

# Fully General Online Imitation Learning

**Michael K. Cohen**

*Department of Engineering Science  
University of Oxford  
Future of Humanity Institute  
Oxford, UK OX1 3PJ*

MICHAEL.COHEN@ENG.OX.AC.UK

**Marcus Hutter**

*DeepMind  
Department of Computer Science  
Australian National University  
Acton, ACT, Australia 2601*

MARCUS.HUTTER@ANU.EDU.AU

**Neel Nanda**

*Independent*

NEELNANDA27@GMAIL.COM

**Editor:** Joelle Pineau

## Abstract

In imitation learning, imitators and demonstrators are policies for picking actions given past interactions with the environment. If we run an imitator, we probably want events to unfold similarly to the way they would have if the demonstrator had been acting the whole time. In general, one mistake during learning can lead to completely different events. In the special setting of environments that restart, existing work provides formal guidance in how to imitate so that events unfold similarly, but outside that setting, no formal guidance exists. We address a fully general setting, in which the (stochastic) environment and demonstrator never reset, not even for training purposes, and we allow our imitator to learn online from the demonstrator. Our new conservative Bayesian imitation learner underestimates the probabilities of each available action, and queries for more data with the remaining probability. Our main result: if an event would have been unlikely had the demonstrator acted the whole time, that event’s likelihood can be bounded above when running the (initially totally ignorant) imitator instead. Meanwhile, queries to the demonstrator rapidly diminish in frequency. If any such event qualifies as “dangerous”, our imitator would have the notable distinction of being relatively “safe”.

**Keywords:** Bayesian Sequence Prediction, Imitation Learning, Active Learning, General Environments

## 1. Introduction

Supervised learning of independent and identically distributed data is often practiced in two phases: training and deployment. This separation makes less sense if the learner’s predictions affect the distribution of future contexts for prediction, since the deployment phase could lose all resemblance to the training phase. When a program’s output changes its future percepts, we often call its output “actions”. Supervised learning in that regime is commonly called “imitation learning”, where labels are the actions of a “demonstrator”

(Syed and Schapire, 2010). Our agent, acting in a general environment that responds to its actions, tries to pick actions according to the same distribution as a demonstrator.

Even in imitation learning, where it is understood that actions can change the distribution of contexts that the agent will face, it is common to separate a training phase from a deployment phase. This assumes away the possibility that the distribution of contexts will shift significantly upon deployment and render the training data increasingly irrelevant. Here, we present an online imitation learner that is robust to this possibility.

The obvious downside is that the training never ends. The agent can always make queries for more data, but importantly, it does this with diminishing probability. It transitions smoothly from a mostly-training phase to a mostly-deployed phase. Our agent also handles totally general stochastic environments (environments serve new contexts for the agent to act in) and totally general stochastic demonstrator policies. No finite-state-Markov-style stationarity assumption is required for either. The lack of assumptions about the environment is a mundane point, because imitation learners don't have to learn the dynamics of the environment, but the lack of assumptions on the prediction target—the demonstrator's policy—makes these results highly non-trivial. The only assumption is that the demonstrator's policy belongs to some known countable class of possibilities. Moreover, stochasticity makes single-elimination-style learning (Gold, 1967) impossible.

For demonstrator policies this general, we present formal results that are unthinkable in the train-then-deploy paradigm. The  $\ell_1$  distance between the imitator and demonstrator policies converges to 0 in mean cube, when conditioned on a high-probability event (Theorem 3). And Theorem 2 shows that the event has high probability. Conditioned on the same high-probability event, we bound the KL divergence from imitator to demonstrator (Theorem 4), and we upper bound the probability of an arbitrary event under the imitator's policy, given a low probability of occurrence under the demonstrator's policy (Theorem 5). Instead of having a finite training phase, our agent's query probability converges to 0 in mean cube (Theorem 1). Without Theorems 1 and 2, the remaining theorems would be uninteresting; they would be easily fulfilled by an imitator that always queried the demonstrator, or they would apply only rarely.

Our imitator maintains a posterior over demonstrator models. At each timestep, it takes the top few demonstrator models in the posterior, in a way that depends on a scalar parameter  $\alpha$ . Then, for each action, it considers the minimum over those models of the probability that the demonstrator picks that action. The imitator samples an action according to those probabilities, and if no action is sampled (since model disagreement makes the probabilities to sum to less than 1), it defers to the demonstrator.

We review theoretical developments in imitation learning in Section 2, define our formal setting in Section 3, define our imitation learner in Section 4, and illustrate it with a toy example in Section 5. We state key formal results in Section 6, and we outline our proof technique and introduce necessary notation in Section 7. Section 8 presents lemmas and intermediate results, and Section 9 presents proofs and proof ideas of our key results, but most of the proofs appear in Appendix B. Appendix A collects notation and definitions.

## 2. Related Work

Recall that a key difficulty of imitation learning over supervised learning is the removal of a standard i.i.d. assumption. However, all existing formal work in imitation learning studies repeated finite episodes of length  $T$ ; even though the dynamics are not i.i.d. from timestep to timestep within an episode, the agent learns from a sequence of episodes that are, as a whole, independent and identically distributed. Thus, the scope of existing formal work is limited to environments that “restart”. A driving agent that gets housed in a new car every time it crashes (or gets hopelessly lost) enjoys a “restarting” environment, whereas a driving agent with only one car to burn does not. If we can accurately simulate a non-restarting environment, then training the imitator in simulation (using existing formal methods) could indeed prepare it to act in a non-restarting one. The viability of this approach depends on the environment; for many, we simply cannot simulate them with enough accuracy. For example, consider imitating a sales rep at a software company, interfacing with potential clients over email. For a real potential client, a relationship cannot be rebooted, and no simulation could anticipate the many diverse needs of clients.

In the context of restarting environments, Syed and Schapire (2010) reduce the problem of predicting a demonstrator’s behavior to i.i.d. classification. The only assumption about the demonstrator is that the value of its policy as a function of state is arbitrarily well approximated by the value of a deterministic policy, which is only slightly weaker than assuming the demonstrator is deterministic itself. They make no assumptions about the environment, other than that we can access identical copies of it repeatedly. They show that if a classifier guessing the demonstrator’s actions has an error rate of  $\varepsilon$ , then the value of the imitator’s policy that uses the classifier is within  $O(\sqrt{\varepsilon})$  of the demonstrator.

Judah et al. (2014) improve the label complexity of Syed and Schapire’s (2010) reduction by actively deciding when to query the demonstrator, instead of simply observing  $N$  full episodes before acting. Making the same assumptions as that paper, and also assuming a realizable hypothesis class with a finite VC dimension, they attempt to reduce the number of queries before the agent can act for a whole episode on its own with an error rate less than  $\varepsilon$ . Letting  $T$  be the length of an episode, compared to Syed and Schapire’s (2010)  $O(T^3/\varepsilon)$  labels, they achieve  $O(T \log(T^3/\varepsilon))$ .

Ross and Bagnell (2010) also reduce the problem to classification. In a trivial reduction, the imitator observes the demonstrator act from the distribution of states induced by the demonstrator policy. In this reduction, if the classifier has an error rate of  $\varepsilon$  per action on the demonstrator’s state distribution, the error rate of the imitator on its own distribution is at most  $T^2\varepsilon$ , where  $T$  is again the length of the episode. Their main contribution is to introduce a cleverer training regime for the classifier to reduce this bound to  $T\varepsilon$  in environments with approximate recoverability.

Ross et al. (2011) reduce imitation learning to something else: a no-regret online learner, for which the average error rate over its lifetime approaches 0, even with a potentially changing loss function. With access to an online learner with average regret  $O(1/N_{\text{predictions}})$ , they construct an imitation learner with regret of the same order. Unlike Syed and Schapire (2010) and Judah et al. (2014), they make no assumption that the demonstrator is arbitrarily well-approximated by a deterministic policy. Unlike Judah et al. (2014), they do not assume a realizable hypothesis class with a finite VC dimension. And unlike the Ross and

Bagnell (2010) (for their main contribution), they do not assume approximate recoverability. They do still assume that we can repeatedly access identical copies of the environment, and the loss function used for their measurement of regret must be bounded. To achieve a regret of order  $O(1/N_{\text{predictions}})$  with probability at least  $1 - \delta$ , they require  $O(T^2 \log(1/\delta))$  observations of the demonstrator.

There is a great deal of empirical study of imitation learning, given the practical applications, which Hussein et al. (2017) review. We call a few specific experiments to the reader’s attention, since they resemble our work in taking an active approach to querying, with an eye to risk aversion, not just label efficiency; they find it works. First, Brown et al. (2018, 2020) consider a context where the imitator can, at any time, ask the demonstrator how it would act in any of finitely many states. These imitators focus on states that they assign higher value at risk. Those papers and the following all show strong label efficiency alongside limited loss. Zhang and Cho (2017) assume some method of predicting the error of an imitator in the process of learning, and they query for help when it is above some threshold. Otherwise, their imitator follows Ross et al.’s (2011) construction. In their paper, the function that predicts the imitator’s error is learned from hand-picked features of a dataset. Menda et al. (2019) query much more extensively, but like Zhang and Cho (2017), they don’t always act on the demonstrator’s suggestion, in order to sample a more diverse set of states. Unlike Zhang and Cho (2017), they *do* act on it when the imitator’s action deviates enough from the demonstrator’s (given some hand-designed distance metric over the action space). They also defer to the demonstrator when there is sufficient disagreement among an ensemble of imitators. They find their imitator is more robust. Hoque et al. (2021) note that in many contexts, it is more convenient for the demonstrator to be queried a few times successively, rather than spread out over a long time. They modify Zhang and Cho’s (2017) approach: the imitator starts querying when the estimated error exceeds the same threshold, but it continues querying until it returns below a lower threshold. At the cost of more total queries, it requires fewer query-periods. Like the formal work, all these experiments regard environments that restart.

Adjacent to pure imitation learning (trying to pick the same actions as a demonstrator would), there is also work on trying to act in pursuit of the same goals as a demonstrator (which must be inferred), or matching only some outcomes of the demonstrator policy, like the expectation of some given set of features. For a review of some work in this area, see Adams et al. (2022).

### 3. Preliminaries

Let  $a_t \in \mathcal{A}$  and  $o_t \in \mathcal{O}$  be the action and observation at timestep  $t \in \mathbb{N}$ . Let  $q_t \in \{0, 1\}$  denote whether the imitator ( $q_t = 0$ ) or demonstrator ( $q_t = 1$ ) selects  $a_t$ . Let  $\mathcal{H} = \{0, 1\} \times \mathcal{A} \times \mathcal{O}$ , and let  $h_t = (q_t, a_t, o_t) \in \mathcal{H}$ . Let  $h_{<t} = (h_0, h_1, \dots, h_{t-1})$ .  $\mathcal{X}^n = \times_{i=1}^n \mathcal{X}$  denotes the set of  $n$ -tuples of elements of  $\mathcal{X}$ , and  $\mathcal{X}^* = \bigcup_{n=0}^{\infty} \mathcal{X}^n$  is the Kleene-star operator, which denotes all tuples of elements of  $\mathcal{X}$ .

Let  $\pi : \mathcal{H}^* \rightsquigarrow \{0, 1\} \times \mathcal{A}$ , and  $\rightsquigarrow$  denotes that  $\pi$  gives a distribution over  $\{0, 1\} \times \mathcal{A}$ .  $\epsilon$  will denote the empty string; it is the element of  $\mathcal{H}^0$ .  $\pi$  is called a policy, and will typically be written  $\pi(q_t a_t \mid h_{<t})$ .  $\pi(a_t \mid h_{<t})$  denotes the marginal distribution over the action. Let  $\mu : \mathcal{H}^* \times \{0, 1\} \times \mathcal{A} \rightsquigarrow \mathcal{O}$ .  $\mu$  is called the environment, and will typically

be written  $\mu(o_t | h_{<t}q_t a_t)$ . Note from this construction that an environment and a policy may qualitatively change over time—instead of being stationary with respect to the latest timestep, they can depend on the whole history.

Much formal work in imitation learning and reinforcement learning involves defining environments in terms of their Markov states and how one transitions through them. The defining property of a state is that that future is independent of the past conditioned on the state. For those more comfortable in that framework, our state space here is  $\mathcal{H}^*$ , so the Markov property trivial: the state is the whole history, so indeed, the future is independent of the history, when conditioned on the history. The point of the Markov Decision Process formalism is that when the state space is finite (or compact, with relevant functions of it being continuous), more tractable inference algorithms become available, but we do not assume finiteness or any structure in the state space. For finite histories denoted  $h_{<t}$ , the reader could mentally substitute  $s_t$ , this being the state at time  $t$ , but the infinite history  $h_{<\infty}$ , which appears in some proofs, has no standard notational analog.

Speaking of which, let  $\mathcal{H}^\infty$  be the set of infinite strings of elements of  $\mathcal{H}$ . Let  $P_\mu^\pi$  be the probability measure over  $\mathcal{H}^\infty$  where query records and actions are sampled from  $\pi$ , and observations are sampled from  $\mu$ . The event space is the standard sigma algebra over cylinder sets  $\sigma(\{\{h_{<t}h_{t:\infty} : h_{t:\infty} \in \mathcal{H}^\infty\} : h_{<t} \in \mathcal{H}^*\})$ . In a stochastic process, a cylinder set is the set of all possible futures given a particular past.

Let  $\Pi$  be a finite or countable set of policies, and for  $\pi \in \Pi$ , let  $w(\pi) > 0$  be a prior weight assigned to  $\pi$ , such that  $\sum_{\pi \in \Pi} w(\pi) = 1$ . This represents the imitator’s initial belief distribution over the demonstrator’s policy. For convenience, let  $\Pi$  only contain policies which assign zero probability to  $q_t = 0$ , since demonstrator models may as well be convinced that the demonstrator is picking the action.

**Example 1 ((Linear-Time) Computable Policies)** *The requirement that  $\Pi$  be countable is not restrictive in theory. Suppose  $\Pi$  is the set of programs that compute a policy (in linear time). These can be easily enumerated, and the prior  $w$  can be set  $\propto 2^{-\text{program length}}$  (Kraft, 1949; Hutter, 2005).*

Given the near absence of constraints, the choice of model class might pique philosophical interest. There are multiple logics with differing powers that we could plausibly use to represent programs, including “programs” higher in the arithmetic hierarchy. In general, the choice of programming language would change programs’ relative length, and there are no clear desiderata when choosing a language. So Example 1 does not appear to offer an approach to solving the Problem of Priors (Talbot, 2016). The option to restrict to linear-time programs is a marginally more practical possibility that might escape most philosophical discussions.

#### 4. Imitation

Let  $w(\pi | h_{<t})$  be the posterior weight after observing  $h_{<t}$  that demonstrator-chosen actions were sampled from  $\pi$ . That is,

$$w(\pi | h_{<t}) \propto w(\pi) \prod_{k<t:q_k=1} \pi(q_k a_k | h_{<k}) \tag{1}$$

normalized such that  $\sum_{\pi \in \Pi} w(\pi \mid h_{<t}) = 1$ . Ranking the policies by posterior weight, let  $\pi_n^{h_{<t}}$  be the one with the  $n^{\text{th}}$  largest posterior weight  $w(\pi \mid h_{<t})$ , breaking ties arbitrarily. Now let  $\Pi_{h_{<t}}^\alpha$  be the set of policies with posterior weights at least  $\alpha$  times the sum of the posterior weights of policies that are at least as likely as it; that is,

$$\Pi_{h_{<t}}^\alpha := \{\pi_n^{h_{<t}} \in \Pi : w(\pi_n^{h_{<t}} \mid h_{<t}) \geq \alpha \sum_{m \leq n} w(\pi_m^{h_{<t}} \mid h_{<t})\} \quad (2)$$

This is the set of policies the imitator takes seriously. The imitator is designed to be robust to policies in this set, so smaller  $\alpha$  will make it more robust. Let  $\pi^d$  denote the demonstrator’s policy, defined such that  $\pi^d(q_t = 1 \mid h_{<t}) = 1$  for all values of  $h_{<t}$ . As later results suggest,  $\alpha$  should be set a few orders of magnitude below  $w(\pi^d)$ ; since  $\pi^d$  is probably unknown to the programmers, or else there would be no need for imitation learning,  $w(\pi^d)$  will have to be estimated. The imitator’s policy  $\pi_\alpha^i$  is defined in the next two equations:

$$\pi_\alpha^i(0, a \mid h_{<t}) := \min_{\pi' \in \Pi_{h_{<t}}^\alpha} \pi'(1, a \mid h_{<t}) \quad (3)$$

The 0 on the l.h.s. means the imitator is picking the action itself instead of deferring to the demonstrator, and the 1 on the r.h.s. means this is the probability of the demonstrator model  $\pi'$  picking that same action.

The imitator uses the leftover probability to query. Let  $\theta_q(h_{<t}) := 1 - \sum_{a \in \mathcal{A}} \pi_\alpha^i(0, a \mid h_{<t})$ .  $\theta_q$  is the probability with which the imitator queries the demonstrator to have it pick the action. Thus,

$$\pi_\alpha^i(1, a \mid h_{<t}) := \theta_q(h_{<t}) \pi^d(1, a \mid h_{<t}) \quad (4)$$

One can see that  $q_t$  records whether the demonstrator was involved in selecting the action. Using the model class and prior from Example 1, the time-complexity constraint makes  $\pi_\alpha^i$  computable.

Conservatism with respect to probability estimates is a core technical innovation of our work. Taking the minimum over a set of models with high posterior weights is an approach to conservatism inspired by Cohen and Hutter’s (2020) pessimistic agent. The pessimistic agent, unlike ours, is a reinforcement learner, but it is also designed to keep certain (risky) events unlikely. By underestimating probabilities, the imitator only acts if it is sure the demonstrator might act that way.

We will also consider hypothetical imitator policies if the demonstrator policy were something else; for an arbitrary demonstrator policy  $\pi$ , let  $\hat{\pi}_\alpha$  denote the corresponding imitator policy, so  $\pi_\alpha^i = (\hat{\pi}_\alpha^d)_\alpha$ . This paper will investigate the probability distribution  $P_\mu^{\pi_\alpha^i}$  and compare it to  $P_\mu^{\pi^d}$ .

## 5. Toy Example

We now walk through a toy example, in which our imitation learner has about a half-million demonstrator models in its model class  $\Pi$ . We begin by defining  $\Pi$ . The action space  $\mathcal{A}$  of the demonstrator is  $\text{null} \cup \{0, 1\}^4$ . The observation space  $\mathcal{O}$  is  $\{\text{“”}, 1, 2, 3\}$ . A demonstrator model  $\pi \in \Pi$  defined by is a 12-tuple of the elements  $\{1/3, 2/3, 1\}$ . When the latest observation is 1, 2, or 3, let  $x$  be the 1<sup>st</sup> - 4<sup>th</sup>, 5<sup>th</sup> - 8<sup>th</sup>, or 9<sup>th</sup> - 12<sup>th</sup> elements of 12-tuple.

Then, the demonstrator model outputs four bits that are Bernoulli distributed according to each of the four elements of  $x$ . All demonstrator models output `null` when the latest observation is `""`. The true demonstrator also takes the form of such a demonstrator model. Each observation is randomly sampled; it is 1 with probability 1/4, 2 with probability 1/16, 3 with probability 1/64, and otherwise `""`.

Let’s give some flavor to this example. The demonstrator does client relations for a high-end travel agency with very fussy clients. The demonstrator gets a feel for her clients, and for any given night that a client needs a restaurant recommendation, the demonstrator sends a Boolean 4-tuple to the restaurant team, who identifies a suitable restaurant. The observation tells the demonstrator which of the three clients needs a recommendation, if any. The first bit of the Boolean 4-tuple tells the restaurant team whether the restaurant should have lots of vegetarian options, the second bit: should it have a Michelin star, the third: should it have unfamiliar local specialties, and the fourth: should it be Instagrammable. Why is the demonstrator stochastic? Many clients want a variety of styles of restaurants from night to night. The demonstrator couldn’t write down the exact probabilities that she is using to generate these Boolean vectors; she goes off of intuition. If we run an imitator that only sometimes asks the demonstrator for help, we can free up some of the demonstrator’s time.

Unfortunately, in this toy environment, the fussy clients sometimes quit. Each client has a 4-tuple of probabilities that they would like their Boolean vector sampled from (conveniently in  $\{1/3, 2/3, 1\}^4$ ). If it becomes clear that this is not how their Boolean vectors are being sampled, they quit. (“Becoming clear” is operationalized as follows:  $H_1$  is the hypothesis that their restaurant recommendations are being sampled correctly;  $H_2$  is the hypothesis that some other 4-tuple in  $\{1/3, 2/3, 1\}^4$  is producing their restaurant recommendations. If, given the set of all restaurant recommendations they have gotten, the likelihood ratio of  $H_2$  exceeds 100, the client quits. Note that this happens if an element is ever `False` when it was supposed to be `True` with probability 1; some clients demand Michelin stars.) When recommendations are made by the demonstrator, who always correctly intuits the client’s desired distribution of restaurants, clients hardly ever quit. We would like clients to hardly ever quit even when the imitator frequently takes over.

For an imitator with  $\alpha = 1e-14$ , Figure 1 shows how often it has to query the demonstrator to pick the restaurant features. Recommendations are random, and this is only one run. Running it with 20 different random seeds, the number of queries required is  $486.75 \pm 52.63$  (out of  $2^{15}$  timesteps), and no client ever quit. Returning to run depicted in Figure 1, Table 1 works through an example of the posterior and the imitator’s behavior. The code for this toy example can be found at <https://tinyurl.com/imitation-toy-example>.

## 6. Results

For the whole of the paper, we assume:

**Assumption 1 (Realizability)**  $\pi^d \in \Pi$ .

That is, the imitator can conceive of the demonstrator. There may be some interesting results in the setting of approximate realizability, where  $\exists \pi \in \Pi$  such that  $\pi \approx_\epsilon \pi^d$  in some sense, but that is out of our scope here.

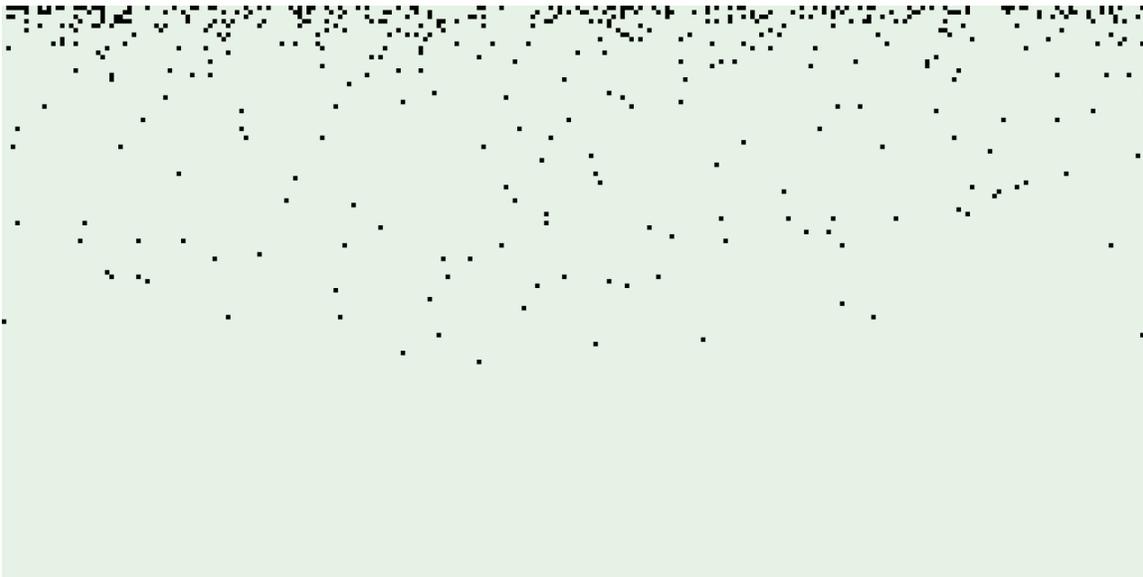


Figure 1: Timesteps when the imitator queries.  $2^{15}$  timesteps are shown, with black representing a query, and green representing the imitator acting unassisted. Pixels are to be read like text, left to right, top to bottom. In the accompanying code, a random seed of 0 is used to generate this image.

We now state and discuss our key results before turning to selected proofs. Our first is that the imitator’s query probability converges to 0 in mean cube. This result

- renders its resemblance to the demonstrator non-trivial, since always querying would yield perfect correspondence,
- is desirable in its own right if demonstrator access is a limited resource,
- and is instrumental in proving the remaining results, since low query probability implies little model disagreement.

**Theorem 1 (Limited Querying)**

$$\mathbb{E}_{\mu^i}^{\pi_\alpha} \left[ \sum_{t=0}^{\infty} \theta_q(h_{<t})^3 \right] \leq |\mathcal{A}| \alpha^{-3} (24w(\pi^d)^{-1} + 12)$$

The in mean cube bound allows infinite querying, but it diminishes in frequency, or else the expectation of an infinite sum of cubed probabilities would not be finite. Since we query under uncertainty, both querying and uncertainty diminish in tandem; this is a theme for active learners in general. Error bounds in Bayesian prediction and MAP prediction tend to be  $\Theta(\log(w(\text{truth})^{-1}))$  and  $\Theta(w(\text{truth})^{-1})$  respectively, so theoretically, our case resembles the MAP one. The cubic dependence on  $\alpha$  is unfortunate, and subsequent results inherit them; the only path we found to proving a bound was fairly circuitous, and we are unsure whether this dependence can be improved.

1/3	2/3	1		
<b>-0.0000</b>	-44.0000	-inf		
-50.0000	<b>-0.0000</b>	-inf		
<b>0.0000</b>	-78.0000	-inf		
-46.0000	<b>-0.0000</b>	-inf		
			$p([\text{False}, \text{True}, \text{False}, \text{True}] \mid \text{client 2})$	Model
-18.0000	<b>-0.0000</b>	-inf	0.11111	(..., 2/3, 1, 2/3, 1, ...)
-69.7384	-25.7384	<b>-0.0000</b>	0.14815	(..., 2/3, 1/3, 1/3, 1, ...)
<b>-0.0000</b>	-22.0000	-inf	0.22222	(..., <b>2/3, 1, 1/3, 1, ...</b> )
-69.7384	-25.7384	<b>-0.0000</b>	0.29630	(..., 1/3, 2/3, 1/3, 1, ...)
			0.44444	(..., 1/3, 1, 1/3, 1, ...)
<b>-0.3219</b>	-2.3219	-inf		
<b>-0.0056</b>	-8.0056	-inf		
<b>-0.0000</b>	-16.0000	-inf		
-28.5303	-10.5303	<b>-0.0010</b>		

Table 1: *Left:*  $\log_2$  posterior at timestep 1000 for the run depicted in Figure 1. The posterior decomposes into posterior probabilities for each of 12 features. Each block is a client, each row is a feature, and each entry is the log posterior probability that the demonstrator picks **True** for that feature with probability 1/3, 2/3, or 1, respectively. To get the posterior for a whole demonstrator model, as in Equation 1, add the independent posteriors for each element in the 12-tuple of the demonstrator model. The posterior weight on the truth is in bold for each feature; that is, the true demonstrator for this run is (1/3, 2/3, 1/3, 2/3, 2/3, 1, 1/3, 1, 1/3, 1/3, 1/3, 1). *Right:* At timestep 1000, with  $\alpha = 1e-14$ , we have many top models, as defined in Equation 2. The first column is a list of probabilities that different top models assign to the outcome [False True False True] for client 2. The second column contains examples of top models that assign those probabilities to the outcome [False True False True] for client 2, with the true model in bold. Recall a demonstrator model is defined by a 12-tuple, but the only relevant elements for client 2 are 5-8. All these models have posterior weight large enough to make it into the top set. Thus, the probability the imitator picks [False True False True] for client 2 is 0.11111, the minimum probability shown, as per Equation 3.

Our remaining results show that the imitator resembles the demonstrator on one condition:  $\pi^d \in \Pi_{h < t}^\alpha$ . Recall that  $\Pi_{h < t}^\alpha$  is a set of top demonstrator models that the imitator takes seriously, and  $\pi^d$  is the true demonstrator model. Low model disagreement implies high accuracy when the truth is one of those models, and recall that our querying regime promises low model disagreement within finite time.

Fortunately, this condition has high probability for  $\alpha \ll w(\pi^d)$ .

**Theorem 2 (Top Models Contain Truth)**  $P_\mu^{\pi^i}(\forall t : \pi^d \in \Pi_{h < t}^\alpha) \geq 1 - \alpha w(\pi^d)^{-1}$

Let  $E$  be the event  $\forall t : \pi^d \in \Pi_{h < t}^\alpha$ , so the true demonstrator policy is always in the top set. The high probability of  $E$  is mainly of interest in the context of subsequent results that depend on it. For instance, conditioned on  $E$ , the imitator, when picking its own actions, converges to the demonstrator in mean cube.

**Theorem 3 (Predictive Convergence)** For  $\alpha < w(\pi^d)$ ,

$$\mathbb{E}_{\mu}^{\pi_{\alpha}^i} \left[ \sum_{t=0}^{\infty} \left( \sum_{a \in \mathcal{A}} |\pi_{\alpha}^i(0, a | h_{<t}) - \pi^d(1, a | h_{<t})| \right)^3 \middle| E \right] \leq \frac{|\mathcal{A}| \alpha^{-3} (24w(\pi^d)^{-1} + 12)}{1 - \alpha w(\pi^d)^{-1}}$$

This theorem finally justifies our calling  $\pi_{\alpha}^i$  an “imitator”, since the policy converges to that of the demonstrator. Existing literature on imitation learning does little to suggest that imitators exist in non-restarting environments. This result shows that they do, at least in a high-probability sense. Note that the denominator is the probability of  $E$ , which will be nearly 1 for appropriate choice of  $\alpha$ . The requirement that  $\alpha < w(\pi^d)$  has important consequence: when  $\alpha$  is set appropriately, the bounds in this theorem and Theorem 1 are effectively quartic in  $w(\pi^d)^{-1}$ . We do not know if a better rate is possible under additional assumptions. It is even possible that stronger results are available without additional assumptions, and we simply failed to identify them. We think this is a ripe area for research.

We argue informally that this disappointing dependence can be mitigated in some circumstances. By pre-training with  $N$  consecutive demonstrator queries and calling the posterior at that point the new “prior” for the purposes of our analysis, the “prior” on  $w(\pi^d)$  could usually be made quite large, unless most demonstrator models behave extremely similarly for the first  $N$  steps. Consider an extreme case: many models of comparable weight almost agree with the true model, except one disagrees at  $t = 1$ , one at  $t = 2$ , etc. In this case, the posterior on the truth increases very slightly every step, as models are excluded one by one. If, on the other hand, half of demonstrator models confidently predict one action, and half confidently predict another, the posterior on the truth will likely nearly double in one step. So to the extent that a large fraction of models in  $\Pi$  disagree with  $\pi^d$  within the first  $N$  steps, the posterior on the truth would increase exponentially following pre-training. That said, the quartic dependence on  $w(\pi^d)^{-1}$  in the worst case is a weakness of our approach.

Any pair of these first three results would be uninteresting on their own, but jointly, they show that with high probability, the imitator converges to the demonstrator with limited querying.

Our stronger results below apply when the environment and demonstrator policy do not depend on the query record. This means that whatever action is taken, the effect does not depend on whether the imitator chose it or the demonstrator did. We would like events to unfold similarly when we replace the demonstrator with the imitator, but this is impossible if the environment discriminates between them. Indeed, if the environment treats identical actions differently depending on whether they were selected by imitator or demonstrator, it’s unclear what imitation accomplishes. We define fairness formally in Section 9.

In a fair setting, we bound the KL divergence between  $P_{\mu}^{\pi_{\alpha}^i}$  and  $P_{\mu}^{\pi^d}$ , the first meaning that actions are picked according to our imitation policy, and the second meaning that all actions are picked by the demonstrator. The objective of imitation is most easily characterized as outputting demonstrator-like actions, but the purpose of imitation learning is for events to unfold similarly. Small errors in the limit do not guarantee that property; this result is only possible with small errors for the imitator’s whole lifetime.

**Theorem 4 (KL Bound)** *Suppose that  $\mu$  and  $\pi^d$  are fair, and  $\alpha < w(\pi^d)$ . Letting the two probability measures below be restricted to  $(\mathcal{A} \times \mathcal{O})^t$  (that is, marginalizing over the query record, and considering only the first  $t$  timesteps),*

$$\text{KL}_t \left( \mathbb{P}_{\mu^\alpha}^i(\cdot | E) \parallel \mathbb{P}_{\mu}^{\pi^d}(\cdot | E) \right) \leq \frac{\alpha^{-1} |\mathcal{A}|^{1/3} (24w(\pi^d)^{-1} + 12)^{1/3}}{(1 - \alpha/w(\pi^d))^2} t^{2/3} - \log(1 - \alpha/w(\pi^d))$$

Notably,  $\text{KL}_t/t \rightarrow 0$  in the limit. The direction of the divergence resembles the variational objective (with the ground truth on the right). Thus, there may be some events that only the demonstrator would cause, but no events that only the imitator would. This consequence is made explicit in our final result.

We construct an upper bound for the probability of an event given the probability of the event if the demonstrator were acting the whole time. This bound is mainly of interest for “bad” events.

**Theorem 5 (Preserving Unlikeliness)** *Fix  $t$ . Let  $B \subset (\mathcal{A} \times \mathcal{O})^t$  be a (bad) event, and extending  $B$  to the outcome space  $(\{0, 1\} \times \mathcal{A} \times \mathcal{O})^t = \mathcal{H}^t$ , let  $D = B \cap E$ . Then, for fair  $\mu$  and  $\pi^d$ ,*

$$\mathbb{P}_{\mu^\alpha}^i(D) \leq \frac{t^2 s_\alpha}{\left( \log \frac{t^2 s_\alpha}{27 \mathbb{P}_{\mu}^{\pi^d}(B)} - 3 \log \log \left( 1 + \frac{t^{2/3} s_\alpha^{1/3}}{3 \mathbb{P}_{\mu}^{\pi^d}(B)^{1/3}} \right) \right)^3}$$

where  $s_\alpha = |\mathcal{A}| \alpha^{-3} (24w(\pi^d)^{-1} + 12)$ .

That is, as  $\mathbb{P}_{\mu}^{\pi^d}(B)^{-1} \rightarrow \infty$ ,  $\mathbb{P}_{\mu^\alpha}^i(D)^{-1} \rightarrow \infty$  at least polylogarithmically. If an event would have been extremely unlikely under the demonstrator’s policy, a similar event is unlikely when running the imitator.

Whereas existing work on imitation learners attempts to be robust to a bounded loss function, our Preserving Unlikeliness Theorem is relevant even in the absence of a uniform bound on badness. In the real world, to quote Theon Greyjoy, “It can always be worse”. But some bounds on badness are possible: we tolerate one-in-ten-chance events; they happen, and we get on with it. One-in-a-hundred-chance events can be meaningfully worse. But in a world largely governed by humans, we keep most truly devastating events below even a 1% chance. It’s hard to apply similar bounds to the badness of one-in-a-billion-chance events, and in general, as the probability gets smaller, a loss function should countenance steadily larger losses. When an event goes from a 1% to a 2% chance, we should be much less concerned than if it went from  $10^{-9}$  to 1%. In the extreme, if an event has probability 0 under a demonstrator’s policy, there might be an arbitrarily good reason for that. Whereas the bounded loss functions of all existing work ignore this effect, our Theorem 5 does not.

The main weaknesses of our results are what they require: a model class that includes the truth and a good choice of  $\alpha$ . Setting  $\alpha$  well requires estimating  $w(\pi^d)$ , something we cannot offer general guidance on; it would depend entirely on the exact nature of the prior. And realistically, in many contexts, the realizability assumption is infeasible. There will always be mismatch between a computational model of a demonstrator and the true demonstrator. We hope this paper opens the door for other research into relaxing the realizability assumption. Plausibly, if the best approximation in  $\Pi$  of  $\pi^d$  produces certain bad events with low probability, then the imitator will too.

## 7. Roadmap and Notation for the Proof of Theorem 1

Much of the work of this paper is to prove Theorem 1. In this section, we state a theorem on which it depends, and we introduce the mathematical objects required to prove it.

The imitator queries when the top few demonstrator models disagree, so we bound the errors that those models can make over the agent’s lifetime. We first must establish a finite bound on the errors of such models in ordinary Bayesian sequence prediction. We define that here.

Let  $\mathcal{X}$  be an arbitrary finite alphabet. Let  $\nu$  be a probability measure over  $\mathcal{X}^\infty$  with the event space generated by the cylinder sets  $\{\{x_{<t}x_{t:\infty} \mid x_{t:\infty} \in \mathcal{X}^\infty\} \mid x_{<t} \in \mathcal{X}^*\}$ . Let  $\mathcal{M}$  be a countable set of such probability measures, and let  $w(\nu)$  be a prior weight over these measures such that  $\sum_{\nu \in \mathcal{M}} w(\nu) = 1$ . Let  $x_{<t} \in \mathcal{X}^t$ , let  $\nu(x_{<t})$  denote the probability that the infinite sequence begins with  $x_{<t}$ , and let  $\nu(x \mid x_{<t}) = \nu(x_{<t}x) / \nu(x_{<t})$ . Let  $\mu \in \mathcal{M}$  be a the “true” measure; that is, in formal results, we will let  $x_{<\infty}$  be sampled from  $\mu$ .

Let  $\nu_n^{x_{<t}}$  be the measure with the  $n^{\text{th}}$  largest posterior weight after observing  $x_{<t}$ ; that is, order  $\mathcal{M}$  to be non-increasing in  $w(\nu)\nu(x_{<t})$ , breaking ties arbitrarily, and take the  $n^{\text{th}}$ . (Ties between any pair should be broken consistently for different  $t$ ). Let the posterior  $w(\nu \mid x_{<t}) \propto w(\nu)\nu(x_{<t})$ , normalized to sum to 1. Let  $\mathcal{M}_n^{x_{<t}}$  be the set of the top  $n$  measures, and let  $w(\mathcal{M}_n^{x_{<t}} \mid x_{<t}) = \sum_{m \leq n} w(\nu \mid x_{<t})$ .

Recall a model belongs to the imitator’s top set if its posterior weight is at least  $\alpha$  times the sum of the posterior weights of the models that are at least as good. Thus, we define

$$\phi_n^{x_{<t}} := \frac{w(\nu_n^{x_{<t}} \mid x_{<t})}{w(\mathcal{M}_n^{x_{<t}} \mid x_{<t})} \tag{5}$$

So if  $\phi_n^{x_{<t}} \leq \alpha$ , then  $\nu_n^{x_{<t}}$  can be considered a “top model” in the same sense that is relevant to our imitation learner.

Our key result on which Theorem 1 is based shows that taking the minimum over predictions in the top measures converges to the truth, and the “missing probability” converges to 0.

### Theorem 6 (Top Model Convergence)

$$\begin{aligned} (i) \quad & \mathbb{E}_\mu \sum_{t=0}^\infty \sum_{x \in \mathcal{X}} \left[ \mu(x \mid x_{<t}) - \min_{n: \phi_n^{x_{<t}} > \alpha} \nu_n^{x_{<t}}(x \mid x_{<t}) \right]^2 \leq \alpha^{-3}(24w(\mu)^{-1} + 12) \\ (ii) \quad & \mathbb{E}_\mu \sum_{t=0}^\infty \left[ 1 - \sum_{x \in \mathcal{X}} \min_{n: \phi_n^{x_{<t}} > \alpha} \nu_n^{x_{<t}}(x \mid x_{<t}) \right]^2 \leq |\mathcal{X}| \alpha^{-3}(24w(\mu)^{-1} + 12) \end{aligned}$$

This is perfectly analogous to the way the imitator predicts actions: taking the minimum over the top models for which  $\phi_n^{x_{<t}} > \alpha$ . The difference is that in this sequence prediction setting, all observations are informative about the true measure, whereas the imitator rarely sees the demonstrator act.

To prove Theorem 6, we show that a posterior-weighted mixture over  $\mathcal{M}_n^{x_{<t}}$  converges to the truth, and if  $\phi_n^{x_{<t}} > \alpha$ , then each constituent must as well. This posterior-weighted mixture is called  $\rho_n^{\text{stat}}$ . We define it here alongside other estimators that will be used in the proof of  $\rho_n^{\text{stat}}$ ’s convergence. First,

$$\rho_n^{\text{stat}}(x \mid x_{<t}) := \frac{\sum_{\nu \in \mathcal{M}_n^{x_{<t}}} w(\nu) \nu(x_{<t}x)}{\sum_{\nu \in \mathcal{M}_n^{x_{<t}}} w(\nu) \nu(x_{<t})} \quad (6)$$

$\rho_n^{\text{stat}}$  resembles a maximum a posteriori estimate, but instead mixes over the top few. We call it a satis magnam a posteriori estimate (SMAP). We will show  $\rho_n^{\text{stat}}$  converges to  $\rho_n$ , which converges to  $\rho_n^{\text{norm}}$ , which converges to  $\mu$ .  $\rho_n$  and  $\rho_n^{\text{norm}}$  are alternative SMAP estimators.

$\rho_n$  is not a measure, as the numerator below sums over a different set than the denominator. It sums over the top measures *after* observing  $x$ :

$$\rho_n(x \mid x_{<t}) := \frac{\sum_{\nu \in \mathcal{M}_n^{x_{<t}x}} w(\nu) \nu(x_{<t}x)}{\sum_{\nu \in \mathcal{M}_n^{x_{<t}}} w(\nu) \nu(x_{<t})} \quad (7)$$

The definition appears more natural when considering a whole sequence:

$$\rho_n(x_{<t}) = \sum_{\nu \in \mathcal{M}_n^{x_{<t}}} w(\nu) \nu(x_{<t}) \quad (8)$$

Since  $\sum_{x \in \mathcal{X}} \rho_n(x \mid x_{<t})$  may not be 1, we construct the measure  $\rho_n^{\text{norm}}$  by normalizing:

$$\rho_n^{\text{norm}}(x \mid x_{<t}) := \frac{\rho_n(x \mid x_{<t})}{\sum_{x' \in \mathcal{X}} \rho_n(x' \mid x_{<t})} = \frac{\rho_n(x_{<t}x)}{\sum_{x' \in \mathcal{X}} \rho_n(x_{<t}x')} \quad (9)$$

Our  $\rho_n$ ,  $\rho_n^{\text{norm}}$ , and  $\rho_n^{\text{stat}}$  are closely inspired by Poland and Hutter (2005), who constructed (in our notation)  $\rho_1$ ,  $\rho_1^{\text{norm}}$ , and  $\rho_1^{\text{stat}}$ . Finally, we define the full Bayes-mixture measure

$$\xi(x_{<t}) := \sum_{\nu \in \mathcal{M}} w(\nu) \nu(x_{<t}) = \rho_\infty^{\text{stat}}(x_{<t}) = \rho_\infty(x_{<t}) = \rho_\infty^{\text{norm}}(x_{<t}) \quad (10)$$

We state those relationships without proof for the reader's interest; they are not used in our results.

## 8. General Sequence Prediction Results

This section organizes the proof of Theorem 6 into lemmas, some of which are proven here and some in Appendix B. We begin with elementary relations between  $\xi$ ,  $\rho_n$ ,  $\rho_n^{\text{norm}}$ , and  $\rho_n^{\text{stat}}$ .

$$\xi(x_{<t}) \geq \rho_n(x_{<t}) \quad (11)$$

$$\rho_n(x_{<t}) \geq w(\mu) \mu(x_{<t}) \quad (12)$$

$$\rho_n(x \mid x_{<t}) \geq \rho_n^{\text{norm}}(x \mid x_{<t}) \quad (13)$$

$$\rho_n(x \mid x_{<t}) \geq \rho_n^{\text{stat}}(x \mid x_{<t}) \quad (14)$$

Inequalities 11 and 12 follow directly from Equation 8. Inequality 13 follows because

$$\rho_n(x_{<t}) = \max_{\mathcal{M}' \subset \mathcal{M}: |\mathcal{M}'|=i} \sum_{\nu \in \mathcal{M}'} w(\nu) \nu(x_{<t}) = \max_{\mathcal{M}' \subset \mathcal{M}: |\mathcal{M}'|=i} \sum_{\nu \in \mathcal{M}'} w(\nu) \sum_{x \in \mathcal{X}} \nu(x_{<t}x)$$

$$\leq \sum_{x \in \mathcal{X}} \max_{\mathcal{M}' \subset \mathcal{M}: |\mathcal{M}'|=i} \sum_{\nu \in \mathcal{M}'} w(\nu) \nu(x_{<t}x) = \sum_{x \in \mathcal{X}} \rho_n(x_{<t}x) \quad (15)$$

so  $\rho_n$  assigns too much probability mass. Inequality 14 follows because

$$\rho_n^{\text{stat}}(x | x_{<t}) = \frac{\sum_{\nu \in \mathcal{M}_n^{x_{<t}x}} w(\nu) \nu(x_{<t}x)}{\sum_{\nu \in \mathcal{M}_n^{x_{<t}}} w(\nu) \nu(x_{<t})} \leq \frac{\sum_{\nu \in \mathcal{M}_n^{x_{<t}x}} w(\nu) \nu(x_{<t}x)}{\rho_n(x_{<t})} = \frac{\rho_n(x_{<t}x)}{\rho_n(x_{<t})} \quad (16)$$

which holds because  $\mathcal{M}_n^{x_{<t}x}$  is chosen to maximize the numerator.

Our first lemma bounds the normalizing factor for  $\rho_n$ , allowing us to show in our next lemma that it converges to both  $\rho_n^{\text{norm}}$  and  $\rho_n^{\text{stat}}$ .

**Lemma 7**

$$0 \leq \mathbb{E}_\mu \sum_{t=0}^{\infty} \frac{\sum_{x \in \mathcal{X}} \rho_n(x_{<t}x)}{\rho_n(x_{<t})} - 1 \leq w(\mu)^{-1}$$

**Proof idea**  $\rho_n$  is bounded above and below by measures, save a multiplicative constant (Inequalities 11 and 12), so  $\rho_n$  converges to being a measure, in that  $\sum_{x \in \mathcal{X}} \rho_n(x | x_{<t}) \rightarrow 1$ .  
□

**Proof** All terms in the sum are non-negative, by Inequality 15. Recall  $\epsilon$  denotes the empty string—the element of  $\mathcal{X}^0$ . Justifications of the upcoming lettered equations follow below the block.

$$\begin{aligned} & \mathbb{E}_\mu \sum_{t=0}^{N-1} \frac{\sum_{x \in \mathcal{X}} \rho_n(x_{<t}x)}{\rho_n(x_{<t})} - 1 \\ &= \sum_{t=0}^{N-1} \sum_{x_{<t} \in \mathcal{X}^t} \mu(x_{<t}) \frac{\sum_{x \in \mathcal{X}} \rho_n(x_{<t}x) - \rho_n(x_{<t})}{\rho_n(x_{<t})} \\ &\stackrel{(a)}{\leq} \sum_{t=0}^{N-1} \sum_{x_{<t} \in \mathcal{X}^t} w(\mu)^{-1} \left[ \sum_{x \in \mathcal{X}} \rho_n(x_{<t}x) - \rho_n(x_{<t}) \right] \\ &\stackrel{(b)}{=} w(\mu)^{-1} \left[ \sum_{x_{<N} \in \mathcal{X}^N} \rho_n(x_{<N}) - \rho_n(\epsilon) \right] \\ &\stackrel{(c)}{\leq} w(\mu)^{-1} \sum_{x_{<N} \in \mathcal{X}^N} \xi(x_{<N}) = w(\mu)^{-1} \end{aligned} \quad (17)$$

where (a) follows from Inequality 12, (b) cancels terms that are added then subtracted, and (c) follows from Inequality 11. ■

Recall we are trying to show  $\rho_n^{\text{stat}} \rightarrow \rho_n \rightarrow \rho_n^{\text{norm}} \rightarrow \mu$ . The following lemma gives two of those links.

**Lemma 8**

$$\begin{aligned}
 (i) \quad & \mathbb{E}_\mu \sum_{t=0}^{\infty} \sum_{x \in \mathcal{X}} |\rho_n(x | x_{<t}) - \rho_n^{\text{stat}}(x | x_{<t})| \leq w(\mu)^{-1} \\
 (ii) \quad & \mathbb{E}_\mu \sum_{t=0}^{\infty} \sum_{x \in \mathcal{X}} |\rho_n(x | x_{<t}) - \rho_n^{\text{norm}}(x | x_{<t})| \leq w(\mu)^{-1}
 \end{aligned}$$

**Proof**

$$\begin{aligned}
 & \mathbb{E}_\mu \sum_{t=0}^{\infty} \sum_{x \in \mathcal{X}} |\rho_n(x | x_{<t}) - \rho_n^{\text{stat}}(x | x_{<t})| \stackrel{(a)}{=} \mathbb{E}_\mu \sum_{t=0}^{\infty} \sum_{x \in \mathcal{X}} \rho_n(x | x_{<t}) - \rho_n^{\text{stat}}(x | x_{<t}) = \\
 & \mathbb{E}_\mu \sum_{t=0}^{\infty} \frac{\sum_{x \in \mathcal{X}} \rho_n(x_{<t}x)}{\rho_n(x_{<t})} - 1 \stackrel{(b)}{\leq} w(\mu)^{-1} \tag{18}
 \end{aligned}$$

where (a) follows from Inequality 14 and (b) follows from Lemma 7. The proof is identical for  $\rho_n^{\text{norm}}$ , except now (a) follows from Inequality 13.  $\blacksquare$

Given Lemma 8, the final link in showing  $\rho_n^{\text{stat}}$  converges to  $\mu$  is to show that  $\rho_n^{\text{norm}}$  does.

**Lemma 9** *Recalling  $\nu(\cdot | x_{<t})$  is a measure over  $\mathcal{X}$ ,*

$$\mathbb{E}_\mu \sum_{t=0}^{\infty} \text{KL}(\mu(\cdot | x_{<t}) \| \rho_n^{\text{norm}}(\cdot | x_{<t})) \leq w(\mu)^{-1} + \log w(\mu)^{-1}$$

**Proof idea** The KL divergence telescopes over timesteps. The  $\log w(\mu)^{-1}$  term comes from a gap between  $\mu$  and  $\rho_n$ , and the  $w(\mu)^{-1}$  term comes from a gap between  $\rho_n$  and  $\rho_n^{\text{norm}}$ .  $\square$

We can now show that  $\rho_n^{\text{stat}}$  converges to  $\mu$ , an independently interesting and novel result in SMAP estimation.

**Theorem 10 (SMAP Convergence)**

$$\mathbb{E}_\mu \sum_{t=0}^{\infty} \sum_{x \in \mathcal{X}} (\rho_n^{\text{stat}}(x | x_{<t}) - \mu(x | x_{<t}))^2 \leq 6w(\mu)^{-1} + 3$$

**Proof idea**  $\rho_n^{\text{stat}}$  is close to  $\rho_n$  in an  $\ell_1$  sense, and likewise for  $\rho_n$  and  $\rho_n^{\text{norm}}$ , and  $\rho_n^{\text{norm}}$  is close to  $\mu$  in an  $\ell_2$  squared sense, since  $\ell_2^2 \leq \text{KL}$ . Finally, for a vector  $v \in [-1, 1]^n$ ,  $\|v\|_2^2 \leq \|v\|_1$ , so  $\ell_1$  proximity implies  $\ell_2$  proximity as well.  $\square$

By applying Theorem 10 to the very similar measures  $\rho_n^{\text{stat}}$  and  $\rho_{n-1}^{\text{stat}}$ , whose only difference is that the former contains  $\nu_n^{x_{<t}}$  in its mixture, we arrive at our final result in the general sequence prediction setting.

**Theorem 6 (Top Model Convergence)**

$$\begin{aligned}
 (i) \quad & \mathbb{E}_\mu \sum_{t=0}^{\infty} \sum_{x \in \mathcal{X}} \left[ \mu(x | x_{<t}) - \min_{n: \phi_n^{x_{<t}} > \alpha} \nu_n^{x_{<t}}(x | x_{<t}) \right]^2 \leq \alpha^{-3} (24w(\mu)^{-1} + 12) \\
 (ii) \quad & \mathbb{E}_\mu \sum_{t=0}^{\infty} \left[ 1 - \sum_{x \in \mathcal{X}} \min_{n: \phi_n^{x_{<t}} > \alpha} \nu_n^{x_{<t}}(x | x_{<t}) \right]^2 \leq |\mathcal{X}| \alpha^{-3} (24w(\mu)^{-1} + 12)
 \end{aligned}$$

**Proof idea**  $\rho_n^{\text{stat}}$  is a weighted average of  $\nu_m^{x_{<t}}$  for  $m \leq n$ , so convergence results for  $\rho_n^{\text{stat}}$  and  $\rho_{n-1}^{\text{stat}}$  are leveraged for  $\nu_n^{x_{<t}}$ 's convergence.  $\phi_n^{x_{<t}} > \alpha$  ensures the weights in the weighted average aren't too small, and that we only need to consider the top  $\lfloor 1/\alpha \rfloor$  models.  $\square$

## 9. Key Proofs

We now prove our bound on the query probability, we define fairness, and we prove our bound on the probabilities of bad events.

**Theorem 1 (Limited Querying)**

$$\mathbb{E}_\mu^{\pi_\alpha^i} \left[ \sum_{t=0}^{\infty} \theta_q(h_{<t})^3 \right] \leq |\mathcal{A}| \alpha^{-3} (24w(\pi^d)^{-1} + 12)$$

**Proof idea** The sort of model mismatch bounded by Theorem 6 (ii) is the basis for the definition of  $\theta_q$ . Theorem 6 (ii) bounds model mismatch on *observed* data, and data is only observed with probability  $\theta_q$ , so with an extra factor of  $\theta_q$  on the l.h.s., we go from an in mean square bound to a weaker in mean cube bound.  $\square$

**Proof** Recall the agent considers a set of possible policies  $\Pi$  that includes the true demonstrator policy  $\pi^d$ , and assigns a strictly positive prior  $w(\pi)$  to each policy in  $\Pi$ . Recall  $\mathbb{P}_\mu^\pi$  is a probability measure over  $(\{0, 1\} \times \mathcal{A} \times \mathcal{O})^\infty = \mathcal{H}^\infty$ . Now we construct a class of measures over  $\mathcal{H}^\infty$ : let  $\mathcal{M} := \{\mathbb{P}_\mu^{\hat{\pi}_\alpha} : \pi \in \Pi\}$  (see the last paragraph of Section 4 for the definition of  $\hat{\pi}_\alpha$ ), and let  $w(\mathbb{P}_\mu^{\hat{\pi}_\alpha}) := w(\pi)$ . Let  $w(\mathbb{P}_\mu^{\hat{\pi}_\alpha} | h_{<t}) := w(\mathbb{P}_\mu^{\hat{\pi}_\alpha}) \mathbb{P}_\mu^{\hat{\pi}_\alpha}(h_{<t})$ . It follows straightforwardly from the definitions of the posterior that  $w(\mathbb{P}_\mu^{\hat{\pi}_\alpha} | h_{<t}) = w(\pi | h_{<t})$ ,  $w(\mathbb{P}_\mu^{\hat{\pi}_\alpha} | h_{<t}q_t) = w(\pi | h_{<t}q_t)$ , and  $w(\mathbb{P}_\mu^{\hat{\pi}_\alpha} | h_{<t}q_t a_t) = w(\pi | h_{<t}q_t a_t)$ , since all measures in  $\mathcal{M}$  assign the probabilities identically to actions after  $q_t = 0$ , and to observations.

Instead of saying  $\mathcal{M}$  contains measures over  $\mathcal{X}^\infty$ , we generalize slightly, and say that  $\mathcal{M}$  contains measures over  $\times_{k=0}^{\infty} \mathcal{X}_k$ . For  $k \equiv 0 \pmod{3}$ ,  $\mathcal{X}_k = \{0, 1\}$ , for  $k \equiv 1 \pmod{3}$ ,  $\mathcal{X}_k = \mathcal{A}$ , and for  $k \equiv 2 \pmod{3}$ ,  $\mathcal{X}_k = \mathcal{O}$ . With  $\nu_n^{x_{<k}}$  and  $\phi_n^{x_{<k}}$  as defined before, we can apply Theorem 6 (i) to the class  $\mathcal{M}$ , after a trivial extension from fixed  $\mathcal{X}$  to variable  $\mathcal{X}_k$ . Checking the definitions is enough to verify that  $\{\nu_n^{x_{<k}} : \phi_n^{x_{<k}} > \alpha\}$  is exactly the set  $\{\mathbb{P}_\mu^{\hat{\pi}_\alpha} : \pi \in \Pi_{h_{<t}}^\alpha\}$ , where  $h_j = (q_j, a_j, o_j) = (x_{3j}, x_{3j+1}, x_{3j+2})$ , and  $t = \lfloor (k+1)/3 \rfloor$ . In short, for this  $\mathcal{M}$ , sequence prediction errors can only come from errors predicting actions after querying, since that's when models differ, so we can use Theorem 6 to bound the latter. Recalling that  $\mathbb{P}_\mu^{\pi_\alpha^i}$  is the true probability measure,

$$\begin{aligned}
 \alpha^{-3}(24w(\pi^d)^{-1} + 12) &= \alpha^{-3}(24w(\mathbb{P}_\mu^{\pi_\alpha^i})^{-1} + 12) \\
 &\stackrel{(a)}{\geq} \mathbb{E}_\mu^{\pi_\alpha^i} \sum_{k=0}^{\infty} \sum_{x \in \mathcal{X}_k} \left[ \mathbb{P}_\mu^{\pi_\alpha^i}(x | x_{<k}) - \min_{i: \phi_n^{x_{<k}} > \alpha} \nu_n^{x_{<k}}(x | x_{<k}) \right]^2 \\
 &\stackrel{(b)}{=} \mathbb{E}_\mu^{\pi_\alpha^i} \sum_{t=0}^{\infty} \sum_{q \in \{0,1\}} \left[ \mathbb{P}_\mu^{\pi_\alpha^i}(q | h_{<t}) - \min_{\pi \in \Pi_{h_{<t}}^\alpha} \mathbb{P}_\mu^{\hat{\pi}_\alpha}(q | h_{<t}) \right]^2 + \\
 &\quad \sum_{a \in \mathcal{A}} \left[ \mathbb{P}_\mu^{\pi_\alpha^i}(a | h_{<t}q_t) - \min_{\pi \in \Pi_{h_{<t}}^\alpha} \mathbb{P}_\mu^{\hat{\pi}_\alpha}(a | h_{<t}q_t) \right]^2 + \\
 &\quad \sum_{o \in \mathcal{O}} \left[ \mathbb{P}_\mu^{\pi_\alpha^i}(o | h_{<t}q_t a_t) - \min_{\pi \in \Pi_{h_{<t}}^\alpha} \mathbb{P}_\mu^{\hat{\pi}_\alpha}(o | h_{<t}q_t a_t) \right]^2 \\
 &\stackrel{(c)}{=} \mathbb{E}_\mu^{\pi_\alpha^i} \sum_{t=0}^{\infty} \sum_{a \in \mathcal{A}} \left[ \mathbb{P}_\mu^{\pi_\alpha^i}(a | h_{<t}q_t) - \min_{\pi \in \Pi_{h_{<t}}^\alpha} \mathbb{P}_\mu^{\hat{\pi}_\alpha}(a | h_{<t}q_t) \right]^2 \\
 &= \mathbb{E}_\mu^{\pi_\alpha^i} \sum_{t=0}^{\infty} \sum_{q \in \{0,1\}} \mathbb{P}_\mu^{\pi_\alpha^i}(q | h_{<t}) \sum_{a \in \mathcal{A}} \left[ \mathbb{P}_\mu^{\pi_\alpha^i}(a | h_{<t}q) - \min_{\pi \in \Pi_{h_{<t}}^\alpha} \mathbb{P}_\mu^{\hat{\pi}_\alpha}(a | h_{<t}q) \right]^2 \\
 &\stackrel{(d)}{=} \mathbb{E}_\mu^{\pi_\alpha^i} \sum_{t=0}^{\infty} \mathbb{P}_\mu^{\pi_\alpha^i}(1 | h_{<t}) \sum_{a \in \mathcal{A}} \left[ \mathbb{P}_\mu^{\pi_\alpha^i}(a | h_{<t}1) - \min_{\pi \in \Pi_{h_{<t}}^\alpha} \mathbb{P}_\mu^{\hat{\pi}_\alpha}(a | h_{<t}1) \right]^2 \\
 &\stackrel{(e)}{\geq} \mathbb{E}_\mu^{\pi_\alpha^i} \sum_{t=0}^{\infty} \theta_q(h_{<t}) |\mathcal{A}| \left[ |\mathcal{A}|^{-1} \sum_{a \in \mathcal{A}} \mathbb{P}_\mu^{\pi_\alpha^i}(a | h_{<t}1) - \min_{\pi \in \Pi_{h_{<t}}^\alpha} \mathbb{P}_\mu^{\hat{\pi}_\alpha}(a | h_{<t}1) \right]^2 \\
 &= |\mathcal{A}|^{-1} \mathbb{E}_\mu^{\pi_\alpha^i} \sum_{t=0}^{\infty} \theta_q(h_{<t}) \left[ 1 - \sum_{a \in \mathcal{A}} \min_{\pi \in \Pi_{h_{<t}}^\alpha} \frac{\hat{\pi}_\alpha(1, a | h_{<t})}{\hat{\pi}_\alpha(1 | h_{<t})} \right]^2 \\
 &\stackrel{(f)}{=} |\mathcal{A}|^{-1} \mathbb{E}_\mu^{\pi_\alpha^i} \sum_{t=0}^{\infty} \theta_q(h_{<t}) \left[ 1 - \sum_{a \in \mathcal{A}} \min_{\pi \in \Pi_{h_{<t}}^\alpha} \frac{\theta_q(h_{<t}) \pi(1, a | h_{<t})}{\theta_q(h_{<t})} \right]^2 \\
 &\stackrel{(g)}{=} |\mathcal{A}|^{-1} \mathbb{E}_\mu^{\pi_\alpha^i} \sum_{t=0}^{\infty} \theta_q(h_{<t}) [\theta_q(h_{<t})]^2 \tag{19}
 \end{aligned}$$

where (a) follows from Theorem 6, (b) groups triples  $(x_{3t}, x_{3t+1}, x_{3t+2})$  into  $h_t$ , (c) follows because all  $\mathbb{P}_\mu^{\hat{\pi}_\alpha} \in \mathcal{M}$  give identical conditional probabilities as  $\mathbb{P}_\mu^{\pi_\alpha^i}$  on queries and observations, (d) follows because all  $\mathbb{P}_\mu^{\hat{\pi}_\alpha} \in \mathcal{M}$  give identical conditional probabilities as  $\mathbb{P}_\mu^{\pi_\alpha^i}$  for actions that follow  $q_t = 0$ , (e) follows from Jensen's Inequality, (f) follows from the definition of  $\hat{\pi}_\alpha$ , and (g) follows from the definition of  $\theta_q(h_{<t})$ . Rearranging Inequality 19 gives the theorem.  $\blacksquare$

Recall that Theorem 3 bounds the error between  $\pi_\alpha^i$  and  $\pi^d$ , conditioned on the event  $E$ .

**Proof idea of Theorem 3** Conditioned on  $\pi^d \in \Pi_{h_{<t}}^\alpha$ , it follows that  $\theta_q \geq$  the  $\ell_1$  norm between  $\pi^d$  and  $\pi_\alpha^i$ . Then we apply Theorem 1.  $\square$

Our remaining theorems apply when the environment and demonstrator policy are fair. Roughly, they are fair if they do not have access to the imitator's internals.

**Definition 11 (Fair)** *An environment  $\mu : \mathcal{H}^* \times \{0, 1\} \times \mathcal{A} \rightsquigarrow \mathcal{O}$  is fair if it does not depend on the query record; that is,  $\mu(\cdot | h_{<t}q_t a_t)$  is not a function of  $q_k$  for  $k \leq t$ . A demonstrator policy  $\pi^d : \mathcal{H}^* \rightsquigarrow \{0, 1\} \times \mathcal{A}$  is likewise fair if  $\pi^d(\cdot | h_{<t})$  is not a function of  $q_k$  for  $k < t$ .*

Theorems 4 and 5 rest on the following crux: if  $\pi^d \in \Pi_{h_{<t}}^\alpha$ , then  $\pi_\alpha^i(0, a | h_{<t}) \leq \pi^d(a | h_{<t})$ . Since  $\pi_\alpha^i(1, a | h_{<t}) = \theta_q(h_{<t})\pi^d(a | h_{<t})$ , we have  $\pi_\alpha^i(a | h_{<t}) \leq (1 + \theta_q(h_{<t}))\pi^d(a | h_{<t})$ . Thus, we have a multiplicative bound relating  $\pi_\alpha^i$  and  $\pi^d$ , and it decreases to 1.

**Theorem 5 (Preserving Unlikeliness)** *Fix  $t$ . Let  $B \subset (\mathcal{A} \times \mathcal{O})^t$  be a (bad) event, and extending  $B$  to the outcome space  $(\{0, 1\} \times \mathcal{A} \times \mathcal{O})^t = \mathcal{H}^t$ , let  $D = B \cap E$ . Then, for fair  $\mu$  and  $\pi^d$ ,*

$$\mathbb{P}_\mu^{\pi_\alpha^i}(D) \leq \frac{t^2 s_\alpha}{\left( \log \frac{t^2 s_\alpha}{27 \mathbb{P}_\mu^{\pi^d}(B)} - 3 \log \log \left( 1 + \frac{t^{2/3} s_\alpha^{1/3}}{3 \mathbb{P}_\mu^{\pi^d}(B)^{1/3}} \right) \right)^3}$$

where  $s_\alpha = |\mathcal{A}| \alpha^{-3} (24w(\pi^d)^{-1} + 12)$ .

**Proof idea**  $\mathbb{P}_\mu^{\pi_\alpha^i}(B \cap E) / \mathbb{P}_\mu^{\pi^d}(B)$  increases by a factor of at most  $1 + \theta_q$  per timestep. While the expectation of  $\theta_q^3$  is summable, the expectation of  $\sum_t \theta_q$  grows as  $O(t^2)$ , hence that dependence in the bound. The final difficulty is that our bound on the query probability only applies in expectation, but a pathological and unlikely event  $B$  could describe a case where querying is much more prolonged than expected. Thus, we do not prove a nice bound on the ratio  $\mathbb{P}_\mu^{\pi_\alpha^i}(B \cap E) / \mathbb{P}_\mu^{\pi^d}(B)$ . Instead, since smaller  $\mathbb{P}_\mu^{\pi^d}(B)$  allows more pathology, our bound on  $\mathbb{P}_\mu^{\pi_\alpha^i}(B \cap E)$  is only polylogarithmic in  $\mathbb{P}_\mu^{\pi^d}(B)$ .  $\square$

**Proof** If  $\pi^d \in \Pi_{h_{<t}}^\alpha$ , then  $\pi_\alpha^i(0, a | h_{<t}) \leq \pi^d(a | h_{<t})$ , and of course  $\pi_\alpha^i(1, a | h_{<t}) = \theta_q(h_{<t})\pi^d(a | h_{<t})$ , so

$$\pi_\alpha^i(a | h_{<t}) \leq (1 + \theta_q(h_{<t}))\pi^d(a | h_{<t}) \quad (20)$$

Thus, for fair  $\mu$  and  $\pi^d$ , for  $h_{<t} \in E$ ,

$$\frac{\mathbb{P}_\mu^{\pi_\alpha^i}(h_{<t})}{\mathbb{P}_\mu^{\pi^d}(h_{<t})} \leq \prod_{k=0}^{t-1} [1 + \theta_q(h_{<k})] \quad (21)$$

It follows from Theorem 1 that

$$\mathbb{E}_\mu^{\pi_\alpha^i} \left[ \sum_{k=0}^{t-1} \theta_q(h_{<k})^3 \middle| D \right] \leq \frac{s_\alpha}{\mathbb{P}_\mu^{\pi_\alpha^i}(D)} \quad (22)$$

By the same derivation as in Inequality 59, we can thus bound the sum

$$\mathbb{E}_\mu^{\pi_\alpha^i} \left[ \sum_{k=0}^{t-1} \theta_q(h_{<k}) \middle| D \right] \leq t^{2/3} \left( \frac{s_\alpha}{P_\mu^{\pi_\alpha^i}(D)} \right)^{1/3} \quad (23)$$

Now, applying Inequality 20 repeatedly,

$$\begin{aligned} & \mathbb{E}_\mu^{\pi_\alpha^i} \left[ \prod_{k=0}^{t-1} (1 + \theta_q(h_{<k}))^{-1} \middle| D \right] \\ &= \frac{\sum_{h_{<t} \in D} P_\mu^{\pi_\alpha^i}(h_{<t}) \prod_{k=0}^{t-1} (1 + \theta_q(h_{<k}))^{-1}}{\sum_{h_{<t} \in D} P_\mu^{\pi_\alpha^i}(h_{<t})} \\ &\leq \frac{\sum_{h_{<t-1} \in E} \sum_{h_{t-1} \in \mathcal{H}: h_{<t} \in B} P_\mu^{\pi_\alpha^i}(h_{<t}) \prod_{k=0}^{t-1} (1 + \theta_q(h_{<k}))^{-1}}{\sum_{h_{<t} \in D} P_\mu^{\pi_\alpha^i}(h_{<t})} \\ &= \frac{\sum_{h_{<t-1} \in E} \left[ P_\mu^{\pi_\alpha^i}(h_{<t-1}) \prod_{k=0}^{t-2} (1 + \theta_q(h_{<k}))^{-1} \right] \sum_{h_{t-1} \in \mathcal{A} \times \mathcal{O}: h_{<t} \in B} P_\mu^{\pi_\alpha^i}(h_{t-1} \setminus | h_{<t-1}) (1 + \theta_q(h_{<t-1}))^{-1}}{\sum_{h_{<t} \in D} P_\mu^{\pi_\alpha^i}(h_{<t})} \\ &\stackrel{(a)}{\leq} \frac{\sum_{h_{<t-1} \in E} \left[ P_\mu^{\pi_\alpha^i}(h_{<t-1}) \prod_{k=0}^{t-2} (1 + \theta_q(h_{<k}))^{-1} \right] \sum_{h_{t-1} \in \mathcal{A} \times \mathcal{O}: h_{<t} \in B} P_\mu^{\pi^d}(h_{t-1} \setminus | h_{<t-1})}{\sum_{h_{<t} \in D} P_\mu^{\pi_\alpha^i}(h_{<t})} \\ &\stackrel{(b)}{\leq} \frac{\sum_{h_{<t-2} \in E} \left[ P_\mu^{\pi_\alpha^i}(h_{<t-2}) \prod_{k=0}^{t-3} (1 + \theta_q(h_{<k}))^{-1} \right] \sum_{h_{t-2} h_{t-1} \in (\mathcal{A} \times \mathcal{O})^2: h_{<t} \in B} P_\mu^{\pi^d}(h_{t-2} \setminus h_{t-1} \setminus | h_{<t-2})}{\sum_{h_{<t} \in D} P_\mu^{\pi_\alpha^i}(h_{<t})} \\ &\stackrel{(c)}{\leq} \frac{\sum_{h_{<t} \in B} P_\mu^{\pi^d}(h_{<t})}{\sum_{h_{<t} \in D} P_\mu^{\pi_\alpha^i}(h_{<t})} = \frac{P_\mu^{\pi^d}(B)}{P_\mu^{\pi_\alpha^i}(D)} \quad (24) \end{aligned}$$

where (a) follows from Inequality 20 since  $h_{<t-1} \in E$  (note the change from  $\pi_\alpha^i$  to  $\pi^d$ ), (b) iterates the previous three lines, and (c) iterates the logic down to 0.

Now we bound the expectation

$$\begin{aligned} \mathbb{E}_\mu^{\pi_\alpha^i} \left[ \prod_{k=0}^{t-1} (1 + \theta_q(h_{<k}))^{-1} \middle| D \right] &\stackrel{(a)}{\geq} \prod_{k=0}^{t-1} (1 + \mathbb{E}_\mu^{\pi_\alpha^i} [\theta_q(h_{<k}) \mid D])^{-1} \\ &= \exp \left( - \sum_{k=0}^{t-1} \log \left( 1 + \mathbb{E}_\mu^{\pi_\alpha^i} [\theta_q(h_{<k}) \mid D] \right) \right) \\ &\geq \exp \left( - \sum_{k=0}^{t-1} \mathbb{E}_\mu^{\pi_\alpha^i} [\theta_q(h_{<k}) \mid D] \right) \\ &\stackrel{(b)}{\geq} e^{-t^{2/3} s_\alpha^{1/3} P_\mu^{\pi_\alpha^i}(D)^{-1/3}} \quad (25) \end{aligned}$$

where (a) follows from Jensen's Inequality (one can easily show the Hessian of  $\prod_i 1/(1+x_i)$  is positive semidefinite for  $x \succ 0$ ), and (b) follows from Inequality 23. Solving for  $P_\mu^{\pi_\alpha^i}(D)$

in terms of  $P_\mu^{\pi^d}(B)$ , we get

$$P_\mu^{\pi^i}(D) \leq \frac{t^2 s_\alpha}{27W\left(\frac{t^{2/3}s_\alpha^{1/3}}{3P_\mu^{\pi^d}(B)^{1/3}}\right)^3} \quad (26)$$

where  $W$  is the Lambert- $W$  function, defined by the property  $W(z)e^{W(z)} = z$ . A property of the Lambert- $W$  function—that  $W(z) \geq \log z - \log \log(1+z)$ —yields the theorem:

$$P_\mu^{\pi^i}(D) \leq \frac{t^2 s_\alpha}{\left(\log \frac{t^2 s_\alpha}{27 P_\mu^{\pi^d}(B)} - 3 \log \log \left(1 + \frac{t^{2/3} s_\alpha^{1/3}}{3 P_\mu^{\pi^d}(B)^{1/3}}\right)\right)^3}$$

One can easily verify this inequality by supposing the opposite and showing that it violates Inequality 25, but we omit this.  $\blacksquare$

## 10. Conclusion

We present the first formal results for an imitation learner in a setting where the environment does not reset. We present the first formal results for an imitation learner that do not depend on a bounded loss assumption. We present the first finite error bounds for an agent acting in general environments; existing results only regard limiting behavior (although existing work considers reinforcement learning, a harder problem than imitation learning). If we would like to have an artificial agent imitate, with particular concern for keeping unlikely events unlikely, this is the first theory of how to do it.

## Acknowledgments

This work was supported by the Leverhulme Centre for the Future of Intelligence, Australian Research Council Discovery Projects DP150104590, the Oxford-Man Institute, and the Berkeley Existential Risk Initiative. Thank you to Jan Leike for encouraging us to write a paper on imitation learning.

## References

- Stephen Adams, Tyler Cody, and Peter A Beling. A survey of inverse reinforcement learning. *Artificial Intelligence Review*, pages 1–40, 2022.
- Daniel Brown, Scott Niekum, and Marek Petrik. Bayesian robust optimization for imitation learning. *Advances in Neural Information Processing Systems*, 33:2479–2491, 2020.
- Daniel S Brown, Yuchen Cui, and Scott Niekum. Risk-aware active inverse reinforcement learning. In *Conference on Robot Learning*, pages 362–372. PMLR, 2018.
- Michael K Cohen and Marcus Hutter. Pessimism about unknown unknowns inspires conservatism. In *Conference on Learning Theory*, pages 1344–1373. PMLR, 2020.

- R Durrett. *Probability: Theory and Examples*. Cambridge University Press, Cambridge, 2010.
- E Mark Gold. Language identification in the limit. *Information and control*, 10(5):447–474, 1967.
- Ryan Hoque, Ashwin Balakrishna, Carl Putterman, Michael Luo, Daniel S Brown, Daniel Seita, Brijen Thananjeyan, Ellen Novoseller, and Ken Goldberg. Lazydagger: Reducing context switching in interactive imitation learning. In *2021 IEEE 17th International Conference on Automation Science and Engineering (CASE)*, pages 502–509. IEEE, 2021.
- 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.
- Marcus Hutter. *Universal Artificial Intelligence: Sequential Decisions based on Algorithmic Probability*. Springer, Berlin, 2005. ISBN 3-540-22139-5. doi: 10.1007/b138233.
- Kshitij Judah, Alan P Fern, Thomas G Dietterich, and Prasad Tadepalli. Active imitation learning: Formal and practical reductions to iid learning. *Journal of Machine Learning Research*, 15(120):4105–4143, 2014.
- Leon Gordon Kraft. *A device for quantizing, grouping, and coding amplitude-modulated pulses*. PhD thesis, Massachusetts Institute of Technology, 1949.
- Kunal Menda, Katherine Driggs-Campbell, and Mykel J Kochenderfer. Ensembledagger: A bayesian approach to safe imitation learning. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5041–5048. IEEE, 2019.
- Jan Poland and Marcus Hutter. Asymptotics of discrete MDL for online prediction. *IEEE Transactions on Information Theory*, 51(11):3780–3795, 2005. ISSN 0018-9448. doi: 10.1109/TIT.2005.856956.
- Stéphane Ross and Drew Bagnell. Efficient reductions for imitation learning. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 661–668, 2010.
- Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635, 2011.
- Umar Syed and Robert E Schapire. A reduction from apprenticeship learning to classification. *Advances in neural information processing systems*, 23:2253–2261, 2010.
- William Talbott. Bayesian Epistemology. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Winter 2016 edition, 2016.
- Jiakai Zhang and Kyunghyun Cho. Query-efficient imitation learning for end-to-end simulated driving. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31, 2017.

## Appendix A. Notation and Definitions

Notation	Meaning
<i>Preliminary Notation</i>	
$\mathcal{A}, \mathcal{O}$	the finite action/observation spaces
$a_t, o_t$	$\in \mathcal{A}, \mathcal{O}$ ; the action and observation at timestep $t$
$q_t$	$\in \{0, 1\}$ ; indicates whether the demonstrator is queried at time $t$
$\mathcal{H}$	$\{0, 1\} \times \mathcal{A} \times \mathcal{O}$
$h_t$	$(q_t, a_t, o_t)$ ; the interaction history in the $t^{\text{th}}$ timestep
$h_{<t}$	$(h_1, \dots, h_{t-1})$
$h_t^\setminus$	$(a_t, o_t)$
$\epsilon$	the empty history
$\pi$	policy stochastically mapping $\mathcal{H}^* \rightsquigarrow \{0, 1\} \times \mathcal{A}$
$\mu$	environment stochastically mapping $\mathcal{H}^* \times \{0, 1\} \times \mathcal{A} \rightsquigarrow \mathcal{O}$
$\mathbb{P}_\nu^\pi$	a probability measure over histories with actions sampled from $\pi$ and observations sampled from $\nu$
$\mathbb{E}_\nu^\pi$	the expectation when the interaction history is sampled from $\mathbb{P}_\nu^\pi$
$w(\pi)$	(positive) prior weight that the policy $\pi$ is the demonstrator's
$w(\pi \mid h_{<t})$	posterior weight on the policy $\pi$ ; $\propto w(\pi) \prod_{k<t: q_k=1} \pi(q_k a_k \mid h_{<k})$
<i>Imitation Learner Definition</i>	
$\alpha$	$\in (0, 1]$ ; lower values mean the imitator better resembles the demonstrator, but queries longer
$\Pi_{h_{<t}}^\alpha$	set of top models; $\{\pi_n^{h_{<t}} \in \Pi : w(\pi_n^{h_{<t}} \mid h_{<t}) \geq \alpha \sum_{m<n} w(\pi_m^{h_{<t}} \mid h_{<t})\}$
$\pi^d$	the demonstrator's policy
$\pi_\alpha^i$	the imitator's policy; $\pi_\alpha^i(0, a \mid h_{<t}) = \min_{\pi' \in \Pi_{h_{<t}}^\alpha} \pi'(1, a \mid h_{<t})$ , and $\pi_\alpha^i(1, a \mid h_{<t}) = \theta_q(h_{<t}) \pi^d(1, a \mid h_{<t})$
$\theta_q(h_{<t})$	the query probability; $1 - \sum_{a \in \mathcal{A}} \pi_\alpha^i(0, a \mid h_{<t})$
$\hat{\pi}_\alpha$	the imitator policy defined with respect to an arbitrary demonstrator $\pi$ , not the real demonstrator $\pi^d$
<i>General Sequence Prediction</i>	
$\mathcal{X}$	finite alphabet
$x_{<t}$	an element of $\mathcal{X}^t$
$\mathcal{M}$	countable set of measures over $\mathcal{X}^\infty$
$w(\nu)$	prior weight on $\nu \in \mathcal{M}$
$w(\nu \mid x_{<t})$	posterior weight on $\nu \in \mathcal{M}$
$\xi$	$\xi(x_{<t}) = \sum_{\nu \in \mathcal{M}} w(\nu) \nu(x_{<t})$
$\rho_n$	$\rho_n(x_{<t}) = \max_{\mathcal{M}' \subset \mathcal{M}:  \mathcal{M}' =i} \sum_{\nu \in \mathcal{M}'} w(\nu) \nu(x_{<t})$
$\rho_n^{\text{norm}}$	like $\rho_n$ , but normalized to be a measure $\rho_n^{\text{norm}}(x \mid x_{<t}) = \rho_n(x \mid x_{<t}) / \sum_{x' \in \mathcal{X}} \rho_n(x' \mid x_{<t})$
$\mathcal{M}_n^{x_{<t}}$	$\text{argmax}_{\mathcal{M}' \subset \mathcal{M}:  \mathcal{M}' =i} \sum_{\nu \in \mathcal{M}'} w(\nu) \nu(x_{<t})$
$\rho_n^{\text{stat}}$	a mixture over the top $i$ models, sorted by posterior weight $\rho_n^{\text{stat}}(x \mid x_{<t}) = \sum_{\nu \in \mathcal{M}_n^{x_{<t}}} w(\nu) \nu(x_{<t}x) / \sum_{\nu \in \mathcal{M}_n^{x_{<t}}} w(\nu) \nu(x_{<t})$
$\phi_n^{x_{<t}}$	$w(\nu_n^{x_{<t}} \mid x_{<t}) / w(\mathcal{M}_n^{x_{<t}} \mid x_{<t})$

**Appendix B. Omitted Proofs**

**Lemma 9** Recalling  $\nu(\cdot | x_{<t})$  is a measure over  $\mathcal{X}$ ,

$$\mathbb{E}_\mu \sum_{t=0}^{\infty} \text{KL}(\mu(\cdot | x_{<t}) \| \rho_n^{\text{norm}}(\cdot | x_{<t})) \leq w(\mu)^{-1} + \log w(\mu)^{-1}$$

**Proof** The KL divergence is non-negative, so we bound an arbitrary finite sum.

$$\begin{aligned}
 & \mathbb{E}_\mu \sum_{t=0}^{N-1} \text{KL}(\mu(\cdot | x_{<t}) \| \rho_n^{\text{norm}}(\cdot | x_{<t})) \\
 &= \sum_{t=0}^{N-1} \mathbb{E}_\mu \sum_{x_t \in \mathcal{X}} \mu(x_t | x_{<t}) \log \frac{\mu(x_t | x_{<t})}{\rho_n^{\text{norm}}(x_t | x_{<t})} \\
 &\stackrel{(a)}{=} \sum_{t=0}^{N-1} \sum_{x_{<t} \in \mathcal{X}^t} \mu(x_{<t}) \sum_{x_t \in \mathcal{X}} \mu(x_t | x_{<t}) \left[ \log \frac{\mu(x_t | x_{<t})}{\rho_n(x_t | x_{<t})} + \log \frac{\sum_{x' \in \mathcal{X}} \rho_n(x_{<t} x')}{\rho_n(x_{<t})} \right] \\
 &\stackrel{(b)}{\leq} w(\mu)^{-1} + \sum_{t=0}^{N-1} \sum_{x_{<t} \in \mathcal{X}^t} \mu(x_{<t}) \sum_{x_t \in \mathcal{X}} \mu(x_t | x_{<t}) \log \frac{\mu(x_t | x_{<t})}{\rho_n(x_t | x_{<t})} \\
 &= w(\mu)^{-1} + \sum_{t=0}^{N-1} \sum_{x_{<t} \in \mathcal{X}^t} \mu(x_{<t}) \sum_{x_t \in \mathcal{X}} \mu(x_t | x_{<t}) \left[ \log \frac{\mu(x_{<t} x_t)}{\rho_n(x_{<t} x_t)} - \log \frac{\mu(x_{<t})}{\rho_n(x_{<t})} \right] \\
 &= w(\mu)^{-1} + \sum_{t=0}^{N-1} \sum_{x_{<t} \in \mathcal{X}^t} \mu(x_{<t}) \left[ \sum_{x_t \in \mathcal{X}} \mu(x_t | x_{<t}) \log \frac{\mu(x_{<t} x_t)}{\rho_n(x_{<t} x_t)} - \log \frac{\mu(x_{<t})}{\rho_n(x_{<t})} \right] \\
 &= w(\mu)^{-1} + \sum_{t=0}^{N-1} \left[ \sum_{x_{\leq t} \in \mathcal{X}^{t+1}} \mu(x_{\leq t}) \log \frac{\mu(x_{\leq t})}{\rho_n(x_{\leq t})} - \sum_{x_{<t} \in \mathcal{X}^t} \mu(x_{<t}) \log \frac{\mu(x_{<t})}{\rho_n(x_{<t})} \right] \\
 &\stackrel{(c)}{=} w(\mu)^{-1} + \sum_{x_{<N} \in \mathcal{X}^N} \mu(x_{<N}) \log \frac{\mu(x_{<N})}{\rho_n(x_{<N})} - \mu(\epsilon) \log \frac{\mu(\epsilon)}{\rho_n(\epsilon)} \\
 &\stackrel{(d)}{\leq} w(\mu)^{-1} + \sum_{x_{<N} \in \mathcal{X}^N} \mu(x_{<N}) \log w(\mu)^{-1} = w(\mu)^{-1} + \log w(\mu)^{-1} \tag{27}
 \end{aligned}$$

where (a) follows from the definition of  $\rho_n^{\text{norm}}$  in Equation 9, (b) follows from Lemma 7 and the fact that  $\log x \leq x - 1$ , (c) cancels like terms, and (d) follows from Inequality 12.  $\blacksquare$

**Theorem 10 (SMAP Convergence)**

$$\mathbb{E}_\mu \sum_{t=0}^{\infty} \sum_{x \in \mathcal{X}} (\rho_n^{\text{stat}}(x | x_{<t}) - \mu(x | x_{<t}))^2 \leq 6w(\mu)^{-1} + 3$$

**Proof** We abbreviate  $w(\mu)^{-1}$  as  $c$ . Let  $[N] := (0, \dots, N-1)$ . We define an  $N|\mathcal{X}|$ -dimensional random vector depending on the infinite sequence  $x_{<\infty}$ :

$$\overrightarrow{\nu_1; \nu_2^N} := (\nu_1(x | x_{<t}) - \nu_2(x | x_{<t}))_{t \in [N], x \in \mathcal{X}} \quad (28)$$

In this notation, we aim to show  $\mathbb{E}_\mu \|\overrightarrow{\rho_n^{\text{stat.}}; \mu^N}\|_2^2 \leq 6c + 3$ . Lemma 8 (i) and (ii) become

$$\mathbb{E}_\mu \|\overrightarrow{\rho_n; \rho_n^{\text{stat.}}}\|_1 \leq c \quad (29)$$

$$\mathbb{E}_\mu \|\overrightarrow{\rho_n; \rho_n^{\text{norm.}}}\|_1 \leq c \quad (30)$$

Therefore,

$$\mathbb{E}_\mu \|\overrightarrow{\rho_n^{\text{stat.}}; \rho_n^{\text{norm.}}}\|_1 \leq 2c \quad (31)$$

Since each element in this vector is in  $[-1, 1]$ , squaring them makes the magnitude no larger, so

$$\mathbb{E}_\mu \|\overrightarrow{\rho_n^{\text{stat.}}; \rho_n^{\text{norm.}}}\|_2^2 \leq 2c \quad (32)$$

The KL divergence is larger than the sum of the squares of the probability differences (proven, for example, in (Hutter, 2005, §3.9.2)), so Lemma 9 implies

$$\mathbb{E}_\mu \|\overrightarrow{\rho_n^{\text{norm.}}; \mu^N}\|_2^2 \leq c + \log c \quad (33)$$

By the triangle inequality,

$$\|\overrightarrow{\rho_n^{\text{stat.}}; \mu^N}\|_2 \leq \|\overrightarrow{\rho_n^{\text{stat.}}; \rho_n^{\text{norm.}}}\|_2 + \|\overrightarrow{\rho_n^{\text{norm.}}; \mu^N}\|_2 \quad (34)$$

so

$$\|\overrightarrow{\rho_n^{\text{stat.}}; \mu^N}\|_2^2 \leq \|\overrightarrow{\rho_n^{\text{stat.}}; \rho_n^{\text{norm.}}}\|_2^2 + \|\overrightarrow{\rho_n^{\text{norm.}}; \mu^N}\|_2^2 + 2\|\overrightarrow{\rho_n^{\text{stat.}}; \rho_n^{\text{norm.}}}\|_2 \|\overrightarrow{\rho_n^{\text{norm.}}; \mu^N}\|_2 \quad (35)$$

and because  $\mathbb{E}[XY] \leq \sqrt{\mathbb{E}[X^2] \mathbb{E}[Y^2]}$  (the Cauchy–Schwarz Inequality),

$$\mathbb{E}_\mu \|\overrightarrow{\rho_n^{\text{stat.}}; \mu^N}\|_2^2 \leq 2c + (c + \log c) + 2\sqrt{2c(c + \log c)} < 6c + 3 \quad (36)$$

■

We name the measure with the  $i^{\text{th}}$  largest posterior weight

$$\nu_n^{x_{<t}} := \mathcal{M}_n^{x_{<t}} \setminus \mathcal{M}_{i-1}^{x_{<t}} \quad (37)$$

with the posterior weight formally defined  $w(\nu | x_{<t}) := \frac{w(\nu)\nu(x_{<t})}{\xi(x_{<t})}$ , and  $w(\mathcal{M}' | x_{<t}) := \sum_{\nu \in \mathcal{M}'} w(\nu | x_{<t})$ . Now, we let

$$\phi_n^{x_{<t}} := \frac{w(\nu_n^{x_{<t}} | x_{<t})}{w(\mathcal{M}_n^{x_{<t}} | x_{<t})} \quad (38)$$

**Theorem 6 (Top Model Convergence)**

$$\begin{aligned}
 (i) \quad & \mathbb{E}_\mu \sum_{t=0}^{\infty} \sum_{x \in \mathcal{X}} \left[ \mu(x | x_{<t}) - \min_{n: \phi_n^{x_{<t}} > \alpha} \nu_n^{x_{<t}}(x | x_{<t}) \right]^2 \leq \alpha^{-3} (24w(\mu)^{-1} + 12) \\
 (ii) \quad & \mathbb{E}_\mu \sum_{t=0}^{\infty} \left[ 1 - \sum_{x \in \mathcal{X}} \min_{n: \phi_n^{x_{<t}} > \alpha} \nu_n^{x_{<t}}(x | x_{<t}) \right]^2 \leq |\mathcal{X}| \alpha^{-3} (24w(\mu)^{-1} + 12)
 \end{aligned}$$

**Proof**  $\rho_n^{\text{stat}}(x | x_{<t})$  is a weighted average of  $\nu_j^{x_{<t}}(x | x_{<t})$  for  $j \leq i$ :

$$\begin{aligned}
 \rho_n^{\text{stat}}(x | x_{<t}) &= \frac{\sum_{\nu \in \mathcal{M}_n^{x_{<t}}} w(\nu) \nu(x_{<t} | x)}{\sum_{\nu \in \mathcal{M}_n^{x_{<t}}} w(\nu) \nu(x_{<t})} \\
 &= \frac{\sum_{\nu \in \mathcal{M}_n^{x_{<t}}} w(\nu) \nu(x_{<t}) \nu(x | x_{<t})}{\sum_{\nu \in \mathcal{M}_n^{x_{<t}}} w(\nu) \nu(x_{<t})} \\
 &= \frac{\sum_{\nu \in \mathcal{M}_n^{x_{<t}}} w(\nu | x_{<t}) \xi(x_{<t}) \nu(x | x_{<t})}{\sum_{\nu \in \mathcal{M}_n^{x_{<t}}} w(\nu | x_{<t}) \xi(x_{<t})} \\
 &= \sum_{\nu \in \mathcal{M}_n^{x_{<t}}} \frac{w(\nu | x_{<t})}{w(\mathcal{M}_n^{x_{<t}} | x_{<t})} \nu(x | x_{<t}) \\
 &= \sum_{j=1}^i \frac{w(\nu_j^{x_{<t}} | x_{<t})}{w(\mathcal{M}_n^{x_{<t}} | x_{<t})} \nu_j^{x_{<t}}(x | x_{<t}) \tag{39}
 \end{aligned}$$

Trivially,

$$\nu_1^{x_{<t}}(x | x_{<t}) = \rho_1^{\text{stat}}(x | x_{<t}) \tag{40}$$

but for  $i > 1$ , we would like to express  $\nu_n^{x_{<t}}$  in terms of  $\rho_n^{\text{stat}}$  and  $\rho_{n-1}^{\text{stat}}$ :

$$\rho_n^{\text{stat}}(x | x_{<t}) = \frac{w(\mathcal{M}_{i-1}^{x_{<t}} | x_{<t})}{w(\mathcal{M}_n^{x_{<t}} | x_{<t})} \rho_{n-1}^{\text{stat}}(x | x_{<t}) + \frac{w(\nu_n^{x_{<t}} | x_{<t})}{w(\mathcal{M}_n^{x_{<t}} | x_{<t})} \nu_n^{x_{<t}}(x | x_{<t}) \tag{41}$$

Thus,

$$\nu_n^{x_{<t}}(x | x_{<t}) = \frac{w(\mathcal{M}_n^{x_{<t}} | x_{<t})}{w(\nu_n^{x_{<t}} | x_{<t})} \rho_n^{\text{stat}}(x | x_{<t}) - \frac{w(\mathcal{M}_{i-1}^{x_{<t}} | x_{<t})}{w(\nu_n^{x_{<t}} | x_{<t})} \rho_{n-1}^{\text{stat}}(x | x_{<t}) \tag{42}$$

Since  $\frac{w(\mathcal{M}_n^{x_{<t}} | x_{<t})}{w(\nu_n^{x_{<t}} | x_{<t})} - \frac{w(\mathcal{M}_{i-1}^{x_{<t}} | x_{<t})}{w(\nu_n^{x_{<t}} | x_{<t})} = 1$ , for  $i > 1$ ,

$$\begin{aligned}
 \nu_n^{x_{<t}}(x | x_{<t}) - \mu(x | x_{<t}) &= \frac{w(\mathcal{M}_n^{x_{<t}} | x_{<t})}{w(\nu_n^{x_{<t}} | x_{<t})} [\rho_n^{\text{stat}}(x | x_{<t}) - \mu(x | x_{<t})] - \\
 &\quad \frac{w(\mathcal{M}_{i-1}^{x_{<t}} | x_{<t})}{w(\nu_n^{x_{<t}} | x_{<t})} [\rho_{n-1}^{\text{stat}}(x | x_{<t}) - \mu(x | x_{<t})] \tag{43}
 \end{aligned}$$

Recall

$$\phi_n^{x_{<t}} := \frac{w(\nu_n^{x_{<t}} | x_{<t})}{w(\mathcal{M}_n^{x_{<t}} | x_{<t})}$$

Since  $w(\mathcal{M}_{i-1}^{x_{<t}} | x_{<t}) \leq w(\mathcal{M}_n^{x_{<t}} | x_{<t})$ , we have

$$(\phi_n^{x_{<t}})^2 [\nu_n^{x_{<t}}(x | x_{<t}) - \mu(x | x_{<t})]^2 \leq 2 [\rho_n^{\text{stat}}(x | x_{<t}) - \mu(x | x_{<t})]^2 + 2 [\rho_{n-1}^{\text{stat}}(x | x_{<t}) - \mu(x | x_{<t})]^2 \quad (44)$$

Now we consider all measures  $\nu_n^{x_{<t}}$  for which  $\phi_n^{x_{<t}} > \alpha$ .

$$\mathbb{E}_\mu \sum_{t=0}^{N-1} \sum_{i: \phi_n^{x_{<t}} > \alpha} \sum_{x \in \mathcal{X}} [\nu_n^{x_{<t}}(x | x_{<t}) - \mu(x | x_{<t})]^2 \leq 2\alpha^{-2} \mathbb{E}_\mu \sum_{t=0}^{N-1} \sum_{i: \phi_n^{x_{<t}} > \alpha} \sum_{x \in \mathcal{X}} [\rho_n^{\text{stat}}(x | x_{<t}) - \mu(x | x_{<t})]^2 + [\rho_{n-1}^{\text{stat}}(x | x_{<t}) - \mu(x | x_{<t})]^2 \quad (45)$$

Now we note that  $\{n : \phi_n^{x_{<t}} > \alpha\} \subset \{n : n < \alpha^{-1}\}$ , since  $w(\nu_n^{x_{<t}} | x_{<t}) \leq w(\nu_j^{x_{<t}} | x_{<t})$  for  $i > j$ . Thus,

$$\begin{aligned} & \mathbb{E}_\mu \sum_{t=0}^{N-1} \sum_{n: \phi_n^{x_{<t}} > \alpha} \sum_{x \in \mathcal{X}} [\nu_n^{x_{<t}}(x | x_{<t}) - \mu(x | x_{<t})]^2 \\ & \leq 2\alpha^{-2} \mathbb{E}_\mu \sum_{t=0}^{N-1} \sum_{i: i < \alpha^{-1}} \sum_{x \in \mathcal{X}} [\rho_n^{\text{stat}}(x | x_{<t}) - \mu(x | x_{<t})]^2 + [\rho_{n-1}^{\text{stat}}(x | x_{<t}) - \mu(x | x_{<t})]^2 \\ & = 2\alpha^{-2} \sum_{i: i < \alpha^{-1}} \mathbb{E}_\mu \sum_{t=0}^{N-1} \sum_{x \in \mathcal{X}} [\rho_n^{\text{stat}}(x | x_{<t}) - \mu(x | x_{<t})]^2 + [\rho_{n-1}^{\text{stat}}(x | x_{<t}) - \mu(x | x_{<t})]^2 \\ & \leq 2\alpha^{-2} \sum_{i: i < \alpha^{-1}} 2(6w(\mu)^{-1} + 3) \leq \alpha^{-3}(24w(\mu)^{-1} + 12) \end{aligned} \quad (46)$$

Considering only a subset of these conditional-probability-errors,

$$\begin{aligned} & \mathbb{E}_\mu \sum_{t=0}^{N-1} \sum_{x \in \mathcal{X}} \left[ \mu(x | x_{<t}) - \min_{n: \phi_n^{x_{<t}} > \alpha} \nu_n^{x_{<t}}(x | x_{<t}) \right]^2 \leq \\ & \mathbb{E}_\mu \sum_{t=0}^{N-1} \sum_{n: \phi_n^{x_{<t}} > \alpha} \sum_{x \in \mathcal{X}} [\nu_n^{x_{<t}}(x | x_{<t}) - \mu(x | x_{<t})]^2 \leq \alpha^{-3}(24w(\mu)^{-1} + 12) \end{aligned} \quad (47)$$

This completes the proof of (i). Finally, with  $\mathbb{U}$  being the uniform distribution,

$$\begin{aligned} & \mathbb{E}_\mu \sum_{t=0}^{N-1} \left[ 1 - \sum_{x \in \mathcal{X}} \min_{n: \phi_n^{x_{<t}} > \alpha} \nu_n^{x_{<t}}(x | x_{<t}) \right]^2 \\ & = \mathbb{E}_\mu \sum_{t=0}^{N-1} \left[ \sum_{x \in \mathcal{X}} \mu(x | x_{<t}) - \min_{n: \phi_n^{x_{<t}} > \alpha} \nu_n^{x_{<t}}(x | x_{<t}) \right]^2 \\ & = \mathbb{E}_\mu \sum_{t=0}^{N-1} \left[ |\mathcal{X}| \mathbb{E}_{x \sim \mathbb{U}(\mathcal{X})} \mu(x | x_{<t}) - \min_{n: \phi_n^{x_{<t}} > \alpha} \nu_n^{x_{<t}}(x | x_{<t}) \right]^2 \end{aligned}$$

$$\begin{aligned}
 &\stackrel{(a)}{\leq} |\mathcal{X}|^2 \mathbb{E}_\mu \sum_{t=0}^{N-1} \mathbb{E}_{x \sim \mathbb{U}(\mathcal{X})} \left[ \mu(x | x_{<t}) - \min_{n: \phi_n^{x_{<t}} > \alpha} \nu_n^{x_{<t}}(x | x_{<t}) \right]^2 \\
 &= |\mathcal{X}| \mathbb{E}_\mu \sum_{t=0}^{N-1} \sum_{x \in \mathcal{X}} \left[ \mu(x | x_{<t}) - \min_{n: \phi_n^{x_{<t}} > \alpha} \nu_n^{x_{<t}}(x | x_{<t}) \right]^2 \\
 &\stackrel{(b)}{\leq} |\mathcal{X}| \alpha^{-3} (24w(\mu)^{-1} + 12)
 \end{aligned} \tag{48}$$

where (a) follows from Jensen's Inequality, and (b) follows from Theorem 6 (i), which completes the proof of (ii).  $\blacksquare$

**Theorem 2 (Top Models Contain Truth)**  $\mathbb{P}_{\mu^\alpha}^i(\forall t : \pi^d \in \Pi_{h_{<t}}^\alpha) \geq 1 - \alpha w(\pi^d)^{-1}$

**Proof** Since,  $w(\pi^d | h_{<t}) > \alpha \implies \pi^d \in \Pi_{h_{<t}}^\alpha$ , we show  $\mathbb{P}_{\mu^\alpha}^i(\forall t : w(\pi^d | h_{<t}) > \alpha) \geq 1 - \alpha w(\pi^d)^{-1}$ . First we show that  $z_t = w(\pi^d | h_{<t})^{-1}$  is a non-negative  $\mathbb{P}_{\mu^\alpha}^i$ -supermartingale.

First, suppose  $q_{t+1} = 0$ . In this case,  $z_{t+1} = z_t$ , because the posterior weight is only updated when the demonstrator picks an action. Now suppose  $q_{t+1} = 1$ .

$$\begin{aligned}
 \mathbb{E}_{\mu^\alpha}^i[z_{t+1} | h_{<t}1] &\stackrel{(a)}{=} \sum_{a_t \in \mathcal{A}: \pi^d(a_t | h_{<t}) > 0} \pi^d(a_t | h_{<t}) w(\pi^d | h_{<t}1 a_t)^{-1} \\
 &\stackrel{(b)}{=} \sum_{a_t \in \mathcal{A}: \pi^d(a_t | h_{<t}) > 0} \pi^d(a_t | h_{<t}) \frac{\sum_{\pi \in \Pi} w(\pi | h_{<t}) \pi(a_t | h_{<t})}{w(\pi^d | h_{<t}) \pi^d(a_t | h_{<t})} \\
 &\stackrel{(c)}{\leq} \sum_{a_t \in \mathcal{A}} \frac{\sum_{\pi \in \Pi} w(\pi | h_{<t}) \pi(a_t | h_{<t})}{w(\pi^d | h_{<t})} \\
 &= z_t \sum_{\pi \in \Pi} w(\pi | h_{<t}) \sum_{a_t \in \mathcal{A}} \pi(a_t | h_{<t}) = z_t
 \end{aligned}$$

where (a) follows because  $a_t \sim \pi^d$  when  $q_t = 1$ , (b) follows from Bayes' rule—the formula for posterior updating, and (c) follows from cancelling, and adding non-negative terms to the sum.

Since  $w(\pi^d | h_{<t})^{-1}$  is a non-negative supermartingale, by the supermartingale convergence theorem (Durrett, 2010, Thm. 5.4.2),

$$\mathbb{P}_{\mu^\alpha}^i(\exists t : w(\pi^d | h_{<t})^{-1} \geq \alpha^{-1}) \leq \alpha w(\pi^d)^{-1} \tag{49}$$

so

$$\mathbb{P}_{\mu^\alpha}^i(\forall t : w(\pi^d | h_{<t}) > \alpha) \geq 1 - \alpha w(\pi^d)^{-1} \tag{50}$$

which implies

$$\mathbb{P}_{\mu^\alpha}^i(\forall t : \pi^d \in \Pi_{h_{<t}}^\alpha) \geq 1 - \alpha w(\pi^d)^{-1} \tag{51}$$

$\blacksquare$

**Theorem 3 (Predictive Convergence)** For  $\alpha < w(\pi^d)$ ,

$$\mathbb{E}_{\mu^\alpha}^{\pi^\alpha} \left[ \sum_{t=0}^{\infty} \left( \sum_{a \in \mathcal{A}} |\pi_\alpha^i(0, a | h_{<t}) - \pi^d(1, a | h_{<t})| \right)^3 \middle| E \right] \leq \frac{|\mathcal{A}| \alpha^{-3} (24w(\pi^d)^{-1} + 12)}{1 - \alpha w(\pi^d)^{-1}}$$

**Proof** Recall  $\pi_\alpha^i(0, a | h_{<t}) = \min_{\pi \in \Pi_{h_{<t}}^\alpha} \pi(1, a | h_{<t})$ , so if  $\pi^d \in \Pi_{h_{<t}}^\alpha$ , then  $\pi_\alpha^i(0, a | h_{<t}) \leq \pi^d(1, a | h_{<t})$ . Thus, in that case,

$$\begin{aligned} \sum_{a \in \mathcal{A}} |\pi_\alpha^i(0, a | h_{<t}) - \pi^d(1, a | h_{<t})| &= \sum_{a \in \mathcal{A}} \pi^d(1, a | h_{<t}) - \pi_\alpha^i(0, a | h_{<t}) \leq \\ &1 - \sum_{a \in \mathcal{A}} \pi_\alpha^i(0, a | h_{<t}) = \theta_q(h_{<t}) \end{aligned} \quad (52)$$

The rest follows easily:

$$\begin{aligned} &\mathbb{E}_{\mu^\alpha}^{\pi^\alpha} \left[ \sum_{t=0}^{\infty} \left( \sum_{a \in \mathcal{A}} |\pi_\alpha^i(0, a | h_{<t}) - \pi^d(1, a | h_{<t})| \right)^3 \middle| \forall t : \pi^d \in \Pi_{h_{<t}}^\alpha \right] \\ &\leq \mathbb{E}_{\mu^\alpha}^{\pi^\alpha} \left[ \sum_{t=0}^{\infty} \theta_q(h_{<t})^3 \middle| \forall t : \pi^d \in \Pi_{h_{<t}}^\alpha \right] \\ &\stackrel{(a)}{\leq} \mathbb{E}_{\mu^\alpha}^{\pi^\alpha} \left[ \sum_{t=0}^{\infty} \theta_q(h_{<t})^3 \right] / \mathbb{P}_{\mu^\alpha}^{\pi^\alpha}(\forall t : \pi^d \in \Pi_{h_{<t}}^\alpha) \\ &\stackrel{(b)}{\leq} \frac{|\mathcal{A}| \alpha^{-3} (24w(\pi^d)^{-1} + 12)}{1 - \alpha w(\pi^d)^{-1}} \end{aligned} \quad (53)$$

where (a) follows because  $\theta_q$  is non-negative, and (b) follows from Equation 51 and Theorem 1 (as long as  $\alpha < w(\pi^d)$ ).  $\blacksquare$

**Lemma 12** For  $a \in \mathcal{A}$ , let  $0 \leq i_a \leq d_a$ , and let  $\sum_{a \in \mathcal{A}} d_a = 1$ . Let  $\theta_q = 1 - \sum_{a \in \mathcal{A}} i_a$ . Then,

$$\Delta := \sum_{a \in \mathcal{A}} (i_a + \theta_q d_a) \log \frac{i_a + \theta_q d_a}{d_a} \leq \theta_q$$

**Proof**

$$\begin{aligned} \sum_{a \in \mathcal{A}} (i_a + \theta_q d_a) \log \frac{i_a + \theta_q d_a}{d_a} &= \sum_{a \in \mathcal{A}} (i_a + \theta_q d_a) \log \left( \frac{i_a}{d_a} + \theta_q \right) \leq \\ &\left( \sum_{a \in \mathcal{A}} i_a + \theta_q \sum_{a \in \mathcal{A}} d_a \right) \log(1 + \theta_q) = (1 - \theta_q + \theta_q) \log(1 + \theta_q) \leq \theta_q \end{aligned} \quad (54)$$

$\blacksquare$

For the remaining proofs, we sometimes consider the restriction of probability measures over  $\mathcal{H}^\infty$  to  $(\mathcal{A} \times \mathcal{O})^\infty$ ; that is, we marginalize over the query record. For a history  $h_{<t} = q_0 a_0 o_0 \dots q_{t-1} a_{t-1} o_{t-1}$ , let  $h_{<t}^\setminus$  denote  $a_0 o_0 \dots a_{t-1} o_{t-1}$ . We define the  $t$ -step KL divergence as follows:

$$\text{KL}_t(P \parallel Q) := \sum_{h_{<t}^\setminus \in (\mathcal{A} \times \mathcal{O})^t} P(h_{<t}^\setminus) \log \frac{P(h_{<t}^\setminus)}{Q(h_{<t}^\setminus)} \quad (55)$$

**Theorem 4 (KL Bound)** *Suppose that  $\mu$  and  $\pi^d$  are fair, and  $\alpha < w(\pi^d)$ . Letting the two probability measures below be restricted to  $(\mathcal{A} \times \mathcal{O})^t$  (that is, marginalizing over the query record, and considering only the first  $t$  timesteps),*

$$\text{KL} \left( \mathbb{P}_{\mu^\alpha}^{\pi^\alpha}(\cdot \mid E) \parallel \mathbb{P}_{\mu}^{\pi^d}(\cdot \mid E) \right) \leq \frac{\alpha^{-1} |\mathcal{A}|^{1/3} (24w(\pi^d)^{-1} + 12)^{1/3}}{(1 - \alpha/w(\pi^d))^2} t^{2/3} - \log(1 - \alpha/w(\pi^d))$$

**Proof** We begin by restricting attention to a particular timestep  $t$ . Recall  $\pi_\alpha^i(0, a \mid h_{<t}) = \min_{\pi' \in \Pi_{h_{<t}}^\alpha} \pi'(1, a \mid h_{<t})$ . We abbreviate this quantity  $i_a$ . We also let  $d_a$  denote  $\pi^d(1, a \mid h_{<t})$ . Note that when  $\pi^d \in \Pi_{h_{<t}}^\alpha$ ,

$$i_a \leq d_a \quad (56)$$

Recall that the query probability  $\theta_q = 1 - \sum_{a \in \mathcal{A}} i_a$ , and the marginalized probability  $\pi_\alpha^i(a \mid h_{<t}) = i_a + \theta_q d_a$ . Assuming  $h_{<k}$  satisfies  $E$ , let

$$\Delta_k := \text{KL}_1 \left( \pi_\alpha^i(\cdot \mid h_{<k}) \parallel \pi^d(\cdot \mid h_{<k}) \right) = \sum_{a \in \mathcal{A}} (i_a + \theta_q d_a) \log \frac{i_a + \theta_q d_a}{d_a} \quad (57)$$

By Lemma 12,  $\Delta_k \leq \theta_q$ .

Now, we write the  $t$ -step KL divergence  $\text{KL}_t$  as a sum of the expectation of 1-step KL divergences. We'll abbreviate a measure  $\mathbb{P}(\cdot \mid E)$  as  ${}^E\mathbb{P}$ .

$$\begin{aligned} \text{KL}_t \left( {}^E\mathbb{P}_{\mu^\alpha}^{\pi^\alpha} \parallel {}^E\mathbb{P}_{\mu}^{\pi^d} \right) &= \mathbb{E}_{h_{<t} \sim {}^E\mathbb{P}_{\mu^\alpha}^{\pi^\alpha}} \log \frac{{}^E\mathbb{P}_{\mu^\alpha}^{\pi^\alpha}(h_{<t}^\setminus)}{{}^E\mathbb{P}_{\mu}^{\pi^d}(h_{<t}^\setminus)} \\ &\stackrel{(a)}{\leq} \mathbb{E}_{h_{<t} \sim {}^E\mathbb{P}_{\mu^\alpha}^{\pi^\alpha}} \log \frac{\mathbb{P}_{\mu^\alpha}^{\pi^\alpha}(h_{<t}^\setminus) / \mathbb{P}_{\mu^\alpha}^{\pi^\alpha}(E)}{\mathbb{P}_{\mu}^{\pi^d}(h_{<t}^\setminus) / \mathbb{P}_{\mu}^{\pi^d}(E)} \\ &= \mathbb{E}_{h_{<t} \sim {}^E\mathbb{P}_{\mu^\alpha}^{\pi^\alpha}} \log \frac{\mathbb{P}_{\mu^\alpha}^{\pi^\alpha}(h_{<t}^\setminus)}{\mathbb{P}_{\mu}^{\pi^d}(h_{<t}^\setminus)} + \log \frac{\mathbb{P}_{\mu}^{\pi^d}(E)}{\mathbb{P}_{\mu^\alpha}^{\pi^\alpha}(E)} \\ &\leq \mathbb{E}_{h_{<t} \sim {}^E\mathbb{P}_{\mu^\alpha}^{\pi^\alpha}} \log \frac{\mathbb{P}_{\mu^\alpha}^{\pi^\alpha}(h_{<t}^\setminus)}{\mathbb{P}_{\mu}^{\pi^d}(h_{<t}^\setminus)} - \log \mathbb{P}_{\mu^\alpha}^{\pi^\alpha}(E) \\ &=: \mathbb{E}_{h_{<t} \sim {}^E\mathbb{P}_{\mu^\alpha}^{\pi^\alpha}} \log \frac{\mathbb{P}_{\mu^\alpha}^{\pi^\alpha}(h_{<t}^\setminus)}{\mathbb{P}_{\mu}^{\pi^d}(h_{<t}^\setminus)} + C_\alpha \\ &\stackrel{(b)}{=} C_\alpha + \mathbb{E}_{h_{<t} \sim {}^E\mathbb{P}_{\mu^\alpha}^{\pi^\alpha}} \sum_{k=0}^{t-1} \log \frac{\mathbb{P}_{\mu^\alpha}^{\pi^\alpha}(h_k^\setminus \mid h_{<k})}{\mathbb{P}_{\mu}^{\pi^d}(h_k^\setminus \mid h_{<k})} \end{aligned}$$

$$\begin{aligned}
 &= C_\alpha + \sum_{k=0}^{t-1} \mathbb{E}_{h_{<k} \sim E P_\mu^{\pi_\alpha^i}} \mathbb{E}_{h_k \sim E P_\mu^{\pi_\alpha^i}(\cdot | h_{<k})} \log \frac{P_\mu^{\pi_\alpha^i}(h_k^\setminus | h_{<k})}{P_\mu^{\pi^d}(h_k^\setminus | h_{<k})} \\
 &= C_\alpha + \sum_{k=0}^{t-1} \mathbb{E}_{h_{<k} \sim E P_\mu^{\pi_\alpha^i}} \sum_{h_k^\setminus \in \mathcal{A} \times \mathcal{O}} E P_\mu^{\pi_\alpha^i}(h_k^\setminus | h_{<k}) \log \frac{P_\mu^{\pi_\alpha^i}(h_k^\setminus | h_{<k})}{P_\mu^{\pi^d}(h_k^\setminus | h_{<k})} \\
 &\leq C_\alpha + \sum_{k=0}^{t-1} \mathbb{E}_{h_{<k} \sim E P_\mu^{\pi_\alpha^i}} \sum_{h_k^\setminus \in \mathcal{A} \times \mathcal{O}} \frac{P_\mu^{\pi_\alpha^i}(h_k^\setminus | h_{<k})}{P_\mu^{\pi_\alpha^i}(E)} \log \frac{P_\mu^{\pi_\alpha^i}(h_k^\setminus | h_{<k})}{P_\mu^{\pi^d}(h_k^\setminus | h_{<k})} \\
 &= C_\alpha + \sum_{k=0}^{t-1} \mathbb{E}_{h_{<k} \sim E P_\mu^{\pi_\alpha^i}} \frac{1}{P_\mu^{\pi_\alpha^i}(E)} \text{KL}_1 \left( P_\mu^{\pi_\alpha^i}(\cdot | h_{<k}) \parallel P_\mu^{\pi^d}(\cdot | h_{<k}) \right) \\
 &= C_\alpha + \frac{1}{P_\mu^{\pi_\alpha^i}(E)} \sum_{k=0}^{t-1} \mathbb{E}_{h_{<k} \sim E P_\mu^{\pi_\alpha^i}} \text{KL}_1 \left( \pi_\alpha^i(\cdot | h_{<k}) \parallel \pi^d(\cdot | h_{<k}) \right) \\
 &\stackrel{(c)}{\leq} -\log P_\mu^{\pi_\alpha^i}(E) + \frac{1}{P_\mu^{\pi_\alpha^i}(E)} E \mathbb{E}_\mu^{\pi_\alpha^i} \sum_{k=0}^{t-1} \theta_q(h_{<k}) \\
 &\leq -\log P_\mu^{\pi_\alpha^i}(E) + \frac{1}{P_\mu^{\pi_\alpha^i}(E)^2} \mathbb{E}_\mu^{\pi_\alpha^i} \sum_{k=0}^{t-1} \theta_q(h_{<k}) \tag{58}
 \end{aligned}$$

where (a) follows from  $h_{<t}$  satisfying  $E$  with  $E P_\mu^{\pi_\alpha^i}$ -prob. 1, (b) follows because  $\mu$  and  $\pi^d$  are fair, and (c) follows from Equation 57 and Lemma 12.

Finally,

$$\begin{aligned}
 \mathbb{E}_\mu^{\pi_\alpha^i} \sum_{k=0}^{t-1} \theta_q(h_{<k}) &= t \mathbb{E}_{k \sim \mathcal{U}([t])} \mathbb{E}_\mu^{\pi_\alpha^i} \theta_q(h_{<k}) \\
 &= t \left( \left( \mathbb{E}_{k \sim \mathcal{U}([t])} \mathbb{E}_\mu^{\pi_\alpha^i} \theta_q(h_{<k}) \right)^3 \right)^{1/3} \\
 &\stackrel{(a)}{\leq} t \left( \mathbb{E}_{k \sim \mathcal{U}([t])} \mathbb{E}_\mu^{\pi_\alpha^i} \theta_q(h_{<k})^3 \right)^{1/3} \\
 &= t \left( \frac{1}{t} \sum_{k=0}^{t-1} \mathbb{E}_\mu^{\pi_\alpha^i} \theta_q(h_{<k})^3 \right)^{1/3} \\
 &\stackrel{(b)}{\leq} t^{2/3} |\mathcal{A}|^{1/3} \alpha^{-1} (24w(\pi^d)^{-1} + 12)^{1/3} \tag{59}
 \end{aligned}$$

where (a) follows from Jensen's Inequality, and (b) follows from Theorem 1. Combining this with Inequality 58, and recalling  $P_\mu^{\pi_\alpha^i}(E) \geq 1 - \alpha/w(\pi^d)$ , we have

$$\text{KL}_t \left( E P_\mu^{\pi_\alpha^i} \parallel E P_\mu^{\pi^d} \right) \leq \frac{\alpha^{-1} |\mathcal{A}|^{1/3} (24w(\pi^d)^{-1} + 12)^{1/3}}{(1 - \alpha/w(\pi^d))^2} t^{2/3} - \log(1 - \alpha/w(\pi^d)) \tag{60}$$

■