

Asymptotically Optimal Agents

Tor Lattimore¹ and Marcus Hutter^{1,2}

¹ Research School of Computer Science
Australian National University

² ETH Zürich

{tor.lattimore,marcus.hutter}@anu.edu.au

Abstract. Artificial general intelligence aims to create agents capable of learning to solve arbitrary interesting problems. We define two versions of asymptotic optimality and prove that no agent can satisfy the strong version while in some cases, depending on discounting, there does exist a non-computable weak asymptotically optimal agent.

Keywords: Rational agents, sequential decision theory, artificial general intelligence, reinforcement learning, asymptotic optimality, general discounting.

1 Introduction

The dream of artificial general intelligence is to create an agent that, starting with no knowledge of its environment, eventually learns to behave optimally. This means it should be able to learn chess just by playing, or Go, or how to drive a car or mow the lawn, or any task we could conceivably be interested in assigning it.

Before considering the existence of universally intelligent agents, we must be precise about what is meant by optimality. If the environment and goal are known, then subject to computation issues, the optimal policy is easy to construct using an expectimax search from sequential decision theory [13]. However, if the true environment is unknown then the agent will necessarily spend some time exploring, and so cannot immediately play according to the optimal policy. Given a class of environments, we suggest two definitions of asymptotic optimality for an agent.

1. An agent is strongly asymptotically optimal if for every environment in the class it plays optimally in the limit.
2. It is weakly asymptotic optimal if for every environment in the class it plays optimally *on average* in the limit.

The key difference is that a strong asymptotically optimal agent must eventually stop exploring, while a weak asymptotically optimal agent may explore forever, but with decreasing frequency.

In this paper we consider the (non-)existence of weak/strong asymptotically optimal agents in the class of all deterministic computable environments. The restriction to deterministic is for the sake of simplicity and because the results

for this case are already sufficiently non-trivial to be interesting. The restriction to computable is more philosophical. The Church-Turing thesis is the unprovable hypothesis that anything that can intuitively be computed can also be computed by a Turing machine. Applying this to physics leads to the strong Church-Turing thesis that the universe is computable (possibly stochastically computable, i.e. computable when given access to an oracle of random noise). Having made these assumptions, the largest interesting class then becomes the class of computable (possibly stochastic) environments.

In [7], Hutter conjectured that his universal Bayesian agent, AIXI, was weakly asymptotically optimal in the class of all computable stochastic environments. Unfortunately this was recently shown to be false in [14], where it is proven that no Bayesian agent (with a static prior) can be weakly asymptotically optimal in this class.¹ The key idea behind Orseau’s proof was to show that AIXI eventually stops exploring. This is somewhat surprising because it is normally assumed that Bayesian agents solve the exploration/exploitation dilemma in a principled way. This result is a bit reminiscent of Bayesian (passive induction) inconsistency results [3, 4], although the details of the failure are very different.

We extend the work of [14], where only Bayesian agents are considered, to show that non-computable weak asymptotically optimal agents do exist in the class of deterministic computable environments for some discount functions (including geometric), but not for others. We also show that no asymptotically optimal agent can be computable, and that for all “reasonable” discount functions there does not exist a strong asymptotically optimal agent.

The weak asymptotically optimal agent we construct is similar to AIXI, but with an exploration component similar to ϵ -learning for finite state Markov decision processes or the UCB algorithm for bandits. The key is to explore sufficiently often and deeply to ensure that the environment used for the model is an adequate approximation of the true environment. At the same time, the agent must explore infrequently enough that it actually exploits its knowledge. Whether or not it is possible to get this balance right depends, somewhat surprisingly, on how forward looking the agent is (determined by the discount function). That it is sometimes not possible to explore enough to learn the true environment without damaging even a weak form of asymptotic optimality is surprising and unexpected.

Note that the exploration/exploitation problem is well-understood in the Bandit case [1, 2] and for (finite-state stationary) Markov decision processes [15]. In these restrictive settings, various satisfactory optimality criteria are available. In this work, we do not make any assumptions like Markov, stationary, ergodicity, or else besides computability of the environment. So far, no satisfactory optimality definition is available for this general case.

2 Notation and Definitions

We use similar notation to [7, 14] where the agent takes actions and the environment returns an observation/reward pair.


¹ Or even the class of computable deterministic environments.

Strings. A finite string a over alphabet \mathcal{A} is a finite sequence $a_1 a_2 a_3 \cdots a_{n-1} a_n$ with $a_i \in \mathcal{A}$. An infinite string ω over alphabet \mathcal{A} is an infinite sequence $\omega_1 \omega_2 \omega_3 \cdots$. \mathcal{A}^n , \mathcal{A}^* and \mathcal{A}^∞ are the sets of strings of length n , strings of finite length, and infinite strings respectively. Let x be a string (finite or infinite) then substrings are denoted $x_{s:t} := x_s x_{s+1} \cdots x_{t-1} x_t$ where $s, t \in \mathbb{N}$ and $s \leq t$. Strings may be concatenated. Let $x, y \in \mathcal{A}^*$ of length n and m respectively, and $\omega \in \mathcal{A}^\infty$. Then define $xy := x_1 x_2 \cdots x_{n-1} x_n y_1 y_2 \cdots y_{m-1} y_m$ and $x\omega := x_1 x_2 \cdots x_{n-1} x_n \omega_1 \omega_2 \omega_3 \cdots$. Some useful shorthands,

$$x_{<t} := x_{1:t-1} \qquad yx_{<t} := y_1 x_1 y_2 x_2 \cdots y_{t-1} x_{t-1}. \tag{1}$$

The second of these is ambiguous with concatenation, so wherever $yx_{<t}$ appears we assume the interleaving definition of (1) is intended. For example, it will be common to see $yx_{<t}y_t$, which represents the string $y_1 x_1 y_2 x_2 y_3 x_3 \cdots y_{t-1} x_{t-1} y_t$. For binary strings, we write $\#1(a)$ and $\#0(a)$ to mean the number of 0's and number of 1's in a respectively.

Environments and Optimality. Let \mathcal{Y} , \mathcal{O} and $\mathcal{R} \subset \mathbb{R}$ be action, observation and reward spaces respectively. Let $\mathcal{X} = \mathcal{O} \times \mathcal{R}$. An agent interacts with an environment as illustrated in the diagram on the right. First, the agent takes an action, upon which it receives a new observation/reward pair. The agent then takes



another action, receives another observation/reward pair, and so-on indefinitely. The goal of the agent is to maximise its discounted rewards over time. In this paper we consider only deterministic environments where the next observation/reward pair is determined by a function of the previous actions, observations and rewards.

Definition 1 (Deterministic Environment). A deterministic environment μ is a function $\mu : (\mathcal{Y} \times \mathcal{X})^* \times \mathcal{Y} \rightarrow \mathcal{X}$ where $\mu(yx_{<t}y_t) \in \mathcal{X}$ is the observation/reward pair given after action y_t is taken in history $yx_{<t}$. Wherever we write x_t we implicitly assume $x_t = (o_t, r_t)$ and refer to o_t and r_t without defining them. An environment μ is computable if there exists a Turing machine that computes it.

Note that since environments are deterministic the next observation need not depend on the previous observations (only actions). We choose to leave the dependence as the proofs become clearer when both the action and observation sequence is more visible.

Assumption 1. \mathcal{Y} and \mathcal{O} are finite, $\mathcal{R} = [0, 1]$.

Definition 2 (Policy). A policy π is a function from a history to an action $\pi : (\mathcal{Y} \times \mathcal{X})^* \rightarrow \mathcal{Y}$.

As expected, a policy π and environment μ can interact with each other to generate a play-out sequence of action/reward/observation tuples.

Definition 3 (Play-out Sequence). We define the play-out sequence $\mathbf{y}^{\mu,\pi} \in (\mathcal{Y} \times \mathcal{X})^\infty$ inductively by $y_k^{\mu,\pi} := \pi(\mathbf{y}_{<k}^{\mu,\pi})$ and $x_k^{\mu,\pi} := \mu(\mathbf{y}_{<k}^{\mu,\pi}, y_k^{\mu,\pi})$.

We need to define the value of a policy π in environment μ . To avoid the possibility of infinite rewards, we will use discounted values. While it is common to use only geometric discounting, we have two reasons to allow arbitrary time-consistent discount functions.

1. Geometric discounting has a constant effective horizon, but we feel agents should be allowed to use a discount function that leads to a growing horizon. This is seen in other agents, such as humans, who generally become less myopic as they grow older. See [5] for an overview of generic discounting.
2. The existence of asymptotically optimal agents depends critically on the effective horizon of the discount function.

Definition 4 (Discount Function). A regular discount function $\gamma \in \mathbb{R}^\infty$ is a vector satisfying $\gamma_k \geq 0$ and $0 < \sum_{t=k}^\infty \gamma_t < \infty$ for all $k \in \mathbb{N}$.

The first condition is natural for any definition of a discount function. The second condition is often cited as the purpose of a discount function (to prevent infinite utilities), but economists sometimes use non-summable discount functions, such as hyperbolic. The second condition also guarantees the agent cares about the infinite future, and is required to make asymptotic analysis interesting. We only consider discount functions satisfying all three conditions. In the following, let

$$\Gamma_t := \sum_{i=t}^\infty \gamma_i \qquad H_t(p) := \min_{h \in \mathbb{N}} \left\{ h : \frac{1}{\Gamma_t} \sum_{k=t}^{t+h} \gamma_k > p \right\}.$$

An infinite sequence of rewards starting at time t , $r_t, r_{t+1}, r_{t+2}, \dots$ is given a value of $\frac{1}{\Gamma_t} \sum_{i=t}^\infty \gamma_i r_i$. The term $\frac{1}{\Gamma_t}$ is a normalisation term to ensure that values scale in such a way that they can still be compared in the limit. A discount function is computable if there exists a Turing machine computing it. All well known discount functions, such as geometric, fixed horizon and hyperbolic are computable. Note that $H_t(p)$ exists for all $p \in [0, 1)$ and represents the effective horizon of the agent. After $H_t(p)$ time-steps into the future, starting at time t , the agent stands to gain/lose at most $1 - p$.

Definition 5 (Values and Optimal Policy). The value of policy π when starting from history $\mathbf{y}_{<t}^{\mu,\pi}$ in environment μ is $V_\mu^\pi(\mathbf{y}_{<t}^{\mu,\pi}) := \frac{1}{\Gamma_t} \sum_{k=t}^\infty \gamma_k r_k^{\mu,\pi}$. The optimal policy π_μ^* and its value V_μ^* are defined $\pi_\mu^*(\mathbf{y}_{<t}) := \arg \max_\pi V_\mu^\pi(\mathbf{y}_{<t})$ and $V_\mu^*(\mathbf{y}_{<t}) := V_\mu^{\pi_\mu^*}(\mathbf{y}_{<t})$.

Assumption 1 combined with Theorem 6 in [9] guarantees the existence of π_μ^* . Note that the normalisation term $\frac{1}{\Gamma_t}$ does not change the policy, but is used to ensure that values scale appropriately in the limit. For example, when discounting geometrically we have, $\gamma_t = \gamma^t$ for some $\gamma \in (0, 1)$ and so $\Gamma_t = \frac{\gamma^t}{1-\gamma}$ and $V_\mu^\pi(\mathbf{y}_{<t}^{\mu,\pi}) = (1 - \gamma) \sum_{k=t}^\infty \gamma^{k-t} r_k^{\mu,\pi}$.

Definition 6 (Asymptotic Optimality). Let $\mathcal{M} = \{\mu_0, \mu_1, \dots\}$ be a finite or countable set of environments and γ be a discount function. A policy π is a strong asymptotically optimal policy in (\mathcal{M}, γ) if

$$\lim_{n \rightarrow \infty} [V_{\mu}^*(y_{<n}^{\mu, \pi}) - V_{\mu}^{\pi}(y_{<n}^{\mu, \pi})] = 0, \text{ for all } \mu \in \mathcal{M}. \tag{2}$$

It is a weak asymptotically optimal policy if

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{t=1}^n [V_{\mu}^*(y_{<t}^{\mu, \pi}) - V_{\mu}^{\pi}(y_{<t}^{\mu, \pi})] = 0, \text{ for all } \mu \in \mathcal{M}. \tag{3}$$

Strong asymptotic optimality demands that the value of a *single* policy π converges to the value of the optimal policy π_{μ}^* for *all* μ in the class. This means that in the limit, a strong asymptotically optimal policy will obtain the maximum value possible in that environments.

Weak asymptotic optimality is similar, but only requires the *average* value of the policy π to converge to the average value of the optimal policy. This means that a weak asymptotically optimal policy can still make infinitely many bad mistakes, but must do so for only a fraction of the time that converges to zero. Strong asymptotic optimality implies weak asymptotic optimality.

While the definition of strong asymptotic optimality is rather natural, the definition of weak asymptotic optimality appears somewhat more arbitrary. The purpose of the average is to allow the agent to make a vanishing fraction of serious errors over its (infinite) life-time. We believe this is a necessary condition for an agent to learn the true environment. Of course, it would be possible to insist that the agent make only $o(\log n)$ serious errors rather than $o(n)$, which would make a stronger version of weak asymptotic optimality. Our choice is the weakest notion of optimality of the above form that still makes sense, which turns out to be already too strong for some discount rates.

Note that for both versions of optimality an agent would be considered optimal if it actively undertook a policy that led it to an extremely bad “hell” state from which it could not escape. Since the state cannot be escaped, its policy would then coincide with the optimal policy and so it would be considered optimal. Unfortunately, this problem seems to be an unavoidable consequence of learning algorithms in non-ergodic environments in general, including the currently fashionable PAC algorithms for arbitrary finite Markov decision processes.

3 Non-existence of Asymptotically Optimal Policies

We present the negative theorem in three parts. The first shows that, at least for computable discount functions, there does not exist a strong asymptotically optimal policy. The second shows that any weak asymptotically optimal policy must be incomputable while the third shows that there exist discount functions for which even incomputable weak asymptotically optimal policies do not exist.

Theorem 7. Let \mathcal{M} be the class of all deterministic computable environments and γ a computable discount function, then:

1. There does not exist a strong asymptotically optimal policy in (\mathcal{M}, γ) .
2. There does not exist a computable weak asymptotically optimal policy in (\mathcal{M}, γ) .
3. If $\gamma_k := \frac{1}{k(k+1)}$ then there does not exist a weak asymptotically optimal policy in (\mathcal{M}, γ) .

Part 1 of Theorem 7 says there is no strong asymptotically optimal policy in the class of all computable deterministic environments when the discount function is computable. It is likely there exist non-computable discount functions for which there are strong asymptotically optimal policies. Unfortunately the discount functions for which this is true are likely to be somewhat pathological and not realistic.

Given that strong asymptotic optimality is too strong, we should search for weak asymptotically optimal policies. Part 2 of Theorem 7 shows that any such policy is necessarily incomputable. This result features no real new ideas and relies on the fact that you can use a computable policy to hand-craft a computable environment in which it does very badly [10]. In general this approach fails for incomputable policies because the hand-crafted environment will then not be computable. Note that this does not rule out the existence of a stochastically computable weak asymptotically optimal policy.

It turns out that even weak asymptotic optimality is too strong for some discount functions. Part 3 of Theorem 7 gives an example discount function for which no such policy (computable or otherwise) exists. In the next section we introduce a weak asymptotically optimal policy for geometric (and may be extended to other) discounting. Note that $\gamma_k = \frac{1}{k(k+1)}$ is an example of a discount function where $H_t(p) = \Omega(t)$. It is also analytically easy to work with.

All negative results are proven by contradiction, and follow the same basic form.

1. Assume π is a computable/arbitrary weak/strong asymptotically optimal.
2. Therefore π is weak/strong asymptotically optimal in μ for some particular μ .
3. Construct ν , which is indistinguishable from μ under π , but where π is not weak/strong asymptotically optimal in ν .

It is worth remarking that for all counter-examples, the set of observations \mathcal{O} is empty and so all results apply also to bandits. Space does not permit us to present the proofs of part 1 and 2, which in any case are substantially easier than part 3.

Proof (Theorem 7, Part 3). Recall $\gamma_k = \frac{1}{k(k+1)}$ and so $\Gamma_t = \frac{1}{t}$. Now let $\mathcal{Y} = \{up, down\}$ and $\mathcal{O} = \emptyset$. Define μ by

$$\mu(y_{r < t} y_t) = \begin{cases} \frac{1}{2} & \text{if } y_t = up \\ \frac{1}{2} - \epsilon & \text{if } y_t = down \end{cases}$$

where $\epsilon \in (0, \frac{1}{2})$ will be chosen later. As before, $V_\mu^*(y_{r < t}) = \frac{1}{2}$. Assume π is weakly asymptotically optimal. Therefore

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{t=1}^n V_\mu^\pi(y_{r < t}^{\mu, \pi}) = \frac{1}{2}. \tag{4}$$

We show by contradiction that π cannot explore (take action *down*) too often. Assume there exists an infinite time-sequence t_1, t_2, t_3, \dots such that $\pi(y_{r < t}^{\mu, \pi}) = \textit{down}$ for all $t \in \bigcup_{i=1}^\infty [t_i, 2t_i]$. Then for $t \in [t_i, \frac{3}{2}t_i]$ we have

$$V_\mu^\pi(y_{r < t}^{\mu, \pi}) \equiv \frac{1}{\Gamma_t} \sum_{k=t}^\infty \gamma_k r_k^{\mu, \pi} \leq t \left[\left(\frac{1}{2} - \epsilon\right) \sum_{k=t}^{2t_i} \gamma_k + \frac{1}{2} \sum_{k=2t_i+1}^\infty \gamma_k \right] \tag{5}$$

$$= \frac{1}{2} - \epsilon \left[1 - \frac{t}{2t_i + 1} \right] < \frac{1}{2} - \frac{\epsilon}{4} \tag{6}$$

where (5) is the definition of the value function and the previous assumption and definition of μ . (6) by algebra and since $t \in [t_i, \frac{3}{2}t_i]$. Therefore

$$\frac{1}{2t_i} \sum_{t=1}^{2t_i} V_\mu^\pi(y_{r < t}^{\mu, \pi}) < \frac{1}{2t_i} \left[\sum_{t=1}^{t_i-1} \frac{1}{2} + \sum_{t=t_i}^{\frac{3}{2}t_i-1} \left(\frac{1}{2} - \frac{\epsilon}{4}\right) + \sum_{t=\frac{3}{2}t_i}^{2t_i} \frac{1}{2} \right] = \frac{1}{2} - \frac{1}{16}\epsilon. \tag{7}$$

The first inequality follows from (6) and because the maximum value of any play-out sequence in μ is $\frac{1}{2}$. The second by algebra. Therefore $\liminf_{n \rightarrow \infty} \frac{1}{n} \sum_{t=1}^n V_\mu^\pi(y_{r < t}^{\mu, \pi}) < \frac{1}{2} - \frac{1}{16}\epsilon < \frac{1}{2}$, which contradicts (4). Therefore there does not exist a time-sequence $t_1 < t_2 < t_3 < \dots$ such that $\pi(y_{r < t}^{\mu, \pi}) = \textit{down}$ for all $t \in \bigcup_{i=1}^\infty [t_i, 2t_i]$.

So far we have shown that π cannot “explore” for t consecutive time-steps starting at time-step t , infinitely often. We now construct an environment similar to μ where this is required. Choose T to be larger than the last time-step t at which $y_s^{\mu, \pi} = \textit{down}$ for all $s \in [t, 2t]$. Define ν by

$$\nu(y_{r < t} y_t) = \begin{cases} \mu(y_{r < t} y_t) & \text{if } t < T \\ \frac{1}{2} & \text{if } y_t = \textit{down} \text{ and there does not exist } t' \geq T \\ & \text{such that } y_s = \textit{down} \forall s \in [t', 2t'] \\ 1 & \text{if } y_t = \textit{down} \text{ and exists } t' \geq T \text{ such that } 2t' < t \text{ and} \\ & y_s = \textit{down} \forall s \in [t', 2t'] \\ \frac{1}{2} - \epsilon & \text{otherwise} \end{cases}$$

Now we compare the values in environment ν of π and π_ν^* at times $t \geq T$. Since π does not take action *down* for t consecutive time-steps at any time after T , it never “unlocks” the reward of 1 and so $V_\nu^\pi(y_{r < t}^{\nu, \pi}) \leq \frac{1}{2}$. Now let $\tilde{\pi}(y_{r < t}) = \textit{down}$ for all $y_{r < t}$. Therefore, for $t \geq 2T$,

$$V_\nu^{\tilde{\pi}}(\mathbf{y}_{<t}^{\nu,\pi}) \equiv \frac{1}{\Gamma_t} \sum_{k=t}^\infty \gamma_k r_k^{\nu,\tilde{\pi}} \geq t \left[\left(\frac{1}{2} - \epsilon\right) \sum_{k=t}^{2t-1} \gamma_k + \sum_{k=2t}^\infty \gamma_k \right] \tag{8}$$

$$= t \left[\left(\frac{1}{2} - \epsilon\right) \left(\frac{1}{t} - \frac{1}{2t}\right) + \frac{1}{2t} \right] = \frac{3}{4} - \frac{1}{2}\epsilon \tag{9}$$

where (8) follows by the definition of ν and $\tilde{\pi}$. (9) by the definition of γ_k and algebra. Finally, setting $\epsilon = \frac{1}{4}$ gives $V_\nu^{\tilde{\pi}}(\mathbf{y}_{<t}^{\nu,\pi}) \geq \frac{5}{8} = \frac{1}{2} + \frac{1}{8}$. Since $V_\nu^* \geq V_\nu^{\tilde{\pi}}$, we get $V_\nu^*(\mathbf{y}_{<t}^{\nu,\pi}) - V_\nu^\pi(\mathbf{y}_{<t}^{\nu,\pi}) \geq V_\nu^{\tilde{\pi}}(\mathbf{y}_{<t}^{\nu,\pi}) - V_\nu^\pi(\mathbf{y}_{<t}^{\nu,\pi}) \geq \frac{1}{8}$. Therefore $\limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{t=1}^n [V^*(\mathbf{y}_{<t}^{\nu,\pi}) - V_\nu^\pi(\mathbf{y}_{<t}^{\nu,\pi})] \geq \frac{1}{8}$, and so π is not weakly asymptotically optimal. \square

We believe it should be possible to generalise the above to computable discount functions with $H_t(p) > c_p t$ with $c_p > 0$ for infinitely many t , but the proof will likely be messy.

4 Existence of Weak Asymptotically Optimal Policies

In the previous section we showed there did not exist a strong asymptotically optimal policy (for most discount functions) and that any weak asymptotically optimal policy must be incomputable. In this section we show that a weak asymptotically optimal policy exists for geometric discounting (and is, of course, incomputable).

The policy is reminiscent of ϵ -exploration in finite state MDPs (or UCB for bandits) in that it spends most of its time exploiting the information it already knows, while still exploring sufficiently often (and for sufficiently long) to detect any significant errors in its model.

The idea will be to use a model-based policy that chooses its current model to be the first environment in the model class (all computable deterministic environments) consistent with the history seen so far. With increasing probability it takes the best action according to this policy, while still occasionally exploring randomly. When it explores it always does so in bursts of increasing length.

Definition 8 (History Consistent). *A deterministic environment μ is consistent with history $\mathbf{y}_{<t}$ if $\mu(\mathbf{y}_{<k} y_k) = x_k$, for all $k < t$.*

Definition 9 (Weak Asymptotically Optimal Policy). *Let $\mathcal{Y} = \{0,1\}$ and $\mathcal{M} = \{\mu_1, \mu_2, \mu_3, \dots\}$ be a countable class of deterministic environments. Define a probability measure P on \mathcal{B}^∞ inductively by, $P(z_n = 1 | z_{<n}) := \frac{1}{n}$, for all $z_{<n} \in \mathcal{B}^{n-1}$. Now let $\chi \in \mathcal{B}^\infty$ be sampled from P and define $\bar{\chi}, \dot{\chi}^h \in \mathcal{B}^\infty$ by*

$$\bar{\chi}_k := \begin{cases} 1 & \text{if } k \in \bigcup_{i:\chi_i=1} [i, i + \log i] \\ 0 & \text{otherwise} \end{cases} \quad \dot{\chi}_k^h := \begin{cases} 0 & \text{if } \bar{\chi}_{k:k+h} = 0^{h+1} \\ 1 & \text{otherwise} \end{cases}$$

Next let ψ be sampled from the uniform measure (each bit of ψ is independently sampled from a Bernoulli $1/2$ distribution) and define a policy π by,

$$\pi(y_{<t}) := \begin{cases} \pi_{\nu_t}^*(y_{<t}^{\pi, \mu}) & \text{if } \bar{\chi}_t = 0 \\ \psi_t & \text{otherwise} \end{cases} \tag{10}$$

where $\nu_t = \mu_{i_t}$ with $i_t = \min \{i : \mu_i \text{ consistent with history } y_{<t}^{\pi, \mu}\} < \infty$. Note that i_t is always finite because there exists an i such that $\mu_i = \mu$, in which case μ_i is necessarily consistent with $y_{<t}^{\pi, \mu}$.

Intuitively, $\chi_k = 1$ at time-steps when the agent will explore for $\log k$ time-steps. $\bar{\chi}_k = 1$ if the agent is exploring at time k and ψ_k is the action taken if exploring at time-step k . $\dot{\chi}$ will be used later, with $\dot{\chi}_k^h = 1$ if the agent will explore at least once in the interval $[k, k + h]$. If the agent is not exploring then it acts according to the optimal policy for the first consistent environment in \mathcal{M} .

Theorem 10. *Let $\gamma_t = \gamma^t$ with $\gamma \in (0, 1)$ (geometric discounting) then the policy defined in Definition 9 is weakly asymptotically optimal in the class of all deterministic computable environments with probability 1.*

Some remarks:

1. That $\mathcal{Y} = \{0, 1\}$ is only convenience, rather than necessity. The policy is easily generalised to arbitrary finite \mathcal{Y} .
2. π is essentially a stochastic policy. With some technical difficulties it is possible to construct an equivalent deterministic policy. This is done by choosing χ to be any P -Martin-Löf random sequence and ψ to be a sequence that is Martin-Löf random w.r.t to the uniform measure. The theorem then holds for *all* deterministic environments. The proof is somewhat delicate and may not extend nicely to stochastic environments. For an introduction to Kolmogorov complexity and Martin-Löf randomness, see [12]. For a reason why the stochastic case may not go through as easily, see [8].
3. The policy defined in Definition 9 is not computable for two reasons. First, because it relies on the stochastic sequences χ and ψ . Second, because the operation of finding the first environment consistent with the history is not computable.² We do not know if there exists a weak asymptotically optimal policy that is computable when given access to a random number generator (or if it is given χ and ψ).
4. The bursts of exploration are required for optimality. Without them it will be possible to construct counter-example environments similar to those used in part 3 of Theorem 7.

Before the proof we require some more definitions and lemmas. Easier proofs are omitted due to space limitations.

² The class of computable environments is not recursively enumerable [12].

Definition 11 (*h*-Difference). Let μ and ν be two environments consistent with history $\mathbf{y}_{<t}$, then μ is *h*-different to ν if there exists $\mathbf{y}_{t:t+h}$ satisfying

$$\begin{aligned} y_k &= \pi_\mu^*(\mathbf{y}_{<k}) \text{ for all } k \in [t, t+h], \\ x_k &= \mu(\mathbf{y}_{<k}y_k) \text{ for all } k \in [t, t+h], \\ x_k &\neq \nu(\mathbf{y}_{<k}y_k) \text{ for some } k \in [t, t+h]. \end{aligned}$$

Intuitively, μ is *h*-different to ν at history $\mathbf{y}_{<t}$ if playing the optimal policy for μ for *h* time-steps makes ν inconsistent with the new history. Note that *h*-difference is *not* symmetric.

Lemma 12. If $a_n \in [0, 1]$ and $\limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n a_i = \epsilon$ and $\alpha \in \mathcal{B}^\infty$ is an indicator sequence with $\alpha_i := \mathbb{1}[a_i \geq \epsilon/4]$,³ then $\prod_{i=1}^\infty [1 - \frac{\alpha_i}{i}] = 0$.

Lemma 13. Let a_1, a_2, a_3, \dots be a sequence with $a_n \in [0, 1]$ for all *n*. The following properties of χ are true with probability 1.

1. For any *h*, $\limsup_{n \rightarrow \infty} \frac{1}{n} \#1(\dot{\chi}_{1:n}^h) = 0$.
2. If $\limsup \frac{1}{n} \sum_{i=1}^n a_i = \epsilon > 0$ and $\alpha_i := \mathbb{1}[a_i > \epsilon/2]$ then $\alpha_i = \chi_i = 1$ for infinitely many *i*.

Proof. 1. Let $i \in \mathbb{N}$, $\epsilon > 0$ and E_i^ϵ be the event that $\#1(\dot{\chi}_{1:2^i}^h) > 2^i \epsilon$. Using the definition of $\dot{\chi}^h$ to compute the expectation $\mathbf{E}[\#1(\dot{\chi}_{1:2^i}^h)] < i(i+1)h$ and applying the Markov inequality gives that $P(E_i^\epsilon) < i(i+1)h2^{-i}/\epsilon$. Therefore $\sum_{i \in \mathbb{N}} P(E_i^\epsilon) < \infty$. Therefore the Borel-Cantelli lemma gives that E_i^ϵ occurs for only finitely many *i* with probability 1. We now assume that $\limsup_{n \rightarrow \infty} \frac{1}{n} \#1(\dot{\chi}_{1:n}^h) > 2\epsilon > 0$ and show that E_i^ϵ must occur infinitely often. By the definition of \limsup and our assumption we have that there exists a sequence n_1, n_2, \dots such that $\#1(\dot{\chi}_{1:n_i}^h) > 2n_i \epsilon$ for all $i \in \mathbb{N}$. Let $n^+ := \min_{k \in \mathbb{N}} \{2^k : 2^k \geq n\}$ and note that $\#1(\dot{\chi}_{1:n^+}^h) > n_i^+ \epsilon$, which is exactly $E_{\log n_i^+}^\epsilon$. Therefore there exist infinitely many *i* such that E_i^ϵ occurs and so $\limsup_{n \rightarrow \infty} \frac{1}{n} \#1(\dot{\chi}_{1:n}^h) = 0$ with probability 1.

2. The probability that $\alpha_i = 1 \implies \chi_i = 0$ for all $i \geq T$ is $P(\alpha_i = 1 \implies \chi_i = 0 \forall i \geq T) = \prod_{i=T}^\infty (1 - \frac{\alpha_i}{i}) =: p = 0$, by Lemma 12. Therefore the probability that $\alpha_i = \chi_i = 1$ for only finitely many *i* is zero. Therefore there exists infinitely many *i* with $\alpha_i = \chi_i = 1$ with probability 1, as required. \square

Lemma 14 (Approximation Lemma). Let π_1 and π_2 be policies, μ an environment and $h \geq H_t(1 - \epsilon)$. Let $\mathbf{y}_{<t}$ be an arbitrary history and $\mathbf{w}_{t:t+h}^{\mu, \pi_i}$ be the future action/observation/reward triples when playing policy π_i . If $\mathbf{w}_{t:t+h}^{\pi_1, \mu} = \mathbf{w}_{t:t+h}^{\pi_2, \mu}$ then $|V_{\mu}^{\pi_1}(\mathbf{y}_{<t}) - V_{\mu}^{\pi_2}(\mathbf{y}_{<t})| < \epsilon$.

Recall that π_μ^* and π_ν^* are the optimal policies in environments μ and ν respectively (see Definition 5).

³ $\mathbb{1}[\textit{expression}] = 1$ if *expression* is true and 0 otherwise.

Lemma 15 (*h*-difference). *If $|V_\mu^{\pi^\mu}(y_{<t}^{\pi,\mu}) - V_\mu^{\pi^\nu}(y_{<t}^{\pi,\mu})| > \epsilon$ then μ is $H_t(1 - \epsilon)$ -different to ν on $y_{<t}^{\pi,\mu}$.*

We are now ready to prove the main theorem.

Proof (Theorem 10). Let π be the policy defined in Definition 9 and μ be the true (unknown) environment. Recall that $\nu_t = \mu_{i_t}$ with $i_t = \min \{i : \mu_i \text{ consistent with history } y_{<t}^{\pi,\mu}\}$ is the first model consistent with the history $y_{<t}^{\pi,\mu}$ at time t and is used by π when not exploring. First we claim there exists a $T \in \mathbb{N}$ and environment ν such that $\nu_t = \nu$ for all $t \geq T$. Two facts,

1. If μ_i is inconsistent with history $y_{<t}^{\pi,\mu}$ then it is also inconsistent with $y_{<t+h}^{\pi,\mu}$ for all $h \in \mathbb{N}$.
2. μ is consistent with $y_{<t}^{\pi,\mu}$ for all π, t .

By 1) we have that the sequence i_1, i_2, i_3, \dots is monotone increasing. By 2) we have that the sequence is bounded by i with $\mu_i = \mu$. The claim follows since any bounded monotone sequence of natural numbers converges in finite time. Let $\nu := \nu_\infty$ be the environment to which $\nu_1, \nu_2, \nu_3, \dots$ converges to. Note that ν must be consistent with history $y_{<t}^{\mu,\pi}$ for all t . We now show by contradiction that the optimal policy for ν is weakly asymptotically optimal in environment μ . Suppose it were not, then

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{t=1}^n \left[V_\mu^*(y_{<t}^{\pi,\mu}) - V_\mu^{\pi^\nu}(y_{<t}^{\pi,\mu}) \right] = \epsilon > 0. \tag{11}$$

Let $\alpha \in \mathcal{B}^\infty$ be defined by $\alpha_t := 1$ if and only if,

$$\left[V_\mu^*(y_{<t}^{\pi,\mu}) - V_\mu^\pi(y_{<t}^{\pi,\mu}) \right] \geq \epsilon/4. \tag{12}$$

By Lemma 13 there exists (with probability one) an infinite sequence t_1, t_2, t_3, \dots for which $\chi_k = \alpha_k = 1$. Intuitively we should view time-step t_k as the start of an ‘‘exploration’’ phase where the agent explores for $\log t_k$ time-steps. Let $h := H_{t_k}(1 - \epsilon/4) = \lceil \log(\epsilon/4) / \log \gamma \rceil$, which importantly is independent of t_k (for geometric discounting). Since $\log t_k \rightarrow \infty$ we will assume that $\log t_k \geq h$ for all t_k . Therefore $\bar{\chi}_i = 1$ for all $i \in \bigcup_{k=1}^\infty [t_k, t_k + h]$. Therefore by the definition of π , $\pi(y_{<i}^{\pi,\mu}) = \psi_i$ for $i \in \bigcup_{k=1}^\infty [t_k, t_k + h]$. By Lemma 15 and Equation (12), μ is h -different to ν on history $y_{<t_k}^{\pi,\mu}$. This means that if there exists a k such that π plays according to the optimal policy for μ on all time-steps $t \in [t_k, t_k + h]$ then ν will be inconsistent with the history $y_{1:t_k+h}^{\mu,\pi}$, which is a contradiction. We now show that π does indeed play according to the optimal policy for μ for all time-steps $t \in [t_k, t_k + h]$ for at least one k . Formally, we show the following holds with probability 1 for some k .

$$\psi_i \equiv \pi(y_{<i}^{\pi,\mu}) = \pi_\mu^*(y_{<i}^{\pi,\mu}), \text{ for all } i \in [t_k, t_k + h]. \tag{13}$$

Recall that $\psi \in \mathcal{B}^\infty$ where $\psi_i \in \mathcal{B}$ is identically independently distributed according to a Bernoulli($\frac{1}{2}$) distribution. Therefore $P(\psi_i = \pi_\mu^*(y_{<i}^{\pi,\mu})) = \frac{1}{2}$. Therefore

$p := P(\psi_i = \pi_\mu^*(y_{<i}^{\pi,\mu}) \forall i \in [t_k, t_k + h]) = \prod_{i=t_k}^{t_k+h} P(\psi_i = \pi_\mu^*(y_{<i}^{\pi,\mu})) = 2^{-h-1} > 0$ and $P(\forall k \exists i \in [t_k, t_k + h] \text{ with } \psi_i \neq \pi_\mu^*(y_{<i}^{\pi,\mu})) = \prod_{k=1}^\infty (1 - p) = 0$. Therefore Equation (13) is satisfied for some k with probability 1 and so Equation (11) leads to a contradiction. Therefore

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{t=1}^n [V_\mu^*(y_{<t}^{\pi,\mu}) - V_\mu^{\pi_\nu^*}(y_{<t}^{\pi,\mu})] = 0. \tag{14}$$

We have shown that the optimal policy for ν has similar μ -values to the optimal policy for μ . We now show that π acts according to π_ν^* sufficiently often that it too has values close to those of the optimum policy for the true environment, μ . Let $\epsilon > 0$, $h := H_t(1 - \epsilon)$ and $t \geq T$. If $\dot{\chi}_t^h = 0$ then by the definition of π and the approximation lemma we obtain

$$\left| V_\mu^{\pi_\nu^*}(y_{<t}^{\pi,\mu}) - V_\mu^\pi(y_{<t}^{\pi,\mu}) \right| < \epsilon. \tag{15}$$

Therefore

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{t=1}^n \left| V_\mu^{\pi_\nu^*}(y_{<t}^{\pi,\mu}) - V_\mu^\pi(y_{<t}^{\pi,\mu}) \right| \leq \limsup_{n \rightarrow \infty} \frac{1}{n} \left[\sum_{t=1}^{T-1} 1 + \sum_{t=T}^n [\dot{\chi}_t^h(1 - \epsilon) + \epsilon] \right] \tag{16}$$

$$= \epsilon + (1 - \epsilon) \limsup_{n \rightarrow \infty} \frac{1}{n} \#1(\dot{\chi}_{T:n}^h) \tag{17}$$

$$= \epsilon \tag{18}$$

where (16) follows since values are bounded in $[0, 1]$ and Equation (15). (17) follows by algebra. (18) by part 1 of Lemma 13. By sending $\epsilon \rightarrow 0$,

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{t=1}^n [V_\mu^{\pi_\nu^*}(y_{<t}^{\pi,\mu}) - V_\mu^\pi(y_{<t}^{\pi,\mu})] = 0. \tag{19}$$

Finally, combining Equations (14) and (19) gives the result. □

We expect this theorem to generalise without great difficulty to discount functions satisfying $H_t(p) < c_p \log(t)$ for all p . There will be two key changes. First, extend the exploration time to some function $E(t)$ with $E(t) \in O(H_p(t))$ for all p . Second, modify the probability of exploration to ensure that Lemma 13 remains true.

5 Discussion

Summary. Part 1 of Theorem 7 shows that no policy can be strongly asymptotically optimal for any computable discount function. The key insight is that strong asymptotic optimality essentially implies exploration must eventually

cease. Once this occurs, the environment can change without the agent discovering the difference and the policy will no longer be optimal.

A weaker notion of asymptotic optimality, that a policy be optimal on average in the limit, turns out to be more interesting. Part 2 of Theorem 7 shows that no weak asymptotically optimal policy can be computable. We should not be surprised by this result. Any computable policy can be used to construct a computable environment in which that policy does very badly. Note that by computable here we mean deterministic and computable. There may be computable stochastic policies that are weakly asymptotically optimal, but we feel this is unlikely.

Part 3 of Theorem 7, shows that even weak asymptotically optimal policies need not exist if the discount function is sufficiently far-sighted. On the other hand, Theorem 10 shows that weak asymptotically optimal policies do exist for some discount rates, in particular, for the default geometric discounting. These non-trivial and slightly surprising result shows that choice of discount function is crucial to the existence of weak asymptotically optimal policies. Where weak asymptotically optimal policies do exist, they must explore infinitely often and in increasing contiguous bursts of exploration where the length of each burst is dependent on the discount function.

Consequences. It would appear that Theorem 7 is problematic for artificial general intelligence. We cannot construct incomputable policies, and so we cannot construct weak asymptotically optimal policies. However, this is not as problematic as it may seem. There are a number of reasonable counter arguments:

1. We may be able to make stochastically computable policies that are asymptotically optimal. If the existence of true random noise is assumed then this would be a good solution.
2. The counter-example environment constructed in part 2 of Theorem 7 is a single environment roughly as complex as the policy itself. Certainly, if the world were adversarial this would be a problem, but in general this appears not to be the case. On the other hand, if the environment is a learning agent itself, this could result in a complexity arms race without bound. There may exist a computable weak asymptotically optimal policy in some extremely large class of environments. For example, the algorithm of Section 4 is stochastically computable when the class of environments is recursively enumerable and contains only computable environments. A natural (and already quite large) class satisfying these properties is finite-state Markov decision processes with $\{0, 1\}$ -valued transition functions and rational-valued rewards.
3. While it is mathematically pleasant to use asymptotic behaviour to characterise optimal general intelligent behaviour, in practise we usually care about more immediate behaviour. We expect that results, and even (parameter free) formal definitions of intelligence satisfying this need will be challenging, but worthwhile.
4. Accept that even weak asymptotic optimality is too strong and find something weaker, but still useful.

Relation to AIXI. The policy defined in Section 4 is not equivalent to AIXI [7], which is also incomputable. However, if the computable environments in \mathcal{M} are ordered by complexity then it is likely the two will be quite similar. The key difference is the policy defined in this paper will continue to explore whereas it was shown in [14] that AIXI eventually ceases exploration in some environments and some histories. We believe, and a proof should not be too hard, that AIXI will become weakly asymptotically optimal if an exploration component is added similarly as in Section 4.

We now briefly compare the self-optimising property in [6] to strong asymptotic optimality. A policy π is self-optimising in a class \mathcal{M} if $\lim_{t \rightarrow \infty} [V_{\mu}^*(y_{x_{<t}}) - V_{\mu}^{\pi}(y_{x_{<t}})] = 0$ for any infinite history $y_{x_{1:\infty}}$ and for all $\mu \in \mathcal{M}$. This is similar to strong asymptotic optimality, but convergence must be on all histories, rather than the histories actually generated by π . This makes the self-optimising property a substantially stronger form of optimality than strong asymptotic optimality. It has been proven that if there exists self-optimising policy for a particular class, then AIXI is also self-optimising in that class [6].

It is possible to define a weak version of the self-optimising property by insisting that $\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{t=1}^n [V_{\mu}^*(y_{x_{<t}}) - V_{\mu}^{\pi}(y_{x_{<t}})] = 0$ for all $y_{x_{1:\infty}}$ and all $\mu \in \mathcal{M}$. It can then be proven that the existence of a weak self-optimising policy would imply that AIXI were also weakly self-optimising. However, the policy defined in Section 4 cannot be modified to have the weak self-optimising property. It must be allowed to choose its actions itself. This is consistent with the work in [14] which shows that AIXI cannot be weakly asymptotically optimal, and so cannot be weak self-optimising either.

Discounting. Throughout this paper we have assumed rewards to be discounted according to a summable discount function. A very natural alternative to discounting, suggested in [11], is to restrict interest to environments satisfying $\sum_{k=1}^{\infty} r_k^{\mu, \pi} \leq 1$. Now the goal of the agent is simply to maximise summed rewards. In this setting it is easy to see that the positive theorem is lost while all negative ones still hold! This is unfortunate, as discounting presents a major philosophical challenge. How to choose a discount function?

Assumptions/Limitations. Assumption 1 ensures that \mathcal{Y} and \mathcal{O} are finite. All negative results go through for countable \mathcal{Y} and \mathcal{O} . The optimal policy of Section 4 may not generalise to countable \mathcal{Y} . We have also assumed bounded reward and discrete time. The first seems reasonable while the second allows for substantially easier analysis. Additionally we have only considered deterministic computable environments. The stochastic case is unquestionably interesting. We invoked Church thesis to assert that computable stochastic environments are essentially the largest class of interesting environments.

Many of our Theorems are only applicable to computable discount functions. All well-known discount function in use today are computable. However [7] has suggested $\gamma_t = 2^{-K(t)}$, where $K(t)$ is the (incomputable) prefix Kolmogorov complexity of t , may have nice theoretical properties.

Acknowledgements. We thank Laurent Orseau, Wen Shao and reviewers for valuable feedback on earlier drafts and the Australian Research Council for support under grant DP0988049.

References

- [1] Auer, P., Cesa-Bianchi, N., Fischer, P.: Finite-time analysis of the multiarmed bandit problem. *Machine Learning* 47, 235–256 (2002)
- [2] Berry, D.A., Fristedt, B.: *Bandit Problems: Sequential Allocation of Experiments*. Chapman and Hall, London (1985)
- [3] Diaconis, P., Freedman, D.: On inconsistent Bayes estimates of location. *The Annals of Statistics* 14(1), 68–87 (1986)
- [4] Diaconis, P., Freedman, D.: On the consistency of Bayes estimates. *The Annals of Statistics* 14(1), 1–26 (1986)
- [5] Frederick, S., Oewenstein, G.L., O’Donoghue, T.: Time discounting and time preference: A critical review. *Journal of Economic Literature* 40(2) (2002)
- [6] Hutter, M.: Self-optimizing and Pareto-optimal policies in general environments based on Bayes-mixtures. In: Kivinen, J., Sloan, R.H. (eds.) *COLT 2002. LNCS (LNAI)*, vol. 2375, pp. 364–379. Springer, Heidelberg (2002)
- [7] Hutter, M.: *Universal Artificial Intelligence: Sequential Decisions based on Algorithmic Probability*. Springer, Berlin (2004)
- [8] Hutter, M., Muchnik, A.A.: On semimeasures predicting Martin-Löf random sequences. *Theoretical Computer Science* 382(3), 247–261 (2007)
- [9] Lattimore, T., Hutter, M.: Time consistent discounting. In: Kivinen, J., Szepesvári, C., Ukkonen, E., Zeugmann, T. (eds.) *Algorithmic Learning Theory. LNCS*, vol. 6925, pp. 384–398. Springer, Heidelberg (2011)
- [10] Legg, S.: Is there an elegant universal theory of prediction? In: Balcázar, J.L., Long, P.M., Stephan, F. (eds.) *ALT 2006. LNCS (LNAI)*, vol. 4264, pp. 274–287. Springer, Heidelberg (2006)
- [11] Legg, S., Hutter, M.: Universal intelligence: A definition of machine intelligence. *Minds & Machines* 17(4), 391–444 (2007)
- [12] Li, M., Vitanyi, P.: *An Introduction to Kolmogorov Complexity and Its Applications*, 3rd edn. Springer, Heidelberg (2008)
- [13] Norvig, P., Russell, S.J.: *Artificial Intelligence: A Modern Approach*, 2nd edn. Prentice Hall Series in Artificial Intelligence. Prentice Hall, Englewood Cliffs (2003)
- [14] Orseau, L.: Optimality issues of universal greedy agents with static priors. In: Hutter, M., Stephan, F., Vovk, V., Zeugmann, T. (eds.) *ALT 2010. LNCS*, vol. 6331, pp. 345–359. Springer, Heidelberg (2010)
- [15] Strehl, A.L., Littman, M.L.: An analysis of model-based interval estimation for Markov decision processes. *Journal of Computer and System Sciences* 74(8), 1309–1331 (2008)