

# UNDERSTANDING THE TRAINING AND GENERALIZATION OF PRE-TRAINED TRANSFORMER FOR SEQUENTIAL DECISION MAKING

## ABSTRACT

In this paper, we consider the supervised pre-trained transformer for a class of sequential decision-making problems. The class of considered problems is a subset of the general formulation of reinforcement learning in that there is no transition probability matrix; though seemingly restrictive, the subset class of problems covers bandits, dynamic pricing, and newsvendor problems as special cases. Such a structure enables the use of optimal actions/decisions in the pre-training phase, and the usage also provides new insights for the training and generalization of the pre-trained transformer. We first note the training of the transformer model can be viewed as a performative prediction problem, and the existing methods and theories largely ignore or cannot resolve an out-of-distribution issue. We propose a natural solution that includes the transformer-generated action sequences in the training procedure, and it enjoys better properties both numerically and theoretically. The availability of the optimal actions in the considered tasks also allows us to analyze the properties of the pre-trained transformer as an algorithm and explains why it may lack exploration and how this can be automatically resolved. Numerically, we categorize the advantages of pre-trained transformers over the structured algorithms such as UCB and Thompson sampling into three cases: (i) it better utilizes the prior knowledge in the pre-training data; (ii) it can elegantly handle the misspecification issue suffered by the structured algorithms; (iii) for short time horizon such as  $T \leq 50$ , it behaves more greedy and enjoys much better regret than the structured algorithms designed for asymptotic optimality.

## 1 INTRODUCTION

In recent years, transformer-based models [48] have achieved great success in a wide range of tasks such as natural language processing [46, 12], computer vision [20], and also reinforcement learning (RL) [13]. In particular, an offline RL problem can be framed as a sequence prediction problem [19] that predicts the optimal/near-optimal/human action based on the observed history. [13, 29] pioneer this approach with the Decision Transformer (DT), which leverages the auto-regressive nature of transformers to maximize the likelihood of trajectories in offline datasets. In this paper, we focus on a subset class of reinforcement learning problems, which we call sequential decision-making problems, that has no transition probability matrix compared to a general RL problem. This subset of problems are general enough to capture a wide range of applications including stochastic and linear bandit problems [35], dynamic pricing [16, 15], and newsvendor problem [37, 9]. Most importantly, the structure brings two benefits: (i) pre-training – the optimal action is either in closed form or efficiently solvable and thus the optimal action can be used as the prediction target in the pre-training data and (ii) understanding – the setting is more aligned with the existing study [53, 22] on the in-context learning ability of the transformer which gives more insight to the working machinery of transformer on decision-making tasks. Moreover, as a side product, our paper extends pre-trained transformers to the application contexts of business, economics, and operations research which are less studied by the existing works.

To summarize, we work on the following aspects in this paper:

- The pre-training procedure (Section 3): We mathematically formulate the learning setup of supervised pre-trained transformer for sequential decision-making problems, and propose a standardized procedure to generate the pre-training data. In addition, for the training process, we identify an out-of-distribution (OOD) issue largely ignored or unsolved by the existing literature. Specifically, the pre-training data involves two streams of actions – one is the actions in the history (features),

054  
055  
056  
057  
058  
059  
060  
061  
062  
063  
064  
065  
066  
067  
068  
069  
070  
071  
072  
073  
074  
075  
076  
077  
078  
079  
080  
081  
082  
083  
084  
085  
086  
087  
088  
089  
090  
091  
092  
093  
094  
095  
096  
097  
098  
099  
100  
101  
102  
103  
104  
105  
106  
107

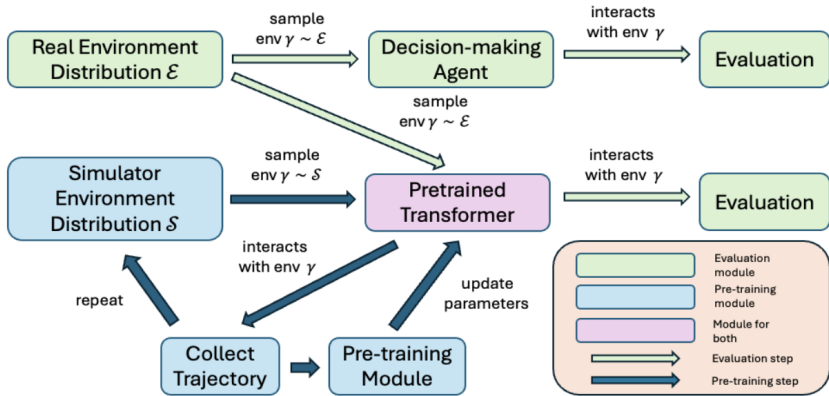


Figure 1: Comparison between the pre-trained transformer framework and traditional sequential decision-making methods (such as the structured algorithm of UCB and Thompson sampling). For traditional methods, the decision-making agent (or policy) interacts directly with a single environment sampled from the real environment distribution, focusing on exploration and exploitation within that specific environment. In contrast, the pre-trained transformer is trained across multiple environments sampled from a simulator distribution. During pre-training, the transformer collects trajectories and updates its parameters by interacting with diverse environments. Once pre-trained, the transformer functions as an algorithm and can be applied effectively in the real environment just as the structured algorithms. But unlike the structured algorithms, the pre-trained transformer leverages the huge amount of pre-training data sampled from the simulation environment.

and the other is the action to predict (target). [38] show that when both streams are generated by one structured algorithm (say UCB or Thompson sampling), the pre-training of the transformer will result in an imitation learning procedure that imitates the structured algorithm at best. If one desires to obtain a better algorithm than the structured algorithm, like using the optimal action [36] as the target to predict in our case, this will cause the OOD issue (reflected by an  $\Theta(\exp T)$  factor in the regret bound of [38]). To resolve this matter, we propose an algorithmic solution by injecting the transformer-generated action sequence into the pre-training data, and this also draws a connection with the problem of performative prediction [43].

- Understanding transformers as decision makers (Section 4): The structure of sequential decision-making problems enables the usage of optimal actions in pre-training data, and as noted earlier and will be exemplified throughout our paper, this brings benefits for both training and theoretical understanding. First, we connect the pre-trained transformer with the existing literature studying in-context learning [53, 56, 6] and we establish the pre-trained transformer as a near Bayes-optimal predictor for the optimal action in a Bayesian sense. Next, we show that (i) why the transformer model exhibits a lack of exploration and (ii) how it achieves a paramount or better performance than the structured algorithms. Other than these, an important aspect is that the problem structure of sequential decision-making renders the setup of our pre-trained transformer different from the existing works that also involve interactions with the underlying environment under the general RL setup. For example, the Online Decision Transformer (ODT) [57] lets the transformers interact with the underlying environment during the training phase, but the interactions are limited to a single environment and are designed to explore better trajectories or policies within that environment than the given offline data. We defer a more detailed comparison between the existing usage of pre-trained transformers on general RL problems and our framework to Appendix A.2.

- Numerically, when and why the pre-trained transformer is better than the structured algorithms (Section 5): While there are many structured algorithms such as UCB and Thompson sampling for each of the decision-making problems that the transformer model is applied to, the question is whether and when the transformer can achieve a better performance than these algorithms. We use numerical experiments to attribute the advantage of the pre-trained transformer into three scenarios: (i) it better utilizes the prior knowledge in the pre-training data; (ii) it can elegantly handle the misspecification issue suffered by the structured algorithms; (iii) for short time horizon such as

108  $T \leq 50$  or 100, it behaves more greedy and enjoys much better regret than the structured algorithms  
 109 which are designed for asymptotic optimality.

110 We defer more discussions on the related works to Section A.  
 111

## 112 2 PROBLEM SETUP

113 In this section, we introduce a general formulation for sequential decision-making and describe  
 114 the setup of supervised pre-training. An agent (decision maker) makes decisions in a sequential  
 115 environment. The full procedure is described by the sequence

$$116 (X_1, a_1, o_1, X_2, a_2, o_2, \dots, X_T, a_T, o_T).$$

117 Here  $T$  is the horizon. At each time  $t = 1, \dots, T$ , the agent first observes a context vector  $X_t \in$   
 118  $\mathcal{X} \subset \mathbb{R}^d$  and takes the action/decision  $a_t \in \mathcal{A}$ , and then observes an outcome  $o_t \in \mathcal{O} \subset \mathbb{R}^k$ . Define  
 119 the history

$$120 H_t := (X_1, a_1, o_1, \dots, X_{t-1}, a_{t-1}, o_{t-1}, X_t),$$

121 and the agent’s decision  $a_t$  is based on  $H_t$ , i.e., non-anticipatory. In this light, we can model the  
 122 decision as a function of  $H_t$ ,

$$123 a_t = \text{TF}_\theta(H_t)$$

124 where  $\theta$  encapsulates all the parameters of the function. That is, the agent learns from the past  
 125 interactions with the underlying *environment* – tuples of  $(X_t, a_t, o_t)$ ’s to optimize future decisions,  
 126 and this process that transfers the past knowledge to the future is described by the function  $\text{TF}_\theta$ . The  
 127 function  $\text{TF}_\theta$  can take as inputs a variable-length sequence ( $H_t$  can be of variable length), which is  
 128 the case for the generative pre-trained transformers [46].  
 129  
 130  
 131

### 132 2.1 ENVIRONMENT AND PERFORMANCE METRICS

133 The function  $\text{TF}_\theta$  describes how the actions  $a_t$ ’s are generated. Now we describe how  $X_t$ ’s and  $o_t$ ’s  
 134 are generated. Assume the context  $X_t$ ’s are i.i.d. and  $o_t$  follows a conditional distribution given  $X_t$   
 135 and  $a_t$

$$136 X_t \stackrel{\text{i.i.d.}}{\sim} \mathbb{P}_\gamma(\cdot), o_t \stackrel{\text{i.i.d.}}{\sim} \mathbb{P}_\gamma(\cdot | X_t, a_t) \quad (1)$$

137 where  $\gamma$  encapsulates all the parameters for these two probability distributions. In particular, we  
 138 emphasize that the parameter  $\gamma$  is unknown to the agent.

139 At each time  $t$ , the agent collects a random reward  $R_t = R(X_t, a_t)$  which depends on  $X_t$ ,  $a_t$ ,  
 140 and also some possible exogenous randomness. The observation  $o_t$  includes  $R_t$  as its first coordi-  
 141 nate. The expected reward is defined by  $r(X_t, a_t) = \mathbb{E}[R(X_t, a_t) | X_t, a_t]$ . Specifically, this reward  
 142 function  $r(\cdot, \cdot)$  is unknown to the agent and it also depends on the underlying environment  $\gamma$ .  
 143  
 144

145 The performance of the agent, namely, the decision function  $\text{TF}_\theta$ , is usually measured by the notion  
 146 of *regret*, which is defined by

$$147 \text{Regret}(\text{TF}_\theta; \gamma) := \mathbb{E} \left[ \sum_{t=1}^T r(X_t, a_t^*) - r(X_t, a_t) \right] \quad (2)$$

148 where the action  $a_t^*$  is the optimal action that maximizes the reward

$$149 a_t^* := \arg \max_{a \in \mathcal{A}} r(X_t, a) \quad (3)$$

150 and the optimization assumes the knowledge of the underlying environment, i.e., the parameter  $\gamma$ .  
 151 In the regret definition equation 2, the expectation is taken with respect to the underlying probability  
 152 distribution equation 1 and possible randomness in the agent’s decision function  $\text{TF}_\theta$ . We use the  
 153 arguments  $\text{TF}_\theta$  and  $\gamma$  to reflect the dependency of the regret on the decision function  $\text{TF}_\theta$  and the  
 154 environment  $\gamma$ . Also, we note that the benchmark is defined as a dynamic oracle where  $a_t^*$  maybe  
 155 differ over time. Regret is not a perfect performance measure for the problem, and we also do not  
 156 believe the algorithm design for sequential decision-making problems should purely aim for getting  
 157 a better regret bound. Despite this, it serves as a good monitor for the effectiveness of the underlying  
 158 learning and the optimization procedure.  
 159  
 160  
 161

## 2.2 SUPERVISED PRE-TRAINING

In the classic study of these sequential decision-making problems, the common paradigm is to assume some structure for the underlying environment and/or the reward function, and to design algorithms accordingly. These algorithms are often known as online or reinforcement learning algorithms and are usually hard to combine with prior knowledge such as pre-training data. Comparatively, the supervised pre-training formulates the decision-making problem as a supervised learning prediction task that predicts the optimal or a near-optimal action, and thus it can utilize the vast availability of pre-training data that can be generated from simulated environments.

Specifically, the supervised pre-training first constructs a pre-training dataset

$$\mathcal{D}_{\text{PT}} := \left\{ \left( H_1^{(i)}, a_1^{(i)*} \right), \left( H_2^{(i)}, a_2^{(i)*} \right), \dots, \left( H_T^{(i)}, a_T^{(i)*} \right) \right\}_{i=1}^n.$$

In particular, we assume the parameter  $\gamma$  (see equation 1) that governs the generation of  $X_t$  and  $o_t$  are generated from some environment distribution  $\mathcal{P}_\gamma$ . Then, we first generate  $n$  environments denoted by  $\gamma_1, \dots, \gamma_n$ . Then there are two important aspects with regard to the generation of  $\mathcal{D}_{\text{PT}}$ :

- Action  $a_t^{(i)*}$ : For this paper, we generate  $a_t^{(i)*}$  as the optimal action specified by equation 3 and this requires the knowledge of the underlying environment  $\gamma_i$ . A common property of all these sequential decision-making problems is that the optimal action can be easily or often analytically computed. Other option is also possible, say,  $a_t^{(i)*}$  can be generated based on some expert algorithms such as Thompson sampling or UCB algorithms. In this case, the action  $a_t^{(i)*} = \text{Alg}(H_t^{(i)})$  where  $\text{Alg}(\cdot)$  represents the algorithm that maps from  $H_t^{(i)}$  to the action without utilizing knowledge of  $\gamma_i$ .

- History  $H_t^{(i)} = \left( X_1^{(i)}, a_1^{(i)}, o_1^{(i)}, \dots, X_{t-1}^{(i)}, a_{t-1}^{(i)}, o_{t-1}^{(i)}, X_t^{(i)} \right)$ . Note that  $X_t^{(i)}$ 's and  $o_t^{(i)}$ 's are generated based on the parameter  $\gamma_i$  following equation 1. The action  $a_t^{(i)}$  is generated based on some pre-specified decision function  $f$  in a recursive manner:

$$a_t^{(i)} = f \left( H_t^{(i)} \right), \quad H_{t+1}^{(i)} = \left( H_t^{(i)}, a_t^{(i)}, o_t^{(i)}, X_{t+1}^{(i)} \right) \quad (4)$$

for  $t = 1, \dots, T$ .

Formally, we define a distribution  $\mathcal{P}_{\gamma, f}$  that generates  $H_t$  and  $a_t^*$  as

$$X_\tau \sim \mathbb{P}_\gamma(\cdot), \quad a_\tau = f(H_\tau), \quad o_\tau \sim \mathbb{P}_\gamma(\cdot | X_\tau, a_\tau), \quad H_{\tau+1} = (H_\tau, a_\tau, o_\tau, X_{\tau+1}). \quad (5)$$

for  $\tau = 1, \dots, t-1$  and  $a_t^*$  is specified by equation 3. The parameter  $\gamma$  and the function  $f$  jointly parameterize the distribution  $\mathcal{P}_{\gamma, f}$ . The parameter  $\gamma$  governs the generation of  $X_t$  and  $o_t$  just as in equation 1. The decision function  $f$  maps from the history to the action and it governs the generation of  $a_\tau$ 's in the sequence  $H_t$ . In this way, the pre-training data is generated with the following flow:

$$\mathcal{P}_\gamma \rightarrow \gamma_i \rightarrow \mathcal{P}_{\gamma_i, f} \rightarrow \left\{ \left( H_1^{(i)}, a_1^{(i)*} \right), \dots, \left( H_T^{(i)}, a_T^{(i)*} \right) \right\}$$

where  $f$  is a prespecified decision function used to generate the data.

Based on the dataset  $\mathcal{D}_{\text{PT}}$ , the supervised pre-training refers to the learning of the decision function  $\text{TF}_\theta$  through minimizing the following empirical loss

$$\hat{\theta} := \arg \min_{\theta} \frac{1}{nT} \sum_{i=1}^n \sum_{t=1}^T l \left( \text{TF}_\theta \left( H_t^{(i)} \right), a_t^{(i)*} \right) \quad (6)$$

where the loss function  $l(\cdot, \cdot) : \mathcal{A} \times \mathcal{A} \rightarrow \mathbb{R}$  specifies the prediction loss. Thus the key idea of supervised pre-training is to formulate the sequential decision-making task as a *sequence modeling* task that predicts the optimal action  $a_t^{(i)*}$  given the history  $H_t^{(i)}$ .

In the test phase, a new environment parameter  $\gamma$  is generated from  $\mathcal{P}_\gamma$  and we apply the learned decision function  $\text{TF}_{\hat{\theta}}$ . Specifically, the procedure is described by

$$X_\tau \sim \mathbb{P}_\gamma(\cdot), \quad a_\tau = \text{TF}_{\hat{\theta}}(H_\tau), \quad o_\tau \sim \mathbb{P}_\gamma(\cdot | X_\tau, a_\tau), \quad H_{\tau+1} = (H_\tau, a_\tau, o_\tau, X_{\tau+1}). \quad (7)$$

The only difference between the test dynamics equation 7 and the training dynamics equation 5 is that the action  $a_t$  (or  $a_\tau$ ) is generated by  $\text{TF}_{\hat{\theta}}$  instead of the pre-specified decision function  $f$ . In this

way, we can write  $(H_t, a_t^*) \sim \mathcal{P}_{\gamma, \text{TF}_\theta}$  for the test phase. Accordingly, we can define the expected loss of the decision function  $\text{TF}_\theta$  under the environment specified by  $\gamma$  as

$$L(\text{TF}_\theta; \gamma) := \mathbb{E}_{(H_t, a_t^*) \sim \mathcal{P}_{\gamma, \text{TF}_\theta}} \left[ \sum_{t=1}^T l(\text{TF}_\theta(H_t), a_t^*) \right] \quad (8)$$

where the expectation is taken with respect to  $X_t$ ,  $o_t$  and the possible randomness in  $a_t$ . As before,  $a_t^*$  are the optimal action defined by equation 3 and obtained from the knowledge of  $\gamma$ .

We defer discussions on how various sequential decision-making problems can be formulated under the general setup above and how the pre-training loss is related to the regret to Section B.

### 3 PRE-TRAINING AND GENERALIZATION

We first point out a subtle point of the supervised pre-training setting presented in the previous section. Unlike supervised learning, the expected loss equation 8 may not be equal to the expectation of the empirical loss equation 6. As noted earlier, this is because in the sequence  $H_t$ , the action  $a_t = \text{TF}_{\hat{\theta}}(H_t)$  as in equation 7 during the test phase as opposed to that  $a_t = f(H_t)$  as in equation 5 for the pretraining data  $\mathcal{D}_{\text{PT}}$ . Taking expectation with respect to  $\gamma$  in equation 8, we can define the expected loss as

$$L(\text{TF}_\theta) := \mathbb{E}_{\gamma \sim \mathcal{P}_\gamma} [L(\text{TF}_\theta; \gamma)].$$

However, the expectation of the empirical loss equation 6 is

$$L_f(\text{TF}_\theta) := \mathbb{E}_{\gamma \sim \mathcal{P}_\gamma} [L_f(\text{TF}_\theta; \gamma)] \quad (9)$$

where

$$L_f(\text{TF}_\theta; \gamma) := \mathbb{E}_{(H_t, a_t^*) \sim \mathcal{P}_{\gamma, f}} \left[ \sum_{t=1}^T l(\text{TF}_\theta(H_t), a_t^*) \right].$$

Here  $f$  is the pre-specified decision function used in generating the pre-training data  $\mathcal{D}_{\text{PT}}$ . We note that the discrepancy between  $L(\text{TF}_\theta)$  and  $L_f(\text{TF}_\theta)$  is caused by the difference in the generations of the samples  $(H_t, a_t^*)$ , namely,  $\mathcal{P}_{\gamma, \text{TF}_\theta}$  versus  $\mathcal{P}_{\gamma, f}$ ; which essentially reduces to how the actions  $a_t$ 's are generated in  $H_t$ . In the training phase, this is generated based on  $f$ , while measuring the test loss requires that the actions to be generated are based on the transformer  $\text{TF}_\theta$ . This discrepancy is somewhat inevitable because when generating the training data, there is no way we can know the final parameter  $\hat{\theta}$ . Consequently, there is no direct guarantee on the loss for  $L(\text{TF}_{\hat{\theta}}; \gamma)$  or  $L(\text{TF}_{\hat{\theta}})$  with  $\hat{\theta}$  being the final learned parameter. This out-of-distribution (OOD) issue motivates the design of our algorithm.

Algorithm 1 describes our algorithm to resolve the above-mentioned OOD issue between training and test. It consists of two phases, an early training phase and a mixed training phase. For the early training phase, the data are generated from the distribution  $\mathcal{P}_{\gamma_i, f}$ , while for the mixed training phase, the data are generated from both  $\mathcal{P}_{\gamma_i, f}$  and  $\mathcal{P}_{\gamma_i, \text{TF}_{\theta_t}}$  with a ratio controlled by  $\kappa$ . The decision function  $f$  can be chosen as a uniform distribution over the action space as in the literature [13, 54, 36], or other more complicated options (details on the  $f$  used in our experiments are provided in Appendix E.1.2). When there is no mixed phase, i.e.,  $M_0 = M$ , the algorithm recovers that of [36]. Intuitively, the benefits of involving the data generated from  $\mathcal{P}_{\gamma_i, \text{TF}_{\theta_t}}$  is to make the sequences  $H_t^{(i)}$ 's in the training data closer to the ones generated from the transformer. For the mixed training phase, a proportion of the data is still generated from the original  $\mathcal{P}_{\gamma_i, f}$  because the generation from  $\mathcal{P}_{\gamma_i, \text{TF}_{\theta_t}}$  costs relatively more time. For our numerical experiments, we choose  $\kappa = 1/3$ .

As far as we know, this OOD matter is largely ignored in the existing empirical works on supervised pre-training transformers. Theoretically, [36] directly assumes the pre-trained transformer learns (for discrete action space) the optimal decision function  $\text{Alg}^*$  defined in the next section and develops their theoretical results accordingly. [38] show that when (i)  $a_t^{(i)*}$ 's in the training data are no longer the optimal actions  $a_t^*$  in our setting but generated from some algorithm  $\text{Alg}$  and (ii) the decision function  $f = \text{Alg}$  in the distribution  $\mathcal{P}_{\gamma_i, f}$ , the pre-trained transformer  $\text{TF}_{\hat{\theta}} \rightarrow \text{Alg}$  as the number of training samples goes to infinity. We note that this only implies that the transformer model can *imitate* an existing algorithm  $\text{Alg}$ . In other words, setting  $f = \text{Alg}$  brings the theoretical

**Algorithm 1** Supervised pre-training transformers for sequential decision making

**Require:** Iterations  $M$ , early training phase  $M_0 \in [1, M]$ , number of training sequences per iteration  $n$ , ratio  $\kappa \in [0, 1]$ , pre-training decision function  $f$ , prior distribution  $\mathcal{P}_\gamma$ , initial parameter  $\theta_0$

- 1: Initialize  $\theta_1 = \theta_0$
- 2: **for**  $m = 1, \dots, M_0$  **do**  
     %% Early training phase
- 3:     Sample  $\gamma_1, \gamma_2, \dots, \gamma_n$  from  $\mathcal{P}_\gamma$  and  $S_i = \left\{ \left( H_1^{(i)}, a_1^{(i)*} \right), \dots, \left( H_T^{(i)}, a_T^{(i)*} \right) \right\}$  from  $\mathcal{P}_{\gamma_i, f}$
- 4:     Optimize over the generated dataset  $\mathcal{D}_m = \{S_i\}_{i=1}^n$  and obtain the updated parameter  $\theta_{m+1}$
- 5: **end for**
- 6: **for**  $m = M_0 + 1, \dots, M$  **do**  
     %% Mixed training phase
- 7:     Sample  $\gamma_1, \gamma_2, \dots, \gamma_n$  from  $\mathcal{P}_\gamma$
- 8:     For  $i = 1, \dots, \kappa n$ , sample  $S_i = \left\{ \left( H_1^{(i)}, a_1^{(i)*} \right), \dots, \left( H_T^{(i)}, a_T^{(i)*} \right) \right\}$  from  $\mathcal{P}_{\gamma_i, f}$
- 9:     For  $i = \kappa n, \dots, n$ , sample  $S_i = \left\{ \left( H_1^{(i)}, a_1^{(i)*} \right), \dots, \left( H_T^{(i)}, a_T^{(i)*} \right) \right\}$  from  $\mathcal{P}_{\gamma_i, \text{TF}_{\theta_m}}$
- 10:     Optimize over the generated dataset  $\mathcal{D}_m = \{S_i\}_{i=1}^n$  and obtain the updated parameter  $\theta_{m+1}$
- 11: **end for**
- 12: **Return:**  $\hat{\theta} = \theta_{M+1}$

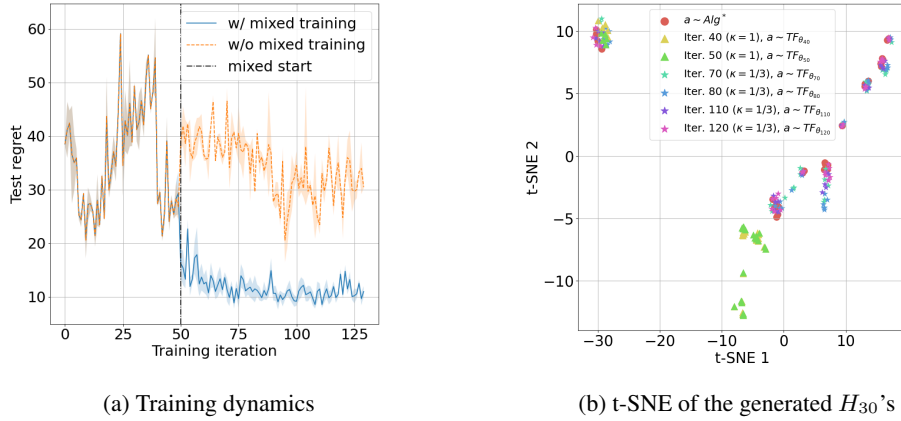


Figure 2: (a) Training dynamics. Orange:  $M_0 = M = 130$ . Blue:  $M_0 = 50$  and  $M = 130$ . It shows the effectiveness of injecting/mixing the transformer-generated sequence into the training procedure. (b) A visualization of the  $H_t$  with  $a_\tau$ 's in  $H_t$  generated from various  $\text{TF}_{\theta_m}$ . For each  $\text{TF}_{\theta_m}$ , we generate 30 sequences. The decision function  $\text{Alg}^*$  is defined in the next section. We observe (i) there is a shift over time in terms of the transformer-generated action sequence, and thus the training should adaptively focus more on the recently generated sequence like the design in Algorithm 1; (ii) the action sequence gradually gets closer to the optimal decision function  $\text{Alg}^*$ . The experiment setups are deferred to Appendix E.4.

benefits of consistency, but it excludes the possibility of using the optimal  $a_t^*$  in the pre-training and thus makes it impossible for the pre-trained transformer  $\text{TF}_{\hat{\theta}}$  to perform better than the algorithm  $\text{Alg}$ . When  $f \neq \text{Alg}$ , the OOD issue is reflected by the distribution ratio in Definition 5 of [38]; the ratio can be as large as  $\Theta(\exp T)$  and thus renders the regret bound to be a vacuous one on the order of  $\Theta(\exp T)$ . Compared to these works, our algorithm provides an algorithmic approach to resolve the OOD issue.

Algorithm 1 from an optimization perspective can be viewed as an iterative procedure to optimize

$$\mathbb{E}_{\gamma \sim \mathcal{P}_\gamma} \left[ \mathbb{E}_{(H_t, a_t^*) \sim \kappa \mathcal{P}_{\gamma, f} + (1-\kappa) \mathcal{P}_{\gamma, \text{TF}_\theta}} \left[ \sum_{t=1}^T l(\text{TF}_\theta(H_t), a_t^*) \right] \right]$$

where both the input data pair  $(H_t, a_t^*)$  and the prediction function involve the parameter  $\theta$ . This falls into the paradigm of performative prediction [43, 28, 39] where the prediction model may affect the data generated to be predicted. In the literature of performative prediction, a critical matter is the instability issue which may cause the parameter  $\theta_m$  to oscillate and not converge. [43] shows that the matter can be solved with strong conditions such as smoothness and strong convexity on the objective function. However, we do not encounter this instability in our numerical experiment, and we make an argument as the following claim. That is, when the underlying function class, such as  $\{TF_\theta : \theta \in \Theta\}$  is rich enough, one does not need to worry about such instability.

**Claim 3.1.** The instability of a performative prediction algorithm and the resultant non-convergence does not happen when the underlying prediction function is rich enough.

The claim will be further reinforced by the optimal decision function  $\text{Alg}^*$  and Proposition 4.1 in the next section. We also defer more discussions on the claim to Section D.1. While it brings a peace of mind from the optimization perspective, the following proposition justifies the mixed training procedure from a statistical generalization perspective.

**Proposition 3.2.** Suppose  $\hat{\theta}$  is determined by equation 6 where  $\kappa n$  data sequences are from  $\mathcal{P}_{\gamma, f}$  and  $(1 - \kappa)n$  data sequences are from  $\mathcal{P}_{\gamma, TF_{\hat{\theta}}}$  for some parameter  $\tilde{\theta}$ . For a Lipschitz and bounded loss  $l$ , the following generalization bound holds with probability at least  $1 - \delta$ ,

$$L(TF_{\hat{\theta}}) \leq \hat{L}(TF_{\hat{\theta}}) + \sqrt{\frac{\text{Comp}(\{TF_\theta : \theta \in \Theta\})}{nT}} \\ + O\left(\kappa T \mathbb{E}_{\gamma \sim \mathcal{P}_\gamma} [W_1(\mathcal{P}_{\gamma, f}, \mathcal{P}_{\gamma, TF_{\hat{\theta}}})] + (1 - \kappa) T \mathbb{E}_{\gamma \sim \mathcal{P}_\gamma} [W_1(\mathcal{P}_{\gamma, TF_{\tilde{\theta}}}, \mathcal{P}_{\gamma, TF_{\hat{\theta}}})] + \sqrt{\frac{\log(\frac{1}{\delta})}{nT}}\right)$$

where  $\text{Comp}(\cdot)$  denotes some complexity measure and  $W_1(\cdot, \cdot)$  is the Wasserstein-1 distance. The notation  $O(\cdot)$  omits the boundedness parameters.

The proof of Proposition 3.2 is not technical and follows the standard generalization bound [40] together with the OOD analysis [8]. We hope to use the result to provide more insights into the training procedure. The parameter  $\tilde{\theta}$  can be interpreted as close to  $\hat{\theta}$  as in the design of our algorithm. This will result in the term related to the distance between  $\mathcal{P}_{\gamma, TF_{\tilde{\theta}}}$  and  $\mathcal{P}_{\gamma, TF_{\hat{\theta}}}$  small. At the two ends of the spectrum, when  $\kappa = 1$ , it will result in a large discrepancy between the training and the test data and thus the right-hand-side will explode; when  $\kappa = 0$ , this will result in the tightest bound but it will cause more computational cost in the training procedure.

## 4 LEARNED DECISION FUNCTION AS AN ALGORITHM

In the previous section, we discuss the training and generalization aspects of the pre-trained transformer. Now, we return to the perspective of sequential decision-making and discuss the properties of the decision function when it is used as an algorithm.

In the ideal case, we define the Bayes-optimal decision function as follow

$$\text{Alg}^*(H) := \arg \min_{a \in \mathcal{A}} \mathbb{E}[l(a, a^*(H, \gamma)) | H] \\ = \arg \min_{a \in \mathcal{A}} \int_{\gamma} l(a, a^*(H, \gamma)) \mathbb{P}(H | \gamma) d\mathcal{P}_\gamma$$

where  $\mathcal{P}_\gamma$  is the environment distribution that generates the environment parameter  $\gamma$ , the optimal action  $a^*(H, \gamma)$  depends on both the history  $H$  and the environment parameter  $\gamma$  (as defined in equation 3), and  $\mathbb{P}(H | \gamma)$  is proportional to the likelihood of observing the history  $H$  under the environment  $\gamma$

$$\mathbb{P}(H | \gamma) \propto \prod_{t=1}^{\tau} \mathbb{P}_\gamma(X_t) \cdot \prod_{t=1}^{\tau-1} \mathbb{P}_\gamma(o_t | X_t, a_t)$$

where  $H = (X_1, a_1, o_1, \dots, X_{\tau-1}, a_{\tau-1}, o_{\tau-1}, X_\tau)$  for some  $\tau \in \{1, \dots, T\}$  and the distributions  $\mathbb{P}_\gamma(\cdot)$  and  $\mathbb{P}_\gamma(\cdot | X_t, a_t)$  are as in equation 1. In the definition, the integration is with respect to

the environment distribution  $\mathcal{P}_\gamma$ , and each possible environment  $\gamma$  is weighted with the likelihood  $\mathbb{P}(H|\gamma)$ .

The following proposition relates  $\text{Alg}^*(H)$  with the pre-training objective which justifies why it is called the *optimal* decision function. The loss function  $L_f(\text{Alg})$  is defined as equation 9; recall that the decision function  $f$  is used to generate the actions in  $H_t$  and the decision function  $\text{Alg}$  is used to predict  $a_t^*$ .

**Proposition 4.1.** *The following holds for any distribution  $\mathcal{P}_{\gamma,f}$ :*

$$\text{Alg}^*(\cdot) \in \arg \min_{\text{Alg} \in \mathcal{F}} L_f(\text{Alg}) = \mathbb{E}_{\gamma \sim \mathcal{P}_\gamma} \left[ \mathbb{E}_{(H_t, a_t^*) \sim \mathcal{P}_{\gamma,f}} \left[ \sum_{t=1}^T l(\text{Alg}(H_t), a_t^*) \right] \right]$$

where  $\mathcal{F}$  is the family of all measurable functions (on a properly defined space that handles variable-length inputs). In particular, we can choose the decision function  $f$  properly to recover the test loss  $L(\text{TF}_\theta)$  or the expected training loss  $L_f(\text{TF}_\theta)$  in Section 3.

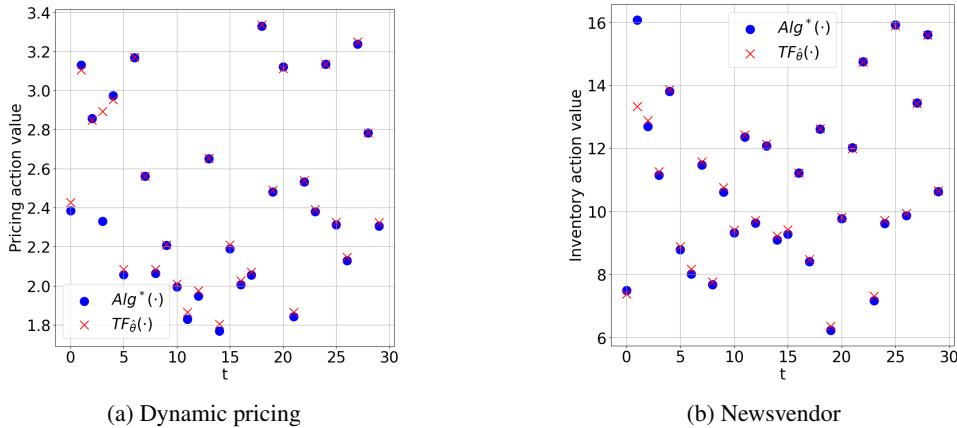


Figure 3: The pre-trained transformer  $\text{TF}_\theta$  well matches the optimal decision function  $\text{Alg}^*$ . For both figures, we plot decision trials for  $\text{Alg}^*$  and  $\text{TF}_\theta$ . The optimal actions change over time because  $X_t$ 's are different for different time  $t$ . The experiment setup is deferred to Appendix E.4.

Proposition 4.1 states that the function  $\text{Alg}^*$  is one minimizer of the expected loss under any distribution  $\mathcal{P}_{\gamma,f}$ . To see this,  $\text{Alg}^*$  is defined in a pointwise manner for every possible  $H$ . The probability  $\mathcal{P}_{\gamma,f}$  defines a distribution over the space of  $H$ . The pointwise optimality of  $\text{Alg}^*$  easily induces the optimality of  $\text{Alg}^*$  for any  $\mathcal{P}_{\gamma,f}$ . Furthermore, this implies that  $\text{Alg}^*$  is also the optimal solution to the expected test loss and the expected training loss. This means that if the transformer class of  $\text{TF}_\theta$  is rich enough to cover  $\text{Alg}^*$ , and with an infinite amount of data, the supervised pre-training will result in  $\text{Alg}^*$  at best. Thus we can interpret the property of the supervised pre-trained transformer by  $\text{Alg}^*$  as a proxy. Such a Bayes-optimal function and the Bayesian perspective have been discussed under various settings in the in-context learning literature [53, 56, 30].

**Example 4.2.**  $\text{Alg}^*$  behaves as (a) posterior sampling under the cross-entropy loss; (b) posterior averaging under the squared loss; (c) posterior median under the absolute loss; where the posterior distribution is with respect to  $\gamma$  and it is defined by  $\mathcal{P}(\gamma|H) \propto \mathbb{P}(H|\gamma) \cdot \mathcal{P}_\gamma(\gamma)$ .

We defer more details for Example 4.2 to Appendix D.4. Part (a) has been directly employed as Assumption 1 in [36] for their theoretical developments. Indeed, part (a) corresponds to a discrete action space such as multi-armed bandits while part (b) and part (c) correspond to the continuous action space such as linear bandits, pricing, and newsvendor problems.

**Proposition 4.3.** *There exist a linear bandits problem instance and a dynamic pricing problem instance, i.e., an environment distribution  $\mathcal{P}_\gamma$ , such that the optimal decision function learned under the squared loss incurs a linear regret for every  $\gamma$ , i.e.,  $\text{Regret}(\text{Alg}^*; \gamma) = \Omega(T)$  for every  $\gamma$ .*

Proposition 4.3 states a negative result that the optimal decision function  $\text{Alg}^*$ , albeit being optimal in a prediction sense (for predicting the optimal action), does not serve as a good algorithm



for sequential decision-making problems. To see the intuition, the decision function is used recursively in the test phase (see equation 7). While it gives the best possible prediction, it does not conduct *exploration* which is critical for the concentration of the posterior measure. This highlights one key difference between the sequential decision-making problem and the regression problem [22] when applying the pre-trained transformer. For the regression problem, the in-context samples (in analogy context-action-observation tuples in  $H_t$ ) are i.i.d. generated, and this results in a convergence/concentration of the posterior to the true model/environment  $\gamma$  [23]. Yet for sequential decision-making problems, the dynamics equation 7 induces a dependency throughout  $H_t$ , and it may result in a non-concentration behavior for the posterior. And this is the key for constructing the linear-regret example in Proposition 4.3 and the examples are also inspired from and share the same spirit as the discussions in [24]. Such lack of exploration does not happen for discrete action space [36]; this is because for discrete action space, the output  $\text{TF}_{\hat{\theta}}$  gives the posterior distribution as a distribution over the action space. The distribution output will lead to a randomized action and thus automatically perform exploration.

**Proposition 4.4.** *Consider a finite prior  $\mathcal{P}_{\gamma}$  supported on  $\Gamma = \{\gamma_1, \dots, \gamma_k\}$ . Suppose there exists constants  $\Delta_{\text{Exploit}}$  and  $\Delta_{\text{Explore}}$  such that for any  $\gamma \in \Gamma$  and  $H_t$ ,  $\text{TF}_{\hat{\theta}}$  and  $\text{Alg}^*(H_t)$  satisfies*

- $\mathbb{E}[r(X_t, \text{Alg}^*(H_t)) - r(X_t, \text{TF}_{\hat{\theta}}(H_t)) | H_t] \leq \Delta_{\text{Exploit}}$ , where the expectation is taken with respect to the possible randomness in  $\text{Alg}^*$  and  $\text{TF}_{\hat{\theta}}$ .
- *The KL divergence*

$$\text{KL}(\mathbb{P}_{\gamma}(\cdot | X_t, \text{TF}_{\hat{\theta}}(H_t)) \| \mathbb{P}_{\gamma'}(\cdot | X_t, \text{TF}_{\hat{\theta}}(H_t))) \geq \Delta_{\text{Explore}}$$

for any  $\gamma' \in \Gamma$ , and the log-likelihood ratio between  $\mathbb{P}_{\gamma}(\cdot | X_t, \text{TF}_{\hat{\theta}}(H_t))$  and  $\mathbb{P}_{\gamma'}(\cdot | X_t, \text{TF}_{\hat{\theta}}(H_t))$  is  $C_{\sigma^2} \Delta_{\text{Explore}}$ -sub-Gaussian for any  $\gamma' \in \Gamma$ .

Then under conditions on the boundedness of the reward and environment parameters, we have

$$\text{Regret}(\text{TF}_{\hat{\theta}}; \gamma) \leq O\left(\Delta_{\text{Exploit}}T + \frac{\log k}{\Delta_{\text{Explore}}}\right)$$

for any  $\gamma \in \Gamma$ , where the notation  $O(\cdot)$  omits logarithm terms in  $T$  and the parameters in the boundedness assumptions.

There are simple remedies to the linear regret behavior of the Bayes-optimal decision function, for example, using an  $\epsilon$ -greedy version of  $\text{Alg}^*$  (and  $\text{TF}_{\hat{\theta}}$ ). Numerically, we do not observe such behavior for any experiment even without any further modification of the pre-trained transformer  $\text{TF}_{\hat{\theta}}$ . Proposition 4.4 provides a simple explanation for why one does not have to be too pessimistic. Intuitively, the pre-trained model  $\text{TF}_{\hat{\theta}}$  can be viewed as the optimal decision function  $\text{Alg}^*$  plus some random noise, and the random noise actually plays the role of exploration and helps the posterior to concentration. Proposition 4.4 materializes such an intuition and it involves two conditions. The first condition is that  $\text{TF}_{\hat{\theta}}$  behaves closely to  $\text{Alg}^*$  and the (reward) deviation between these two is captured by  $\Delta_{\text{Exploit}}$ . This condition can be justified by Figure 3. The second condition characterizes the KL divergence between the conditional distribution of the observation  $o_t$  for two environments: the true testing environment  $\gamma$  and any other one  $\gamma' \in \Gamma$ . Note that this KL divergence critically depends on  $\text{TF}_{\hat{\theta}}$  because the ability to distinguish between two environments depends on the intensity of exploration of the underlying algorithm. These conditions work more like a stylized model to understand the behavior of the transformer. It is a largely simplified model; say, the parameter  $\Delta_{\text{Explore}}$  should have a dependency on  $t$  for many online learning algorithms such as UCB or Thompson sampling. Nevertheless, the conditions give us a perspective to understand how the transformer works as an online algorithm. In particular,  $\text{TF}_{\hat{\theta}}$  deviates from  $\text{Alg}^*$  by an amount of  $\Delta_{\text{Exploit}}$  in terms of the reward, and this deviation accumulates over time into the first term in the regret bound. On the other hand, this deviation also gives an exploration whose intensity is measured by  $\Delta_{\text{Explore}}$ , and the second term captures the regret caused during the concentration of the posterior onto the true environment. In Appendix D.7, we provide some further justifications for these conditions. We also note that the finiteness of  $\Gamma$  is not essential. The KL divergence condition can be modified for a continuous  $\Gamma$  and combined with a covering argument to obtain the bound accordingly.

## 5 NUMERICAL EXPERIMENTS AND DISCUSSIONS

The following figure summarizes the experiments for our model against benchmarks for each task. We note that the prior environment distribution  $\mathcal{P}_\gamma$  takes an infinite support which means the environment  $\gamma$  in the test phase is different from the environment  $\gamma_i$ 's used for training with probability 1. This provides a fair and possibly more challenging setup to examine the ability of the transformer.




Figure 4: The average out-of-sample regret of  $\text{TF}_\delta$  against benchmark algorithms (see details in Appendix E.3) calculated based 100 runs. The numbers in the legend bar are the final regret at  $t = 100$ . The shaded area in the plots indicates the 90% (empirical) confidence interval for the regrets. The prior distribution  $\mathcal{P}_\gamma$  is continuous (infinitely many possible  $\gamma$ ). The problem dimension: number of arms for MAB = 20, dimension of linear bandits = 2, dimension of  $X_t$  for pricing = 6, dimension of  $X_t$  for newsvendor = 4.

We make the following observations on the numerical experiment. The transformer achieves a better performance than the structured algorithms for each task. This result is impressive to us in that the transformer is not replicating/imitating the performance of these benchmark algorithms but it discovers a new and more effective algorithm. This in fact shows the effectiveness of using the optimal actions  $a_t^*$  as the target variable for training the transformer. Moreover, we attribute the advantage of the transformer against the benchmark algorithms to two aspects: prior knowledge from pre-training data and more greedy exploitation. First, the transformer well utilizes the pre-training data and can be viewed as tailored for the pre-training distribution. Comparatively, all the benchmark algorithms are cold-start. Second, the benchmark algorithms are all designed in ways to cater for asymptotic optimality and thus may sacrifice short-term rewards. Comparatively, the transformer behaves more greedy; just like  $\text{Alg}^*$  is a decision function that does not explicitly encourage exploration. And such greedy behavior brings better reward when  $T$  is small such as  $T \leq 100$ .

We provide more numerical experiments to Section C. In Section C.1, we present additional experiments on training dynamics, including an ablation study on  $f$  and  $\kappa$  as well as further results extended from Figure 2 (a). Section C.2 provides more results regarding the alignment of  $\text{TF}_\delta$  and  $\text{Alg}^*$ . Section C.3 compares the performance of  $\text{TF}_\delta$  and  $\text{Alg}^*$ , showing that  $\text{TF}_\delta$  can outperform  $\text{Alg}^*$  in certain scenarios. In Section C.4, we demonstrate that  $\text{TF}_\delta$  can be a solution to model misspecifications. We also evaluate the performance of  $\text{TF}_\delta$  under environments with varying levels of problem complexity (Section C.5), longer testing horizons (Section C.6), and the case of model misspecification when the test environment is sampled from a distribution different from pre-training (i.e., its OOD performance) (Section C.7). Finally, we conduct an ablation study on the model architecture, model size, and task complexity in Section C.8 and Section C.9.

## 6 CONCLUSION

In this paper, we study the pre-trained transformer for sequential decision-making problems. Such problems enable using the optimal action as the target variable for training the transformer. We mathematically formulate the setup of the training pipeline and propose a way to utilize simulation environments to generate pre-training data. We investigate the training and generation aspect of the transformer and identify an OOD issue that is largely ignored by the existing literature. Theoretically, we interpret how the transformer works as an online algorithm through the lens of the Bayes-optimal decision function. The numerical experiments are encouraging in that the trained transformer learns a decision function significantly better than all the benchmark algorithms.

## REFERENCES

- 540  
541  
542 [1] Shipra Agrawal and Navin Goyal. Thompson sampling for contextual bandits with linear  
543 payoffs. In *International conference on machine learning*, pp. 127–135. PMLR, 2013.
- 544  
545 [2] Shipra Agrawal and Randy Jia. Learning in structured mdps with convex cost functions: Im-  
546 proved regret bounds for inventory management. *Operations Research*, 70(3):1646–1664,  
547 2022.
- 548  
549 [3] Anurag Ajay, Yilun Du, Abhi Gupta, Joshua Tenenbaum, Tommi Jaakkola, and Pulkit  
550 Agrawal. Is conditional generative modeling all you need for decision-making? *arXiv preprint*  
551 *arXiv:2211.15657*, 2022.
- 552  
553 [4] Ekin Akyürek, Dale Schuurmans, Jacob Andreas, Tengyu Ma, and Denny Zhou. What  
554 learning algorithm is in-context learning? investigations with linear models. *arXiv preprint*  
555 *arXiv:2211.15661*, 2022.
- 556  
557 [5] Anirudhan Badrinath, Yannis Flet-Berliac, Allen Nie, and Emma Brunskill. Waypoint trans-  
558 former: Reinforcement learning via supervised learning with intermediate targets. *Advances*  
559 *in Neural Information Processing Systems*, 36, 2024.
- 560  
561 [6] Yu Bai, Fan Chen, Huan Wang, Caiming Xiong, and Song Mei. Transformers as statisticians:  
562 Provable in-context learning with in-context algorithm selection. *Advances in neural informa-*  
563 *tion processing systems*, 36, 2024.
- 564  
565 [7] Gah-Yi Ban and Cynthia Rudin. The big data newsvendor: Practical insights from machine  
566 learning. *Operations Research*, 67(1):90–108, 2019.
- 567  
568 [8] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jen-  
569 nifer Wortman Vaughan. A theory of learning from different domains. *Machine learning*,  
570 79:151–175, 2010.
- 571  
572 [9] Omar Besbes and Alp Muharremoglu. On implications of demand censoring in the newsvendor  
573 problem. *Management Science*, 59(6):1407–1424, 2013.
- 574  
575 [10] Omar Besbes and Assaf Zeevi. On the (surprising) sufficiency of linear models for dynamic  
576 pricing with demand learning. *Management Science*, 61(4):723–739, 2015.
- 577  
578 [11] David Brandfonbrener, Alberto Bietti, Jacob Buckman, Romain Laroche, and Joan Bruna.  
579 When does return-conditioned supervised learning work for offline reinforcement learning?  
580 *Advances in Neural Information Processing Systems*, 35:1542–1553, 2022.
- 581  
582 [12] Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece  
583 Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. Sparks of artificial general  
584 intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*, 2023.
- 585  
586 [13] Lili Chen, Kevin Lu, Aravind Rajeswaran, Kimin Lee, Aditya Grover, Misha Laskin, Pieter  
587 Abbeel, Aravind Srinivas, and Igor Mordatch. Decision transformer: Reinforcement learning  
588 via sequence modeling. *Advances in neural information processing systems*, 34:15084–15097,  
589 2021.
- 590  
591 [14] Wei Chu, Lihong Li, Lev Reyzin, and Robert Schapire. Contextual bandits with linear payoff  
592 functions. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence*  
593 *and Statistics*, pp. 208–214. JMLR Workshop and Conference Proceedings, 2011.
- [15] Maxime C Cohen, Ilan Lobel, and Renato Paes Leme. Feature-based dynamic pricing. *Man-*  
*agement Science*, 66(11):4921–4943, 2020.
- [16] Arnoud V Den Boer. Dynamic pricing and learning: historical origins, current research, and  
new directions. *Surveys in operations research and management science*, 20(1):1–18, 2015.
- [17] Arnoud V Den Boer and Bert Zwart. Simultaneously learning and optimizing using controlled  
variance pricing. *Management science*, 60(3):770–783, 2014.

- 594 [18] Jingying Ding, Woonghee Tim Huh, and Ying Rong. Feature-based inventory control with  
595 censored demand. *Manufacturing & Service Operations Management*, 2024.  
596
- 597 [19] Yiming Ding, Carlos Florensa, Pieter Abbeel, and Mariano Phielipp. Goal-conditioned imita-  
598 tion learning. *Advances in neural information processing systems*, 32, 2019.
- 599 [20] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai,  
600 Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly,  
601 et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv*  
602 *preprint arXiv:2010.11929*, 2020.
- 603 [21] Scott Emmons, Benjamin Eysenbach, Ilya Kostrikov, and Sergey Levine. Rvs: What is essen-  
604 tial for offline rl via supervised learning? *arXiv preprint arXiv:2112.10751*, 2021.  
605
- 606 [22] Shivam Garg, Dimitris Tsipras, Percy S Liang, and Gregory Valiant. What can transformers  
607 learn in-context? a case study of simple function classes. *Advances in Neural Information*  
608 *Processing Systems*, 35:30583–30598, 2022.
- 609 [23] Subhashis Ghosal, Jayanta K Ghosh, and Aad W Van Der Vaart. Convergence rates of posterior  
610 distributions. *Annals of Statistics*, pp. 500–531, 2000.  
611
- 612 [24] J Michael Harrison, N Bora Keskin, and Assaf Zeevi. Bayesian dynamic pricing policies:  
613 Learning and earning under a binary prior distribution. *Management Science*, 58(3):570–586,  
614 2012.
- 615 [25] S Hochreiter. Long short-term memory. *Neural Computation MIT-Press*, 1997.  
616
- 617 [26] Woonghee Tim Huh and Paat Rusmevichientong. A nonparametric asymptotic analysis of  
618 inventory planning with censored demand. *Mathematics of Operations Research*, 34(1):103–  
619 123, 2009.
- 620 [27] Woonghee Tim Huh, Ganesh Janakiraman, John A Muckstadt, and Paat Rusmevichientong.  
621 An adaptive algorithm for finding the optimal base-stock policy in lost sales inventory systems  
622 with censored demand. *Mathematics of Operations Research*, 34(2):397–416, 2009.
- 623 [28] Zachary Izzo, Lexing Ying, and James Zou. How to learn when data reacts to your model:  
624 performative gradient descent. In *International Conference on Machine Learning*, pp. 4641–  
625 4650. PMLR, 2021.  
626
- 627 [29] Michael Janner, Qiyang Li, and Sergey Levine. Offline reinforcement learning as one big  
628 sequence modeling problem. *Advances in neural information processing systems*, 34:1273–  
629 1286, 2021.
- 630 [30] Hong Jun Jeon, Jason D Lee, Qi Lei, and Benjamin Van Roy. An information-theoretic analysis  
631 of in-context learning. *arXiv preprint arXiv:2401.15530*, 2024.  
632
- 633 [31] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon  
634 Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural  
635 language models. *arXiv preprint arXiv:2001.08361*, 2020.
- 636 [32] N Bora Keskin and Assaf Zeevi. Dynamic pricing with an unknown demand model: Asymp-  
637 totically optimal semi-myopic policies. *Operations research*, 62(5):1142–1167, 2014.  
638
- 639 [33] Tze Leung Lai and Herbert Robbins. Iterated least squares in multiperiod control. *Advances*  
640 *in Applied Mathematics*, 3(1):50–73, 1982.
- 641 [34] Michael Laskin, Luyu Wang, Junhyuk Oh, Emilio Parisotto, Stephen Spencer, Richie Steiger-  
642 wald, DJ Strouse, Steven Hansen, Angelos Filos, Ethan Brooks, et al. In-context reinforcement  
643 learning with algorithm distillation. *arXiv preprint arXiv:2210.14215*, 2022.
- 644 [35] Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.  
645
- 646 [36] Jonathan Lee, Annie Xie, Aldo Pacchiano, Yash Chandak, Chelsea Finn, Ofir Nachum, and  
647 Emma Brunskill. Supervised pretraining can learn in-context reinforcement learning. *Ad-*  
*vances in Neural Information Processing Systems*, 36, 2024.

- 648 [37] Retsef Levi, Georgia Perakis, and Joline Uichanco. The data-driven newsvendor problem:  
649 New bounds and insights. *Operations Research*, 63(6):1294–1306, 2015.  
650
- 651 [38] Licong Lin, Yu Bai, and Song Mei. Transformers as decision makers: Provable in-context  
652 reinforcement learning via supervised pretraining. *arXiv preprint arXiv:2310.08566*, 2023.  
653
- 654 [39] Celestine Mendler-Düner, Juan Perdomo, Tijana Zrnic, and Moritz Hardt. Stochastic opti-  
655 mization for performative prediction. *Advances in Neural Information Processing Systems*,  
656 33:4929–4939, 2020.
- 657 [40] Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of machine learn-*  
658 *ing*. MIT press, 2018.  
659
- 660 [41] Samuel Müller, Noah Hollmann, Sebastian Pineda Arango, Josif Grabocka, and Frank Hutter.  
661 Transformers can do bayesian inference. In *International Conference on Learning Represen-*  
662 *tations*, 2022. URL <https://openreview.net/forum?id=KSugKcbNf9>.
- 663 [42] Chanwoo Park, Xiangyu Liu, Asuman Ozdaglar, and Kaiqing Zhang. Do llm agents have  
664 regret? a case study in online learning and games. *arXiv preprint arXiv:2403.16843*, 2024.  
665
- 666 [43] Juan Perdomo, Tijana Zrnic, Celestine Mendler-Düner, and Moritz Hardt. Performative pre-  
667 diction. In *International Conference on Machine Learning*, pp. 7599–7609. PMLR, 2020.
- 668 [44] Sheng Qiang and Mohsen Bayati. Dynamic pricing with demand covariates. *arXiv preprint*  
669 *arXiv:1604.07463*, 2016.  
670
- 671 [45] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language  
672 understanding by generative pre-training. 2018.  
673
- 674 [46] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al.  
675 Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- 676 [47] Daniel J Russo, Benjamin Van Roy, Abbas Kazerouni, Ian Osband, Zheng Wen, et al. A tutorial  
677 on thompson sampling. *Foundations and Trends® in Machine Learning*, 11(1):1–96, 2018.  
678
- 679 [48] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez,  
680 Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information*  
681 *processing systems*, 30, 2017.
- 682 [49] Johannes Von Oswald, Eyvind Niklasson, Ettore Randazzo, João Sacramento, Alexander  
683 Mordvintsev, Andrey Zhmoginov, and Max Vladymyrov. Transformers learn in-context by  
684 gradient descent. In *International Conference on Machine Learning*, pp. 35151–35174. PMLR,  
685 2023.  
686
- 687 [50] Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48.  
688 Cambridge university press, 2019.
- 689 [51] Hanzhao Wang, Kalyan Talluri, and Xiaocheng Li. On dynamic pricing with covariates. *arXiv*  
690 *preprint arXiv:2112.13254*, 2021.  
691
- 692 [52] Yueh-Hua Wu, Xiaolong Wang, and Masashi Hamaya. Elastic decision transformer. *Advances*  
693 *in Neural Information Processing Systems*, 36, 2024.  
694
- 695 [53] Sang Michael Xie, Aditi Raghunathan, Percy Liang, and Tengyu Ma. An explanation of in-  
696 context learning as implicit bayesian inference. *arXiv preprint arXiv:2111.02080*, 2021.
- 697 [54] Fred Weiyang Zhang, Jiaxin Ye, and Zhuoran Yang. In-context multi-armed bandits via su-  
698 pervised pretraining. In *NeurIPS 2023 Foundation Models for Decision Making Workshop*,  
699 2023.  
700
- 701 [55] Huanan Zhang, Xiuli Chao, and Cong Shi. Closing the gap: A learning algorithm for lost-sales  
inventory systems with lead times. *Management Science*, 66(5):1962–1980, 2020.

- [56] Yufeng Zhang, Fengzhuo Zhang, Zhuoran Yang, and Zhaoran Wang. What and how does in-context learning learn? bayesian model averaging, parameterization, and generalization. *arXiv preprint arXiv:2305.19420*, 2023.
- [57] Qinqing Zheng, Amy Zhang, and Aditya Grover. Online decision transformer. In *international conference on machine learning*, pp. 27042–27059. PMLR, 2022.

## A RELATED WORKS

### A.1 LITERATURE REVIEW

In this section, we discuss more related works which complement our discussions in Section 1. In particular, we elaborate the literature in three streams.

#### **Understanding transformer and in-context learning.**

A remarkable characteristic of pre-trained transformers is their ability to perform in-context learning (ICL). Once pre-trained on a vast corpus, transformers can solve new tasks with just a few examples, without updating their parameters. ICL has captured the attention of the theoretical machine learning community, leading to considerable efforts into understanding ICL from different theoretical perspectives [53, 4, 49, 38]. A stream of literature that is related to our work explains the transformer’s behavior as Bayesian inference. See [53, 56] for the Bayesian argument in regression context, [41] for approximating a large set of posteriors, and [36] for the Bayesian analysis in decision-making. Another stream of literature aims to develop generalization bound or other types of convergence results. See [30, 6] for such results in the regression context and [38] in the decision-making context.

#### **Pre-trained transformer for RL.**

The earliest effort of applying transformers to reinforcement learning lies in the area of offline RL [29, 13]. By autoregressively maximizing the likelihood of trajectories in the offline dataset, this paradigm essentially converts offline RL to a supervised learning problem. When evaluating the policy, the actions are sampled according to the likelihood of PT conditioned on the so called “return-to-go”, which is commonly chosen to be the cumulative reward of “good” trajectories in the dataset. By doing so, PT retrieves information from trajectories in training data with similar return-to-go values and performs imitation learning to high-reward offline trajectories. Based on this observation, several works further examine the validity of return-to-go conditioning [21, 3, 11] and propose alternative methods that improve this approach or extend to other settings [34, 5, 57, 52]. For example, one of the drawbacks of the return-to-go conditioning is that it often fails in trajectory-stitching when offline data comes from sub-optimal policies. One recent work [42] finds that pre-trained LLMs, although trained by text data, can behave like a follow-the-perturbed-leader (FTPL) algorithm in bandits and games. It further introduces a regret-based training loss, through which transformers can be trained in an unsupervised way.

#### **Sequential decision-making.**

We carry out experiments involving some stylized sequential decision-making models including dynamic pricing and inventory management. Dynamic pricing entails setting prices to discern the underlying revenue function, usually defined as the product of demand and price [16]. Algorithms that purely exploit based on historical data tend to be suboptimal and have been found converging to non-optimal prices with a positive probability [33, 32]. Therefore, random exploration is essential for a more accurate understanding of the demand function and to avoid settling on suboptimal prices. See [24, 17, 10, 15] for algorithms that tackle this exploration-exploitation trade-off and achieving near-optimal regret under different settings.

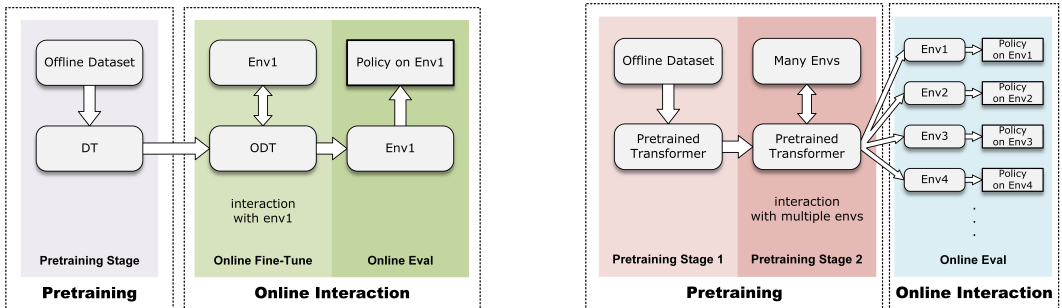
For inventory management, demand estimation is different from that in dynamic pricing due to the censored demand [26]. When demand exceeds inventory levels, the unmet demand is not observed, leading to lost sales. The optimal decision-making policy must balance exploration and exploitation, which involves occasionally setting inventory levels higher than usual to gather more demand data, but not so high that the additional costs of exploration become prohibitive. Regarding near-optimal

756 policies, various works [27, 55, 2] have developed algorithms that achieve an  $O(\sqrt{T})$  regret bound  
 757 with different assumptions about lead times.  
 758

759 A.2  $TF_{\hat{\theta}}$  V.S. ONLINE DECISION TRANSFORMER  
 760

761 In this section, we compare  $TF_{\hat{\theta}}$  with online decision transformer (ODT) [57]. As shown in Figure  
 762 5, although both approaches involve a pre-training phase and an (online) interaction phase, they are  
 763 indeed very different and designed to address different problems. Specifically:  
 764

- 765 • ODT interacts with a specific testing environment to fine-tune the policy, whereas  $TF_{\hat{\theta}}$   
 766 (with Algorithm 1) interacts with different environments and does the learning in the pre-  
 767 training stage. As a result, ODT is tailored to optimize performance in a single test environ-  
 768 ment, while  $TF_{\hat{\theta}}$  learns various policies across different environments during pre-training,  
 769 enabling it to perform well in online evaluations across different environments.
- 770 • ODT has been shown to underperform in stochastic environments. Although some RL environ-  
 771 ments—such as Atari and Gym—are relatively less stochastic, sequential decision-  
 772 making often involves more stochasticity. This is why ODT has rarely been applied  
 773 to sequential decision-making problems like bandit problems. More specifically, ODT,  
 774 along with its precursor Decision Transformer (DT) [13], belongs to the class of “return-  
 775 conditioned supervised learning” methods, which are not suitable even for offline reinforc-  
 776 ement learning in stochastic environments. In such cases, the transformers (ODT or DT) is  
 777 prone to learning from a “survival bias”, as demonstrated in [11], leading to suboptimal  
 778 performance.



789 (a) Online Decision Transformer (ODT)

790 (b)  $TF_{\hat{\theta}}$  with Algorithm 1 (ours)

791 Figure 5: Comparison between training frameworks of Online Decision Transformer (ODT) [57]  
 792 and  $TF_{\hat{\theta}}$ . The online interactions of ODT is in the (single) test environment to explore a (single)  
 793 good policy in the same environment, while our proposed training framework (Algorithm 1) has  
 794 online interactions in multiple environments to mitigate the OOD issue in the pre-training data (as  
 795 discussed in Section 3), and  $TF_{\hat{\theta}}$  can handle different test environments.

796 We further conduct an experiment comparing our framework with ODT as shown in Figure 6 to  
 797 highlight their differences. Specifically, we use a noiseless linear bandits task (i.e., no reward noise)  
 798 with dimension  $d = 2$  and a horizon of  $T = 20$ , where the test environments are entirely unseen  
 799 during the training phase, including the online fine-tuning phase for ODT. It is important to empha-  
 800 size that this problem setting differs from the one for which ODT is designed (as shown in Figure  
 801 5(a), where the test environment is identical to the environment used during online fine-tuning), and  
 802 aligns with the setting we study in this work. We set fair configurations for training ODT and  $TF_{\hat{\theta}}$   
 803 (which will be specified later). As expected, Figure 6 demonstrates that ODT fails in this setting,  
 804 performing no better than a random policy that uniformly samples actions.

805 **Setup:** for both ODT and  $TF_{\hat{\theta}}$ , we set the transformer architecture with 4 layers, 4 heads, 512 em-  
 806 bedding dimensions, and a 20-length context window, which are the same as in [57]. For the offline  
 807 dataset, we sample 100,000 environments by the same method in our paper, and then use Thompson  
 808 sampling to generate trajectories for each environment. And for the online interactions, the envi-  
 809 ronments are sampled from the same distribution as the offline data. The testing environments are  
 still sampled from the same distribution. We set the offline training phase (pre-training stage for

810  
811  
812  
813  
814  
815  
816  
817  
818  
819  
820  
821  
822  
823  
824  
825  
826  
827  
828  
829  
830  
831  
832  
833  
834  
835  
836  
837  
838  
839  
840  
841  
842  
843  
844  
845  
846  
847  
848  
849  
850  
851  
852  
853  
854  
855  
856  
857  
858  
859  
860  
861  
862  
863

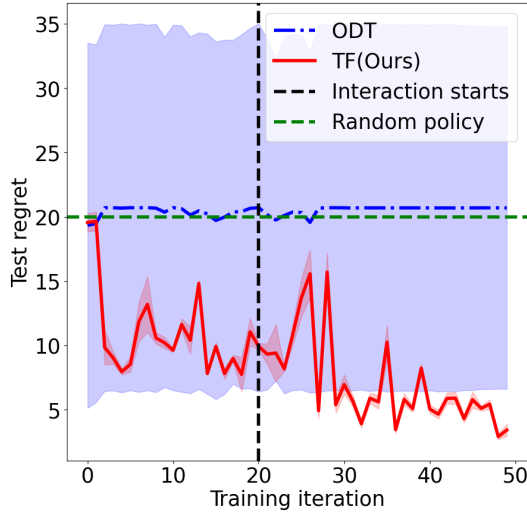


Figure 6: We test the ODT and our method in a same linear bandits task, where test environments are unseen before the evaluation phase. The ODT fails in the test environments, which are different from the online fine-tuning phase’s.

ODT and stage 1 for  $TF_{\hat{\theta}}$ ) with 20 iterations (each iteration has 500 updates/gradient descents on parameters), and set the online interaction phase (online fine-tuning for ODT and stage 2 for  $TF_{\hat{\theta}}$ ) with 30 iterations (each iteration has 50 updates/gradient descents on parameters). The other training parameters (e.g., learning rate) are set the same as in [57] except the return-to-go (RTG): we set the eval RTG as 55, which is the maximum RTG observed in the offline dataset, and online RTG as 110, which doubles the eval RTG as used in [57].

## B PROBLEM EXAMPLES FOR THE GENERAL SETUP IN SECTION 2

Throughout the paper, we make the following boundedness assumption.

**Assumption B.1.** Assume the context space  $\mathcal{X}$ , the action space  $\mathcal{A}$ , and the observation space  $\mathcal{O}$  is bounded (under the Euclidean norm) by  $D$ .

We provide a few examples as special cases of the framework in Section 2. We hope these examples make the actions, observations, and context in the general setup more tangible.

- Stochastic multi-armed bandits : There is no context,  $X_t = \text{null}$  for all  $t$ . The action  $a_t \in \mathcal{A} = \{1, 2, \dots, k\}$  denotes the index of the arm played at time  $t$ . The observation (also the random reward) is generated following the distribution  $P_{a_t}$ . The parameter  $\gamma$  encapsulates the distributions  $P_1, \dots, P_k$ .
- Linear bandits: There is no context,  $X_t = \text{null}$  for all  $t$ . The action  $a_t \in \mathcal{A} \in \mathbb{R}^d$  is selected from some pre-specified domain  $\mathcal{A}$ . The random reward  $R(X_t, a_t) = w^\top a_t + \epsilon_t$  where  $\epsilon_t$  is some noise random variable and  $w$  is a vector unknown to the agent. The observation  $o_t = R(X_t, a_t)$  and the expected reward  $r(X_t, a_t) = \mathbb{E}[R(X_t, a_t)|X_t, a_t] = w^\top a_t$ . The parameter  $\gamma$  encapsulates the vector  $w$  and the noise distribution.
- Dynamic pricing: The context vector  $X_t$  describes the market-related information at time  $t$ . Upon the reveal of  $X_t$ , the agent takes the action  $a_t$  as the pricing decision. Then the agent observes the demand  $d_t = d(X_t, a_t) + \epsilon_t$  where  $\epsilon_t$  is some noise random variable and  $d(\cdot, \cdot)$  is a demand function unknown to the agent. The random reward  $R_t(X_t, a_t) = d_t \cdot a_t$  is the revenue collected from the sales at time  $t$  (which equals the demand times the price), and the observation  $o_t = d_t$ . The expected reward  $r(X_t, a_t) = \mathbb{E}[R(X_t, a_t)|X_t, a_t] = d(X_t, a_t) \cdot a_t$ . The parameter  $\gamma$  governs the generation of  $X_t$ ’s and also contains information about the demand function  $d$  and the noise distribution.



- Newsvendor problem: As the dynamic pricing problem, the context vector  $X_t$  describes the market-related information at time  $t$ . Upon its reveal, the agent takes the action  $a_t$  which represents the number of inventory prepared for the sales at time  $t$ . Then the agent observes the demand  $d_t = d(X_t, a_t) + \epsilon_t$  where  $\epsilon_t$  is some noise random variable and  $d(\cdot, \cdot)$  is a demand function unknown to the agent. The observation  $o_t = d_t$ , and the random reward is the negative cost  $R(X_t, a_t) = -h \cdot (a_t - d_t)^+ - l \cdot (d_t - a_t)^+$  where  $(\cdot)^+$  is the positive-part function,  $h$  is the left-over cost, and  $l$  is the lost-sale cost. The expected reward  $r(X_t, a_t) = \mathbb{E}[R(X_t, a_t)|X_t, a_t]$  with the expectation taken with respect to  $\epsilon_t$ . The parameter  $\gamma$  governs the generation of  $X_t$ 's and encodes the demand function  $f$  and the noise distribution.

For the general problem setup in Section 2, the following proposition relates the pre-training and the regret through the lens of loss function.

**Proposition B.2** (Surrogate property). *We say the loss function  $l(\cdot, \cdot)$  satisfies the surrogate property if there exists a constant  $C > 0$  such that*

$$\text{Regret}(f; \gamma) \leq C \cdot L(f; \gamma)$$

holds for any  $f$  and  $\gamma$ . We have

- The cross-entropy loss satisfies the surrogate property for the multi-armed bandits problem.
- The squared loss satisfies the surrogate property for the dynamic pricing problem.
- The absolute loss satisfies the surrogate property for the linear bandits and the newsvendor problem.

Proposition B.2 gives a first validity of the supervised pre-training approach. The surrogate property is a property of the loss function and it is not affected by the distribution or the decision function  $\text{TF}_\theta$ . Specifically, it upper bounds the regret with the action prediction loss. It is also natural in that a closer prediction of the optimal action gives a smaller regret. One implication is that the result guides the choice of the loss function for different underlying problems.

## B.1 PROOF OF PROPOSITION B.2

To prove Proposition B.2, we assume the following mild conditions. First, for all these problems, the noise random variable  $\epsilon_t$  has mean 0, is i.i.d. across time, and is independent of the action  $a_t$ . Second, for the linear bandits problem, the  $L_\infty$  norm of the unknown parameter  $w$  is bounded by  $D$ . Third, for the dynamic pricing problem, we assume the demand function is differentiable with respect to the action coordinate, and both the absolute value of the first- and second-order derivatives is bounded by  $D$ . Further details regarding the problem formulations can be found in Appendix E.2.

*Proof.* Since both the regret  $\text{Regret}(f; \gamma)$  and the prediction error  $L(f; \gamma)$  are the sums of single-step regret and prediction error over the horizon, it is sufficient to prove the single-step surrogate property of the loss function. Therefore, in the following proof, we will focus on proving the single-step surrogate property and omit the subscript  $t$  for simplicity.

- For the multi-arm bandits problem, without loss of generality, we assume  $\mathcal{A} = \{a_1, \dots, a_J\}$  and the optimal arm is  $a_1$ . For  $1 \leq i \leq J$ , let  $\mu_j$  be the mean of the reward distribution for arm  $j$ , and define  $\Delta_j = \mu_1 - \mu_j$  as the reward gap for arm  $j$ . Define  $\Delta_{\max} = \max\{\Delta_1, \dots, \Delta_J\}$  as the maximum action gap. Under cross-entropy loss, we denote the output distribution of decision function  $f$  is  $(p_1, \dots, p_J)$ . Then the single-step regret can be bounded by:

$$\mathbb{E}[r(X, a^*) - r(X, f(H)) | H] = \mu_1 - \sum_{j=1}^J p_j \mu_j = \sum_{j=2}^J p_j \Delta_j \leq \Delta_{\max}(1 - p_1) \leq -\Delta_{\max} \cdot \log(p_1),$$

where the last step is by the inequality  $1 - x \leq -\log(x)$  for  $0 < x \leq 1$ . Then by the definition of cross-entropy loss, we finish the proof.

- For the dynamic pricing problem, we have

$$r(X, a^*) - r(X, f(H)) = \nabla_a r(X, a^*)(a^* - f(H)) + \nabla_a^2 r(X, \tilde{a}) |f(H) - a^*|^2 \leq D |f(H) - a^*|^2$$

where  $\tilde{a}$  is in the line of  $f(H)$  and  $a^*$ . Here, the first step is by the Taylor expansion, and the second step is by the first order condition of  $a^*$  and the assumption. Then by noting  $|f(H) - a^*|^2$  is the square loss, we finish the proof.

- For the linear bandits problem, we have

$$r(X, a^*) - r(X, f(H)) = w^\top a^* - w^\top f(H) \leq \|w\|_\infty \|f(H) - a^*\|_1 \leq D \|f(H) - a^*\|_1,$$

where the second step is by the Holder's inequality and the last step is by the assumption. Then by noting  $\|f(H) - a^*\|_1$  is the absolute loss, we finish the proof.

- For the newsvendor problem, we have

$$\begin{aligned} r(X, a^*) - r(X, f(H)) &\leq \max\{h, l\} \mathbb{E}_X [|a^* - d(X) - f(H) + d(X)|] \\ &= \max\{h, l\} |a^* - f(H)| \end{aligned}$$

where the first line is because the (random) reward function is  $\max\{h, l\}$ -Lipschitz in  $a - d(X)$ . Then by noting  $|f(H) - a^*|$  is the absolute loss, we finish the proof.

□

## C MORE EXPERIMENTS

### C.1 MORE EXPERIMENTS ON TRAINING DYNAMICS

#### C.1.1 ABLATION STUDY ON $f$ AND $\kappa$

In this section, we present ablation studies to explore the impact of two key factors: the decision function  $f$  used for generating pre-training data, and the mix ratio  $\kappa$  applied in Algorithm 1. Figure 7 summarizes the results, where we pre-train and test models in multi-armed bandits problems defined in Appendix E.2, and only tune  $f$  or  $\kappa$  during the pre-training and keep all other parameters the same as in Appendix E.1.2.

- **Effect of  $f$ :** We evaluate the influence of different decision functions  $f$  on the performance of the transformer in a multi-armed bandit task. Three types of decision functions are considered: the one we designed, as defined in Appendix E.1.2, the UCB algorithm for multi-armed bandits [35], and a 50/50 mixture of both. The training dynamics for each function are shown in Figure 7a. We observe that all three approaches achieve similar final performance once training ends. While it is possible that some decision functions could lead to faster convergence or better performance after pre-training, we did not encounter any training failures due to the choice of  $f$ .
- **Effect of  $\kappa$ :** We also conduct an experiment to numerically investigate the impact of different mix ratios  $\kappa$ . The results are presented in Figure 7b. We test four values of  $\kappa$  (10%, 33%, 66%, 90%), keeping all other hyperparameters constant for a multi-armed bandit task. The findings show that transformer's performance remains robust across different values of  $\kappa$ , except when  $\kappa = 0.9$ , where performance declines. This supports our earlier discussion in Algorithm 1, highlighting the importance of mixing a non-trivial amount of self-generated data (i.e., non-trivial value for  $1 - \kappa$ ).

#### C.1.2 ADDITIONAL RESULTS ON FIGURE 2 (A)

In addition to Figure 2 (a), we further evaluate the effectiveness of Algorithm 1 across various tasks and architectures, as illustrated in Figure 8.

Figure 8 presents comparisons of the transformer's testing regret (see definition and implementation in Appendix E.1) with and without self-generated data (i.e., using Algorithm 1 or not) during the training, colored by blue and orange respectively. The comparisons are conducted on (a) linear

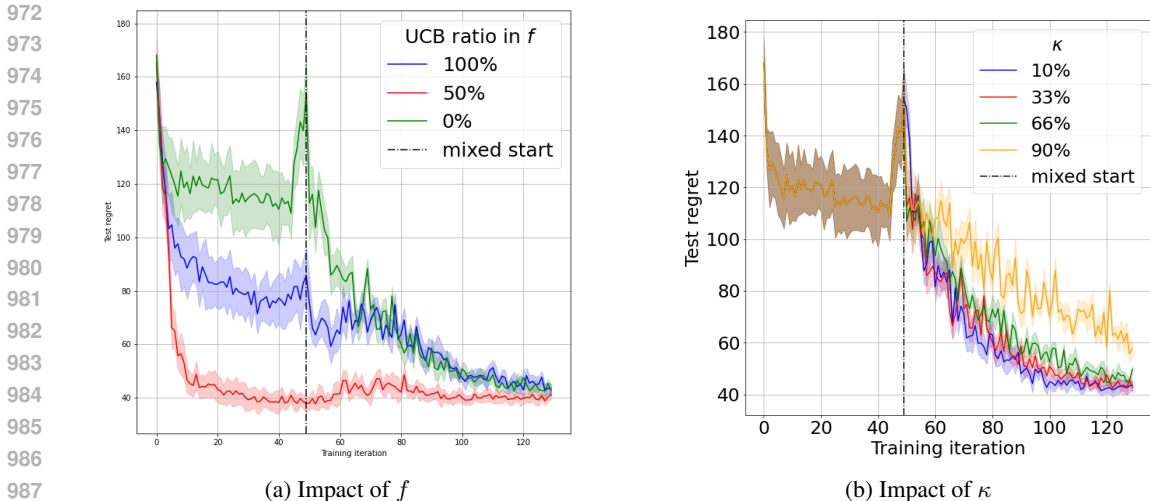


Figure 7: (a) The effect of  $f$  (decision function for generating the pre-training data), where we mix the UCB algorithm [35] and decision function designed in Appendix E.1.2 with different ratios to create the pre-training data. (b) The effect of  $\kappa$  (ratio of samples generated by  $f$  in the mixed training phase).

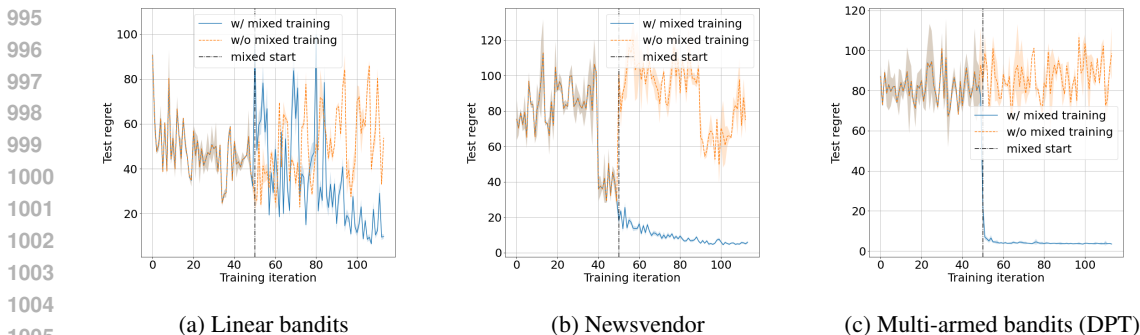


Figure 8: Effectiveness of Algorithm 1. The values are averaged out-of-sample regrets based on 128 runs, where the shaded area indicates the standard deviation.

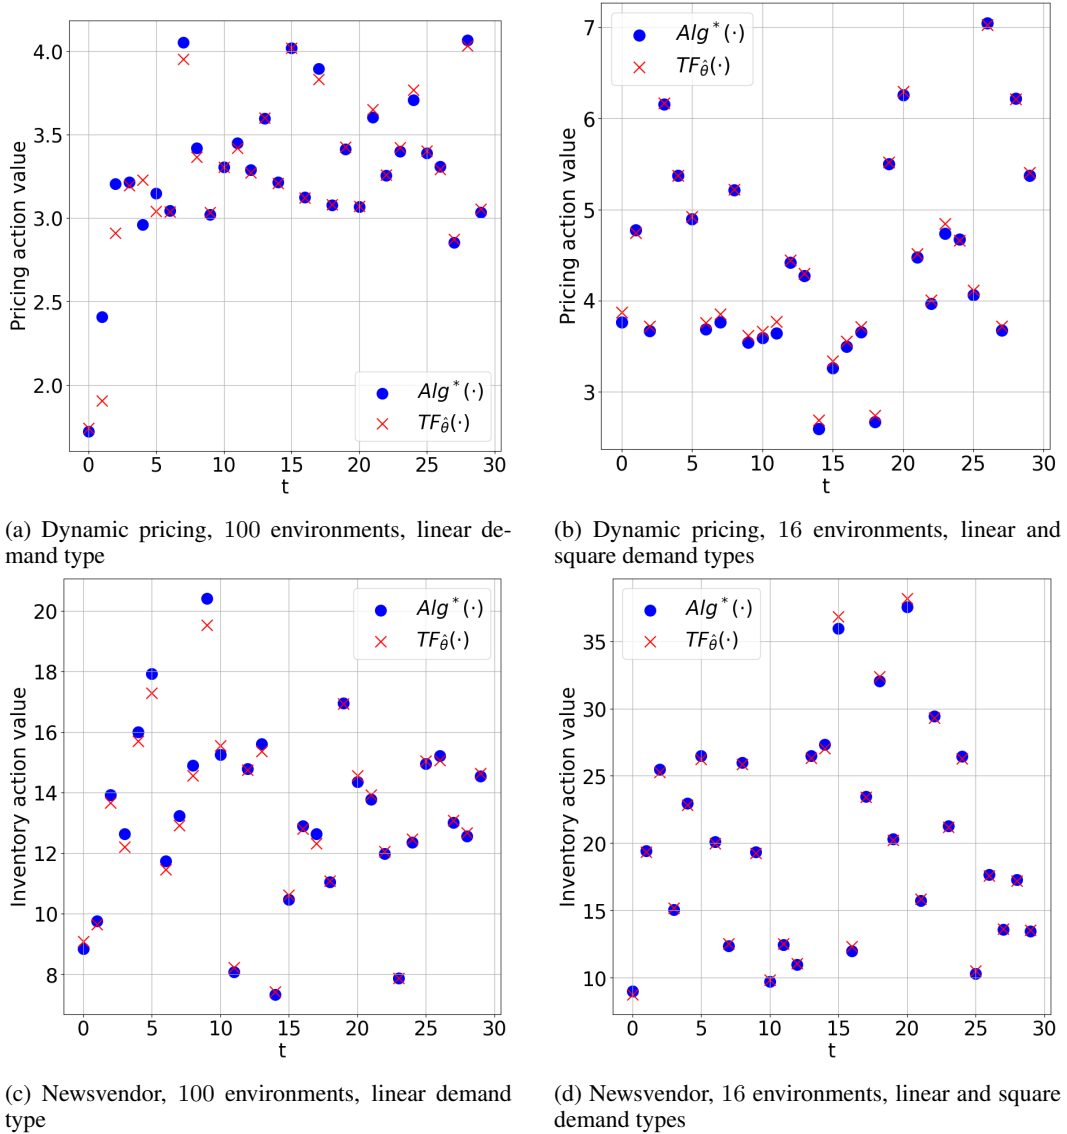
bandits and (b) newsvendor tasks. Additionally, (c) includes a similar comparison conducted in a model with different architecture, the Decision Pre-trained Transformer (DPT) [36], on the multi-armed bandits task. The results consistently demonstrate the same pattern observed in Figure 2 (a): incorporating transformer-generated data significantly reduces the testing regret loss at each iteration  $m > 50$  compared to not using it. This shows the effectiveness of Algorithm 1 (i) across different tasks and (ii) for various model architectures.

**Setup.** In Figure 8 (a) and (b), we consider two tasks: a linear bandits task (with 2-dimensional action space) and a newsvendor task (with 4-dimensional contexts). Both tasks have infinite support of the prior environment distribution  $\mathcal{P}_\gamma$  (i.e., infinite possible environments in both the pre-training and testing). The details of the tasks can be found in Appendix E.2. For each task, we independently train two transformer models by Algorithm 1. For both models, the training parameters are identical except for the  $M_0$ : the first curve (blue, thick line) uses  $M_0 = 50$ , while the second curve (orange, dashed line) uses  $M_0 = 130$ , meaning no transformer-generated data is utilized during training. All other parameters follow the configuration detailed in Appendix E.1. The figure shows the mean testing regret at each training iteration across 128 randomly sampled environments from  $\mathcal{P}_\gamma$ , with the shaded areas representing the standard deviation. We conduct the same experiment on the multi-armed bandits task (with infinite possible environments and 20 arms) for a different architecture for

1026 sequential decision making, Decision Pre-trained Transformer (DPT) [36], in Figure 8 (c). We refer  
 1027 to the original paper for details on the DPT’s architecture.

1028  
 1029 C.2 MORE EXPERIMENTS ON MATCHINGS OF  $TF_{\hat{\theta}}$  AND  $Alg^*$   
 1030

1031 This section provides more results regarding the matching of  $TF_{\hat{\theta}}$  and  $Alg^*$  across different tasks:  
 1032 Figure 9 provides more examples comparing decisions made by  $TF_{\hat{\theta}}$  and  $Alg^*$  at sample path levels,  
 1033 while Figure 10 compares them at population levels.  
 1034



1072 Figure 9: Examples to compare the actions from the transformer  $TF_{\hat{\theta}}$  and the optimal decision  
 1073 function  $Alg^*$ .  
 1074

1075 We plot the actions generated from  $TF_{\hat{\theta}}$  and  $Alg^*$  on the same environment with the same contexts  
 1076 in Figure 9 across dynamic pricing and newsvendor tasks. Each subfigure has a different number  
 1077 of possible environments or possible demand types (which are shared in both the pre-training and  
 1078 testing phases). The population-level differences  $\{TF_{\hat{\theta}}(H_t) - Alg^*(H_t)\}_{t=1}^T$  are presented in Figure  
 1079 10. We observe that across different tasks: (1)  $TF_{\hat{\theta}}$  nearly matches  $Alg^*$ , but (2) the matchings are  
 not perfect. The actions from  $TF_{\hat{\theta}}$  and  $Alg^*$  are not exactly the same, and their differences increase

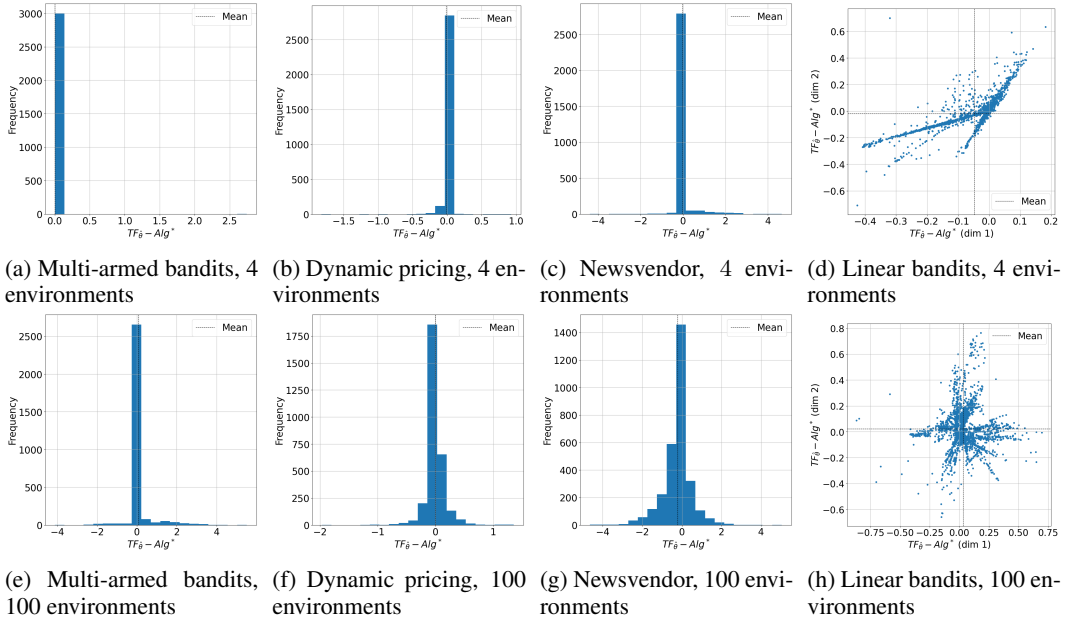


Figure 10:  $\text{TF}_{\hat{\theta}}(H_t) - \text{Alg}^*(H_t)$  across different tasks with various numbers of possible environments. More possible environments lead to a harder decision-making problem.

as the underlying problems become more complex, as seen by comparing the first and second rows in Figure 10. This latter observation numerically supports the conditions in Proposition 4.4.

**Setup.** For Figure 9, each subfigure is based on a sampled environment with a sampled sequence of contexts  $\{X_t\}_{t=1}^{30}$  from the corresponding task. For tasks in (a),(c), the data generation process follows the description detailed in Appendix E.2, and the architecture of  $\text{TF}_{\hat{\theta}}$  and the definition of  $\text{Alg}^*$  can be found in Appendix E.1 and Appendix E.3 respectively. For tasks in (b), (d), there are 16 environments included in the support of  $\mathcal{P}_{\gamma}$ , where 8 environments have demand functions of the linear type and the other 8 have demand functions of the square type (see Appendix E.2 for the definitions of these two types). Both the pre-training and testing samples are drawn from  $\mathcal{P}_{\gamma}$ , i.e., from these 16 environments. The remaining setups for these tasks follow the methods provided in Appendix E.2.

For each task shown in Figure 10, we generate 100 sequences of  $\{\text{TF}_{\hat{\theta}}(H_t) - \text{Alg}^*(H_t)\}_{t=1}^{100}$ , i.e., each subfigure is based on 10000 samples of  $\text{TF}_{\hat{\theta}}(H_t) - \text{Alg}^*(H_t)$ . For dynamic pricing and newsvendor, actions are scalars and thus  $\text{TF}_{\hat{\theta}}(H_t) - \text{Alg}^*(H_t)$  is also a scalar; For multi-armed bandits, the action space is discrete and we use the reward difference associated with the chosen actions (by  $\text{TF}_{\hat{\theta}}(H_t)$  or  $\text{Alg}^*(H_t)$ ) as the value of  $\text{TF}_{\hat{\theta}}(H_t) - \text{Alg}^*(H_t)$ . We use histograms to summarize the samples from them. For linear bandits, the actions are 2-dimensional and thus we use the scatters to show  $\text{TF}_{\hat{\theta}}(H_t) - \text{Alg}^*(H_t)$  in Figure 10 (d) or (h).

### C.3 REGRET COMPARISON WITH $\text{ALG}^*$

In this subsection, we study the behavior of  $\text{TF}_{\hat{\theta}}$  by comparing regret performances on  $\text{TF}_{\hat{\theta}}$  and  $\text{Alg}^*$ , which utilizes the posterior distribution  $\mathcal{P}(\gamma|H)$  of environments. As shown in Section 4,  $\text{Alg}^*$ 's are defined as the Bayes-optimal decision functions.

Algorithm 2 presents a general framework for the Bayes-optimal decision functions/algorithms in Example 4.2: posterior averaging, sampling, and median. For simplicity of notation, we omit the subscript  $t$  and only consider a finite space  $\Gamma$  of environments, while the algorithms can be easily extended to the infinite case. For the posterior median algorithm, we assume the action space  $\mathcal{A} \subseteq \mathbb{R}$ .

Figure 11 presents the testing regret and action suboptimality for  $\text{TF}_{\hat{\theta}}$  compared to the Bayes-optimal decision functions/algorithms. Generally, all algorithms exhibit good regret performance. However,

**Algorithm 2** Posterior averaging/sampling/median at time  $t$ 


---

**Require:** Posterior probability  $\mathcal{P}(\gamma|H)$  and optimal action  $a_{\gamma}^*$  for each environment  $\gamma \in \Gamma$ , algorithm  $\text{Alg} \in \{\text{posterior averaging, posterior sampling, posterior median}\}$

```

%% Posterior averaging
1: if Alg=posterior averaging then
2:    $a = \sum_{\gamma \in \Gamma} \mathcal{P}(\gamma|H) \cdot a_{\gamma}^*$ 
%% Posterior sampling
3: else if Alg=posterior sampling then
4:    $a = a_{\tilde{\gamma}}^*$ , where  $\tilde{\gamma} \sim \mathcal{P}(\cdot|H)$ 
%% Posterior median
5: else if Alg=posterior median then
6:   Sort and index the environments as  $\{\gamma_i\}_{i=1}^{|\Gamma|}$  such that the corresponding optimal actions are
   in ascending order:  $a_{\gamma_1}^* \leq a_{\gamma_2}^* \leq \dots \leq a_{\gamma_{|\Gamma|}}^*$ .
7:   Choose  $a = \min \left\{ a_{\gamma_i}^* \mid \sum_{i'=1}^i \mathcal{P}(\gamma_{i'}|H) \geq 0.5 \right\}$ 
8: end if
9: Return:  $a$ 

```

---

we observe that  $\text{TF}_{\hat{\theta}}$  is (i) superior to posterior sampling in multi-armed bandits, as shown in (a), and (ii) outperforms posterior averaging in linear bandits, as shown in (c).

Specifically, (b) shows that posterior sampling has higher action suboptimality during the initial time steps compared to  $\text{TF}_{\hat{\theta}}$ . This suggests that the exploration inherent in the sampling step of posterior sampling can introduce additional regret, which might be unnecessary for simpler problems. Conversely, (d) indicates that the decisions from posterior averaging do not converge to the optimal action  $a_i^*$  as quickly as those from  $\text{TF}_{\hat{\theta}}$  and posterior sampling. This suggests that posterior averaging may be too greedy, thus failing to sufficiently explore the environment. This second observation also numerically supports Proposition 4.3.

These observations illustrate that  $\text{TF}_{\hat{\theta}}$  can outperform Bayes-optimal decision functions in certain scenarios. This further demonstrates that the advantage of  $\text{TF}_{\hat{\theta}}$  over benchmark algorithms is not solely due to its use of prior knowledge about the task. Rather,  $\text{TF}_{\hat{\theta}}$  discovers a new decision rule that achieves better short-term regret than oracle posterior algorithms. It can be more greedy than posterior sampling while exploring more than posterior averaging.

**Setup.** We consider a multi-armed bandits task (with 20 arms) and linear bandits task (with 2-dimensional actions) with 4 environments. The results are based on 100 runs. The details of tasks with their specific oracle posterior algorithms are provided in Appendix E.2 and Appendix E.3.

#### C.4 $\text{TF}_{\hat{\theta}}$ AS A SOLUTION TO MODEL MISSPECIFICATION

Most sequential decision-making algorithms, including all the benchmark algorithms in our experiments, rely on structural or model assumptions about the underlying tasks. For instance, demand functions in pricing and newsvendor problems are typically assumed to be linear in context. Applying these algorithms in misspecified environments, where these assumptions do not hold, can lead to degraded performance.

In contrast,  $\text{TF}_{\hat{\theta}}$  offers a potential solution to model misspecifications. By generating pre-training samples from all possible types of environments (e.g., newsvendor models with both linear and non-linear demand functions),  $\text{TF}_{\hat{\theta}}$  leverages its large capacity to make near-optimal decisions across different types. Figure 12 demonstrates the performance of  $\text{TF}_{\hat{\theta}}$  trained on tasks with two types of demand functions: linear and square (definitions provided in Appendix E.2). It compares  $\text{TF}_{\hat{\theta}}$  with benchmark algorithms designed solely for linear demand cases. The performance is tested in environments where demand functions (i) can be either type with equal probability (first column); (ii) are strictly square type (second column); and (iii) are strictly linear type (third column). The benchmark algorithms face model misspecification issues in the first two cases.

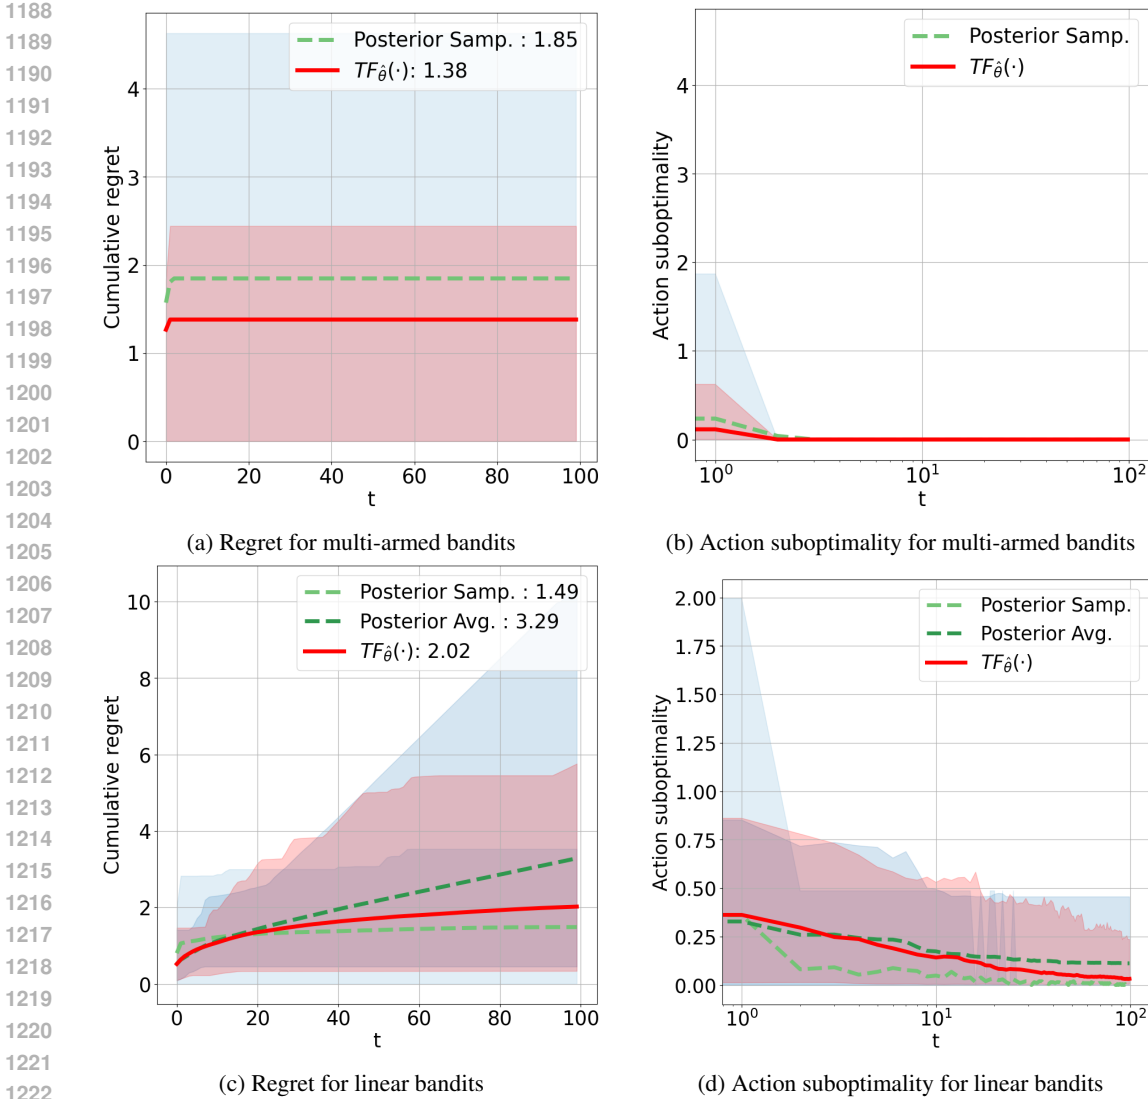


Figure 11: Performances of  $TF_{\hat{\theta}}$  and Bayes-optimal decision functions. The numbers in the legend bar are the final regret at  $t = 100$  and the shaded areas indicate the 90% (empirical) confidence intervals.

For both pricing and newsvendor tasks,  $TF_{\hat{\theta}}$  consistently outperforms all benchmarks across the three scenarios, especially in the first two cases. This result highlights the potential of  $TF_{\hat{\theta}}$  to effectively handle model misspecifications.

**Setup.** For both the pricing and newsvendor tasks, during the pre-training phase of  $TF_{\hat{\theta}}$ , the pre-training data are generated from two types of demand functions: linear and square, with each type having a half probability. The support of  $\mathcal{P}_{\gamma}$  is infinite, meaning the parameters in the demand functions are not restricted. The context dimension is 6 for the pricing task and 4 for the newsvendor task. The generating details and definitions of these two types of demands are provided in Appendix E.2. Additionally, when generating data using  $TF_{\theta_m}$  in Algorithm 1, the sampled environment also has an equal probability of having a linear or square demand type.

Each subfigure in Figure 12 is based on 100 runs. We consider three different cases during testing: (i) the environment can have either the square or linear demand type (each with half probability); (ii) the environment can only have the linear demand type; and (iii) the environment can only have the square demand type. In the first case, the testing environment is sampled in the same way as the

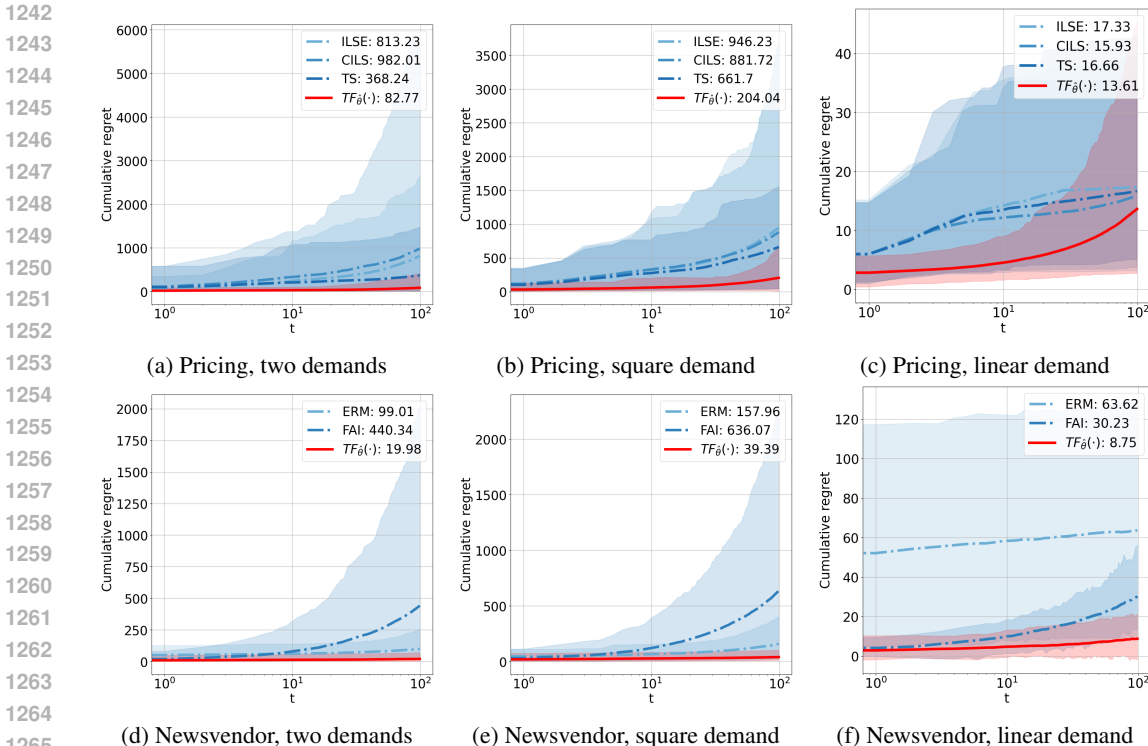


Figure 12: The average out-of-sample regret on dynamic pricing (first row) and newsvendor (second row) with two possible demand types: linear and square. The numbers in the legend bar are the final regret at  $t = 100$  and the shaded areas indicate the 90% (empirical) confidence intervals. The details of benchmarks can be found in Appendix E.3.

pre-training environment, while in the other two cases, we only include the sampled environments with the specified demand type until the total testing samples reach 100.

### C.5 ADDITIONAL RESULTS ON $TF_{\hat{\theta}}$ PERFORMANCES

This subsection presents further results on the performance of  $TF_{\hat{\theta}}$ . Specifically, we compare  $TF_{\hat{\theta}}$  with benchmark algorithms on (i) simple tasks, which include only 4 possible environments in both pre-training and testing samples (Figure 13), and (ii) more complex tasks, which involve either 100 possible environments (Figure 14 (a), (b)) or 16 environments with two possible types of demand functions (Figure 14 (c), (d)).


Figures 13 and 14 illustrate the consistently and significantly superior performance of  $TF_{\hat{\theta}}$  and  $Alg^*$  across all tasks compared to the benchmark algorithms. These results underscore the advantage of leveraging prior knowledge about the tested environments.

**Setup.** All figures are based on 100 runs. We consider four tasks with 4 environments (i.e., the support of  $\mathcal{P}_{\gamma}$  contains only 4 environments) in Figure 13 and two tasks with 100 environments in Figure 14 (a), (b). The setup for these tasks follows the generation methods provided in Appendix E.2.

For the tasks in Figure 14 (c), (d), there are 16 environments included in the support of  $\mathcal{P}_{\gamma}$ , where 8 environments have demand functions of the linear type and the other 8 have demand functions of the square type (see Appendix E.2 for the definitions of these two types). Both the pre-training and testing samples are drawn from  $\mathcal{P}_{\gamma}$ , i.e., from these 16 environments. The remaining setups for these tasks follow the methods provided in Appendix E.2.



1296  
1297  
1298  
1299  
1300  
1301  
1302  
1303




1304  
1305  
1306  
1307  
1308  
1309

(a) Multi-armed bandits (b) Linear bandits (c) Dynamic pricing (d) Newsvendor

Figure 13: The average out-of-sample regret on tasks with simpler environments, where each task only has 4 possible environments. The numbers in the legend bar are the final regret at  $t = 100$  and the shaded areas indicate the 90% (empirical) confidence intervals. The details of benchmarks can be found in Appendix E.3.

1310  
1311  
1312  
1313  
1314  
1315  
1316  
1317



1318  
1319  
1320  
1321

(a) Multi-armed bandits, 100 environments, linear demand (b) Linear bandits, 100 environments, linear demand (c) Dynamic pricing, 16 environments, 2 demand types (d) Newsvendor, 16 environments, 2 demand types

1322  
1323  
1324  
1325  
1326

Figure 14: The average out-of-sample regret on harder tasks: with 100 environments in (a) multi-armed bandits and (b) linear bandits; or with 16 environments, 2 demand types in (c) dynamic pricing and (d) newsvendor. The numbers in the legend bar are the final regret at  $t = 100$  and the shaded areas indicate the 90% (empirical) confidence intervals. The details of benchmarks can be found in Appendix E.3

1327  
1328  
1329  
1330

C.6 PERFORMANCE ON LONGER TESTING HORIZON

1331  
1332  
1333  
1334  
1335  
1336

Our paradigm can be generalized to a longer testing horizon (beyond the length of the pre-training data) by introducing a *context window*. Specifically, we define a context window size  $W$  and limit the input to the last  $\min\{W, t\}$  timesteps of the input sequence  $H_t$  to predict  $a_t^*$  during both the pre-training and testing phases. By considering only the latest  $\min\{W, t\}$  timesteps, we enable the model to handle testing sequences of any length, even when  $t > T$ . This technique is commonly used in the literature (e.g., Chen et al. [13]).

1337  
1338  
1339  
1340  
1341  
1342  
1343  
1344  
1345

In this part, we evaluate  $TF_{\hat{\theta}}$ 's generalization ability when the testing horizon is longer than that seen during pre-training. As for a shorter testing horizon,  $TF_{\hat{\theta}}$  can simply run until the end of the horizon, and previous results demonstrate strong performances in such cases, particularly when the horizon is less than 100. Figure 15 shows the performance when the testing horizon is extended from 100 (pre-training horizon) to 200 in the newsvendor problem. The results demonstrate that  $TF_{\hat{\theta}}$  generalizes well in this extended horizon scenario, even in the censored demand setting where exploration is essential. Its actions remain nearly optimal beyond  $t = 100$ , resulting in lower regret compared to benchmark algorithms. For this experiment, all the pre-training and testing setups follow Appendix E, except the testing horizon length.

1346  
1347  
1348  
1349

C.7 OUT-OF-DISTRIBUTION PERFORMANCE

In this section, we evaluate  $TF_{\hat{\theta}}$ 's generalization ability under out-of-distribution (OOD) conditions, using dynamic pricing problems as the test case.

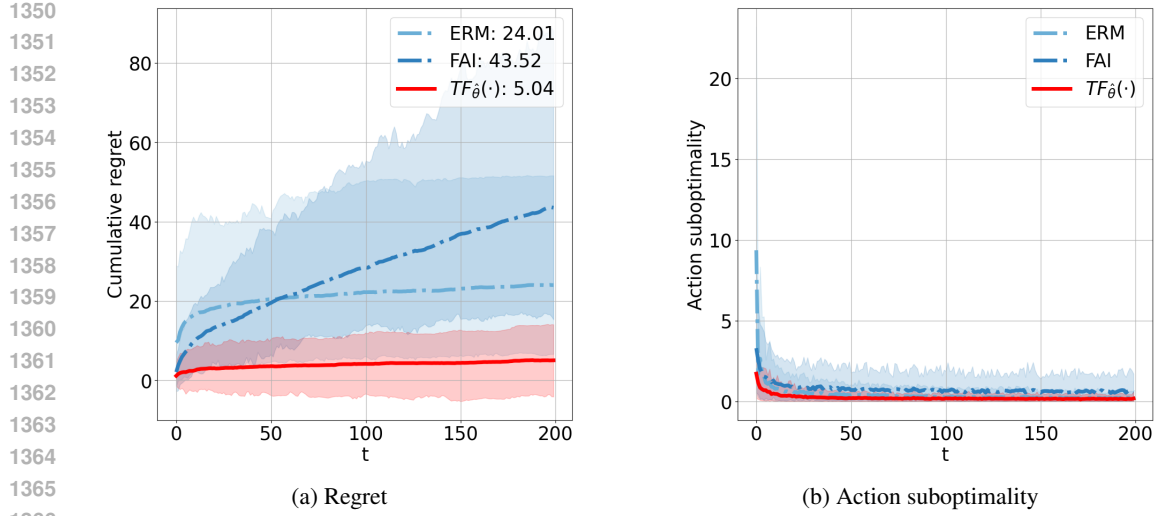


Figure 15: Performance under horizon generalization in newsvendor problems, where the pre-training samples have a horizon of 100 but the testing samples have a horizon of 200. It shows the average out-of-sample regret (first row) and action suboptimality, i.e.,  $|a_t^* - \text{Alg}(H_t)|$ , (second row) of  $TF_{\hat{\theta}}$  against benchmark algorithms. The numbers in the legend bar are the final regret at  $t = 200$ .

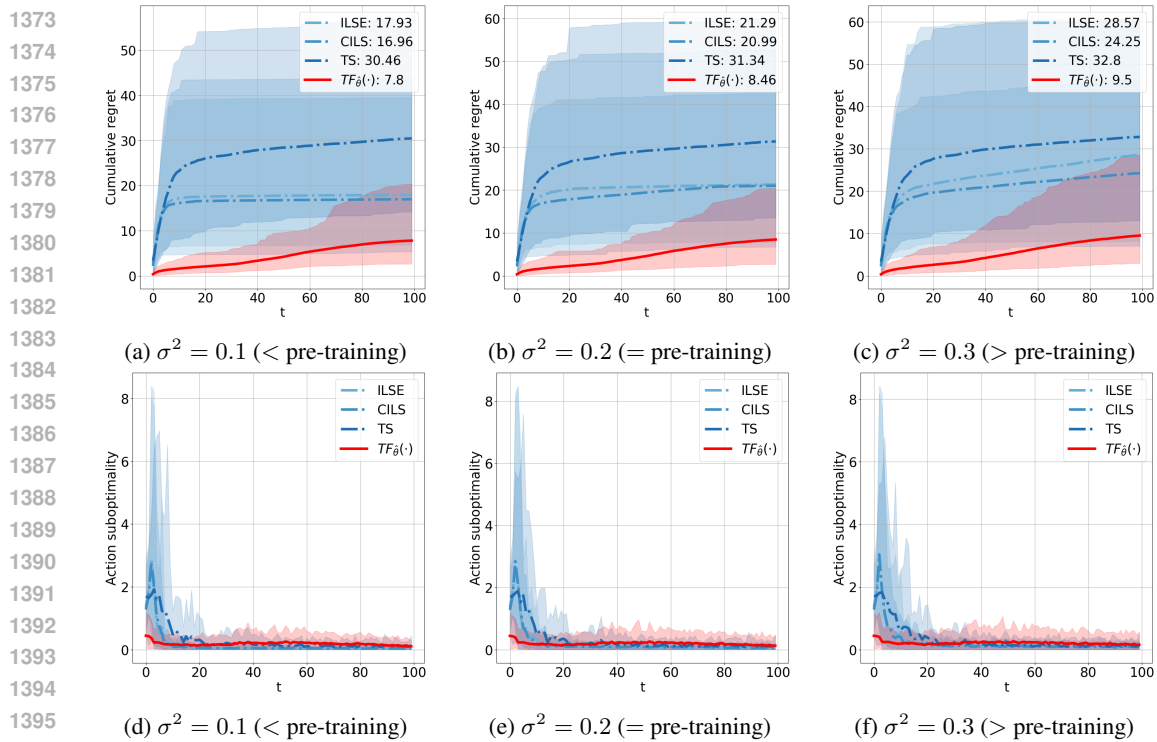


Figure 16: OOD performance under different testing noise variances, which may deviate from the pre-training variance of  $\sigma^2 = 0.2$ . It shows the average out-of-sample regret (first row) and action suboptimality, i.e.,  $|a_t^* - \text{Alg}(H_t)|$ , (second row) of  $TF_{\hat{\theta}}$  against benchmark algorithms. The numbers in the legend bar are the final regret at  $t = 100$ .

We test  $TF_{\hat{\theta}}$  in three different problem settings with varying testing noise variances while keeping the pre-training variance as 0.2 for all of them:  $\sigma^2 = 0.2$ , which matches the noise variance used

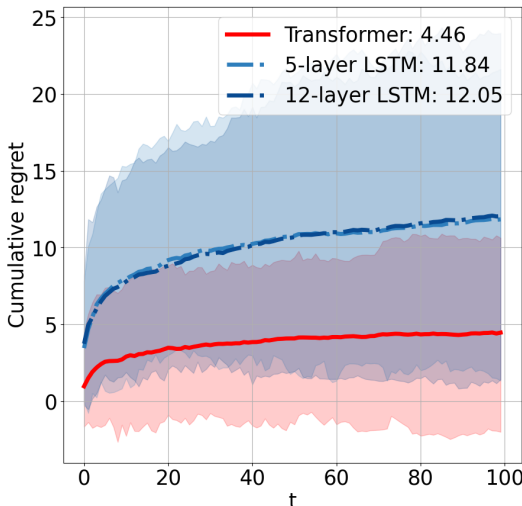
1404 in pre-training;  $\sigma^2 = 0.1$ , which is smaller than that used in pre-training; and  $\sigma^2 = 0.3$ , which is  
 1405 larger. Figure 16 provides the results. We do not observe any significant sign of failure in  $\text{TF}_{\hat{\theta}}$ 's  
 1406 OOD performance: across all three settings, the benchmark algorithms consistently incur higher regret  
 1407 than  $\text{TF}_{\hat{\theta}}$ . Although  $\text{TF}_{\hat{\theta}}$  shows some variation in mean regret across OOD settings (specifically,  
 1408 lower regret when  $\sigma^2 = 0.2$  and higher regret when  $\sigma^2 = 0.3$ ), this is partially due to the varying  
 1409 “difficulty” of the underlying tasks. As expected, higher noise variances, which need more data for  
 1410 accurately estimating the demand function compared to lower variance cases, lead to worse perform-  
 1411 ance for all algorithms tested. However, these variations should not be interpreted as evidence  
 1412 of  $\text{TF}_{\hat{\theta}}$  failing to handle OOD issues. In fact, all benchmark algorithms show similar performance  
 1413 fluctuations under these conditions, further demonstrating  $\text{TF}_{\hat{\theta}}$ 's good OOD performances. For this  
 1414 experiment, all the pre-training and testing setups follow Appendix E, except the noise variance in  
 1415 the testing samples.

### 1416 C.8 IMPACT OF NETWORK ARCHITECTURE

1418 Here we investigate the effect of different network architectures on performance. Specifically, we  
 1419 replace the transformer/GPT-2 architecture used in  $\text{TF}_{\hat{\theta}}$  with Long Short-Term Memory (LSTM)  
 1420 [25].

1421 We maintain the overall architecture of  $\text{TF}_{\hat{\theta}}$ , as described in Appendix E.1 except for replacing  
 1422 the transformer/GPT-2 module with LSTM. We evaluate two LSTM variants: a 5-layer LSTM and  
 1423 a 12-layer LSTM (for comparison, the tested transformer also has 12 layers). Typically, LSTM  
 1424 architectures should be shallower than transformers in practice, which is why we also include the  
 1425 5-layer version. Other hyperparameters like embedding space dimension are kept the same across  
 1426 all models. We pre-train and test these models on dynamic pricing problems (see Appendix E.2 for  
 1427 the experimental setup), and the pre-training procedure is identical to the one outlined in Appendix  
 1428 3.

1429 Figure 17 presents the testing regret for the three models. The results indicate no significant dif-  
 1430 ference between the 5-layer and 12-layer LSTM models. However, the transformer consistently  
 1431 performs better than both LSTM variants. This outcome aligns with the transformer’s superior per-  
 1432 formance in other domains, such as natural language processing, and suggests that transformers are  
 1433 more effective for sequential decision making tasks as well.



1434  
1435  
1436  
1437  
1438  
1439  
1440  
1441  
1442  
1443  
1444  
1445  
1446  
1447  
1448  
1449  
1450  
1451  
1452 Figure 17: Comparison of the average out-of-sample regret between using the Transformer and  
 1453 LSTM architectures.

### 1455 C.9 ABLATION STUDY ON MODEL SIZE AND TASK COMPLEXITY

1456 In this section, we explore the performance changes when tuning both model size and task complex-  
 1457 ity. Specifically, we adjust the number of layers to control the model size and the problem dimension

(the dimension of the context) in a dynamic pricing task to control the task complexity, where the number of unknown parameters is twice the context dimension. Since the optimal reward may scale differently across dimensions, we provide a relatively fair comparison by evaluating the advantage of  $\text{TF}_{\hat{\theta}}$  relative to the best-performing benchmark algorithm. This advantage is defined as the reward improvement rate of  $\text{TF}_{\hat{\theta}}$  compared to the best benchmark algorithm (the one among ILSE, CILS, and TS that achieves the highest reward). Thus, a positive value indicates  $\text{TF}_{\hat{\theta}}$  outperforms all benchmark algorithms, and a larger value indicates a better performance of  $\text{TF}_{\hat{\theta}}$ . We test three model sizes (4, 8, and 12 layers) and three problem dimensions (4, 10, and 20), while keeping other hyperparameters the same. The results are presented in Figure 18.

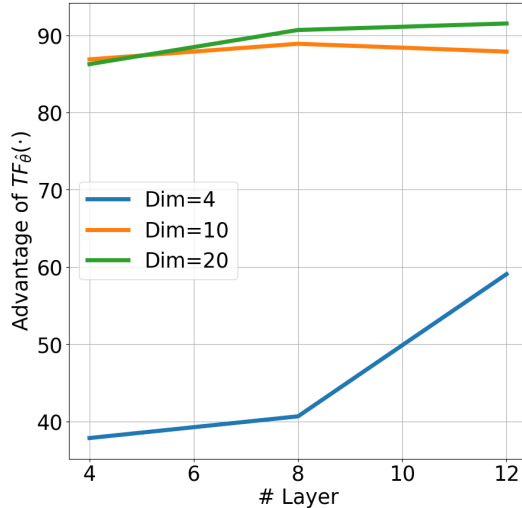


Figure 18: The advantage of  $\text{TF}_{\hat{\theta}}$  compared to the best benchmark algorithm in dynamic pricing, across different model size (number of layers) and problem complexity (problem dimensions).

From Figure 18, we observe the following: (i)  $\text{TF}_{\hat{\theta}}$  consistently outperforms the benchmark algorithms, and this advantage grows as the complexity of the problem increases. (ii) Increasing model size consistently improves the performance of  $\text{TF}_{\hat{\theta}}$  across all problem dimensions, indicating that larger models are generally preferred. This highlights  $\text{TF}_{\hat{\theta}}$ 's superior ability to handle more complex tasks and these observations align with the scaling law [31] for large language models.

## D PROOFS

### D.1 DISCUSSIONS ON CLAIM 3.1

Consider the following setup where a feature-target pair  $(X, Y)$  is generated from some distribution  $\mathcal{P}_{X,Y}(\theta)$ . Then we consider the minimization of the following loss

$$\min_{\theta} \mathbb{E}_{\mathcal{P}_{X,Y}(\theta)}[l(f_{\theta}(X), Y)]$$

for some loss function  $l : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ .

Under this formulation, both the data generation distribution  $\mathcal{P}_{X,Y}(\theta)$  and the prediction function  $f_{\theta}$  are parameterized by  $\theta$ . This describes the *performative prediction* problem [43].

In addition, we note that in our pre-trained transformer setting, the joint distribution can be factorized in

$$\mathcal{P}_{X,Y}(\theta) = \mathcal{P}_X(\theta) \cdot \mathcal{P}_{Y|X}.$$

That is, the parameter  $\theta$  only induces a covariate shift by affecting the marginal distribution of  $X$ , but the conditional distribution  $\mathcal{P}_{Y|X}$  remains the same for all the  $\theta$ . This exactly matches our setting of pre-trained transformer; to see this, different decision functions,  $f$  or  $\text{TF}_{\theta_m}$ , only differ in terms of the generation of the actions  $a_{\tau}$ 's in  $H_t$  (for  $\tau = 1, \dots, t - 1$ ) which is the features  $X$  in this formulation but the optimal action  $a_t^*$  will only be affected by  $X_t$  in  $H_t$ .

Also, under this factorization, we define the Bayes-optimal estimator as

$$f^*(X) = \min_y \mathbb{E}_{\mathcal{P}_{Y|X}}[l(y, Y)|X].$$

Then it is easy to note that when  $f^* = f_{\theta^*}$  for some  $\theta^*$ , and have

$$\theta^* \in \arg \min_{\theta} \mathbb{E}_{\mathcal{P}_{X,Y}(\theta')} [l(f_{\theta}(X), Y)]$$

for any  $\theta'$ .

In this light, the oscillating behavior of the optimization algorithms in [43] will not happen. Because for all the data generation distribution  $\mathcal{P}_{X,Y}(\theta')$ , they all point to one optimal  $\theta^*$ .

Back to the context of the pre-trained transformer, such a nice property is contingent on two factors: (i) the transformer function class is rich enough to cover  $f^*$ ; (ii) there are infinitely many samples/we can use the expected loss. Also, the above argument is connected to the argument of Proposition 4.1.

## D.2 PROOF OF PROPOSITION 3.2

As noted in the proposition, let  $b$  denote the bound of the loss function, and let the loss function  $l(\text{TF}_{\theta}(H), a)$  is  $D$ -Lipschitz with respect to  $H$  and  $a$ . We remark that the constant  $D$  could be very large for the transformer model; the bound derived here serves more to illustrate the relationship between the generalization error and various problem parameters, but not really as an empirical bound to predict the test error of the underlying transformer.

Step 1. We aim to show with probability at least  $1 - h$ ,

$$\begin{aligned} \kappa L_f(\text{TF}_{\hat{\theta}}) + (1 - \kappa) L_{\text{TF}_{\hat{\theta}}}(\text{TF}_{\hat{\theta}}) &\leq \frac{1}{nT} \sum_{i=1}^n \sum_{t=1}^T l(\text{TF}_{\hat{\theta}}(H_t^{(i)}), a_t^{(i)*}) \\ &\quad + \sqrt{\frac{\text{Comp}(\{\text{TF}_{\theta} : \theta \in \Theta\})}{nT}} + \sqrt{\frac{2b^2}{nT} \log\left(\frac{4}{h}\right)}, \end{aligned}$$

where  $\text{Comp}(\{\text{TF}_{\theta} : \theta \in \Theta\})$  is some inherent constant describing the complexity of  $\{\text{TF}_{\theta} : \theta \in \Theta\}$ . Note that when samples follow the generation process (actions are generated by  $f$ )

$$\mathcal{P}_{\gamma} \rightarrow \gamma_i \rightarrow \mathcal{P}_{\gamma_i, f} \rightarrow \left\{ \left( H_1^{(i)}, a_1^{(i)*} \right), \dots, \left( H_T^{(i)}, a_T^{(i)*} \right) \right\},$$

and

$$L_f(\text{TF}_{\theta}) = \mathbb{E}_{\gamma \sim \mathcal{P}_{\gamma}} [L_f(\text{TF}_{\theta}; \gamma)] = \mathbb{E}_{\gamma \sim \mathcal{P}_{\gamma}} \left[ \mathbb{E}_{(H_t, a_t^*) \sim \mathcal{P}_{\gamma, f}} \left[ \sum_{t=1}^T l(\text{TF}_{\theta}(H_t), a_t^*) \right] \right].$$

Since  $l(\text{TF}_{\theta}(\cdot), \cdot)$  is bounded by  $b$ , we can apply the McDiarmid's inequality (see Theorem D.8 in [40] for more details) with constant  $\frac{2b}{nT}$  to  $\sup_{\theta \in \Theta} \left| \frac{1}{nT} \sum_{i=1}^n \sum_{t=1}^T l(\text{TF}_{\theta}(H_t^{(i)}), a_t^{(i)*}) - L_f(\text{TF}_{\theta}) \right|$  such that with probability at least  $1 - h$ , we have

$$\begin{aligned} &\sup_{\theta \in \Theta} \left| \frac{1}{nT} \sum_{i=1}^n \sum_{t=1}^T l(\text{TF}_{\theta}(H_t^{(i)}), a_t^{(i)*}) - L_f(\text{TF}_{\theta}) \right| \\ &\leq \mathbb{E}_{(H_t, a_t^*) \sim \mathcal{P}_{\gamma, f}} \left[ \sup_{\theta \in \Theta} \left| \frac{1}{nT} \sum_{i=1}^n \sum_{t=1}^T l(\text{TF}_{\theta}(H_t^{(i)}), a_t^{(i)*}) - \mathbb{E}_{(H_t, a_t^*) \sim \mathcal{P}_{\gamma, f}} \left[ \sum_{t=1}^T l(\text{TF}_{\theta}(H_t), a_t^*) \right] \right| \right] \quad (10) \end{aligned}$$

$$+ \sqrt{\frac{2b^2}{nT} \log\left(\frac{2}{h}\right)}$$

$$\leq 2 \mathbb{E}_{(H_t, a_t^*) \sim \mathcal{P}_{\gamma, f}, w_{it}} \left[ \sup_{\theta \in \Theta} \left| \frac{1}{nT} \sum_{i=1}^n \sum_{t=1}^T w_{it} l(\text{TF}_{\theta}(H_t^{(i)}), a_t^{(i)*}) \right| \right] + \sqrt{\frac{2b^2}{nT} \log\left(\frac{2}{h}\right)}$$

$$\leq \sqrt{\frac{\text{Comp}(\{\text{TF}_{\theta} : \theta \in \Theta\})}{nT}} + \sqrt{\frac{2b^2}{nT} \log\left(\frac{2}{h}\right)}, \quad (11)$$

where the second inequality uses the symmetric tricks (see Theorem 3.3 in [40] for more details), and  $w_{it}$ , so-called Rademacher variables are independent uniform random variables taking values in  $\{-1, 1\}$ .

Similarly, when samples follow the generation process (actions are generated by  $\text{TF}_{\hat{\theta}}$ )

$$\mathcal{P}_{\gamma} \rightarrow \gamma_i \rightarrow \mathcal{P}_{\gamma_i, \text{TF}_{\hat{\theta}}} \rightarrow \left\{ \left( H_1^{(i)}, a_1^{(i)*} \right), \dots, \left( H_T^{(i)}, a_T^{(i)*} \right) \right\},$$

by using the same tricks, we can have that with probability at least  $1 - h$ ,

$$\sup_{\theta \in \Theta} \left| \frac{1}{nT} \sum_{i=1}^n \sum_{t=1}^T l \left( \text{TF}_{\theta} \left( H_t^{(i)} \right), a_t^{(i)*} \right) - L_{\text{TF}_{\hat{\theta}}}(\text{TF}_{\theta}) \right| \leq \sqrt{\frac{\text{Comp}(\{\text{TF}_{\theta} : \theta \in \Theta\})}{nT}} + \sqrt{\frac{2b^2}{nT} \log \left( \frac{2}{h} \right)}. \quad (12)$$

Since  $\hat{\theta}$  is determined by equation 6 where  $\kappa n$  data sequences are from  $\mathcal{P}_{\gamma, f}$  and  $(1 - \kappa)n$  data sequences are from  $\mathcal{P}_{\gamma, \text{TF}_{\hat{\theta}}}$ , by taking the uniform bound of equations (11) and (12), we have with probability at least  $1 - h$ ,

$$\begin{aligned} \kappa L_f(\text{TF}_{\hat{\theta}}) + (1 - \kappa) L_{\text{TF}_{\hat{\theta}}}(\text{TF}_{\hat{\theta}}) &\leq \frac{1}{nT} \sum_{i=1}^n \sum_{t=1}^T l \left( \text{TF}_{\hat{\theta}} \left( H_t^{(i)} \right), a_t^{(i)*} \right) \\ &\quad + \sqrt{\frac{\text{Comp}(\{\text{TF}_{\theta} : \theta \in \Theta\})}{nT}} + \sqrt{\frac{2b^2}{nT} \log \left( \frac{4}{h} \right)}. \end{aligned}$$

Step 2. Now we show

$$L(\text{TF}_{\hat{\theta}}) \leq L_f(\text{TF}_{\hat{\theta}}) + TD \cdot \mathbb{E}_{\gamma \sim \mathcal{P}_{\gamma}} [W_1(\mathcal{P}_{\gamma, f}, \mathcal{P}_{\gamma, \text{TF}_{\hat{\theta}}})],$$

$$L(\text{TF}_{\hat{\theta}}) \leq L_f(\text{TF}_{\hat{\theta}}) + TD \cdot \mathbb{E}_{\gamma \sim \mathcal{P}_{\gamma}} [W_1(\mathcal{P}_{\gamma, \text{TF}_{\hat{\theta}}}, \mathcal{P}_{\gamma, \text{TF}_{\hat{\theta}}})],$$

where  $W_1(\cdot, \cdot)$  is the Wasserstein-1 distance.

By Kantorovich-Rubinstein's theorem, for distributions  $\mathcal{D}_1$  and  $\mathcal{D}_2$ , we have

$$W_1(\mathcal{D}_1, \mathcal{D}_2) = \sup \{ |\mathbb{E}_{X \sim \mathcal{D}_1} g(X) - \mathbb{E}_{X \sim \mathcal{D}_2} g(X)| \mid g : \mathbb{R}^p \rightarrow \mathbb{R}, g \text{ is 1-Lipschitz} \}.$$

Since  $l(\text{TF}_{\theta}(H), a)$  is  $D$ -Lipschitz with respect to  $H$  and  $a$  for all  $\theta \in \Theta$ , we have

$$\begin{aligned} L(\text{TF}_{\hat{\theta}}; \gamma) - L_f(\text{TF}_{\hat{\theta}}; \gamma) &= \mathbb{E}_{(H_t, a_t^*) \sim \mathcal{P}_{\gamma, \text{TF}_{\hat{\theta}}}} \left[ \sum_{t=1}^T l(\text{TF}_{\hat{\theta}}(H_t), a_t^*) \right] - \mathbb{E}_{(H_t, a_t^*) \sim \mathcal{P}_{\gamma, f}} \left[ \sum_{t=1}^T l(\text{TF}_{\hat{\theta}}(H_t), a_t^*) \right] \\ &\leq TD \cdot W_1(\mathcal{P}_{\gamma, f}, \mathcal{P}_{\gamma, \text{TF}_{\hat{\theta}}}), \end{aligned}$$

which implies that  $L(\text{TF}_{\hat{\theta}}) \leq L_f(\text{TF}_{\hat{\theta}}) + TD \cdot \mathbb{E}_{\gamma \sim \mathcal{P}_{\gamma}} [W_1(\mathcal{P}_{\gamma, f}, \mathcal{P}_{\gamma, \text{TF}_{\hat{\theta}}})]$ . Similarly, we have

$$L(\text{TF}_{\hat{\theta}}) \leq L_f(\text{TF}_{\hat{\theta}}) + TD \cdot \mathbb{E}_{\gamma \sim \mathcal{P}_{\gamma}} [W_1(\mathcal{P}_{\gamma, \text{TF}_{\hat{\theta}}}, \mathcal{P}_{\gamma, \text{TF}_{\hat{\theta}}})].$$

Therefore, by combining step 1 and step 2, we can conclude that with probability at least  $1 - h$ ,

$$\begin{aligned} L(\text{TF}_{\hat{\theta}}) &\leq \hat{L}(\text{TF}_{\hat{\theta}}) + \sqrt{\frac{\text{Comp}(\{\text{TF}_{\theta} : \theta \in \Theta\})}{nT}} \\ &\quad + \kappa TD \cdot \mathbb{E}_{\gamma \sim \mathcal{P}_{\gamma}} [W_1(\mathcal{P}_{\gamma, f}, \mathcal{P}_{\gamma, \text{TF}_{\hat{\theta}}})] \\ &\quad + (1 - \kappa) TD \cdot \mathbb{E}_{\gamma \sim \mathcal{P}_{\gamma}} [W_1(\mathcal{P}_{\gamma, \text{TF}_{\hat{\theta}}}, \mathcal{P}_{\gamma, \text{TF}_{\hat{\theta}}})] + \sqrt{\frac{2b^2}{nT} \log \left( \frac{4}{h} \right)}, \end{aligned}$$

which completes the proof.

### D.3 PROOF OF PROPOSITION 4.1

*Proof.* The proof can be done with the definition of  $\text{Alg}^*$ . Also, we refer more discussion to Section D.1.  $\square$

## D.4 PROOF OF EXAMPLE 4.2

*Proof.* Since  $\text{Alg}^*$  is a function of any possible history  $H_t$ , we omit the subscript  $t$  for the optimal actions for notation simplicity, and use  $a^{(\gamma)^*}$  as the optimal action of  $\gamma$ .

- **Posterior sampling under the cross-entropy loss.** We assume that  $\mathcal{A} = \{a_1, \dots, a_J\}$  and the output of  $\text{Alg}^*$  is a probability vector  $(p_1, \dots, p_J)$  such that  $\sum_{j=1}^J p_j = 1$ , meaning the probability of choosing each action. Then

$$\begin{aligned} \text{Alg}^*(H) &= \arg \min_{\sum_{j=1}^J p_j = 1} \int_{\gamma} -\log(p_{a^{(\gamma)^*}}) \mathbb{P}(H|\gamma) d\mathcal{P}_{\gamma} \\ &= \arg \min_{\sum_{j=1}^J p_j = 1} \sum_{j=1}^J (-\log(p_j)) \int_{\{\gamma: a^{(\gamma)^*} = j\}} \mathbb{P}(H|\gamma) d\mathcal{P}_{\gamma}. \end{aligned}$$

By noting  $\int_{\{\gamma: a^{(\gamma)^*} = j\}} \mathbb{P}(H|\gamma) d\mathcal{P}_{\gamma}$  is the probability such that action  $j$  is the optimal action conditional on  $H$  (i.e., the posterior distribution of the optimal action), we are indeed minimizing the cross-entropy of the posterior distribution of the optimal action relative to the decision variables  $p_j$ . Thus, the optimal solution  $p_j^* = \int_{\{\gamma: a^{(\gamma)^*} = j\}} \mathbb{P}(H|\gamma) d\mathcal{P}_{\gamma}$  for each  $j$  and  $\text{Alg}^*(H_t)$  behaves as the posterior sampling.

- **Posterior averaging under the squared loss.** By definition,

$$\begin{aligned} \text{Alg}^*(H) &= \arg \min_{a \in \mathcal{A}} \int_{\gamma} \|a - a^{(\gamma)^*}\|_2^2 \mathbb{P}(H|\gamma) d\mathcal{P}_{\gamma} \\ &= \frac{\int_{\gamma} a^{(\gamma)^*} \mathbb{P}(H|\gamma) d\mathcal{P}_{\gamma}}{\int_{\gamma} \mathbb{P}(H|\gamma) d\mathcal{P}_{\gamma}} \\ &= \mathbb{E}_{\gamma}[a^{(\gamma)^*} | H] \end{aligned}$$

where the second line is by the first order condition.

- **Posterior median under the absolute loss.** By definition,

$$\text{Alg}^*(H) = \arg \min_{a \in \mathcal{A}} \int_{\gamma} |a - a^{(\gamma)^*}| \mathbb{P}(H|\gamma) d\mathcal{P}_{\gamma}.$$

Then by zero-subgradient condition, we can conclude that  $\text{Alg}^*(H)$  satisfies

$$\int_{\{\gamma: \text{Alg}^*(H) \leq a^{(\gamma)^*}\}} \mathbb{P}(H|\gamma) d\mathcal{P}_{\gamma} = \int_{\{\gamma: \text{Alg}^*(H) \geq a^{(\gamma)^*}\}} \mathbb{P}(H|\gamma) d\mathcal{P}_{\gamma}.$$

Divide both sides of this equation by  $\int_{\gamma} \mathbb{P}(H|\gamma) d\mathcal{P}_{\gamma}$ , we can conclude that

$$\int_{\{\gamma: \text{Alg}^*(H) \leq a^{(\gamma)^*}\}} d\mathcal{P}(\gamma|H) = \int_{\{\gamma: \text{Alg}^*(H) \geq a^{(\gamma)^*}\}} d\mathcal{P}(\gamma|H)$$

where  $\mathcal{P}(\gamma|H_t)$  is the posterior distribution of  $\gamma$ . Hence  $\text{Alg}^*(H)$  is the posterior median.

□

## D.5 PROOF OF PROPOSITION 4.3

*Proof. For the linear bandits problem,* we consider the following example. Suppose we have two environments,  $\gamma_1, \gamma_2$ , each with standard normal distributed noise and unknown parameters  $w_{\gamma_1} = (1, 0)$  and  $w_{\gamma_2} = (0, 1)$ , respectively. Let the action space be  $\mathcal{A} = [-1, 1] \times [-1, 1]$ . The optimal actions for the two environments are  $a^{(1)^*} = (1, 0)$ ,  $a^{(2)^*} = (0, 1)$ . We assume the prior distribution  $\mathcal{P}_{\gamma}$  is given by  $\mathcal{P}_{\gamma}(\gamma_1) = \mathcal{P}_{\gamma}(\gamma_2) = \frac{1}{2}$ .

1674 Then if  $\text{Alg}^*$  is the posterior averaging, we have  
1675

$$1676 \quad a_t = \text{Alg}^*(H_t) = \mathcal{P}_\gamma(\gamma_1|H_t) \cdot a^{(1)*} + \mathcal{P}_\gamma(\gamma_2|H_t) \cdot a^{(2)*},$$

1677 where the posterior distribution  $\mathcal{P}_\gamma(\gamma_i|H_t)$  is given by  
1678

$$1679 \quad \mathcal{P}_\gamma(\gamma_i|H_t) = \frac{\prod_{\tau=1}^{t-1} \mathbb{P}_{\gamma_i}(o_\tau|a_\tau)}{\sum_{i'=1}^2 \prod_{\tau=1}^{t-1} \mathbb{P}_{\gamma_{i'}}(o_\tau|a_\tau)}, \quad i = 1, 2.$$

1680 We now use induction to show  $\mathcal{P}_\gamma(\gamma_i|H_t) = \frac{1}{2}$  for any  $t \geq 1$  and  $i = 1, 2$ , where  $H_t$  can be  
1681 generated by either  $\gamma_1$  or  $\gamma_2$ . This results in  $a_t = \frac{1}{2}a^{(1)*} + \frac{1}{2}a^{(2)*} = (\frac{1}{2}, \frac{1}{2})$ , which does not change  
1682 with respect to  $t$  and will cause regret linear in  $T$ .

1683 Step 1. Since  $\mathcal{P}_\gamma(\gamma_i|H_0) = \mathcal{P}_\gamma(\gamma_i) = \frac{1}{2}$ , we have  $a_1 = \frac{1}{2}a^{(1)*} + \frac{1}{2}a^{(2)*} = (\frac{1}{2}, \frac{1}{2})$ , and the  
1684 conclusion holds for  $t = 1$ .  
1685

1686 Step 2. Now assume the conclusion holds for  $t$ , i.e.,  $\mathcal{P}_\gamma(\gamma_i|H_t) = \frac{1}{2}$ , and  $a_t = (\frac{1}{2}, \frac{1}{2})$ . Since  
1687  $w_{\gamma_1}^T a_t = w_{\gamma_2}^T a_t = \frac{1}{2}$ , we have  $\mathbb{P}_{\gamma_i}(o_t|a_t) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{(o_t - \frac{1}{2})^2}{2}\right)$  for  $i = 1, 2$ . Observe that  
1688

$$1689 \quad \mathcal{P}_\gamma(\gamma_i|H_{t+1}) = \frac{\mathbb{P}_{\gamma_i}(o_{t+1}|a_{t+1})\mathcal{P}_\gamma(\gamma_i|H_t)}{\sum_{i'=1}^2 \mathbb{P}_{\gamma_{i'}}(o_{t+1}|a_{t+1})\mathcal{P}_\gamma(\gamma_{i'}|H_t)} = \frac{1}{2},$$

1690 which implies  $a_{t+1} = (\frac{1}{2}, \frac{1}{2})$ , and the conclusion holds for  $t + 1$ .  
1691

1692 Thus, the conclusion holds for all  $t \geq 1$ . Then the regret is  
1693

$$1694 \quad \text{Regret}(\text{Alg}^*; \gamma_i) = \mathbb{E} \left[ \sum_{t=1}^T r(X_t, a_t^*) - r(X_t, a_t) \right] = \frac{1}{2}T$$

1695 for  $i = 1, 2$ .  
1696

1697 **For the dynamic pricing problem**, we can construct a similar example. Suppose we have two  
1698 environments without context  $X_t$ , denoted by  $\gamma_1, \gamma_2$ . The demands  $o_t$  of them are set to be  $o_t =$   
1699  $2 - a_t + \epsilon_t$  and  $o_t = \frac{4}{5} - \frac{1}{5} \cdot a_t + \epsilon_t$  respectively, where  $\epsilon_t \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1)$  and  $a_t$  is the price.  
1700 Accordingly, the optimal actions are then  $a^{(1)*} = 1$  and  $a^{(2)*} = 2$ .  
1701

1702 Then if  $\text{Alg}^*$  is the posterior averaging, we have  
1703

$$1704 \quad a_t = \text{Alg}^*(H_t) = \mathcal{P}_\gamma(\gamma_1|H_t) \cdot a^{(1)*} + \mathcal{P}_\gamma(\gamma_2|H_t) \cdot a^{(2)*},$$

1705 where the posterior distribution  $\mathcal{P}_\gamma(\gamma_i|H_t)$  is given by  
1706

$$1707 \quad \mathcal{P}_\gamma(\gamma_i|H_t) = \frac{\prod_{\tau=1}^{t-1} \mathbb{P}_{\gamma_i}(o_\tau|a_\tau)}{\sum_{i'=1}^2 \prod_{\tau=1}^{t-1} \mathbb{P}_{\gamma_{i'}}(o_\tau|a_\tau)}, \quad i = 1, 2.$$

1708 with  $\mathbb{P}_{\gamma_1}(o_\tau|a_\tau)$  as the normal distribution  $\mathcal{N}(2 - a_t, 1)$  and  $\mathbb{P}_{\gamma_2}(o_\tau|a_\tau)$  as  $\mathcal{N}(\frac{4-a_t}{5}, 1)$ .  
1709

1710 Observe that  $a_1 = \frac{1}{2}a^{(1)*} + \frac{1}{2}a^{(2)*} = \frac{3}{2}$  satisfies  $2 - a_1 = \frac{4-a_1}{5}$ . Following a similar analysis as in  
1711 the linear bandits example, we can conclude that  $\mathcal{P}_\gamma(\gamma_i|H_t) = \frac{1}{2}$  for any  $t$  and  $i = 1, 2$ , where  $H_t$   
1712 can be generated by either  $\gamma_1$  or  $\gamma_2$ . Therefore,  $a_t = \frac{3}{2}$  for all  $t \geq 1$  and the regret is  
1713

$$1714 \quad \text{Regret}(\text{Alg}^*; \gamma_i) = \mathbb{E} \left[ \sum_{t=1}^T r(X_t, a_t^*) - r(X_t, a_t) \right] = \frac{1}{4}T \cdot \mathbb{I}\{\gamma_i = \gamma_1\} + \frac{1}{20}T \cdot \mathbb{I}\{\gamma_i = \gamma_2\},$$

1715 for  $i = 1, 2$ .  
1716

1717

□



## D.6 PROOF OF PROPOSITION 4.4

We make the following additional assumptions:

- There exists a constant  $\bar{r}$  such that  $\sup_{x \in \mathcal{X}, a \in \mathcal{A}, \gamma \in \Gamma} r(x, a; \gamma) \leq \bar{r}$ , where we use  $r(x, a; \gamma)$  to denote the reward function of environment  $\gamma$ .
- There exists a constant  $C_r$  such that  $\text{Alg}^*$  satisfies

$$\mathbb{E}[r(X_t, a_t^*) - r(X_t, \text{Alg}^*(H_t)) | H_t] \leq C_r \sum_{\gamma_i \neq \gamma} \mathcal{P}(\gamma_i | H_t),$$

where the expectation on the left side is taken with respect to the possible randomness in  $\text{Alg}^*$ .

- $\mathcal{P}_\gamma$  is a uniform distribution over  $\Gamma$ .

*Proof.* We first compute a concentration inequality for  $\sum_{\tau=1}^t \log \left( \frac{\mathbb{P}_\gamma(o_\tau | X_\tau, a_\tau)}{\mathbb{P}_{\gamma'}(o_\tau | X_\tau, a_\tau)} \right)$ . By the Bernstein-type concentration bound for a martingale difference sequence (Theorem 2.19 in [50]), under the given conditions, we have for any  $\gamma' \in \Gamma$ ,  $t > 0$  and  $1 > \delta > 0$  with probability  $1 - \delta$ ,

$$\sum_{\tau=1}^t \log \left( \frac{\mathbb{P}_\gamma(o_\tau | X_\tau, a_\tau)}{\mathbb{P}_{\gamma'}(o_\tau | X_\tau, a_\tau)} \right) - \mathbb{E} \left[ \log \left( \frac{\mathbb{P}_\gamma(o_\tau | X_\tau, a_\tau)}{\mathbb{P}_{\gamma'}(o_\tau | X_\tau, a_\tau)} \right) \right] \geq -\sqrt{2C_{\sigma^2} \Delta_{\text{Explore}} t \log(1/\delta)},$$

where

$$\sum_{\tau=1}^t \mathbb{E} \left[ \log \left( \frac{\mathbb{P}_\gamma(o_\tau | X_\tau, a_\tau)}{\mathbb{P}_{\gamma'}(o_\tau | X_\tau, a_\tau)} \right) \right] = \sum_{\tau=1}^t \text{KL}(\mathbb{P}_\gamma(\cdot | X_\tau, a_\tau) \| \mathbb{P}_{\gamma'}(\cdot | X_\tau, a_\tau)) \geq t \Delta_{\text{Explore}}.$$

We think about two situations:

- If  $\sqrt{t \Delta_{\text{Explore}}} \geq 2\sqrt{2C_{\sigma^2} \log(1/\delta)}$ , then

$$\begin{aligned} \sum_{\tau=1}^t \log \left( \frac{\mathbb{P}_\gamma(o_\tau | X_\tau, a_\tau)}{\mathbb{P}_{\gamma'}(o_\tau | X_\tau, a_\tau)} \right) &\geq t \Delta_{\text{Explore}} - \frac{1}{2} t \Delta_{\text{Explore}} \\ &= \frac{1}{2} t \Delta_{\text{Explore}} \\ &> \frac{1}{2} t \Delta_{\text{Explore}} - 4C_{\sigma^2} \log(1/\delta). \end{aligned}$$

- If  $\sqrt{t \Delta_{\text{Explore}}} < 2\sqrt{2C_{\sigma^2} \log(1/\delta)}$ , then

$$\begin{aligned} \sum_{\tau=1}^t \log \left( \frac{\mathbb{P}_\gamma(o_\tau | X_\tau, a_\tau)}{\mathbb{P}_{\gamma'}(o_\tau | X_\tau, a_\tau)} \right) &> t \Delta_{\text{Explore}} - 4C_{\sigma^2} \log(1/\delta) \\ &\geq \frac{1}{2} t \Delta_{\text{Explore}} - 4C_{\sigma^2} \log(1/\delta). \end{aligned}$$

Thus with probability  $1 - \delta$ , we have

$$\sum_{\tau=1}^t \log \left( \frac{\mathbb{P}_\gamma(o_\tau | X_\tau, a_\tau)}{\mathbb{P}_{\gamma'}(o_\tau | X_\tau, a_\tau)} \right) > \frac{1}{2} t \Delta_{\text{Explore}} - 4C_{\sigma^2} \log(1/\delta).$$

Now we apply the union bound to  $\gamma_i$  for  $i = 1, \dots, k$  and for  $t = 1, \dots, T$ : with probability  $1 - 1/T$ , for all  $i = 1, \dots, k$ , and  $t = 1, \dots, T$

$$\mathcal{P}(\gamma_i | H_t) \leq \mathcal{P}(\gamma | H_t) \cdot \exp \left( -\frac{1}{2} t \Delta_{\text{Explore}} + 4C_{\sigma^2} \log(1/\delta) \right).$$

Now we are ready to prove the regret bound. We first decompose the regret:

$$\begin{aligned} \text{Regret}(\text{TF}_{\hat{\theta}}; \gamma) &= \sum_{t=1}^T \mathbb{E}[r(X_t, \text{Alg}^*(H_t)) - r(X_t, \text{TF}_{\hat{\theta}}(H_t))] + \sum_{t=1}^T \mathbb{E}[r(X_t, a_t^*) - r(X_t, \text{Alg}^*(H_t))] \\ &\leq \Delta_{\text{Exploit}} T + C_r \mathbb{E} \left[ \sum_{t=1}^T \sum_{\gamma_i \neq \gamma} \mathcal{P}(\gamma_i | H_t) \right]. \end{aligned}$$

Since  $\mathcal{P}(\gamma_i | H_t) + \mathcal{P}(\gamma | H_t) \leq 1$ , we have with probability  $1 - 1/T$ , for all  $i = 1, \dots, k$ , and  $t = 1, \dots, T$ ,

$$\mathcal{P}(\gamma_i | H_t) \leq \frac{1}{1 + \exp\left(\frac{1}{2}t\Delta_{\text{Explore}} - 4C_{\sigma^2} \log(kT^2)\right)}.$$

Let  $t_0 = \left\lceil \frac{16C_{\sigma^2} \log(kT^2)}{\Delta_{\text{Explore}}} \right\rceil$ , then we have with probability  $1 - 1/T$ , for all  $i = 1, \dots, k$ , and  $t = 1, \dots, T$ ,

$$\mathcal{P}(\gamma_i | H_t) \leq \frac{1}{1 + \exp\left(\frac{1}{2}t\Delta_{\text{Explore}} - 4C_{\sigma^2} \log(kT^2)\right)} < \exp\left(-\frac{1}{4}t\Delta_{\text{Explore}}\right).$$

Thus, we can conclude that

$$\begin{aligned} \mathbb{E} \left[ \sum_{t=1}^T \sum_{\gamma_i \neq \gamma} \mathcal{P}(\gamma_i | H_t) \right] &< 1 + \frac{16C_{\sigma^2} \bar{r} \log(kT^2)}{\Delta_{\text{Explore}}} + k \mathbb{E} \left[ \sum_{t=t_0}^T \exp\left(-\frac{1}{4}t\Delta_{\text{Explore}}\right) \right] \\ &= O\left(\frac{\log k}{\Delta_{\text{Explore}}}\right) \end{aligned}$$

and finish the regret bound.  $\square$

#### D.7 EXAMPLE/JUSTIFICATION FOR THE CONDITIONS IN PROPOSITION 4.4

We use a set of linear bandits problems as an example to illustrate Proposition 4.4. Consider a set of linear bandits problems with  $o_t = R_t = w^\top a_t + \epsilon_t$ , where  $\epsilon_t$  is i.i.d. from a standard Normal distribution. Recall the environment parameter  $\gamma = w$  here. We assume it has dimension  $d$  and both  $\mathcal{A}$  and  $\Gamma$  are bounded with respect to the  $L_2$  norm. Further, we assume  $\text{Alg}^*(H_t)$  is posterior averaging and for each time  $t$ ,

$$\text{TF}_{\hat{\theta}}(H_t) = \text{Alg}^*(H_t) + \Delta_t,$$

where  $\Delta_t$  follows a uniform distribution in  $[-1, 1]^d$ , i.i.d. across time and independent of  $H_t$ ,  $a_t$ , and  $\gamma$ . This follows the numerical observations in Figure 3.

In this setting, we verify the conditions in Proposition 4.4. To simplify notations, we denote  $\tilde{a}_t^* = \text{Alg}^*(H_t)$ .

**First condition.** Since

$$\mathbb{E}[r(X_t, \text{Alg}^*(H_t)) - r(X_t, \text{TF}_{\hat{\theta}}(H_t)) | H_t] = \mathbb{E}[\gamma^\top \tilde{a}_t^* - \gamma^\top \tilde{a}_t^* - \gamma^\top \Delta_t | H_t] = 0,$$

we have  $\Delta_{\text{Exploit}} = 0$ .

**Second condition.** Since  $\epsilon_t$  is i.i.d. from a standard Normal distribution, we have for  $o_t = \gamma^\top a_t + \epsilon_t$  and  $a_t = \text{TF}_{\hat{\theta}}(H_t)$ ,

$$\begin{aligned} \log\left(\frac{\mathbb{P}_{\gamma}(o_t | X_t, a_t)}{\mathbb{P}_{\gamma'}(o_t | X_t, a_t)}\right) &= \frac{(o_t - \gamma'^\top a_t)^2 - (o_t - \gamma^\top a_t)^2}{2} \\ &= \frac{((\gamma - \gamma')^\top a_t)^2}{2} - (\gamma - \gamma')^\top a_t \epsilon_t. \end{aligned}$$

Thus, the KL divergence, which is the expectation of the above term with respect to  $\epsilon_t$  and  $\Delta_t$  (which are independent of  $H_t$ ), is

$$\mathbb{E} \left[ \frac{((\gamma - \gamma')^\top (\tilde{a}_t^* + \Delta_t))^2}{2} \middle| H_t \right] \geq \frac{1}{6} \|\gamma - \gamma'\|_2^2.$$

And further, for the “noise” term  $(\gamma - \gamma')^\top a_t \epsilon_t$ , since its variance is bounded by  $O(\|\gamma - \gamma'\|_2^2)$  (due to the assumptions on the boundedness of  $a_t$  and that  $\epsilon_t$  follows a standard Normal distribution), we can set  $\Delta_{\text{Explore}} = \min_{\gamma' \in \Gamma} \|\gamma - \gamma'\|_2^2$ .

Further, since

$$\mathbb{E}[r(X_t, a_t^*) - r(X_t, \text{Alg}^*(H_t)) | H_t] \leq \bar{r} \sum_{\gamma_i \neq \gamma} \mathcal{P}(\gamma_i | H_t),$$

we can set  $C_r = \bar{r}$ .

And thus, we can conclude the regret bound is  $\text{Regret}(\text{TF}_{\hat{\theta}}, \gamma) = O\left(\frac{\log k}{\min_{\gamma' \in \Gamma} \|\gamma - \gamma'\|_2^2}\right)$ .

## E DETAILS FOR NUMERICAL EXPERIMENTS

### E.1 TRANSFORMER ARCHITECTURE AND ALGORITHM 1 IMPLEMENTATION

#### E.1.1 TRANSFORMER FOR SEQUENTIAL DECISION MAKING

We adopt the transformer architecture for regression tasks from [22] to solve sequential decision-making tasks, with modifications tailored to our specific setting.

**Prompt.** At time  $t$ , the prompt consists of two types of elements derived from the history  $H_t$ : (i) the “feature” elements, which stack the context  $X_\tau \in \mathbb{R}^d$  at each time  $\tau \leq t$  and the observation  $o_{\tau-1} \in \mathbb{R}^k$  from  $\tau - 1$  (with  $o_0$  set as a  $k$ -dimensional zero vector in our experiments), i.e.,  $\{(X_\tau, o_{\tau-1})\}_{\tau=1}^t \subset \mathbb{R}^{d+k}$ . These elements serve as the “features” in the prediction; and (ii) the “label” elements, which are the actions  $\{a_\tau\}_{\tau=1}^{t-1}$  in  $H_t$ . Therefore, the prompt contains  $2t - 1$  elements in total.

**Architecture.** The transformer is based on the GPT-2 family [45]. It uses two learnable linear transformations to map each “feature” and “label” element into vectors in the embedding space respectively (similar to tokens in language models). These vectors are processed through the GPT-2’s attention mechanism, resulting in a vector that encapsulates relevant contextual information. This vector then undergoes another learnable linear transformation, moving from the embedding space to the action  $\mathcal{A}$ , ultimately resulting in the prediction of  $a_t^*$ . In our experiments, the GPT-2 model has 12 layers, 16 attention heads, and a 256-dimensional embedding space. The experiments are conducted on 2 A100 GPUs with `DistributedDataParallel` method of Pytorch.

#### E.1.2 ALGORITHM 1 IMPLEMENTATION AND PRE-TRAINING DETAILS

During pre-training, we use the AdamW optimizer with a learning rate of  $10^{-4}$  and a weight decay of  $10^{-4}$ . The dropout rate is set to 0.05. For implementing Algorithm 1 in our experiments, we set the total iterations  $M = 130$  with an early training phase of  $M_0 = 50$ . The number of training sequences per iteration is  $n = 1500 \times 64$ , which are randomly split into 1500 batches with a batch size of 64. The transformer’s parameter  $\theta$  is optimized to minimize the averaged loss of each batch. As suggested in Proposition B.2, we select the cross-entropy loss for the multi-armed bandits, squared loss for the dynamic pricing, and absolute loss for the linear bandits and newsvendor problem.

To reduce computation costs from sampling histories during the early training phase, instead of sampling a new dataset  $D_m$  in each iteration (as described in Algorithm 1), we initially sample  $10^6$  data samples before the pre-training phase as an approximation of  $\mathcal{P}_{\gamma, f}$ . Each sample in the batch is uniformly sampled from these  $10^6$  data samples.

To reduce transformer inference costs during the mixed training phase, we reset the number of training sequences to  $n = 15 \times 64$  per iteration and set the ratio  $\kappa = 1/3$ , meaning  $10 \times 64$  samples are from  $\mathcal{P}_{\gamma, \text{TF}_{\theta_m}}$  and  $5 \times 64$  samples are from  $\mathcal{P}_{\gamma, f}$  (sampled from the pool of pre-generated  $10^6$

1890 samples as before). Consequently, we adjust the number of batches from 1500 to 50 to fit the smaller  
 1891 size of training samples, keeping the batch size at 64.

1892 **Decision Function  $f$  in Pre-Training.** To mitigate the OOD issue mentioned in Section 3, we aim  
 1893 for the decision function  $f$  to approximate  $\text{Alg}^*$  or  $\text{TF}_{\hat{\theta}}$ . However, due to the high computation cost  
 1894 of  $\text{Alg}^*$  and the unavailability of  $\text{TF}_{\hat{\theta}}$ , we set  $f(H_t) = a_t^* + \epsilon'_t$ , where  $\epsilon'_t$  is random noise simulating  
 1895 the suboptimality of  $\text{Alg}^*$  or  $\text{TF}_{\hat{\theta}}$  and independent of  $H_t$ . As we expect such suboptimality to  
 1896 decrease across  $t$ , we also reduce the influence of  $\epsilon'_t$  across  $t$ . Thus, in our experiments,  $\epsilon'_t$  is defined  
 1897 as:

$$1898 \epsilon'_t \sim \begin{cases} 0 & \text{w.p. } \max\{0, 1 - \frac{2}{\sqrt{t}}\}, \\ \text{Unif}[-1, 1] & \text{w.p. } \min\{1, \frac{2}{\sqrt{t}}\}, \end{cases}$$

1901 where for the multi-armed bandits we replace  $\text{Unif}[-1, 1]$  by  $\text{Unif}\{-2, -1, 1, 2\}$  and for the linear  
 1902 bandits with dimension  $d > 1$  we apply the uniform distribution of the set  $[-1, -1]^d$ . We further  
 1903 project  $f(H_t)$  into  $\mathcal{A}$  when  $f(H_t) \notin \mathcal{A}$ .

1904 **Curriculum.** Inspired by Garg et al. [22] in the regression task, we apply curriculum training to  
 1905 potentially speed up Algorithm 1. This technique uses “simple” task data initially and gradually  
 1906 increases task complexity during the training. Specifically, we train the transformer on samples with  
 1907 a smaller horizon  $\tilde{T} = 20$  (generated by truncating samples from  $\mathcal{P}_{\gamma, f}$  or  $\mathcal{P}_{\gamma, \text{TF}_{\theta_m}}$ ) at the beginning  
 1908 and gradually increase the sample horizon to the target  $T = 100$ . We apply this curriculum in both  
 1909 the early training phase ( $m \leq 50$ ) and the mixed training phase ( $m > 50$ ). The last 30 iterations  
 1910 focus on training with  $\tilde{T} = 100$  using non-truncated samples from  $\mathcal{P}_{\gamma, \text{TF}_{\theta}}$  to let the transformer fit  
 1911 more on the non-truncated samples. The exact setting of  $\tilde{T}$  is as follows:

$$1912 \tilde{T} = \begin{cases} 20 \times (m \% 10 + 1) & \text{when } m = 1, \dots, 50, \\ 20 \times (m \% 10 - 4) & \text{when } m = 51, \dots, 100, \\ 100 & \text{when } m = 100, \dots, 130. \end{cases}$$

## 1916 E.2 ENVIRONMENT GENERATION

1917 Throughout the paper, we set the horizon  $T = 100$  for all tasks. The environment distribution  $\mathcal{P}_{\gamma}$  is  
 1918 defined as follows for each task:

- 1921 • **Stochastic multi-armed bandit:** We consider the number of actions/arms  $k = 20$ . The  
 1922 environment parameter  $\gamma$  encapsulates the expected reward  $r_a$  of each arm  $a = 1, \dots, 20$ ,  
 1923 where the reward of arm  $a$  is sampled from  $\mathcal{N}(r_a, 0.2)$ . Thus, the environment distribution  
 1924  $\mathcal{P}_{\gamma}$  is defined by the (joint) distribution of  $(r_1, \dots, r_{20})$ . We set  $r_a \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$  for each  
 1925 action  $a$ . The optimal arm is  $a^* = \arg \max_a r_a$ .
- 1926 • **Linear bandits:** We consider the dimension of actions/arms  $d = 2$  and  $\mathcal{A} = \mathcal{B}(0, 1)$ , i.e.,  
 1927 a unit ball centered at the origin (with respect to the Euclidean norm). The environment  
 1928 parameter  $\gamma$  encapsulates the (linear) reward function’s parameter  $w \in \{\tilde{w} \in \mathbb{R}^2 : \|\tilde{w}\|_2 = 1\}$ ,  
 1929 where the reward of arm  $a$  is sampled from  $\mathcal{N}(w^\top a, 0.2)$ . Thus, the environment  
 1930 distribution  $\mathcal{P}_{\gamma}$  is defined by the distribution of  $w$ . We set  $w$  to be uniformly sampled from  
 1931 the unit sphere, and the corresponding optimal arm being  $a^* = w$ .
- 1932 • **Dynamic pricing:** We set the noise  $\epsilon_t \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 0.2)$  and the context  $X_t$  to be sampled i.i.d.  
 1933 and uniformly from  $[0, 5/2]^d$ , with  $d = 6$  as the dimension of contexts, and  $\mathcal{A} = [0, 30]$ .  
 1934 We consider the linear demand function family  $d(X_t, a_t) = w_1^\top X_t - w_2^\top X_t \cdot a_t + \epsilon_t$ , where  
 1935  $w_1, w_2 \in \mathbb{R}^d$  are the demand parameters. Thus, the environment distribution  $\mathcal{P}_{\gamma}$  is defined  
 1936 by the (joint) distribution of  $(w_1, w_2)$ . We set  $w_1$  to be uniformly sampled from  $[1/2, 3/2]^6$   
 1937 and  $w_2$  to be uniformly sampled from  $[1/20, 21/20]^6$ , independent of  $w_1$ . The optimal  
 1938 action  $a_t^* = \frac{w_1^\top X_t}{2 \cdot w_2^\top X_t}$ . We also consider a square demand function family, which will be  
 1939 specified later, to test the transformer’s performance on a mixture of different demand tasks.
- 1940 • **Newsvendor problem:** We set the context  $X_t$  to be sampled i.i.d. and uniformly from  
 1941  $[0, 3]^d$  with  $d = 4$ , and  $\mathcal{A} = [0, 30]$ . We consider the linear demand function family  
 1942  $d(X_t, a_t) = w^\top X_t + \epsilon_t$ , where  $w \in \mathbb{R}^2$ . The environment parameter  $\gamma$  encapsulates (i)  
 1943 the upper bound  $\bar{\epsilon}$  of the noise  $\epsilon_t$ , where  $\epsilon_t \stackrel{\text{i.i.d.}}{\sim} \text{Unif}(0, \bar{\epsilon})$ ; (ii) the left-over cost  $h$  (with

the lost-sale cost  $l$  being 1); and (iii) the demand parameter  $w$ . Accordingly, we set (i)  $\bar{\epsilon} \sim \text{Unif}[1, 10]$ ; (ii)  $h \sim \text{Unif}[1/2, 2]$ ; and (iii)  $w$  uniformly sampled from  $[0, 3]^2$ . The optimal action can be computed by  $a_t^* = w^\top X_t + \frac{\bar{\epsilon}}{1+h}$  [18], which is indeed the  $\frac{1}{1+h}$  quantile of the the random variable  $w^\top X_t + \epsilon_t$ . We also consider a square demand function family, which will be specified later, to test the transformer’s performance on a mixture of different demand tasks.

**Finite Pool of Environments.** To study the behavior and test the performance of the transformer on finite possible environments (e.g., to see if and how  $\text{TF}_{\hat{\theta}}$  converges to  $\text{Alg}^*$ ), we also consider a finite pool of environments as the candidates of  $\gamma$ . Specifically, for some tasks we first sample finite environments (e.g.,  $\gamma_1, \dots, \gamma_4$ ) i.i.d. following the sampling rules previously described. We then set the environment distribution  $\mathcal{P}_\gamma$  as a uniform distribution over the pool of sampled environments (e.g.,  $\{\gamma_1, \dots, \gamma_4\}$ ). We should notice this pool of environments does not restrict the context generation (if any): in both the pre-training and testing, contexts are generated following the rules previously described and are independent of the given finite pool.

**Tasks with Two Demand Types.** To test the transformer’s performance on tasks with a mixture of different demand types, we also consider the square demand function family for the dynamic pricing and newsvendor tasks. Specifically, besides the linear demand function family, we also consider  $d^{\text{sq}}(X_t, a_t) = (w_1^\top X_t)^2 - (w_2^\top X_t) \cdot a_t + \epsilon_t$  in dynamic pricing and  $d^{\text{sq}}(X_t, a_t) = (w^\top X_t)^2 + \epsilon_t$  in the newsvendor problem. Thus, with a slight abuse of notation, we can augment  $\gamma$  such that it parameterizes the type of demand (linear or square). In experiments related to multi-type demands, the probability of each type being sampled is  $1/2$ , while the distributions over other parameters of the environment remain the same as in the linear demand case.

### E.3 BENCHMARK ALGORITHMS

#### E.3.1 MULTI-ARMED BANDITS

- Upper Confidence Bound (UCB) [35]: Given  $H_t$ , the action  $a_t$  is chosen by  $a_t = \arg \max_{a \in \mathcal{A}} \hat{r}_a + \frac{\sqrt{2 \log T}}{\min\{1, n_a\}}$ , where  $n_a$  is the number of pulling times of arm  $a$  before time  $t$  and  $\hat{r}_a$  is the empirical mean reward of  $a$  based on  $H_t$ .
- Thomspon sampling (TS) [47]: Given  $H_t$ , the action  $a_t$  is chosen by  $a_t = \arg \max_{a \in \mathcal{A}} \tilde{r}_a$ , where each  $\tilde{r}_a \sim \mathcal{N}\left(\hat{r}_a, \frac{\sqrt{2 \log T}}{\min\{1, n_a\}}\right)$  is the sampled reward of arm  $a$  and  $\hat{r}_a, n_a$  follow the same definitions as in UCB.
- $\text{Alg}^*$ : Given  $H_t$  from a task with a pool of  $|\Gamma|$  environments  $\{\gamma_1, \dots, \gamma_{|\Gamma|}\}$ , the action  $a_t$  is chosen by the posterior sampling defined in Algorithm 2. The posterior distribution can be computed by

$$\mathcal{P}(\gamma_i | H_t) = \frac{\exp(-\frac{1}{\sigma^2} \sum_{\tau=1}^{t-1} (o_\tau - r_{a_\tau}^i)^2)}{\sum_{i'=1}^{|\Gamma|} \exp(-\frac{1}{\sigma^2} \sum_{\tau=1}^{t-1} (o_\tau - r_{a_\tau}^{i'})^2)},$$

where  $r_a^i$  is the expected reward of  $a$  in environment  $\gamma_i$  and  $\sigma^2$  is the variance of the noise (which equals to 0.2 in our experiments).

#### E.3.2 LINEAR BANDITS

- LinUCB [14]: Given  $H_t$ , we define  $\Sigma_t = \sum_{\tau=1}^{t-1} a_\tau a_\tau^\top + \sigma^2 I_d$ , where  $\sigma^2$  is the variance of the reward noise. The action  $a_t$  is chosen by  $a_t \in \arg \max_{a \in \mathcal{A}} \hat{w}_t^\top a + \sqrt{2 \log T} \|a\|_{\Sigma_t^{-1}}$ , where  $\hat{w}_t = \Sigma_t^{-1} (\sum_{\tau=1}^{t-1} o_\tau \cdot a_\tau)$  is the current estimation of  $w$  from  $H_t$ .
- LinTS [1]: Given  $H_t$ , the action  $a_t$  is chosen by  $a_t = \arg \max_{a \in \mathcal{A}} \tilde{w}_t^\top a$ , where  $\tilde{w}_t \sim \mathcal{N}(\hat{w}_t, \sqrt{2 \log T} \Sigma_t^{-1})$  as the sampled version of  $w$  and  $\hat{w}_t, \Sigma_t$  follow the same definitions in LinUCB.
- $\text{Alg}^*$ : Given  $H_t$  from a task with a pool of  $|\Gamma|$  environments  $\{\gamma_1, \dots, \gamma_{|\Gamma|}\}$ , the action  $a_t$  is chosen by the posterior median defined in Algorithm 2. The posterior distribution can be

1998 computed by

$$1999 \mathcal{P}(\gamma_i|H_t) = \frac{\exp(-\frac{1}{\sigma^2} \sum_{\tau=1}^{t-1} (o_\tau - w_i^\top a_\tau)^2)}{2000 \sum_{i'=1}^{|\Gamma|} \exp(-\frac{1}{\sigma^2} \sum_{\tau=1}^{t-1} (o_\tau - w_{i'}^\top a_\tau)^2)},$$

2002 where  $w_i$  is the reward function parameter in environment  $\gamma_i$  and  $\sigma^2$  is the variance of the  
2003 noise (which equals to 0.2 in our experiments).

### 2005 E.3.3 DYNAMIC PRICING

2006 All the benchmark algorithms presented below assume the demand model belongs to the linear  
2007 demand function family as defined in Appendix E.2 (which can be mis-specified when we deal with  
2008 dynamic pricing problems with two demand types).

- 2010 • Iterative least square estimation (ILSE) [44]: At each  $t$ , it first estimates the unknown  
2011 demand parameters  $(w_1, w_2)$  by applying a ridge regression based on  $H_t$ , and chooses  $a_t$  as  
2012 the optimal action according to the estimated parameters and the context  $X_t$ . Specifically,  
2013 we denote  $(\hat{w}_{1,t}, -\hat{w}_{2,t}) = \Sigma_t^{-1} (\sum_{\tau=1}^{t-1} o_\tau \cdot z_\tau)$  as the estimation of  $(w_1, -w_2)$  through  
2014 a ridge regression, where  $z_\tau = (X_\tau, a_\tau \cdot X_\tau)$  serves as the “feature vector” and  $\Sigma_t =$   
2015  $\sum_{\tau=1}^{t-1} z_\tau z_\tau^\top + \sigma^2 I_d$  ( $\sigma^2$  is the variance of the demand noise). Then the action is chosen as  
2016 the optimal one by treating  $(\hat{w}_{1,t}, \hat{w}_{2,t})$  as the true parameter:  $a_t = \frac{\hat{w}_{1,t}^\top X_t}{2\hat{w}_{2,t}^\top X_t}$ .
- 2018 • Constrained iterated least squares (CILS) [32]: It is similar to the ILSE algorithm except  
2019 for the potential explorations when pricing. Specifically, we follow the notations in ILSE  
2020 and denote  $\hat{a}_t = \frac{\hat{w}_{1,t}^\top X_t}{2\hat{w}_{2,t}^\top X_t}$  as the optimal action by treating  $(\hat{w}_{1,t}, \hat{w}_{2,t})$  as the true parameter  
2021 (the chosen action of ILSE), and denote  $\bar{a}_{t-1}$  as the empirical average of the chosen actions  
2022 so far. Then the CILS chooses

$$2024 a_t = \begin{cases} \bar{a}_{t-1} + \text{sgn}(\delta_t) \frac{t^{-\frac{1}{4}}}{10}, & \text{if } |\delta_t| < \frac{1}{10} t^{-\frac{1}{4}}, \\ \hat{a}_t, & \text{otherwise,} \end{cases}$$

2027 where  $\delta_t = \hat{a}_t - \bar{a}_{t-1}$ . The intuition is that if the tentative price  $\hat{a}_t$  stays too close to the  
2028 history average, it will introduce a small perturbation around the average as price experi-  
2029 mentation to encourage parameter learning.

- 2030 • Thomson sampling for pricing (TS) [51]: Like the ILSE algorithm, it also runs a regres-  
2031 sion to estimate  $(w_1, w_2)$  based on the history data, while the chosen action is the optimal  
2032 action of a sampled version of parameters. Specifically, we follow the notations in ILSE  
2033 and then TS chooses

$$2034 a_t = \frac{\tilde{\alpha}}{2 \cdot \tilde{\beta}},$$

2036 where  $(\tilde{\alpha}, \tilde{\beta}) \sim \mathcal{N}((\hat{w}_{1,t}^\top X_t, \hat{w}_{2,t}^\top X_t), \tilde{\Sigma}_t^{-1}) \in \mathbb{R}^2$  are the sampled parameters of the “in-  
2037 tercept” and “price coefficient” in the linear demand function and

$$2039 \tilde{\Sigma}_t = \left( \begin{pmatrix} X_t & \mathbf{0} \\ \mathbf{0} & X_t \end{pmatrix}^\top \Sigma_t^{-1} \begin{pmatrix} X_t & \mathbf{0} \\ \mathbf{0} & X_t \end{pmatrix} \right)^{-1}$$

2042 is the empirical covariance matrix given  $X_t$ .

- 2043 •  $\text{Alg}^*$ : Given  $H_t$  from a task with a pool of  $|\Gamma|$  environments  $\{\gamma_1, \dots, \gamma_{|\Gamma|}\}$ , the action  
2044  $a_t$  is chosen by the posterior averaging defined in Algorithm 2. To compute the posterior  
2045 distribution, we follow the notations in ILSE and denote  $w = (w_1, w_2)$  as the stacked  
2046 vector of parameters, then the posterior distribution is

$$2048 \mathcal{P}(\gamma_i|H_t) = \frac{\exp(-\frac{1}{\sigma^2} \sum_{\tau=1}^{t-1} (o_\tau - w_i^\top z_\tau)^2)}{2049 \sum_{i'=1}^{|\Gamma|} \exp(-\frac{1}{\sigma^2} \sum_{\tau=1}^{t-1} (o_\tau - w_{i'}^\top z_\tau)^2)},$$

2050 where  $w_i$  is the demand function parameter in environment  $\gamma_i$  and  $\sigma^2$  is the variance of the  
2051 noise (which equals to 0.2 in our experiments).

### E.3.4 NEWSVENDOR

All the benchmark algorithms presented below assume the demand model belongs to the linear demand function family as defined in Appendix E.2 (which can be mis-specified when we deal with newsvendor problems with two demand types).

- Empirical risk minimization (ERM) [7]: Since the optimal action  $a_t^*$  is the  $\frac{1}{1+h}$  quantile of the random variable  $w^\top X_t + \epsilon_t$  [18], ERM conducts a linear quantile regression based on the observed contexts and demands  $\{(X_\tau, o_\tau)\}_{\tau=1}^{t-1}$  to predict the  $\frac{1}{1+h}$  quantile on  $X_t$ .
- Feature-based adaptive inventory algorithm (FAI) [18]: FAI is an online gradient descent style algorithm aiming to minimize the cost  $\sum_{t=1}^T h \cdot (a_t - o_t)^+ + l \cdot (o_t - a_t)^+$  (we set  $l = 1$  in our experiments). Specifically, it chooses  $a_t = \tilde{w}_t^\top X_t$ , where

$$\tilde{w}_t = \begin{cases} \tilde{w}_{t-1} - \frac{h}{\sqrt{t}} \cdot x_{t-1}, & \text{if } o_{t-1} < a_{t-1}, \\ \tilde{w}_{t-1} + \frac{l}{\sqrt{t}} \cdot x_{t-1}, & \text{otherwise,} \end{cases}$$

is the online gradient descent step and  $\tilde{w}_0$  can be randomly sampled in  $[0, 1]^d$ .

- $\text{Alg}^*$ : Given  $H_t$  from a task with a pool of  $|\Gamma|$  environments  $\{\gamma_1, \dots, \gamma_{|\Gamma|}\}$ , the action  $a_t$  is chosen by the posterior median as shown in Algorithm 2. To compute the posterior distribution, we denote  $\bar{\epsilon}_\gamma$  and  $\beta_\gamma$  as the noise upper bound and demand function parameter of  $\gamma$  at  $\tau \leq t - 1$ , and define the event  $\mathcal{E}_{\gamma, \tau} = \{0 \leq o_\tau - \beta_\gamma^\top X_\tau \leq \bar{\epsilon}\}$  to indicate the feasibility of environment  $\gamma$  from  $(X_\tau, o_\tau)$ , and denote  $\bar{\mathcal{E}}_{\gamma, t} = \bigcap_{\tau=1}^{t-1} \mathcal{E}_{\gamma, \tau}$  to indicate the feasibility at  $t$ . Then the posterior distribution of the underlying environment is

$$\mathcal{P}(\gamma_i | H_t) = \frac{\mathbb{1}_{\bar{\mathcal{E}}_{\gamma_i, t}} \cdot \bar{\epsilon}_{\gamma_i}^{1-t}}{\sum_{i'=1}^{|\Gamma|} \mathbb{1}_{\bar{\mathcal{E}}_{\gamma_{i'}, t}} \cdot \bar{\epsilon}_{\gamma_{i'}}^{1-t}}.$$

## E.4 DETAILS FOR FIGURES

### E.4.1 DETAILS FOR FIGURE 2

**Figure 2 (a).** In Figure 2 (a), we independently implement Algorithm 1 to train two transformer models on a dynamic pricing task, which has an infinite support of the prior environment distribution  $\mathcal{P}_\gamma$  and 6-dimensional contexts (more details can be found in Appendix E.2). For both models, the training parameters are identical except for the  $M_0$  in Algorithm 1: the first curve (blue, thick line) uses  $M_0 = 50$ , while the second curve (orange, dashed line) uses  $M_0 = 130$ , meaning no transformer-generated data is utilized during training. All other parameters follow the configuration detailed in Appendix E.1. The figure shows the mean testing regret at each training iteration across 128 randomly sampled environments, with the shaded area representing the standard deviation. We further provide more results in Appendix C.1.2.

**Figure 2 (b).** In Figure 2 (b), we consider a dynamic pricing task with a pool of 8 linear demand functions and has 6-dimensional contexts. We generate 30 sequences from  $\text{TF}_{\theta_m}$  across different training iterations for this task: before using the transformer-generated data in the training ( $m = 40, 50$ ) and after using such data ( $m = 70, 80, 110, 120$ ). All generated sequences share the same context sequence  $\{X_t\}_{t=1}^T$  to control the effect of contexts on the chosen actions and the resulting observations. We stack the first 20 actions and observations into a single sample point and use t-SNE method to visualize these high-dimension points as shown in Figure 2 (b). We further apply the same method to generate points from  $\text{Alg}^*$  to study the behavior of TF.

We can observe that: (i) compared to the points from TF which have not been trained on self-generated data (i.e.,  $\text{TF}_{\theta_{40}}, \text{TF}_{\theta_{50}}$ ), the points from the trained ones are closer to the  $\text{Alg}^*$ 's, i.e., the expected decision rule of a well-trained TF; (ii) When being trained with more self-generated data, the points from  $\text{TF}_{\theta_m}$  get closer to the  $\text{Alg}^*$ 's. These observations suggest that utilizing the self-generated data (like in Algorithm 1) can help mitigate the OOD issue as discussed in Section 3: the data from  $\text{TF}_{\theta_m}$  can get closer to the target samples from  $\text{Alg}^*$  during the training and thus the pre-training data is closer to the testing data.

2106 E.4.2 DETAILS FOR FIGURE 3.  
2107

2108 For Figure 3, subfigure (a) is based on a dynamic pricing task with 6-dimensional contexts and (b)  
2109 is based on a newsvendor task with 4-dimensional contexts. Both tasks have a pool of 4 environ-  
2110 ments, i.e., the support of  $\mathcal{P}_\gamma$  only contains 4 environments, with linear demand function type. Each  
2111 subfigure is based on a sampled environment with a sampled sequence of contexts  $\{X_t\}_{t=1}^{30}$  from  
2112 the corresponding task. For each task, the data generation process follows the description detailed  
2113 in Appendix E.2, and the architecture of  $\text{TF}_\beta$  and the definition of  $\text{ALG}^*$  can be found in Appendix  
2114 E.1 and Appendix E.3 respectively. We further provide more results and discussions in Appendix  
2115 C.2.

2116  
2117  
2118  
2119  
2120  
2121  
2122  
2123  
2124  
2125  
2126  
2127  
2128  
2129  
2130  
2131  
2132  
2133  
2134  
2135  
2136  
2137  
2138  
2139  
2140  
2141  
2142  
2143  
2144  
2145  
2146  
2147  
2148  
2149  
2150  
2151  
2152  
2153  
2154  
2155  
2156  
2157  
2158  
2159