



HAL
open science

Model selection for behavioral learning data and applications to contextual bandits

Julien Aubert, Louis Köhler, Luc Lehéricy, Giulia Mezzadri, Patricia Reynaud-Bouret

► **To cite this version:**

Julien Aubert, Louis Köhler, Luc Lehéricy, Giulia Mezzadri, Patricia Reynaud-Bouret. Model selection for behavioral learning data and applications to contextual bandits. 28th International Conference on Artificial Intelligence and Statistics (AISTATS), May 2025, Mai Khao, Thailand. hal-04946461

HAL Id: hal-04946461

<https://hal.science/hal-04946461v1>

Submitted on 18 Feb 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Model selection for behavioral learning data and applications to contextual bandits

Julien Aubert¹

Giulia Mezzadri³

Louis Köhler²

Patricia Reynaud-Bouret¹

Luc Lehéricy¹

¹Université Côte d’Azur, CNRS, LJAD, France

²Université Côte d’Azur, LJAD, France

³Columbia University, Cognition and Decision Lab, United States

Abstract

Learning for animals or humans is the process that leads to behaviors better adapted to the environment. This process highly depends on the individual that learns and is usually observed only through the individual’s actions. This article presents ways to use this individual behavioral data to find the model that best explains how the individual learns. We propose two model selection methods: a general hold-out procedure and an AIC-type criterion, both adapted to non-stationary dependent data. We provide theoretical error bounds for these methods that are close to those of the standard i.i.d. case. To compare these approaches, we apply them to contextual bandit models and illustrate their use on both synthetic and experimental learning data in a human categorization task.

1 INTRODUCTION

1.1 Modeling Learning

From a behavioral perspective, learning can be defined as “a process by which an organism benefits from experience so that its future behavior is better adapted to its environment” (Rescorla, 1988). What if we want to model the learning process of an individual based only on the observations of its actions? That is, given an experiment of a learning task and some behavioral data of an individual performing that task, we want to find a model that explains the behavior and actions

of that individual. This question lies in the computational modeling of behavioral data (also referred to as cognitive modeling) (Wilson and Collins, 2019; Farrell and Lewandowsky, 2018; Collins and Shenhav, 2022; Mezzadri, 2020; Busemeyer and Diederich, 2010).

By its very nature, learning is a non stationary and dependent process. An extensive theory on statistical estimation and model selection exists for independent and stationary data (see Section 1.4). On the other hand, there are very few theoretical studies of statistical methods for non stationary and non independent data: in (Aubert et al., 2023), the properties of the MLE are studied for the **Exp3** model (Auer et al., 2002) on learning data; in (Aubert et al., 2024), a very general model selection procedure is presented that can be applied to non stationary data but works with restrictive assumptions on the models.

In this work, we provide a new oracle inequality for a hold-out procedure specifically designed for non-stationary data that offers flexibility for a wide range of non-stationary learning scenarios. However this method does not encompass parameter estimation in each model. Hence we also provide a AIC-like penalized log-likelihood estimation. The corresponding theoretical result is an oracle inequality derived by applying (Aubert et al., 2024). It requires assumptions on the parametrization on the models that hold in particular for contextual bandits.

1.2 Contextual Bandits as Models of Learning

A contextual bandit algorithm (Auer et al., 2002; Lattimore and Szepesvári, 2020) is a decision-making framework where, at each time step, the learner observes contextual information, selects an action and receives feedback in the form of a reward based on both the chosen action and the context. Unlike the standard multi-armed bandit problem, contextual bandits

incorporate additional information (the context) from the environment to guide action selection, allowing the model to adapt its policy dynamically based on the current situation. The usual goal of a reinforcement algorithm is to find an optimal policy that maximizes cumulative rewards. Contextual bandits have broad applications across machine learning, including recommendation systems, healthcare decision-making, and personalized medicine (Bouneffouf and Rish, 2019).

In this work, we use contextual bandits as models of learning. They are gaining popularity in cognitive psychology (Lan and Baraniuk, 2016; Schulz et al., 2018) because they provide a simple yet effective framework for understanding how individuals adjust their decision based on past experiences and present contextual cues. Many traditional cognitive models, such as the Component Cue (Gluck and Bower, 1988) or the ALCOVE models (Kruschke, 1992), address similar problems and can be reinterpreted as contextual bandits.

One key advantage in using contextual bandits as models of learning is their flexibility—the “context” can be almost anything (see Section 5 for an example), making the framework highly adaptable to a wide range of learning situations. Additionally, their ability to efficiently balance exploration and exploitation is critical in uncertain environments. Their relatively loose modeling approach allows for tractable representations of complex behavioral data which makes them easier to implement than traditional cognitive models.

1.3 Contributions

In what follows, n is the number of observed choices and actions during the learning experiment.

- In Section 3, we show that for any finite family of models \mathcal{M} , a hold-out estimator satisfies an oracle inequality with an $\mathcal{O}((\log n + \log |\mathcal{M}|)/n)$ error bound, regardless of the nature of the models of learning.
- In Section 4, we consider an AIC-type criterion with a possibly infinite countable number of parametric models built from contextual bandits algorithms. We show an oracle inequality with an $\mathcal{O}(\log(n)/n)$ error bound and explain the pros and cons of the AIC-type criterion vs the hold-out procedure.
- Section 5 is devoted to numerical illustrations of both methods on both synthetic and experimental learning data in a categorization task. Here the models in competition are contextual bandits (see Appendix B for more details).
- We present examples of bandits algorithms for which assumptions of Section 4 are satisfied in Appendix A.
- In Appendix C, we discuss how bandits with ex-

pert advice can be used to model metalearning (Binz et al., 2023), which refers to the processes by which an individual acquires knowledge about its own learning abilities, strategies, and preferences. We provide the details to obtain model selection methods and guarantees for metalearning as well.

- The complete proofs of the theoretical results are given in Appendix D.

1.4 Related Work

Hold-out estimators are commonly applied in cognitive modeling for learning data (Mezzadri et al., 2022a; James et al., 2023) for arbitrating between models. Theoretical analysis of hold-out procedures in the literature generally assumes the data stationary and independent (Massart, 2007; Arlot and Lerasle, 2016; Arlot and Celisse, 2010). There exist some limited results for time-dependent data (Opsomer et al., 2001), but these involve very different approaches from ours. A key challenge in our context is that the training set is not independent from the validation set, which prevents the use of techniques such as V -fold cross-validation.

Section 4 is similar in design to the framework of Castellan (2003) for selecting the best histogram for density estimation or more generally to non asymptotic model selection (Massart, 2007). The main difference is that we are in a non stationary and non independent framework.

In Section 4, unlike standard approaches in reinforcement learning (RL) (Bubeck et al., 2012), our goal is not to improve regret or to develop algorithms optimizing rewards as in (Dimakopoulou et al., 2017; Foster et al., 2019; Pacchiano et al., 2020); instead, we aim to understand how an individual learns. We select the contextual bandit model that best fits an individual’s learning curve from their learning data, without assuming the individual understands the context-action relationship. Thus, we seek the most realistic model rather than an optimal one. This theoretical statistical problem was first studied in (Aubert et al., 2023). Our method goes further and accommodates misspecified models.

At first sight, this problem may seem similar to the perspectives of Imitation Learning (IL) and Inverse Reinforcement Learning (IRL). Our work consists in trying to reproduce the learning curve of an expert (the individual under observation). However, IL methods use data from an expert who has already mastered the task (Hussein et al., 2017) and learn to perform it by copying them, while IRL (Arora and Doshi, 2021) aims to infer an underlying reward function based on the expert’s observed behavior across multiple trajectories. In contrast, in the field of cognition, experi-

menters control the reward function and seek to infer an individual’s behavior based on a single learning trajectory.

Our work is close to Hüyük et al. (2022); Schulz et al. (2015, 2018) who estimate how a learner’s behavior evolves with Bayesian models, however they focus on reward estimation. Therefore we cannot compare their method (whose output is a reward function) with our method (whose output is a policy).

2 FRAMEWORK AND NOTATIONS

The methodology we use is similar to the one in (Wilson and Collins, 2019; Farrell and Lewandowsky, 2018; Daw, 2011). It starts by considering a family of models that are relevant for the experiment and observed behavior: each model provides a family of candidate distributions for the sequence of choices made by the individual. For each model, model fitting is usually done by maximum likelihood estimation (MLE). Finally, the best cognitive model is selected by hold-out or an Akaike-type information criterion (AIC). In the particular case of the hold-out procedure, we focus on the selection step, disregarding how the fitted models are obtained so that our results hold regardless of the method used to extract one distribution per model.

2.1 Notations

Given two integers $s \leq t$ and a sequence $(a_s)_{s \in \mathbb{Z}}$, we define $a_s^t = (a_s, \dots, a_t)$, with a_s^t being the empty sequence when $s > t$. Let \mathbb{N}^* denote the set of positive integers, and for any $n \in \mathbb{N}^*$, we write $[n] = \{1, \dots, n\}$. Lastly, we denote the natural logarithm by \log .

Let $n \geq 3$ be an integer. We observe the sequence of actions $(A_t)_{1 \leq t \leq n}$ defined on a Polish measure space $(\Omega, \mathcal{F}, \mu)$ and adapted to a filtration $(\mathcal{F}_t)_{1 \leq t \leq n}$. We denote the corresponding probability measure by \mathbb{P} and the expectation by \mathbb{E} .

Our goal is to estimate the successive conditional densities of A_t w.r.t. μ . If Ω is discrete and μ is the counting measure, these densities are given by the sequence

$$p_t^*(\cdot) = \mathbb{P}(A_t = \cdot \mid \mathcal{F}_{t-1}), \quad \forall t \in [n].$$

In the following discussion, we will assume that μ is a fixed measure on Ω . For general measured spaces $(\Omega, \mathcal{F}, \mu)$, we assume that the conditional density of A_t given \mathcal{F}_{t-1} with respect to μ exists, denoting it by $p_t^*(\cdot)$. Consequently, for all $x \in \Omega$, the sequence $(p_t^*(x))_{1 \leq t \leq n}$ is predictable with respect to the filtration. Let $p^* = (p_t^*)_{t \in [n]}$ represent the vector of all successive conditional densities.

2.2 Some Examples of Filtrations

The Filtration Depends Only on Past Observations. A simple scenario consists in assuming that an action at time t depends only on the history of past actions. In this case, the filtration is defined by $\mathcal{F}_t = \sigma(A_1^t)$ for $t \geq 1$, and \mathcal{F}_0 is the trivial σ -algebra with no prior information. Under this setup, the true density at time t , p_t^* , is the conditional density of A_t given the past actions A_1^{t-1} . For $t = 1$, p_1^* is the deterministic density of the first action since no prior actions have been taken.

This can occur as soon as the environment of the individual is fixed (see for instance (Rescorla, 1988; Barron et al., 2015)). Cognitive tasks in such situations usually look like a standard bandit problem, where the individual tries to optimize its reward by pulling arms. Such tasks have been used in humans for studies on addiction (Bouneffouf et al., 2017) or in rodents with the Skinner box (Skinner, 2019).

A classical model for such a task is the **Gradient Bandit** algorithm (Sutton, 2018; Mei et al., 2024), with learning rate θ . In this algorithm 1, the agent chooses an action A_t among a finite set of actions $[K]$ where K is a positive integer, and receives a reward $g_{A_t, t} > 0$ from the environment drawn from some probability distribution. The algorithm then uses policy gradient methods to directly optimize the probabilities of selecting each action.

Algorithm 1 Gradient Bandit (Mei et al., 2023)

Inputs: $n, \theta > 0, K \in \mathbb{N}^*$.

Initialization: $p_{\theta, 1} = (\frac{1}{K}, \dots, \frac{1}{K})$.

for $t \in [n]$ **do**

Draw an action $A_t \sim p_t$ and receive a reward $g_{A_t, t} > 0$.

Update for all $a \in [K]$,

$$p_{\theta, t}(a) = \frac{\exp\left(-\theta \sum_{s \in [t]} \hat{g}_{a, s}^\theta\right)}{\sum_{b \in [K]} \exp\left(-\theta \sum_{s \in [t]} \hat{g}_{b, s}^\theta\right)}$$

where $\hat{g}_{b, s}^\theta = (\mathbf{1}_{A_s = b} - p_{\theta, s}(b)) g_{A_s, s}$.

A model $(p_\theta)_{\theta \in \Theta}$ is said to be *well-specified* if there exist θ^* such that the true distribution $p^* = p_{\theta^*}$. Properties of the MLE for a well-specified model in this particular learning framework have been explored in (Aubert et al., 2023).

The Filtration Depends on Past Observations and Additional Variables. A more comprehensive approach to modeling learning (Wilson and Collins,

2019; Marr, 2010; Collins and Shenhav, 2022) involves the assumption that a learner’s decisions are influenced not only by past choices but also by observable contextual variables. Before making a decision at time t , the individual has access to context information X_t . The natural filtration in this framework is defined as $\mathcal{F}_{t-1} = \sigma(A_1^{t-1}, X_1^t)$, which captures both the learner’s previous actions A_1^{t-1} and all context variables X_1^t available up to time t .

2.3 Partial Log-likelihood

Given a sequence of distributions $p = (p_t)_{t \in [n]}$, we use as criterion the partial log-likelihood

$$\ell_n(p) = \sum_{t=N}^n \log p_t(A_t) \quad (1)$$

where the sum starts at $N = 1$ in the AIC-type framework, and some $N > 1$ for the hold-out procedure (in which the first $N - 1$ actions are used to estimate the distribution in each model and the last $n - N + 1$ to select a model). For the example of Section 2.2 where $\mathcal{F}_t = \sigma(A_1^t)$ for $t \geq 1$ and \mathcal{F}_0 is the trivial σ -algebra, $\ell_n(p)$ is exactly the log-likelihood $\log p(A_1^n)$. For the second example of Section 2.2, we use the term “partial” log-likelihood as per Cox (1975), because we are only interested in modeling the action process A_1^n and not the entire vector of observations $(X_1, A_1, \dots, X_n, A_n)$.

2.4 Stochastic Risk Function

Classical approaches (Massart, 2007; Spokoiny, 2012, 2017) typically define risk using an expectation of the contrast. For example, in i.i.d. scenarios, the log-likelihood is inherently linked to the Kullback-Leibler (KL) divergence between the estimated and the true distributions.

In line with Aubert et al. (2024), we introduce the stochastic risk function $\mathbf{K}_{N,n}$, defined as follows. For $1 \leq N \leq n$ and any sequence of conditional densities $p = (p_t)_{1 \leq t \leq n}$, we have:

$$\mathbf{K}_{N,n}(p) = \frac{1}{n - N + 1} \sum_{t=N}^n \mathbb{E} \left[\log \frac{p_t^*(A_t)}{p_t(A_t)} \middle| \mathcal{F}_{t-1} \right].$$

When $N = 1$, we simply write $\mathbf{K}_n(p)$. This expression represents the empirical mean of the conditional KL divergence, as the term

$$\mathbb{E} \left[\log \frac{p_t^*(A_t)}{p_t(A_t)} \middle| \mathcal{F}_{t-1} \right]$$

is a predictable quantity corresponding to the KL divergence between the distributions with densities p_t^* and p_t w.r.t. μ , conditionally to \mathcal{F}_{t-1} .

When $\mathcal{F}_t = \sigma(A_1^t)$ for $t \geq 1$ and \mathcal{F}_0 is the trivial σ -algebra, the quantity $(n - N + 1)\mathbb{E}[\mathbf{K}_{N,n}(p)]$ precisely equals the KL divergence between the distributions defined by p^* and p .

Similarly, we define the empirical mean of the Hellinger distance, denoted $\mathbf{H}_{N,n}^2$, as:

$$\mathbf{H}_{N,n}^2(p) = \frac{1}{n - N + 1} \sum_{t=N}^n \mathbb{E} [\mathbf{H}^2(p_t^*(A_t), p_t(A_t)) | \mathcal{F}_{t-1}]$$

with $\mathbf{H}^2(\cdot, \cdot)$ the squared Hellinger distance.

3 HOLD-OUT PROCEDURE

In this section, we assume to have access to a family of sequences of conditional densities $(p^m)_{m \in \mathcal{M}}$, where \mathcal{M} is a finite set with at least two elements ($|\mathcal{M}| \geq 2$), as is commonly encountered in hold-out procedures (Massart, 2007, Chapter 8). Each model of learning m is characterized by a sequence of conditional densities $p^m = (p_t^m)_{t \in [n]}$, where p_t^m serves as a candidate for approximating the true conditional density p_t^* . Specifically, under model $m \in \mathcal{M}$, $p_t^m(\cdot)$ represents some conditional density of A_t given the past information \mathcal{F}_{t-1} . The hold-out estimator \hat{m} is defined as follows. Let $n > N \geq 1$ and select

$$\hat{m} \in \arg \max_{m \in \mathcal{M}} \sum_{t=N}^n \log p_t^m(A_t).$$

Theorem 1. *Assume that for all $m \in \mathcal{M}$, for all $x \in \Omega$,*

- *for all $t \in [N]$, $p_t^m(x)$ is \mathcal{F}_{N-1} -measurable*
- *for all $t \in \{N, \dots, n\}$, $p_t^m(x)$ is \mathcal{F}_{t-1} -measurable.*

For all $\diamond > 1$, there exists $c > 0$ such that

$$\mathbb{E}[\mathbf{H}_{N,n}^2(p^{\hat{m}}) | \mathcal{F}_{N-1}] \leq \diamond \inf_{m \in \mathcal{M}} \mathbb{E}[\mathbf{K}_{N,n}(p^m) | \mathcal{F}_{N-1}] + c \frac{\log(n - N + 1) + \log |\mathcal{M}|}{n - N + 1}.$$

This result holds for arbitrary p^m as long as they are predictable w.r.t. $(\mathcal{F}_t)_t$. In particular, it allows, as usual for hold-out, to take $p^m = p_{\theta^m}^m$, where $\hat{\theta}^m \in \arg \max_{\theta^m \in \Theta^m} \sum_{t=1}^{N-1} \log p_{\theta^m, t}^m(A_t)$, for some family of models $(\Theta^m)_{m \in \mathcal{M}}$. This result is new, since there are no hold-out results for non-independent and non-stationary data up to our knowledge. However, it is the analog of Theorem 8.9 in (Massart, 2007) for this learning framework, adding only a multiplicative factor $\log n$ in the error bound. It justifies the use of

hold-out procedures to model learning data in cognitive experiments such as (Mezzadri et al., 2022a; James et al., 2023), using classical cognitive models as Alcové (Kruschke, 1992), Component-Cue (Gluck and Bower, 1988) or Activity-based Credit Assignment (see (James et al., 2023) and the references therein).

4 AIC-TYPE CRITERION

The hold-out estimator does not require prior knowledge of the underlying family of densities, but in practice (Wilson and Collins, 2019), it is more natural to use procedures that allow for both parameter estimation of a model and model selection, though they rely on structural assumptions about the models (Massart, 2007; Spokoiny, 2012; Aubert et al., 2024). In this section, we provide a new application of (Aubert et al., 2024, Theorem 1) to the contextual setting defined in the second example of Section 2.2. To simplify the framework, we consider a finite set of actions $[K]$. The filtration is therefore $\mathcal{F}_t = \sigma(A_1^{t-1}, X_1^t)$, where X_t is the context at time t that belongs to some context space \mathcal{X} . To stress out the dependency with the context at time t , we write the true unknown probability of picking an action A_t given X_t at time t as $p_t^*(A_t|X_t)$.

We model learning through *partition-based contextual bandit* algorithms (Lattimore and Szepesvári, 2020, Chapter 18), which allow for a straightforward comparison to the hold-out procedure in both theoretical (Section 4.2) and empirical settings (Section 5). This approach is linked to the classical problem of selecting the best partition when constructing a histogram (Castellan, 2003; Massart, 2007). This modeling also offers insight into how learners use contextual information to make decisions. While we focus on a simple scenario for the sake of simplicity, Appendix C extends this result to more complex bandit models for meta-learning.

4.1 Partition-based Contextual Bandits: An Example of Parametric Models

Partition-based contextual bandits (Lattimore and Szepesvári, 2020, Chapter 18) assume that the individual partitions the context space \mathcal{X} into disjoint cells C . This often occurs when the individual is familiar with the contexts and has developed a personal understanding of them. This allows the individual to learn a new task only by updating one elementary and non-contextual bandit algorithm (like **Gradient Bandit**), denoted **CellBandit**(C), in each context cell C .

Being able to select the partition used by the individual among several candidates is important to understand its behavior. For instance, in a categorization

task where contexts are objects (see Section 5), this approach reveals the perceived similarity between objects by the learner thanks to the estimated partition.

Formally, let $g_t = (g_{1,t}, \dots, g_{K,t}) \in [0, 1]^K$ represent the vector of losses (or rewards) at time t , which models the feedback from the environment. We do not impose specific assumptions on how losses are generated, except that g_t must be \mathcal{F}_{t-1} -measurable. Contexts X_t may also be generated independently of past actions or may depend on them.

Each model $m \in \mathcal{M}$ corresponds to a partition \mathcal{P}_m of \mathcal{X} into D_m cells. The model is parameterized by a vector $\theta^m = (\theta_C)_{C \in \mathcal{P}_m}$, where each **CellBandit**(C) uses a procedure characterized by a parameter θ_C , such as the learning rate in the **Gradient Bandit**. The resulting candidate for p^* is $p_{\theta^m}^m = (p_{\theta^m, t}^m)_{t \in [n]}$.

Under model m , each **CellBandit**(C) is updated each time $X_t \in C$ and its decisions depend only on the contexts and actions within the set $F_t(C) = \{s \in [t] : X_s \in C\}$, with cardinality $T_t^C = |F_t(C)|$. We denote the action distribution at time t for **CellBandit**(C) with parameter θ_C as $\pi_{C, T_t^C}^{\theta_C}$ (see Algorithm 2). Thus, for all $t \in [n]$ and $a \in [K]$:

$$\begin{aligned} p_{\theta^m, t}^m(a|X_t) &= \mathbb{P}_{\theta^m}^m(A_t = a | \mathcal{F}_{t-1}) \\ &= \sum_{C \in \mathcal{P}_m} \pi_{C, T_t^C}^{\theta_C}(a) \mathbf{1}_{X_t \in C}. \end{aligned} \quad (2)$$

Algorithm 2 Partition-based contextual bandit for model m (Lattimore and Szepesvári, 2020)

Inputs: partition \mathcal{P}_m of the context space \mathcal{X} , parameters $\theta^m = (\theta_C)_{C \in \mathcal{P}_m} \in \Theta^m = \otimes_{C \in \mathcal{P}_m} \Theta_C$, with Θ_C compact parametric set.

Initialization: For all $C \in \mathcal{P}_m$, for all $a \in [K]$, $\pi_{C, 1}^{\theta_C}(a) = 1/K$.

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

Learner observes context $X_t \in \mathcal{X}$ and finds $C \in \mathcal{P}_m$ such that $X_t \in C$.

Learner plays **CellBandit**(C) with parameter θ_C and samples action $A_t \sim \pi_{C, T_t^C}^{\theta_C}$.

Learner observes loss $g_{A_t, t}$ and updates the probability distribution $\pi_{C, T_t^C}^{\theta_C}$ in **CellBandit**(C).

First, we need to assume that the probabilities do not vanish.

Assumption 1. There exists $\varepsilon > 0$ and an integer $T_\varepsilon \geq 2$, such that, almost surely,

$$\forall t \leq T_\varepsilon, \forall x \in \mathcal{X}, \forall a \in [K], \quad p_t^*(a|x) \geq \varepsilon \quad (3)$$

and that for all $m \in \mathcal{M}$ and all $C \in \mathcal{P}_m$, the $\text{CellBandit}(C)$ satisfies, for all parameter $\theta_C \in \Theta_C$

$$\forall t \leq T_\varepsilon, \forall a \in [K], \pi_{C, T_t}^{\theta_C}(a) \geq \varepsilon. \quad (4)$$

Assumption 1 is relevant because once the true probability $p_t^*(a|x)$ of picking an action becomes too small and the learner stops making mistakes, further improvement in parameter estimation is no longer possible. This is emphasized in (Aubert et al., 2023), where the authors show a counterexample demonstrating that when the probability of error (picking the wrong action) is too low, the quality of estimation cannot be enhanced.

Let $\text{pen} : \mathcal{M} \rightarrow \mathbb{R}_+$ be a penalty function. For each $m \in \mathcal{M}$, let $\hat{\theta}^m \in \arg \max_{\theta^m \in \Theta^m} \ell_{T_\varepsilon}(p_{\hat{\theta}^m}^m)$ be a MLE of model m , with ℓ defined as in (1), and select a model \hat{m} that minimizes the penalized log-likelihood stopped at T_ε :

$$\hat{m} \in \arg \min_{m \in \mathcal{M}} \left(-\frac{\ell_{T_\varepsilon}(p_{\hat{\theta}^m}^m)}{T_\varepsilon} + \text{pen}(m) \right). \quad (5)$$

To prove oracle inequalities, we need a smoothness assumption on the parametrization of each $\text{CellBandit}(C)$ which can be extended to p^m in Proposition 2. Assumption 2 is standard for model selection and parameter estimation (Massart, 2007; Spokoiny, 2012).

Assumption 2. With the notation of Assumption 1, there exists $L_\varepsilon > 0$ such that, almost surely, for all $m \in \mathcal{M}$, all $C \in \mathcal{P}_m$, for all $\delta_C, \theta_C \in \Theta_C$, for all $t \leq T_\varepsilon$,

$$\sup_{a \in [K]} \left| \log \left(\frac{\pi_{C, T_t}^{\delta_C}(a)}{\pi_{C, T_t}^{\theta_C}(a)} \right) \right| \leq L_\varepsilon \|\delta_C - \theta_C\|_2. \quad (6)$$

Proposition 2. Assume that p^m is a partition-based contextual bandit as in (2) and Algorithm 2 and that there exists T_ε such that for all $C \in \mathcal{P}_m$, $\text{CellBandit}(C)$ satisfies (4) and (6). Then, almost surely, for all $\theta^m, \delta^m \in \Theta^m$, for all $t \leq T_\varepsilon$, for all $x \in \mathcal{X}$, for all $a \in [K]$, $p_{\theta^m, t}^m(a, x) \geq \varepsilon$ and

$$\sup_{a \in [K]} \left| \log \left(\frac{p_{\delta^m, t}^m(a|x)}{p_{\theta^m, t}^m(a|x)} \right) \right| \leq L_\varepsilon \sup_{C \in \mathcal{P}_m} \|\delta_C - \theta_C\|_2.$$

Finally, we make the following assumption.

Assumption 3. The number of parameters of all CellBandit procedures are uniformly bounded, so that $d = \sup_{m \in \mathcal{M}} \sup_{C \in \mathcal{P}_m} \dim(\Theta_C)$ is finite.

With these assumptions, one obtains the following result.

Theorem 3. Let \mathcal{M} be a countable set, and for each $m \in \mathcal{M}$, consider a partition-based contextual bandit model $\{p_{\theta^m}^m, \theta^m \in \Theta^m\}$ (see Algorithm 2 and (2)). Let R and r be such that all coordinates $\theta_{i,C}$'s of $\theta_C \in \Theta_C$, for $C \in \mathcal{P}_m$ and $m \in \mathcal{M}$, satisfy $r \leq \theta_{i,C} \leq R$ and let $A_\varepsilon = L_\varepsilon \sqrt{d}(R-r) + 2 \log(\varepsilon^{-1})$. Let $\Sigma_\varepsilon = \log(A_\varepsilon) \sum_{m \in \mathcal{M}} e^{-D_m} < +\infty$. Under Assumptions 1, 2 and 3, for any $\diamond > 1$, there exist $c, c' > 0$ such that the following holds: if for all $m \in \mathcal{M}$,

$$\text{pen}(m) \geq c A_\varepsilon^2 \log(\varepsilon^{-1}) \log(T_\varepsilon A_\varepsilon)^2 \frac{D_m}{T_\varepsilon},$$

then,

$$\begin{aligned} \mathbb{E}[\mathbf{K}_{T_\varepsilon}(p_{\hat{\theta}^{\hat{m}}}^{\hat{m}})] &\leq \\ \mathbb{E} \left[\diamond \inf_{m \in \mathcal{M}} \left(\inf_{\theta^m \in \Theta^m} \mathbf{K}_{T_\varepsilon}(p_{\theta^m}^m) + 2 \text{pen}(m) \right) \right] &+ c' A_\varepsilon \Sigma_\varepsilon \log(\varepsilon^{-1}) \frac{\log(T_\varepsilon)}{T_\varepsilon}. \end{aligned}$$

The set of models \mathcal{M} can be any set as long as it satisfies $\sum_{m \in \mathcal{M}} e^{-D_m} < +\infty$. This condition is standard in the classic model selection literature (Massart, 2007). For our application to partitions of \mathcal{X} , we may take any subset of the family of all possible partitions, but the number of partitions with a given dimension should not grow too fast with the dimension: for instance, it is impractical to take all possible partitions of \mathcal{X} since we incur a cost $\Sigma_\varepsilon/T_\varepsilon$ in the error bound.

Theorem 3 is an application of Aubert et al. (2024). However, despite the generality of their results, Aubert et al. (2024) show mainly applications to the classical settings of time-homogeneous models. We go further by providing a ready-to-use version for partition-based contextual bandits.

This result closely resembles the model selection “à la Birgé-Massart” (Massart, 2007, Section 7.4), with a bias-variance compromise and a penalty close to the variance in D_m/T_ε , up to additional logarithmic terms, $\log^2 T_\varepsilon$ in the penalty and $\log T_\varepsilon$ in the residual error. It follows from the general result of Aubert et al. (2024), applicable to dependent non-stationary data, though verifying its assumptions can be tedious. Partition-based contextual bandits easily meet these assumptions, such as those in (4) and (6) for the CellBandit (see Section A for examples).

Studying the properties of MLE typically requires regularity assumptions on the parameterization of the distribution (differentiability of the likelihood (Spokoiny, 2012), for instance). In particular, Aubert et al. (2024) assume a Lipschitz or Hölder condition on the parameterization. This condition is reflected in the case of a partition-based contextual bandit by Assumption 2.

This assumption is arduous to check for contextual bandits because errors in the parameters compound over time. In Appendix A, we prove that the **Exp3-IX** and **Gradient bandit** algorithms satisfy them. In addition, we also apply the result of Aubert et al. (2024) in the metalearning setup (Appendix C) for the Exp4 algorithm.

4.2 Method Comparison and Limitations

The hold-out method, while assumption-free regarding the family of densities, relies on two losses—Hellinger distance and KL divergence—which are standard in model selection, although they are in general not equivalent (see (Massart, 2007), Theorem 7.11).

Due to the data’s strong dependencies, the hold-out procedure performs a single split at $t = N$ unlike classical cross-validation. A trade-off is needed: N must be large enough to accurately perform model fitting but not too large to leave enough data for the model selection step. This approach is unsuited when the individual learns differently over time, such as when they change their behavior between the start and the end of the learning process (e.g. change the underlying partition in partition-based contextual bandits).

Unlike the hold-out, the AIC-type approach does not require sample splitting and performs well in practice (Section 5). However, its oracle inequality holds only for data from the time interval $[T_\varepsilon]$ (see Appendix A). This restriction is actually quite reasonable in practice, since nothing can be estimated about the learning process once it reaches a stage where it makes no errors (see also the negative results of Aubert et al. (2023)).

Both model selection methods rely on a hyperparameter N (the split position) or c (the penalty constant), which lack theoretical guidelines and must be empirically tuned. In the classical i.i.d. framework, Arlot and Celisse (2010) provide guidance on sample splitting in V-fold cross-validation. However, for our scenario—an individual learning task—there is no theoretical recommendation for choosing N for the hold-out criterion. We need to adjust N numerically as illustrated in Figure 2b.

For the AIC-type criterion, the penalty involves the constant c which is unknown beforehand and requires numerical calibration as well (see Section 5). We could either use the hold-out procedure from Section 3 or apply heuristics like the dimension jump method or slope heuristics (Baudry et al., 2012; Arlot, 2019) to determine c . In Section 5, we use simulations for this calibration, as shown in Figure 2c.

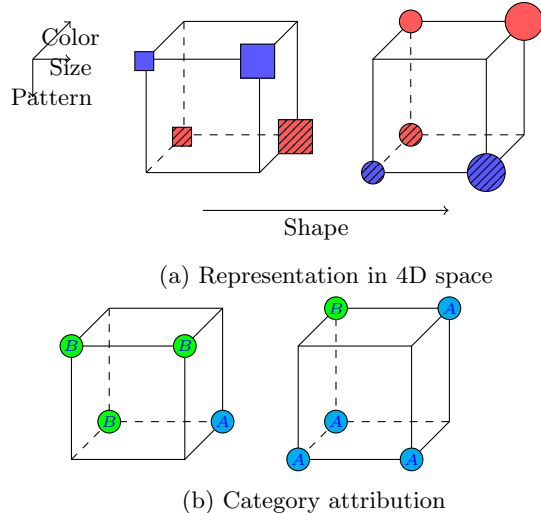


Figure 1: Experiment presentation: classic 5-4 category structure, widely used in cognition (Medin and Schaffer, 1978). In 1a, the 9 objects to classify represented in a 4D space with respect to their attributes: Color, Size, Filling Pattern, and Shape. In 1b, by position in the 4D space, the category attribution (A or B).

5 NUMERICAL ILLUSTRATIONS

We consider an experiment on the following categorization task: learners have to classify nine objects in two categories A and B in a sequential way with feedback after each individual choice. Figure 1 presents the objects and the classification rule the learners have to learn. It is a quite difficult task that has been experimented for instance in (Mezzadri et al., 2022b), where the learners needed about 300 trials to learn the classification rule.

For modeling, we fix the reward: 1 if the learner finds the good category and 0 in the other case. We focus on six models (detailed in Table 1 and Figure 4 of the Appendix) with clear cognitive interpretations. Each model represents a partition of the object space, where a **CellBandit** procedure is applied within each part. For example, the **OnePerItem** model is the most complex, with the finest partition, where each element of the object space forms its own subset.

On Synthetic Data. All the simulations were performed with $n = 500$ and for **Gradient Bandit** as **CellBandit**, for all the 6 models. The way synthetic data are generated can be found in Appendix B.4.

Figures 2b and 2c show that the hold-out is almost systematically outperformed by the penalized MLE. Both struggle to identify the most complex model **OnePerItem**, preferring simpler alternatives.

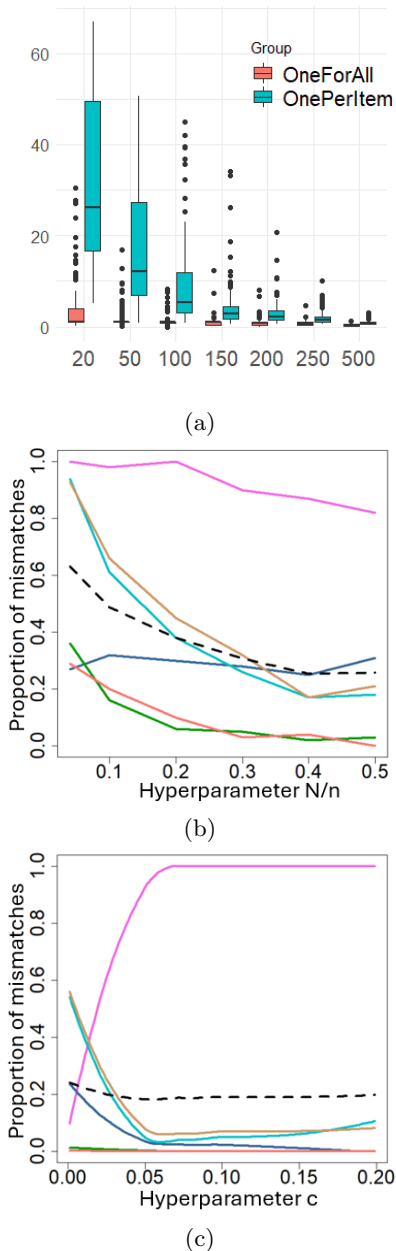


Figure 2: Errors of the procedures as a function of the tuning parameters. In 2a, average of the $|\hat{\theta}_C - \theta_C|/\theta_C$ over all cells C in model **OneForAll** and **OnePerItem** for the data generated respectively by the same models, where $\hat{\theta}_C$ is the MLE with likelihood truncated at N (in abscissa). In 2b and 2c, percentage of mismatch between \hat{m} and the simulated model over 100 simulations. The colors for each model are the ones given in Figure 3 whereas the average error on the models in the dash line. In 2b, for the hold-out estimator as a function of N/n . In 2c, for the penalized MLE with $\text{pen}(m) = c \log(n)^2 D_m/n$, as a function of c .

On the choice of hyperparameters. The hyperparameters (c for the penalized MLE and the proportion N/n in the hold-out) can be chosen either by minimizing the average of the misclassification errors in each model or by minimizing the maximum of these errors (2b and 2c). Given the nature of the plot, we chose the latter option. This led to the choice $c=0.012$ and $N/n=0.5$. We provide, in Appendix B.2, additional plots with two different choices of c : one that minimizes the average of the errors, and the other that corresponds to the AIC, that is $c \log(n)^2 = 1$.

The models in competition do not have the same complexity. In particular, **OnePerItem** has a much larger number of parameters than the other models. The penalty function disincentivizes models with large numbers of parameters, which explains why **OnePerItem** tends to be less selected in the simulations.

However, **OnePerItem** is the model which, in real life, means that participants have learned by heart without trying to find a rule (such as treating similar shapes in the same way). It has an important meaning, especially because Mezzadri et al. (2022b) have proven in various cognitive experiments, including this one, that some participants do not reason by rules, but by heart.

The proportion of mismatches for each model are reported in Figure 3a for the hold-out and 3b for the penalized MLE. Both methods manage to recover the true model with less than 35% of mistakes, except for the model **OnePerItem**, for which only the penalized MLE is able to achieve a successful match more than 60% of the time. The models that are confused the most are the ones that are able to correctly learn the categorization, that is **ByPatternExc**, **ByShapeExc** and **OnePerItem**.

Each method has a small bias: the penalized estimator prefers the **OnePerItem** model and has a tendency to select it even when it is wrong (Figure 3b), whereas the hold-out favors **ByShapeExc** (3a).

On Real Data. The data have been collected in (Mezzadri et al., 2022b)¹ and we focus only on the learning data. We use only the 176 participants that needed at least $n = 100$ trials. In Figure 3d, we see that most of participants are attributed one of the 3 models able to learn. The most frequent is **OnePerItem** (about 70% for the penalized MLE) and this percentage is larger than the one obtained on simulation, probably meaning that a significant proportion of the participants do not see the division along the dimensions **Shape** or **Pattern**. It would be interesting for

¹We refer the reader to (Mezzadri et al., 2022b) for precise description of the task as well as the ethics agreement.

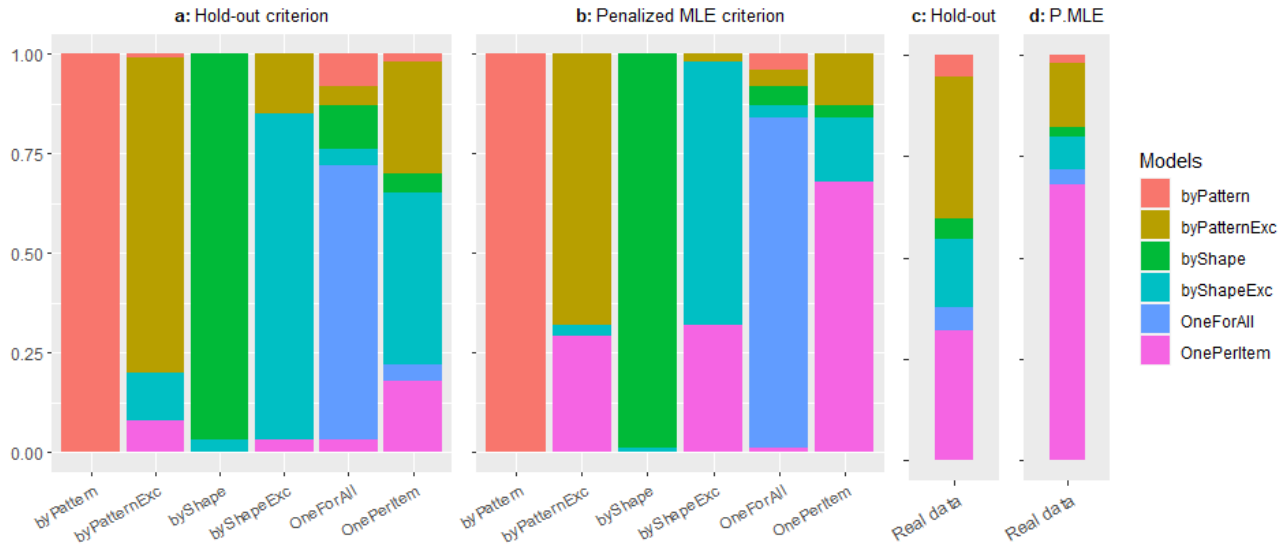


Figure 3: Distribution of the model choices. In **a**, hold-out with $N = 250$ over 100 simulations. In **b**, penalized MLE with $c = 0.012$ over 100 simulations. In **c**, hold-out on the data recorded in (Mezzadri et al., 2022b)–176 participants. In **d**, penalized MLE on the same experimental data.

further study to see if this is linked to the presentation order of the objects, as it has been proved for *Alcove* and *Component Cue* in (Mezzadri et al., 2022b).

6 CONCLUSION

We proposed two selection methods of models of learning that satisfy oracle inequalities: the hold-out method selects the best estimator in any finite family of estimators up to logarithmic terms; the AIC-like method selects the best trade-off between the model bias and the number of parameters of this model, under some parametrization assumptions. In all cases, the statistical challenge is that an individual’s learning data are dependent and non stationary, so that classical statistical model selection results do not apply. Future work involves a refinement of the model class to capture more complex learning phenomena (see as premisses, the Appendix C on metalearning).

Acknowledgements

This work was supported by the French government, through CNRS, the UCA *Jedi* and 3IA Côte d’Azur Investissements d’Avenir managed by the National Research Agency (ANR-15 IDEX-01 and ANR-19-P3IA-0002), directly by the ANR project ChaMaNe (ANR-19-CE40-0024-02) and by the interdisciplinary Institute for Modeling in Neuroscience and Cognition (NeuroMod).

References

- Sylvain Arlot. Minimal penalties and the slope heuristics: a survey. *J. SFdS*, 160(3):158–168, 2019. ISSN 2102-6238.
- Sylvain Arlot and Alain Celisse. A survey of cross-validation procedures for model selection. *Statistics Surveys*, 4:40–79, 2010.
- Sylvain Arlot and Matthieu Lerasle. Choice of V for V -fold cross-validation in least-squares density estimation. *J. Mach. Learn. Res.*, 17:Paper No. 208, 50, 2016. ISSN 1532-4435,1533-7928.
- Saurabh Arora and Prashant Doshi. A survey of inverse reinforcement learning: Challenges, methods and progress. *Artificial Intelligence*, 297:103500, 2021.
- Julien Aubert, Luc Lehéricy, and Patricia Reynaud-Bouret. On the convergence of the MLE as an estimator of the learning rate in the exp3 algorithm. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 1244–1275. PMLR, 7 2023.
- Julien Aubert, Luc Lehéricy, and Patricia Reynaud-Bouret. General oracle inequalities for a penalized log-likelihood criterion based on non-stationary data. Working paper or preprint, 5 2024.
- Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. The nonstochastic multiarmed bandit problem. *SIAM journal on computing*, 32(1): 48–77, 2002.

- Andrew B Barron, Eileen A Hebets, Thomas A Cleland, Courtney L Fitzpatrick, Mark E Hauber, and Jeffrey R Stevens. Embracing multiple definitions of learning. *Trends in neurosciences*, 38(7):405–407, 2015.
- Jean-Patrick Baudry, Cathy Maugis, and Bertrand Michel. Slope heuristics: overview and implementation. *Statistics and Computing*, 22:455–470, 2012.
- Alina Beygelzimer, John Langford, Lihong Li, Lev Reyzin, and Robert Schapire. Contextual bandit algorithms with supervised learning guarantees. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 19–26. JMLR Workshop and Conference Proceedings, 2011.
- Marcel Binz, Ishita Dasgupta, Akshay K Jagadish, Matthew Botvinick, Jane X Wang, and Eric Schulz. Meta-learned models of cognition. *Behavioral and Brain Sciences*, pages 1–38, 2023.
- Djallel Bouneffouf and Irina Rish. A survey on practical applications of multi-armed and contextual bandits. *arXiv preprint arXiv:1904.10040*, 2019.
- Djallel Bouneffouf, Irina Rish, and Guillermo A. Cecchi. Bandit models of human behavior: Reward processing in mental disorders. *CoRR*, abs/1706.02897, 2017.
- Sébastien Bubeck, Nicolo Cesa-Bianchi, et al. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends® in Machine Learning*, 5(1):1–122, 2012.
- Jerome R Busemeyer and Adele Diederich. *Cognitive modeling*. Sage, 2010.
- Gwénaëlle Castellán. Density estimation via exponential model selection. *IEEE Trans. Inform. Theory*, 49(8):2052–2060, 2003. ISSN 0018-9448,1557-9654.
- Dean S. Clark. Short proof of a discrete gronwall inequality. *Discrete Applied Mathematics*, 16(3):279–281, 1987. ISSN 0166-218X.
- Anne G.E. Collins and Amitai Shenhav. Advances in modeling learning and decision-making in neuroscience. *Neuropsychopharmacology*, 47(1):104–118, 2022.
- David R Cox. Partial likelihood. *Biometrika*, 62(2): 269–276, 1975.
- Nathaniel D. Daw. Trial-by-trial data analysis using computational models. *Decision making, affect, and learning: Attention and performance XXIII*, 23(1), 2011.
- Maria Dimakopoulou, Zhengyuan Zhou, Susan Athey, and Guido Imbens. Estimation considerations in contextual bandits. *arXiv preprint arXiv:1711.07077*, 2017.
- Simon Farrell and Stephan Lewandowsky. *Computational modeling of cognition and behavior*. Cambridge University Press, 2018.
- Dylan J Foster, Akshay Krishnamurthy, and Haipeng Luo. Model selection for contextual bandits. *Advances in Neural Information Processing Systems*, 32, 2019.
- Bolin Gao and Lacra Pavel. On the properties of the softmax function with application in game theory and reinforcement learning. *arXiv preprint arXiv:1704.00805*, 2017.
- Mark A Gluck and Gordon H Bower. Evaluating an adaptive network model of human learning. *Journal of memory and Language*, 27(2):166–195, 1988.
- Niels Richard Hansen, Patricia Reynaud-Bouret, and Vincent Rivoirard. Lasso and probabilistic inequalities for multivariate point processes. *Bernoulli*, 21(1), 2 2015.
- Christian Houdré and Patricia Reynaud-Bouret. Exponential inequalities for u-statistics of order two with constants. *Stochastic Inequalities and Applications. Progress in Probability*, 56, 01 2002.
- Ahmed Hussein, Mohamed Medhat Gaber, Eyad Elyan, and Chrisina Jayne. Imitation learning: A survey of learning methods. *ACM Computing Surveys (CSUR)*, 50(2):1–35, 2017.
- Alihan Hüyük, Daniel Jarrett, and Mihaela van der Schaar. Inverse contextual bandits: Learning how behavior evolves over time. In *International Conference on Machine Learning*, pages 9506–9524. PMLR, 2022.
- Ashwin James, Patricia Reynaud-Bouret, Giulia Mezzadri, Francesca Sargolini, Ingrid Bethus, and Alexandre Muzy. Strategy inference during learning via cognitive activity-based credit assignment models. *Scientific reports*, 13(1):9408, 2023.
- John K. Kruschke. Alcove: an exemplar-based connectionist model of category learning. *Psychological Review*, 99:22–44, 1992.
- Andrew S. Lan and Richard G. Baraniuk. A contextual bandits framework for personalized learning action selection. In *EDM*, pages 424–429, 2016.
- Tor Lattimore and Csaba Szepesvári. *Bandit Algorithms*. Cambridge University Press, 2020.
- David Marr. *Vision: A computational investigation into the human representation and processing of visual information*. MIT press, 2010.
- Pascal Massart. *Concentration inequalities and model selection: Ecole d’Eté de Probabilités de Saint-Flour XXXIII-2003*. Springer, 2007.

- Douglas L Medin and Marguerite M Schaffer. Context theory of classification learning. *Psychological review*, 85(3):207, 1978.
- Jincheng Mei, Zixin Zhong, Bo Dai, Alekh Agarwal, Csaba Szepesvari, and Dale Schuurmans. Stochastic gradient succeeds for bandits. In *International Conference on Machine Learning*, pages 24325–24360. PMLR, 2023.
- Jincheng Mei, Zixin Zhong, Bo Dai, Alekh Agarwal, Csaba Szepesvari, and Dale Schuurmans. Stochastic gradient succeeds for bandits, 2024.
- Giulia Mezzadri. *Statistical inference for categorization models and presentation order*. Phd thesis, Université Côte d’Azur, LJAD, France, 2020.
- Giulia Mezzadri, Thomas Laloë, Fabien Mathy, and Patricia Reynaud-Bouret. Hold-out strategy for selecting learning models: application to categorization subjected to presentation orders. *Journal of Mathematical Psychology*, 109:102691, 2022a.
- Giulia Mezzadri, Patricia Reynaud-Bouret, Thomas Laloë, and Fabien Mathy. Investigating interactions between types of order in categorization. *Scientific Reports*, 12(1):21625, 2022b.
- Gergely Neu. Explore no more: Improved high-probability regret bounds for non-stochastic bandits. *Advances in Neural Information Processing Systems*, 28, 2015.
- Jean Opsomer, Yuedong Wang, and Yuhong Yang. Nonparametric Regression with Correlated Errors. *Statistical Science*, 16(2):134 – 153, 2001.
- Aldo Pacchiano, My Phan, Yasin Abbasi Yadkori, Anup Rao, Julian Zimmert, Tor Lattimore, and Csaba Szepesvari. Model selection in contextual stochastic bandit problems. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 10328–10337. Curran Associates, Inc., 2020.
- Robert A Rescorla. Behavioral studies of pavlovian conditioning. *Annual review of neuroscience*, 1988.
- Herbert Robbins and Sutton Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951.
- Eric Schulz, Emmanouil Konstantinidis, and Maarten Speekenbrink. Learning and decisions in contextual multi-armed bandit tasks. In *CogSci*, 2015.
- Eric Schulz, Emmanouil Konstantinidis, and Maarten Speekenbrink. Putting bandits into context: How function learning supports decision making. *Journal of experimental psychology: learning, memory, and cognition*, 44(6):927, 2018.
- Burrhus Frederic Skinner. *The behavior of organisms: An experimental analysis*. BF Skinner Foundation, 2019.
- Vladimir Spokoiny. Parametric estimation. Finite sample theory. *Ann. Statist.*, 40(6):2877–2909, 2012. ISSN 0090-5364,2168-8966.
- Vladimir Spokoiny. Penalized maximum likelihood estimation and effective dimension. *Ann. Inst. Henri Poincaré Probab. Stat.*, 53(1):389–429, 2017. ISSN 0246-0203,1778-7017.
- Richard S Sutton. Reinforcement learning: An introduction. *A Bradford Book*, 2018.
- Robert C. Wilson and Anne G.E. Collins. Ten simple rules for the computational modeling of behavioral data. *Elife*, 8:e49547, 2019.

Checklist

1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes] The notations are described in Section 2. The set of models considered are also described in this section and specified for the bandit application in Section 4. Assumptions for the penalized log-likelihood criterion are also provided in this section. The algorithms that satisfy these assumptions are provided in Section A of the Appendix.
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes] Details about running time and complexity are available in Section B of the Appendix.
 - (c) Anonymized source code, with specification of all dependencies, including external libraries. [Yes] The code for the experiments is available in the zip file attached to the submission. The simulation of data according to each model is precisely explained in Section B of the Appendix, and the code to produce is given. The codes to compute both estimators (hold-out and penalized MLE) are given and commented, also in the supplementary material. The real data are taken from a published paper (Mezzadri et al., 2022b) and we asked the authors of (Mezzadri et al., 2022b) to provide us with these data. We don’t think it is possible to make these data public because the ethics agreement that has been signed by the authors of (Mezzadri et al., 2022b), prior to data collection, might not include

this possibility. However, this application to real data is mainly to prove that this can be done in practice. Since there is not a truth to be compared to in these data and since even cross validation is hard to perform, this real dataset cannot be used as a benchmark to compare methods anyway.

2. For any theoretical claim, check if you include:
 - (a) Statements of the full set of assumptions of all theoretical results. [Yes]
 - (b) Complete proofs of all theoretical results. [Yes]
 - (c) Clear explanations of any assumptions. [Yes] We stated precisely all our theoretical results with assumptions that are clearly referenced. We give intuition about how each assumption is used. The proofs are given in the supplementary material.
3. For all figures and tables that present empirical results, check if you include:
 - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Yes] All codes to generate synthetic data and to perform penalized MLE and hold-out are provided, so that all the numerical part on the calibration of both methods can be faithfully reproduced. Only the real dataset, as explained before, cannot be given for reproduction (Figure 3).
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes] The whole purpose of our numerical study in Section 5 is to explain the choice of hyperparameters (such as the splitting in the hold-out of the calibration of the constant c in the penalty).
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes] We have run our simulations on 600 independent simulated learners and we show with a boxplot (Figure 2a) and mismatch proportion graphs (Figure 2b, 2c and 3) the proportion of erroneous selections. This cannot be done on real data, since each real participant to the categorization task is unique.
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Yes] It is not central in our analysis so it is just mentioned in the supplementary material in Appendix B
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
 - (a) Citations of the creator If your work uses existing assets. [Yes] We clearly stated that the real data come from (Mezzadri et al., 2022b). The code has been developed by us solely, using classical packages in R that are clearly mentioned in the code and supplementary material.
 - (b) The license information of the assets, if applicable. [Not Applicable]
 - (c) New assets either in the supplemental material or as a URL, if applicable. [Not Applicable] We do not provide new packages associated to our results.
 - (d) Information about consent from data providers/curators. [Yes] Authors from (Mezzadri et al., 2022b) accepted that we use their datasets to perform our simulations.
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable] This work is theoretical. The methods that are validated theoretically here have already been in use in practice for a long time (see for instance the rules to follow for cognitive modeling in (Wilson and Collins, 2019)) and so the expected societal impact of the present work is negligible.
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
 - (a) The full text of instructions given to participants and screenshots. [No] We did not collect data for the present article but used data collected for (Mezzadri et al., 2022b), a work that is already published. In this article, all the details about instructions are given and we do not think it makes sense to reproduce it here for our illustration. We only kept the main description of the task so that the readers can understand what was done.
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [No] The data collection done for (Mezzadri et al., 2022b) had the approval of the local ethic committee as mentioned in their article. Here we do not feel necessary to reproduce this here but rather point towards (Mezzadri et al., 2022b) for additional information about the task and its ethic agreement.

- (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [No] In (Mezzadri et al., 2022b), the details about wages are given.

SUPPLEMENTARY MATERIAL

This document contains all the additional material for the article.

A EXAMPLES OF CELLBANDITS

In this section we provide examples of `CellBandit` satisfying (4) and (6). All the algorithms below are written for a cell C and a `CellBandit`(C) parameterized by $\theta_C \in \Theta_C$ compact subset of \mathbb{R}^d such that $R \geq \sup_{\theta_C \in \Theta_C} \|\theta_C\|_\infty$.

A.1 Example 1: Exp3-IX

This algorithm is a generalization of `Exp3` and was introduced in (Neu, 2015). Following (Aubert et al., 2023), we write `Exp3-IX` with parameters decreasing as a square root of the sample size to ensure a good MLE estimation of the parameters. Note in addition that, for `Exp3` and its variants, it is well known that sublinear convergence of the regret occurs when the learning rate η and the exploration term γ are decreasing as a square root of the sample size. This renormalization ensures that the learner is able to learn at a good pace and at the same time be robust to changes in the environment.

Algorithm 3 Exp3-IX(Neu, 2015) as a `CellBandit`(C)

Inputs: n (Sample size), $\theta_C = (\eta, \gamma) \in \Theta_C$ (Parameter), K (Number of actions).

Initialization: $\pi_{C,1}^{\theta_C} = (\frac{1}{K}, \dots, \frac{1}{K})$.

for $t \in F_n(C)$, the set of times where $X_s \in C$, **do**

Draw an action $A_t \sim \pi_{C,T_t}^{\theta_C}$ and receive a loss $g_{A_t,t} \in [0, 1]$.

Update for all $a \in [K]$,

$$\pi_{C,T_t+1}^{\theta_C}(a) = \frac{\exp\left(-\frac{\eta}{\sqrt{n}} \sum_{s \in F_t(C)} \hat{g}_{a,s}^{\theta_C}\right)}{\sum_{b \in [K]} \exp\left(-\frac{\eta}{\sqrt{n}} \sum_{s \in F_t(C)} \hat{g}_{b,s}^{\theta_C}\right)} \quad \text{where } \hat{g}_{b,s}^{\theta_C} = \frac{g_{b,s}}{\gamma/\sqrt{n} + \pi_{C,T_s}^{\theta_C}(b)} \mathbf{1}_{A_s=b}$$

In this case, $\Theta_C \subset \mathbb{R}^2$. When $\gamma = 0$, we recover the classical `Exp3` algorithm, studied from the MLE point of view in (Aubert et al., 2023). Note that while $g_{A_t,t}$ is observed and known, the estimated loss $\hat{g}_{b,s}^\theta$ depends on the parameterization. The following result shows that one can choose `Exp3-IX` as a `CellBandit` in the partition-based contextual bandits to perform partition selection.

Proposition 4. *Let $\varepsilon \in (0, 1/K)$ and let $\Theta_C \subset [0, R]^2$ with $R > 0$. Then `Exp3-IX` can be a `CellBandit`(C) with parameterization $\theta_C \in \Theta_C$ that satisfies (4) and (6), as soon as*

$$T_\varepsilon = \left\lceil \left(\frac{1}{K} - \varepsilon \right) \frac{\sqrt{n}}{R} \right\rceil \wedge n \quad \text{and} \quad L_\varepsilon = \frac{\sqrt{R^2/n + \varepsilon^2}}{\varepsilon^3 R} e^{1/\varepsilon^2}.$$

This shows that one can apply Theorem 3 with `Exp3-IX` as `CellBandit` as long as we stop using observations after \sqrt{n} time steps. The dependence in ε is not very critical, since it has been proved at least for `Exp3` in (Aubert et al., 2023), that in practice, we may take ε quite large (non-vanishing) with almost no impact on T_ε . This is a good thing since the theoretical dependency of L_ε in ε is quite pessimistic.

Limitations. This algorithm considers the horizon n fixed in order to renormalize the parameterization. From Proposition 4, it follows that Theorem 3 holds when only the first \sqrt{n} observations are used in the MLE, but this in no way means that the estimator will perform poorly when based on all data. Taking \sqrt{n} observations compounds on the usual issue that if the number of cells is large, only a small amount of data may be available for each cell, making estimation difficult.

A.2 Example 2: Gradient Bandit

`Gradient Bandit` is another possible algorithm. We still choose for similar reason a parameterization in η/\sqrt{n} , which echoes the Robbins-Monro conditions (Robbins and Monro, 1951) even if (Mei et al., 2023) proved convergence in a stochastic bandit framework even for non renormalized parameters.

Algorithm 4 Gradient Bandit (Mei et al., 2023) as a CellBandit

Inputs: n (Sample size), $\theta_C \in [r, R]$ (Parameter), K (Number of actions).

Initialization: $\pi_{C,1}^{\theta_C} = (\frac{1}{K}, \dots, \frac{1}{K})$.

for $t \in F_n(C)$ **do**

Draw an action $A_t \sim \pi_{C,T_t}^{\theta_C}$ and receive a reward $g_{A_t,t} \in [0, 1]$.

Update for all $a \in [K]$,

$$\pi_{C,T_t}^{\theta_C}(a) = \frac{\exp\left(-\frac{\theta_C}{\sqrt{n}} \sum_{s \in F_t(C)} \hat{g}_{a,s}^{\theta_C}\right)}{\sum_{b \in [K]} \exp\left(-\frac{\theta_C}{\sqrt{n}} \sum_{s \in F_t(C)} \hat{g}_{b,s}^{\theta_C}\right)} \quad \text{where } \hat{g}_{b,s}^{\theta_C} = \left(\mathbf{1}_{A_s=b} - \pi_{C,T_s}^{\theta_C}(b)\right) g_{A_s,s}$$

Proposition 5. Let $\varepsilon \in (0, 1)$ and let $\Theta_C \subset [0, R]^2$ with $R > 0$. Then, Gradient Bandit can be a CellBandit(C) with parameterization $\theta_C \in \Theta_C$ that satisfies (4) and (6), as soon as

$$T_\varepsilon := \left\lceil \log\left(\sqrt{\frac{1}{K\varepsilon}}\right) \frac{\sqrt{n}}{R} \right\rceil \wedge n \quad \text{and} \quad L_\varepsilon = \frac{\sqrt{2} \log\left(\sqrt{\frac{1}{K\varepsilon}}\right)}{R\varepsilon \sqrt{K\varepsilon}}.$$

This theoretical result has the same interpretation as before: the theoretical guarantees of Theorem 3 with Gradient Bandit as CellBandit hold when we stop using observations after \sqrt{n} time steps. In practice, we can use the observations up to time n (see Section 5).

B CODE AND DATA DESCRIPTION

The code for reproducing the figures in the article is available at the following link: <https://github.com/JulienAubert3/ContextualBandits>.

B.1 Contextual Bandit Based Models

Table 1: Description of models and their learning abilities

Model	Number of cells	Description of the cells	Learns categorization
OneForAll	1	One giant cell	No
ByShape	2	One for circles, one for squares	Partly
ByPattern	2	One for striped items, one for plain items	Partly
ByShapeExc	4	Cells from ByShape model with exceptions isolated	Yes
ByPatternExc	4	Cells from ByPattern model with exceptions isolated	Yes
OnePerItem	9	One cell for each item	Yes

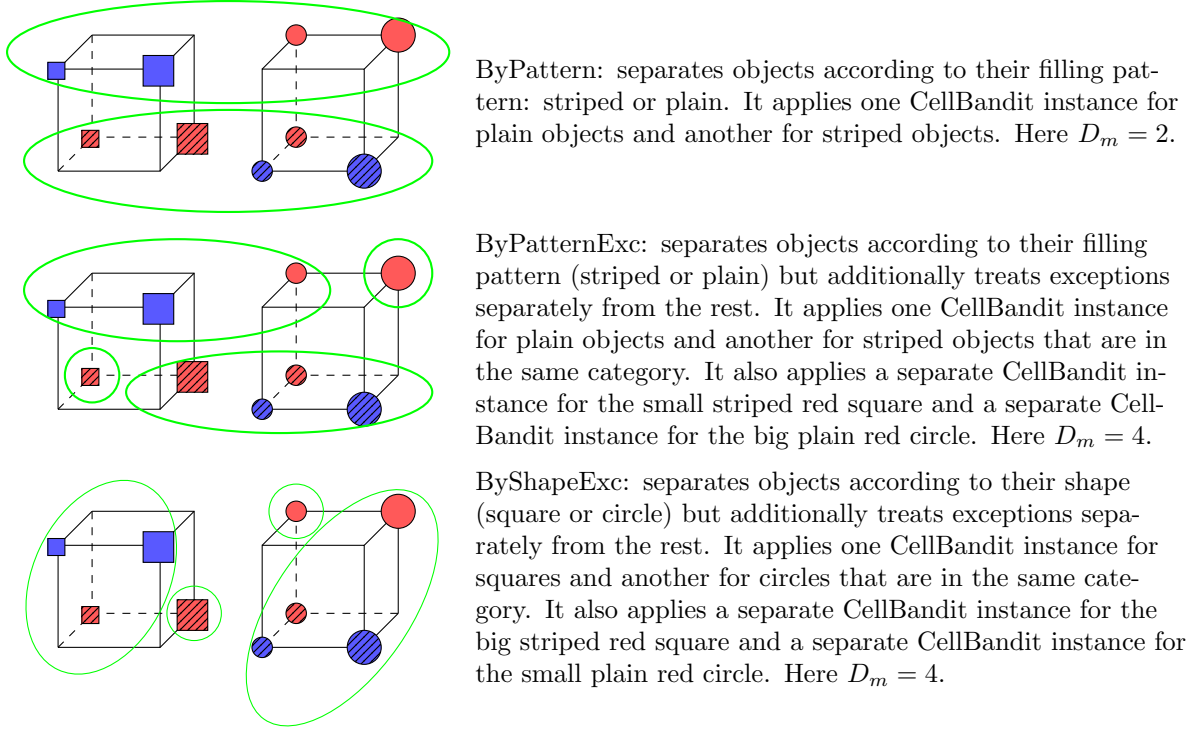


Figure 4: Design of models ByPattern, ByPatternExc and ByShapeExc. Exceptions are determined by the categorization rule in Figure 1b of the article.

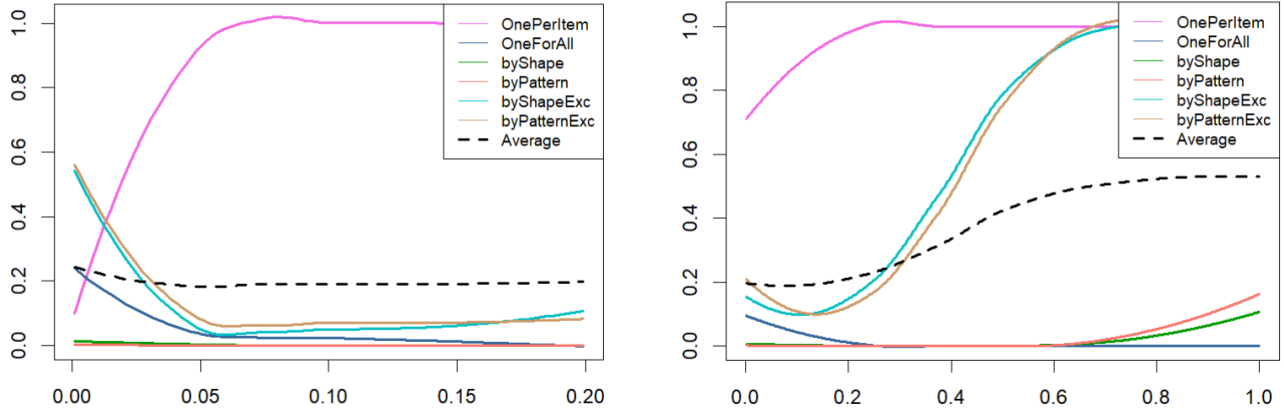


Figure 5: Errors of the penalized log-likelihood criterion as a function of the tuning parameter c , where c is such that $\text{pen}(m) = c \log(n)^2 D_m / n$ is a function of c . Both figures show the percentage of mismatches between \hat{m} and the simulated model over 100 simulations. The same simulations were used for both figures. The evolution of errors as a function of c is logical in relation to the value of D_m .

B.2 Additional Plots

In this section, we provide additional plots (like Figure 3) for two values of the tuning parameter c .

- The value of c minimizing the average of all errors for the penalized log likelihood criterion. As noted in Figure 2c, this corresponds to the choice $c = 0.053$. This choice strongly penalizes the model OnePerItem, which is almost never selected, even when it should be.
- The value of c corresponding to the AIC criterion, that is: $c \log(n)^2 = 1$.

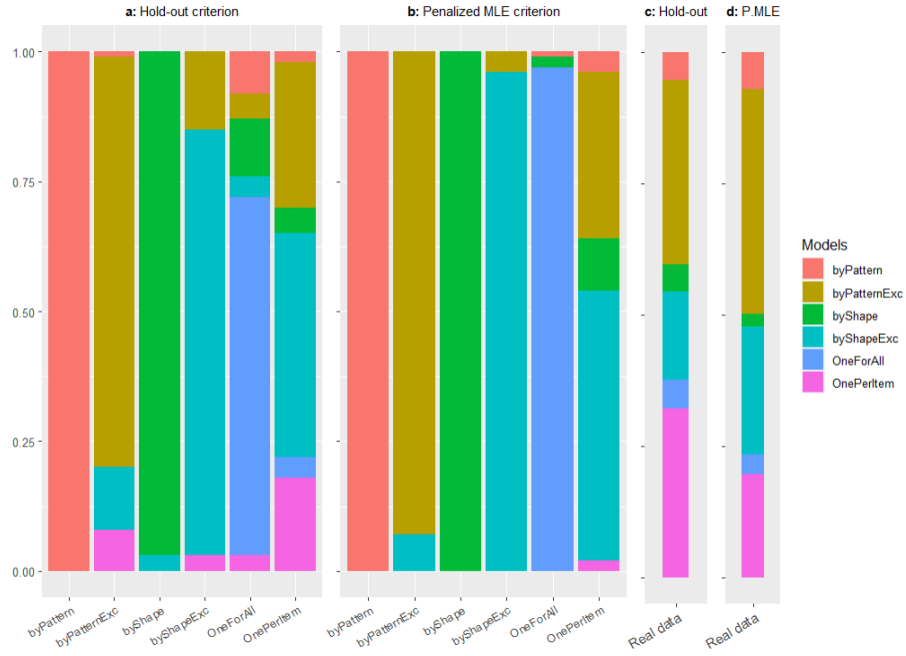


Figure 6: Penalized log-likelihood procedure and Hold-out on synthetic and real data, for the choice $N = 250$ and $c = 0.53$.

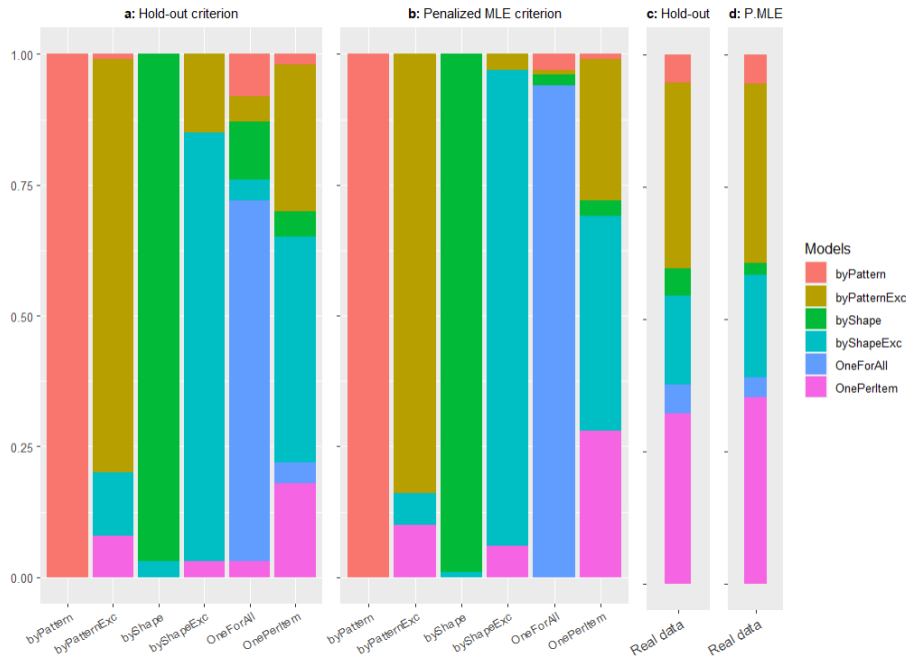


Figure 7: Penalized log-likelihood procedure and Hold-out on synthetic and real data, for the choice $N = 250$ and $c \log(n)^2 = 1$.

B.3 Difference Between The Empirical and Theoretical Threshold

We have not been able to run the simulation with **Exp3-IX**. Indeed, as also shown practically in (Aubert et al., 2023) for the simple **Exp3** case, the probabilities $\pi_{C,T_C}^{\theta_C}$ can go to zero extremely fast. When the individual learns over an horizon $n = 500$, only $\sqrt{n} = 22$ observations would be usable and the estimations even of just the MLE is unreliable. So all the simulations were performed with $n = 500$ and for **Gradient Bandit** as **CellBandit**, for all the 6 models described in Table 1. The way synthetic data are generated can be found in Appendix B.4.

Figure 2a shows that despite the conservative theoretical bound given in Proposition 5 with T_ϵ of order \sqrt{n} , **Gradient Bandit** provides good results when the MLE is applied to all n data points. The truncation at $\sqrt{n} \simeq 20$ required in the theoretical results does not seem necessary in practice, and actually looks suboptimal for **Gradient Bandit**.

B.4 Details on Numerical Illustrations

In this section, we give details on the numerical illustrations of Section 5. The images were obtained using the `ggplot2` package of R. Two types of analyses were conducted, on synthetic data and on real data.

On Synthetic Data. The simulations of the synthetic data helped us calibrate the tuning parameters choices for the hold-out and the penalized log-likelihood procedure. In Section 3, the parameter N must be calibrated for choosing the correct training data sample size. In Section 4, as said in the *Limitations*, since the constant c in the penalty term is not known a priori, it must be calibrated as well. To do this, we follow the guidelines of Wilson and Collins (2019). The procedure is as follows.

- 1) Sample size: $n = 500$. It is of the same order of magnitude as real data.
- 2) Objects generation: periodic sequence of the nine objects repeated through the n trials. We generate a sequence of objects following the same structure as in (Mezzadri et al., 2022b). Due to the periodic pattern, each object is therefore seen roughly the same number of times for all time t .
- 3) Actions generation: for each model in Table 1, we generated 100 sequences of actions called synthetic agents with respect to the procedure given in Algorithm 2 with **Gradient Bandit** as **CellBandit**. The parameters θ_C we used were the same for each model and the same for each cell, equal to $0.03 \times \sqrt{n}$, except for the **OnePerItem** model where we changed slightly the values of the parameter in each cell to make the model identifiable. For $m = \text{OnePerItem}$, we took $\theta^m = ((0.03/10 + k \times 0.007) \times \sqrt{n})_{k \in \{0, \dots, 8\}}$ following the same order of presentation of the sequence of objects defined earlier.
- 4) Parameters estimation: we then fitted each of the six models on all the synthetic agents generated data, and we estimated the associated parameters using (MLE) and the package `DEoptim` in R with range $(0, 1)$ for the parameters θ_C/\sqrt{n} and with the default parameters and a maxiter value equal to 20. We then computed the log-likelihood associated to the estimated parameters. We did this for the likelihood stopped at time $N \in \{25, 50, 100, 150, 200, 250, n\}$.
 - With such data, we were able to plot Figure 2a and Figure 2b with the hold-out criterion defined in Section 3. In Figure 2a, we computed the average error made in each cell by the model fitting of the same model that generated the data. For the Figure 2b, we simply counted the number of times each model verified the hold-out criterion for all the synthetic agents and for each model that generated the data.
 - With the log-likelihood stopped at time n for the estimated parameters, we were able to plot Figure 2c according to the penalized log-likelihood criterion defined in (5). In the same way we counted the number of times each model satisfied the penalized log-likelihood criterion for all the synthetic agents and for each model that generated the data.
- 5) Choice of the parameters N and c for the real data: Given the results of Figure 2b and Figure 2c, we chose to use N to be equal to half of the data length and $c = 0.012$ to account for a reasonable error for model **OnePerItem**, even if in average $c = 0.04$ gives better results. With this data, we were able plot the two first chart of Figure 3.

On Real Data. For the real experimental data, here is the process we followed.

- 1) Sample size: dependent on each individual, the average data sample size is 300.
- 2) Objects and Actions: we collected for each individual their objects sequence and associated choices.
- 3) For each individual, we fitted the 6 models and estimated the parameters associated to each model. To perform hold-out and penalized log-likelihood model selection, we used the parameters N and c chosen thanks to the synthetic data. With this data, we were able to plot Figure 3.

B.5 About the Code and the Data

In this section, we give explanations about the code and data (e.g. computation time, link between code and data). All the data, code and images used are provided in the zip file associated to submission, called `ContextualBanditsCode`. We run all the simulations in R and used the following packages: `DEoptim`, `crayon`, `magrittr`, `dplyr`, `tidyr`, `ggplot2`, `gridExtra`.

For the sample size we chose, all the simulations can run on a PC in a reasonable time of execution (detailed hereafter). Overall, computing the different data and running the code took approximately 6 hours excluding the time needed for the real data. The biggest file is 373 kilobytes. The PC we used was a Gigabyte - AORUS 15G XC, with processor: Intel(R) Core(TM) i7-10870H CPU 2.20GHz, 2208 MHz, 8 cores, 16 logical processors.

On Real Data. As mentioned earlier, we could not provide the experimental data used in (Mezzadri et al., 2022b), since they have already been published in another paper and we do not want to break the ethic agreement. We can only provide the results and estimated data resulting from the experimental data. Note however that the procedures to obtain the following RData files are the same as for the synthetic data which we detail later. The three RData files on the real data are `realdatamle`, `realdata_holdout_trainingset`, `realdata_holdout_testingset`.

- `realdatamle` is a list of estimators and associated log-likelihood for each model and each individual.
- `realdata_holdout_trainingset` is a list of estimators and associated log-likelihood on the first half of the sample for each model and each individual.
- `realdata_holdout_testingset` is a list of log-likelihood on the testing part of the sample for each individual and each model with parameters estimated in `realdata_holdout_trainingset`.

On Synthetic Data. All the synthetic data obtained in the other files can be computed by running the code `ContextualbanditsCodebis`. The code is commented and starts with a list of functions which are necessary to run the different procedures. In the code, we explain how the different procedures lead to the following list of files. We have commented with `#` the parts of the code that would modify the files so that running the code now would give the same images as the ones used in the article. If one wants to generate new data, one should uncomment these lines of code. However, we advise the reader that some of the procedures take a certain time, and would recommend not to do so. We detail hereafter the content of the different csv and RData files and the time it took to run them.

- To begin with, we generate a csv file called `databis_500.csv` of 500 trials and associated list of objects in the file `synthetic_data`.
- In the same `synthetic_data` file we create the different model files and within each of them generate 100 csv files of actions, rewards, and objects according to the procedure described in B.4. This procedure takes around 5 minutes. Then, we begin to compute the MLE for each of the synthetic data csv file.
- `Datalikelihood100agents6modeletabis500horizon` is a nested list of estimators, associated log-likelihood stopped at time n for each model fitted to the data of all the synthetic agents. Computing these data took approximately 2 hours .
- `holdoutbis100agents6models_horizon_20` is the same nested list of estimators but computed on a log-likelihood stopped at time $N = 20$. Computing these data took approximately 10 minutes .

- `holdoutbis100agents6models_horizon_50` is the same nested list of estimators but computed on a log-likelihood stopped at time $N = 50$. Computing these data took approximately 20 minutes .
- `holdoutbis100agents6models_horizon_100` is the same nested list of estimators but computed on a log-likelihood stopped at time $N = 100$. Computing these data took approximately 30 minutes .
- `holdoutbis100agents6models_horizon_150` is the same nested list of estimators but computed on a log-likelihood stopped at time $N = 150$. Computing these data took approximately 40 minutes .
- `holdoutbis100agents6models_horizon_200` is the same nested list of estimators but computed on a log-likelihood stopped at time $N = 200$. Computing these data took approximately 50 minutes .
- `holdoutbis100agents6models_horizon_250` is the same nested list of estimators but computed on a log-likelihood stopped at time $N = 250$. Computing these data took approximately 1 hour .
- `alldataholdoutbis` is a nested list of errors on estimation for the training data and log-likelihood function on the testing data for all synthetic agents, all models and all training data sample size $N \in \{20, 50, 100, 150, 200, 250\}$. Computing these data took approximately 10 minutes .

C METALEARNING

By looking at the experiment in Section 5, it is hard to believe that learners start directly with a model like `ByPatternExc`. It is more likely that they start with a model like `ByPattern` and realize that there are too many exceptions, so that they progressively end up with `ByPatternExp`. One way to model this progressive switch from one strategy to the other is to use bandits with expert advice. In this framework, there is a finite set E of randomized policies called experts, $(\xi_{j,t}(\cdot))_{t \in [n]}$, probabilities over the set of actions $[k]$, that are modeling the different strategies the learner might have. No assumptions are made here on the way experts compute their randomized predictions: they might be the result of contextual bandits like `ByPattern` or more generally any kind of computations that depend on the learner’s past choices. `Exp4` (see Algorithm 5) is an adaptation of `Exp3` to this case (see (Lattimore and Szepesvári, 2020) for regret convergence and variants such as `Exp4.P` (Beygelzimer et al., 2011)).

Algorithm 5 `Exp4` (Bubeck et al., 2012)

Inputs: n (Sample size), $\theta \in [r, R]$ (Parameter), K (Number of actions), E (Set of experts).

Initialization: $q_{E,1}^\theta$ uniform distribution over the experts E .

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

Receive experts advice $a \mapsto \xi_{j,t}(a)$ probability distribution over $[K]$ for all j .

Draw an action $A_t \sim \pi_{E,t}^\theta(\cdot) = \sum_{j \in E} q_{E,t}^\theta(j) \xi_{j,t}(\cdot)$ and receive a reward $g_{A_t,t} \in [0, 1]$.

Update for all $j \in E$,

$$q_{E,t+1}^\theta(j) = \frac{\exp\left(-\frac{\theta}{\sqrt{n}} \sum_{s=1}^t \hat{y}_{j,s}^\theta\right)}{\sum_{i \in E} \exp\left(-\frac{\theta}{\sqrt{n}} \sum_{s=1}^t \hat{y}_{i,s}^\theta\right)} \quad \text{with} \quad \hat{y}_{i,s}^\theta = \sum_{a \in [K]} \xi_{i,t}(a) \frac{g_{a,s}}{\pi_{E,s}^\theta(a)} \mathbf{1}_{A_s=a}$$

In this setting, a model m is defined by a finite set E_m representing the different experts/strategies the learner is learning from. Since there is only one parameter by model (namely $\theta \in [r, R]$), the penalty plays no role, nor the calibration of c . So there is no need for hold-out and one can prove that the model with the smallest log-likelihood on the first $T_\epsilon \sim \sqrt{n}$ time steps satisfies an oracle inequality if \mathcal{M} is finite, as well as $|F| := \max_{m \in \mathcal{M}} |E_m|$. This shows that one can select the set E_m of strategies which is the closest to reality among the sets of strategies that are put in competition.

Limitations. The only limitation with this approach is that we need at first to know the eventual parameters of each strategy. Again we could split the data in a hold-out fashion to make the injection of estimated parameters possible. However, it would be then nearly impossible to correctly estimate the parameters of strategies that are

not used at the beginning of the learning. We refer to (Binz et al., 2023) for other methods in meta-learning for cognition.

Since we work in a more general setting and not simply with contexts, we assume that we observe a process $(A_t)_{1 \leq t \leq n}$ adapted to a general filtration $(\mathcal{F}_t)_{1 \leq t \leq n}$ where for all $t \in [n]$, $A_t \in [K]$. In particular, for all $t \in [n]$, \mathcal{F}_t is generated by the past actions (A_1, \dots, A_t) and any other additional variable which might be observed or not – such as a context at time $t + 1$ for instance. We write, for all $a \in [K]$, and all $t \in [n]$

$$p_t^*(a) = \mathbb{P}(A_t = a | \mathcal{F}_{t-1})$$

the true conditional distribution we wish to estimate.

Additionally, we consider the family of models $\{\pi_{E_m}^{\theta^m} = (\pi_{E_m,t}^{\theta^m})_{t \in [n]}, m \in \mathcal{M}\}$ where \mathcal{M} is a finite set, $\theta^m \in [r, R]$, and for all $m \in \mathcal{M}$, $(\pi_{E_m,t}^{\theta^m})_{t \in [n]}$ is the sequence of mixtures of probability distributions over actions defined recursively in Algorithm 5 for the finite set E_m . Each model m is thus defined by a set of experts $(\xi_{j,t}(\cdot))_{j \in E_m, t \in [n]}$ where for all $m \in \mathcal{M}$, $t \in [n]$, $\xi_{j,t}$ can be any probability distribution over arms $[K]$ as long as it is measurable with respect to \mathcal{F}_{t-1} .

Let $|F| := \max_{m \in \mathcal{M}} |E_m|$. The goal is to select the set E_m of policies – that we see as learning strategies – with which the individual learns to learn. This approach is again based on partial log-likelihood $\ell_n(\pi_{E_m}^{\theta^m})$ of the observations (A_1, \dots, A_n) defined by

$$\ell_n(\pi_{E_m}^{\theta^m}) = \sum_{t=1}^n \log \left(\pi_{E_m,t}^{\theta^m}(A_t) \right). \quad (7)$$

To achieve a model selection result, we need the following assumption on the policies given by the experts.

Assumption 4. There exists $\rho > 0$, such that almost surely, for all $m \in \mathcal{M}$, for all $t \in [n]$ and all $i \in [K]$, $\sum_{j \in E_m} \xi_{j,t}(i) \geq \rho$.

Then, with Assumption 4, we can deduce a result similar to Propositions 4 and 5 because of the very structure of Algorithm 5 which mimics Exp3.

Proposition 6. Assume Assumption 4 holds. Let ρ be the associated constant. Let $\varepsilon \in (0, \rho/|F|)$, and let

$$T_\varepsilon = \left\lfloor \left(\frac{1}{|F|} - \frac{\varepsilon}{\rho} \right) \frac{\sqrt{n}}{R} \right\rfloor \wedge n \quad \text{and} \quad L_\varepsilon = \frac{1}{R\varepsilon^2} \exp \left(\frac{1}{\varepsilon^2} \right).$$

Then, for all $t \in [T_\varepsilon]$, for all $m \in \mathcal{M}$, $\theta^m, \delta^m \in [r, R]$, for all $k \in [K]$,

$$\pi_{E_m,t}^{\theta^m}(k) \geq \varepsilon \quad \text{and} \quad \sup_{k \in [K]} \left| \log \left(\frac{\pi_{E_m,t}^{\theta^m}(k)}{\pi_{E_m,t}^{\delta^m}(k)} \right) \right| \leq L_\varepsilon |\theta^m - \delta^m|.$$

Finally, we still assume that the true distribution is bounded away from 0 (as in (3)).

Assumption 5. Assume that Assumption 4 holds. Let ε and T_ε be the constants of Proposition 6. Assume that

$$\forall t \leq T_\varepsilon, \forall a \in [K], p_t^*(a) \geq \varepsilon.$$

Assumptions 4 and 5 allow us to verify Assumptions 1 and 2 of (Aubert et al., 2024). As for Section 4, it is thus possible to put into competition different sets of experts. Let $A_\varepsilon = L_\varepsilon(R - r) + 2 \log(\varepsilon^{-1})$. Since all the models have the same dimension, there is no penalty term to account for. So the term Σ_ε in Theorem 3 becomes $\log(A_\varepsilon)|\mathcal{M}|e^{-1}$. The result of (Aubert et al., 2024, Corollary 2) states that for any $\diamond > 1$, there exist constants c, c' such that,

$$\begin{aligned} \mathbb{E} \left[\mathbf{K}_{T_\varepsilon}(p^*, \hat{\pi}_{E_m}^{\hat{\theta}^m}) \right] &\leq \mathbb{E} \left[\diamond \inf_{m \in \mathcal{M}} \left(\inf_{\theta \in \Theta^{D_m}} \mathbf{K}_{T_\varepsilon}(p^*, \pi_{E_m}^\theta) \right) \right] + \frac{c}{\kappa} A_\varepsilon^2 \log(\varepsilon^{-1}) \log(T_\varepsilon A_\varepsilon)^2 \frac{1}{T_\varepsilon} \\ &\quad + \frac{18e^{-1}c'}{\kappa} A_\varepsilon \log(A_\varepsilon) |\mathcal{M}| \log(\varepsilon^{-1}) \frac{\log(T_\varepsilon)}{T_\varepsilon}. \end{aligned}$$

D PROOFS

In all of the following, we use the conditional Kullback-Leibler divergence D_{KL} between $p_t^*(\cdot)$ and $p_t^m(\cdot)$:

$$D_{\text{KL}}(p_t^*(A_t), p_t^m(A_t)) = \mathbb{E} \left[\log \frac{p_t^*(A_t)}{p_t^m(A_t)} \middle| \mathcal{F}_{t-1} \right].$$

D.1 Proof of Section 3

Proof of Theorem 1. In the sequel to avoid losing track of important dependence, we denote: for any $t \in [n]$, $m \in \mathcal{M}$, $x_t \mapsto p_t^m(x_t | \mathcal{F}_{t-1})$ the density of A_t conditionally to \mathcal{F}_{t-1} , and likewise for $p_t^*(x_t | \mathcal{F}_{t-1})$.

Consider the following functions, defined for all $t \in [n]$, $x_t^i \in \mathcal{X}^t$, $m \in \mathcal{M}$ by

$$\begin{cases} g_t^m(A_t, \mathcal{F}_{t-1}) = -\frac{1}{2} \log \left(\frac{p_t^m(A_t | \mathcal{F}_{t-1})}{p_t^*(A_t | \mathcal{F}_{t-1})} \right) \\ f_t^m(A_t, \mathcal{F}_{t-1}) = -\log \left(\frac{p_t^*(A_t | \mathcal{F}_{t-1}) + p_t^m(A_t | \mathcal{F}_{t-1})}{2p_t^*(A_t | \mathcal{F}_{t-1})} \right). \end{cases}$$

For any family $h = (h_t)_{t \in [n]}$ of functions A_t that may depend on the past, that is $h_t(A_t, \mathcal{F}_{t-1})$, we write

$$\begin{cases} \tilde{\nu}(h) = \frac{1}{n - N + 1} \sum_{t=N}^n (h_t(A_t, \mathcal{F}_{t-1}) - \mathbb{E}[h_t(A_t, \mathcal{F}_{t-1}) | \mathcal{F}_{t-1}]), \\ \tilde{P}(h) = \frac{1}{n - N + 1} \sum_{t=N}^n h_t(A_t, \mathcal{F}_{t-1}), \\ \tilde{C}(h) = \frac{1}{n - N + 1} \sum_{t=N}^n \mathbb{E}[h_t(A_t, \mathcal{F}_{t-1}) | \mathcal{F}_{t-1}]. \end{cases}$$

Let $m \in \mathcal{M}$. From the definition of \hat{m} ,

$$\tilde{P}(g^{\hat{m}}) \leq \tilde{P}(g^m).$$

Since, $f^{\hat{m}}(A_t, \mathcal{F}_{t-1}) \leq g^{\hat{m}}(A_t, \mathcal{F}_{t-1})$ by concavity of \log , it holds that

$$\tilde{\nu}(f^{\hat{m}}) + \tilde{C}(f^{\hat{m}}) = \tilde{P}(f^{\hat{m}}) \leq \tilde{P}(g^m) = \tilde{\nu}(g^m) + \tilde{C}(g^m).$$

That is

$$\tilde{\nu}(f^{\hat{m}} - f^m) + \tilde{C}(f^{\hat{m}}) \leq \tilde{\nu}(g^m - f^m) + \tilde{C}(g^m).$$

Let $U_m = \tilde{\nu}(g^m - f^m)$, then

$$\tilde{C}(f^{\hat{m}}) \leq \tilde{C}(g^m) - \tilde{\nu}(f^{\hat{m}} - f^m) + U_m. \quad (8)$$

Note that U_m is centered. For $m' \in \mathcal{M}$, let $M_N^{m'} = 0$, and for $t \geq N + 1$, let

$$M_t^{m'} = - \sum_{s=N}^{t-1} \left(f_s^{m'}(A_s, \mathcal{F}_{s-1}) - f_s^m(A_s, \mathcal{F}_{s-1}) - \mathbb{E} \left[f_s^{m'}(A_s, \mathcal{F}_{s-1}) - f_s^m(A_s, \mathcal{F}_{s-1}) \middle| \mathcal{F}_{s-1} \right] \right).$$

Then, $(M_t^{m'})_{t \geq N}$ is an $(\mathcal{F}_t)_{t \geq N}$ -martingale. For $\ell \geq 2$, let $B_N^\ell = 0$, and for $t \geq N + 1$, let

$$B_t^\ell := \sum_{s=N}^{t-1} \mathbb{E} \left[\left(M_{s+1}^{m'} - M_s^{m'} \right)^\ell \middle| \mathcal{H}_s \right].$$

For $t \in \{N, \dots, n-1\}$, note that

$$|M_{t+1}^{m'} - M_t^{m'}| \leq 2 \int \left| f_t^{m'}(x_t, \mathcal{F}_{t-1}) - f_t^m(x_t, \mathcal{F}_{t-1}) \right| \frac{d\delta_{A_t}(x_t) + p_t^*(x_t | \mathcal{F}_{t-1}) d\mu(x_t)}{2},$$

so that, by convexity of $x \mapsto |x|^\ell$,

$$|M_{t+1}^{m'} - M_t^{m'}| \leq 2^\ell \int \left| f_t^{m'}(x_t, \mathcal{F}_{t-1}) - f_t^m(x_t, \mathcal{F}_{t-1}) \right|^\ell \frac{d\delta_{A_t}(x_t) + p_t^*(x_t|\mathcal{F}_{t-1})d\mu(x_t)}{2}.$$

Thus,

$$\begin{aligned} B_t^\ell &\leq 2^\ell \sum_{s=N}^{t-1} \int \mathbb{E} \left[\left| f_s^{m'}(x_s, \mathcal{F}_{s-1}) - f_s^m(x_s, \mathcal{F}_{s-1}) \right|^\ell \frac{d\delta_{X_s}(x_s) + p_s^*(x_s|\mathcal{F}_{s-1})d\mu(x_s)}{2} \middle| \mathcal{F}_{s-1} \right] \\ &= 2^\ell \sum_{s=N}^{t-1} \int \left| f_s^{m'}(x_s, \mathcal{F}_{s-1}) - f_s^m(x_s, \mathcal{F}_{s-1}) \right|^\ell p_s^*(x_s|\mathcal{F}_{s-1})d\mu(x_s) \\ &= 2^\ell \sum_{s=N}^{t-1} \int \left| \log \left(\frac{p_s^*(x_s|\mathcal{F}_{s-1}) + p_s^{m'}(x_s|\mathcal{F}_{s-1})}{p_s^*(x_s|\mathcal{F}_{s-1}) + p_s^m(x_s|\mathcal{F}_{s-1})} \right) \right|^\ell p_s^*(x_s|\mathcal{F}_{s-1})d\mu(x_s). \end{aligned} \quad (9)$$

We now need the following Lemma to continue.

Lemma 7. (Massart, 2007, Lemma 7.26) For all $\ell \geq 2$ and all $x > 0$,

$$\frac{|\log(x)|^\ell}{\ell!} \leq \frac{9}{64} \left(x - \frac{1}{x} \right)^2.$$

Proof. The complete Lemma and proof of the Lemma can be found in (Massart, 2007). \square

Applying Lemma 7 to $x = \sqrt{\frac{p_s^*(x_s|\mathcal{F}_{s-1}) + p_s^{m'}(x_s|\mathcal{F}_{s-1})}{p_s^*(x_s|\mathcal{F}_{s-1}) + p_s^m(x_s|\mathcal{F}_{s-1})}}$ leads to, for all x_s ,

$$\begin{aligned} &\left| \log \left(\frac{p_s^*(x_s|\mathcal{F}_{s-1}) + p_s^{m'}(x_s|\mathcal{F}_{s-1})}{p_s^*(x_s|\mathcal{F}_{s-1}) + p_s^m(x_s|\mathcal{F}_{s-1})} \right) \right|^\ell \\ &\leq \frac{9}{64} 2^\ell \ell! \frac{\left(p_s^m(x_s|\mathcal{F}_{s-1}) - p_s^{m'}(x_s|\mathcal{F}_{s-1}) \right)^2}{(p_s^m(x_s|\mathcal{F}_{s-1}) + p_s^*(x_s|\mathcal{F}_{s-1})) (p_s^*(x_s|\mathcal{F}_{s-1}) + p_s^{m'}(x_s|\mathcal{F}_{s-1}))}. \end{aligned}$$

Plugging this in Equation (9) leads to

$$|B_t^\ell| \leq \frac{9}{4} 2^{2(\ell-2)} \ell! \sum_{s=N}^{t-1} \int \frac{\left(p_s^m(x_s|\mathcal{F}_{s-1}) - p_s^{m'}(x_s|\mathcal{F}_{s-1}) \right)^2 p_s^*(x_s|\mathcal{F}_{s-1})}{(p_s^m(x_s|\mathcal{F}_{s-1}) + p_s^*(x_s|\mathcal{F}_{s-1})) (p_s^*(x_s|\mathcal{F}_{s-1}) + p_s^{m'}(x_s|\mathcal{F}_{s-1}))} d\mu(x_s).$$

For all $x, y, z \geq 0$,

$$(\sqrt{x} + \sqrt{y})^2 z \leq (z + y)(z + x),$$

therefore, with $z = p_s^*(x_s|\mathcal{F}_{s-1})$, $x = p_s^m(x_s|\mathcal{F}_{s-1})$ and $y = p_s^{m'}(x_s|\mathcal{F}_{s-1})$,

$$|B_t^\ell| \leq \frac{9}{4} 4^{\ell-2} \ell! \sum_{s=N}^{t-1} \int \left(\sqrt{p_s^m(x_s|\mathcal{F}_{s-1})} - \sqrt{p_s^{m'}(x_s|\mathcal{F}_{s-1})} \right)^2 d\mu(x_s) \leq \frac{1}{2} 4^{\ell-2} \ell! V_t^{m'}, \quad (10)$$

where

$$\begin{aligned} V_t^{m'} &:= \frac{9}{2} \sum_{s=N}^{t-1} \int \left(\sqrt{p_s^m(x_s|\mathcal{F}_{s-1})} - \sqrt{p_s^{m'}(x_s|\mathcal{F}_{s-1})} \right)^2 d\mu(x_s) \\ &= 9 \sum_{s=N}^{t-1} H \left(p_s^m(A_s|\mathcal{F}_{s-1}), p_s^{m'}(A_s|\mathcal{F}_{s-1}) \right)^2 \end{aligned} \quad (11)$$

where H is the Hellinger distance between the two probability distributions $p_s^m(A_s|\mathcal{F}_{s-1})$ and $p_s^{m'}(A_s|\mathcal{F}_{s-1})$. Lemma 3.3 of (Houdré and Reynaud-Bouret, 2002) gives that for all $\lambda > 0$,

$$\mathcal{E}_t = \exp\left(\lambda M_t^{m'} - \sum_{\ell \geq 2} \frac{\lambda^\ell}{\ell!} B_t^\ell\right)$$

is a supermartingale and that in particular $\mathbb{E}(\mathcal{E}_{n+1}) \leq 1$. By Equation (10), for $\lambda \in (0, 1/4)$,

$$\sum_{\ell \geq 2} \frac{\lambda^\ell}{\ell!} B_t^\ell \leq \frac{\lambda^2}{2} \sum_{\ell \geq 2} (4\lambda)^{\ell-2} V_t^{m'} = \frac{\lambda^2}{2(1-4\lambda)} V_t^{m'}.$$

Let $\Psi(\lambda) = \frac{\lambda^2}{2(1-4\lambda)}$ for $\lambda \in (0, 1/4)$. Then,

$$\mathbb{E}\left[e^{\lambda M_{n+1}^{m'} - \Psi(\lambda) V_{n+1}^{m'}} \middle| \mathcal{F}_{N-1}\right] \leq 1.$$

By Markov's inequality, for all $x \geq 0$ and $\lambda \in (0, 1/4)$,

$$\mathbb{P}\left(M_{n+1}^{m'} \geq V_{n+1}^{m'} \frac{\Psi(\lambda)}{\lambda} + \frac{x}{\lambda} \middle| \mathcal{F}_{N-1}\right) \leq e^{-x}. \quad (12)$$

Therefore, for all $x, u \geq 0$ and $\lambda \in (0, 1/4)$,

$$\mathbb{P}\left(M_{n+1}^{m'} \geq \frac{\Psi(\lambda)}{\lambda} u + \frac{x}{\lambda} \quad \text{and} \quad V_{n+1}^{m'} \leq u \middle| \mathcal{F}_{N-1}\right) \leq e^{-x}.$$

To choose the optimal λ , we use Lemma 2 from (Hansen et al., 2015).

Lemma 8. (Hansen et al., 2015, Lemma 2) *Let a, b and x be positive constants and let us consider on $(0, 1/b)$,*

$$g(\xi) = \frac{a\xi}{1-b\xi} + \frac{x}{\xi}.$$

Then $\min_{\xi \in (0, 1/b)} g(\xi) = 2\sqrt{ax} + bx$ and the minimum is achieved in $\xi(a, b, x) = \frac{\sqrt{x}}{\sqrt{xb} + \sqrt{a}}$.

For $a = \frac{u}{2}$ and $b = 4$, Lemma 8 shows that for all $x, u \geq 0$,

$$\mathbb{P}\left(M_{n+1}^{m'} \geq \sqrt{2ux} + 4x \quad \text{and} \quad V_{n+1}^{m'} \leq u \middle| \mathcal{F}_{N-1}\right) \leq e^{-x}.$$

Let us use a peeling argument similar to (Hansen et al., 2015):

Lemma 9. *Let X, V be real-valued random variables and α, b, v, w be positive numbers such that $V \in [w, v]$ a.s. and such that for all $x \geq 0$ and $u \in [w, v]$,*

$$\mathbb{P}(X \geq \sqrt{ux} + bx \quad \text{and} \quad (1+\alpha)^{-1}u \leq V \leq u) \leq e^{-x},$$

then for any $x \geq 0$,

$$\mathbb{P}(X \geq \sqrt{(1+\alpha)Vx} + bx) \leq \left(1 + \frac{\log(v/w)}{\log(1+\alpha)}\right) e^{-x}.$$

Proof. Let $v_0 = w$, $v_{d+1} = (1+\alpha)v_d$, and D the smallest integer such that $v_D \geq v$. For all $d \in [D]$ and $x \geq 0$,

$$\mathbb{P}(X \geq \sqrt{v_d x} + bx \quad \text{and} \quad v_{d-1} \leq V \leq v_d) \leq e^{-x}.$$

In particular, since $V \geq v_{d-1} = (1 + \alpha)^{-1} v_d$ on this event,

$$\mathbb{P} \left(X \geq \sqrt{(1 + \alpha)Vx + bx} \quad \text{and} \quad v_{d-1} \leq V \leq v_d \right) \leq e^{-x}.$$

Taking the union bound,

$$\mathbb{P} \left(X \geq \sqrt{(1 + \alpha)Vx + bx} \right) \leq De^{-x},$$

and by definition $D \leq \frac{\log(v/w)}{\log(1 + \alpha)} + 1$. \square

We may apply Lemma 9 to $X = M_{n+1}^{m'}$ and $b = 4$. Since $V_{n+1}^{m'}$ does not have an obvious lower bound, we consider $V = 2(V_{n+1}^{m'} + \beta)$ for some $\beta > 0$ to be chosen later. We may therefore take $w = 2\beta$. For the upper bound v on V , by (11), since the Hellinger distance is upper bounded by 1, we may take $v = 2(\beta + 9(n - N + 1))$. With these choices, for any $\beta, \alpha, x > 0$,

$$\mathbb{P} \left(M_{n+1}^{m'} \geq \sqrt{2(1 + \alpha)(V_{n+1}^{m'} + \beta)x + 4x} \mid \mathcal{F}_{N-1} \right) \leq \left(\frac{\log \left(\frac{9(n-N+1)}{\beta} + 1 \right)}{\log(1 + \alpha)} + 1 \right) e^{-x}.$$

For $\alpha = \sqrt{2}$,

$$\mathbb{P} \left(M_{n+1}^{m'} \geq \sqrt{5(V_{n+1}^{m'} + \beta)x + 4x} \mid \mathcal{F}_{N-1} \right) \leq \left(2 \log \left(\frac{9(n-N+1)}{\beta} + 1 \right) + 1 \right) e^{-x}. \quad (13)$$

By definition of $V_{n+1}^{m'}$ and the triangle inequality,

$$\begin{aligned} V_{n+1}^{m'} &= 9 \sum_{s=N}^n H \left(p_s^m(A_s | \mathcal{F}_{s-1}), p_s^{m'}(A_s | \mathcal{F}_{s-1}) \right)^2 \\ &\leq 18 \sum_{s=N}^n \left(H \left(p_s^*(A_s | \mathcal{F}_{s-1}), p_s^m(A_s | \mathcal{F}_{s-1}) \right)^2 + H \left(p_s^*(A_s | \mathcal{F}_{s-1}), p_s^{m'}(A_s | \mathcal{F}_{s-1}) \right)^2 \right). \end{aligned} \quad (14)$$

We now use (Massart, 2007, Lemma 7.23) giving a connection between the Kullback-Leibler divergence D_{KL} and the Hellinger distance H .

Lemma 10. (Massart, 2007, Lemma 7.23) *Let P and Q be some probability measures. Then,*

$$D_{\text{KL}} \left(P, \frac{P + Q}{2} \right) \geq (2 \log 2 - 1) H^2(P, Q).$$

Moreover, whenever $P \ll Q$,

$$2H^2(P, Q) \leq D_{\text{KL}}(P, Q).$$

Since $\frac{18}{2 \log(2) - 1} \leq 48$,

$$\begin{aligned} V_{n+1}^{m'} &\leq 48 \sum_{s=N}^n \left(D_{\text{KL}} \left(p_s^*(A_s | \mathcal{F}_{s-1}), \frac{p_s^*(A_s | \mathcal{F}_{s-1}) + p_s^{m'}(A_s | \mathcal{F}_{s-1})}{2} \right) + \frac{1}{2} D_{\text{KL}}(p_s^*(A_s | \mathcal{F}_{s-1}), p_s^m(A_s | \mathcal{F}_{s-1})) \right) \\ &=: 9W_n^{m'}. \end{aligned} \quad (15)$$

Let $\beta = 9(n - N + 1)y^2$, where $y > 0$ is to be chosen later. Replacing x by $x + \log(|\mathcal{M}|)$ leads to

$$\begin{aligned} \mathbb{P} \left(\frac{M_{n+1}^{m'}}{n - N + 1} \geq 3 \sqrt{5 \left(\frac{W_n^{m'}}{n - N + 1} + y^2 \right) \frac{x + \log(|\mathcal{M}|)}{n - N + 1}} + 4 \frac{x + \log(|\mathcal{M}|)}{n - N + 1} \mid \mathcal{F}_{N-1} \right) \\ \leq (2 \log(y^{-2} + 1) + 1) e^{-(x + \log(|\mathcal{M}|))}. \end{aligned}$$

Let $\kappa_1 \in (0, 1/(8\sqrt{5})]$, then, using $2\sqrt{ab} \leq \kappa_1 a + \kappa_1^{-1} b$ and taking $y^2 = \frac{x + \log(|\mathcal{M}|)}{(n-N+1)\log 2} \geq \frac{1}{n-N+1}$ since $x \geq 0$ and $|\mathcal{M}| \geq 2$,

$$\begin{aligned} \mathbb{P} \left(\frac{M_{n+1}^{m'}}{n-N+1} \geq \frac{3\sqrt{5}}{2} \kappa_1 \frac{W_n^{m'}}{n-N+1} + \left(4 + \frac{3\sqrt{5}}{2} \left(\frac{\kappa_1}{\log 2} + \kappa_1^{-1} \right) \right) \frac{x + \log(|\mathcal{M}|)}{n-N+1} \middle| \mathcal{F}_{N-1} \right) \\ \leq (2 \log(n-N+2) + 1) e^{-(x + \log(|\mathcal{M}|))}. \end{aligned} \quad (16)$$

By the union bound on all $m' \in \mathcal{M}$, the previous inequality holds with probability at least $1 - (2 \log(n-N+2) + 1)e^{-x}$ for all $m' \in \mathcal{M}$. It holds in particular for \hat{m} . Recall with (8) that,

$$\begin{aligned} \frac{1}{n-N+1} \sum_{s=N}^n \text{D}_{\text{KL}} \left(p_s^*(A_s | \mathcal{F}_{s-1}), \frac{p_s^*(A_s | \mathcal{F}_{s-1}) + p_s^{\hat{m}}(A_s | \mathcal{F}_{s-1})}{2} \right) - U_m \\ \leq \frac{1}{2(n-N+1)} \sum_{s=N}^n \text{D}_{\text{KL}} (p_s^*(A_s | \mathcal{F}_{s-1}), p_s^m(A_s | \mathcal{F}_{s-1})) + \frac{M_{n+1}^{\hat{m}}}{n-N+1}. \end{aligned} \quad (17)$$

Plugging (15) and (16) in (17) leads to, conditionally on \mathcal{F}_{N-1} , with probability at least $1 - (2 \log(n-N+2) + 1)e^{-x}$

$$\begin{aligned} \frac{(1 - C_{\kappa_1})}{n-N+1} \sum_{s=N}^n \text{D}_{\text{KL}} \left(p_s^*(A_s | \mathcal{F}_{s-1}), \frac{p_s^*(A_s | \mathcal{F}_{s-1}) + p_s^{\hat{m}}(A_s | \mathcal{F}_{s-1})}{2} \right) - U_m \\ \leq \frac{(1 + C_{\kappa_1})}{n-N+1} \sum_{s=N}^n \frac{1}{2} \text{D}_{\text{KL}} (p_s^*(A_s | \mathcal{F}_{s-1}), p_s^m(A_s | \mathcal{F}_{s-1})) + C'_{\kappa_1} \frac{x + \log(|\mathcal{M}|)}{n-N+1}, \end{aligned}$$

where

- $C_{\kappa_1} = 8\sqrt{5}\kappa_1$,
- $C'_{\kappa_1} = 4 + \frac{3\sqrt{5}}{2} \left(\frac{\kappa_1}{\log 2} + \kappa_1^{-1} \right) \leq 13\kappa_1^{-1} = \frac{104\sqrt{5}}{C_{\kappa_1}}$,

Integrating on $x \geq 0$ and noting that $\mathbb{E}[U_m | \mathcal{F}_{N-1}] = 0$ leads to, for all $m \in \mathcal{M}$,

$$\begin{aligned} \frac{(1 - C_{\kappa_1})}{n-N+1} \mathbb{E} \left[\sum_{s=N}^n \text{D}_{\text{KL}} \left(p_s^*(A_s | \mathcal{F}_{s-1}), \frac{p_s^*(A_s | \mathcal{F}_{s-1}) + p_s^{\hat{m}}(A_s | \mathcal{F}_{s-1})}{2} \right) \middle| \mathcal{F}_{N-1} \right] \\ \leq \frac{(1 + C_{\kappa_1})}{n-N+1} \mathbb{E} \left[\frac{1}{2} \sum_{s=N}^n \text{D}_{\text{KL}} (p_s^*(A_s | \mathcal{F}_{s-1}), p_s^m(A_s | \mathcal{F}_{s-1})) \middle| \mathcal{F}_{N-1} \right] \\ + \frac{104\sqrt{5}}{C_{\kappa_1}} \frac{2 \log(n-N+2) + 1 + \log(|\mathcal{M}|)}{n-N+1}. \end{aligned}$$

Pick any κ_1 so that $C_{\kappa_1} \in (0, 1]$. And finally, with (Massart, 2007, Lemma 7.23),

$$(2 \log 2 - 1) \text{H}^2(p_s^*(A_s | \mathcal{F}_{s-1}), p_s^{\hat{m}}(A_s | \mathcal{F}_{s-1})) \leq \text{D}_{\text{KL}} \left(p_s^*(A_s | \mathcal{F}_{s-1}), \frac{p_s^*(A_s | \mathcal{F}_{s-1}) + p_s^{\hat{m}}(A_s | \mathcal{F}_{s-1})}{2} \right).$$

□

D.2 Proof of Section 4

Proof of Proposition 2. The proof is straightforward with the definition of $p_{\theta^m, t}^m$ in (2). Let $\theta^m, \delta^m \in \Theta^m$, $t \leq T_\varepsilon$, $x \in \mathcal{X}$, and $k \in [K]$.

$$p_{\theta^m, t}^m(k|x) = \sum_{C \in \mathcal{P}_m} \pi_{C, T_t^C}^{\theta^m}(k) \mathbf{1}_{x \in C} \geq \sum_{C \in \mathcal{P}_m} \varepsilon \mathbf{1}_{x \in C} = \varepsilon.$$

For the second part of the proof, it holds that, almost surely, for all $t \leq T_\varepsilon$

$$\left| \log \left(\frac{p_{\delta^m, t}^m(k|x)}{p_{\theta^m, t}^m(k|x)} \right) \right| = \sum_{C \in \mathcal{P}_m} \left| \log \left(\frac{\pi_{C, T_t^C}^{\delta_C}(k)}{\pi_{C, T_t^C}^{\theta_C}(k)} \right) \right| \mathbf{1}_{x \in C} \leq L_\varepsilon \sum_{C \in \mathcal{P}_m} \|\delta_C - \theta_C\|_2 \mathbf{1}_{x \in C} \leq L_\varepsilon \sup_{C \in \mathcal{P}_m} \|\delta_C - \theta_C\|_2.$$

□

Proof of Theorem 3. Our goal is to apply (Aubert et al., 2024).

Assumption 1 of (Aubert et al., 2024) is satisfied for $n = T_\varepsilon$ since with Proposition 2, there exists $\varepsilon > 0$ such that a.s., for all $t \in [T_\varepsilon]$, for all $k \in [K]$, $p_t^*(k|X_t) \in [\varepsilon, 1]$ and for all $m \in \mathcal{M}$ and all $\theta^m \in \Theta^{D_m}$, $p_{\theta^m, t}^m(k|X_t) \in [\varepsilon, 1]$.

Assumption 2 of (Aubert et al., 2024) is satisfied since with Proposition 2, there exists a positive constants L_ε such that a.s., for all $t \in [T_\varepsilon]$, for all $m \in \mathcal{M}$ and all $\delta^m, \theta^m \in \Theta^{D_m}$,

$$\sup_{k \in [K]} \left| \log \left(\frac{p_{\delta^m, t}^m(k|X_t)}{p_{\theta^m, t}^m(k|X_t)} \right) \right| \leq L_\varepsilon \sup_{C \in \mathcal{P}_m} \|\delta_C - \theta_C\|_2$$

and by Assumption, for all $\theta^m, \delta^m \in \Theta^{D_m}$

$$\sup_{C \in \mathcal{P}_m} \|\delta_C - \theta_C\|_2 \leq \sqrt{d}(R - r).$$

Note in particular that the Lipschitz constant in Proposition 2 does not depend on m .

Assumptions 3 and 4 in (Aubert et al., 2024) are always satisfied because the set of actions $[K]$ is finite.

Setting $A_\varepsilon = L_\varepsilon \sqrt{d}(R - r) + 2 \log(\varepsilon^{-1})$, Corollary 2 in (Aubert et al., 2024) simply reads as follows. There exist positive numerical constants C and C' such that the following holds. Assume that

$$\Sigma_\varepsilon = \log(A_\varepsilon) \sum_{m \in \mathcal{M}} e^{-D_m} < +\infty.$$

Let $\kappa \in (0, 1]$. If for all $m \in \mathcal{M}$,

$$\text{pen}(m) \geq \frac{C}{\kappa} A_\varepsilon^2 \log(\varepsilon^{-1}) \log(T_\varepsilon A_\varepsilon)^2 \frac{D_m}{T_\varepsilon},$$

then,

$$\begin{aligned} & \frac{1 - \kappa}{T_\varepsilon} \sum_{t=1}^{T_\varepsilon} \mathbb{E} \left[\text{D}_{\text{KL}} \left(p_t^*(A_t|X_t), p_{\theta^m, t}^m(A_t|X_t) \right) \right] \\ & \leq \mathbb{E} \left[\inf_{m \in \mathcal{M}} \left((1 + \kappa) \inf_{\theta \in \Theta^m} \frac{1}{T_\varepsilon} \sum_{t=1}^{T_\varepsilon} \text{D}_{\text{KL}} \left(p_t^*(A_t|X_t), p_{\theta, t}^m(A_t|X_t) \right) + 2 \text{pen}(m) \right) \right] \\ & \quad + \frac{36C'}{\kappa} A_\varepsilon \Sigma_\varepsilon \log(\varepsilon^{-1}) \frac{\log(T_\varepsilon)}{T_\varepsilon}. \end{aligned}$$

To conclude take for instance $\kappa = 1/2$, hence the result. □

D.3 Proof of Section A.1

Proof of Proposition 4. Let $\theta_C \in \Theta_C$. Write $\theta_{C, n} = (\eta_{C, n}, \gamma_{C, n}) = \theta_C / \sqrt{n} \in \Theta$. To ease the notations in the proof, we remove the C from the notations. θ_C becomes θ and $\theta_{C, n}$ becomes θ_n . In the same way, $\theta_n = (\eta_n, \gamma_n)$ now.

Let $t \in F_n(C)$. Then,

$$\begin{aligned} \pi_{C, T_t^C+1}^\theta(k) &= \frac{\pi_{C, T_t^C}^\theta(k) e^{-\eta_n g_{k,t}/(\gamma_n + \pi_{C, T_t^C}^\theta(k))} \mathbf{1}_{A_t=k}}{(1 - \pi_{C, T_t^C}^\theta(k)) + \pi_{C, T_t^C}^\theta(k) e^{-\eta_n g_{k,t}/(\gamma_n + \pi_{C, T_t^C}^\theta(k))}} \\ &\quad + \sum_{\substack{j=1 \\ j \neq k}}^K \frac{\pi_{C, T_t^C}^\theta(k) \mathbf{1}_{A_t=j}}{(1 - \pi_{C, T_t^C}^\theta(j)) + \pi_{C, T_t^C}^\theta(j) e^{-\eta_n g_{j,t}/(\gamma_n + \pi_{C, T_t^C}^\theta(j))}}. \end{aligned}$$

For any $q \in [0, 1]$, since $g_{k,t} \in [0, 1]$, $1 - q + qe^{-\eta_n g_{k,t}/(q+\gamma)} \leq 1$. Therefore,

$$\pi_{C, T_t^C+1}^\theta(k) \geq \pi_{C, T_t^C}^\theta(k) e^{-\eta_n g_{k,t}/(\gamma_n + \pi_{C, T_t^C}^\theta(k))} \mathbf{1}_{A_t=k} + \sum_{\substack{j=1 \\ j \neq k}}^K \pi_{C, T_t^C}^\theta(k) \mathbf{1}_{A_t=j}.$$

Since $e^{-\eta_n g_{k,t}/(\gamma_n + \pi_{C, T_t^C}^\theta(k))} \leq 1$,

$$\begin{aligned} &\pi_{C, T_t^C+1}^\theta(k) \\ &\geq \pi_{C, T_t^C}^\theta(k) e^{-\eta_n g_{k,t}/(\gamma_n + \pi_{C, T_t^C}^\theta(k))} \mathbf{1}_{A_t=k} + \sum_{\substack{j=1 \\ j \neq k}}^K \pi_{C, T_t^C}^\theta(k) e^{-\eta_n g_{k,t}/(\gamma_n + \pi_{C, T_t^C}^\theta(k))} \mathbf{1}_{A_t=j} \\ &= \pi_{C, T_t^C}^\theta(k) e^{-\eta_n g_{k,t}/(\gamma_n + \pi_{C, T_t^C}^\theta(k))} \geq \pi_{C, T_t^C}^\theta(k) e^{-\eta_n/(\gamma_n + \pi_{C, T_t^C}^\theta(k))} \end{aligned}$$

since $g_{k,t} \in [0, 1]$. Then,

$$\pi_{C, T_t^C+1}^\theta(k) \geq \pi_{C, T_t^C}^\theta(k) \left(1 - \frac{\eta_n g_{k,t}}{\gamma_n + \pi_{C, T_t^C}^\theta(k)} \right) \geq \pi_{C, T_t^C}^\theta(k) - \eta_n.$$

Summing for all $s \in F_t(C)$, since $\pi_{k,1}^\theta = \frac{1}{K}$,

$$\pi_{C, T_t^C}^\theta(k) \geq \frac{1}{K} - \eta_n T_t^C.$$

Note that $T_t^C \leq t \leq T_\varepsilon$. Since, $T_\varepsilon \leq \left\lceil \left(\frac{1}{K} - \varepsilon \right) \frac{\sqrt{n}}{R} \right\rceil$, for all $t \leq T_\varepsilon$ and $1 \leq k \leq K$,

$$\varepsilon \leq \frac{1}{K} - \frac{R}{\sqrt{n}} T_\varepsilon \leq \frac{1}{K} - \eta_n T_\varepsilon \leq \frac{1}{K} - \eta_n t \leq \pi_{C, T_t^C}^\theta(k).$$

For the second part of the proof, let $\theta = (\eta, \gamma), \theta' = (\eta', \gamma') \in \Theta_C$. For $t \geq 2$ Let $h_{j,t}^\theta = \eta_n \sum_{s \in F_t(C)} \hat{g}_{j,s}^\theta$. Then $\pi_{C, T_t^C}^\theta = \text{softmax}(h_{\cdot,t}^\theta)$. The function softmax is 1-Lipschitz with respect to the $\|\cdot\|_2$ -norm in \mathbb{R}^K (see (Gao and Pavel, 2017) for a proof). Therefore,

$$\|\pi_{C, T_t^C}^\theta - \pi_{C, T_t^C}^{\theta'}\|_2 \leq \|h_{\cdot,t}^\theta - h_{\cdot,t}^{\theta'}\|_2.$$

Since $g_{j,s} \in [0, 1]$, by the triangle inequality

$$\|\pi_{C, T_t^C}^\theta - \pi_{C, T_t^C}^{\theta'}\|_2 \leq \sum_{s \in F_t(C)} \left\| \left(\frac{\eta_n}{\gamma_n + \pi_{C, T_s^C}^\theta(\cdot)} - \frac{\eta'_n}{\gamma'_n + \pi_{C, T_s^C}^{\theta'}(\cdot)} \right) \mathbf{1}_{A_s=\cdot} \right\|_2.$$

Again, using the triangle inequality,

$$\begin{aligned} \|\pi_{C, T_t^C}^\theta - \pi_{C, T_t^C}^{\theta'}\|_2 &\leq \sum_{s \in F_t(C)} \left\| \left(\frac{\eta_n}{\gamma_n + \pi_{C, T_s^C}^\theta} - \frac{\eta'_n}{\gamma'_n + \pi_{C, T_s^C}^\theta} \right) \mathbf{1}_{A_s=\cdot} \right\|_2 \\ &\quad + \sum_{s \in F_t(C)} \left\| \left(\frac{\eta'_n}{\gamma'_n + \pi_{C, T_s^C}^\theta} - \frac{\eta'_n}{\gamma'_n + \pi_{C, T_s^C}^{\theta'}} \right) \mathbf{1}_{A_s=\cdot} \right\|_2. \end{aligned} \tag{18}$$

For $1 \geq q \geq \varepsilon$, let $f : (x_1, x_2) \in [0, R_n] \times [0, R_n] \mapsto \frac{x_1}{x_2 + q}$ where $R_n = \frac{R}{\sqrt{n}}$. The function f is continuously differentiable, and

$$\nabla f = \frac{1}{(x_2 + q)^2} \begin{pmatrix} x_2 + q \\ -x_1 \end{pmatrix}.$$

The ℓ^2 -norm of the gradient can be upper bounded by

$$\|\nabla f\|_2 \leq \frac{1}{\varepsilon^2} \sqrt{R_n^2 + \varepsilon^2} =: c_\varepsilon$$

By the mean value theorem, for all $k \in [K]$

$$\left| \frac{\eta_n}{\gamma_n + \pi_{C, T_s^C}^\theta(k)} - \frac{\eta'_n}{\gamma'_n + \pi_{C, T_s^C}^\theta(k)} \right| \leq c_\varepsilon \|\theta_n - \theta'_n\|_2.$$

As a result,

$$\begin{aligned} \left\| \left(\frac{\eta_n}{\gamma_n + \pi_{C, T_s^C}^\theta} - \frac{\eta'_n}{\gamma'_n + \pi_{C, T_s^C}^\theta} \right) \mathbf{1}_{A_s=} \right\|_2^2 &= \sum_{k=1}^K \left(\frac{\eta_n}{\gamma_n + \pi_{C, T_s^C}^\theta(k)} - \frac{\eta'_n}{\gamma'_n + \pi_{C, T_s^C}^\theta(k)} \right)^2 \mathbf{1}_{A_s=k} \\ &\leq c_\varepsilon^2 \|\theta_n - \theta'_n\|_2^2 \sum_{k=1}^K \mathbf{1}_{A_s=k} \\ &= c_\varepsilon^2 \|\theta_n - \theta'_n\|_2^2 \end{aligned}$$

Therefore,

$$\begin{aligned} \sum_{s \in F_t(C)} \left\| \left(\frac{\eta_n}{\gamma_n + \pi_{C, T_s^C}^\theta} - \frac{\eta'_n}{\gamma'_n + \pi_{C, T_s^C}^\theta} \right) \mathbf{1}_{A_s=} \right\|_2 \\ \leq \sum_{s \in F_t(C)} c_\varepsilon \|\theta_n - \theta'_n\|_2 = T_t^C c_\varepsilon \|\theta_n - \theta'_n\|_2 \leq T_\varepsilon c_\varepsilon \|\theta_n - \theta'_n\|_2. \end{aligned} \quad (19)$$

For $(\eta, \gamma) \in [0, R_n] \times [0, R_n]$, let $g : q \in [\varepsilon, 1] \mapsto \frac{\eta}{\gamma + q}$. The function g is continuously differentiable, and

$$0 \leq f'(q) = \frac{\eta}{(\gamma + q)^2} \leq \frac{R_n}{\varepsilon^2}.$$

By the mean value theorem, for all $k \in [K]$,

$$\left| \frac{\eta'_n}{\gamma'_n + \pi_{C, T_s^C}^\theta(k)} - \frac{\eta'_n}{\gamma'_n + \pi_{C, T_s^C}^{\theta'}(k)} \right| \leq \frac{R_n}{\varepsilon^2} \left| \pi_{C, T_s^C}^\theta(k) - \pi_{C, T_s^C}^{\theta'}(k) \right|.$$

Therefore,

$$\begin{aligned} \left\| \left(\frac{\eta'_n}{\gamma'_n + \pi_{C, T_s^C}^\theta} - \frac{\eta'_n}{\gamma'_n + \pi_{C, T_s^C}^{\theta'}} \right) \mathbf{1}_{A_s=} \right\|_2^2 &= \sum_{k=1}^K \left(\frac{\eta'_n}{\gamma'_n + \pi_{C, T_s^C}^\theta(k)} - \frac{\eta'_n}{\gamma'_n + \pi_{C, T_s^C}^{\theta'}(k)} \right)^2 \mathbf{1}_{A_s=k} \\ &\leq \frac{R_n^2}{\varepsilon^4} \sum_{k=1}^K \left| \pi_{C, T_s^C}^\theta(k) - \pi_{C, T_s^C}^{\theta'}(k) \right|^2 \mathbf{1}_{A_s=k} \\ &\leq \frac{R_n^2}{\varepsilon^4} \left(\left\| \pi_{C, T_s^C}^\theta - \pi_{C, T_s^C}^{\theta'} \right\|_2 \right)^2. \end{aligned}$$

Thus,

$$\sum_{s \in F_t(C)} \left\| \left(\frac{\eta'_n}{\gamma'_n + \pi_{C, T_s^C}^\theta} - \frac{\eta'_n}{\gamma'_n + \pi_{C, T_s^C}^{\theta'}} \right) \mathbf{1}_{A_s=} \right\|_2 \leq \frac{R_n}{\varepsilon^2} \sum_{s \in F_t(C)} \left\| \pi_{C, T_s^C}^\theta - \pi_{C, T_s^C}^{\theta'} \right\|_2 \quad (20)$$

Plugging Equations (19) and (20) in Equation (18)

$$\|\pi_{C,T_t^C}^\theta - \pi_{C,T_t^C}^{\theta'}\|_2 \leq T_\varepsilon c_\varepsilon \|\theta_n - \theta'_n\|_2 + \frac{R_n}{\varepsilon^2} \sum_{s \in F_t(C)} \|\pi_{C,T_s^C}^\theta - \pi_{C,T_s^C}^{\theta'}\|_2,$$

Using the discrete Gronwall Lemma (Clark, 1987) leads to, for all $t \leq T_\varepsilon$,

$$\begin{aligned} \|\pi_{C,T_t^C}^\theta - \pi_{C,T_t^C}^{\theta'}\|_2 &\leq T_\varepsilon c_\varepsilon \|\theta_n - \theta'_n\|_2 \prod_{s \in F_t(C)} \left(1 + \frac{R_n}{\varepsilon^2}\right) \\ &\leq T_\varepsilon c_\varepsilon \|\theta_n - \theta'_n\|_2 \exp\left(\frac{R_n T_\varepsilon}{\varepsilon^2}\right). \end{aligned}$$

But, since $\frac{1}{K} - \varepsilon \leq 1$,

$$T_\varepsilon \leq \left(\frac{1}{K} - \varepsilon\right) \frac{\sqrt{n}}{R} \leq \frac{\sqrt{n}}{R}.$$

Therefore, $R_n T_\varepsilon \leq 1$ and for $t \leq T_\varepsilon$,

$$\|\pi_{C,T_t^C}^\theta - \pi_{C,T_t^C}^{\theta'}\|_2 \leq \frac{c_\varepsilon}{R} e^{1/\varepsilon^2} \|\theta - \theta'\|_2.$$

To conclude note that log is $1/\varepsilon$ -Lipschitz on $[\varepsilon, 1]$ and that

$$\sup_{k \in [K]} \left| \log \left(\frac{\pi_{C,T_t^C}^{\delta_C}(k)}{\pi_{C,T_t^C}^\theta(k)} \right) \right| \leq \frac{1}{\varepsilon} \|\pi_{C,T_t^C}^\theta - \pi_{C,T_t^C}^{\theta'}\|_2.$$

□

D.4 Proof of Section A.2

Proof of Proposition 5. We take the same notations as the previous section. The updated probability can be written

$$\pi_{C,T_t^C+1}^\theta(k) = \frac{\pi_{C,T_t^C}^\theta(k) e^{\theta_n (\mathbf{1}_{A_t=k} - \pi_{C,T_t^C}^\theta(k)) g_{A_t,t}}}{\pi_{C,T_t^C}^\theta(A_t) e^{\theta_n (1 - \pi_{C,T_t^C}^\theta(A_t)) g_{A_t,t}} + \sum_{j \neq A_t} \pi_{C,T_t^C}^\theta(j) e^{-\theta_n \pi_{C,T_t^C}^\theta(j) g_{A_t,t}}}.$$

Therefore,

$$\begin{aligned} \pi_{C,T_t^C+1}^\theta(k) &\geq \frac{\pi_{C,T_t^C}^\theta(k) e^{\theta_n (\mathbf{1}_{A_t=k} - \pi_{C,T_t^C}^\theta(k)) g_{A_t,t}}}{\pi_{C,T_t^C}^\theta(A_t) e^{\theta_n (1 - \pi_{C,T_t^C}^\theta(A_t)) g_{A_t,t}} + \sum_{j \neq A_t} \pi_{C,T_t^C}^\theta(j)} \\ &\geq \frac{\pi_{C,T_t^C}^\theta(k) e^{-\theta_n}}{\pi_{C,T_t^C}^\theta(A_t) e^{\theta_n} + 1 - \pi_{C,T_t^C}^\theta(A_t)} \\ &\geq \pi_{C,T_t^C}^\theta(k) e^{-2\theta_n} \end{aligned}$$

where

- the first inequality holds because $e^{-\theta_n \pi_{C,T_t^C}^\theta(j) g_{A_t,t}} \leq 1$,
- the second inequality holds because $g_{j,t} \in (0, 1)$, for $j \in [K]$,
- the last inequality holds because $\pi_{C,T_t^C}^\theta(A_t) e^{\theta_n} + 1 - \pi_{C,T_t^C}^\theta(A_t) \leq e^{\theta_n}$.

Thus, for all $t \leq T_\varepsilon$,

$$\pi_{C, T_t^C}^\theta(k) \geq \frac{1}{K} e^{-2\theta_n T_t^C}.$$

Since $T_t^C \leq t$ and since by definition, $T_\varepsilon \leq \log\left(\sqrt{\frac{1}{K\varepsilon}}\right) \frac{\sqrt{n}}{R}$, it holds that for all $t \leq T_\varepsilon$,

$$\pi_{C, T_t^C}^\theta(k) \geq \frac{1}{K} e^{-2\theta_n t} \geq \frac{1}{K} e^{-2R_n t} \geq \frac{1}{K} e^{-2\frac{R_n}{\sqrt{n}} \log\left(\sqrt{\frac{1}{K\varepsilon}}\right) \frac{\sqrt{n}}{R}} \geq \frac{1}{K} e^{-\log\left(\frac{1}{K\varepsilon}\right)} \geq \varepsilon.$$

For the second part of the proof, for $t \geq 2$ and $j \in [K]$, let $h_{j, T_t^C}^\theta = \theta_n \sum_{s \in F_t(C)} \hat{g}_{j, s}^\theta$. Then $\pi_{C, T_t^C}^\theta = \text{softmax}(h_{\cdot, t}^\theta)$. The function softmax is 1-Lipschitz with respect to the $\|\cdot\|_2$ -norm in \mathbb{R}^K (see (Gao and Pavel, 2017) for a proof). Therefore,

$$\|\pi_{C, T_t^C}^\delta - \pi_{C, T_t^C}^\theta\|_2 \leq \|h_{\cdot, t}^\delta - h_{\cdot, t}^\theta\|_2.$$

Then,

$$\begin{aligned} \|h_{\cdot, t}^\delta - h_{\cdot, t}^\theta\|_2 &\leq |\delta_n - \theta_n| \sum_{s \in F_t(C)} \|\hat{g}_{j, s}^\delta\|_2 + \theta_n \sum_{s \in F_t(C)} g_{A_s, s} \|\pi_{C, T_s^C}^\delta - \pi_{C, T_s^C}^\theta\|_2 \\ &\leq \sqrt{2} T_t^C |\delta_n - \theta_n| + \theta_n \sum_{s \in F_t(C)} \|\pi_{C, T_s^C}^\delta - \pi_{C, T_s^C}^\theta\|_2 \\ &\leq \sqrt{2} T_\varepsilon |\delta_n - \theta_n| + \theta_n \sum_{s \in F_t(C)} \|\pi_{C, T_s^C}^\delta - \pi_{C, T_s^C}^\theta\|_2. \end{aligned}$$

where

- the first inequality holds because of the triangle inequality,
- the second inequality holds because for all $j \in [K]$, $g_{j, s} \in [0, 1]$ and

$$\|\hat{g}_{j, s}^\delta\|_2^2 = (1 - \pi_{C, T_s^C}^\delta(A_s))^2 + \sum_{j \neq A_s} (\pi_{C, T_s^C}^\delta(j))^2 \leq 2,$$

- the last inequality holds because $T_t^C \leq T_\varepsilon$.

By the discrete Gronwall Lemma (Clark, 1987), for all $t \leq T_\varepsilon$

$$\|\pi_{C, T_t^C}^\delta - \pi_{C, T_t^C}^\theta\|_2 \leq \sqrt{2} |\delta_n - \theta_n| T_\varepsilon \prod_{s \in F_t(C)} (1 + \theta_n) \leq \sqrt{2} |\delta_n - \theta_n| T_\varepsilon e^{\theta_n T_\varepsilon}.$$

Since $T_\varepsilon \leq \left\lceil \log\left(\sqrt{\frac{1}{K\varepsilon}}\right) \frac{\sqrt{n}}{R} \right\rceil$, $\theta_n T_\varepsilon \leq R_n T_\varepsilon \leq \log\left(\sqrt{\frac{1}{K\varepsilon}}\right)$, therefore,

$$\|\pi_{C, T_t^C}^\delta - \pi_{C, T_t^C}^\theta\|_2 \leq \frac{\sqrt{2} \log\left(\sqrt{\frac{1}{K\varepsilon}}\right)}{R} |\delta - \theta|.$$

Finally, log is $\frac{1}{\varepsilon}$ -Lipschitz on $[\varepsilon, 1]$. Thus, for all $k \in [K]$, $t \leq T_\varepsilon$,

$$\left| \log\left(\frac{\pi_{C, T_t^C}^\delta(k)}{\pi_{C, T_t^C}^\theta(k)}\right) \right| \leq \frac{1}{\varepsilon} |\pi_{C, T_t^C}^\delta(k) - \pi_{C, T_t^C}^\theta(k)| \leq \frac{\sqrt{2} \log\left(\sqrt{\frac{1}{K\varepsilon}}\right)}{R\varepsilon} |\delta - \theta|.$$

□

D.5 Proof of Section C

Let us recall that $|F| = \max_{m \in \mathcal{M}} |E_m|$. For this Section, we drop the dependence m of the model and simply write E and θ generic set of policies and parameter in $[r, R]$.

Proof of Proposition 6. For any $\theta \in [r, R]$, write $\theta_n = \theta/\sqrt{n}$. Assume that Assumption 4 holds. Let's write $q_{E,t+1}^\theta$ as

$$q_{E,t+1}^\theta(j) = \frac{q_{E,t}^\theta(j) e^{-\theta_n \hat{y}_{j,t}^\theta}}{\sum_{i \in E} q_{E,t}^\theta(i) e^{-\theta_n \hat{y}_{i,t}^\theta}}$$

Since $q_{E,t}^\theta$ is a probability distribution over the experts,

$$\sum_{i \in E} q_{E,t}^\theta(i) e^{-\theta_n \hat{y}_{i,t}^\theta} \leq \sum_{i \in E} q_{E,t}^\theta(i) = 1.$$

Therefore, $q_{E,t+1}^\theta(j) \geq q_{E,t}^\theta(j) e^{-\theta_n \hat{y}_{j,t}^\theta}$. By definition,

$$\hat{y}_{j,t}^\theta = \sum_{k=1}^K \xi_{j,t}(k) \frac{g_{k,t}}{\pi_{E,t}^\theta(k)} \mathbf{1}_{A_t=k} = \xi_{j,t}(A_t) \frac{g_{A_t,t}}{\pi_{E,t}^\theta(A_t)}.$$

Using that $e^{-x} \geq 1 - x$ for any $x \geq 0$, leads to

$$q_{E,t+1}^\theta(j) \geq q_{E,t}^\theta(j) \left(1 - \theta_n \xi_{j,t}(A_t) \frac{g_{A_t,t}}{\pi_{E,t}^\theta(A_t)} \right).$$

Since $g_{A_t,t} \in [0, 1]$ and $q_{E,t}^\theta(j) \xi_{j,t}(A_t) \leq \pi_{E,t}^\theta(A_t)$,

$$q_{E,t+1}^\theta(j) \geq q_{E,t}^\theta(j) - \theta_n.$$

Summing for all s from 1 to t ,

$$q_{E,t}^\theta(j) \geq \frac{1}{|E|} - \theta_n t.$$

Since

$$T_\varepsilon = \left\lfloor \left(\frac{1}{|F|} - \frac{\varepsilon}{\rho} \right) \frac{\sqrt{n}}{R} \right\rfloor \wedge n \quad \text{and} \quad |F| \geq |E|,$$

it holds that,

$$T_\varepsilon \leq \left\lfloor \left(\frac{1}{|E|} - \frac{\varepsilon}{\rho} \right) \frac{\sqrt{n}}{R} \right\rfloor \wedge n.$$

Therefore, for all $t \leq T_\varepsilon$,

$$q_{E,t}^\theta(j) \geq \frac{1}{|E|} - \frac{R}{\sqrt{n}} t \geq \frac{1}{|E|} - \frac{R}{\sqrt{n}} \left(\frac{1}{|E|} - \frac{\varepsilon}{\rho} \right) \frac{\sqrt{n}}{R} = \frac{\varepsilon}{\rho}.$$

Finally, for all $t \leq T_\varepsilon$, for all $k \in [K]$,

$$\pi_{E,t}^\theta(k) = \sum_{j \in E} q_{E,t}^\theta(j) \xi_{j,t}(k) \geq \frac{\varepsilon}{\rho} \sum_{j \in E} \xi_{j,t}(k) = \varepsilon.$$

For the second part of the proof, let $\eta, \delta \in [r, R]$, and write $\eta_n = \eta/\sqrt{n}$ and likewise for δ_n . For $t \geq 2$, let $g_{j,t}^\eta = \eta_n \sum_{s=1}^{t-1} \hat{y}_{j,s}^\eta$. Then, $q_{E,t}^\eta = \text{softmax}(g_t^\eta)$ where $g_t^\eta = (g_{j,t}^\eta)_{j \in E}$. Since the function softmax is 1-Lipschitz with respect to the $\|\cdot\|_2$ -norm in $\mathbb{R}^{|E|}$,

$$\|q_{E,t}^\eta - q_{E,t}^\delta\|_2 \leq \|g_t^\eta - g_t^\delta\|_2.$$

Therefore, by the triangle inequality,

$$\begin{aligned} \|q_{E,t}^\eta - q_{E,t}^\delta\|_2 &\leq \sum_{s=1}^{t-1} \|\eta_n \hat{y}_{j,s}^\eta - \delta_n \hat{y}_{j,s}^\delta\|_2 \\ &\leq \sum_{s=1}^{t-1} (\|\eta_n - \delta_n\| \|\hat{y}_{j,s}^\eta\|_2 + \delta_n \|\hat{y}_{j,s}^\eta - \hat{y}_{j,s}^\delta\|_2). \end{aligned} \quad (21)$$

Since $\xi_{j,t}$ is a probability distribution,

$$\begin{aligned} \|\hat{y}_{j,t}^\eta\|_2^2 &= \sum_{j \in E} (\hat{y}_{j,t}^\eta)^2 = \sum_{j \in E} \left(\sum_{k=1}^K \xi_{j,t}(k) \frac{g_{k,t}}{\pi_{E,t}^\eta(k)} \mathbf{1}_{A_t=k} \right)^2 \\ &= \sum_{j \in E} \left(\xi_{j,t}(A_t) \frac{g_{A_t,t}}{\pi_{E,t}^\eta(A_t)} \right)^2 \leq |E| \left(\frac{g_{A_t,t}}{\pi_{E,t}^\eta(A_t)} \right)^2. \end{aligned}$$

Since $g_{A_t,t}^\eta \in [0, 1]$ and $\pi_{E,t}^\eta(A_t) \geq \varepsilon$ for all $t \leq T_\varepsilon$,

$$\|\hat{y}_{j,t}^\eta\|_2 \leq \frac{\sqrt{|E|}}{\varepsilon}. \quad (22)$$

Similarly,

$$\begin{aligned} \|\hat{y}_{j,s}^\eta - \hat{y}_{j,s}^\delta\|_2^2 &= \sum_{j \in E} \left(\sum_{k=1}^K \xi_{j,t}(k) g_{k,t} \left(\frac{1}{\pi_{E,t}^\eta(k)} - \frac{1}{\pi_{E,t}^\delta(k)} \right) \mathbf{1}_{A_t=k} \right)^2 \\ &\leq \sum_{j \in E} \left(\sum_{k=1}^K \xi_{j,t}(k) \left(\frac{1}{\pi_{E,t}^\eta(k)} - \frac{1}{\pi_{E,t}^\delta(k)} \right) \mathbf{1}_{A_t=k} \right)^2. \end{aligned}$$

Thus, for all $t \leq T_\varepsilon$,

$$\|\hat{y}_{j,t}^\eta - \hat{y}_{j,t}^\delta\|_2^2 \leq \frac{1}{\varepsilon^4} \sum_{j \in E} \left(\sum_{k=1}^K \xi_{j,t}(k) (\pi_{E,t}^\eta(k) - \pi_{E,t}^\delta(k)) \mathbf{1}_{A_t=k} \right)^2.$$

Since $\xi_{j,t}(k) \leq 1$,

$$\|\hat{y}_{j,t}^\eta - \hat{y}_{j,t}^\delta\|_2 \leq \frac{\sqrt{|E|}}{\varepsilon^2} \|\pi_{E,t}^\eta - \pi_{E,t}^\delta\|_2. \quad (23)$$

Injecting (22) and (23) in (21) leads to

$$\|q_{E,t}^\eta - q_{E,t}^\delta\|_2 \leq \frac{\sqrt{|E|}}{\varepsilon} |\eta_n - \delta_n| (t-1) + \delta_n \frac{\sqrt{|E|}}{\varepsilon^2} \sum_{s=1}^{t-1} \|\pi_{E,s}^\eta - \pi_{E,s}^\delta\|_2.$$

On the other hand,

$$\begin{aligned} \|\pi_{E,t}^\eta - \pi_{E,t}^\delta\|_2^2 &= \sum_{k=1}^K (\pi_{k,t}^\eta - \pi_{k,t}^\delta)^2 = \sum_{k=1}^K \left(\sum_{j \in E} \xi_{j,t}(k) (q_{E,t}^\eta(j) - q_{E,t}^\delta(j)) \right)^2 \\ &\leq \sum_{k=1}^K \sum_{j \in E} \xi_{j,t}(k)^2 \sum_{j \in E} (q_{E,t}^\eta(j) - q_{E,t}^\delta(j))^2 \\ &\leq |E| \|q_{E,t}^\eta - q_{E,t}^\delta\|_2^2 \\ &\leq |F| \|q_{E,t}^\eta - q_{E,t}^\delta\|_2^2 \end{aligned}$$

where

- the first inequality holds by the Cauchy-Schwarz inequality,
- the second inequality holds because $\xi_{j,t}$ is a probability distribution over the actions set $[K]$,
- the last inequality holds because $|E| \leq |F|$

Therefore,

$$\|\pi_{E,t}^\eta - \pi_{E,t}^\delta\|_2 \leq \frac{|F|}{\varepsilon^2} \left(\varepsilon |\eta_n - \delta_n| (t-1) + \delta_n \sum_{s=1}^{t-1} \|\pi_{E,s}^\eta - \pi_{E,s}^\delta\|_2 \right).$$

Using the discrete Gronwall Lemma, for all $t \leq T_\varepsilon$,

$$\|\pi_{E,t}^\eta - \pi_{E,t}^\delta\|_2 \leq \frac{|F|}{\varepsilon} |\eta_n - \delta_n| T_\varepsilon \prod_{s=1}^{t-1} \left(1 + \frac{|F|}{\varepsilon^2} \delta_n \right) \leq \frac{|F|}{\varepsilon} |\eta_n - \delta_n| T_\varepsilon \exp \left(\frac{|F|}{\varepsilon^2} \delta_n T_\varepsilon \right).$$

If Assumption 4 is satisfied, then since $\delta_n \leq R_n = \frac{R}{\sqrt{n}}$ and $T_\varepsilon \leq \left(\frac{1}{|F|} - \frac{\varepsilon}{\rho} \right) \frac{\sqrt{n}}{R}$,

$$\|\pi_{E,t}^\eta - \pi_{E,t}^\delta\|_2 \leq \frac{|F|}{R\varepsilon} \left(\frac{1}{|F|} - \frac{\varepsilon}{\rho} \right) \exp \left(\frac{|F|}{\varepsilon^2} \left(\frac{1}{|F|} - \frac{\varepsilon}{\rho} \right) \right) |\eta - \delta|.$$

To conclude note that $x \rightarrow \ln(x)$ is $1/\varepsilon$ -Lipschitz on $[\varepsilon, +\infty)$. □