# Private memory confers no advantage 

Samuel Allen Alexander ${ }^{1[0000-0002-7930-110 X]}$<br>${ }^{1}$ The U.S. Securities and Exchange Commission, samuelallenalexander@gmail.com


#### Abstract

Mathematicians and software developers use the word "function" very differently, and yet, sometimes, things that are in practice implemented using the software developer's "function", are mathematically formalized using the mathematician's "function". This mismatch can lead to inaccurate formalisms. We consider a special case of this meta-problem. Various kinds of agents might, in actual practice, make use of private memory, reading and writing to a memory-bank invisible to the ambient environment. In some sense, we humans do this when we silently subvocalize thoughts about the actions we are taking (at least when the environment we're in is too primitive to probe the contents of our brains). Mathematical function formalizations of agents often ignore this ability. We show that in a general agent-environment framework (of which reinforcement learning is a special case), in a technical sense, such private memories do not enable qualitatively different agent behavior.


## 1 Introduction

There are many different ways to formalize the interaction between an agent and an environment. At a high level, the idea is simple: agent and environment take turns. On the agent's turn, the agent takes an action. On the environment's turn, the environment generates a percept for the agent to see. But precisely how to formalize this depends on various questions about things like what the agent and environment can remember, whether or not agent and/or environment can have an element of randomness, and so on. For example, if we are only interested in memoryless agents who see nothing but the most recent environmental percept, then we can formalize an agent as a function which takes a percept as input and outputs an action (or an action probability distribution). Examples of such agents include the so-called policies in the Stable Baselines3 reinforcement learning package [7]. On the other hand, if we are interested in agents who remember the entire interaction so far, then we should formalize an agent as a function which takes a history as input, not just a single percept. Examples of such agents include the agents in the Legg-Hutter formalization of reinforcement learning $[4,6]$.

And when formalizing agents in the latter way, there are still questions about what needs to be included in the history which is input into the agent-function. For example, if the agent's actions are completely deterministic, then a history need only include the sequence of percepts seen so far. From these, the agent's actions can be inferred, provided the agent is deterministic. It is superfluous
(at least ignoring computational efficiency concerns) for a deterministic agent to remember its own past actions, provided that it remembers the percepts that prompted those actions: the actions themselves, being deterministic, are determined by the percepts. But, if the agent is non-deterministic (i.e., if agents are functions that output action probability distributions rather than just actions), then the history which we input into the agent should include both the past environmental percepts and also the agent's own past actions. Past actions can not, generally, be inferred from past environmental percepts, if said actions had an element of randomness to them.

In addition to memories of past percepts and past actions, we can also imagine agents who have additional private memory. For example, we might imagine that the agent has access to a piece of scratch paper stored outside of the environment itself (and thus invisible to the environment) on which the agent can write notes to its own future self. The human hippocampus is not, in reality, an example of such a private memory bank, because, in reality, the human hippocampus is part of the environment (and thus vulnerable to being observed or modified by the outside world). But if, as simplifying assumption, the human agent is modelled as having his brain perfectly isolated, then the human hippocampus does become a place for such private memories to be stored. In such an idealized model, the human agent can silently think to herself things like: "I'm going to see what happens when I push this button. If it hurts me, then I won't push any more buttons like it in the future." And then, later on, the human agent has memory of such internal chatter that she made in the past.

In this paper, we consider agents who have memory of past environmental percepts and of their own past actions. We are concerned with the question: does private memory (in addition to memory of past actions and past environmental percepts) confer any advantage? We will show that the answer is "No," in the following sense. We will show that if $\pi$ is any agent with private memory in addition to memory of the past agent-environment interaction, then there is an agent $\pi^{\prime}$ with only memory of past agent-environment interaction (and no private memory), such that for every finite initial history $h$ of percepts and actions, the probably of $h$ occurring when $\pi$ interacts with a given environment is exactly equal to the probably of $h$ occurring when $\pi^{\prime}$ interacts with that environment. This is trivial if private memories are written deterministically ( $\pi^{\prime}$ can then reconstruct $\pi$ 's private memories by simulating $\pi$ on the past percepts and actions which $\pi^{\prime}$ remembers). But if private memories have an element of randomness to them, then this result is much less obvious.

Although private memory confers no advantage in the above formal sense, there are other senses in which is might be considered to confer advantage. First of all, the $\pi^{\prime}$ described above is not guaranteed to have as good of a runtime complexity as the original $\pi$, and indeed, the $\pi^{\prime}$ which we will construct in our proof will generally be much more expensive than $\pi$. Second, the ability to use private memory might simplify the task of actually implementing a given agent. And thirdly, our proof only applies to traditional environments, not to environments capable of simulating the agent itself, as in [1] or [3].

## 2 Formal definitions

Throughout the paper, we fix a nonempty set $\mathcal{O}$ of percepts, a nonempty set $\mathcal{A}$ of actions, and a nonempty set $\mathcal{M}$ of private memories. We assume $\mathcal{O}, \mathcal{A}$, and $\mathcal{M}$ are pairwise disjoint.

Definition 1. If $X$ is any set, by a probability distribution on $X$ we mean a function $f: X \rightarrow[0,1]$ such that $\sum_{x \in X} f(x)=1$. For each $x \in X$, we say that $f$ assigns probability $f(x)$ to $x$. We write $\Delta(X)$ for the set of all probability distributions on $X$.

### 2.1 Agents without private memory

## Definition 2. (Histories)

- By an environment-seen history, we mean a sequence $\left\langle o_{1}, a_{1}, \ldots, o_{n}, a_{n}\right\rangle$, where each $o_{i} \in \mathcal{O}$ and each $a_{i} \in \mathcal{A}$. We also consider the empty sequence $\rangle$ to be an environment-seen history.
- By an agent-seen history, we mean a sequence $\left\langle o_{1}, a_{1}, \ldots, o_{n-1}, a_{n-1}, o_{n}\right\rangle$, where each $o_{i} \in \mathcal{O}$ and each $a_{i} \in \mathcal{A}$. Note that for each percept $o \in \mathcal{O}$, the length-1 sequence $\langle o\rangle$ is an agent-seen history.
- By a history we mean an environment-seen history or an agent-seen history.

Definition 3. (Environments) By an environment we mean a function $\mu$ which takes as input an environment-seen history $h$, and outputs a probability distribution $\mu(h) \in \Delta(\mathcal{O})$. If $h$ is an environment-seen history and $o \in \mathcal{O}$, we write $\mu(o \mid h)$ for $(\mu(h))(o)$ (the probability of o according to $\mu(h))$.

Definition 4. (Deterministic agents without private memory) By a deterministic agent without private memory we mean a function $\pi$ which takes as input an agent-seen history $h$, and outputs an action $\pi(h) \in \mathcal{A}$.

Definition 5. (Non-deterministic agents without private memory) By a nondeterministic agent without private memory we mean a function $\pi$ which takes as input an agent-seen history $h$, and outputs a probability distribution $\pi(h) \in$ $\Delta(\mathcal{A})$. If $h$ is an agent-seen history and $a \in \mathcal{A}$, we write $\pi(a \mid h)$ for $(\pi(h))(a)$ (the probability of a according to $\pi(h)$ ).

When a deterministic agent $\pi$ without private memory interacts with an environment $\mu$, we imagine a sequence $o_{1}, a_{1}, o_{2}, a_{2}, \ldots$ chosen randomly as follows:

- $o_{1}$ is chosen randomly using the probability distribution $\mu(\rangle)$.
$-a_{1}=\pi\left(\left\langle o_{1}\right\rangle\right)$.
$-o_{2}$ is chosen randomly using the probability distribution $\mu\left(\left\langle o_{1}, a_{1}\right\rangle\right)$.
$-a_{2}=\pi\left(\left\langle o_{1}, a_{1}, o_{2}\right\rangle\right)$.
- And so on.

When a non-deterministic agent $\pi$ without private memory interacts with an environment $\mu$, we imagine a sequence $o_{1}, a_{1}, o_{2}, a_{2}, \ldots$ chosen randomly as follows:

- $o_{1}$ is chosen randomly using the probability distribution $\mu(\rangle)$.
$-a_{1}$ is chosen randomly using the probability distribution $\pi\left(\left\langle o_{1}\right\rangle\right)$.
$-o_{2}$ is chosen randomly using the probability distribution $\mu\left(\left\langle o_{1}, a_{1}\right\rangle\right)$.
- $a_{2}$ is chosen randomly using the probability distribution $\mu\left(\left\langle o_{1}, a_{1}, o_{2}\right\rangle\right)$.
- And so on.

But instead of reasoning about these randomly-generated sequences, it is easier to reason about the probability of individual finite histories. In the following definition and throughout the rest of the paper, $\frown$ denotes concatenation.

Definition 6. (The probability of a history) Suppose $\mu$ is an environment and $\pi$ is a (deterministic or non-deterministic) agent without private memory. For every history $h$, let $P_{\mu}^{\pi}(h)$ be the probability that $h$ would be an initial sequence of the sequence $o_{1}, a_{1}, \ldots$ randomly generated from letting $\pi$ interact with $\mu$ as described above.

Some authors, such as [5], would write $P(h)$ or a variation thereof for $P_{\mu}^{\pi}(h)$, if $\pi$ and $\mu$ are clear from context.

### 2.2 Agents with private memory

Definition 7. By an agent-seen private-memory-augmented history we mean a sequence $\left\langle o_{1}, a_{1}, m_{1}, \ldots, o_{n-1}, a_{n-1}, m_{n-1}, o_{n}\right\rangle$ where each $o_{i} \in \mathcal{O}$, each $a_{i} \in \mathcal{A}$, and each $m_{i} \in \mathcal{M}$. Note that for each percept $o \in \mathcal{O}$, the length-1 sequence $\langle o\rangle$ is an agent-seen private-memory-augmented history.

Definition 8. (Deterministic agents with deterministic private memory) By a deterministic agent with deterministic private memory we mean a function $\pi$ which takes as input a private-memory-augmented agent-seen history $h$, and outputs an action-memory pair $\pi(h) \in \mathcal{A} \times \mathcal{M}$.

Definition 9. (Deterministic agents with non-deterministic private memory) By a deterministic agent with non-deterministic private memory we mean a function $\pi$ which takes as input a private-memory-augmented agent-seen history $h$, and outputs $\pi(h) \in \mathcal{A} \times \Delta(\mathcal{M})$.

Definition 10. (Non-deterministic agents with deterministic private memory) By a non-deterministic agent with deterministic private memory we mean a function $\pi$ which takes as input an agent-seen history $h$, and outputs $\pi(h) \in$ $\Delta(\mathcal{A}) \times \mathcal{M}$.

Definition 11. (Non-deterministic agents with non-deterministic private memory) By a non-deterministic agent with non-deterministic private memory we mean a function $\pi$ which takes as input an agent-seen history $h$, and outputs $\pi(h) \in \Delta(\mathcal{A} \times \mathcal{M})$.

When a deterministic agent $\pi$ with deterministic private memory interacts with an environment $\mu$, we imagine a sequence $o_{1}, a_{1}, m_{1}, o_{2}, a_{2}, m_{2}, \ldots$ chosen randomly as follows:

- $o_{1}$ is chosen randomly using the probability distribution $\mu(\rangle)$.
$-\left(a_{1}, m_{1}\right)=\pi\left(\left\langle o_{1}\right\rangle\right)$.
- $o_{2}$ is chosen randomly using the probability distribution $\mu\left(\left\langle o_{1}, a_{1}\right\rangle\right)$ (note the absense of $\left.m_{1}\right)$.
$-\left(a_{2}, m_{2}\right)=\pi\left(\left\langle o_{1}, a_{1}, m_{1}, o_{2}\right\rangle\right)$.
- $o_{3}$ is chosen randomly using the probability distribution $\mu\left(\left\langle o_{1}, a_{1}, o_{2}, a_{2}\right\rangle\right)$ (note the absense of $m_{1}, m_{2}$ ).
$-\left(a_{3}, m_{3}\right)=\pi\left(\left\langle o_{1}, a_{1}, m_{1}, o_{2}, a_{2}, m_{2}, o_{3}\right\rangle\right)$.
- And so on.

When a deterministic agent $\pi$ with non-deterministic private memory interacts with an environment $\mu$, we imagine a sequence $o_{1}, a_{1}, m_{1}, o_{2}, a_{2}, m_{2}, \ldots$ chosen randomly as follows:

- $o_{1}$ is chosen randomly using the probability distribution $\mu(\rangle)$.
$-a_{1}=a$ and $m_{1}$ is chosen randomly using the probability distribution $f$, where $(a, f)=\pi\left(\left\langle o_{1}\right\rangle\right)$.
$-o_{2}$ is chosen randomly using the probability distribution $\mu\left(\left\langle o_{1}, a_{1}\right\rangle\right)$ (note the absense of $m_{1}$ ).
$-a_{2}=a$ and $m_{2}$ is chosen randomly using the probability distribution $f$, where $(a, f)=\pi\left(\left\langle o_{1}, a_{1}, m_{1}, o_{2}\right\rangle\right)$.
$-o_{3}$ is chosen randomly using the probability distribution $\mu\left(\left\langle o_{1}, a_{1}, o_{2}, a_{2}\right\rangle\right)$ (note the absense of $m_{1}, m_{2}$ ).
$-a_{3}=a$ and $m_{3}$ is chosen randomly using the probability distribution $f$, where $(a, f)=\pi\left(\left\langle o_{1}, a_{1}, m_{1}, o_{2}, a_{2}, m_{2}, o_{3}\right\rangle\right)$.
- And so on.

When a non-deterministic agent $\pi$ with deterministic private memory interacts with an environment $\mu$, we imagine a sequence $o_{1}, a_{1}, m_{1}, o_{2}, a_{2}, m_{2}, \ldots$ chosen randomly as follows:

- $o_{1}$ is chosen randomly using the probability distribution $\mu(\rangle)$.
$-m_{1}=m$ and $a_{1}$ is chosen randomly using the probability distribution $f$, where $(f, m)=\pi\left(\left\langle o_{1}\right\rangle\right)$.
$-o_{2}$ is chosen randomly using the probability distribution $\mu\left(\left\langle o_{1}, a_{1}\right\rangle\right)$ (note the absense of $m_{1}$ ).
- $m_{2}=m$ and $a_{2}$ is chosen randomly using the probability distribution $f$, where $(f, m)=\pi\left(\left\langle o_{1}, a_{1}, m_{1}, o_{2}\right\rangle\right)$.
- $o_{3}$ is chosen randomly using the probability distribution $\mu\left(\left\langle o_{1}, a_{1}, o_{2}, a_{2}\right\rangle\right)$ (note the absense of $m_{1}, m_{2}$ ).
$-m_{3}=m$ and $a_{3}$ is chosen randomly using the probability distribution $f$, where $(f, m)=\pi\left(\left\langle o_{1}, a_{1}, m_{1}, o_{2}, a_{2}, m_{2}, o_{3}\right)\right.$.
- And so on.

When a non-deterministic agent $\pi$ with non-deterministic private memory interacts with an environment $\mu$, we imagine a sequence $o_{1}, a_{1}, m_{1}, o_{2}, a_{2}, m_{2}, \ldots$ chosen randomly as follows:

- $o_{1}$ is chosen randomly using the probability distribution $\mu(\rangle)$.
$-\left(a_{1}, m_{1}\right)$ is chosen randomly using the probability distribution $\pi\left(\left\langle o_{1}\right\rangle\right)$.
$-o_{2}$ is chosen randomly using the probability distribution $\mu\left(\left\langle o_{1}, a_{1}\right\rangle\right)$ (note the absense of $m_{1}$ ).
- $\left(a_{2}, m_{2}\right)$ is chosen randomly using the probability distribution

$$
\pi\left(\left\langle o_{1}, a_{1}, m_{1}, o_{2}\right\rangle\right)
$$

$-o_{3}$ is chosen randomly using the probability distribution $\mu\left(\left\langle o_{1}, a_{1}, o_{2}, a_{2}\right\rangle\right)$ (note the absense of $m_{1}, m_{2}$ ).

- $\left(a_{3}, m_{3}\right)$ is chosen randomly using the probability distribution

$$
\pi\left(\left\langle o_{1}, a_{1}, m_{1}, o_{2}, a_{2}, m_{2}, o_{3}\right\rangle\right)
$$

- And so on.

But instead of reasoning about these randomly-generated sequences, it is easier to reason about the probability of individual finite histories.

### 2.3 Conditional probabilities

Definition 12. (Conditional probabilities of histories) Suppose $\mu$ is an environment, $\pi$ is an agent (deterministic or non-deterministic, with or without deterministic or non-deterministic private memory), and $h$ is a history.

- Let $P_{\mu}^{\pi}(h)$ be the probability that when $\pi$ interacts with $\mu, h$ is an initial segment of the percept-action sequence randomly generated using $\pi$ and $\mu$ as described above.
- Let $P^{\pi}(h)$ be the probability that $h$ would be an initial segment of the perceptaction sequence randomly generated as described above if $\pi$ were to interact with some environment, on the condition that, in said random generation process, the percepts in $h$ are always selected.
- Let $P_{\mu}(h)$ be the probability that $h$ would be an initial segment of the perceptaction sequence randomly generated as described above if some agent were to interact with $\mu$, on the condition that, in said random generation process, the actions in $h$ are always selected.

Lemma 1. For any $\mu, \pi, h$ as in Definition 12, $P_{\mu}^{\pi}(h)=P^{\pi}(h) P_{\mu}(h)$.
Proof. By the multiplication rule of probability.

## 3 Private memory confers no advantage

Definition 13. (Agent equivalence) Let $\pi$ be an agent (deterministic or nondeterministic, with or without deterministic or non-deterministic private memory). Let $\rho$ be an agent (deterministic or non-deterministic, with or without deterministic or non-deterministic private memory). We say $\pi \sim \rho$ if the following requirement holds: for every environment $\mu$, for every history $h, P_{\mu}^{\pi}(h)=P_{\mu}^{\rho}(h)$.

Theorem 1. (Private memory confers no advantage)

1. If $\pi$ is a deterministic agent with deterministic private memory, there exists a deterministic agent $\pi^{\prime}$ without private memory, such that $\pi^{\prime} \sim \pi$.
2. If $\pi$ is a deterministic agent with non-deterministic private memory, there exists a non-deterministic agent $\pi^{\prime}$ without private memory, such that $\pi^{\prime} \sim$ $\pi$.
3. If $\pi$ is a non-deterministic agent with (deterministic or non-deterministic) private memory, there exists a non-deterministic agent $\pi^{\prime}$ without private memory, such that $\pi^{\prime} \sim \pi$.

Proof. Fix some $a_{0} \in \mathcal{A}$.
First, we will prove (2) and (3) together. Then, we will prove (1).
(2) and (3). Define $\pi^{\prime}$ so that for each agent-seen history $h$ and each $a \in \mathcal{A}$,

$$
\pi^{\prime}(a \mid h)= \begin{cases}\frac{P^{\pi}(h \frown a)}{P^{\pi}(h)} & \text { if } P^{\pi}(h) \neq 0 \\ 1 & \text { if } P^{\pi}(h)=0 \text { and } a=a_{0} \\ 0 & \text { if } P^{\pi}(h)=0 \text { and } a \neq a_{0}\end{cases}
$$

Claim: $\pi^{\prime}$ is a non-deterministic agent without private memory. The only thing required to show is that $\pi^{\prime}$ outputs probability distributions. Thus, we must show that for any agent-seen history $h, \sum_{a \in \mathcal{A}} \pi^{\prime}(a \mid h)=1$, and for all $a \in \mathcal{A}, 0 \leq \pi^{\prime}(a \mid h) \leq 1$. The latter inequalities are easy to show by induction. To show $\sum_{a \in \mathcal{A}} \pi^{\prime}(a \mid h)=1$, it suffices to consider the case $P^{\pi}(h) \neq 0$ (the other case is trivial). Suppose a percept-action sequence is randomly generated by letting $\pi$ interact with some environment, and that the percepts and actions so generated initially match $h$. The next action must be some action in $\mathcal{A}$. Thus $\sum_{a \in \mathcal{A}} P^{\pi}(h \frown a)=P^{\pi}(h)$. Therefore

$$
\sum_{a \in \mathcal{A}} \pi^{\prime}(a \mid h)=\sum_{a \in \mathcal{A}} \frac{P^{\pi}(h \frown a)}{P^{\pi}(h)}=\frac{1}{P^{\pi}(h)} P^{\pi}(h)=1
$$

as desired, proving the claim.
It remains to show $\pi^{\prime} \sim \pi$. Let $\mu$ be any environment. We will show by induction on $h$ that for every history $h, P_{\mu}^{\pi^{\prime}}(h)=P_{\mu}^{\pi}(h)$.

Case 1: $h=\langle \rangle$. Then $P_{\mu}^{\pi^{\prime}}(h)=P_{\mu}^{\pi}(h)=1$.

Case 2: $h=h_{0} \frown o$ for some $o \in \mathcal{O}$. Then

$$
\begin{array}{rlr}
P_{\mu}^{\pi^{\prime}}(h) & =P_{\mu}^{\pi^{\prime}}\left(h_{0} \frown o\right) & \\
& =P_{\mu}^{\pi^{\prime}}\left(h_{0}\right) \mu\left(o \mid h_{0}\right) & \text { (Multiplicative rule) } \\
& =P_{\mu}^{\pi}\left(h_{0}\right) \mu\left(o \mid h_{0}\right) & \text { (Induction) } \\
& =P_{\mu}^{\pi}\left(h_{0} \frown o\right)=P_{\mu}^{\pi}(h) . & \text { (Multiplicative rule) }
\end{array}
$$

Case 3: $h=h_{0} \frown a$ for some $a \in \mathcal{A}$.
Subcase 3.1: $P^{\pi}\left(h_{0}\right)=0$. Then also $P^{\pi}\left(h_{0} \frown a\right)=0$, since one cannot randomly generate initial percepts and actions $h_{0} \frown a$ without generating initial percepts and actions $h_{0}$ first. Thus $P_{\mu}^{\pi}\left(h_{0} \frown a\right)=0$ by Lemma 1. Thus

$$
\begin{array}{rlr}
P_{\mu}^{\pi^{\prime}}(h) & =P_{\mu}^{\pi^{\prime}}\left(h_{0} \frown a\right) & \\
& =P_{\mu}^{\pi^{\prime}}\left(h_{0}\right) \pi^{\prime}\left(a \mid h_{0}\right) & \text { (Multiplicative rule) } \\
& =P_{\mu}^{\pi}\left(h_{0}\right) \pi^{\prime}\left(a \mid h_{0}\right) & \text { (Induction) }  \tag{Induction}\\
& =0 \pi^{\prime}\left(a \mid h_{0}\right)=0=P_{\mu}^{\pi}\left(h_{0} \frown a\right)=P_{\mu}^{\pi}(h) . &
\end{array}
$$

Subcase 3.2: $P^{\pi}\left(h_{0}\right) \neq 0$. Then

$$
\begin{array}{rlr}
P_{\mu}^{\pi^{\prime}}(h) & =P_{\mu}^{\pi^{\prime}}\left(h_{0} \frown a\right) \\
& =P_{\mu}^{\pi^{\prime}}\left(h_{0}\right) \pi^{\prime}\left(a \mid h_{0}\right) & \\
& =P_{\mu}^{\pi}\left(h_{0}\right) \pi^{\prime}\left(a \mid h_{0}\right) & \text { (Multiplicative rule) } \\
& =P^{\pi}\left(h_{0}\right) P_{\mu}\left(h_{0}\right) \pi^{\prime}\left(a \mid h_{0}\right) & \text { (Induction) } \\
& =P^{\pi}\left(h_{0}\right) P_{\mu}\left(h_{0}\right) P^{\pi}\left(h_{0} \frown a\right) / P^{\pi}\left(h_{0}\right) & (\text { Lemma 1) } \\
& =P_{\mu}\left(h_{0}\right) P^{\pi}\left(h_{0} \frown a\right) & \left(\text { Definition of } \pi^{\prime}\right) \\
& =P_{\mu}\left(h_{0} \frown a\right) P^{\pi}\left(h_{0} \frown a\right) & \left(\text { Clearly } P_{\mu}\left(h_{0} \frown a\right)=P_{\mu}\left(h_{0}\right)\right) \\
& =P_{\mu}^{\pi}\left(h_{0} \frown a\right)=P_{\mu}^{\pi}(h) . & (\text { Lemma 1) }
\end{array}
$$

(1) Let $\pi^{\prime}$ be the deterministic agent without private memory, defined so that for every agent-seen history $h$ :

- If $P^{\pi}(h) \neq 0$ then $\pi^{\prime}(h)=a$ where $(a, m)=\pi(h) \in \mathcal{A} \times \mathcal{M}$.
- If $P^{\pi}(h)=0$ then $\pi^{\prime}(h)=a_{0}$.

Adopt the following convention: for every action $a$ and agent-seen history $h$, write $\pi(a \mid h)$ for 1 if $\pi(h)=a$ or 0 if $\pi(h) \neq a$. Since $\pi$ is deterministic, clearly if $P^{\pi}(h) \neq 0$ then $P^{\pi}(h)=1$. In that case, if $(a, m)=\pi(h)$, then clearly $P^{\pi}(h \frown$ $a)=1$, and by definition $\pi^{\prime}(h)=a$, so $\pi^{\prime}(a \mid h)=1=P^{\pi}(h \frown a) / P^{\pi}(h)$. On the other hand, if $P^{\pi}(h)=0$, then by definition $\pi^{\prime}(h)=a_{0}$, so that $\pi^{\prime}(a \mid h)=1$ if $a=a_{0}$ and $\pi^{\prime}(a \mid h)=0$ if $a \neq a_{0}$. So altogether,

$$
\pi^{\prime}(a \mid h)= \begin{cases}\frac{P^{\pi}(h \frown a)}{P^{\pi}(h)} & \text { if } P^{\pi}(h) \neq 0 \\ 1 & \text { if } P^{\pi}(h)=0 \text { and } a=a_{0} \\ 0 & \text { if } P^{\pi}(h)=0 \text { and } a \neq a_{0}\end{cases}
$$

exactly as in the proof of (2) and (3) above. Thus, an identical argument as above shows that $\pi^{\prime} \sim \pi$.

Note that the $\pi^{\prime}$ constructed above is much more computationally expensive than $\pi$. If $\mathcal{M}$ is finite, then computing $P^{\pi}(h)$ (in general) "merely" requires computing $\pi$ on a number of private-memory-augmented histories which is exponential in the length of $h$. If $\mathcal{M}$ is infinite, then computing $P^{\pi}(h)$ (in general) requires computing $\pi$ infinitely many times and summing infinite series ${ }^{1}$.

## 4 The reinforcement learning special case and generalization of prior work

One special case of agents and environments is reinforcement learning (or RL), in which the percepts generated from an environment include numerical rewards. In RL, agents are considered to perform better or worse depending on their ability or inability to maximize average rewards across the whole space of environments (or some suitable subspace thereof).

At the beginning of Section 2 we assumed a nonempty set $\mathcal{O}$ of percepts. All the results of the paper continue to hold if additional structure is imposed on those percepts, e.g., if we further require that each percept include a numerical reward. Thus, this paper's results automatically apply to RL. Let us impose exactly that requirement for the remainder of this section. For each $o \in \mathcal{O}$, let $r(o)$ denote the numerical reward included in $o$.

The following notation is standard in RL.
Definition 14. If $\mu$ is an environment and $\pi$ is an agent (deterministic or non-deterministic), with or without (deterministic or non-deterministic) private memory, let $V_{\mu}^{\pi}$ denote the expected total reward $\pi$ would obtain from $\mu$ (i.e., the expected value of the sum $r\left(o_{1}\right)+r\left(o_{2}\right)+\cdots$ if we generate actions $a_{1}, a_{2}, \ldots$, percepts $o_{1}, o_{2}, \ldots$, and possibly private memories $m_{1}, m_{2}, \ldots$, as in Section 2), assuming this expected value exists. If not, $V_{\mu}^{\pi}$ is undefined.

For an example in which $V_{\mu}^{\pi}$ might be undefined, take $\mu$ to be an environment which, every other turn, assigns probability $100 \%$ to a percept with reward 1 , and every other turn, assigns probability $100 \%$ to a percept with reward -1 . Any agent interacting with $\mu$ would obtain expected total reward $1-1+1-1+\cdots$, a divergent series.

Proposition 1. Suppose $\pi, \rho$ are as in Definition 13. If $\pi \sim \rho$ then for every environment $\mu, V_{\mu}^{\pi}=V_{\mu}^{\rho}$ (and the left-hand side is defined iff the right-hand side is defined).

[^0]Proof. For each agent $\sigma$ (of whatever kind), let $V_{\mu, n}^{\sigma}$ denote the expected total reward that $\sigma$ would obtain after interacting with $\mu$ until $n$ percepts are generated. This can be computed by considering all possible histories terminating in an $n$th percept: for each such history $h$, multiply the probability $P_{\mu}^{\sigma}(h)$ of that history by its total reward $r(h)$ (defined as the sum of $r(o)$ where $o$ ranges over the percepts in $h$ ), to obtain the expected total reward contributed by $h$ to $V_{\mu, h}^{\sigma}$. Thus,

$$
V_{\mu, n}^{\pi}=\sum_{h} P_{\mu}^{\pi}(h) r(h)
$$

where $h$ varies over the set of histories terminating in an $n$th percept. By identical reasoning, $V_{\mu, h}^{\rho}=\sum_{h} P_{\mu}^{\rho}(h) r(h)$. But $\pi \sim \rho$, so each $P_{\mu}^{\pi}(h)=P_{\mu}^{\rho}(h)$. Thus each $V_{\mu, n}^{\pi}=V_{\mu, n}^{\rho}$. Taking the limit as $n \rightarrow \infty$ proves the proposition.

Proposition 1 and Theorem 1 together justify this paper's title: private memory confers no advantage. Not only do equivalent agents have the same probability of resulting in any given history, but if different percepts come with different numerical rewards, then Proposition 1 shows that equivalent agents obtain the same expected total reward from each environment. Thus, as numerically measured by RL rewards, equivalent agents have the same exact numerical performance. Theorem 1 shows that for any agent which makes use of private memories, there is an equivalent agent which does not make use of private memories. Thus, private memories do not enable out-performance of agents without private memories.

This work generalizes the main result of [2]. The authors of that paper considered sequences $\vec{\pi}=\left(\pi_{1}, \ldots, \pi_{n}\right)$ of non-deterministic RL agents (without private memory), along with weights $\vec{w}=\left(w_{1}, \ldots, w_{n}\right)$ (each $w_{i}>0$, with sum $w_{1}+\cdots+w_{n}=1$ ). The authors showed that for every such weighted distribution, there is a so-called mixture agent (a non-deterministic agent without private memory) $\vec{w} \cdot \vec{\pi}$ with the property that for every environment $\mu$, $V_{\mu}^{\vec{w} \cdot \vec{\pi}}=\vec{w} \cdot\left(V_{\mu}^{\pi_{1}}, \ldots, V_{\mu}^{\pi_{n}}\right)$. Expressed aphoristically: "the performance of the weighted mixture is the weighted mixture of the performances". But it is trivial to construct such an agent if the agent is allowed to have non-deterministic private memories. Assume $|\mathcal{M}|=n$, say $\mathcal{M}=\left\{m_{1}, \ldots, m_{n}\right\}$. Let $\pi$ be the nondeterministic agent, with non-deterministic private memory, defined as follows (where $h$ is a private-memory-augmented history).

- If $h=\langle o\rangle$ then let $\pi(h)$ assign to each pair $\left(a, m_{i}\right) \in \mathcal{A} \times \mathcal{M}$ the probability $\pi\left(a, m_{i} \mid h\right)=w_{i} \pi_{i}(a \mid h)$. In plain English: with probability $w_{i}$, act as $\pi_{i}$ and remember (using private memory $m_{i}$ ) having done so.
- Otherwise, $h$ is of the form $\left(o, a, m_{i}\right) \frown h_{0}$ for some $o \in \mathcal{O}, a \in \mathcal{A}, m_{i} \in \mathcal{M}$. Let $h^{-}$be the history obtained by deleting all private memories from $h$. For each pair $\left(a, m_{j}\right) \in \mathcal{A} \times \mathcal{M}$, let $\pi(h)$ assign the probability $\pi\left(a, m_{j} \mid h\right)=$ $\pi_{i}\left(a \mid h^{-}\right)$if $j=i$, or probability $\pi\left(a, m_{j} \mid h\right)=0$ otherwise. In plain English: act as $\pi_{i}$, where $i$ is the agent which you previously remembered (and repeat said memory).

Informally, for its initial turn, $\pi$ randomly chooses one of the agents $\pi_{1}, \ldots, \pi_{n}$ (using the given weights), plays as that agent on that turn, and commits (using private memory) to play as that agent thereafter. Clearly for every environment $\mu, V_{\mu}^{\pi}=\vec{w} \cdot\left(V_{\mu}^{\pi_{1}}, \ldots, V_{\mu}^{\pi_{n}}\right)$. Proposition 1 and Theorem 1 together show there is a non-deterministic agent $\pi^{\prime}$ without private memory with the same performance. Thus, the main result from [2] is a special case of this paper.

## 5 Further generalization

The results of this paper could be further generalized to agents who, in response to input $h$ :

- (Pre-rationalizing) First generate a memory $m$ (depending on $h$ ), and then generate an action which may depend on both $h$ and $m$.
- (Post-rationalizing) First generate an action $a$ (depending on $h$ ), and then generate a memory which may depend on both $h$ and $a$.
- (Pre-and-post-rationalizing) First generate a memory $m_{p r e}$ (depending on $h$ ), then an action $a$ (depending on $h$ and $m_{\text {pre }}$ ), and finally a memory $m_{\text {post }}$ (depending on $h, m_{\text {pre }}$, and $a$ ).

The detailed formalization, however, becomes very verbose, so we leave it, and the corresponding version of Theorem 1 , to the reader.

## 6 Conclusion

Mathematicians and software developers understand the word "function" in different ways. The software developer's "function" can, in the process of computing an output, do various things which the mathematician's function cannot do, for example, reading from or writing to a persistent memory-bank. We typically use the mathematician's "function" when we formalize certain things, because the mathematician's function is easier to reason about and prove things about. But the mismatch could, a priori, lead to inaccurate formalisms. In this paper, we considered agents, with or without such "private memory", in a very general agent-environment framework, and showed that in a technical sense, every agent with such private memory is equivalent to one without. Thus, in some sense, the software developer's agent's ability to use private memory confers no advantage. This partially justifies formalizing said agents as mathematical functions. Of course, software developers' agents can also do various other things that a mathematician's agent cannot do (such as query the system clock, access the computer's camera, or even surf the internet). At least, they can in principlethey might not do so very often in practice. Perhaps it is an advantage of the mathematical formalism that, by formalizing agents as mathematical functions, we automatically rule out things like agents who check the system clock or surf the internet in order to compute their outputs. We also rule out private memory, but as this paper shows, that's not a big loss.

## Acknowledgements

We gratefully acknowledge Arthur Paul Pedersen, Len Du, Oscar Martinez, and the reviewers for feedback and discussion.

## References

1. Samuel Allen Alexander, Michael Castaneda, Kevin Compher, and Oscar Martinez. Extending environments to measure self-reflection in reinforcement learning. Journal of Artificial General Intelligence, 13(1):1-24, 2022.
2. Samuel Allen Alexander, David Quarel, Len Du, and Marcus Hutter. Universal agent mixtures and the geometry of intelligence. In AISTATS. PMLR, 2023.
3. James Henry Bell, Linda Linsefors, Caspar Oesterheld, and Joar Skalse. Reinforcement learning in Newcomblike environments. In NeurIPS, 2021.
4. Marcus Hutter. Universal artificial intelligence: Sequential decisions based on algorithmic probability. Springer, 2004.
5. Marcus Hutter. Discrete MDL predicts in total variation. Advances in Neural Information Processing Systems, 22, 2009.
6. Shane Legg and Marcus Hutter. Universal intelligence: A definition of machine intelligence. Minds and machines, 17(4):391-444, 2007.
7. Antonin Raffin, Ashley Hill, Maximilian Ernestus, Adam Gleave, Anssi Kanervisto, and Noah Dormann. Stable Baselines3. https://github.com/DLR-RM/stablebaselines3, 2019.

[^0]:    ${ }^{1}$ And if $\pi$ is sophisticated enough that the machinery on which $\pi$ is run somehow experiences consciousness when $\pi$ is computed on various inputs, then there could potentially even be ethical implications about plugging so many inputs into $\pi$ in order to compute $\pi^{\prime}$. This was perhapse foreshadowed by Nietzsche's doctrine of the eternal return.

