




Modeling the effects of environmental and perceptual uncertainty using deterministic reinforcement learning dynamics with partial observability

Wolfram Barfuss ^{1,2,*} and Richard P. Mann ²

¹*Institute for Theoretical Physics, University of Tübingen, 72076 Tübingen, Germany*

²*Department of Statistics, School of Mathematics, University of Leeds, Leeds LS2 9JT, United Kingdom*

 (Received 25 August 2021; revised 14 January 2022; accepted 24 February 2022; published 21 March 2022)

Assessing the systemic effects of uncertainty that arises from agents' partial observation of the true states of the world is critical for understanding a wide range of scenarios, from navigation and foraging behavior to the provision of renewable resources and public infrastructures. Yet previous modeling work on agent learning and decision-making either lacks a systematic way to describe this source of uncertainty or puts the focus on obtaining optimal policies using complex models of the world that would impose an unrealistically high cognitive demand on real agents. In this work we aim to efficiently describe the emergent behavior of biologically plausible and parsimonious learning agents faced with partially observable worlds. Therefore we derive and present deterministic reinforcement learning dynamics where the agents observe the true state of the environment only partially. We showcase the broad applicability of our dynamics across different classes of partially observable agent-environment systems. We find that partial observability creates unintuitive benefits in several specific contexts, pointing the way to further research on a general understanding of such effects. For instance, partially observant agents can learn better outcomes faster, in a more stable way, and even overcome social dilemmas. Furthermore, our method allows the application of dynamical systems theory to partially observable multiagent learning. In this regard we find the emergence of catastrophic limit cycles, a critical slowing down of the learning processes between reward regimes, and the separation of the learning dynamics into fast and slow directions, all caused by partial observability. Therefore, the presented dynamics have the potential to become a formal, yet practical, lightweight and robust tool for researchers in biology, social science, and machine learning to systematically investigate the effects of interacting partially observant agents.

DOI: [10.1103/PhysRevE.105.034409](https://doi.org/10.1103/PhysRevE.105.034409)

I. INTRODUCTION

We do not observe the world as it is but instead as our limited sensory and cognitive apparatus perceives it. There are always elements of the world that are hidden from us, such as the detailed physical state of our environment and the internal states of other agents. As such uncertainty is a fundamental feature of life. To be more specific, we might not know what will happen (*stochastic uncertainty*), what currently is (*state uncertainty*), and what others are going to do (*strategic uncertainty*), among other forms of uncertainty [1–3]. In common with other animals, we must learn and make decisions amid this uncertainty using the limited cognitive resources available to us. So must everybody else.

Given the cognitive demands of fully integrating all sources of uncertainty when learning from experience and making decisions, real agents must employ methods of bounded rationality [4] that use cognitive resources efficiently to obtain acceptable solutions in a timely manner [5]. As such, evolutionary game theory [6] takes into account *strategic uncertainty* by assuming that other agents are not perfectly rational but instead by allowing agents to adapt to each other sequentially, with relatively successful strategies be-

ing reinforced and less successful strategies selected against. Tools and methods from evolutionary game theory have also been used successfully to formally study the dynamics of multiagent reinforcement learning [7,8]. Börgers and Sarin [9] established the formal relationship between the learning behavior of one of the most basic reinforcement learning schemes, Cross learning [10], and the replicator dynamics of evolutionary game theory. Since then, this approach of evolutionary reinforcement learning dynamics has been extended to stateless Q -learning [11,12], regret-minimization [13], and temporal-difference learning [14], as well as discrete-time dynamics [15], continuous strategy spaces [16], and extensive-form games [17]. This learning dynamic approach offers a formal, lightweight, and deterministically reproducible way to gain improved, descriptive insights into the emerging multiagent learning behavior.

Apart from strategic uncertainty, representing *stochastic uncertainty*, i.e., uncertainty about what will happen in the form of probabilistic events within the environment, requires foremost the presence of an environment. Recent years have seen a growing interest to move evolutionary and learning dynamics in stateless games to changing environments. Here the term environment can mean external fluctuations [18,19], a varying population density [20,21], spatial network structure [22,23], or coupled systems out of evolutionary and environmental dynamics. Coupled systems may further be

*wolfram.barfuss@uni-tuebingen.de

categorized into those with continuous environmental state spaces [24–28] or discrete ones [14,29–31]. We will focus on learning dynamics in stochastic games [14,29] which encode stochastic uncertainty via action-dependent transition probabilities between environmental states.

However, all dynamics discussed so far are either applicable only to stateless environments, assume that agents do not tailor their response to the current environmental state, or if they do, assume that agents observe the true states of the environment perfectly. Yet often in real-world settings state observations are noisy and incomplete. Thus, they are lacking a systematic way to describe interacting agents under *state uncertainty*.

In this work, we relax the assumption of perfect observations and introduce deterministic reinforcement learning dynamics for partially observable environments. With the derived dynamics we are able to study the idealized reinforcement learning behavior in a wide range of environmental classes, from partially observable Markov decision processes [POMDPs, 32], decentralized (Dec-)POMDPs [33], and fully general partially observable stochastic games [34].

Note that while a great deal of works on partially observable decision domains is of normative nature, ours is descriptive. For the normative agenda, agents are often enriched with, e.g., generative models and belief-state representations [32,33], abstractions [35], or predictive state representations [36] in order to learn optimal policies in partially observable decision domains. Also the economic value of signal is often studied by asking how fully rational agents optimally deal with a specific form of state uncertainty [37]. However, such techniques can become computationally extremely expensive [38]. It is unlikely that biological agents perform those elaborate calculations [39] and the focus on unboundedly rational game equilibria lacks a dynamic perspective [40], making it unable to answer which equilibrium (of the often many) the agents select.

Instead, this work takes a dynamical systems perspective on individual learning agents employing the widely occurring principle of temporal-difference reinforcement learning [41] in which the agents simply treat their observations as if they were the true states of the environment. Temporal-difference learning is not only a computational technique [42], it also occurs in biological agents through the dopamine reward-prediction error signal [43,44]. We focus on agents which employ either so-called memoryless policies, at which they choose their actions based solely on their current observation [45] or they use a short and fixed history of current and past observations and actions to base the current action on. This has the advantage of being simple to act on [46] and they are easy to realize at no or little additional computational cost.

To highlight the broad applicability of our dynamics we study the emerging learning behavior across five partially observable environment classes. We find a variety of effects caused by partial observability which generally depend on the environment and its representation. For instance, partial observability can lead to better learning outcomes faster in a single-agent renewable resource harvesting task, stabilize a chaotic learning process in a multistate zero-sum game and even overcome social dilemmas. Compared to fully obser-

vant agents, partially observant learning often requires more exploration and less weight on future rewards to obtain the favorable learning outcomes. Furthermore, our method allows the application of dynamical systems theory to partially observable multiagent learning. We find that partial observability can cause the emergence of catastrophic limit cycles, a critical slowing down of the learning processes between reward regimes and the separation of the learning dynamics into fast and slow eigendirections. We hope that the presented dynamics become a practical, lightweight, and robust tool to systematically investigate the effect of uncertainty of interacting agents.

II. BACKGROUND

A. Partially observable stochastic games

1. Definition

The game $G = \langle N, \mathcal{S}, \underline{\mathcal{A}}, \underline{\mathcal{O}}, T, \underline{R}, \underline{Q} \rangle$ is a stochastic game with $N \in \mathbb{N}$ agents. The environment consists of $Z \in \mathbb{N}$ states $\mathcal{S} = (S_1, \dots, S_Z)$. In each state s , each agent $i \in \{1, \dots, N\}$ has $M \in \mathbb{N}$ available actions $\mathcal{A}^i = (A_1^i, \dots, A_M^i)$ to choose from. $\underline{\mathcal{A}} = \otimes_i \mathcal{A}^i$ is the joint-action set and agents choose their actions simultaneously. A joint action is denoted by $\underline{a} = (a^1, \dots, a^N) \in \underline{\mathcal{A}}$. With $\underline{a}^{-i} = (a^1, \dots, a^{i-1}, a^{i+1}, \dots, a^N)$ we denote the joint action except agent i 's. We chose an identical number of actions for all states and all agents out of notational convenience. Throughout this paper, we restrict ourselves to ergodic environments without absorbing states.

The transition function $T : \mathcal{S} \times \underline{\mathcal{A}} \times \mathcal{S} \rightarrow [0, 1]$ determines the probabilistic state changes. $T(s, \underline{a}, s')$ is the transition probability from current state s to next state s' under joint action \underline{a} .

The reward function $\underline{R} : \mathcal{S} \times \underline{\mathcal{A}} \times \mathcal{S} \rightarrow \mathbb{R}^N$ maps the triple of current state s , joint action \underline{a} and next state s' to an immediate reward scalar for each agent. $R^i(s, \underline{a}, s')$ is the reward agent i receives.

Instead of observing the states $s \in \mathcal{S}$ directly, each agent i observes one of $Q \in \mathbb{N}$ observations $\mathcal{O}^i = (\mathcal{O}_1^i, \dots, \mathcal{O}_Q^i)$ according to the observation functions $\mathcal{O}^i : \mathcal{S} \times \underline{\mathcal{O}}^i \rightarrow [0, 1]$. $\mathcal{O}^i(s, o)$ is the probability that agent i observes observation $o \in \mathcal{O}^i$ given that the environment is in state $s \in \mathcal{S}$. $\underline{\mathcal{O}} = \otimes_i \mathcal{O}^i$ is the joint observation set and $\underline{\mathcal{O}} = \otimes_i \mathcal{O}^i : \mathcal{S} \times \underline{\mathcal{O}} \rightarrow [0, 1]^N$ is the joint observation function. We chose an identical number of observations for all agents out of notational convenience. By construction, this observation function can model both noisy state observations ($Q = Z$) and hidden states ($Q < Z$).

2. Policies

We consider agents that choose their actions probabilistically according to their memoryless policy $X^i : \mathcal{O}^i \times \mathcal{A}^i \rightarrow [0, 1]$. $X^i(o^i, a^i)$ is the probability that agent i chooses action a^i given that it observed observation o^i . We denote the joint policy by $\underline{X} = \underline{X}(\underline{o}, \underline{a}) = \otimes_i X^i(o^i, a^i) : \underline{\mathcal{O}} \times \underline{\mathcal{A}} \rightarrow [0, 1]^N$.

3. Histories

Besides memoryless policies we also consider policies with fixed histories \mathcal{H}_h of type h . The type h is composed

of $\mathbf{h} = \mathbf{h}_o \times \mathbf{h}_a$ with $\mathbf{h}_o \in \mathbb{N}$ and $\mathbf{h}_a \in \mathbb{N}^N$. \mathbf{h}_o represents how many of current and past observations are to be used to encode the histories. Likewise, \mathbf{h}_a represents how many past actions of each agent are to be encoded in the histories. For example, the default memoryless policy is of type $\mathbf{h} = (1, \mathbf{0})$. Practically, histories induce an embedding of the game into a larger state space at which the histories \mathcal{H}_h correspond to the larger state set and transitions, rewards and observations are adjusted accordingly.

B. Temporal-difference reinforcement learning

Temporal-difference Q -learning is one of the most widely studied reinforcement learning processes [42,43,47]. Agents successively improve their evaluations of the quality of the available actions. Originally developed under the assumption that agents can observe the true Markov state of the environment, we here present the basic temporal-difference Q -learning algorithm in the more general formulation, where agents use observations instead of states. When observations exactly map onto the states, the original algorithm is recovered.

At time step t agent i evaluates action a^i at observation o^i to be of quality $Q_t^i(o^i, a^i)$. Those state-action values $Q_t^i(o^i, a^i)$ are then updated after selecting action a_t^i after observing observation o_t^i according to

$$Q_{t+1}^i(o_t^i, a_t^i) = Q_t^i(o_t^i, a_t^i) + \alpha \cdot \delta_t^i(o_t^i, a_t^i), \quad (1)$$

with the temporal-difference error

$$\delta_t^i(o_t^i, a_t^i) := (1 - \gamma)r_t^i + \gamma \max_b Q_t^i(o_{t+1}^i, b) - Q_t^i(o_t^i, a_t^i). \quad (2)$$

The *discount factor* parameter $\gamma \in [0, 1)$ regulates how much the agent cares for future rewards. The *learning rate* parameter $\alpha \in (0, 1)$ regulates how much new information is used for an observation-action-value update. For the sake of simplicity, we assume identical parameters across agents throughout this paper and therefore do not equip parameters with agent indices. The variable r_t^i refers to the immediate reward at time step t . Note that the $(1 - \gamma)$ prefactor in front of the reward occurs when we assume that agents aim to maximize a return defined as $G_t^i = (1 - \gamma) \sum_{k=0}^{\infty} \gamma^k r_{t+k}^i$ [14]. This leads the values to be on the same scale as the rewards.

Agents select actions based on the current observation-action values $Q_t^i(o^i, a^i)$ balancing exploitation (i.e., selecting the action of maximum quality) and exploration (i.e., selecting lower quality actions in order to learn more about the environment). We here use the widely used Boltzmann policy function. The probability of choosing action a^i under observation o^i is

$$X_t^i(o^i, a^i) = \frac{e^{\beta Q_t^i(o^i, a^i)}}{\sum_{b \in \mathcal{A}^i} e^{\beta Q_t^i(o^i, b)}}, \quad (3)$$

where the *intensity of choice* parameter β controls the exploration-exploitation trade-off. Throughout this paper, we are interested in the idealized learning process with fixed parameters α , β , and γ throughout learning and evaluating a policy.

III. DERIVATION

In this section we derive the deterministic reinforcement learning dynamics under partial observability in discrete time. As classic evolutionary dynamics operate in the theoretical limit of an infinite population, the learning dynamics are derived by considering an infinite memory batch [48,49]. A learning dynamic update of the current policy uses policy-averages instead of individual samples. Thus, we need to construct the policy-average temporal-difference error $\bar{\delta}^i$ to be inserted in the update for the joint policy,

$$X_{t+1}^i(o^i, a^i) = \frac{X_t^i(o^i, a^i) \cdot \exp[\alpha \beta \bar{\delta}^i(o^i, a^i)]}{\sum_b X_t^i(o^i, b) \cdot \exp[\alpha \beta \bar{\delta}^i(o^i, b)]}. \quad (4)$$

Equation (4) can be derived by combining Eqs. (1) and (3). The bar on top of δ^i indicates implicitly that $\bar{\delta}^i$ depends fully on the current joint policy \underline{X}_t . Computing $\bar{\delta}^i(o^i, a^i)$ involves averaging over policies, environmental transitions, and observations for the first two terms of the temporal-difference error [Eq. (2)], the immediate rewards and the qualities of the next observation. The quality of the current observation, $Q_t^i(o_t^i, a_t^i)$, becomes $\beta^{-1} \ln X_t^i(o^i, a^i)$ in the average temporal-difference error and serves as regularization term. This can be derived by inverting Eq. (3) and realizing that the dynamics induced by Eq. (4) are invariant under additive transformations which are constant in actions.

A. Beliefs

The challenge is that the rewards $R^i(s, \underline{a}, s')$ in the stochastic game model depend on the true states, not on the observations of the agents. Thus, in order to obtain the average observation-action rewards $\bar{R}^i(o^i, a^i)$, we need a mapping from observations to states. The observation function is a mapping from states to observations. With Bayes's rule,

$$\bar{B}^i(o^i, s) = \frac{O^i(s, o^i) \bar{P}(s)}{\sum_s O^i(s, o^i) \bar{P}(s)}, \quad (5)$$

we can transform the observation function into a belief function, following the rules of probability. $\bar{B}^i(o^i, s)$ is the belief of agent i (or simply the probability) that the environment is in state s when it observed observation o^i .

The only problem is how to obtain the policy-average stationary state distribution $\bar{P}(s)$. $\bar{P}(s)$ is the left-eigenvector of the average transition matrix $\bar{T}(s, s')$ where the entry $\bar{T}(s, s')$ denotes the probability of transitioning from state s to state s' . This matrix could be obtained as $\bar{T}(s, s') = \prod_j \sum_{a^j} \bar{Y}^j(s, a^j) T(s, \underline{a}, s')$ if we had the probability for each agent j to choose action a^j in state s , $\bar{Y}^j(s, a^j)$. However, we assumed that agents condition their actions only on observations, $X^j(o^j, a^j)$. Yet whenever the environment is in state s , agent j observes observation o^j with probability $O^j(s, o^j)$ and then chooses action a^j with probability $X^j(o^j, a^j)$. Thus, with

$$\bar{Y}^j(s, a^j) := \sum_{o^j \in \mathcal{O}^j} O^j(s, o^j) X^j(o^j, a^j), \quad (6)$$

we can average out the observation and obtain the policy-average state policies $\bar{Y}^j(s, a^j)$. Note that $\bar{Y}^j(s, a^j)$ are proper conditional probabilities, which can be seen by applying \sum_{a^j} to both sides of Eq. (6). With $\bar{Y}^j(s, a^j)$ we can then

compute the policy-average transition matrix $\bar{T}(\underline{s}, \underline{s})$, its left-eigenvector, the stationary state distribution $\bar{P}(s)$, and thus, the policy-average belief of agent i that the environment is in state s when it observed observation o^j , $\bar{B}^i(o^j, s)$.

B. Rewards

Whenever agent i observes observation o^j , with probability $\bar{B}^i(o^j, s)$ the environment is in state s where all other agents $j \neq i$ behave according to $\bar{Y}^j(s, a^j)$, the environment transitions to a next state s' with probability $T(s, \underline{a}, s')$, and agent i receives the reward $R^i(s, \underline{a}, s')$. Mathematically, the policy-average reward for action a^i under observation o^j reads

$$\bar{R}^i(o^j, a^i) := \sum_s \sum_{a^j} \sum_{s'} \prod_{j \neq i} \bar{B}^i(o^j, s) \bar{Y}^j(s, a^j) \times T(s, \underline{a}, s') R^i(s, \underline{a}, s'). \quad (7)$$

C. Qualities

Second, the policy-average of the quality of the next observation ($\max_b Q_t^i(o_{t+1}^j, b)$) in Eq. (2) is computed by averaging over all states, all actions of the other agents, next states, and next observations. Whenever agent i observes observation o^j , the environment is in state s with probability $\bar{B}^i(o^j, s)$. There, all other agents $j \neq i$ choose their action a^j with probability $\bar{Y}^j(s, a^j)$. Consequently, the environment transitions to the next state s' with probability $T(s, \underline{a}, s')$. At s' , the agent observes observation o' with probability $O^i(s', o')$ and estimates the quality to be of value $\max_b \bar{Q}^i(o', b)$. Mathematically, we write

$$\max_b \bar{Q}^i(o^j, a^i) := \sum_s \sum_{a^j} \sum_{s'} \sum_{o'} \prod_{j \neq i} \bar{B}^i(o^j, s) \bar{Y}^j(s, a^j) \times T(s, \underline{a}, s') O^i(s', o') \max_b \bar{Q}^i(o', b). \quad (8)$$

Here we replace the quality estimates $Q_t^i(o^j, a^i)$, which evolve in time t [Eq. (1)], with the policy-average observation-action quality $\bar{Q}^i(o^j, a^i)$, which is the expected discounted sum of future rewards from executing action a^i at observation o^j and then following along the joint policy \underline{X} . It is obtained by a discount factor weighted average of the current policy-average reward $\bar{R}^i(o^j, a^i)$ and the policy-average observation quality of the next observation $\bar{V}^i(o')$,

$$\bar{Q}^i(o^j, a^i) = (1 - \gamma) \bar{R}^i(o^j, a^i) + \gamma \sum_{o' \in \mathcal{O}^i} \bar{T}^i(o^j, a^i, o') \bar{V}^i(o'). \quad (9)$$

Here $\bar{T}^i(o^j, a^i, o')$ is agent i 's policy-average transition probability of observing observation o' at the next time step given it observed observation o^j at the current time step and chose action a^i . It is computed by averaging over all states, next states and all actions of the other agents. Whenever agent i observes observation o^j and selects action a^i , the environment is in state s with probability $\bar{B}^i(o^j, s)$, where all other agents $j \neq i$ select action a^j with probability $\bar{Y}^j(s, a^j)$. Consequently, the environment will transition to the next state s' with probability $T(s, \underline{a}, s')$ which is observed with probability $O^i(s', o')$ as o'

by agent i . Mathematically, we write

$$\bar{T}^i(o^j, a^i, o') = \sum_s \sum_{a^j} \sum_{s'} \prod_{j \neq i} \bar{B}^i(o^j, s) \bar{Y}^j(s, a^j) \times T(s, \underline{a}, s') O^i(s', o'). \quad (10)$$

Further at Eq. (9), $\bar{V}^i(o')$ is the policy-average observation quality, i.e., the expected discounted sum of future rewards from observation o' and then following along the joint policy \underline{X} . They are computed via matrix inversion according to

$$\bar{V}^i(o) = (1 - \gamma) [\underline{1}_Q - \gamma \bar{T}^i(o, o)]^{-1} \bar{R}(o). \quad (11)$$

This equation is a direct conversion of the Bellman equation $\bar{V}^i(o) = (1 - \gamma) \bar{R}(o) + \gamma \sum_{o'} \bar{T}^i(o, o') \bar{V}^i(o')$, which expresses that the value of the current observation is the discount factor weighted average of the current reward and the value of the next observation. Underlined observation variables indicate that the corresponding object is a vector or matrix and $\underline{1}_Q$ is a Q -by- Q identity matrix.

$\bar{T}^i(o, o')$ denotes the policy-averaged transition matrix for agent i . The entry $\bar{T}^i(o^j, o')$ indicates the probability that agent i will observe observation o' after observing observation o^j at the previous time step, given all agents follow the joint policy \underline{X} . We compute them by averaging over all states, all actions from all agents and all next states,

$$\bar{T}^i(o^j, o') = \sum_s \sum_{a^j} \sum_{s'} \prod_j \bar{B}^i(o^j, s) \bar{Y}^j(s, a^j) \times T(s, \underline{a}, s') O^i(s', o'). \quad (12)$$

For any observation o^j , $\bar{B}^i(o^j, s)$ is the probability to be in state s , where all agents j act according to $\bar{Y}^j(s, a^j)$. Therefore, the environment transitions with probability $T(s, \underline{a}, s')$ from state s to the next state s' , which is observed by agent i as observation o' with probability $O^i(s', o')$. Note that $\bar{T}^i(o, o')$ is a proper probabilistic matrix. This can be seen by applying $\sum_{o'}$ to both sides of Eq. (12).

Further in Eq. (11), $\bar{R}^i(o^j)$ denotes the policy-average reward agent i obtains from observation o^j . We compute them by averaging over all states, all actions from all agents and all next states. Whenever agent i observes observation o^j , the environment is in state s with probability $\bar{B}^i(o^j, s)$. Here, all agents j choose action a^j with probability $\bar{Y}^j(s, a^j)$. Hence, the environment transitions to the next state s' with probability $T(s, \underline{a}, s')$ and agent i receives the reward $R^i(s, \underline{a}, s')$,

$$\bar{R}^i(o^j) := \sum_s \sum_{a^j} \sum_{s'} \prod_j \bar{B}^i(o^j, s) \bar{Y}^j(s, a^j) \times T(s, \underline{a}, s') R^i(s, \underline{a}, s'). \quad (13)$$

Note that the quality $\max_b \bar{Q}^i(o^j, a^i)$ depends on o^j and a^i although it is the policy-averaged maximum observation-action value of the next observation.

D. Temporal-difference error

All together, the policy-average temporal-difference error, to be inserted into Eq. (4), reads

$$\bar{\delta}^i(o^j, a^i) = (1 - \gamma) \bar{R}^i(o^j, a^i) + \gamma \max_b \bar{Q}^i(o^j, a^i) - \frac{\ln X^i(o^j, a^i)}{\beta}. \quad (14)$$

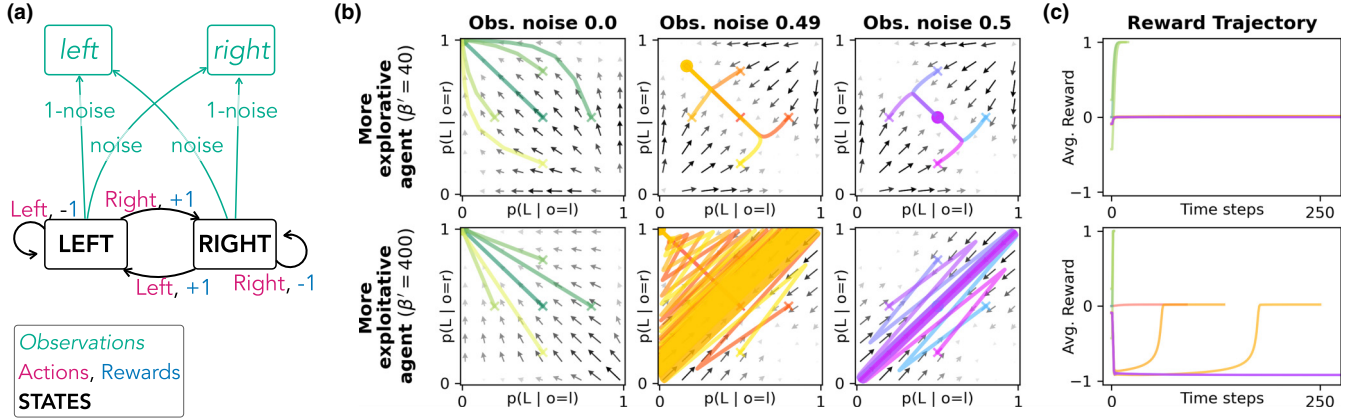


FIG. 1. Deterministic learning dynamics in a simple coordination task. The agent has to anticominate its actions on the environmental state, which it observes through a noisy channel (a). Learning trajectories are computed for three observational noise levels (0.0, 0.49, 0.5) and two intensities of choice, a low one [$\beta' = \beta/(1 - \gamma) = 40$], shown in the top row, and a high one ($\beta' = 400$), shown in the bottom row. Panel (b) shows the corresponding policy spaces in which the agent’s probability of choosing Left, given the agent perceived the environment to be in the left state, is plotted on the x axes; and the agent’s probability of choosing Left, given the agent perceived the environment to be in the right state, is plotted on the y axes. Five individual trajectories, whose initial policies were centered around the center of the policy space, are plotted in color. Arrows in gray indicate the flow of the learning dynamical system. Panel (c) shows the corresponding reward trajectories. Remaining hyperparameters were set as $\alpha = 0.01$, $\gamma = 0.9$. Partial observability can cause the learning to enter low-rewarding limit cycles under high intensity of choice.

IV. EXPERIMENTS

We study the emerging learning dynamics across five test environments: three single-agent decision problems and two multiagent games. Three environments will cover noisy observations, the other two focus on a reduced observation space, where a given observation is consistent with multiple true states of the world. As one evaluation metric we use the average reward, $\sum_s \bar{P}(s)\bar{R}^i(s)$, where $\bar{P}(s)$ is the stationary state distribution and $\bar{R}^i(s) = \sum_{a^j} \sum_{s'} \prod_j \bar{Y}^j(s, a^j)T(s, \underline{a}, s')R^i(s, \underline{a}, s')$ is the average reward for each state given the current policy \underline{X} (see Sec. III). We defined a learning trajectory as having converged if the norm between old and updated policy [according to Eq. (4)] is below 10^{-5} . Since we defined the return with the $(1 - \gamma)$ prefactor we also consider a scaled version of the intensity of choice parameter $\beta = \beta'/(1 - \gamma)$ for some experiments. Doing so preserves the ratio of exploration and exploitation in the temporal-difference error [Eq. (14)] under changes in the discount factor γ .

A. Simple coordination task

1. Environment description

The first environment is a simple coordination task in which the agent must move between the left and right environmental state in order to obtain a maximum reward of 1. Coordinating which of the two available actions (Left, Right) to choose from is complicated by observational noise ν , letting the agent perceive the correct state only with probability $1 - \nu$ [Fig. 1(a)]. This environment is adapted from Singh *et al.* [45].

2. Results

Figure 1 shows how partial observability can cause the deterministic learning dynamics to enter low-rewarding limit

cycles under a high intensity of choice. Often, learning a policy involves a trade-off between the amount of reward from that policy and the amount of time required learning it. In the simple coordination task with perfect observation, a high intensity of choice can speed up the learning process by a factor of 6. The trajectories with $\beta' = 40$ require about 18 time steps to arrive at the optimal policy with average reward 1 (green lines, top row), the trajectories with $\beta' = 400$ require only 3 time steps (green lines, bottom row). Thus, a high intensity of choice is clearly preferable under perfect observation.

With fully uninformative observations [observational noise level $\nu = 0.5$, Fig. 1(b), third column] a more explorative agent (i.e., lower intensity of choice, top row) has an advantage. From all initial policies, it takes the agent about 580 time steps to learn to fully randomize its actions. This yields an average reward of zero and is also the optimal memoryless policy [45]. The more exploitative agent (bottom row) on the other hand enters a limit cycle between choosing Left and Right almost deterministically, irrespective of its current observations. Thus, while choosing Left the agent is trapped in the LEFT state, obtaining an average reward of -1 . While choosing Right, the agent is trapped in the RIGHT state also obtaining an average reward of -1 . The positive reward obtained through the move between states is neglected, since the derived dynamics consider the theoretical limit of an infinite memory batch [48]. This can also be interpreted as a complete separation of the interaction timescale and the adaptation timescale [14,49]. The agent experiences an infinite amount of negative reward during interaction and only one single positive reward after the policy adaptation. It will be interesting to reexamine this scenario under relaxed conditions when interaction and adaption timescales are not completely separated in future work (cf. Ref. [50]).

When observations are almost completely noisy, yet still contain some information about the true environmental state, the more exploitative agent learns a slightly more rewarding policy faster [Fig. 1(b), second column]. An observational noise level of $\nu = 0.49$ means that of 100 times being in the LEFT environmental state, the agent will observe on average *left* 51 times and *right* 49 times. Here, from all initial policies, the more explorative agent (top row) converges to a fixed point in the upper left part of the policy space in about 600 time steps, i.e., slower than under completely noisy observations. This policy yields an average reward of about 0.013. The more exploitative agent (bottom row) learns on an interesting transient resembling the limit cycle of the fully uninformative case, yet manages to converge to the deterministic policy in the upper left of the policy space. This yields an average reward of about 0.02 and takes at most 250 times, depending on the initial policy. This is still distinctly faster than the more explorative agent.

Overall, it is interesting to observe how partial observability caused a well-known dynamical-systems phenomenon in the learning dynamics, which can explain the abrupt improvements in the reward trajectories [Fig. 1(c), bottom]: the separation of the dynamics into a fast eigendirection along the diagonal from the bottom left to the top right of the policy space and a slow eigendirection perpendicular to that [51]. The slow eigendirection corresponds to a coordinated policy where the agent's observation is decisive for its actions. Along the fast eigendirection the agent's policy is independent of its observations. The more explorative agent moves along these axes whereas the more exploitative agent overrides. Yet as long as there is some information in the observations about the environmental state, the more exploitative agent learns better policies faster.

So far we examined only memoryless policies, i.e., policies that condition their choice of action only on the current observation. If the more exploitative agent ($\beta' = 400$) is able to condition its choice of action not only on the current observations but also on its last action, then it learns the optimal policy with an average reward of 1 in at most only two time steps, even under fully uninformative state observations (Fig. 2). The agent learns to alternate between Left and Right. This learned policy and even the whole learning dynamics do not depend on the state-observation, as shown by the straight line trajectories and corresponding dynamical flow arrows in Fig. 2.

As a consistency check we compare the derived deterministic learning dynamics with partial observability to a sample-batch reinforcement learning algorithm, as detailed in Ref. [49] (Fig. 3). The batch learning algorithm collects observation and reward experiences inside a batch of size K while keeping its policy fixed before it then updates its policy using the whole information of the collected batch. This is a widely occurring principle for improved data efficiency and learning stability [52] and is used for example in memory-replay [53] and model-based reinforcement learning [54]. Figure 3 shows that our deterministic theory describes such batch learning approaches well under large batch sizes. Yet the calculation time of the deterministic dynamics was in the order of 100 times faster than the simulation of the algorithms.

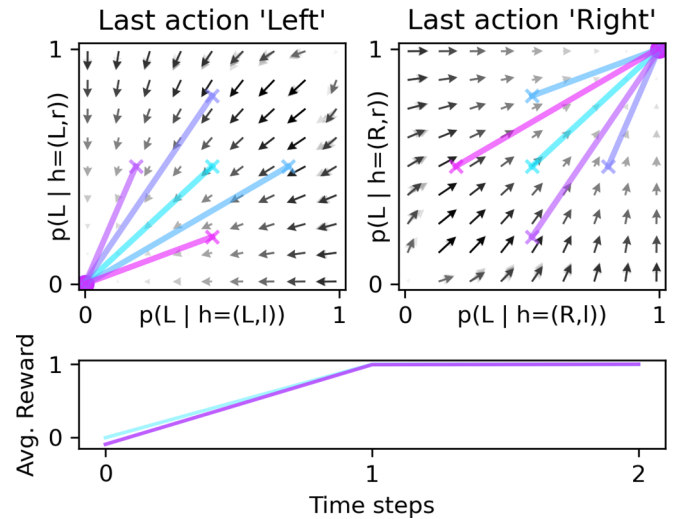


FIG. 2. Deterministic learning dynamics with history in a simple coordination task. Same situation as in Fig. 1 with high intensity of choice $\beta' = 400$ and observational noise level 0.5 but here the agent remembers and conditions its action not only on the last observation but also on its last action. Thus, the left (right) panel shows a projection of the learning dynamics, given the last action was Left (Right). The x axes show the probabilities of choosing Left, given the last environmental observation was *left*. The y axes show the probabilities of choosing Left, given the last environmental observation was *right*. The agent learns to alternate between Left and Right, which yield the highest reward possible, in (at most) only two time steps.

B. Navigation task

1. Environment description

The next environment is the single-agent navigation task adapted from Parr and Russell's Grid World [55]. It consists of 11 states: 6 observations, 4 actions, and 1 agent [Fig. 4(a)].




FIG. 3. Comparison between deterministic learning dynamics and batch-learning algorithm in a simple coordination task with noise level 0.4. The learning trajectory from an initial policy of choosing Left with 30% given both observations is shown for a batch reinforcement learning algorithm with batch sizes 128 (blue; sample size 25) and 1024 (cyan; sample size 5) as well as for the deterministic theory (red); in policy space (left) and versus learning steps (right). On the right only every fourth learning step of the batch learners mean policy values are plotted. Error bars indicate the standard deviation. Hyperparameters are $\alpha = 0.01$, $\beta' = 40$, and $\gamma = 0.09$. The deterministic learning dynamics are well approximated by the batch-learning algorithm.




FIG. 4. Deterministic learning dynamics in a navigation task. Learning trajectories under both partial (dashed lines) and full observability (straight lines) are plotted for various hyperparameter combinations from 15 random initial policies each. Panel (a) shows the grid world. The trajectories of the policies’ action probabilities are projected into each observation or state, such that a deterministic policy toward one direction appears at the edge of that direction in the center. Stochastic policies appear inside the patches. Panel (b) shows the corresponding reward trajectories. In panel (c), two hyperparameter grids show the average reward at convergence for an agent with partial and full observability (with independent color scales for each case). The learning rate was set to $\alpha = 0.01$. In contrast to a fully observant agent, neither a high weight on future rewards (large γ) nor a high intensity of choice (large β') leads to the highest reward for a partially observant agent. Instead, the highest reward depends on the mutual combination of the two hyperparameters.

The agent can move north, south, east, and west. If the agent would move into a wall, then it stays on its current patch. The agent wants to reach the patch in the upper right, which is rewarded by a reward of 1. However, entering the patch below is punished by a reward of -1 . In both cases, the episode ends and the agent begins a new episode on one randomly chosen patch out of the nine other patches. All other state-action combinations yield zero reward. We use this environment to compare the effect of various hyperparameter combinations on the learning behavior of an agent with partial observability to an agent with full observability. Under partial observability the agent can only observe whether or not there is a wall east and west of its current patch. Imagine, for example, a robot equipped only with haptic sensors on its sides or an insect with corresponding antennae. With full observability, the agent can distinguish each grid patch separately.

2. Results

In contrast to a fully observant agent, neither a high weight on future rewards (large γ) nor a high intensity of choice (large β') leads to a large reward for a partially observable agent. Instead, the highest reward depends on the mutual combination of the two hyperparameters. For a hyperparameter combination of $\gamma = 0.99$ and $\beta' = 50$ an agent with full observability quickly learns the optimal policy [Figs. 4(a) and 4(b), light-red straight lines]. Observe also how the light-red straight lines in states (3,2) and (3,3) avoid being close to the penalty state. In contrast, less weight on future rewards (a lower discount factor of $\gamma = 0.4$) and more exploration with $\beta' = 20$ lead to a lower average reward at convergence (light blue straight lines). Observe how the convergence points in policy space (light blue dots) are increasingly farther apart from the optimal policy (light red dots) the more steps the grid cell is away from the goal. However, when we turn to the agent with partial observability, it is the other way around. Here less weight on future rewards and more exploration

lead to a better average reward at convergence (dark colored dashed lines). This result can be explained as follows. In a fully observable Markov decision process there is always an optimal deterministic policy [56]. See how the light red straight lines converges to edge of most grid cells, indicating a deterministic action in that direction [Fig. 4(a)]. Yet policies in a partially observable Markov decision process often require stochasticity [45]. More exploration directly ensures that, although not in a reward-targeted way. Less weight on future rewards might be advantageous under partial observability since too much weight on too distant rewards in the future cannot pay off when there is a fundamental uncertainty about which state the agent occupies or even about what the real states of the environment are. When the environment is only partially observable, anticipating too distant rewards does not have to be beneficial.

A systematic analysis of the hyperparameter grid [Fig. 4(c)] confirms that partial observability requires the right combination of the two hyperparameters in order to learn the highest reward. Under full observability, simply setting a sufficiently high weight on future rewards γ and a sufficiently strong intensity of choice β (i.e., little exploration) leads to the average reward of the optimal policy. In contrast, for partial observability too much farsightedness and too intense exploitation can hurt the performance of the agent. Instead, the optimal reward at convergence is obtained by a more randomly explorative and myopic agent.

With memoryless policies, the average reward of the partially observant agent is smaller by an order of magnitude compared to the fully observant agent. Figure 5 compares the results of the two simplest types of history, i.e., where the agent uses one more piece of information. Thus, the agent conditions its actions either on the current observation and the last action $h = (1, 1)$ or on the current and last observations $h = (2, 0)$. Both types of history are able to obtain a similar maximum average reward, with a slight advantage for history $h = (1, 1)$. Although both are of the simplest type of history




FIG. 5. Hyperparameter grids for the deterministic learning dynamics with history in the navigation task. Two types of histories are compared. The left plots show results for an agent which conditions its action on the current observation and last action $h = (1, 1)$. The plots on the right show results for an agent which conditions its action on the current and last observations $h = (2, 0)$. The top plots show the average reward at convergence, the bottom plots the time steps to convergence, each on the same color scale. Results are averaged over 15 Monte Carlo runs from random initial policies. The learning rate was set to $\alpha = 0.01$. Both types of history obtain similar maximum average reward, but at different hyperparameter combinations.

conceivable the difference in maximum reward between the partially observant agent and fully observant agent is already halved, compared to the partially observant agent without history.

Also, the set of hyperparameter combinations that obtain a high average reward is shifted to the lower right corner of the parameter space where also the fully observant agent obtains its maximum. Interestingly, though, the set of high rewarding hyperparameter combinations is not identical across the two types of history. The action-dependent history $[h = (1, 1)]$ performs best with a high weight on future rewards γ and more exploration, whereas the two-observation history $[h = (2, 0)]$ obtains the highest rewards by more exploitation across a wider range of weights on future rewards γ .

Furthermore, the learner experiences another dynamical systems phenomenon: a critical slowing down of its learning dynamics [51] before a hyperparameter bifurcation into the high rewarding regime (Fig. 5, bottom row). In the area around the hyperparameter regions which obtain high average reward (yellow area in the plots in the top row), the number of time steps it takes the agent to converge is distinctly

higher compared to other hyperparameter regions. Interestingly, this effect is absent in the memoryless learner of Fig. 4 (not shown). Utilizing such dynamical systems phenomena have the potential to improve the efficiency of hyperparameter search.

C. Renewable resource harvesting

1. Environment description

Harvesting a renewable resource is a foundational challenge in environmental economics and the Earth and sustainability science [57–60]. Here we use a standard logistic growth model, in which the (continuous) resource stock $\tilde{s}_{t+1} = \tilde{s}_t + r\tilde{s}_t(1 - \tilde{s}_t/C)$ first regrows exponentially with rate $r \in \mathbb{R}$ until it saturates at capacity $C \in \mathbb{N}$. In order to turn the stock-continuous logistic growth into a state-discrete Markov decision process, we discretize the continuous resource stock into the environmental states $s \in \{0, \dots, C - 1\}$. The agent has three possible actions: *harvest nothing*, *harvest a small amount*, or *harvest a large amount*. What is small and large depends on the maximum amount, Δs_{\max} , the resource regrows from environmental states \mathcal{S} . The small harvest amounts to $(1 - \Delta E)\Delta s_{\max}$, the large harvest amounts to $(1 + \Delta E)\Delta s_{\max}$, with ΔE representing the deviation in the agent’s harvesting effort.

State transitions work as follows: The harvest amount is subtracted from the current stock state s_t . The stock regrows according to the logistic growth equation, yielding a new hypothetical stock \tilde{s}_{t+1} . In order to avoid the complete depletion of the resource, the minimum hypothetical stock yields a value proportional to a base level \tilde{s}_{base} . Since the agent should have an influence on the regrowth of the resource, \tilde{s}_{base} is multiplied by $(1 + \Delta E)$ if the agent chose to *harvest nothing*, by $(1 - \Delta E)$ if the agent chose to *harvest a little*, and by 0 if the agent chose to *harvest a lot*. The resource stock is then discretized by a normal distribution around \tilde{s}_{t+1} with variance σ^2 . The probability mass that lies between stock $s_{t+1} - 0.5$ and $s_{t+1} + 0.5$ gives the probability to transition to the new state s_{t+1} . (For $s_{t+1} = 0$ the lower bound is $-\infty$, for $s_{t+1} = C$ the upper bound is $+\infty$.) Thus, σ represents the level of stochasticity within the environmental dynamics.

The rewards are identical to the harvest amount. Harvesting a lot yields a higher immediate reward than harvesting a little. Except when the resource is degraded, i.e., either the current state s_t or the next state s_{t+1} equals zero, then the rewards are only 10% of the harvest amount. Thus, the agent has always an immediate incentive to harvest more over a little. The optimal policy depends on the weight the agent puts on future rewards (by its discount factor γ).

We use this environment to showcase how partial observability can be used to investigate the effect of different (imperfect) representations of the environment. We focus on representations under which the agent perceives several adjacent states as a single coherent observations. Figures 6(a) and 6(b) illustrate the renewable resource harvesting environment and the investigated observation representations for capacity $C = 5$.

2. Results

We find that inaccurate (reduced complexity) representations of the environment can lead to a better learning outcome

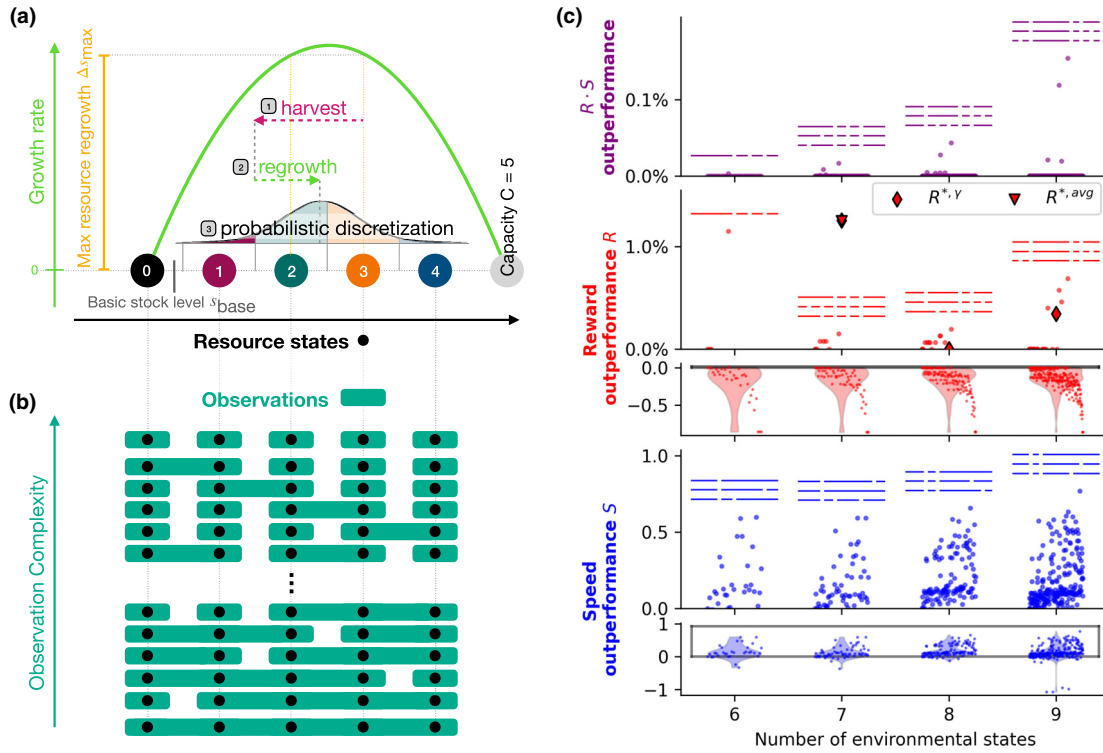


FIG. 6. Deterministic learning dynamics in renewable resource environments. Panel (a) sketches the functioning of the renewable resource harvesting environment. (1) The agent decides on a harvest, which is subtracted from the current environmental stock. (2) The stock regrows according to a logistic function. (3) The stock is discretized by a normal distribution in order to have the number of states equaling the capacity C of the logistic growth function. We set the growth rate $r = 0.8$, the effort deviation $\Delta E = 0.2$, the stock base level $\bar{s}_{\text{base}} = 0.1$, and the environmental stochasticity $\sigma = 0.5$. Panel (b) shows the possible observation spaces (how the environment is represented by the agent)- ordered by decreasing complexity, for a world in which there are five possible true environmental states. In the most complex (at the top) the agent perceives all real states of the world as distinct; in the least complex (at the bottom), the agent makes the same observation regardless of the true state. We investigate all representations where the agent perceives several adjacent states as a single coherent observation. Panels (c) shows the reward out-performance $R = r/r_{\text{ac}} - 1$ (red), the speed out-performance $S = 1 - l/l_{\text{ac}}$ (blue), and the combined reward-speed out-performance $R \cdot S$ (if $R > 0 \wedge S > 0$) (purple) for all possible representations, for the four renewable resource environments with capacities C and likewise number of states, 6–9. Out-performance is measured with respect to the agent which used the accurate representation of the environment and obtained a reward r_{ac} in l_{ac} time steps. For each environment, each dot represents the average of 100 Monte Carlo simulations from random initial policies of a single representation, ordered from the most complex, i.e., the accurate one, on the left to the simplest, i.e., perceiving all states as one, on the right. Violin plots show the distribution of rewards and speed, relative to the agent with the accurate representation. The three top performing representations are shown schematically by the dashed lines. Additionally, the average rewards of the optimal discounted policy $R^{*,\gamma}$ and the optimal average-reward policy $R^{*,\text{avg}}$ are shown. The agent’s discount factor $\gamma = 0.9$, intensity of choice $\beta' = 25$, and learning rate $\alpha = 0.02$. There exist inaccurate representations (partial observation functions) of the environment that lead to a better learning outcome faster compared to the fully observant agent.

faster, when compared to an agent which perceives the environment accurately (at full complexity) (Fig. 6).

In the majority of cases an inaccurate representation of the environment leads to a speed out-performance in the order of 10%, i.e., a smaller number of time steps it takes the learner to converge to a fixed point. Only four representations of the 9-state environment take distinctly longer to converge. Overall, there is a slight tendency that simpler representations lead to faster convergence. Representations (dots) are ordered from the most complex, i.e., the accurate one, on the left to the simplest, i.e., perceiving all states as one, on the right (per environment). All top speed representations (dashed bars) cluster the resource stock 0 and 1 together but separate between stock 1 and 2. In the environments with capacity 8 and

9, a resource stock of 2 is represented completely separate by all top speed representations.

In contrast, the majority of inaccurate representations lead to a worse reward at convergence. Clearly visible by the red dots on the right for each environment, the simpler the representation the worse the performance. Nevertheless, a few representations of intermediate complexity lead to a reward out-performance in the order of 1%. This is remarkable, since Blackwell [61]’s theorem showed that a rational decision maker cannot improve by an inaccurate representation. Of course, our result does not contradict Blackwell, since we investigate a learning process.

To better understand the relationship between the learning process and the rational optimal policies, Table I shows the

TABLE I. Average reward of the optimal average-reward policy $R^{*,\text{avg}}$ and the optimal policy of the discounted reward setting $R^{*,\gamma}$ for the same four renewable resource environments as in Fig. 6. Rewards are also transformed in the same way ($R = r/r_{\text{ac}} - 1$, with r_{ac} being the reward the fully observant agent obtained at convergence).

Env. states	6	7	8	9
Reward $R^{*,\text{avg}}$	0.25	0.013	0.019	0.024
Reward $R^{*,\gamma}$	-0.015	0.013	0	0.003

average reward of the optimal policy $R^{*,\gamma}$ and the average-reward optimal policy $R^{*,\text{avg}}$ relative to the reward obtained by the fully observant agent (shown in Fig. 6 by diamonds and down-triangles). The optimal policy maximizes the state values for each state and depends on the discount factor γ . The average-reward optimal policy maximizes the average reward. Since in this environment $R^{*,\gamma}$ approaches $R^{*,\text{avg}}$ under $\gamma \rightarrow 1$, the rewards between $R^{*,\text{avg}}$ and $R^{*,\gamma}|_{\gamma=0.9}$ represent the rewards more patient or future caring agents could obtain. Thus, the out-performing representations cause the learner to behave as if it were more patient or future-oriented than it actually is (defined by its discount factor γ). However, it is not obvious to identify regularities across the environments between the top rewarding representations. Moreover, Table I shows that the learning process under full observability yields decent results. For the environment with 8 states the learner obtains the exact same reward as the optimal policy. In the environment with 6 states the learner obtains an average reward which is even above the one of the optimal policy.

Taken together, a speed out-performance in the order of 10% multiplied by a reward out-performance in the order of

1% leads to combined speed-reward out-performance in the order of 0.1%. Along the four environments investigated, the magnitude in out-performance is increasing with the number of environmental states. Future work is needed to investigate this effect in larger, more complex resource harvesting environments and also how to obtain those representations which lead to better outcomes faster.

Notably, this result resembles the one by Mark *et al.* [62] who show also that simpler views on the world can be of advantage. However, in their model perceiving the truths comes with a cost which is subtracted from the rewards of the environment. If this cost parameter is sufficiently large, then perceiving the truths cannot pay off by design. We do not model such a cognitive cost of being close to the truths and still find that some inaccurate representations lead to better outcomes faster.

D. Uncertain social dilemma

1. Environment description

The emergence of cooperation in social dilemmas is another key research challenge for evolutionary biology and the social and sustainability sciences [63–65]. We focus on the situation where two agents can either cooperate (C) or defect (D) and either face a Prisoner’s Dilemma or a Stag Hunt game with equal probability [Fig. 7(a)] (cf. Refs. [66,67]). In the pure Prisoner’s Dilemma defection is the Nash equilibrium, which leads to a suboptimal reward for both agents, also known as the tragedy of the commons [68]. In the pure Stag Hunt game, both mutual cooperation and mutual defection are Nash equilibria with the difference that mutual cooperation yields a higher reward than mutual defection for both agents. It is therefore also referred to as a coordination challenge [69].

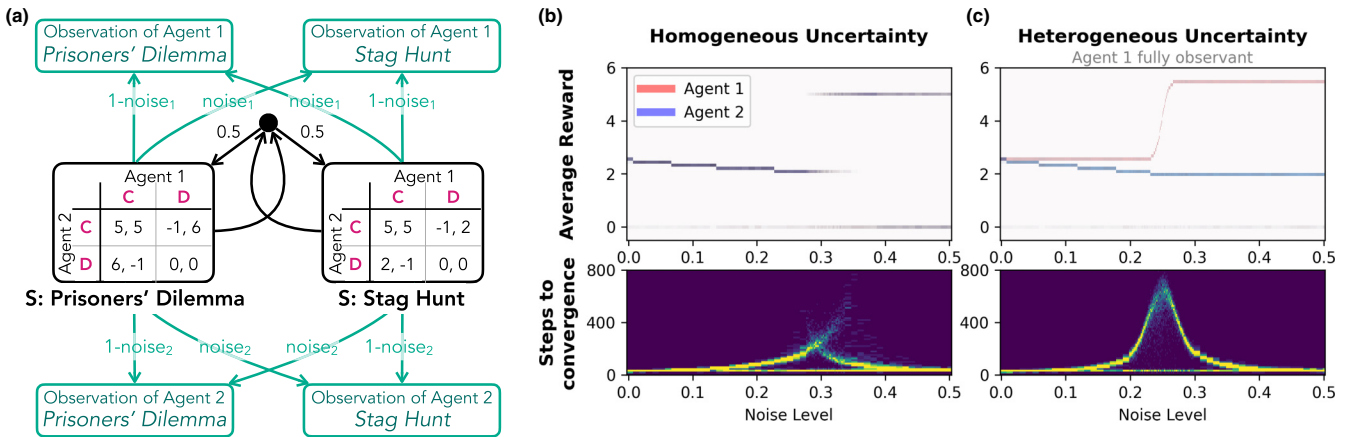


FIG. 7. Deterministic learning dynamics in an uncertain social dilemma. Panel (a) illustrates the environment. Panels (b) and (c) show the average rewards at convergence for agent 1 in red and agent 2 in blue (top row) and the time steps it takes the learners to convergence (bottom row) for various observational noise levels from 0 to 0.5. For each noise level, the plots show a histogram via the color scale. Each histogram results from a Monte Carlo simulation from 100 random initial policies. Panel (b) shows the case of homogeneous uncertainty where both agents’ observations are corrupted equally by noise. In Panel (c) only agent 2 is increasingly unable to observe the environment correctly (heterogeneous uncertainty). The discount factor was set to $\gamma = 0.5$ since future states are independent of the agents actions, which makes the discount factor irrelevant for the learning in this case. Remaining hyperparameters were set to $\alpha = 0.01$ and $\beta' = 50$. Homogeneous uncertainty can overcome the social dilemma through the emergence of a stable, mutually high rewarding fixed point above a critical level of observational noise. Heterogeneous uncertainty, however, leads to reward inequality. In both cases, the transition is accompanied by a critical slowing down of the convergence speed.

Here, we consider the situation when the agents are uncertain about the type of game they are facing at each decision point. Whether we are facing a tragedy or a coordination challenge is relevant for, e.g., the mitigation of human-caused climate change [70]. We investigate two scenarios. Under homogeneous uncertainty [Fig. 7(b)], both agents' observations are blurred by an increasing level of observational noise. Under heterogeneous uncertainty [Fig. 7(c)], only agent 2's observations become noisier. Since the environment is symmetric under exchanging the roles of the agents, it suffices to explore only one heterogeneous uncertainty scenario.

2. Results

Homogeneous uncertainty can overcome the social dilemma through the emergence of a stable, mutually high rewarding fixed point above a critical level of observational noise. Under perfect observation both agents converge to full defection when observing the Prisoner's Dilemma. When observing the Stag Hunt game it depends on the initial joint policy whether the agents converge to mutual defection or mutual cooperation. Reward values are as such that the defective basin of attraction is comparably small [see the light line at an average reward of 0 in Fig. 7(b)]. Increasing the observational noise level from zero under homogeneous uncertainty will first decrease the average reward at convergence. The agents still converge to the perfect observation policy which leads them to defect when they observe the Prisoner's Dilemma but the situation is actually the Stag Hunt. However, increasing observational noise further eventually leads to a bifurcation [Fig. 7(b)]. Mutual cooperation under both observations becomes a stable fixed point. As a consequence both agents obtain an average reward of 5 at convergence. Interestingly, there seems to be a small range of observational noise at which all three rewards 0, ~ 2 , and 5 are supported by equilibria. For large noise levels only the rewards at 0 and 5 are stable.

Thus, we find that the deterministic learning dynamics under homogeneous partial observability are able to converge to mutually more rewarding policies compared to the perfect observation case. The existence of those equilibria is long known in traditional static game theory [66]. Here we show that our derived dynamics are able to serve as a dynamic micro-foundation for those static equilibria. They correspond to fixed points of the derived learning dynamics, and the transitions between equilibria are again accompanied by the dynamical systems phenomenon of a critical slowing down of the convergence speed [Fig. 7(b), bottom].

However, the mutual benefit of uncertainty vanishes when not all agents' observations are uncertain [Fig. 7(c)]. Under slight uncertainty only the reward of the ill-informed agent (Agent 2 in Fig. 7) decreases. After the bifurcation point under large uncertainty, the ill-informed agent converges to full cooperation under both observations, whereas the well-informed agent still defects in the Prisoner's Dilemma which earns it an average reward of even more than 5. The knowledgeable agent exploits the ill-informed and heterogeneous uncertainty leads to reward-inequality between the agents.

Interestingly, Fig. 7 suggests a difference in the type of phase transition between the policy of mediocre reward at low observational noise levels and the policies at high noise

levels. The phase transition under homogeneous uncertainty seems to be discontinuous and shifted toward greater noise levels whereas the transition under heterogeneous uncertainty seems to be continuous. Investigating the relationship between the learning dynamics, free energy equivalents [49] and phase transitions is a promising direction of future work.

E. Uncertain zero-sum competition

1. Environment description

The last environment we use as a test bed is a two-agent, two-state, two-action zero-sum competition, also known as the two-state matching pennies game [71]. It roughly models the situation of penalty kicks between a kicker and a keeper. Both agents can choose between the *left* and the *right* side of the goal. The keeper agent scores one point if it catches the ball (when both agents have chosen the same action); otherwise, the kicker agent receives one point. The two states of the environment encode which agent is the keeper and which one is the kicker. In state *KeepKick* agent 1 is the keeper and agent 2 is the kicker. In the state *KickKeep* it is the other way around. Agents change roles under state transitions, which depend only on agent 1's actions. When agent 1 selects either *left* as keeper or *right* as kicker both agents will change roles. With symmetrical rewards but asymmetrical state transitions, this two-state zero-sum game presents the challenge of coordinating both agents on playing a mixed strategy with equiprobable actions. Similarly as in Secs. IV A and IV D, the agents' observations of the environmental states are obscured by a noise level ν .

2. Results

Figure 8 shows how partial observability can stabilize the learning process. When both agents observe the environment perfectly the learning dynamics are prone to be unstable, either unpredictably chaotic or on periodic orbits and limit cycles [Fig. 8(a)] [14]. The rewards of agent 1 and 2 are circulating around zero. Under a medium observational noise level of $\nu = 0.25$ the learning dynamics are still unstable. Especially the transient dynamics in the policy space [Fig. 8(b), on the right] appear strange. The average reward trajectory looks damped compared to the fully observant agents. Increasing the observational noise further such that the agents perceive the two environmental states (*KeepKick* and *KickKeep*) as a single observation (*KK*), is able to stabilize the learning process. Interestingly, the flow of the learning dynamics is separated into two half circles directed at the upper half of the line at which agent 1 chooses both actions with equal probabilities. As shown by the gray arrows, the circled flow is on a fast timescale compared to the movement downward to the center of the policy space (which is not reached here within 1000 time steps). At this downward movement, both agents play the different roles of kicker and keeper in equal amounts, since only agent 1 is responsible for the state transitions. Any advantage agent 2 gains from deviating from the equiprobable policy as kicker is balanced by the same amount of disadvantage agent 2 loses as keeper. Thus, the rewards for both agents quickly stabilize at zero.

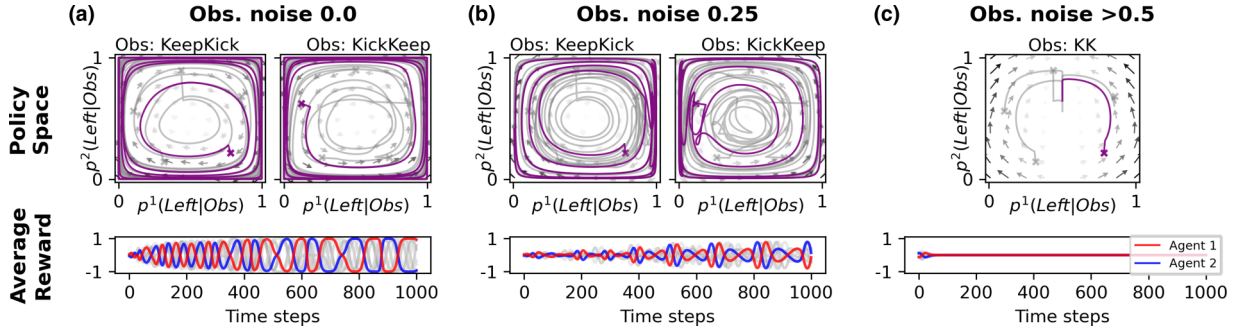


FIG. 8. Deterministic learning dynamics in an uncertain zero-sum competition. Policy spaces and reward trajectories are shown for three different observational noise levels: (a) $\nu = 0.0$, i.e., perfect observation, (b) $\nu = 0.25$, and (c) $\nu > 0.5$, i.e., both states are observed inseparably as one. The probability of choosing action *left*, conditioned on the current observation, is plotted on the *x* axis for agent 1 and on the *y* axis for agent 2. Learning trajectories are shown from five initial policies around the center of the policy spaces. For better visual inspection only one of those trajectories is portrayed in color. Arrows in gray indicate the flow of the learning dynamical system. Hyperparameter were $\alpha = 0.005$, $\beta' = 200$, and $\gamma = 0.9$. Here partial observability is able to stabilize the learning process.

V. DISCUSSION

In this article we analyzed the efficacy of temporal-difference reinforcement learning under irreducible environmental uncertainty. To do so, we introduced deterministic multiagent reinforcement learning dynamics, in which the agents are only partially able to observe the true states of the environment. These dynamics operate in the theoretical limit of an infinite memory batch, make implicit inference about the true states via Bayes's rule and can be well approximated by finite-size batch learning algorithms. This limit allows us to systematically separate the stochasticity of reinforcement learning, resulting from probabilistic environmental dynamics, observations, and decisions, from the environmental uncertainty that originates in the agents' incomplete awareness of the true state space.

Overall, we have shown how these dynamics can serve as a practical, lightweight, deterministically reproducible and robust tool to systematically study the combined effects of *strategic uncertainty*, *stochastic uncertainty*, and *state uncertainty* in collectives of self-learning agents across a wide range of partially observable environment classes.

We have found a variety of effects caused by partial observability, yet general conclusion and recommendations cannot be stated, due to the generality of the partially observable agent-environment setting. Providing agents with only a partial view of the true state of the world might be expected to always result in poorer learning decision-making outcomes. However, we have demonstrated that irreducible environmental uncertainty can instead lead to better learning outcomes, even in a single-agent environment, stabilize the learning process, and overcome social dilemmas in multiagent domains.

Furthermore, our method allows the application of dynamical systems theory to partially observable multiagent learning. We have found that partial observability can cause the emergence of catastrophic limit cycles, within which the agent obtains the worst possible reward. We also found instances where partial observability induces phase transitions between low and high rewarding regimes accompanied by a critical slowing down of the learning processes. Further, we saw partial observability induced separations of the learning

dynamics into fast and slow eigendirections, as well as multi-stability of the learning process.

1. Potential applications

These results may be of use in technological applications of multiagent reinforcement learning, with respect to training regimes, hyperparameter tuning, and the development of novel algorithms. For example, if agents are able to detect that they entered a slow eigendirection, then they can safely increase their learning rate for a faster convergence. Or training regimes and hyperparameter search techniques might be on the lookout for a critical slowing down since this can indicate a phase transition toward high rewarding solutions. With respect to the hyperparameter values required for a decent performance we found across environments that learning with partial observability demands more exploration and less weight on future rewards, compared to fully observant agents. Moreover, the learning with partial observability might depend crucially on the precise combination of the two parameters, whereas without uncertainty both parameters can be tuned fairly independently. The fast computation speed and visualization capabilities of the deterministic learning dynamics approach might be particular suited for the challenge to engineer interpretable and safety-critical learning systems.

We have shown that whether partial observability in the classic principle of temporal-difference learning is advantageous depends on the specific nature of the environment and its representation (cf. ecological rationality) [72,73]. Given that temporal-difference reinforcement learning is a relatively simple and widely effective algorithm, and one which closely matches known features of neurological learning [43,44], this points to a potential evolutionary pressure for agents to develop internal models of the world that do not match the true state space of their environment (cf. Refs. [62,74]). The proposed dynamics are therefore a suitable tool to advance theoretical research in cognitive ecology, which studies how animals acquire, retain, and use information within their ecology, evolution and behavior [75,76]. Within this area, research has begun to ask how agents' may evolve nonveridical or

incomplete representations of the world [62,77,78]; the dynamic model presented here offers a tool to study the effect of nonveridical representations in greater depth.

We also showed that partial observability can lead to better collective outcomes in the case of social dilemmas. The question for the preconditions of cooperation and sustainable behavior presents an important area for deeper investigation [79–81]. Temporal-difference learning is a widespread principle in neuroscience and psychology [44] and there is indeed evidence that humans use a payoff-based learning rule in social dilemmas [82]. The topic of uncertainty is of special relevance in the mitigation of the climate crisis through global cooperation agreements [69,83–86]. Our results highlight the potential for a systematic investigation of mechanisms that incorporate useful uncertainty [87] for learning and adaptive actors. In our examples, the mutual benefit of uncertainty in the social dilemma vanishes when not all agents are likewise ill-informed causing reward-inequality between the agents. This suggests that partial observability as a mechanism for solving social dilemmas may need to be regulated externally (e.g., by authorities that monitor information flow, or as a feature of the environment) rather than something that is likely to be generated as an evolutionary adaptation amongst individuals in competition with each other.

2. Future directions

A promising directions for future work is the integration of *model uncertainty* through an analytical treatment of noisy dynamics (cf. Ref. [50]). The stochastic noise models the finiteness of a reasonable learning algorithm compared to the theoretical limit of the infinite memory batch of the present dynamics. The challenge is that this problem is ill defined and many reasonable learning algorithms exist. Furthermore exciting is the embedding of representation and generalization dynamics into the nonlinear dynamics of learning, acting and environment to study the principles of advantageous representations.

CODE AVAILABILITY

Python code to reproduce all results is available at [88].

ACKNOWLEDGMENTS

This work was supported by UK Research and Innovation Future Leaders Fellowship MR/S032525/1 and the German Federal Ministry of Education and Research (BMBF): Tübingen AI Center, FKZ: 01IS18039A.

-
- [1] M. J. Kochenderfer, *Decision Making under Uncertainty: Theory and Application* (MIT Press, Cambridge, MA, 2015)
 - [2] J. Y. Halpern, *Reasoning about Uncertainty* (MIT Press, Cambridge, MA, 2017).
 - [3] V. A. Marchau, W. E. Walker, P. J. Bloemen, and S. W. Popper, *Decision Making under Deep Uncertainty: From Theory to Practice* (Springer Nature, Berlin, 2019).
 - [4] H. A. Simon, *Models of Bounded Rationality: Empirically Grounded Economic Reason*, Vol. 3 (MIT Press, Cambridge, MA, 1997).
 - [5] T. L. Griffiths, F. Lieder, and N. D. Goodman, Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic, *Top. Cogn. Sci.* **7**, 217 (2015).
 - [6] J. Hofbauer and K. Sigmund, *Evolutionary Games and Population Dynamics* (Cambridge University Press, Cambridge, UK, 1998).
 - [7] D. Bloembergen, K. Tuyls, D. Hennes, and M. Kaisers, Evolutionary dynamics of multi-agent learning: A survey, *J. Artif. Intell. Res.* **53**, 659 (2015).
 - [8] W. Barfuss, Towards a unified treatment of the dynamics of collective learning, in *AAAI Spring Symposium: Challenges and Opportunities for Multi-Agent Reinforcement Learning* (AAAI, Palo Alto, California, USA, 2020).
 - [9] T. Börgers and R. Sarin, Learning through reinforcement and replicator dynamics, *J. Econ. Theory* **77**, 1 (1997).
 - [10] J. G. Cross, A stochastic learning model of economic behavior, *Quart. J. Econ.* **87**, 239 (1973).
 - [11] K. Tuyls, K. Verbeeck, and T. Lenaerts, A selection-mutation model for Q-learning in multi-agent systems, in *International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (2003), pp. 693–700.
 - [12] Y. Sato and J. P. Crutchfield, Coupled replicator equations for the dynamics of learning in multiagent systems, *Phys. Rev. E* **67**, 015206(R) (2003).
 - [13] T. Klos, G. J. Van Ahee, and K. Tuyls, Evolutionary dynamics of regret minimization, in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases* (Springer, Berlin, 2010), pp. 82–96.
 - [14] W. Barfuss, J. F. Donges, and J. Kurths, Deterministic limit of temporal difference reinforcement learning for stochastic games, *Phys. Rev. E* **99**, 043305 (2019).
 - [15] T. Galla and J. D. Farmer, Complex dynamics in learning complicated games, *Proc. Natl. Acad. Sci. U.S.A.* **110**, 1232 (2013).
 - [16] A. Galstyan, Continuous strategy replicator dynamics for multi-agent Q-learning, *Auton. Agents Multi-Agent Syst.* **26**, 37 (2013).
 - [17] F. Panozzo, N. Gatti, and M. Restelli, Evolutionary Dynamics of Q-Learning Over the Sequence form, in *Conference on Artificial Intelligence (AAAI)* (2014), pp. 2034–2040.
 - [18] M. Assaf, M. Mobilia, and E. Roberts, Cooperation Dilemma in Finite Populations Under Fluctuating Environments, *Phys. Rev. Lett.* **111**, 238101 (2013).
 - [19] P. Ashcroft, P. M. Altrock, and T. Galla, Fixation in finite populations evolving in fluctuating environments, *J. R. Soc. Interface.* **11**, 20140663 (2014).
 - [20] C. Hauert, M. Holmes, and M. Doebeli, Evolutionary games and population dynamics: Maintenance of cooperation in public goods games, *Proc. R. Soc. B* **273**, 2565 (2006).
 - [21] C. S. Gokhale and C. Hauert, Eco-evolutionary dynamics of social dilemmas, *Theor. Pop. Biol.* **111**, 28 (2016).
 - [22] C. Gracia-Lázaro, L. M. Floría, J. Gómez-Gardeñes, and Y. Moreno, Cooperation in changing environments: Irreversibility

- in the transition to cooperation in complex networks, *Chaos Solitons Fractals* **56**, 188 (2013).
- [23] A. Szolnoki and X. Chen, Environmental feedback drives cooperation in spatial social dilemmas, *Europhys. Lett.* **120**, 58001 (2017).
- [24] A. Tavoni, M. Schlüter, and S. Levin, The survival of the conformist: Social pressure and renewable resource management, *J. Theor. Biol.* **299**, 152 (2012).
- [25] J. S. Weitz, C. Eksin, K. Paarporn, S. P. Brown, and W. C. Ratcliff, An oscillating tragedy of the commons in replicator dynamics with game-environment feedback, *Proc. Natl. Acad. Sci. U.S.A.* **113**, E7518 (2016).
- [26] X. Chen and A. Szolnoki, Punishment and inspection for governing the commons in a feedback-evolving game, *PLoS Comput. Biol.* **14**, e1006347 (2018).
- [27] A. R. Tilman, J. B. Plotkin, and E. Akçay, Evolutionary games with environmental feedbacks, *Nat. Commun.* **11**, 915 (2020).
- [28] X. Wang and F. Fu, Eco-evolutionary dynamics with environmental feedback: Cooperation in a changing world, *Europhys. Lett.* **132**, 10001 (2020).
- [29] C. Hilbe, Š. Šimsa, K. Chatterjee, and M. A. Nowak, Evolution of cooperation in stochastic games, *Nature (Lond.)* **559**, 246 (2018).
- [30] C. Hauert, C. Saade, and A. McAvoy, Asymmetric evolutionary games with environmental feedback, *J. Theor. Biol.* **462**, 347 (2019).
- [31] Q. Su, A. McAvoy, L. Wang, and M. A. Nowak, Evolutionary dynamics with game transitions, *Proc. Natl. Acad. Sci. U.S.A.* **116**, 25398 (2019).
- [32] M. T. Spaan, Partially observable Markov decision processes, in *Reinforcement Learning: State-of-the-Art* (Springer, Berlin, 2012), pp. 387–414.
- [33] F. A. Oliehoek and C. Amato, *A Concise Introduction to Decentralized POMDPs*, Vol. 1 (Springer, Berlin, 2016).
- [34] E. A. Hansen, D. S. Bernstein, and S. Zilberstein, Dynamic programming for partially observable stochastic games, in *National Conference on Artificial Intelligence (AAAI)* (2004), pp. 709–715.
- [35] R. S. Sutton, E. Rafols, and A. Koop, Temporal abstraction in temporal-difference networks, in *International Conference on Neural Information Processing Systems (NeurIPS)* (2006), pp. 1313–1320.
- [36] M. L. Littman, R. S. Sutton, and S. P. Singh, Predictive representations of state, in *International Conference on Neural Information Processing Systems (NeurIPS)*, Vol. 14 (2001), p. 30.
- [37] A. Bagh and Y. Kusunose, On the economic value of signals, *B.E.J. Theor. Econ.* **20**, 20180191 (2020).
- [38] J. Loch and S. P. Singh, Using eligibility traces to find the best memoryless policy in partially observable markov decision processes, in *International Conference on Machine Learning (ICML)* (1998), pp. 323–331.
- [39] G. Gigerenzer and W. Gaissmaier, Heuristic decision making, *Annu. Rev. Psychol.* **62**, 451 (2011).
- [40] C. Papadimitriou and G. Piliouras, Game dynamics as the meaning of a game, *ACM SIGecom Exch.* **16**, 53 (2019).
- [41] R. S. Sutton, Learning to predict by the methods of temporal differences, *Mach. Learn.* **3**, 9 (1988).
- [42] R. S. Sutton and A. G. Barto, *Reinforcement Learning* (MIT Press, Cambridge, MA, 2018).
- [43] W. Schultz, P. Dayan, and P. R. Montague, A neural substrate of prediction and reward, *Science* **275**, 1593 (1997).
- [44] P. Dayan and Y. Niv, Reinforcement learning: The good, the bad and the ugly, *Curr. Opin. Neurobiol.* **18**, 185 (2008).
- [45] S. P. Singh, T. Jaakkola, and M. I. Jordan, Learning without state-estimation in partially observable markovian decision processes, in *Machine Learning Proceedings 1994* (Elsevier, Amsterdam, 1994), pp. 284–292.
- [46] J. K. Williams and S. P. Singh, Experimental results on learning stochastic memoryless policies for partially observable Markov decision processes, in *International Conference on Neural Information Processing Systems (NeurIPS)* (1998), pp. 1073–1079.
- [47] C. J. Watkins and P. Dayan, Q-learning, *Mach. Learn.* **8**, 279 (1992).
- [48] W. Barfuss, Reinforcement learning dynamics in the infinite memory limit, in *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)* (2020), pp. 1768–1770.
- [49] W. Barfuss, Dynamical systems as a level of cognitive analysis of multi-agent learning, *Neural Comput. Appl.* **34**, 1653 (2022).
- [50] T. Galla, Intrinsic noise in game dynamical learning, *Phys. Rev. Lett.* **103**, 198702 (2009).
- [51] S. H. Strogatz, *Nonlinear Dynamics and Chaos* (CRC Press, Boca Raton, FL, 2018).
- [52] S. Lange, T. Gabel, and M. Riedmiller, Batch reinforcement learning, in *Reinforcement Learning* (Springer, Berlin, 2012), pp. 45–73.
- [53] L.-J. Lin, Self-improving reactive agents based on reinforcement learning, planning and teaching, *Mach. Learn.* **8**, 293 (1992).
- [54] R. S. Sutton, Integrated architectures for learning, planning, and reacting based on approximating dynamic programming, in *Machine Learning Proceedings 1990* (Elsevier, Amsterdam, 1990), pp. 216–224.
- [55] R. Parr and S. Russell, Approximating optimal policies for partially observable stochastic domains, in *International Joint Conference on Artificial Intelligence (IJCAI)* (1995), pp. 1088–1094.
- [56] M. L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming* (John Wiley & Sons, New York, 2014).
- [57] R. Perman, Y. Ma, J. McGilvray, and M. Common, *Natural Resource and Environmental Economics* (Pearson Education, London, 2003).
- [58] E. Lindkvist and J. Norberg, Modeling experiential learning: The challenges posed by threshold dynamics for sustainable renewable resource management, *Ecol. Econ.* **104**, 107 (2014).
- [59] W. Barfuss, J. F. Donges, M. Wiedermann, and W. Lucht, Sustainable use of renewable resources in a stylized social-ecological network model under heterogeneous resource distribution, *Earth Syst. Dyn.* **8**, 255 (2017).
- [60] F. Geier, W. Barfuss, M. Wiedermann, J. Kurths, and J. F. Donges, The physics of governance networks: Critical transitions in contagion dynamics on multilayer adaptive networks with application to the sustainable use of renewable resources, *Eur. Phys. J.: Spec. Top.* **228**, 2357 (2019).
- [61] D. Blackwell, Equivalent comparisons of experiments, *Ann. Math. Statist.* **24**, 265 (1953).
- [62] J. T. Mark, B. B. Marion, and D. D. Hoffman, Natural selection and veridical perceptions, *J. Theor. Biol.* **266**, 504 (2010).

- [63] M. A. Nowak, Five rules for the evolution of cooperation, *Science* **314**, 1560 (2006).
- [64] P. Kolllock, Social dilemmas: The anatomy of cooperation, *Annu. Rev. Sociol.* **24**, 183 (1998).
- [65] W. Barfuss, J. F. Donges, V. V. Vasconcelos, J. Kurths, and S. A. Levin, Caring for the future can turn tragedy into comedy for long-term collective action under risk of collapse, *Proc. Natl. Acad. Sci. U.S.A.* **117**, 12915 (2020).
- [66] P. Levine and J.-P. Ponsard, The values of information in some nonzero sum games, *Int. J. Game Theory* **6**, 221 (1977).
- [67] M. LiCalzi and R. Mühlenbernd, Categorization and cooperation across games, *Games* **10**, 5 (2019).
- [68] G. Hardin, The tragedy of the commons, *Science* **162**, 1243 (1968).
- [69] S. Barrett and A. Dannenberg, Climate negotiations under scientific uncertainty, *Proc. Natl. Acad. Sci. U.S.A.* **109**, 17372 (2012).
- [70] S. Barrett and A. Dannenberg, Tipping versus cooperating to supply a public good, *J. Eur. Econ. Assoc.* **15**, 910 (2017).
- [71] D. Hennes, M. Kaisers, and K. Tuyls, RESQ-learning in stochastic games, in *Adaptive and Learning Agents Workshop at AAMAS (ALA)* (2010).
- [72] P. M. Todd and G. E. Gigerenzer, *Ecological Rationality: Intelligence in the World* (Oxford University Press, Oxford, 2012).
- [73] R. Hertwig, T. J. Pleskac, and T. Pachur, *Taming Uncertainty* (MIT Press, Cambridge, MA, 2019).
- [74] D. D. Hoffman, M. Singh, and C. Prakash, The interface theory of perception, *Psychon. Bull. Rev.* **22**, 1480 (2015).
- [75] S. J. Shettleworth, *Cognition, Evolution, and Behavior* (Oxford University Press, Oxford, 2009).
- [76] R. Dukas and J. M. Ratcliffe, *Cognitive Ecology II* (University of Chicago Press, Chicago, 2009).
- [77] C. Prakash, K. D. Stephens, D. D. Hoffman, M. Singh, and C. Fields, Fitness beats truth in the evolution of perception, *Acta Biotheor.* **69**, 319 (2021).
- [78] R. P. Mann, Optimal use of simplified social information in sequential decision-making, *J. R. Soc. Interface.* **18**, 20210082 (2021).
- [79] W. Barfuss, J. F. Donges, S. J. Lade, and J. Kurths, When optimization for governing human-environment tipping elements is neither sustainable nor safe, *Nat. Commun.* **9**, 2354 (2018).
- [80] F. M. Strnad, W. Barfuss, J. F. Donges, and J. Heitzig, Deep reinforcement learning in world-earth system models to discover sustainable management strategies, *Chaos* **29**, 123122 (2019).
- [81] J. B. Bak-Coleman, M. Alfano, W. Barfuss, C. T. Bergstrom, M. A. Centeno, I. D. Couzin, J. F. Donges, M. Galesic, A. S. Gersick, J. Jacquet *et al.*, Stewardship of global collective behavior, *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2025764118 (2021).
- [82] M. N. Burton-Chellew, H. H. Nax, and S. A. West, Payoff-based learning explains the decline in cooperation in public goods games, *Proc. R. Soc. B* **282**, 20142678 (2015).
- [83] C. D. Kolstad, Systematic uncertainty in self-enforcing international environmental agreements, *J. Environ. Econ. Manage.* **53**, 68 (2007).
- [84] M. Milinski, R. D. Sommerfeld, H.-J. Krambeck, F. A. Reed, and J. Marotzke, The collective-risk social dilemma and the prevention of simulated dangerous climate change, *Proc. Natl. Acad. Sci. U.S.A.* **105**, 2291 (2008).
- [85] C. D. Kolstad and A. Ulph, Uncertainty, learning and heterogeneity in international environmental agreements, *Environ. Resour. Econ.* **50**, 389 (2011).
- [86] E. F. Domingos, J. Grujić, J. C. Burguillo, G. Kirchsteiger, F. C. Santos, and T. Lenaerts, Timing uncertainty in collective risk dilemmas encourages group reciprocation and polarization, *iScience* **23**, 101752 (2020).
- [87] H. H. Nax, S. Baliatti, R. O. Murphy, and D. Helbing, Adding noise to the institution: an experimental welfare investigation of the contribution-based grouping mechanism, *Social Choice Welfare* **50**, 213 (2018).
- [88] <https://github.com/wbarfuss/POLD> and archived at <https://doi.org/10.5281/zenodo.6361994>.