minimum T for which maximum information gain cannot grow linearly.

Definition D.4. (Critical Information Gain). For a fixed constant $c > 0$, the critical information gain is defined as

$$\tilde{\gamma}_i(\lambda, c) \triangleq \min \{k : \gamma_{i,k}(\lambda) \leq ck\}. \quad (249)$$

D.2 Achievable Regret of the MCBO Algorithm

Lemma D.5 (informal). (*Sussex et al., 2023, Theorem 1*) Consider SCMs in which functions f_i are Lipschitz-continuous and bounded in \mathcal{H}_i , noise terms are bounded, the Gaussian process space that satisfies the continuous kernel and calibrated assumption. Then, for all $T \geq 1$, with probability at least $1 - \delta$, the cumulative regret of MCBO is bounded by

$$\mathbb{E}[R(T)] = \mathcal{O} \left(K \beta_T^L d^L N \sqrt{T \gamma_T} \right), \quad (250)$$

where K is associated with Lipschitz constants of the function and kernel space, β_T is the confidence radius, and σ^2 is the noise variance.

While the assumptions adopted in (Sussex et al., 2023) differ from ours and are not directly comparable, this upper bound leads us to draw two primary conclusions.

- First, the regret bound is linearly proportional to the number of nodes N .
- Secondly, the radius of the confidence set associated with estimating the parameters of the Gaussian process (mean and variance) is

$$\beta_T = \tilde{\mathcal{O}} \left(\frac{\sqrt{\gamma_T}}{d} \right), \quad (251)$$

which hinges on the maximum information gain. By comparing (250) and (251), we have the following growth for the regret:

$$\mathbb{E}[R_T] = \mathcal{O} \left(\gamma_T^{\frac{L+1}{2}} \sqrt{T} \right). \quad (252)$$

This regret bound becomes linear in T when γ_T has super-logarithmic growth with T and L is large enough. Specifically, we show this happens for the class specified in (244), which we establish next.

Lemma D.6. (Yang et al., 2020, Lemmas D.5) For the class of function \mathcal{F}_i defined in (244), the maximum information gain is upper bounded as follows.

$$\gamma_{i,T}(\lambda) = \tilde{\mathcal{O}}\left(\lambda^{-\frac{1}{\beta}} T^{\frac{d+1}{\beta+d}}\right). \quad (253)$$

We conclude that the regret of MCBO scales as

$$\mathbb{E}[R(T)] = \tilde{\mathcal{O}}\left(T^{\frac{(L+1)(d+1)}{2(\beta+d)} + \frac{1}{2}}\right). \quad (254)$$

Thus, when $L \geq \frac{\beta+d}{d+1} - 1$, the MCBO algorithm cannot guarantee sub-linear regret.

D.3 Achievable Regret of the GCB Algorithm

Next, we analyze the regret of the GCB algorithms for the same model analyzed in the previous section, i.e., the class of functions specified in (244). We leverage the following result on the eluder dimension and the critical information gain.

Lemma D.7. (Huang et al., 2021, Theorem 3.1) Consider the function class \mathcal{F}_i defined in (244). By setting $\lambda = \frac{\epsilon^2}{4}$, we have the following connection between the eluder dimension and the critical information gain:

$$\dim(\mathcal{F}_i, \epsilon) < \tilde{\gamma}_i\left(\lambda, \log \frac{3}{2}\right). \quad (255)$$

Furthermore, from Lemma D.6 we know that $\gamma_{i,T} = \tilde{\mathcal{O}}\left(T^{\frac{d+1}{\beta+d}}\right)$. Hence, the critical information gain satisfies (Huang et al., 2021):

$$\tilde{\gamma}_i(\lambda, c) = \tilde{\mathcal{O}}\left(\lambda^{-\frac{\beta+d}{\beta(\beta-1)}}\right). \quad (256)$$

In our analysis and design of the algorithm in(24), we have set $\epsilon \geq \frac{1}{T}$, based on which we obtain

$$\dim(\mathcal{F}_i, \frac{1}{T}) < \tilde{\gamma}_i\left(\frac{1}{4T^2}, \log \frac{3}{2}\right) = \tilde{\mathcal{O}}\left(T^{\frac{2(\beta+d)}{\beta(\beta-1)}}\right). \quad (257)$$

Additionally, we have the following upper bound for the covering number.

Lemma D.8. (Yang et al., 2020, Lemmas D.2) Consider the function class \mathcal{F}_i defined in (244). The covering number of this class satisfies:

$$\log \text{cn}_\alpha(\mathcal{F}_i) = \tilde{\mathcal{O}}\left(\epsilon^{-\frac{2}{\beta-1}}\right). \quad (258)$$

Thus, by setting $\alpha = \frac{1}{T}$, we obtain

$$\log \text{cn}_{\frac{1}{T}}(\mathcal{F}_i) = \tilde{\mathcal{O}}\left(T^{\frac{2}{\beta-1}}\right). \quad (259)$$

Therefore, based on (257) and (259), we find the following bound for the GCB algorithms.

Corollary D.9. Consider the function class \mathcal{F}_i defined in (244). The regret upper bounds in Theorem 5.3 becomes

$$\mathbb{E}[R(T)] = \tilde{\mathcal{O}}\left(T^{\frac{2\beta+d}{\beta(\beta-1)} + \frac{1}{2}}\right). \quad (260)$$

We observe that regret is independent of L . This provides the freedom to choose the values of other parameters and L such that we maintain a sublinear growth in T for the regret in (260). On the other hand, having a linear growth for the regret in (254) is not guaranteed. Specifically, this can be readily verified by setting $\beta = d = 10$, for which (260) is sublinear whereas (254) becomes linear for $L \geq 2$.

E Experiments

In this section, we assess the regret performance (average cumulative regret) of the GCB algorithm and its scaling behavior with the time horizon, and the graph parameters d (graph maximum in-degree) and L (maximum causal path length). We assess the performance of the the GCB-TS algorithm on polynomial and neural network SCMs.

Causal graph. We consider a hierarchical graph as illustrated in Figure 3. This graph consists of $L + 1$ layers with the first L layers having d nodes. The nodes between two adjacent layers are fully connected. Finally, the last layer consists of one node (reward node) that is fully connected to nodes in layer L .




Figure 3: Hierarchical graph with degree $d = 3$ and causal length $L = 2$.

Parameter setting. In all experiments, we assume that the noise terms $\{\epsilon_i : i \in [N]\}$ are drawn from a standard Gaussian distribution $\mathcal{N}(0, 1)$. To approximate the Bayesian, we assume that the parameters θ_i for f_i have a Gaussian prior distribution and sample them from the Gaussian prior with small variance at each trail.

- **Polynomial SCMs.** For the causal bandit with polynomial SCMs, we consider the case when $p = 2$, and set the mean of weights’ Gaussian prior to $\frac{1}{d+1}$ to control the scaling behavior of the system. We set the intervention space to $\mathcal{A}_i = \{0, 1\}$.
- **Neural network SCMs.** For the neural network SCMs, we consider the function space of a two-layer neural network with a leaky ReLU activation function with a slope of $\alpha = 0.1$ and set the mean of weights’ Gaussian prior to $\frac{1}{2\sqrt{d+1}}$. Furthermore, We set the intervention space as $\mathcal{A}_i = [0, 1]$.

Algorithm settings. For polynomial SCMs, the least square estimator $\tilde{f}_{i,t}$ in (16) can be solved by solving a linear system, and we sample \mathbf{f} from the posterior $\tilde{f}_{i,t} \sim \mathcal{N}(\tilde{f}_{i,t}, \bar{V}_{i,t}^{-1})$, where $\bar{V}_{i,t}^{-1}$ is a Gram matrix related to the linear system. For neural network SCMs, we use stochastic gradient descent (SGD) to approximate the least square estimator $\tilde{f}_{i,t}$ defined in (16), and we sample \mathbf{f} from the posterior $\tilde{f}_{i,t} \sim \mathcal{N}(\tilde{f}_{i,t}, V_{i,t}^{-1})$ with the Gram matrix $V_{i,t}$ defined as



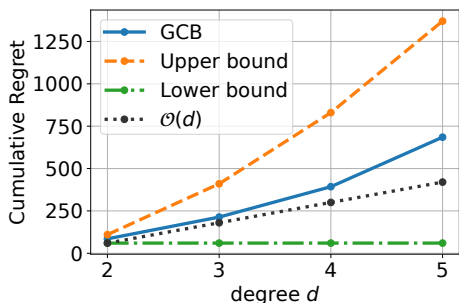

$$V_{i,t} = \sum_{s=1}^t X_{\text{pa}(i)}(s) X_{\text{pa}(i)}^\top(s). \tag{261}$$

We note that a similar sampling strategy is adopted by (Jun et al., 2017; Kveton et al., 2020) for the generalized linear bandit.

Regret performance. In Figure 4 and Figure 5, we present the cumulative regret of GCB-TS algorithm applied to polynomial SCMs and neural network SCMs, respectively. The underlying causal graph is set to be a hierarchical graph with maximum causal path length $L = 2$. Each figure displays the cumulative regret for both $d = 2$ and $d = 3$. It is noticeable that in scenarios with higher complexity, specifically the polynomial SCM of degree $d = 3$, the regret scaling rates fall within the regime $\mathcal{O}(\sqrt{T})$, which is in line with our theoretical results.

Scaling behavior with graph parameters. Figure 6 and Figure 7 illustrate the variations of the cumulative regret versus the graph degree d with the maximum causal path length $L = 2$ under polynomial and neural network SMCs. Furthermore, Figure 8 and Figure 9 display the corresponding variations when $L = 1$. In these figures, we have also included the theoretical lower and upper bounds. In all four plots it is observed that the the scaling trend of the achievable regret are closer to the lower bound than the upper bounds. More specifically, the results in these figures indicate a polynomial behavior in d , which conforms to our theoretical results.

Finally, Figure 10 and Figure 11 elucidate the variations of the cumulative regret versus the maximum causal path length L , given a constant degree of $d = 2$. The emerging patterns across these settings imply an exponential growth of regret with L , conforming with our theoretical findings. Notably, in the polynomial case, the scaling


 Figure 4: Cumulative regret vs T (polynomial)

 Figure 5: Cumulative regret vs T (neural network)

 Figure 6: Cumulative regret vs degree d for $L = 2$ (polynomial)

 Figure 7: Cumulative regret vs degree d for $L = 2$ (neural network)

behavior of L is considerably closer to that of the upper bound, attributed to the complexity of polynomial function spaces with growing value magnitude.

Comparison with baseline algorithms. As the two most relevant existing approaches, we have compared the regret of our algorithm to those of LinSEM-UCB (Varici et al., 2023) and UCB (Agrawal, 1995). In Figure 12, we provide the cumulative regret under a quadratic SCM. Both LinSEM-UCB and UCB algorithms yield linear regrets. LinSEM-UCB encounters failure due to a violation of the linearity of the system, and UCB fails because the reward node’s distribution is heavy-tailed, which is difficult to estimate without information about the causal system. In Figure 13, we compare the results for linear SCMs. LinSEM-UCB, designed and optimized for this case, exhibits the lowest regret, while our algorithm shows a slightly higher regret. This discrepancy arises since our confidence radius is designed for general function classes, and its slightly weaker performance in this special setting is the cost of its universality.




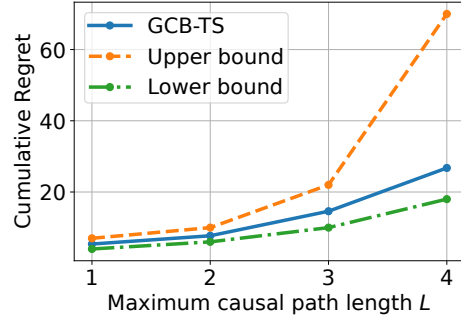
F Extension to do interventions

In this section, we extend the results we provided under soft interventions for do interventions. Consider the same causal system as in Section 2 with the distinction that we apply do interventions. Under applying do intervention $a_i \in \mathcal{A}_i$ on node $i \in [N]$, we have $X_i = a_i$. We also use the convention of setting $a_i = 0$ to generate the observational model on node i , under which the observation relationship is captured by $r_i(X_{\text{pa}(i)}) = f_i(X_{\text{pa}(i)}; 0)$. Hence, the SCMs under do interventions can be restated as

$$X_i = \begin{cases} r_i(X_{\text{pa}(i)}) + \epsilon_i & \text{if } a_i = 0 \text{ (observational model)} \\ a_i & \text{if } a_i \neq 0 \text{ (interventional model)} \end{cases}, \quad \forall i \in [N]. \quad (262)$$

where the functions $\{r_i : i \in [N]\}$ and are *unknown* and f_i and r_i belong to a specified class \mathcal{F}_i and \mathcal{R}_i , respectively.

GCB Algorithms. Our algorithm can be applied to this setting without modifications. The only impact of this change in the algorithms is in the estimators and confidence sets. Since the effects of the intervention are known to be constant, so we only need to estimate and construct confidence sets for the function $r_i \in \mathcal{R}_i$.


 Figure 8: Cumulative regret vs degree d for $L = 1$ (polynomial)

 Figure 9: Cumulative regret vs degree d for $L = 1$ (neural network)

 Figure 10: Cumulative regret vs L for $d = 2$ (polynomial)

 Figure 11: Cumulative regret vs L for $d = 2$ (neural network)

Regret Analysis. The regret guarantee in Theorem 5.3 and Corollary 5.4 still holds. For analyzing the regret guarantees in special SCMs, we provide the following relationships between the eluder dimensions and covering numbers of \mathcal{F}_i and \mathcal{R}_i .

Theorem F.1. *For a given function space \mathcal{F}_i and the associated function space \mathcal{R}_i defined for do intervention, we have the following relationships:*

$$\dim(\mathcal{F}_i, \alpha_i) = \dim(\mathcal{R}_i, \alpha_i) , \quad \text{and} \quad \text{cn}_{\alpha_i}(\mathcal{F}_i) = \text{cn}_{\alpha_i}(\mathcal{R}_i) . \quad (263)$$

Proof. We first provide the relationship for eluder dimensions, ensured by the covering numbers.

Eluder dimension. We start the proof by showing that the longest ordered set in Definition 3.2 does not include the inputs $Z_i = (X_{\text{pa}(i)}, a_i)$ with $a_i \neq 0$ when $\dim(\mathcal{F}_i, \alpha_i) \geq 2$. This is due to the fact that Z_i is always α_i -dependent on other inputs in \mathcal{Z}_i since for any two functions $f, \tilde{f} \in \mathcal{F}_i$ we have

$$f(X_{\text{pa}(i)}; a_i) - \tilde{f}(X_{\text{pa}(i)}; a_i) = a_i - a_i = 0 < \alpha_i . \quad (264)$$

Hence, it is sufficient to prove the next property. Corresponding to any given ordered set $\{Z_i(m) : m \in [n]\}$, where $Z_i(m) = (X_{\text{pa}(i)}(m), a_i(m))$, we define the ordered set $\{\bar{Z}_i(m) : m \in [n]\}$, where $\bar{Z}_i(m) = X_{\text{pa}(i)}(m)$. We show that the following two statements are equivalent for these two ordered sets.

- **S1.** Each element in $\{Z_i(m) : m \in [n]\}$ is α_i -independent of its predecessors with respect to \mathcal{F}_i .
- **S2.** Each element in $\{\bar{Z}_i(m) : m \in [n]\}$ is α_i -independent of its predecessors with respect to \mathcal{R}_i .

This is obvious due to the fact that for any $f \in \mathcal{F}_i$ and $r \in \mathcal{R}_i$ that satisfy $r(\cdot) = f(\cdot; 0)$ we have

$$r(\bar{Z}_i(m)) = f(Z_i(m)) . \quad (265)$$

Thus, based on the above two claims, we conclude that $\dim(\mathcal{F}_i, \alpha_i) = \dim(\mathcal{R}_i, \alpha_i)$.

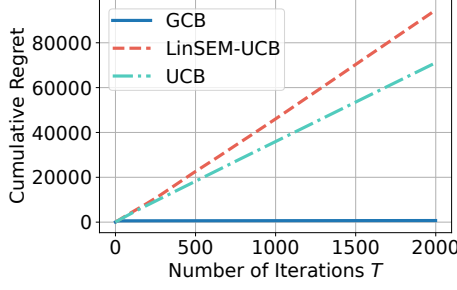


Figure 12: Cumulative regret for different algorithms (polynomial)

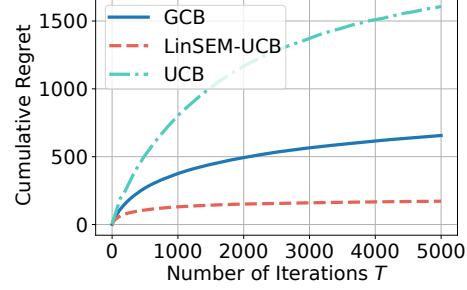


Figure 13: Cumulative regret for different algorithms (linear)

Covering number. We prove this part by showing that for any α_i -cover of \mathcal{F}_i , there exists an α_i -cover of \mathcal{R}_i with an identical cardinality, and vice versa. For this purpose, for any \mathcal{C} that is an α_i -cover of \mathcal{F}_i , we construct $\bar{\mathcal{C}}$ as

$$\bar{\mathcal{C}} = \{r \mid r = f(\cdot; 0) \text{ for } f \in \mathcal{C}\}. \quad (266)$$

Subsequently, for any $r \in \mathcal{R}_i$, we can construct $f \in \mathcal{F}_i$ as

$$f(\cdot; a_i) = \begin{cases} r(\cdot) + \epsilon_i & a_i = 0 \\ a_i & a_i \neq 0 \end{cases}. \quad (267)$$

Since \mathcal{C} is an α_i -cover of \mathcal{F}_i , we can find a function $\tilde{f} \in \mathcal{C}$ such that $\|f - \tilde{f}\|_\infty \leq \alpha_i$. Hence, we find $\tilde{r} = \tilde{f}(\cdot; 0) \in \bar{\mathcal{C}}$ such that

$$\|r - \tilde{r}\|_\infty = \|f - \tilde{f}\|_\infty \leq \alpha. \quad (268)$$

Thus, $\bar{\mathcal{C}}$ is an α_i -cover of \mathcal{R}_i with the same cardinality as \mathcal{C} and we have

$$\text{cn}_{\alpha_i}(\mathcal{F}_i) \geq \text{cn}_{\alpha_i}(\mathcal{R}_i). \quad (269)$$

On the other hand, for any $\bar{\mathcal{C}}$ that is an α_i -cover of \mathcal{R}_i , if we construct \mathcal{C} as

$$\mathcal{C} = \left\{ f \mid f(\cdot, a_i) = \begin{cases} g(\cdot) + \epsilon_i & \text{if } a_i = 0 \\ a_i & \text{Otherwise} \end{cases}, \text{ for } r \in \bar{\mathcal{C}} \right\}. \quad (270)$$

Then for any $f \in \mathcal{F}_i$, we can construct g as $g = f(\cdot, 0)$. Since $\bar{\mathcal{C}}$ is an α_i -cover of \mathcal{R}_i , there exists $\tilde{r} \in \bar{\mathcal{C}}$ such that $\|r - \tilde{r}\| \leq \alpha_i$. Hence we find $\tilde{f} \in \mathcal{C}$ that satisfies

$$\tilde{f}(\cdot, a_i) = \begin{cases} \tilde{r}(\cdot) + \epsilon_i & \text{if } a_i = 0 \\ a_i & \text{Otherwise} \end{cases}. \quad (271)$$

We know $\tilde{f} \in \mathcal{C}$ and

$$\|f - \tilde{f}\|_\infty = \|r - \tilde{r}\|_\infty \leq \alpha_i. \quad (272)$$

Thus, \mathcal{C} is an α_i -cover of \mathcal{F}_i with the same cardinality as $\bar{\mathcal{C}}$ and we have

$$\text{cn}_{\alpha_i}(\mathcal{F}_i) \leq \text{cn}_{\alpha_i}(\mathcal{R}_i). \quad (273)$$

Therefore, by comparing (269) and (273) we conclude

$$\text{cn}_{\alpha_i}(\mathcal{R}_i) = \text{cn}_{\alpha_i}(\mathcal{F}_i). \quad (274)$$

□