
Practical Contextual Bandits with Feedback Graphs

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 While contextual bandit has a mature theory, effectively leveraging different feed-
2 back patterns to enhance the pace of learning remains unclear. Bandits with feed-
3 back graphs, which interpolates between the full information and bandit regimes,
4 provides a promising framework to mitigate the statistical complexity of learning.
5 In this paper, we propose and analyze an approach to contextual bandits with
6 feedback graphs based upon reduction to regression. The resulting algorithms are
7 computationally practical and achieve established minimax rates, thereby reducing
8 the statistical complexity in real-world applications.

9 1 Introduction

10 This paper is primarily concerned with increasing the pace of learning for contextual bandits [Auer
11 et al., 2002, Langford and Zhang, 2007]. While contextual bandits have enjoyed broad applicabil-
12 ity [Bouneffouf et al., 2020], the statistical complexity of learning with bandit feedback imposes a
13 data lower bound for application scenarios [Agarwal et al., 2012]. This has inspired various mitigation
14 strategies, including exploiting function class structure for improved experimental design [Zhu and
15 Mineiro, 2022], and composing with memory for learning with fewer samples [Rucker et al., 2022].
16 In this paper we exploit alternative graph feedback patterns to accelerate learning: intuitively, there is
17 no need to explore a potentially suboptimal action if a presumed better action, when exploited, yields
18 the necessary information.

19 The framework of bandits with feedback graphs is mature and provides a solid theoretical foundation
20 for incorporating additional feedback into an exploration strategy [Mannor and Shamir, 2011, Alon
21 et al., 2015, 2017]. Succinctly, in this framework, the observation of the learner is decided by a
22 directed feedback graph G : when an action is played, the learner observes the loss of every action
23 to which the chosen action is connected. When the graph only contains self-loops, this problem
24 reduces to the classic bandit case. For non-contextual bandits with feedback graphs, [Alon et al.,
25 2015] provides a full characterization on the minimax regret bound with respect to different graph
26 theoretic quantities associated with G according to the type of the feedback graph.

27 However, contextual bandits with feedback graphs have received less attention [Singh et al., 2020,
28 Wang et al., 2021]. Specifically, there is no prior work offering a solution for general feedback graphs
29 and function classes. In this work, we take an important step in this direction by adopting recently
30 developed minimax algorithm design principles in contextual bandits, which leverage realizability
31 and reduction to regression to construct practical algorithms with strong statistical guarantees [Foster
32 et al., 2018, Foster and Rakhlin, 2020, Foster et al., 2020, Foster and Krishnamurthy, 2021, Foster
33 et al., 2021, Zhu and Mineiro, 2022]. Using this strategy, we construct a practical algorithm for
34 contextual bandits with feedback graphs that achieves the optimal regret bound. Moreover, although
35 our primary concern is accelerating learning when the available feedback is more informative than
36 bandit feedback, our techniques also succeed when the available feedback is less informative than
37 bandit feedback, e.g., in spam filtering where some actions generate no feedback. More specifically,
38 our contributions are as follows.

39 **Contributions.** In this paper, we extend the minimax framework proposed in [Foster et al., 2021]
40 to contextual bandits with general feedback graphs, aiming to promote the utilization of different
41 feedback patterns in practical applications. Following [Foster and Rakhlin, 2020, Foster et al., 2021,
42 Zhu and Mineiro, 2022], we assume that there is an online regression oracle for supervised learning
43 on the loss. Based on this oracle, we propose SquareCB.G, the first algorithm for contextual bandits
44 with feedback graphs that operates via reduction to regression (Algorithm 1). Eliding regression regret
45 factors, our algorithm achieves the matching optimal regret bounds for deterministic feedback graphs,
46 with $\tilde{O}(\sqrt{\alpha T})$ regret for strongly observable graphs and $\tilde{O}(d^{\frac{1}{3}}T^{\frac{2}{3}})$ regret for weakly observable
47 graphs, where α and d are respectively the independence number and weakly domination number
48 of the feedback graph (see Section 3.2 for definitions). Notably, SquareCB.G is computationally
49 tractable, requiring the solution to a convex program (Theorem 3.6), which can be readily solved
50 with off-the-shelf convex solvers (Appendix A.3). In addition, we provide closed-form solutions for
51 specific cases of interest (Section 4), leading to a more efficient implementation of our algorithm.
52 Empirical results further showcase the effectiveness of our approach (Section 5).

53 2 Problem Setting and Preliminary

54 Throughout this paper, we let $[n]$ denote the set $\{1, 2, \dots, n\}$ for any positive integer n . We consider
55 the following contextual bandits problem with informed feedback graphs. The learning process goes
56 in T rounds. At each round $t \in [T]$, an environment selects a context $x_t \in \mathcal{X}$, a (stochastic) directed
57 feedback graph $G_t \in [0, 1]^{\mathcal{A} \times \mathcal{A}}$, and a loss distribution $\mathbb{P}_t : \mathcal{X} \rightarrow \Delta([-1, 1]^{\mathcal{A}})$; where \mathcal{A} is the
58 action set with finite cardinality K . For convenience, we use \mathcal{A} and $[K]$ interchangeably for denoting
59 the action set. Both G_t and x_t are revealed to the learner at the beginning of each round t . Then the
60 learner selects one of the actions $a_t \in \mathcal{A}$, while at the same time, the environment samples a loss
61 vector $\ell_t \in [-1, 1]^{\mathcal{A}}$ from $\mathbb{P}_t(\cdot | x_t)$. The learner then observes some information about ℓ_t according
62 to the feedback graph G_t . Specifically, for each action j , she observes the loss of action j with
63 probability $G_t(a_t, j)$, resulting in a realization A_t , which is the set of actions whose loss is observed.
64 With a slight abuse of notation, denote $G_t(\cdot | a)$ as the distribution of A_t when action a is picked.
65 We allow the context x_t , the (stochastic) feedback graphs G_t and the loss distribution $\mathbb{P}_t(\cdot | x_t)$ to be
66 selected by an adaptive adversary. When convenient, we will consider G to be a K -by- K matrix and
67 utilize matrix notation.

68 **Other Notations.** Let $\Delta(K)$ denote the set of all Radon probability measures over a set $[K]$.
69 $\text{conv}(S)$ represents the convex hull of a set S . Denote I as the identity matrix with an appropriate
70 dimension. For a K -dimensional vector v , $\text{diag}(v)$ denotes the K -by- K matrix with the i -th diagonal
71 entry v_i and other entries 0. We use $\mathbb{R}_{\geq 0}^K$ to denote the set of K -dimensional vectors with non-negative
72 entries. For a positive definite matrix $M \in \mathbb{R}^{K \times K}$, we define norm $\|z\|_M = \sqrt{z^\top M z}$ for any vector
73 $z \in \mathbb{R}^K$. We use the $\tilde{O}(\cdot)$ notation to hide factors that are polylogarithmic in K and T .

74 **Realizability.** We assume that the learner has access to a known function class $\mathcal{F} \subset (\mathcal{X} \times \mathcal{A} \mapsto$
75 $[-1, 1])$ which characterizes the mean of the loss for a given context-action pair, and we make the
76 following standard realizability assumption studied in the contextual bandit literature [Agarwal et al.,
77 2012, Foster et al., 2018, Foster and Rakhlin, 2020, Simchi-Levi and Xu, 2021].

78 **Assumption 1 (Realizability).** *There exists a regression function $f^* \in \mathcal{F}$ such that $\mathbb{E}[\ell_{t,a} | x_t] =$
79 $f^*(x_t, a)$ for any $a \in \mathcal{A}$ and across all $t \in [T]$.*

80 Two comments are in order. First, we remark that, similar to [Foster et al., 2020], misspecification can
81 be incorporated while maintaining computational efficiency, but we do not complicate the exposition
82 here. Second, Assumption 1 induces a “semi-adversarial” setting, wherein nature is completely free
83 to determine the context and graph sequences; and has considerable latitude in determining the loss
84 distribution subject to a mean constraint.

85 **Regret.** For each regression function $f \in \mathcal{F}$, let $\pi_f(x_t) := \text{argmin}_{a \in \mathcal{A}} f(x_t, a)$ denote the
86 induced policy, which chooses the action with the least loss with respect to f . Define $\pi^* := \pi_{f^*}$
87 as the optimal policy. We measure the performance of the learner via regret to π^* : $\text{Reg}_{\text{CB}} :=$
88 $\sum_{t=1}^T \ell_{t,a_t} - \sum_{t=1}^T \ell_{t,\pi^*(x_t)}$, which is the difference between the loss suffered by the learner and the
89 one if the learner applies policy π^* .

90 **Regression Oracle** We assume access to an online regression oracle Alg_{Sq} for function class \mathcal{F} ,
 91 which is an algorithm for online learning with squared loss. We consider the following protocol.
 92 At each round $t \in [T]$, the algorithm produces an estimator $\hat{f}_t \in \text{conv}(\mathcal{F})$, then receives a set of
 93 context-action-loss tuples $\{(x_t, a, \ell_{t,a})\}_{a \in A_t}$ where $A_t \subseteq \mathcal{A}$. The goal of the oracle is to accurately
 94 predict the loss as a function of the context and action, and we evaluate its performance via the square
 95 loss $\sum_{a \in A_t} (\hat{f}_t(x_t, a) - \ell_{t,a})^2$. We measure the oracle's cumulative performance via the following
 96 square-loss regret to the best function in \mathcal{F} .

97 **Assumption 2** (Bounded square-loss regret). *The regression oracle Alg_{Sq} guarantees that for any*
 98 *(potentially adaptively chosen) sequence $\{(x_t, a, \ell_{t,a})\}_{a \in A_t, t \in [T]}$ in which $A_t \subseteq \mathcal{A}$,*

$$\sum_{t=1}^T \sum_{a \in A_t} \left(\hat{f}_t(x_t, a) - \ell_{t,a} \right)^2 - \inf_{f \in \mathcal{F}} \sum_{t=1}^T \sum_{a \in A_t} (f(x_t, a) - \ell_{t,a})^2 \leq \mathbf{Reg}_{\text{Sq}}.$$

99 For finite \mathcal{F} , Vovk's aggregation algorithm yields $\mathbf{Reg}_{\text{Sq}} = \mathcal{O}(\log|\mathcal{F}|)$ [Vovk, 1995]. This regret is
 100 dependent upon the scale of the loss function, but this need not be linear with the size of A_t , e.g., the
 101 loss scale can be bounded by 2 in classification problems. See Foster and Krishnamurthy [2021] for
 102 additional examples of online regression algorithms.

103 3 Algorithms and Regret Bounds

104 In this section, we provide our main algorithms and results.

105 3.1 Algorithms via Minimax Reduction Design

106 Our approach is to adapt the minimax formulation of [Foster et al., 2021] to contextual bandits with
 107 feedback graphs. In the standard contextual bandits setting (that is, $G_t = I$ for all t), Foster et al.
 108 [2021] define the *Decision-Estimation Coefficient* (DEC) for a parameter $\gamma > 0$ as $\text{dec}_\gamma(\mathcal{F}) :=$
 109 $\sup_{\hat{f} \in \text{conv}(\mathcal{F}), x \in \mathcal{X}} \text{dec}_\gamma(\mathcal{F}; \hat{f}, x)$, where

$$\begin{aligned} \text{dec}_\gamma(\mathcal{F}; \hat{f}, x) &:= \inf_{p \in \Delta(K)} \text{dec}_\gamma(p, \mathcal{F}; \hat{f}, x) \\ &:= \inf_{p \in \Delta(K)} \sup_{\substack{a^* \in [K] \\ f^* \in \mathcal{F}}} \mathbb{E}_{a \sim p} \left[f^*(x, a) - f^*(x, a^*) - \frac{\gamma}{4} \cdot \left(\hat{f}(x, a) - f^*(x, a) \right)^2 \right]. \end{aligned} \quad (1)$$

110 Their proposed algorithm is as follows. At each round t , after receiving the context x_t , the algorithm
 111 first computes \hat{f}_t by calling the regression oracle. Then, it solves the solution p_t of the minimax
 112 problem defined in Eq. (1) with \hat{f} and x replaced by \hat{f}_t and x_t . Finally, the algorithm samples an
 113 action a_t from the distribution p_t and feeds the observation (x_t, a_t, ℓ_{t,a_t}) to the oracle. Foster et al.
 114 [2021] show that for any value γ , the algorithm above guarantees that

$$\mathbb{E}[\mathbf{Reg}_{\text{CB}}] \leq T \cdot \text{dec}_\gamma(\mathcal{F}) + \frac{\gamma}{4} \cdot \mathbf{Reg}_{\text{Sq}}. \quad (2)$$

115 However, the minimax problem Eq. (1) may not be solved efficiently in many cases. Therefore,
 116 instead of obtaining the distribution p_t which has the exact minimax value of Eq. (1), Foster et al.
 117 [2021] show that any distribution that gives an upper bound C_γ on $\text{dec}_\gamma(p, \mathcal{F}; \hat{f}, x)$ also works and
 118 enjoys a regret bound with $\text{dec}_\gamma(\mathcal{F})$ replaced by C_γ in Eq. (2).

119 To extend this framework to the setting with feedback graph G , we define $\text{dec}_\gamma(\mathcal{F}; \hat{f}, x, G)$ as follows

$$\begin{aligned} \text{dec}_\gamma(\mathcal{F}; \hat{f}, x, G) &:= \inf_{p \in \Delta(K)} \text{dec}_\gamma(p, \mathcal{F}; \hat{f}, x, G) \\ &:= \inf_{p \in \Delta(K)} \sup_{\substack{a^* \in [K] \\ f^* \in \mathcal{F}}} \mathbb{E}_{a \sim p} \left[f^*(x, a) - f^*(x, a^*) - \frac{\gamma}{4} \mathbb{E}_{A \sim G(\cdot|a)} \left[\sum_{a' \in \mathcal{A}} (\hat{f}_t(x, a') - f^*(x, a'))^2 \right] \right]. \end{aligned} \quad (3)$$

Algorithm 1 SquareCB.G. Note [Theorem 3.6](#) provides an efficient implementation of [Eq. \(4\)](#).

Input: parameter $\gamma \geq 4$, a regression oracle \mathbf{Alg}_{Sq}

for $t = 1, 2, \dots, T$ **do**

 Receive context x_t and directed feedback graph G_t .

 Obtain an estimator \hat{f}_t from the oracle \mathbf{Alg}_{Sq} .

 Compute the distribution $p_t \in \Delta(K)$ such that $p_t = \operatorname{argmin}_{p \in \Delta(K)} \overline{\text{dec}}_\gamma(p; \hat{f}_t, x_t, G_t)$, where

$$\begin{aligned} & \overline{\text{dec}}_\gamma(p; \hat{f}_t, x_t, G_t) \\ & := \sup_{\substack{a^* \in [K] \\ f^* \in \Phi}} \mathbb{E}_{a \sim p} \left[f^*(x_t, a) - f^*(x_t, a^*) - \frac{\gamma}{4} \mathbb{E}_{A \sim G_t(\cdot|a)} \left[\sum_{a' \in A} (\hat{f}_t(x_t, a') - f^*(x_t, a'))^2 \right] \right], \quad (4) \end{aligned}$$

 and $\Phi := \mathcal{X} \times [K] \mapsto \mathbb{R}$.

 Sample a_t from p_t and observe $\{\ell_{t,j}\}_{j \in A_t}$ where $A_t \sim G_t(\cdot|a_t)$.

 Feed the tuples $\{(x_t, j, \ell_{t,j})\}_{j \in A_t}$ to the oracle \mathbf{Alg}_{Sq} .

end

120 Compared with [Eq. \(1\)](#), the difference is that we replace the squared estimation error on action a by
121 the expected one on the observed set $A \sim G(\cdot|a)$, which intuitively utilizes more feedbacks from
122 the graph structure. When the feedback graph is the identity matrix, we recover [Eq. \(1\)](#). Based on
123 $\text{dec}_\gamma(\mathcal{F}; \hat{f}, x, G)$, our algorithm SquareCB.G is shown in [Algorithm 1](#). As what is done in [[Foster](#)
124 [et al., 2021](#)], in order to derive an efficient algorithm, instead of solving the distribution p_t with
125 respect to the supremum over $f^* \in \mathcal{F}$, we solve p_t that minimize $\overline{\text{dec}}_\gamma(p; \hat{f}, x_t, G_t)$ ([Eq. \(4\)](#)), which
126 takes supremum over $f^* \in (\mathcal{X} \times [K] \mapsto \mathbb{R})$, leading to an upper bound on $\text{dec}_\gamma(\mathcal{F}; \hat{f}, x_t, G_t)$. Then,
127 we receive the loss $\{\ell_{t,j}\}_{j \in A_t}$ and feed the tuples $\{(x_t, j, \ell_{t,j})\}_{j \in A_t}$ to the regression oracle \mathbf{Alg}_{Sq} .
128 Following a similar analysis to [[Foster et al., 2021](#)], we show that to bound the regret \mathbf{Reg}_{CB} , we
129 only need to bound $\overline{\text{dec}}_\gamma(p_t; \hat{f}_t, x_t, G_t)$.

130 **Theorem 3.1.** *Suppose $\overline{\text{dec}}_\gamma(p_t; \hat{f}_t, x_t, G_t) \leq C\gamma^{-\beta}$ for all $t \in [T]$ and some $\beta > 0$, [Algorithm 1](#)
131 with $\gamma = \max\{4, (CT)^{\frac{1}{\beta+1}} \mathbf{Reg}_{\text{Sq}}^{-\frac{1}{\beta+1}}\}$ guarantees that $\mathbb{E}[\mathbf{Reg}_{\text{CB}}] \leq \mathcal{O}\left(C^{\frac{1}{\beta+1}} T^{\frac{1}{\beta+1}} \mathbf{Reg}_{\text{Sq}}^{\frac{\beta}{\beta+1}}\right)$.*

132 The proof is deferred to [Appendix A](#). In [Section 3.3](#), we give an efficient implementation for solving
133 [Eq. \(4\)](#) via reduction to convex programming.

134 3.2 Regret Bounds

135 In this section, we derive regret bounds for [Algorithm 1](#) when G_t 's are specialized to deterministic
136 graphs, i.e., $G_t \in \{0, 1\}^{A \times A}$. We utilize discrete graph notation $G = ([K], E)$, where $E \subseteq$
137 $[K] \times [K]$; and define $N^{\text{in}}(G, j) \triangleq \{i \in A : (i, j) \in E\}$ as the set of nodes that can observe node
138 j . In this case, at each round t , the observed node set A_t is a deterministic set which contains any
139 node i satisfying $a_t \in N^{\text{in}}(G_t, i)$. In the following, we introduce several graph-theoretic concepts
140 for deterministic feedback graphs [[Alon et al., 2015](#)].

141 **Strongly/Weakly Observable Graphs.** For a directed graph $G = ([K], E)$, a node i is observable
142 if $N^{\text{in}}(G, i) \neq \emptyset$. An observable node is strongly observable if either $i \in N^{\text{in}}(G, i)$ or $N^{\text{in}}(G, i) =$
143 $[K] \setminus \{i\}$, and weakly observable otherwise. Similarly, a graph is observable if all its nodes are
144 observable. An observable graph is strongly observable if all nodes are strongly observable, and
145 weakly observable otherwise. Self-aware graphs are a special type of strongly observable graphs
146 where $i \in N^{\text{in}}(G, i)$ for all $i \in [K]$.

147 **Independent Set and Weakly Dominating Set.** An independence set of a directed graph is a subset
148 of nodes in which no two distinct nodes are connected. The size of the largest independence set of a
149 graph is called its independence number. For a weakly observable graph $G = ([K], E)$, a weakly
150 dominating set is a subset of nodes $D \subseteq [K]$ such that for any node j in G without a self-loop, there
151 exists $i \in D$ such that directed edge $(i, j) \in E$. The size of the smallest weakly dominating set of a

152 graph is called its weak domination number. Alon et al. [2015] show that in non-contextual bandits
 153 with a fixed feedback graph G , the minimax regret bound is $\tilde{\Theta}(\sqrt{\alpha T})$ when G is strongly observable
 154 and $\tilde{\Theta}(d^{\frac{1}{3}} T^{\frac{2}{3}})$ when G is weakly observable, where α and d are the independence number and the
 155 weak domination number of G , respectively.

156 3.2.1 Strongly Observable Graphs

157 In the following theorem, we show the regret bound of Algorithm 1 for strongly observable graphs.

158 **Theorem 3.2** (Strongly observable graphs). *Suppose that the feedback graph G_t is deterministic and*
 159 *strongly observable with independence number no more than α . Then Algorithm 1 guarantees that*
 160 $\overline{\text{dec}}_{\gamma}(p_t; \hat{f}_t, x_t, G_t) \leq \mathcal{O}\left(\frac{\alpha \log(K\gamma)}{\gamma}\right).$

161 In contrast to existing works that derive a closed-form solution of p_t in order to show how large the
 162 DEC can be [Foster and Rakhlin, 2020, Foster and Krishnamurthy, 2021], in our case we prove the
 163 upper bound of $\overline{\text{dec}}_{\gamma}(p_t; \hat{f}_t, x_t, G_t)$ by using the Sion’s minimax theorem and the graph-theoretic
 164 lemma proven in [Alon et al., 2015]. The proof is deferred to Appendix A.1. Combining Theorem 3.2
 165 and Theorem 3.1, we directly have the following corollary:

166 **Corollary 3.3.** *Suppose that G_t is deterministic, strongly observable, and has independence number*
 167 *no more than α for all $t \in [T]$. Algorithm 1 with choice $\gamma = \max\left\{4, \sqrt{\alpha T / \mathbf{Reg}_{\text{Sq}}}\right\}$ guarantees*
 168 *that $\mathbb{E}[\mathbf{Reg}_{\text{CB}}] \leq \tilde{\mathcal{O}}(\sqrt{\alpha T \mathbf{Reg}_{\text{Sq}}})$.*

169 3.2.2 Weakly Observable Graphs

170 For the weakly observable graph, we have the following theorem.

171 **Theorem 3.4** (Weakly observable graphs). *Suppose that the feedback graph G_t is deterministic and*
 172 *weakly observable with weak domination number no more than d . Then Algorithm 1 with $\gamma \geq 16d$*
 173 *guarantees that $\overline{\text{dec}}_{\gamma}(p_t; \hat{f}_t, x_t, G_t) \leq \mathcal{O}\left(\sqrt{\frac{d}{\gamma}} + \frac{\tilde{\alpha} \log(K\gamma)}{\gamma}\right)$, where $\tilde{\alpha}$ is the independence number*
 174 *of the subgraph induced by nodes with self-loops in G_t .*

175 The proof is deferred to Appendix A.2. Similar to Theorem 3.2, we do not derive a closed-form
 176 solution to the strategy p_t but prove this upper bound using the minimax theorem. Combining
 177 Theorem 3.4 and Theorem 3.1, we are able to obtain the following regret bound for weakly observable
 178 graphs, whose proof is deferred to Appendix A.2.

179 **Corollary 3.5.** *Suppose that G_t is deterministic, weakly observable, and has weak domination*
 180 *number no more than d for all $t \in [T]$. In addition, suppose that the independence number of the*
 181 *subgraph induced by nodes with self-loops in G_t is no more than $\tilde{\alpha}$ for all $t \in [T]$. Then, Algorithm 1*
 182 *with $\gamma = \max\{16d, \sqrt{\tilde{\alpha} T / \mathbf{Reg}_{\text{Sq}}}, d^{\frac{1}{3}} T^{\frac{2}{3}} \mathbf{Reg}_{\text{Sq}}^{-\frac{2}{3}}\}$ guarantees that $\mathbb{E}[\mathbf{Reg}_{\text{CB}}] \leq \tilde{\mathcal{O}}(d^{\frac{1}{3}} T^{\frac{2}{3}} \mathbf{Reg}_{\text{Sq}}^{\frac{1}{3}} +$
 183 $\sqrt{\tilde{\alpha} T \mathbf{Reg}_{\text{Sq}}})$.*

184 3.3 Implementation

185 In this section, we show that solving $\text{argmin}_{p \in \Delta(K)} \overline{\text{dec}}_{\gamma}(p; \hat{f}, x, G)$ in Algorithm 1 is equivalent to
 186 solving a convex program, which can be easily and efficiently implemented in practice.

187 **Theorem 3.6.** *Solving $\text{argmin}_{p \in \Delta(K)} \overline{\text{dec}}_{\gamma}(p; \hat{f}, x, G)$ is equivalent to solving the following convex*
 188 *optimization problem.*

$$\begin{aligned} \min_{p \in \Delta(K), z} \quad & p^{\top} \hat{f} + z & (5) \\ \text{subject to} \quad & \forall a \in [K] : \frac{1}{\gamma} \|p - e_a\|_{\text{diag}(G^{\top} p)^{-1}}^2 \leq \hat{f}(x, a) + z, \\ & G^{\top} p \succ 0, \end{aligned}$$

189 where \hat{f} in the objective is a shorthand for $\hat{f}(x, \cdot) \in \mathbb{R}^K$, e_a is the a -th standard basis vector, and \succ
 190 means element-wise greater.

191 The proof is deferred to [Appendix A.4](#). Note that this implementation is not restricted to the deter-
 192 ministic feedback graphs but applies to the general stochastic feedback graph case. In [Appendix A.3](#),
 193 we provide the 20 lines of Python code that solves [Eq. \(5\)](#).

194 4 Examples with Closed-Form Solutions

195 In this section, we present examples and corresponding closed-form solutions of p that make the
 196 value $\overline{\text{dec}}_\gamma(p; \hat{f}, x, G)$ upper bounded by at most a constant factor of $\min_p \overline{\text{dec}}_\gamma(p; \hat{f}, x, G)$. This
 197 offers an alternative to solving the convex program defined in [Theorem 3.6](#) for special (and practically
 198 relevant) cases, thereby enhancing the efficiency of our algorithm. All the proofs are deferred to
 199 [Appendix B](#).

200 **Cops-and-Robbers Graph.** The ‘‘cops-and-robbers’’ feedback graph $G_{\text{CR}} = 11^\top - I$, also known
 201 as the loopless clique, is the full feedback graph removing self-loops. Therefore, G_{CR} is strongly
 202 observable with independence number $\alpha = 1$. Let a_1 be the node with the smallest value of \hat{f} and
 203 a_2 be the node with the second smallest value of \hat{f} . Our proposed closed-form distribution p is only
 204 supported on $\{a_1, a_2\}$ and defined as follows:

$$p_{a_1} = 1 - \frac{1}{2 + \gamma(\hat{f}_{a_2} - \hat{f}_{a_1})}, \quad p_{a_2} = \frac{1}{2 + \gamma(\hat{f}_{a_2} - \hat{f}_{a_1})}. \quad (6)$$

205 In the following proposition, we show that with the construction of p in [Eq. \(6\)](#), $\overline{\text{dec}}_\gamma(p; \hat{f}, x, G_{\text{CR}})$ is
 206 upper bounded by $\mathcal{O}(1/\gamma)$, which matches the order of $\min_p \overline{\text{dec}}_\gamma(p; \hat{f}, x, G)$ based on [Theorem 3.2](#)
 207 since $\alpha = 1$.

208 **Proposition 1.** *When $G = G_{\text{CR}}$, given any \hat{f} , context x , the closed-form distribution p in [Eq. \(6\)](#)
 209 guarantees that $\overline{\text{dec}}_\gamma(p; \hat{f}, x, G_{\text{CR}}) \leq \mathcal{O}\left(\frac{1}{\gamma}\right)$.*

210 **Apple Tasting Graph.** The apple tasting feedback graph $G_{\text{AT}} = \begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix}$ consists of two nodes,
 211 where the first node reveals all and the second node reveals nothing. This scenario was originally
 212 proposed by [Helmbold et al. \[2000\]](#) and recently denoted the spam filtering graph [[van der Hoeven](#)
 213 [et al., 2021](#)]. The independence number of G_{AT} is 1. Let \hat{f}_1 be the oracle prediction for the first node
 214 and let \hat{f}_2 be the prediction for the second node. We present a closed-form solution p for [Eq. \(4\)](#) as
 215 follows:

$$p_1 = \begin{cases} 1 & \hat{f}_1 \leq \hat{f}_2 \\ \frac{2}{4 + \gamma(\hat{f}_1 - \hat{f}_2)} & \hat{f}_1 > \hat{f}_2 \end{cases}, \quad p_2 = 1 - p_1. \quad (7)$$

216 We show that this distribution p satisfies that $\overline{\text{dec}}_\gamma(p; \hat{f}, x, G_{\text{AT}})$ is upper bounded by $\mathcal{O}(1/\gamma)$ in the
 217 following proposition. We remark that directly applying results from [[Foster et al., 2021](#)] cannot lead
 218 to a valid upper bound since the second node does not have a self-loop.

219 **Proposition 2.** *When $G = G_{\text{AT}}$, given any \hat{f} , context x , the closed-form distribution p in [Eq. \(7\)](#)
 220 guarantees that $\overline{\text{dec}}_\gamma(p; \hat{f}, x, G_{\text{AT}}) \leq \mathcal{O}\left(\frac{1}{\gamma}\right)$.*

221 **Inventory Graph.** In this application, the algorithm needs to decide the inventory level in order
 222 to fulfill the realized demand arriving at each round. Specifically, there are K possible chosen
 223 inventory levels $a_1 < a_2 < \dots < a_K$ and the feedback graph G_{inv} has entries $G(i, j) = 1$ for all
 224 $1 \leq j \leq i \leq K$ and $G(i, j) = 0$ otherwise, meaning that picking the inventory level a_i informs
 225 about all actions $a_{j \leq i}$. This is because items are consumed until either the demand or the inventory is
 226 exhausted. The independence number of G_{inv} is 1. Therefore, (very) large K is statistically tractable,
 227 but naively solving the convex program [Eq. \(5\)](#) requires superlinear in K computational cost. We
 228 show in the following proposition that there exists an analytic form of p , which guarantees that
 229 $\overline{\text{dec}}_\gamma(p; \hat{f}, x, G_{\text{inv}})$ can be bounded by $\mathcal{O}(1/\gamma)$.

230 **Proposition 3.** *When $G = G_{\text{inv}}$, given any \hat{f} , context x , there exists a closed-form distribution
 231 $p \in \Delta(K)$ guaranteeing that $\overline{\text{dec}}_\gamma(p; \hat{f}, x, G_{\text{inv}}) \leq \mathcal{O}\left(\frac{1}{\gamma}\right)$, where p is defined as follows: $p_j =$
 232 $\max\left\{\frac{1}{1 + \gamma(\hat{f}_j - \min_i \hat{f}_i)} - \sum_{j' > j} p_{j'}, 0\right\}$ for all $j \in [K]$.*

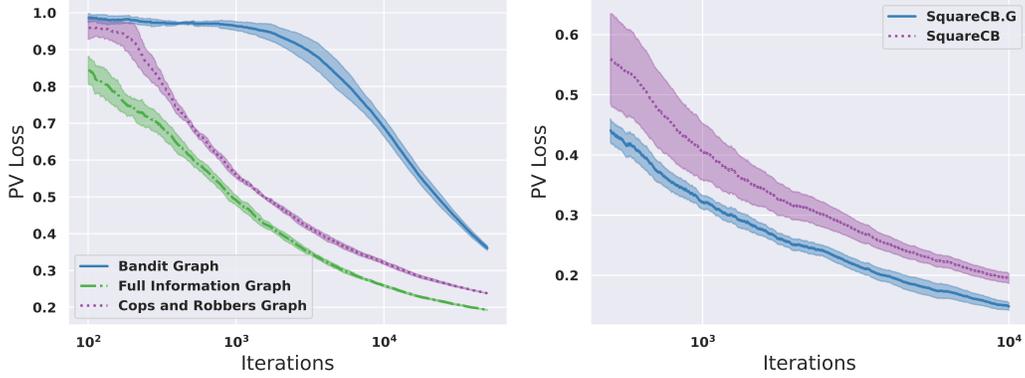


Figure 1: **Left figure:** Performance of SquareCB.G on RCV1 dataset under three different feedback graphs. **Right figure:** Performance comparison between SquareCB.G and SquareCB under random directed self-aware feedback graphs.

233 **Undirected Self-Aware Graph.** For the undirected and self-aware feedback graph G , which means
 234 that G is symmetric and has diagonal entries all 1, we also show that a certain closed-form solution
 235 of p satisfies that $\overline{\text{dec}}_\gamma(p; \hat{f}, x, G)$ is bounded by $\mathcal{O}(\frac{\alpha}{\gamma})$.

236 **Proposition 4.** When G is an undirected self-aware graph, given any \hat{f} , context x , there exists a
 237 closed-form distribution $p \in \Delta(K)$ guaranteeing that $\overline{\text{dec}}_\gamma(p; \hat{f}, x, G) \leq \mathcal{O}(\frac{\alpha}{\gamma})$.

238 5 Experiments

239 In this section, we use empirical results to demonstrate the significant benefits of SquareCB.G
 240 in leveraging the graph feedback structure and its superior effectiveness compared to SquareCB.
 241 Following Foster and Krishnamurthy [2021], we use progressive validation (PV) loss as the evaluation
 242 metric, defined as $L_{\text{pv}}(T) = \frac{1}{T} \sum_{t=1}^T \ell_{t, a_t}$. All the feedback graphs used in the experiments are
 243 deterministic. We run experiments on CPU Intel Xeon Gold 6240R 2.4G and the convex program
 244 solver is implemented via Vowpal Wabbit [Langford et al., 2007].

245 5.1 SquareCB.G under Different Feedback Graphs

246 In this subsection, we show that our SquareCB.G benefits from considering the graph structure
 247 by evaluating the performance of SquareCB.G under three different feedback graphs. We conduct
 248 experiments on RCV1 dataset and leave the implementation details in Appendix C.1.

249 The performances of SquareCB.G under bandit graph, full information graph and cops-and-robbers
 250 graph are shown in the left part of Figure 1. We observe that SquareCB.G performs the best under
 251 full information graph and performs worst under bandit graph. Under the cops-and-robbers graph,
 252 much of the gap between bandit and full information is eliminated. This improvement demonstrates
 253 the benefit of utilizing graph feedback for accelerating learning.

254 5.2 Comparison between SquareCB.G and SquareCB

255 In this subsection, we compare the effectiveness of SquareCB.G with the SquareCB algorithm. To
 256 ensure a fair comparison, both algorithms update the regressor using the same feedbacks based on
 257 the graph. The only distinction lies in how they calculate the action probability distribution. We
 258 summarize the main results here and leave the implementation details in Appendix C.2.

259 5.2.1 Results on Random Directed Self-aware Graphs

260 We conduct experiments on RCV1 dataset using random directed self-aware feedback graphs. Specif-
 261 ically, the diagonal elements are all 1, and off-diagonal entries are drawn from a Bernoulli($3/4$)

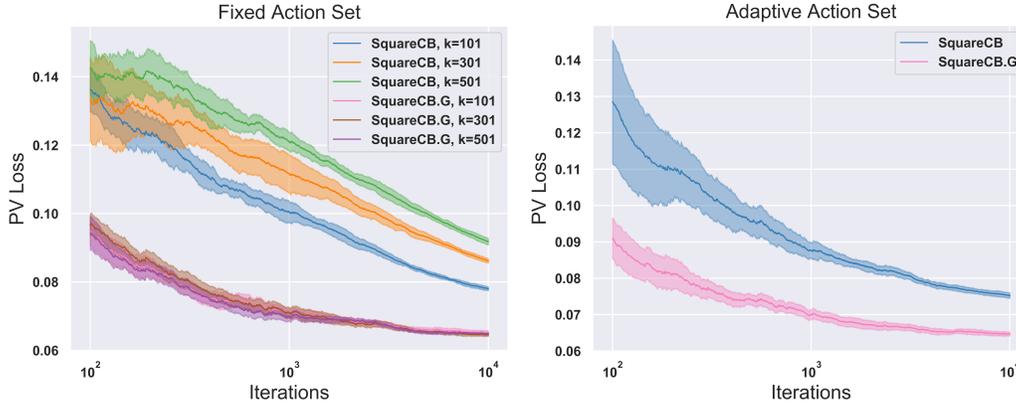


Figure 2: Performance comparison between SquareCB.G and SquareCB on synthetic inventory dataset. **Left figure:** Results under fixed discretized action set. **Right figure:** Results under adaptive discretization of the action set. Both figures show the superiority of SquareCB.G compared with SquareCB.

262 distribution. The results are presented in the right part of Figure 1. Our SquareCB.G consistently
 263 outperforms SquareCB and demonstrates lower variance, particularly when the number of iterations
 264 was small. This is because when there are fewer samples available to train the regressor, it is more
 265 crucial to design an effective algorithm that can leverage the graph feedback information.

266 5.2.2 Results on Synthetic Inventory Dataset

267 In the inventory graph experiments, we create a synthetic inventory dataset and design a loss function
 268 for each inventory level $a_t \in [0, 1]$ with demand $d_t \in [0, 1]$. Since the action set $[0, 1]$ is continuous,
 269 we discretize the action set in two different ways to apply the algorithms.

270 **Fixed discretized action set.** In this setting, we discretize the action set using fixed grid size
 271 $\varepsilon \in \{\frac{1}{100}, \frac{1}{300}, \frac{1}{500}\}$, which leads to a finite action set \mathcal{A} of size $\frac{1}{\varepsilon} + 1$. Note that according
 272 to Theorem 3.2, our regret *does not* scale with the size of the action set (to within polylog factors), as
 273 the independence number is always 1. The results are shown in the left part of Figure 2.

274 We remark several observations from the results. First, our algorithm SquareCB.G outperforms
 275 SquareCB for all choices $K \in \{101, 301, 501\}$. This indicates that SquareCB.G utilizes a bet-
 276 ter exploration scheme and effectively leverages the structure of G_{inv} . Second, we observe that
 277 SquareCB.G indeed does not scale with the size of the discretized action set \mathcal{A} , since under different
 278 discretization scales, SquareCB.G has similar performances and the slight differences are from the
 279 improved approximation error with finer discretization. This matches the theoretical guarantee that
 280 we prove in Theorem 3.2. On the other hand, SquareCB does perform worse when the size of the
 281 action set increases, matching its theoretical guarantee which scales with the square root of the size
 282 of the action set.

283 **Adaptively changing action set.** In this setting, we adaptively discretize the action set $[0, 1]$
 284 according to the index of the current round. Specifically, for SquareCB.G, we uniformly discretize
 285 the action set $[0, 1]$ with size $\lceil \sqrt{t} \rceil$, whose total discretization error is $\mathcal{O}(\sqrt{T})$ due to the Lipschitzness
 286 of the loss function. For SquareCB, to optimally balance the dependency on the size of the action
 287 set and the discretization error, we uniformly discretize the action set $[0, 1]$ into $\lceil t^{\frac{1}{3}} \rceil$ actions. The
 288 results are illustrated in the right part of Figure 2. We can observe that SquareCB.G consistently
 289 outperforms SquareCB by a clear margin.

290 6 Related Work

291 Multi-armed bandits with feedback graphs have been extensively studied. An early example is
 292 the apple tasting problem of Helmbold et al. [2000]. The general formulation was introduced by

293 [Mannor and Shamir \[2011\]](#). [Alon et al. \[2015\]](#) characterized the minimax rates in terms of graph-
294 theoretic quantities. Follow-on work includes relaxing the assumption that the graph is observed
295 prior to decision [[Cohen et al., 2016](#)]; analyzing the distinction between the stochastic and adversarial
296 settings [[Alon et al., 2017](#)]; considering stochastic feedback graphs [[Li et al., 2020](#), [Esposito et al.,](#)
297 [2022](#)]; instance-adaptivity [[Ito et al., 2022](#)]; data-dependent regret bound [[Lykouris et al., 2018](#), [Lee](#)
298 [et al., 2020](#)]; and high-probability regret under adaptive adversary [[Neu, 2015](#), [Luo et al., 2023](#)].

299 The contextual bandit problem with feedback graphs has received relatively less attention. [Wang](#)
300 [et al. \[2021\]](#) provide algorithms for adversarial linear bandits with uninformed graphs and stochastic
301 contexts. However, this work assumes several unrealistic assumptions on both the policy class
302 and the context space and is not comparable to our setting, since we consider the informed graph
303 setting with adversarial context. [Singh et al. \[2020\]](#) study a stochastic linear bandits with informed
304 feedback graphs and are able to improve over the instance-optimal regret bound for bandits derived
305 in [[Lattimore and Szepesvari, 2017](#)] by utilizing the additional graph-based feedbacks.

306 Our work is also closely related to the recent progress in designing efficient algorithms for classic
307 contextual bandits. Starting from [[Langford and Zhang, 2007](#)], numerous works have been done to the
308 development of practically efficient algorithms, which are based on reduction to either cost-sensitive
309 classification oracles [[Dudik et al., 2011](#), [Agarwal et al., 2014](#)] or online regression oracles [[Foster](#)
310 [and Rakhlin, 2020](#), [Foster et al., 2020, 2021](#), [Zhu and Mineiro, 2022](#)]. Following the latter trend, our
311 work assumes access to an online regression oracle and extends the classic bandit problems to the
312 bandits with general feedback graphs.

313 7 Discussion

314 In this paper, we consider the design of practical contextual bandits algorithm with provable guaran-
315 tees. Specifically, we propose the first efficient algorithm that achieves near-optimal regret bound for
316 contextual bandits with general directed feedback graphs with an online regression oracle.

317 While we study the informed graph feedback setting, where the entire feedback graph is exposed to
318 the algorithm prior to each decision, many practical problems of interest are possibly uninformed
319 graph feedback problems, where the graph is unknown at the decision time. It is unclear how to
320 formulate an analogous minimax problem to [Eq. \(1\)](#) under the uninformed setting. One idea is to
321 consume the additional feedback in the online regressor and adjust the prediction loss to reflect this
322 additional structure, e.g., using the more general version of the E2D framework which incorporates
323 arbitrary side observations [[Foster et al., 2021](#)]. [Cohen et al. \[2016\]](#) consider this uninformed setting
324 in the non-contextual case and prove a sharp distinction between the adversarial and stochastic
325 settings: even if the graphs are all strongly observable with bounded independence number, in the
326 adversarial setting the minimax regret is $\Theta(T)$ whereas in the stochastic setting the minimax regret
327 is $\Theta(\sqrt{\alpha T})$. Intriguingly, our setting is semi-adversarial due to realizability of the mean loss, and
328 therefore it is apriori unclear whether the negative adversarial result applies.

329 In addition, bandits with graph feedback problems often present with associated policy constraints,
330 e.g., for the apple tasting problem, it is natural to rate limit the informative action. Therefore, another
331 interesting direction is to combine our algorithm with the recent progress in contextual bandits with
332 knapsack [[Slivkins and Foster, 2022](#)], leading to more practical algorithms.

333 References

- 334 Alekh Agarwal, Miroslav Dudík, Satyen Kale, John Langford, and Robert Schapire. Contextual
335 bandit learning with predictable rewards. In *Artificial Intelligence and Statistics*, pages 19–26.
336 PMLR, 2012.
- 337 Alekh Agarwal, Daniel Hsu, Satyen Kale, John Langford, Lihong Li, and Robert Schapire. Taming
338 the monster: A fast and simple algorithm for contextual bandits. In *International Conference on*
339 *Machine Learning*, pages 1638–1646. PMLR, 2014.
- 340 Noga Alon, Nicolo Cesa-Bianchi, Claudio Gentile, and Yishay Mansour. From bandits to experts: A
341 tale of domination and independence. *Advances in Neural Information Processing Systems*, 26,
342 2013.

- 343 Noga Alon, Nicolo Cesa-Bianchi, Ofer Dekel, and Tomer Koren. Online learning with feedback
344 graphs: Beyond bandits. In *Conference on Learning Theory*, pages 23–35. PMLR, 2015.
- 345 Noga Alon, Nicolo Cesa-Bianchi, Claudio Gentile, Shie Mannor, Yishay Mansour, and Ohad Shamir.
346 Nonstochastic multi-armed bandits with graph-structured feedback. *SIAM Journal on Computing*,
347 46(6):1785–1826, 2017.
- 348 Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. The nonstochastic multiarmed
349 bandit problem. *SIAM journal on computing*, 32(1):48–77, 2002.
- 350 Djallel Bouneffouf, Irina Rish, and Charu Aggarwal. Survey on applications of multi-armed and
351 contextual bandits. In *2020 IEEE Congress on Evolutionary Computation (CEC)*, pages 1–8. IEEE,
352 2020.
- 353 Alon Cohen, Tamir Hazan, and Tomer Koren. Online learning with feedback graphs without the
354 graphs. In *International Conference on Machine Learning*, pages 811–819. PMLR, 2016.
- 355 Miroslav Dudik, Daniel Hsu, Satyen Kale, Nikos Karampatziakis, John Langford, Lev Reyzin, and
356 Tong Zhang. Efficient optimal learning for contextual bandits. In *Proceedings of the Twenty-Seventh
357 Conference on Uncertainty in Artificial Intelligence*, pages 169–178, 2011.
- 358 Emmanuel Esposito, Federico Fusco, Dirk van der Hoeven, and Nicolò Cesa-Bianchi. Learning on
359 the edge: Online learning with stochastic feedback graphs. In *Advances in Neural Information
360 Processing Systems*, 2022.
- 361 Dylan Foster and Alexander Rakhlin. Beyond ucb: Optimal and efficient contextual bandits with
362 regression oracles. In *International Conference on Machine Learning*, pages 3199–3210. PMLR,
363 2020.
- 364 Dylan Foster, Alekh Agarwal, Miroslav Dudik, Haipeng Luo, and Robert Schapire. Practical
365 contextual bandits with regression oracles. In *International Conference on Machine Learning*,
366 pages 1539–1548. PMLR, 2018.
- 367 Dylan J Foster and Akshay Krishnamurthy. Efficient first-order contextual bandits: Prediction,
368 allocation, and triangular discrimination. *Advances in Neural Information Processing Systems*, 34:
369 18907–18919, 2021.
- 370 Dylan J Foster, Claudio Gentile, Mehryar Mohri, and Julian Zimmert. Adapting to misspecification in
371 contextual bandits. *Advances in Neural Information Processing Systems*, 33:11478–11489, 2020.
- 372 Dylan J Foster, Sham M Kakade, Jian Qian, and Alexander Rakhlin. The statistical complexity of
373 interactive decision making. *arXiv preprint arXiv:2112.13487*, 2021.
- 374 David P Helmbold, Nicholas Littlestone, and Philip M Long. Apple tasting. *Information and
375 Computation*, 161(2):85–139, 2000.
- 376 Shinji Ito, Taira Tsuchiya, and Junya Honda. Nearly optimal best-of-both-worlds algorithms for
377 online learning with feedback graphs. In *Advances in Neural Information Processing Systems*,
378 2022.
- 379 John Langford and Tong Zhang. The epoch-greedy algorithm for contextual multi-armed bandits.
380 *Advances in neural information processing systems*, 20(1):96–1, 2007.
- 381 John Langford, Lihong Li, and Alex Strehl. Vowpal wabbit online learning project, 2007.
- 382 Tor Lattimore and Csaba Szepesvari. The end of optimism? an asymptotic analysis of finite-armed
383 linear bandits. In *Artificial Intelligence and Statistics*, pages 728–737. PMLR, 2017.
- 384 Chung-Wei Lee, Haipeng Luo, and Mengxiao Zhang. A closer look at small-loss bounds for bandits
385 with graph feedback. In *Conference on Learning Theory*, pages 2516–2564. PMLR, 2020.
- 386 David D Lewis, Yiming Yang, Tony Russell-Rose, and Fan Li. Rcv1: A new benchmark collection
387 for text categorization research. *Journal of machine learning research*, 5(Apr):361–397, 2004.

- 388 Shuai Li, Wei Chen, Zheng Wen, and Kwong-Sak Leung. Stochastic online learning with probabilistic
389 graph feedback. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34,
390 pages 4675–4682, 2020.
- 391 Haipeng Luo, Hanghang Tong, Mengxiao Zhang, and Yuheng Zhang. Improved high-probability
392 regret for adversarial bandits with time-varying feedback graphs. In *International Conference on*
393 *Algorithmic Learning Theory*, pages 1074–1100. PMLR, 2023.
- 394 Thodoris Lykouris, Karthik Sridharan, and Éva Tardos. Small-loss bounds for online learning with
395 partial information. In *Conference on Learning Theory*, pages 979–986. PMLR, 2018.
- 396 Shie Mannor and Ohad Shamir. From bandits to experts: On the value of side-observations. *Advances*
397 *in Neural Information Processing Systems*, 24, 2011.
- 398 Gergely Neu. Explore no more: Improved high-probability regret bounds for non-stochastic bandits.
399 *Advances in Neural Information Processing Systems*, 28, 2015.
- 400 Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor
401 Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style,
402 high-performance deep learning library. *Advances in neural information processing systems*, 32,
403 2019.
- 404 Mark Rucker, Joran T Ash, John Langford, Paul Mineiro, and Ida Momennejad. Eigen memory tree.
405 *arXiv preprint arXiv:2210.14077*, 2022.
- 406 David Simchi-Levi and Yunzong Xu. Bypassing the monster: A faster and simpler optimal algorithm
407 for contextual bandits under realizability. *Mathematics of Operations Research*, 2021.
- 408 Rahul Singh, Fang Liu, Xin Liu, and Ness Shroff. Contextual bandits with side-observations. *arXiv*
409 *preprint arXiv:2006.03951*, 2020.
- 410 Aleksandrs Slivkins and Dylan Foster. Efficient contextual bandits with knapsacks via regression.
411 *arXiv preprint arXiv:2211.07484*, 2022.
- 412 Dirk van der Hoeven, Federico Fusco, and Nicolò Cesa-Bianchi. Beyond bandit feedback in online
413 multiclass classification. *Advances in Neural Information Processing Systems*, 34:13280–13291,
414 2021.
- 415 Vladimir G Vovk. A game of prediction with expert advice. In *Proceedings of the eighth annual*
416 *conference on Computational learning theory*, pages 51–60, 1995.
- 417 Lingda Wang, Bingcong Li, Huozhi Zhou, Georgios B Giannakis, Lav R Varshney, and Zhizhen Zhao.
418 Adversarial linear contextual bandits with graph-structured side observations. In *Proceedings of*
419 *the AAAI Conference on Artificial Intelligence*, pages 10156–10164, 2021.
- 420 Yinglun Zhu and Paul Mineiro. Contextual bandits with smooth regret: Efficient learning in continu-
421 ous action spaces. In *International Conference on Machine Learning*, pages 27574–27590. PMLR,
422 2022.

423 **Broader Impact**

424 This work is mostly theoretical, and we do not foresee any negative ethical or societal outcomes. Our
 425 algorithms can be applied for many applications with context and partial information feedback such
 426 as recommendation systems.

427 **A Omitted Details in Section 3**

428 **Theorem 3.1.** Suppose $\overline{\text{dec}}_\gamma(p_t; \hat{f}_t, x_t, G_t) \leq C\gamma^{-\beta}$ for all $t \in [T]$ and some $\beta > 0$, *Algorithm 1*
 429 with $\gamma = \max\{4, (CT)^{\frac{1}{\beta+1}} \mathbf{Reg}_{\text{Sq}}^{-\frac{1}{\beta+1}}\}$ guarantees that $\mathbb{E}[\mathbf{Reg}_{\text{CB}}] \leq \mathcal{O}\left(C^{\frac{1}{\beta+1}} T^{\frac{1}{\beta+1}} \mathbf{Reg}_{\text{Sq}}^{\frac{\beta}{\beta+1}}\right)$.

430 *Proof.* Following [Foster et al., 2020], we decompose \mathbf{Reg}_{CB} as follows:

$$\begin{aligned}
 & \mathbb{E}[\mathbf{Reg}_{\text{CB}}] \\
 &= \mathbb{E}\left[\sum_{t=1}^T f^*(x_t, a_t) - \sum_{t=1}^T f^*(x_t, \pi^*(x_t))\right] \\
 &= \mathbb{E}\left[\sum_{t=1}^T \left(f^*(x_t, a_t) - f^*(x_t, \pi^*(x_t)) - \frac{\gamma}{4} \mathbb{E}_{A \sim G_t(\cdot|a_t)} \left[\sum_{a \in A} (\hat{f}_t(x_t, a) - f^*(x_t, a))^2\right]\right)\right] \\
 &\quad + \frac{\gamma}{4} \mathbb{E}\left[\sum_{t=1}^T \mathbb{E}_{A \sim G_t(\cdot|a_t)} \left[\sum_{a \in A} (\hat{f}_t(x_t, a) - f^*(x_t, a))^2\right]\right] \\
 &\leq \mathbb{E}\left[\sum_{t=1}^T \max_{\substack{a^* \in [K] \\ f \in (\mathcal{X} \times [K]) \rightarrow \mathbb{R}}} \mathbb{E}_{a_t \sim p_t} \left[f(x_t, a_t) - f(x_t, a^*) - \frac{\gamma}{4} \mathbb{E}_{A \sim G_t(\cdot|a_t)} \left[\sum_{a \in A} (\hat{f}_t(x_t, a) - f(x_t, a))^2\right]\right]\right] \\
 &\quad + \frac{\gamma}{4} \mathbb{E}\left[\sum_{t=1}^T \mathbb{E}_{A \sim G_t(\cdot|a_t)} \left[\sum_{a \in A} (\hat{f}_t(x_t, a) - f^*(x_t, a))^2\right]\right] \\
 &= \mathbb{E}\left[\sum_{t=1}^T \overline{\text{dec}}_\gamma(p_t; \hat{f}_t, x_t, G_t)\right] + \frac{\gamma}{4} \mathbb{E}\left[\sum_{t=1}^T \mathbb{E}_{A \sim G_t(\cdot|a_t)} \left[\sum_{a \in A} (\hat{f}_t(x_t, a) - f^*(x_t, a))^2\right]\right] \quad (8) \\
 &\leq CT\gamma^{-\beta} + \frac{\gamma}{4} \mathbb{E}\left[\sum_{t=1}^T \mathbb{E}_{A \sim G_t(\cdot|a_t)} \left[\sum_{a \in A} (\hat{f}_t(x_t, a) - f^*(x_t, a))^2\right]\right].
 \end{aligned}$$

431 Next, since $\mathbb{E}[\ell_{t,a} | x_t] = f^*(x_t, a)$ for all $t \in [T]$ and $a \in \mathcal{A}$, we know that

$$\begin{aligned}
 & \mathbb{E}\left[\sum_{t=1}^T \mathbb{E}_{A \sim G_t(\cdot|a_t)} \left[\sum_{a \in A} (\hat{f}_t(x_t, a) - f^*(x_t, a))^2\right]\right] \\
 &= \mathbb{E}\left[\sum_{t=1}^T \mathbb{E}_{A \sim G_t(\cdot|a_t)} \left[\sum_{a \in A} (\hat{f}_t(x_t, a) - \ell_{t,a})^2 - \sum_{a \in A} (f^*(x_t, a) - \ell_{t,a})^2\right]\right] \leq \mathbf{Reg}_{\text{Sq}}, \quad (9)
 \end{aligned}$$

432 where the final inequality is due to [Assumption 2](#).

433 Therefore, we have

$$\mathbb{E}[\mathbf{Reg}_{\text{CB}}] \leq CT\gamma^{-\beta} + \frac{\gamma}{4} \mathbf{Reg}_{\text{Sq}}.$$

434 Picking $\gamma = \max\left\{4, \left(\frac{CT}{\mathbf{Reg}_{\text{Sq}}}\right)^{\frac{1}{\beta+1}}\right\}$, we obtain that

$$\mathbb{E}[\mathbf{Reg}_{\text{CB}}] \leq \mathcal{O}\left(C^{\frac{1}{\beta+1}} T^{\frac{1}{\beta+1}} \mathbf{Reg}_{\text{Sq}}^{\frac{\beta}{\beta+1}}\right).$$

435

□

436 **A.1 Proof of Theorem 3.2**

437 Before proving Theorem 3.2, we first show the following key lemma, which is useful in proving
 438 that $\overline{\text{dec}}_\gamma(p; \hat{f}, x, G)$ is convex for both strongly and weakly observable feedback graphs G . We
 439 highlight that the convexity of $\overline{\text{dec}}_\gamma(p; \hat{f}, x, G)$ is crucial for both proving the upper bound of
 440 $\min_{p \in \Delta(K)} \overline{\text{dec}}_\gamma(p; \hat{f}, x, G)$ and showing the efficiency of Algorithm 1.

441 **Lemma A.1.** Suppose $u, v, x \in \mathbb{R}^d$ with $\langle v, x \rangle > 0$. Then both $g(x) = \frac{\langle u, x \rangle^2}{\langle v, x \rangle}$ and $h(x) = \frac{(1 - \langle u, x \rangle)^2}{\langle v, x \rangle}$
 442 are convex in x .

Proof. The function $f(x, y) = x^2/y$ is convex for $y > 0$ due to

$$\nabla^2 f(x, y) = \frac{2}{y^3} \begin{bmatrix} y \\ -x \end{bmatrix} \begin{bmatrix} y \\ -x \end{bmatrix}^\top \succeq 0.$$

443 By composition with affine functions, both $g(x) = f(\langle u, x \rangle, \langle v, x \rangle)$ and $h(x) = f(1 - \langle u, x \rangle, \langle v, x \rangle)$
 444 are convex. \square

445 **Theorem 3.2** (Strongly observable graphs). Suppose that the feedback graph G_t is deterministic and
 446 strongly observable with independence number no more than α . Then Algorithm 1 guarantees that
 447 $\overline{\text{dec}}_\gamma(p_t; \hat{f}_t, x_t, G_t) \leq \mathcal{O}\left(\frac{\alpha \log(K\gamma)}{\gamma}\right)$.

448 *Proof.* For conciseness, we omit the subscript t . Direct calculation shows that for all $a^* \in [K]$,

$$\begin{aligned} & \mathbb{E}_{a \sim p} \left[f^*(x, a) - f^*(x, a^*) - \frac{\gamma}{4} \sum_{a' \in N^{\text{in}}(G, a)} (\hat{f}(x, a') - f^*(x, a'))^2 \right] \\ &= \sum_{a=1}^K p_a f^*(x, a) - f^*(x, a^*) - \frac{\gamma}{4} \sum_{a=1}^K W_a (\hat{f}(x, a) - f^*(x, a))^2, \end{aligned}$$

449 where $W_a = \sum_{a' \in N^{\text{in}}(G, a)} p_{a'}$. Therefore, taking the gradient over $f^*(x, \cdot)$ and we know that

$$\begin{aligned} & \sup_{f^* \in (\mathcal{X} \times [K] \mapsto \mathbb{R})} \left[\sum_{a=1}^K p_a f^*(x, a) - f^*(x, a^*) - \frac{\gamma}{4} \sum_{a=1}^K W_a (\hat{f}(x, a) - f^*(x, a))^2 \right] \\ &= \sum_{a=1}^K p_a \hat{f}(x, a) - \hat{f}(x, a^*) + \frac{1}{\gamma} \|p - e_{a^*}\|_{\text{diag}(W)^{-1}}^2. \end{aligned}$$

450 Then, denote $\hat{f} \in \mathbb{R}^K$ to be $\hat{f}(x, \cdot)$ and consider the following minimax form:

$$\begin{aligned} & \inf_{p \in \Delta(K)} \sup_{a^* \in \mathcal{A}} \left\{ \sum_{a=1}^K p_a \hat{f}(x, a) - \hat{f}(x, a^*) + \frac{1}{\gamma} \|p - e_{a^*}\|_{\text{diag}(W)^{-1}}^2 \right\} \\ &= \min_{p \in \Delta(K)} \max_{a^* \in \mathcal{A}} \left\{ \sum_{a=1}^K p_a \hat{f}(x, a) - \hat{f}(x, a^*) + \frac{1}{\gamma} \sum_{a \neq a^*} \frac{p_a^2}{W_a} + \frac{1}{\gamma} \frac{(1 - p_{a^*})^2}{W_{a^*}} \right\} \quad (10) \end{aligned}$$

$$= \min_{p \in \Delta_K} \max_{q \in \Delta_K} \left\{ \sum_{a=1}^K (p_a - q_a) \hat{f}_a + \frac{1}{\gamma} \sum_{a=1}^K \frac{p_a^2 (1 - q_a)}{W_a} + \sum_{a=1}^K \frac{q_a (1 - p_a)^2}{\gamma W_a} \right\} \quad (11)$$

$$= \max_{q \in \Delta_K} \min_{p \in \Delta_K} \left\{ \sum_{a=1}^K (p_a - q_a) \hat{f}_a + \frac{1}{\gamma} \sum_{a=1}^K \frac{p_a^2 (1 - q_a)}{W_a} + \sum_{a=1}^K \frac{q_a (1 - p_a)^2}{\gamma W_a} \right\}, \quad (12)$$

451 where the last equality is due to Sion's minimax theorem and the fact that Eq. (10) is convex in
 452 $p \in \Delta(K)$ by applying Lemma A.1 with $u = e_a$ and $v = g_a$ for each $a \in [K]$, where $g_a \in \{0, 1\}^K$
 453 is defined as $g_{a,i} = \mathbb{1}\{(i, a) \in E\}$, $G = ([K], E)$, $\forall i \in [K]$.

- 454 Choose $p_a = (1 - \frac{1}{\gamma})q_a + \frac{1}{\gamma K}$ for all $a \in [K]$. Let S be the set of nodes in $[K]$ that have a self-loop.
 455 Then we can upper bound the value above as follows:

$$\begin{aligned}
& \max_{q \in \Delta(K)} \min_{p \in \Delta(K)} \left\{ \sum_{a=1}^K (p_a - q_a) \hat{f}_a + \frac{1}{\gamma} \sum_{a=1}^K \frac{p_a^2 (1 - q_a)}{W_a} + \sum_{a=1}^K \frac{q_a (1 - p_a)^2}{\gamma W_a} \right\} \\
& \leq \max_{q \in \Delta(K)} \left\{ \frac{2}{\gamma} + \frac{1}{\gamma} \sum_{a=1}^K \frac{\left((1 - \frac{1}{\gamma})q_a + \frac{1}{\gamma K} \right)^2 (1 - q_a) + q_a \left(1 - (1 - \frac{1}{\gamma})q_a - \frac{1}{\gamma K} \right)^2}{W_a} \right\} \\
& \leq \max_{q \in \Delta(K)} \left\{ \frac{2}{\gamma} + \frac{1}{\gamma} \sum_{a=1}^K \frac{2 \left((1 - \frac{1}{\gamma})^2 q_a^2 + \frac{1}{\gamma^2 K^2} \right) (1 - q_a) + q_a \left(1 - (1 - \frac{1}{\gamma})q_a \right)^2}{W_a} \right\} \\
& \leq \max_{q \in \Delta(K)} \left\{ \frac{2}{\gamma} + \frac{2}{\gamma^2} + \frac{1}{\gamma} \sum_{a=1}^K \frac{2q_a^2 (1 - q_a) + 2q_a (1 - q_a)^2 + \frac{2q_a^3}{\gamma^2}}{W_a} \right\} \\
& \hspace{15em} (W_a = \sum_{j \in N^{\text{in}}(G,a)} p_j \geq \frac{1}{\gamma K} \text{ for all } a \in [K]) \\
& \leq \max_{q \in \Delta(K)} \left\{ \frac{2}{\gamma} + \frac{2}{\gamma^2} + \frac{2}{\gamma} \sum_{a=1}^K \frac{q_a (1 - q_a)}{W_a} + \frac{2}{\gamma^3} \sum_{a=1}^K \frac{q_a^3}{W_a} \right\} \\
& = \max_{q \in \Delta(K)} \left\{ \frac{2}{\gamma} + \frac{2}{\gamma^2} + \frac{2}{\gamma} \sum_{a=1}^K \frac{q_a (1 - q_a)}{W_a} + \frac{2}{\gamma^3} \sum_{a \in S} \frac{q_a^3}{W_a} + \frac{2}{\gamma^3} \sum_{a \notin S} \frac{q_a^3}{W_a} \right\} \tag{13} \\
& \leq \max_{q \in \Delta(K)} \left\{ \frac{2}{\gamma} + \frac{2}{\gamma^2} + \frac{2}{\gamma} \sum_{a=1}^K \frac{q_a (1 - q_a)}{W_a} + \frac{2}{\gamma^3} \sum_{a \in S} q_a^2 + \frac{2}{\gamma^3} \sum_{a \notin S} \frac{q_a^3}{\frac{K-1}{\gamma K}} \right\} \\
& \hspace{15em} (\text{if } a \notin S, W_a = 1 - p_a \geq \frac{K-1}{\gamma K}) \\
& \leq \max_{q \in \Delta(K)} \left\{ \frac{8}{\gamma} + \frac{2}{\gamma} \sum_{a=1}^K \frac{q_a (1 - q_a)}{W_a} \right\}. \tag{K \geq 2}
\end{aligned}$$

- 456 Next we bound $\frac{2q_a(1-q_a)}{W_a}$ for each $a \in [K]$. If $a \in [K] \setminus S$, we have $W_a = 1 - p_a$ and

$$\frac{2q_a(1-q_a)}{W_a} \leq \frac{2q_a(1-q_a)}{1 - (1 - \frac{1}{\gamma})q_a - \frac{1}{\gamma K}} \leq \frac{2q_a(1-q_a)}{(1 - \frac{1}{\gamma})(1 - q_a) + \frac{K-1}{\gamma K}} \leq \frac{2}{1 - \frac{1}{\gamma}} q_a \leq 4q_a. \tag{14}$$

- 457 If $a \in S$, we know that

$$\begin{aligned}
\sum_{a \in S} \frac{2q_a(1-q_a)}{W_a} & \leq \sum_{a \in S} \frac{2q_a(1-q_a)}{\sum_{j \in N^{\text{in}}(G,a)} \left((1 - \frac{1}{\gamma})q_j + \frac{1}{\gamma K} \right)} \\
& \leq \frac{\gamma}{\gamma - 1} \sum_{a \in S} \frac{2 \left((1 - \frac{1}{\gamma})q_a + \frac{1}{\gamma K} \right) (1 - q_a)}{\sum_{j \in N^{\text{in}}(G,a)} \left((1 - \frac{1}{\gamma})q_j + \frac{1}{\gamma K} \right)} \\
& \leq 4 \sum_{a \in S} \frac{\left((1 - \frac{1}{\gamma})q_a + \frac{1}{\gamma K} \right)}{\sum_{j \in N^{\text{in}}(G,a)} \left((1 - \frac{1}{\gamma})q_j + \frac{1}{\gamma K} \right)} \leq \mathcal{O}(\alpha \log(K\gamma)), \tag{15}
\end{aligned}$$

- 458 where the last inequality is due to Lemma 5 in Alon et al. [2015]. We include this lemma (Lemma D.1)
 459 for completeness. Combining all the above inequalities, we obtain that

$$\begin{aligned}
& \inf_{p \in \Delta(K)} \sup_{a^* \in \mathcal{A}} \left\{ \sum_{a=1}^K p_a \hat{f}(x, a) - \hat{f}(x, a^*) + \frac{1}{\gamma} \|p - e_{a^*}\|_{\text{diag}(W)^{-1}}^2 \right\} \\
& = \max_{q \in \Delta(K)} \min_{p \in \Delta(K)} \left\{ \sum_{a=1}^K (p_a - q_a) \hat{f}_a + \frac{1}{\gamma} \sum_{a=1}^K \frac{p_a^2 (1 - q_a)}{W_a} + \sum_{a=1}^K \frac{q_a (1 - p_a)^2}{\gamma W_a} \right\}
\end{aligned}$$

$$\leq \max_{q \in \Delta(K)} \left\{ \frac{8}{\gamma} + \frac{2}{\gamma} \sum_{a=1}^K \frac{q_a(1-q_a)}{W_a} \right\} \leq \mathcal{O} \left(\frac{\alpha \log(K\gamma)}{\gamma} \right).$$

460

□

461 A.2 Proof of Theorem 3.4

462 **Theorem 3.4** (Weakly observable graphs). *Suppose that the feedback graph G_t is deterministic and*
 463 *weakly observable with weak domination number no more than d . Then Algorithm 1 with $\gamma \geq 16d$*
 464 *guarantees that $\overline{\text{dec}}_\gamma(p_t; \hat{f}_t, x_t, G_t) \leq \mathcal{O} \left(\sqrt{\frac{d}{\gamma}} + \frac{\tilde{\alpha} \log(K\gamma)}{\gamma} \right)$, where $\tilde{\alpha}$ is the independence number*
 465 *of the subgraph induced by nodes with self-loops in G_t .*

466 *Proof.* Similar to the strongly observable graphs setting, for weakly observable graphs, we know that

$$\begin{aligned} & \overline{\text{dec}}_\gamma(p; \hat{f}, x, G) \\ &= \max_{q \in \Delta_K} \min_{p \in \Delta_K} \left\{ \sum_{a=1}^K (p_a - q_a) \hat{f}_a + \frac{1}{\gamma} \sum_{a=1}^K \frac{p_a^2(1-q_a)}{W_a} + \sum_{a=1}^K \frac{q_a(1-p_a)^2}{\gamma W_a} \right\}. \end{aligned} \quad (16)$$

467 Choose $p_a = (1 - \frac{1}{\gamma} - \eta d)q_a + \frac{1}{\gamma K} + \eta \mathbb{1}\{a \in D\}$ where D with $|D| = d$ is the minimum weak
 468 dominating set of G and $0 < \eta \leq \frac{1}{4d}$ is some parameter to be chosen later. Substituting the form of p
 469 to Eq. (16) and using the fact that $|\hat{f}_a| \leq 1$ for all $a \in [K]$, we can obtain that

$$\begin{aligned} & \overline{\text{dec}}_\gamma(p; \hat{f}, x, G) \\ & \leq \max_{q \in \Delta_K} \left\{ \frac{2}{\gamma} + \eta d + \frac{1}{\gamma} \sum_{a=1}^K \frac{p_a^2(1-q_a)}{W_a} + \sum_{a=1}^K \frac{q_a(1-p_a)^2}{\gamma W_a} \right\}. \end{aligned}$$

470 Then we can upper bound the value above as follows:

$$\begin{aligned} & \overline{\text{dec}}_\gamma(p; \hat{f}, x, G) \\ & \leq \max_{q \in \Delta_K} \left\{ \frac{2}{\gamma} + \eta d + \frac{1}{\gamma} \sum_{a=1}^K \frac{\left((1 - \frac{1}{\gamma} - \eta d)q_a + \frac{1}{\gamma K} + \eta \mathbb{1}\{a \in D\} \right)^2 (1-q_a)}{W_a} \right. \\ & \quad \left. + \sum_{a=1}^K \frac{q_a \left(1 - (1 - \frac{1}{\gamma} - \eta d)q_a \right)^2}{W_a} \right\} \\ & \leq \max_{q \in \Delta_K} \left\{ \frac{2}{\gamma} + \eta d + \frac{1}{\gamma} \sum_{a \notin D} \frac{\left(q_a + \frac{1}{\gamma K} \right)^2 (1-q_a) + q_a \left((1-q_a) + \frac{1}{\gamma} q_a + \eta d q_a \right)^2}{W_a} \right. \\ & \quad \left. + \frac{1}{\gamma} \sum_{a \in D} \frac{\left(q_a + \frac{1}{\gamma K} + \eta \right)^2 (1-q_a) + q_a \left((1-q_a) + \frac{1}{\gamma} q_a + \eta d q_a \right)^2}{W_a} \right\} \\ & \leq \max_{q \in \Delta_K} \left\{ \frac{2}{\gamma} + \eta d + \frac{1}{\gamma} \sum_{a \notin D} \frac{2 \left(q_a^2 + \frac{1}{\gamma^2 K^2} \right) (1-q_a) + 3q_a \left((1-q_a)^2 + \frac{q_a^2}{\gamma^2} + \eta^2 d^2 q_a^2 \right)}{W_a} \right. \\ & \quad \left. + \frac{1}{\gamma} \sum_{a \in D} \frac{3 \left(q_a^2 + \frac{1}{\gamma^2 K^2} + \eta^2 \right) (1-q_a) + 3q_a \left((1-q_a)^2 + \frac{q_a^2}{\gamma^2} + \eta^2 d^2 q_a^2 \right)}{W_a} \right\}. \end{aligned} \quad (17)$$

471 Now consider $a \in D$. If $a \in S$, then we know that $W_a \geq \eta$; Otherwise, we know that this node can
 472 be observed by at least one node in D , meaning that $W_a \geq \eta$. Combining the two cases above, we

473 know that

$$\begin{aligned}
& \frac{1}{\gamma} \sum_{a \in D} \frac{3 \left(q_a^2 + \frac{1}{\gamma^2 K^2} + \eta^2 \right) (1 - q_a) + 3q_a \left((1 - q_a)^2 + \frac{1}{\gamma^2} q_a^2 + \eta^2 d^2 q_a^2 \right)}{W_a} \\
& \leq \frac{3}{\eta\gamma} \sum_{a \in D} \left[\left(q_a^2 + \frac{1}{\gamma^2 K^2} + \eta^2 \right) (1 - q_a) + q_a \left((1 - q_a)^2 + \frac{1}{\gamma^2} q_a^2 + \eta^2 d^2 q_a^2 \right) \right] \\
& \leq \frac{3}{\eta\gamma} \sum_{a \in D} \left[q_a - q_a^2 + \frac{1}{\gamma^2} q_a^3 + \eta^2 d^2 q_a^3 + \frac{1}{\gamma^2 K^2} + \eta^2 \right] \\
& \leq \mathcal{O} \left(\frac{1}{\eta\gamma} + \frac{d\eta}{\gamma} + \frac{1}{\eta\gamma^3 K} \right) \quad (\eta \leq \frac{1}{4d} \text{ and } \gamma \geq 16d) \\
& \leq \mathcal{O} \left(\frac{1}{\eta\gamma} \right), \tag{18}
\end{aligned}$$

474 where the last inequality is because $\eta \leq \frac{1}{4d}$ and $\gamma \geq 16d$. Consider $a \notin D$. Let S_0 be the set of
475 nodes which either have a self loop or can be observed by all the other node. Recall that S represents
476 the set of nodes with a self-loop. Then similar to the derivation of Eq. (13), we know that for $a \in S_0$,

$$\begin{aligned}
& \frac{1}{\gamma} \sum_{a \notin D, a \in S_0} \frac{2 \left(q_a^2 + \frac{1}{\gamma^2 K^2} \right) (1 - q_a) + 3q_a \left((1 - q_a)^2 + \frac{q_a^2}{\gamma^2} + \eta^2 d^2 q_a^2 \right)}{W_a} \\
& \leq \frac{1}{\gamma} \sum_{a \notin D, a \in S_0} \frac{2q_a^2(1 - q_a) + 3q_a \left((1 - q_a)^2 + \frac{q_a^2}{\gamma^2} + \eta^2 d^2 q_a^2 \right)}{W_a} + \mathcal{O} \left(\frac{1}{\gamma^2} + \frac{1}{\eta\gamma^3 K} \right) \\
& \quad (W_a \geq \frac{1}{\gamma K} \text{ if } a \in S \text{ and } W_a \geq \eta \text{ if } a \in [K] \setminus S) \\
& \leq \mathcal{O} \left(\frac{1}{\gamma} \sum_{a \in S_0, a \notin D} \frac{q_a(1 - q_a)}{W_a} + \frac{1}{\gamma^3} \sum_{a \in S, a \notin D} q_a^2 + \frac{1}{\gamma^3} \sum_{a \in S_0, a \notin D \cup S} \frac{q_a^3}{\gamma K} + \frac{1}{\gamma^2} + \frac{1}{\eta\gamma^3 K} \right) \\
& \quad + \mathcal{O} \left(\frac{1}{\gamma} \sum_{a \in S, a \notin D} \eta^2 d^2 q_a^2 + \frac{1}{\gamma} \sum_{a \in S_0, a \notin D \cup S} \frac{\eta^2 d^2 q_a^3}{\eta} \right) \\
& \quad (\text{for } a \in S_0, a \notin S, W_a \geq \max\{\frac{K-1}{\gamma K}, \eta\}) \\
& \leq \mathcal{O} \left(\frac{1}{\gamma} \sum_{a \in S_0, a \notin D} \frac{q_a(1 - q_a)}{W_a} + \frac{1}{\eta\gamma} \right). \tag{19}
\end{aligned}$$

477 For $a \notin S_0$, we know that $W_a \geq \eta$. Therefore,

$$\begin{aligned}
& \frac{1}{\gamma} \sum_{a \notin D \cup S_0} \frac{2 \left(q_a^2 + \frac{1}{\gamma^2 K^2} \right) (1 - q_a) + 3q_a \left((1 - q_a)^2 + \frac{q_a^2}{\gamma^2} + \eta^2 d^2 q_a^2 \right)}{W_a} \\
& \leq \frac{1}{\gamma\eta} \sum_{a \notin D \cup S_0} \left(2 \left(q_a^2 + \frac{1}{\gamma^2 K^2} \right) (1 - q_a) + 3q_a \left((1 - q_a)^2 + \frac{q_a^2}{\gamma^2} + \frac{1}{16} q_a^2 \right) \right) \\
& \leq \frac{1}{\gamma\eta} \sum_{a \notin D \cup S_0} \left(2q_a(1 - q_a) + \frac{1}{\gamma^2 K^2} + \frac{2q_a^3}{\gamma^2} + \frac{3}{16} q_a^3 \right) \\
& \leq \mathcal{O} \left(\frac{1}{\gamma\eta} \right). \tag{20}
\end{aligned}$$

478 Plugging Eq. (18), Eq. (19), and Eq. (20) to Eq. (17), we obtain that

$$\overline{\text{dec}}_\gamma(p; \hat{f}, x, G) \leq \mathcal{O} \left(\frac{1}{\gamma} + \eta d + \frac{1}{\eta\gamma} + \frac{1}{\gamma} \sum_{a \in S_0, a \notin D} \frac{q_a(1 - q_a)}{W_a} \right) \tag{21}$$

479 Consider the last term. If $a \in S_0 \setminus S$, similar to Eq. (14), we know that

$$\frac{q_a(1-q_a)}{W_a} \leq \frac{q_a(1-q_a)}{1 - (1 - \frac{1}{\gamma} - d\eta)q_a - \frac{1}{\gamma K}} \leq \frac{q_a(1-q_a)}{(1 - \frac{1}{\gamma} - \eta d)(1-q_a)} \leq \frac{1}{1 - \frac{1}{\gamma} - \eta d} q_a \leq \mathcal{O}(q_a),$$

480 where the last inequality is due to $\gamma \geq 16d$ and $\eta \leq \frac{1}{4d}$. If $a \in S$, similar to Eq. (15), we know that

$$\begin{aligned} \sum_{a \in S} \frac{q_a(1-q_a)}{W_a} &\leq \sum_{a \in S} \frac{q_a(1-q_a)}{\sum_{j \in N^{\text{in}}(G,a)} ((1 - \frac{1}{\gamma} - \eta d)q_j + \frac{1}{\gamma K})} \\ &\leq \frac{\gamma}{\gamma - 1 - \gamma \eta d} \sum_{a \in S} \frac{((1 - \frac{1}{\gamma} - \eta d)q_a + \frac{1}{\gamma K})(1-q_a)}{\sum_{j \in N^{\text{in}}(G,a)} ((1 - \frac{1}{\gamma} - \eta d)q_j + \frac{1}{\gamma K})} \\ &\leq 2 \sum_{a \in S} \frac{\left((1 - \frac{1}{\gamma} - \eta d)q_a + \frac{1}{\gamma K} \right)}{\sum_{j \in N^{\text{in}}(G,a)} \left((1 - \frac{1}{\gamma} - \eta d)q_j + \frac{1}{\gamma K} \right)} \quad (\gamma \geq 4, \eta \leq \frac{1}{4d}) \\ &\leq \mathcal{O}(\tilde{\alpha} \log(K\gamma)), \end{aligned} \quad (22)$$

481 where the last inequality is again due to Lemma 5 in [Alon et al., 2015] and $\tilde{\alpha}$ is the independence
482 number of the subgraph induced by nodes with self-loops in G . Plugging Eq. (22) to Eq. (21) gives

$$\overline{\text{dec}}_\gamma(p; \hat{f}, x, G) \leq \mathcal{O} \left(\eta d + \frac{1}{\gamma \eta} + \frac{\tilde{\alpha} \log(K\gamma)}{\gamma} \right).$$

483 Picking $\eta = \sqrt{\frac{1}{\gamma d}} \leq \frac{1}{4d}$ proves the result. \square

484 Next, we prove Corollary 3.5 by combining Theorem 3.4 and Theorem 3.1.

485 **Corollary 3.5.** *Suppose that G_t is deterministic, weakly observable, and has weak domination*
486 *number no more than d for all $t \in [T]$. In addition, suppose that the independence number of the*
487 *subgraph induced by nodes with self-loops in G_t is no more than $\tilde{\alpha}$ for all $t \in [T]$. Then, Algorithm 1*
488 *with $\gamma = \max\{16d, \sqrt{\tilde{\alpha}T/\mathbf{Reg}_{\text{Sq}}}, d^{\frac{1}{3}}T^{\frac{2}{3}}\mathbf{Reg}_{\text{Sq}}^{-\frac{2}{3}}\}$ guarantees that $\mathbb{E}[\mathbf{Reg}_{\text{CB}}] \leq \tilde{\mathcal{O}}(d^{\frac{1}{3}}T^{\frac{2}{3}}\mathbf{Reg}_{\text{Sq}}^{\frac{1}{3}} +$
489 $\sqrt{\tilde{\alpha}T\mathbf{Reg}_{\text{Sq}}})$.*

490 *Proof.* Combining Eq. (8), Eq. (9) and Theorem 3.4, we can bound \mathbf{Reg}_{CB} as follows:

$$\mathbb{E}[\mathbf{Reg}_{\text{CB}}] \leq \mathcal{O} \left(\sqrt{\frac{d}{\gamma}T} + \frac{\tilde{\alpha}T \log(K\gamma)}{\gamma} + \gamma \mathbf{Reg}_{\text{CB}} \right).$$

491 Picking $\gamma = \max\{16d, \sqrt{\tilde{\alpha}T/\mathbf{Reg}_{\text{Sq}}}, d^{\frac{1}{3}}T^{\frac{2}{3}}\mathbf{Reg}_{\text{Sq}}^{-\frac{2}{3}}\}$ finishes the proof. \square

492 A.3 Python Solution to Eq. (5)

```

493 def makeProblem(nactions):
494     import cvxpy as cp
495
496     sqrtgammaG = cp.Parameter((nactions, nactions), nonneg=True)
497     sqrtgammafhat = cp.Parameter(nactions)
498     p = cp.Variable(nactions, nonneg=True)
499     sqrtgammaz = cp.Variable()
500     objective = cp.Minimize(sqrtgammafhat @ p + sqrtgammaz)
501     constraints = [
502         cp.sum(p) == 1
503     ] + [
504         cp.sum([ cp.quad_over_lin(eai - pi, vi)
505                 for i, (pi, vi) in enumerate(zip(p, v))
506                 for eai in (1 if i == a else 0,)
507                 ]) <= sqrtgammafhat + sqrtgammaz
508         for v in (sqrtgammaG @ p,)

```

```

510     for a, sqrtgammahata in enumerate(sqrtgammahat)
511     ]
512     problem = cp.Problem(objective, constraints)
513     assert problem.is_dcp(dpp=True) # proof of convexity
514     return problem, sqrtgammaG, sqrtgammahat, p, sqrtgammaz

```

516 This particular formulation multiplies both sides of the constraint in Eq. (5) by $\sqrt{\gamma}$ while scaling the
517 objective by $\sqrt{\gamma}$. While mathematically equivalent to Eq. (5), empirically it has superior numerical
518 stability for large γ . For additional stability, when using this routine we recommend subtracting off the
519 minimum value from \hat{f} , which is equivalent to making the substitutions $\sqrt{\gamma}\hat{f} \leftarrow \sqrt{\gamma}\hat{f} - \sqrt{\gamma}\min_a \hat{f}_a$
520 and $z \leftarrow z + \sqrt{\gamma}\min_a \hat{f}_a$ and then exploiting the $1^\top p = 1$ constraint.

521 A.4 Proof of Theorem 3.6

522 **Theorem 3.6.** Solving $\operatorname{argmin}_{p \in \Delta(K)} \overline{\operatorname{dec}}_\gamma(p; \hat{f}, x, G)$ is equivalent to solving the following convex
523 optimization problem.

$$\begin{aligned}
& \min_{p \in \Delta(K), z} && p^\top \hat{f} + z && (5) \\
& \text{subject to} && \forall a \in [K] : \frac{1}{\gamma} \|p - e_a\|_{\operatorname{diag}(G^\top p)^{-1}}^2 \leq \hat{f}(x, a) + z, \\
& && G^\top p \succ 0,
\end{aligned}$$

524 where \hat{f} in the objective is a shorthand for $\hat{f}(x, \cdot) \in \mathbb{R}^K$, e_a is the a -th standard basis vector, and \succ
525 means element-wise greater.

526 *Proof.* Denote $f^* = f^*(x, \cdot) \in \mathbb{R}^K$. Note that according to the definition of G , we know that
527 $(G^\top p)_i$ denotes the probability that action i 's loss is revealed when the selected action a is sampled
528 from distribution p . Then, we know that

$$\begin{aligned}
& \overline{\operatorname{dec}}_\gamma(p; \hat{f}, x, G) \\
&= \sup_{\substack{a^* \in [K] \\ f^* \in \mathbb{R}^K}} \mathbb{E}_{a_t \sim p} \left[f_{a_t}^* - f_{a^*}^* - \frac{\gamma}{4} \mathbb{E}_{A \sim G(\cdot | a_t)} \left[\sum_{a \in A} (\hat{f}_a - f_a^*)^2 \right] \right] \\
&= \sup_{\substack{a^* \in [K] \\ f^* \in \mathbb{R}^K}} (p - e_{a^*})^\top f^* - \frac{\gamma}{4} \sum_{a \in [K]} \|\hat{f} - f^*\|_{\operatorname{diag}(G^\top p)}^2 \\
&= \sup_{a^* \in [K]} (p - e_{a^*})^\top \hat{f} + \frac{1}{\gamma} \|p - e_{a^*}\|_{\operatorname{diag}(G^\top p)^{-1}}^2 && (G^\top p \succ 0) \\
&= p^\top \hat{f} + \max_{a^* \in [K]} \left\{ \frac{1}{\gamma} \|p - e_{a^*}\|_{\operatorname{diag}(G^\top p)^{-1}}^2 - e_{a^*}^\top \hat{f} \right\},
\end{aligned}$$

529 where the third equality is by picking f^* to be the maximizer and introduces a constraint. Therefore,
530 the minimization problem $\min_{p \in \Delta(K)} \overline{\operatorname{dec}}_\gamma(p; \hat{f}, x, G)$ can be written as the following constrained
531 optimization by variable substitution:

$$\begin{aligned}
& \min_{p \in \Delta(K), z} && p^\top \hat{f} + z \\
& \text{subject to} && \forall a \in [K] : \frac{1}{\gamma} \|p - e_a\|_{\operatorname{diag}(G^\top p)^{-1}}^2 \leq e_a^\top \hat{f} + z, \\
& && G^\top p \succ 0.
\end{aligned}$$

532 The convexity of the constraints follows from Lemma A.1. \square

533 B Omitted Details in Section 4

534 In this section, we provide proofs for Section 4. We define $W_a := \sum_{a' \in N^{\operatorname{in}}(G, a)} p_{a'}$ to be the
535 probability that the loss of action a is revealed when selecting an action from distribution p . Let

536 $\widehat{f} = \widehat{f}(x, \cdot) \in \mathbb{R}^K$ and $f = f(x, \cdot) \in \mathbb{R}^K$. Direct calculation shows that for any $a^* \in [K]$,

$$\begin{aligned} f^* &= \operatorname{argmax}_{f \in \mathbb{R}^K} \mathbb{E}_{a \sim p} \left[f(x, a) - f(x, a^*) - \frac{\gamma}{4} \cdot \sum_{a' \in N^{\text{in}}(G, a)} (\widehat{f}_t(x, a') - f(x, a'))^2 \right] \\ &= \frac{2}{\gamma} \operatorname{diag}(W)^{-1} (p - e_{a^*}) + \widehat{f}. \end{aligned}$$

537 Therefore, substituting f^* into Eq. (4), we obtain that

$$\overline{\operatorname{dec}}_\gamma(p; \widehat{f}, x, G) = \max_{a^* \in [K]} \left\{ \frac{1}{\gamma} \left(\sum_{a=1}^K \frac{p_a^2}{W_a} + \frac{1 - 2p_{a^*}}{W_{a^*}} \right) + \langle p - e_{a^*}, \widehat{f} \rangle \right\}. \quad (23)$$

538 Without loss of generality, we assume the $\min_{i \in [K]} \widehat{f}_i = 0$. This is because shifting \widehat{f} by $\min_{i \in [K]} \widehat{f}_i$
 539 does not change the value of $\langle p - e_{a^*}, \widehat{f} \rangle$. In the following sections, we provide proofs showing that
 540 a certain closed-form of p leads to optimal $\overline{\operatorname{dec}}_\gamma(p; \widehat{f}, x, G)$ up to constant factors for several specific
 541 types of feedback graphs, respectively.

542 B.1 Cops-and-Robbers Graph

543 **Proposition 1.** When $G = G_{\text{CR}}$, given any \widehat{f} , context x , the closed-form distribution p in Eq. (6)
 544 guarantees that $\overline{\operatorname{dec}}_\gamma(p; \widehat{f}, x, G_{\text{CR}}) \leq \mathcal{O}\left(\frac{1}{\gamma}\right)$.

545 *Proof.* We use the following notation for convenience: $p_1 := p_{a_1}$, $p_2 := p_{a_2}$, $\widehat{f}_1 := \widehat{f}_{a_1} = 0$,
 546 $\widehat{f}_2 := \widehat{f}_{a_2}$. For the cops-and-robbers graph and closed-form solution p in Eq. (6), Eq. (23) becomes:

$$\overline{\operatorname{dec}}_\gamma(p; \widehat{f}, x, G_{\text{CR}}) = \max_{a^* \in [K]} \left\{ \frac{1}{\gamma} \left(\frac{p_1^2}{1 - p_1} + \frac{(1 - p_1)^2}{p_1} + \frac{1 - 2p_{a^*}}{W_{a^*}} \right) + \langle p - e_{a^*}, \widehat{f} \rangle \right\}.$$

547 If $a^* \neq a_1$ and $a^* \neq a_2$, we know that

$$\begin{aligned} & \frac{1}{\gamma} \left(\frac{p_1^2}{1 - p_1} + \frac{(1 - p_1)^2}{p_1} + \frac{1 - 2p_{a^*}}{W_{a^*}} \right) + \langle p - e_{a^*}, \widehat{f} \rangle \\ &= \frac{1}{\gamma} \left(\frac{p_1^2}{1 - p_1} + \frac{(1 - p_1)^2}{p_1} + 1 \right) + p_1 \widehat{f}_1 + p_2 \widehat{f}_2 - \widehat{f}_{a^*} \\ &\leq \frac{1}{\gamma} \left(\frac{p_1^2}{1 - p_1} + \frac{(1 - p_1)^2}{p_1} + 1 \right) - p_1 \widehat{f}_2 \quad (\widehat{f}_{a^*} \geq \widehat{f}_2 \geq \widehat{f}_1 = 0) \\ &\leq \frac{1}{\gamma} \left(\frac{1}{1 - p_1} + 1 + 1 \right) - p_1 \widehat{f}_2 \quad (p_1 \in [\frac{1}{2}, 1], p_1 \geq p_2 \in [0, \frac{1}{2}]) \\ &= \frac{1}{\gamma} \left(4 + \gamma \widehat{f}_2 \right) - \left(1 - \frac{1}{2 + \gamma \widehat{f}_2} \right) \widehat{f}_2 \\ &\leq \frac{5}{\gamma}. \end{aligned}$$

548 If $a^* = a_2$, we can obtain that

$$\begin{aligned} & \frac{1}{\gamma} \left(\frac{p_1^2}{1 - p_1} + \frac{(1 - p_1)^2}{p_1} + \frac{1 - 2p_{a^*}}{W_{a^*}} \right) + \langle p - e_{a^*}, \widehat{f} \rangle \\ &= \frac{1}{\gamma} \left(\frac{p_1^2}{1 - p_1} + \frac{(1 - p_1)^2}{p_1} + \frac{1 - 2p_2}{p_1} \right) + p_1 \widehat{f}_1 + p_2 \widehat{f}_2 - \widehat{f}_2 \\ &\leq \frac{1}{\gamma} \left(\frac{p_1^2}{1 - p_1} + \frac{(1 - p_1)^2}{p_1} + \frac{1 - 2(1 - p_1)}{p_1} \right) - p_1 \widehat{f}_2 \quad (\widehat{f}_1 = 0) \\ &\leq \frac{1}{\gamma} \left(\frac{1}{1 - p_1} + 1 + 2 - \frac{1}{p_1} \right) - p_1 \widehat{f}_2 \quad (p_1 \in [\frac{1}{2}, 1], p_2 \in [0, \frac{1}{2}]) \end{aligned}$$

$$\begin{aligned}
&\leq \frac{1}{\gamma} \left(5 + \gamma \widehat{f}_2 \right) - \left(1 - \frac{1}{2 + \gamma \widehat{f}_2} \right) \widehat{f}_2 && (p_1 = \frac{1}{2 + \gamma \widehat{f}_2}) \\
&\leq \frac{6}{\gamma}.
\end{aligned}$$

549 If $a^* = a_1$, we have

$$\begin{aligned}
&\frac{1}{\gamma} \left(\frac{p_1^2}{1 - p_1} + \frac{(1 - p_1)^2}{p_1} + \frac{1 - 2p_{a^*}}{W_{a^*}} \right) + \langle p - e_{a^*}, \widehat{f} \rangle \\
&\leq \frac{1}{\gamma} \left(\frac{p_1^2}{1 - p_1} + \frac{(1 - p_1)^2}{p_1} + \frac{1 - 2p_1}{1 - p_1} \right) + (1 - p_1) \widehat{f}_2 \\
&\leq \frac{1}{\gamma} \left(1 - p_1 + \frac{(1 - p_1)^2}{p_1} \right) + (1 - p_1) \widehat{f}_2 \\
&\leq \frac{1}{\gamma} \left(1 + \frac{1}{2} \right) + \frac{\widehat{f}_2}{2 + \gamma \widehat{f}_2} && (p_1 \in [\frac{1}{2}, 1]) \\
&\leq \frac{3}{\gamma}.
\end{aligned}$$

550 Putting everything together, we prove that $\overline{\text{dec}}_\gamma(p; \widehat{f}, x, G_{\text{CR}}) \leq \frac{6}{\gamma} \leq \mathcal{O}\left(\frac{1}{\gamma}\right)$. \square

551 B.2 Apple Tasting Graph

552 **Proposition 2.** When $G = G_{\text{AT}}$, given any \widehat{f} , context x , the closed-form distribution p in Eq. (7)
553 guarantees that $\overline{\text{dec}}_\gamma(p; \widehat{f}, x, G_{\text{AT}}) \leq \mathcal{O}\left(\frac{1}{\gamma}\right)$.

554 *Proof.* For the apple tasting graph and closed-form solution p in Eq. (7), Eq. (23) becomes:

$$\overline{\text{dec}}_\gamma(p; \widehat{f}, x, G) = \max_{a^* \in [K]} \left\{ \frac{1}{\gamma} \left(p_1 + \frac{(1 - p_1)^2}{p_1} + \frac{1 - 2p_{a^*}}{W_{a^*}} \right) + \langle p - e_{a^*}, \widehat{f} \rangle \right\}.$$

555 Suppose $\widehat{f}_1 = 0$, we know that $p_1 = 1, p_2 = 0$ and

556 1. If $a^* = 1$, we have

$$\frac{1}{\gamma} \left(p_1 + \frac{(1 - p_1)^2}{p_1} + \frac{1 - 2p_{a^*}}{W_{a^*}} \right) + \langle p - e_{a^*}, \widehat{f} \rangle = 0.$$

557 2. If $a^* = 2$, direct calculation shows that

$$\frac{1}{\gamma} \left(p_1 + \frac{(1 - p_1)^2}{p_1} + \frac{1 - 2p_{a^*}}{W_{a^*}} \right) + \langle p - e_{a^*}, \widehat{f} \rangle \leq \frac{2}{\gamma}.$$

558 Suppose $\widehat{f}_2 = 0$, we know that $p_1 = \frac{2}{4 + \gamma \widehat{f}_1}, p_2 = 1 - p_1$ and

559 1. If $a^* = 1$, we have

$$\begin{aligned}
&\frac{1}{\gamma} \left(p_1 + \frac{(1 - p_1)^2}{p_1} + \frac{1 - 2p_{a^*}}{W_{a^*}} \right) + \langle p - e_{a^*}, \widehat{f} \rangle \\
&= \frac{1}{\gamma} \left(p_1 + \frac{(1 - p_1)^2}{p_1} + \frac{1 - 2p_1}{p_1} \right) - (1 - p_1) \widehat{f}_1 \\
&= \frac{2(1 - p_1)^2}{\gamma p_1} - (1 - p_1) \widehat{f}_1 \\
&= \frac{(2 + \gamma \widehat{f}_1)^2}{\gamma(4 + \gamma \widehat{f}_1)} - (1 - p_1) \widehat{f}_1 \\
&\leq \frac{4 + \gamma \widehat{f}_1}{\gamma} + \frac{2 \widehat{f}_1}{4 + \gamma \widehat{f}_1} - \widehat{f}_1 \leq \frac{6}{\gamma}.
\end{aligned}$$

560 2. If $a^* = 2$, direct calculation shows that

$$\frac{1}{\gamma} \left(p_1 + \frac{(1-p_1)^2}{p_1} + \frac{1-2p_{a^*}}{W_{a^*}} \right) + \langle p - e_{a^*}, \hat{f} \rangle = \frac{2p_1}{\gamma} + p_1 \hat{f}_1 \leq \frac{1}{\gamma} + \frac{2\hat{f}_1}{4+\gamma\hat{f}_1} \leq \frac{3}{\gamma}.$$

561 Putting everything together, we prove that $\overline{\text{dec}}_\gamma(p; \hat{f}, x, G_{\text{AT}}) \leq \frac{6}{\gamma} \leq \mathcal{O}\left(\frac{1}{\gamma}\right)$. \square

562 B.3 Inventory Graph

563 **Proposition 3.** When $G = G_{\text{inv}}$, given any \hat{f} , context x , there exists a closed-form distribution
 564 $p \in \Delta(K)$ guaranteeing that $\overline{\text{dec}}_\gamma(p; \hat{f}, x, G_{\text{inv}}) \leq \mathcal{O}\left(\frac{1}{\gamma}\right)$, where p is defined as follows: $p_j =$
 565 $\max\left\{\frac{1}{1+\gamma(\hat{f}_j - \min_i \hat{f}_i)} - \sum_{j' > j} p_{j'}, 0\right\}$ for all $j \in [K]$.

566 *Proof.* Based on the distribution defined above, define $A \subseteq [K]$ to be the set such that for all $i \in A$,
 567 $p_i > 0$ and denote $N = |A|$. We index each action in A by $k_1 < k_2 < \dots < k_N = K$. According
 568 to the definition of p_i , we know that p_i is strictly positive only when $\hat{f}_i < \hat{f}_j$ for all $j > i$ and
 569 specifically, when $p_i > 0$, we know that $W_i = \sum_{j \geq i} p_j = \frac{1}{1+\gamma\hat{f}_i}$ (recall that $\min_i \hat{f}_i = 0$ since we
 570 shift \hat{f}). Therefore, define $W_{k_{N+1}} = 0$ and we know that

$$\begin{aligned} & \overline{\text{dec}}_\gamma(p; \hat{f}, x, G_{\text{inv}}) \\ &= \sum_{i=1}^N p_{k_i} \hat{f}_{k_i} + \frac{1}{\gamma} \sum_{a=1}^K \frac{p_a^2}{W_a} + \max_{a^* \in [K]} \left\{ \frac{1-2p_{a^*}}{\gamma W_{a^*}} - \hat{f}_{a^*} \right\} \\ &\leq \sum_{i=1}^N (W_{k_i} - W_{k_{i+1}}) \hat{f}_{k_i} + \frac{1}{\gamma} + \max_{a^* \in [K]} \left\{ \frac{1-2p_{a^*}}{\gamma W_{a^*}} - \hat{f}_{a^*} \right\} \\ &\leq \frac{2}{\gamma} + \sum_{i=1}^{N-1} \left(\frac{1}{1+\gamma\hat{f}_{k_i}} - \frac{1}{1+\gamma\hat{f}_{k_{i+1}}} \right) \hat{f}_{k_i} + \max_{a^* \in [K]} \left\{ \frac{1-2p_{a^*}}{\gamma W_{a^*}} - \hat{f}_{a^*} \right\} \\ &\leq \frac{3}{\gamma} + \sum_{i=2}^N \frac{\hat{f}_{k_i} - \hat{f}_{k_{i-1}}}{1+\gamma\hat{f}_{k_i}} + \max_{a^* \in [K]} \left\{ \frac{1-2p_{a^*}}{\gamma W_{a^*}} - \hat{f}_{a^*} \right\}. \end{aligned}$$

571 According to Lemma 9 of [Alon et al., 2013] (included as Lemma D.2 for completeness), we know
 572 that

$$\sum_{i=2}^N \frac{\hat{f}_{k_i} - \hat{f}_{k_{i-1}}}{1+\gamma\hat{f}_{k_i}} = \frac{1}{\gamma} \sum_{i=2}^N \frac{\hat{f}_{k_i} - \hat{f}_{k_{i-1}}}{\frac{1}{\gamma} + \hat{f}_{k_i}} \leq \frac{\text{mas}(G_A)}{\gamma} = \frac{1}{\gamma}, \quad (24)$$

573 where G_A is the subgraph of G restricted to node set A and $\text{mas}(G)$ is the size of the maximum
 574 acyclic subgraphs of G . It is direct to see that any subgraph G of G_{inv} has $\text{mas}(G) = 1$.

575 Next, consider the value of $a^* \in [K]$ that maximizes $\frac{1-2p_{a^*}}{\gamma W_{a^*}} - \hat{f}_{a^*}$. If $a^* \leq k_1$, then we know that
 576 $W_{a^*} = 1$ and $\frac{1-2p_{a^*}}{\gamma W_{a^*}} - \hat{f}_{a^*} \leq \frac{1}{\gamma}$. Otherwise, suppose that $k_i < a^* \leq k_{i+1}$ for some $i \in [N-1]$.
 577 According to the definition of p , if $a^* \neq k_{i+1}$ we know that $p_{a^*} = 0$ and

$$\frac{1}{1+\gamma\hat{f}_{a^*}} \leq \sum_{j' > a^*} p_{j'} = W_{k_{i+1}} = W_{a^*}.$$

578 Therefore,

$$\frac{1-2p_{a^*}}{\gamma W_{a^*}} - \hat{f}_{a^*} = \frac{1}{\gamma W_{a^*}} - \hat{f}_{a^*} \leq \frac{1}{\gamma}.$$

579 Otherwise, $W_{a^*} = W_{k_{i+1}}$ and $\frac{1-2p_{a^*}}{\gamma W_{a^*}} - \hat{f}_{a^*} \leq \frac{1}{\gamma W_{k_{i+1}}} - \hat{f}_{k_{i+1}} = \frac{1}{\gamma}$. Combining the two cases
 580 above and Eq. (24), we obtain that

$$\overline{\text{dec}}_\gamma(p; \hat{f}, x, G_{\text{inv}}) \leq \frac{3}{\gamma} + \frac{1}{\gamma} + \frac{1}{\gamma} = \mathcal{O}\left(\frac{1}{\gamma}\right).$$

581 \square

582 B.4 Undirected and Self-Aware Graphs

583 **Proposition 4.** *When G is an undirected self-aware graph, given any \hat{f} , context x , there exists a*
 584 *closed-form distribution $p \in \Delta(K)$ guaranteeing that $\overline{\text{dec}}_\gamma(p; \hat{f}, x, G) \leq \mathcal{O}\left(\frac{\alpha}{\gamma}\right)$.*

585 *Proof.* We first introduce the closed-form of p and then show that $\overline{\text{dec}}_\gamma(p; \hat{f}, x, G) \leq \mathcal{O}\left(\frac{\alpha}{\gamma}\right)$. Specif-
 586 ically, we first sort \hat{f}_a in an increasing order and choose a maximal independent set by choosing
 587 the nodes in a greedy way. Specifically, we pick $k_1 = \text{argmin}_{i \in [K]} \hat{f}_i$. Then, we ignore all the
 588 nodes that are connected to k_1 and select the node a with the smallest \hat{f}_a in the remaining node
 589 set. This forms a maximal independent set $I \subseteq [K]$, which has size no more than α and is also
 590 a dominating set. Set $p_a = \frac{1}{\alpha + \gamma \hat{f}_a}$ for $a \in I \setminus \{k_1\}$ and $p_{k_1} = 1 - \sum_{a \neq k_1, a \in I} p_a$. This is a valid
 591 distribution as we only choose at most α nodes and $p_a \leq 1/\alpha$ for all $a \in I \setminus \{k_1\}$. Now we show
 592 that $\overline{\text{dec}}_\gamma(p; \hat{f}, x, G) \leq \mathcal{O}\left(\frac{\alpha}{\gamma}\right)$. Specifically, we only need to show that with this choice of p , for any
 593 $a^* \in [K]$,

$$\sum_{a=1}^K p_a \hat{f}_a - \hat{f}_{a^*} + \frac{1}{\gamma} \sum_{a=1}^K \frac{p_a^2}{W_a} + \frac{1 - 2p_{a^*}}{\gamma W_{a^*}} \leq \mathcal{O}\left(\frac{\alpha}{\gamma}\right).$$

594 Plugging in the form of p , we know that

$$\begin{aligned} & \sum_{a=1}^K p_a \hat{f}_a - \hat{f}_{a^*} + \frac{1}{\gamma} \sum_{a=1}^K \frac{p_a^2}{W_a} + \frac{1 - 2p_{a^*}}{\gamma W_{a^*}} \\ & \leq \sum_{a \in I \setminus \{k_1\}} \frac{\hat{f}_a}{\alpha + \gamma \hat{f}_a} - \hat{f}_{a^*} + \frac{1 - 2p_{a^*}}{\gamma W_{a^*}} + \frac{1}{\gamma} \quad (p_a \leq W_a \text{ for all } a \in [K]) \\ & \leq \frac{\alpha}{\gamma} - \hat{f}_{a^*} + \frac{1 - 2p_{a^*}}{\gamma W_{a^*}}. \quad (|I| \leq \alpha) \end{aligned}$$

595 If $a^* = k_1$, then we can obtain that $\frac{1 - 2p_{a^*}}{\gamma W_{a^*}} \leq \frac{1}{\gamma W_{k_1}} \leq \frac{\alpha}{\gamma}$ as $p_{k_1} \geq \frac{1}{\alpha}$ according to the definition of
 596 p . Otherwise, note that according to the choice of the maximal independent set I , $W_{a^*} \geq \frac{1}{\alpha + \gamma \hat{f}_{a'}}$ for
 597 some $a' \in I$ such that $\hat{f}_{a'} \leq \hat{f}_{a^*}$. Therefore,

$$-\hat{f}_{a^*} + \frac{1 - 2p_{a^*}}{\gamma W_{a^*}} \leq -\hat{f}_{a^*} + \frac{1}{\gamma W_{a^*}} \leq -\hat{f}_{a^*} + \frac{\alpha + \gamma \hat{f}_{a'}}{\gamma} \leq \frac{\alpha}{\gamma}.$$

598 Combining the two inequalities above together proves the bound. \square

599 C Implementation Details in Experiments

600 C.1 Implementation Details in Section 5.1

601 We conduct experiments on RCV1 [Lewis et al., 2004], which is a multilabel text-categorization
 602 dataset. We use a subset of RCV1 containing 50000 samples and $K = 50$ sub-classes. Therefore,
 603 the feedback graph in our experiment has $K = 50$ nodes. We use the bag-of-words vector of each
 604 sample as the context with dimension $d = 47236$ and treat the text categories as the arms. In each
 605 round t , the learner receives the bag-of-words vector x_t and makes a prediction $a_t \in [K]$ as the
 606 text category. The loss is set to be $\ell_{t, a_t} = 0$ if the sample belongs to the predicted category a_t and
 607 $\ell_{t, a_t} = 1$ otherwise.

608 The function class we consider is the following linear function class:

$$\mathcal{F} = \{f : f(x, a) = \text{Sigmoid}((Mx)_a), M \in \mathbb{R}^{K \times d}\},$$

609 where $\text{Sigmoid}(u) = \frac{1}{1 + e^{-u}}$ for any $u \in \mathbb{R}$. The oracle is implemented by applying online gradient
 610 descent with learning rate η searched over $\{0.1, 0.2, 0.5, 1, 2, 4\}$. As suggested by [Foster and

611 [Krishnamurthy, 2021](#)], we use a time-varying exploration parameter $\gamma_t = c \cdot \sqrt{\alpha t}$, where t is the
612 index of the iteration, c is searched over $\{8, 16, 32, 64, 128\}$, and α is the independence number of
613 the corresponding feedback graph. Our code is built on PyTorch framework [[Paszke et al., 2019](#)]. We
614 run 5 independent experiments with different random seeds and plot the mean and standard deviation
615 value of PV loss.

616 C.2 Implementation Details in [Section 5.2](#)

617 C.2.1 Details for Results on Random Directed Self-aware Graphs

618 We conduct experiments on a subset of RCV1 containing 10000 samples with $K = 10$ sub-classes.
619 Our code is built on Vowpal Wabbit [[Langford and Zhang, 2007](#)]. For SqaureCB, the exploration
620 parameter γ_t at round t is set to be $\gamma_t = c \cdot \sqrt{Kt}$, where t is the index of the round and c is the hyper-
621 parameter searched over set $\{8, 16, 32, 64, 128\}$. The remaining details are the same as described in
622 [Appendix C.1](#).

623 C.2.2 Details for Results on Synthetic Inventory Dataset

624 In this subsection, we introduce more details in the synthetic inventory data construction, loss function
625 constructions, oracle implementation, and computation of the strategy at each round.

626 **Dataset.** In this experiment, we create a synthetic inventory dataset constructed as follows. The
627 dataset includes T data points, the t -th of which is represented as (x_t, d_t) where $x_t \in \mathbb{R}^m$ is the
628 context and d_t is the realized demand given context x_t . Specifically, in the experiment, we choose
629 $m = 100$ and x_t 's are drawn i.i.d from Gaussian distribution with mean 0 and standard deviation 0.1.
630 The demand d_t is defined as

$$d_t = \frac{1}{\sqrt{m}} x_t^\top \theta + \varepsilon_t,$$

631 where $\theta \in \mathbb{R}^m$ is an arbitrary vector and ε_t is a one-dimensional Gaussian random variable with mean
632 0.3 and standard deviation 0.1. After all the data points $\{(x_t, d_t)\}_{t=1}^T$ are constructed, we normalize
633 d_t to $[0, 1]$ by setting $d_t \leftarrow \frac{d_t - \min_{t' \in [T]} d_{t'}}{\max_{t' \in [T]} d_{t'} - \min_{t' \in [T]} d_{t'}}$. In all our experiments, we set $T = 10000$.

634 **Loss construction.** Next, we define the loss at round t when picking the inventory level a_t with
635 demand d_t , which is defined as follows:

$$\ell_{t,a_t} = h \cdot \max\{a_t - d_t, 0\} + b \cdot \max\{d_t - a_t, 0\}, \quad (25)$$

636 where $h > 0$ is the holding cost per remaining items and $b > 0$ is the backorder cost per remaining
637 items. In the experiment, we set $h = 0.25$ and $b = 1$.

638 **Regression oracle.** The function class we use in this experiment is as follows:

$$\mathcal{F} = \{f : f(x, a) = h \cdot \max\{a - (x^\top \theta + \beta), 0\} + b \cdot \max\{x^\top \theta + \beta - a, 0\}, \theta \in \mathbb{R}^m, \beta \in \mathbb{R}\}.$$

639 This ensures the realizability assumption according to the definition of our loss function shown
640 in [Eq. \(25\)](#). The oracle uses online gradient descent with learning rate η searched over
641 $\{0.01, 0.05, 0.1, 0.5, 1\}$.

642 **Calculation of p_t .** To make SquareCB.G more efficient, instead of solving the convex program
643 defined in [Eq. \(5\)](#), we use the closed-form of p_t derived in [Proposition 3](#), which only requires $\mathcal{O}(K)$
644 computational cost and has the same theoretical guarantee (up to a constant factor) as the one enjoyed
645 by the solution solved by [Eq. \(5\)](#). Similar to the case in [Appendix C.1](#), at each round t , we pick
646 $\gamma_t = c \cdot \sqrt{t}$ with c searched over the set $\{0.25, 0.5, 1, 2, 3, 4\}$. Note again that the independence
647 number for inventory graph is 1.

648 We run 8 independent experiments with different random seeds and plot the mean and standard
649 deviation value of PV loss.

650 **D Auxiliary Lemmas**

651 **Lemma D.1** (Lemma 5 in [Alon et al., 2015]). *Let $G = (V, E)$ be a directed graph with $|V| = K$,*
 652 *in which $i \in N^{\text{in}}(G, i)$ for all vertices $i \in [K]$. Assign each $i \in V$ with a positive weight w_i such*
 653 *that $\sum_{i=1}^n w_i \leq 1$ and $w_i \geq \varepsilon$ for all $i \in V$ for some constant $0 < \varepsilon < \frac{1}{2}$. Then*

$$\sum_{i=1}^K \frac{w_i}{\sum_{j \in N^{\text{in}}(G, i)} w_j} \leq 4\alpha(G) \log \frac{4K}{\alpha(G)\varepsilon},$$

654 *where $\alpha(G)$ is the independence number of G .*

655 **Lemma D.2** (Lemma 9 in [Alon et al., 2013]). *Let $G = (V, E)$ be a directed graph with vertex set*
 656 *$|V| = K$, in which $i \in N^{\text{in}}(G, i)$ for all $i \in [K]$. Let p be an arbitrary distribution over $[K]$. Then,*
 657 *we have*

$$\sum_{i=1}^K \frac{p_i}{\sum_{j \in N^{\text{in}}(G, i)} p_j} \leq \text{mas}(G),$$

658 *where $\text{mas}(G)$ is the size of the maximum acyclic subgraphs of G .*