Outsourcing Memory Through Niche Construction

Edward D. Lee[a], **Jessica C. Flack**[b], **and David C. Krakauer**[b]

[a]Complexity Science Hub Vienna, Josefstædter Strasse 39, Vienna, Austria; [b]Santa Fe Institute, 1399 Hyde Park Rd, Santa Fe, NM 87501

This manuscript was compiled on September 1, 2022

Adaptation to changing environments is a universal feature of life. How should adaptation—to include the duration of memory—scale with environmental rate of change given trade-offs in remembering vs. forgetting and active modification of the environment (e.g. niche construction)? We derive a universal scaling law for optimal memory duration as a function of sensory precision with environmental bias and stability. We find the rate of adaptation scales sublinearly with rate of environmental change across a range of environmental volatility. We use this result and an understanding of metabolic constraints on long memory to explore game dynamics favoring outsourcing of memory to the environment through active modification. We predict stabilizing niche construction will evolve when neural tissue is costly, the environment is complex, and it is beneficial to be able to encode a rich repertoire of environmental states.

adaptation | learning | stigmergy | niche construction | scaling

What is the optimal timescale of adaptation—how long should memory of the environment persist given a changing environment? Research in a wide range of fields suggests that there is no single, simple answer to this question when the environment responds to adaptation. Slowly evolving genes co-evolve with quickly evolving culture (1), as in the evolution of dairy-farming facilitated selection of alleles for adult lactase persistence (2). Quickly evolving organisms modify their otherwise slowly changing niches and alter selection pressures (3–5), illustrated by yeast modifying fruit environments to attract Drosophilid flies that enhance yeast propagation (6). Institutions, the product of collective behavior, feedback to influence individual decisions by changing cost landscapes and enhancing cultural transmission (7, 8), illustrated by legislation in support of same-sex marriage that increases the willingness to voice support in the face of risk (9). To gain information about noisy, hidden variables and reduce social uncertainty, error-prone individual pigtailed macaques collectively compute a social power structure that reduces uncertainty about the cost of social interaction (reviewed in (10, 11)). Bacteria quorum sense, controlling group behavior in dynamically complex, changing environments (reviewed in (12)). Individuals, institutions, and firms all adapt to audit targets (Goodhart's Law) and game systems by creating new feedbacks (13–16). In order to undermine competitors, agents can destabilize a system like in the recent Reddit-Gamestop event in which powerful hedge funds are thought to have introduced volatility to markets by manipulating Reddit users to short squeeze other hedge funds (17). Here we develop a synthetic framework for calculating solutions to the problem of adaptation when strategic decisions bear on active modification of shared environments, known as niche construction. We do this by combining information theory, game dynamics, and scaling theory in order to determine how adaptive memory scales in a range of plausible strategic settings.

As a starting point, we relate the rate of adaptation to the rate of discounting of the past. This determines a persistence time for memory that depends on four crucial factors: the variability in the environment (bias), the rate at which the environment meaningfully changes (stability), the capacity agents have to resolve environmental signal (precision), and the rate of agent modification of the environment (feedback). Importantly, we drop the strict assumption of a separation of timescales between agent learning and environmental change in order to formally account for the role of emergent feedback. We allow modification of the environment to be either passive or active, where active modification can be destabilizing (increasing entropy) as well as stabilizing (reducing entropy). We take into account how the precision of an agent's or organism's estimates of environmental state influences its ability to fit the environment given volatility. We describe these features in more detail in "The problem of adaptation." In "Model structure & assumptions," we present the modeling framework. In "Result 1," we explore when long memory is beneficial. In "Result 2," we derive the scaling relationship for optimal memory duration and environmental change. In "Result 3," we derive by way of a back-of-the-envelope calculation the costs of memory using the literature on metabolic scaling. In "Result 4," we introduce game dynamics building on our derivation of the metabolic costs of memory to explore the evolution of active modification and the outsourcing of memory to the environment. In the discussion, we predict when active modification and outsourcing will evolve.

The problem of adaptation

Agents are situated in unpredictable environments. The degree of environmental volatility is a function of the complexity of the environment and its stability. Simple environments are low entropy and highly biased, or take on a few preferred states, whereas complex environments are high entropy and low bias, or take on a multiplicity of states. Stability, on the other hand, is related to fluctuations in the set of preferred states. When environments fluctuate, agents decide how frequently to update their strategies to handle changed circumstances. This is the timescale of adaptation, and it is influenced by the strength with which stored estimates of environmental history influence present estimates (18–21).

Implicitly, the ability of an agent to adapt to the environment is influenced by its perceptual or sensory precision (22). Precision is a function of how sensitive "sensory cells" or units are, how long they sample, and how many cells contribute to the agent's overall estimate of the state of the environment.

In many systems, active agents also intervene on the environment, whereas passive agents adapt but do not modify the environment. As one strategy, active agents can destabilize the

All authors helped develop the initial idea. E.D.L. did the analysis and wrote the code. The authors drafted the manuscript jointly.

The authors declare no competing interests.

[2]To whom correspondence should be addressed. E-mail: edlee@csh.ac.at

environment by depleting resources or introducing variation. The Reddit-Gamestop event is one example of this "entropic niche construction." Another is guerrilla warfare in which a weaker party randomly determines which battles to neglect by allocating zero resources (23). In contrast, active agents can stabilize the environment by buffering against variation (5) or slowing its rate of change to reduce uncertainty about the future (10, 24). A relatively simple example is stigmergy in which trails or routes are consolidated through repeated use (25). More complicated examples include the collective computation of power structures in macaque groups (26) and of foraging subgroup size distributions by spider monkeys (27) in which social structures are developed through communication and decision networks.

The timescale on which active agents intervene—the niche construction timescale—influences the environmental timescale, creating feedback between agent and environment. In particular, stabilization creates a public good expressed as increased environmental predictability that can benefit all. This in turn creates an opportunity for free-riding and introduces strategic dilemmas best understood through game dynamics (explored in Result 4). Games clarify the importance of taking into account the monopolization costs of controlling public goods as well as the opportunity costs of increased precision and memory.

Model structure & assumptions

We summarize the structure of the model in Figure 1. The environment E at time t is described by a probability distribution $p_E(s,t)$ over configurations s, a vector of descriptive properties. The environment has a bias for preferred states that changes on a relatively slow timescale. Here, we represent the state of the environment with a single bit $s \in \{-1, 1\}$, analogous to the location of a resource as a choice between left and right (28–31). In one configuration, the distribution of resources p_E is biased to the left at a given time t, or $p_E(s=-1,t) > p_E(s=1,t)$, such that an agent with matching preference would do better on average than an agent with misaligned preference. In the mirrored configuration, the environment shows a bias of equal magnitude to the right $p_E(s=-1,t) < p_E(s=1,t)$. Such probabilistic bias can be represented as an evolving "field" $h_E(t)$,

$$p_E(s,t) = \frac{1}{2} + \frac{s}{2}\tanh h_E(t), \qquad [1]$$

such that reversal in bias corresponds to flip of sign $h_E(t) \to -h_E(t)$ that naturally embodies a symmetry between left and right. At every time point, the environment has clearly defined bias in one direction or another, determined by setting the external field to either $h_E(t) = -h_0$ or $h_E(t) = h_0$. With probability $1/\tau_E$ per unit time, the bias in the environment reverses such that over time τ_E the environment remains correlated with its past. This formulation yields a stochastic environment whose uncertainty depends on both the rate of fluctuations flipping its statistical bias (i.e. short timescale implies low stability) and the strength of environmental signal given by the bias (i.e. low bias means low signal).

Passive agents sample from the environment to choose an optimal binary action. The precision τ_c with which an agent samples is dependent on the number of sensory cells contributing to the estimate of environmental state, the sensitivity of

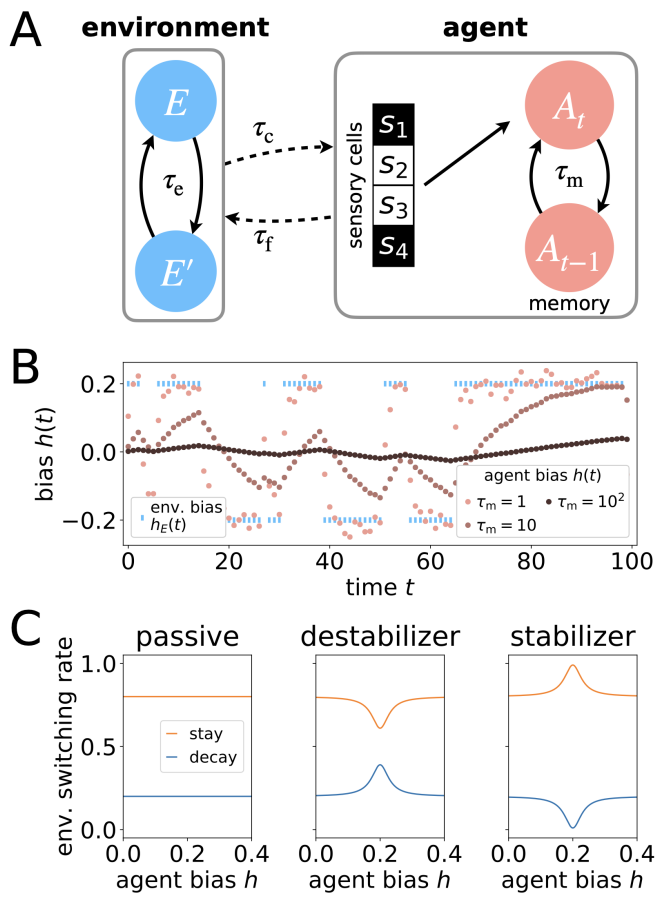

Fig. 1. (A) Overview of framework. Environment E switches configuration on timescale τ_E. The agent measures the current environment through sensory cells with precision τ_c, here worth 4 bits. To obtain an estimate of environment statistics at time t, the agent A_t combines present sensory estimates with memory of previous estimates recorded in an aggregator A_{t-1} (Eq 4) such that memory decays over time τ_m (Eq 5). Coupling with the environment speeds up or slows down environmental change on timescale τ_f (Eq 6). (B) Example trajectories of agents adapting to environmental state $h_E(t)$ with short, medium, and long memory. (C) Rate of environment switching per time step as a function of agent bias h relative to environmental bias $h_E = 0.2$. For passive agents, switching rate does not depend on agent bias. For destabilizers $\alpha = 0.95$, for stabilizers $\alpha = -0.95$. For both, $v = 0.1$ from Eq 6 and environmental timescale $\tau_E = 5$.

those cells, and the number of samples each cell collects. When τ_c is high (either because the sensory cells sampled from the environment for a long time, many sensory cells contributed estimates, or each sensory cell is very sensitive) agents obtain exact measurements of the environment, but a small τ_c corresponds to noisy estimates. The resulting estimate of environmental state $\hat{p}$ thus incurs an error ϵ_{τ_c},

$$\hat{p}(s,t) = p_E(s,t) + \epsilon_{\tau_c}(t). \qquad [2]$$

From this noisy signal, sensory cells obtain an estimate of bias $\hat{h}(t)$, which is related to environmental bias $h_E(t)$ plus measurement noise $\eta_{\tau_c}(t)$,

$$\hat{h}(t) = h_E(t) + \eta_{\tau_c}(t). \qquad [3]$$

In the limit of large precision τ_c and given that the noise in the estimated probabilities $\epsilon_{\tau_c}(t)$ from Eq 2 is binomial distributed, the corresponding error in field $\eta_{\tau_c}(t)$ converges to a Gaussian distribution (see Materials and Methods). Then, at each time

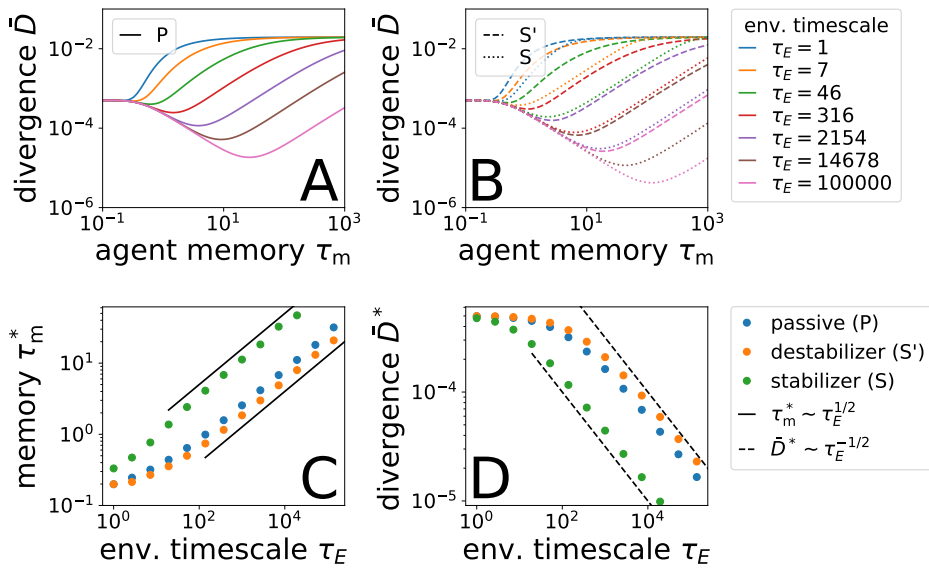

Fig. 2. Divergence $\bar{D}$ as a function of agent memory $\tau_{\rm m}$ and environmental timescale τ_E for (A) passive and (B) active agents including destabilizers S' and stabilizers S. For longer τ_E, agents with longer memory always do better, a pattern emphasized for stabilizers and diminished for destabilizers. (C) Scaling of optimal memory duration $\tau_{\rm m}^*$ with environmental timescale τ_E, corresponding to minima from panels A and B. (D) Divergence at optimal memory duration $\bar{D}^* \equiv \bar{D}(\tau_{\rm m}^*)$. Environmental bias $h_0 = 0.2$.

step the agent's measurement of the environment includes finite-sample noise which is inversely related to precision.

An aggregation algorithm determines how to connect the current measurement with historical ones. This determines the duration of memory by recording the agent's estimate of the state of the environment at the current moment in time $h(t)$ and feeding it to sensory cells at time $t+1$ with some linear weighting $0 \leq \beta \leq 1$ (32),

$$h(t+1) = (1-\beta)\hat{h}(t+1) + \beta h(t). \qquad [4]$$

This estimate is stored in an "aggregator" A_t, and we define $h(0) = 0$. The weight β determines how quickly the previous state of the system is forgotten such that when $\beta = 0$ the agent has no memory and when $\beta = 1$ the agent never forgets its initial state. In between, agent memory decays exponentially with lifetime

$$\tau_{\rm m} \equiv -1/\log\beta. \qquad [5]$$

The weight β that the aggregation algorithm places on the current estimate relative to the stored value gives the timescale of adaptation $\tau_{\rm m}$, or agent memory duration.

The output of this computation is the agent's behavior, $p(s,t)$. We measure the effectiveness of adaptation by quantifying the divergence between a probability vector describing an agent and that of the environment, thereby mathematically capturing how well an agent fits the environment. Measures of divergence, like Kullback-Leibler (KL) divergence, and, more generally, mutual information, have been shown to be natural measures of goodness of fit in evolutionary and learning dynamics from reinforcement learning through to Bayesian inference (33, 34).

As an extension, *active* agents modify the environment's rate of change in a feedback loop. Agents can alter environmental stability, interpreted as changing the probability of switching. We add to the switching rate of the environment $1/\tau_E$ the agent construction rate,

$$\frac{1}{\tau_{\rm f}(t)} \equiv \frac{v^2/\tau_E}{[h(t)-h_E(t)]^2 + v^2}, \qquad [6]$$

such that the probability q that the environment changes at the next point in time is

$$q[h_E(t+1) \neq h_E(t)] = 1/\tau_E + \alpha/\tau_{\rm f}(t). \qquad [7]$$

Eq 6 is written so that it remains normalized for arbitrary v and that the rate gets smaller as the squared distance between agent bias and environmental bias $[h(t)-h_E(t)]^2$ goes to zero. The probability q of the environment switching to the opposite configuration includes weight $\alpha \in (0,1]$ to tune the strength of destabilizers or $\alpha \in [-1,0)$ for stabilizers. This means that for positive α, the rate of switching increases as the agent matches the environment more closely and the opposite for negative α, while the parameter v controls how closely the agent must match the environment to have an effect (i.e. the width of the peak as plotted in Figure 1C). The two types of active agents constitute the basic components of our formulation whose adaptive behavior feeds forward into the timescale of environmental change.* We note that when $1/\tau_{\rm f} = 0$, we obtain *passive* agents that do not modify their environment, thus connecting passive and active agents to one another along a continuum scale.

Putting these elements of adaptation together, as shown in Figure 1A, we obtain a toy learning agent that infers the statistics of a time-varying and stochastic environment.

Result 1: Long memory is beneficial when sensory cells are imprecise & environmental bias is high

The timescale of adaptation represents a balance between the trade-offs of remembering too long and forgetting too fast.

*Note that in this binary example, the new environmental configuration when switching is unique, enforcing a deterministic switch, but that in general there may be a large number of $K \gg 1$ options such that the agent cannot easily guess at the results of environmental fluctuations.

We explore this trade-off by calculating an agent's fit to a changing environment. The fit can be quantified with the KL divergence between environment $p_E(s,t)$ with bias $h_E(t)$ and agent $p(s,t)$,

$$D_{\mathrm{KL}}[p_E||p](t) = \sum_{s\in\{-1,1\}} p_E(s,t)\log_2\left(\frac{p_E(s,t)}{p(s,t)}\right). \quad [8]$$

When the KL divergence is $D_{\mathrm{KL}} = 0$, the agents use optimal bet-hedging, known as "proportional betting," which is important for population growth dynamics (35, 36). Eq 8 is also minimized for Bayesian learners under optimal encoding (37). Assuming agents are playing a set of games in which they must guess the state of the environment at each time step, Eq 8 is the information penalty paid by imperfect compared to perfect agents.

After averaging over many environmental bias switches, we obtain the agent's typical *divergence*,

$$\bar{D} \equiv \lim_{T\to\infty}\frac{1}{T}\sum_{t=0}^{T-1} D_{\mathrm{KL}}[p_E||p](t), \quad [9]$$

The bar notation signals an average over time. Thus, fit improves as $\bar{D}$ decreases.

In Figures 2A and B, we show divergence $\bar{D}(\tau_{\mathrm{m}},\tau_E)$ as a function of the agent's memory τ_{m} given environmental timescale τ_E. In the limiting cases in which an agent has either no memory or has infinite memory, the timescale on which environmental bias changes ultimately has no effect—we observe convergence across all degrees of bias and stability. When an agent has no memory, or $\tau_{\mathrm{m}} = 0$, an agent's ability to match the environment is solely determined by its sensory cells. Low precision τ_{c} leads to large errors on measured environmental bias $h_E(t)$ and large divergence $\bar{D}(\tau_{\mathrm{m}} = 0)$. On the other hand, high precision τ_{c} increases performance and depresses the intercept (Eq 23). At the right hand side of Figure 2A, for large $\tau_{\mathrm{m}} \gg 1$, behavior does not budge from its initial state. Assuming that we start with an unbiased agent such that the transition probability is centered as $q(h) = \delta(h)$, the Dirac delta function, the agent's field is forever fixed at $h = 0$. Then, divergence $\bar{D}(\tau_{\mathrm{m}} = \infty)$ reduces to a fixed value that only depends on environmental bias (Eq 24). In between the two limits of zero and infinite agent memory, the model produces a minimum divergence $\bar{D}(\tau_{\mathrm{m}} = \tau_{\mathrm{m}}^*)$. This indicates the optimal duration of memory τ_{m}^* for a given degree of environmental bias and stability.

The benefits of memory are more substantial for agents with imprecise sensory cells. This benefit is the difference $\bar{D}(\tau_{\mathrm{m}} = 0) - \bar{D}(\tau_{\mathrm{m}} = \tau_{\mathrm{m}}^*)$ as shown in Figure 3A. As one might expect, long memory provides more of a benefit when the present estimate $\hat{p}$ is noisy, ${\tau_{\mathrm{c}}}^{-1}$ is large, and sensory cells are not particularly precise, a deficiency in precision that memory counters by allowing organisms to accumulate information over time. This intuition, however, only applies in the limit of large environmental bias h_0 where the contours of optimal memory flatten and become orthogonal to precision ${\tau_{\mathrm{c}}}^{-1}$. When the bias in the environment is weak, the curved contours show that the benefits of memory come to depend strongly on nontrivial interaction of precision and environmental bias. The complementary plot is the benefit from forgetting, $\bar{D}(\tau_{\mathrm{m}} = \infty) - \bar{D}(\tau_{\mathrm{m}} = \tau_{\mathrm{m}}^*)$ in Figure 3B, which is largely determined by bias h_0. When bias is strong, the costs of estimating the environment inaccurately are large, and it becomes important to forget if sensory cells are imprecise. Thus, our model encapsulates the trade-off between remembering and forgetting both in terms of their absolute benefits as well as the emergence of simple dependence of the respective benefits in the limits of high environmental bias and high sensory precision. An agent has optimally tuned its timescale of adaptation $\tau_{\mathrm{m}} = \tau_{\mathrm{m}}^*$ when it has balanced the implicit costs of tuning to fluctuations against the benefits of fitting bias correctly.

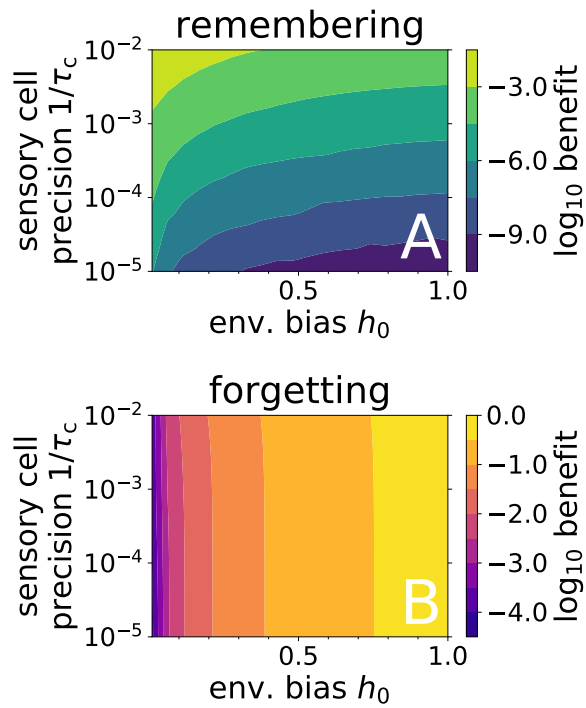

Fig. 3. Benefit from (A) remembering and from (B) forgetting defined as the reduction in divergence at optimal memory duration relative to no memory, $\bar{D}(\tau_{\mathrm{m}} = 0) - \bar{D}(\tau_{\mathrm{m}}^*)$, and optimal memory duration to infinite memory, $\bar{D}(\tau_{\mathrm{m}} = \infty) - \bar{D}(\tau_{\mathrm{m}}^*)$, respectively. We show passive agents given environmental timescale $\tau_E = 10$. All contours must converge (A) when $h_0 = 0$ and (B) when $\tau_{\mathrm{c}} = 0$. Agent-based simulation parameters specified in accompanying code.

Result 2: Sublinear scaling of adaptive time

A finite global optimum for memory duration exists when the environmental timescale is sufficiently long for agents to adapt to environmental signal. At sufficiently large τ_E, we find that optimal memory duration τ_{m}^* scales with the environmental timescale τ_E sublinearly as in Figure 2C. To derive the scaling between optimal memory and environmental timescale, we consider the limit when agent memory is small relative to the environmental timescale $\tau_{\mathrm{m}} \ll \tau_E$ such that the transient time during which the agent is forgetting about the previous environment is short. Under this condition, optimal memory represents a trade-off of poor fit during this transient phase lasting time τ_{m} and the gain from remembering the past when the environment is stable lasting time $\tau_E - \tau_{\mathrm{m}}$. During the transient phase, the agent pays a typical cost at every single time step such that the cost grows linearly with its duration, $\mathcal{C}\tau_{\mathrm{m}}$, for constant $\mathcal{C}$. When the environmental configuration is stable, agent precision is enhanced by a factor of τ_{m} because it effectively averages over that many random samples, or a

gain of $\mathcal{G}\log\tau_{\rm m}$ for constant $\mathcal{G}$. When we weight each term by the fraction of time spent in either transient or stable phases, $\tau_{\rm m}/\tau_E$ and $(\tau_E-\tau_{\rm m})/\tau_E$ respectively, we obtain the trade-off

$$\mathcal{C}\frac{\tau_{\rm m}^2}{\tau_E}-\mathcal{G}\frac{\tau_E-\tau_{\rm m}}{\tau_E}\log\tau_{\rm m}. \qquad [10]$$

At optimal memory $\tau_{\rm m}^*$, Eq 10 will have zero derivative. Keeping only the dominant terms and balancing the resulting equation, we find

$$\tau_{\rm m}^*\sim\tau_E^{1/2}. \qquad [11]$$

This scaling argument aligns with numerical calculation as shown in Figure 2C.

Similarly, we calculate how optimal divergence $\bar{D}^*$ scales with environmental timescale. Assuming that the agent has a good estimate of the environment such that the error in average configuration $\epsilon_{\tau_{\rm c}}(t)$ is small, agent behavior is $p_E(s,t)+\epsilon_{\tau_{\rm c}}(t)$ and $\epsilon_{\tau_{\rm c}}(t)$ is normally distributed. Then, we expand the divergence about $p_E(s,t)$ in Taylor series of error $\epsilon_{\tau_{\rm c}}(t)$ (Materials & Methods). Over a timescale of $\tau_{\rm m}^*$, the precision of this estimate is further narrowed by a factor of $\tau_{\rm m}^*$ such that

$$\bar{D}^*\sim 1/\tau_{\rm m}^*\sim\tau_E^{-1/2}, \qquad [12]$$

Although we do not account for the transient phase, we expect the relation in Eq 12 to dominate in the limit of large τ_E, and our numerical calculations indeed approach the predicted scaling in Figure 2C. In contrast, when environment does not fluctuate, or bias $h_0=0$, agents pay no cost for failing to adapt to new environments and infinite memory is optimal. In this scenario, other physical costs for maintaining large memory would become relevant. Overall, the sublinear scaling between memory duration and rate of environmental change indicates an economy of scale—agents need relatively less memory, or a slower timescale of adaptation, than expected to fit to slowly changing environments than they do to fit to rapidly changing ones.

Result 3: Adaptive vs. metabolic costs

Here we consider how metabolic costs of memory constrain memory duration, and in Result 4 we explore how these constraints influence game dynamics.

We start with the well-documented observation that physical constraints on circulatory networks responsible for energy distribution influence organismal traits including lifespan and size across the animal kingdom from microorganisms to blue whales (38, 39). Metabolic costs matter for brain mass $M_{\rm br}$, which scales with body mass $M_{\rm bo}$ sublinearly, $M_{\rm br}=AM_{\rm bo}^a$, where $a=3/4$ across taxa (within individual taxa it spans the range 0.24 to 0.81 (40)). To account for adaptive cost of memory, we make the simple assumption that brain mass for adaptation is proportional to memory duration. Then,

$$M_{\rm br}\propto N\tau_E, \qquad [13]$$

such that small organisms with short memory, or "mice," have small N and large organisms, "elephants," large N.

Now, we use predictions of allometric scaling theory to relate metabolic rate B to mass, $B\propto M^{1/4}$ (41), and lifespan to body mass, $T\propto M_{\rm bo}^b$ for metabolic exponent $b=1/3$ (42). By relating organism lifetime T to total memory duration (more precisely, episodic memory) from Eq 13, we obtain a relationship between metabolic rate and memory duration, $B\propto N^{\phi}\tau_E^{\phi}$, where $\phi\equiv a/4b$.[†] Note that this scaling is sublinear for biological organisms, $\phi<1$. While adaptive cost decays with τ_E in Figure 4A, metabolism grows as τ_E^{ϕ} as shown in Figure 4B. The contrasting scalings suggest that for "mice" the cost of poor adaptation makes a disproportionate contribution to the life-time energy budget of the organism, consistent with observations in developmental neural growth in wasps (43).

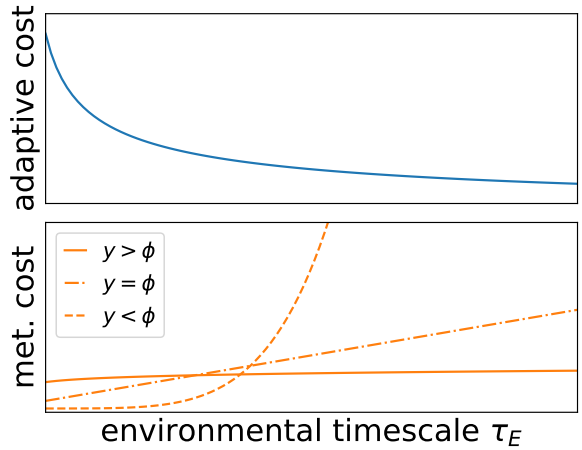

Fig. 4. Schematic of adaptive and metabolic costs with environmental timescale τ_E. (A) Adaptive cost $\bar{D}$ is largest at small τ_E, but (B) metabolic costs are largest for longer-lived organisms with scaling dependent on exponents y and ϕ such as for "elephants" that experience slower environments (Eq 14).

To generalize the previous argument, we consider when larger organisms experience longer environmental timescales; after all, "mice" and "elephants" may experience different local environments. Then, $\tau_E\propto T^y$, where $y\in[0,1]$ to ensure that τ_E and N increase together since $\tau_E^{1/y-1}\propto N$. We now find the relationship between metabolic rate and environmental timescale

$$B\propto\tau_E^{\phi/y}\propto N^{\phi/(1-y)}, \qquad [14]$$

which reduces to the previous case when $y=1$. Such dependence implies that the metabolic cost of adaptation will explode with environmental timescale (and organism lifetime) as y approaches unity or grow slowly and sublinearly when $y=0$ as are contrasted in Figure 4B. More generally, the range of possibilities predict when metabolic costs supporting adaptation dominate over costs of poor adaptation to the environment (44, 45).

Result 4: Niche construction & the outsourcing principle

In Result 3, we showed that long memory has metabolic costs. But outsourcing presents one way of potentially avoiding such costs. Whether ant pheromone trails, food caching, writing, or map-making, niche construction promotes the stability and predictability of the local environment (5, 46), and consequently an organism does not need long memory. Stabilizing

[†] When we use $a=3/4$, we obtain the range $\phi=[5/8,15/16]$, the endpoints depending on whether $b=0.3$ or $b=0.2$ (accounting for taxa-specific variation in a leads to much wider range of $\phi\in[0.2,1.01]$). Thus, we hypothesize that longer environmental timescales lead to increased brain mass and metabolic expenditure with sublinear scaling.

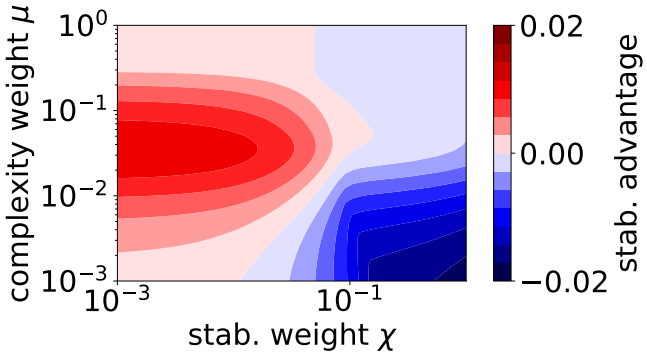

Fig. 5. Comparison of total divergence for stabilizers $\mathcal{D}_S$ and destabilizers $\mathcal{D}_{S'}$, or $\mathcal{D}_{S'} - \mathcal{D}_S$, in a fixed environment and common sensory precision and costs. The difference is between agents poised at optimal memory duration given μ, χ, and β. Small stabilization weight χ favors stabilizers, whereas high monopolization cost μ favors destabilizers. Simulation parameters are specified in accompanying code.

niche construction, however, creates a public good that by reducing uncertainty is useful to all agents, can be exploited by free riders, and leads to a tragedy of the commons (47).

We explore the conditions under which active modification can evolve given the free-rider problem. We introduce stabilizing mutants into a population of passive agents. Assuming other organisms are poorly attuned to regularities in the environment, we expect stabilizing mutants to gain a competitive advantage but only over the short term. In the longer term, an established stabilizer population is susceptible to invasion by free-riders exploiting outsourced memory; said another way, stabilizers slow environmental timescales and reduce divergence for all individuals sharing the environment, but they uniquely pay for stabilization. Thus, as in the classical example of niche construction, the usual "tragedy of the commons" argument makes it an evolutionary dead end (47). It follows that stabilization is only a competitive strategy if individuals can monopolize resulting resources. In the natural world, this could occur through physical encryption (e.g. undetectable pheromones (48)), the erasure of signal (e.g. food caching (49)), or the restriction of social information (e.g. concealment (50)). In this framework, stabilizers modify the environment in a way difficult for alternative strategies to exploit.

To model competition between monopolistic stabilizers and other strategies, we account for the costs of memory, stabilization, and precision. We measure the richness or complexity of memory as

$$H(\tau_{\mathrm{m}}) = \log_2(1 + 1/\tau_{\mathrm{m}}). \qquad [15]$$

Eq 15 can be thought of as a cost of exploring more configurations over a short period time versus agents that are temporally confined.

We define the cost stabilizers pay for niche construction as the extent of change to the environmental switching rate, or the KL divergence between the natural environmental rate $1/\tau_E$ and the time-averaged, modified rate $\langle 1/\tilde{\tau}_E \rangle$,

$$G(1/\tau_E, \langle 1/\tilde{\tau}_E \rangle) = \frac{1}{\tau_E} \log_2 \left(\frac{1/\tau_E}{\langle 1/\tilde{\tau}_E \rangle} \right) + \left(1 - \frac{1}{\tau_E}\right) \log_2 \left(\frac{1 - 1/\tau_E}{1 - \langle 1/\tilde{\tau}_E \rangle} \right). \qquad [16]$$

The quantity G depends implicitly on stabilization strength α because smaller α slows the environment further. For passive agents and destabilizers, $G = 0$ by definition because non-stabilizers fit to τ_E and only stabilizers benefit from the slower timescale with monopolization.

We finally consider the cost of precision, which we assume to be given by the information obtained by the agent from sampling the environment (see SI Appendix C),

$$C(\tau_{\mathrm{c}}) = \log_2 \tau_{\mathrm{c}}. \qquad [17]$$

Putting these costs together with divergence $\bar{D}$, we obtain the total divergence

$$\mathcal{D} = \bar{D} + \mu H + \chi G + \beta C. \qquad [18]$$

Weights $\mu \geq 0$, $\chi \geq 0$, $\beta \geq 0$ represent the relative contribution of these costs. As a result, we can distinguish dominant strategies by comparing total divergence such as between the pair of destabilizer and stabilizer strategies shown in Figure 5. Large μ, or high complexity cost, means that a pure population of stabilizers would be stable to invasion from destabilizers. Whereas for large χ, or heavy stabilization cost, the opposite is true. The generalized measure of adaptive cost in Eq 18, given the weights, carves out regions of agent morphospace along axes of computational cost. This is a morphospace that captures the relative advantage of internal versus external memory that can be thought of as a space of evolutionary outsourcing.

As has often been remarked in relation to evolution, survival is not the same as arrival. We now determine when stabilizer strategies can emerge in this landscape. We start with a pure population of passive agents with stabilization strength $\alpha = 0$ and poised about optimal memory duration $\tau_{\mathrm{m}} = \tau_{\mathrm{m}}^*$ determined by minimizing both divergence $\bar{D}$ and complexity μH. Whether or not stabilizers emerge under mutation and selection can be determined through adaptive dynamics (51–53), that is by inspecting the gradient of the total divergence along the parameters $(\partial_{\tau_{\mathrm{m}}}\mathcal{D}, \partial_\alpha \mathcal{D}, \partial_{\tau_{\mathrm{c}}}\mathcal{D})$, or memory complexity, stabilizer strength, and precision. As we show in SI Appendix D and Eq S16, the gradient terms can be calculated under a set of perturbative approximations. Using local convexity about optimal memory τ_{m}^*, we show that the term $\partial_\alpha \mathcal{D}$ drives passive agents to smaller α and slower timescales; it originates from combining the scaling law from Eq 12 and complexity of memory. The term $\partial_{\tau_{\mathrm{c}}}\mathcal{D}$ shows that precision tends to decrease when the cost gradient $\partial_{\tau_{\mathrm{c}}}(\beta C)$ dominates over $\partial_{\tau_{\mathrm{c}}}\bar{D}$. In this case, the general conditions $\partial_\alpha \mathcal{D} < 0$ and $\partial_{\tau_{\mathrm{c}}}\mathcal{D} < 0$ funnel a passive population towards stabilization and reduced precision.

The conditions on the gradient depend critically on the costs specified. That niche construction is commonplace in the natural and cultural worlds, however, suggests that the costs in Eq 18 should favor the emergence of stabilization strategies. With these costs in place adaptive dynamics that stabilize the environment will slow down environment fluctuations, lengthen optimal τ_{m}^*, and make weak stabilizers less competitive. This moves a population as a whole towards slower timescales, which reduces pressure to monopolize the target environmental variable and can reduce precision. Consequently, stabilizing niche construction may then promote outsourcing of memory by reducing pressure on the organism to invest in neural material to support a long memory. This

I Bias-Stability	II Stability-Precision	III Bias-Precision	IV Integrated
Taxis of larval invertebrates (54)	Seed dormancy/germ banking (55)	Bandit problems (56)	Volatile bandits (57)
Stochastic voting models (58)	Particle swarms (59)	Microbial chemotaxis (60)	Learning changing data sources (61)
Learning changing distributions (62)	Cognitive aging (63)	Speed-accuracy trade-offs (31, 64, 65)	Consensus with link failure (66)
Loss/Change aversion (67)		Optimal foraging (68)	Retinal sensitivity rescaling (19)
		Page Rank consensus (69)	

Table 1. Classification of studies in terms of emphasized factors of bias, stability, and precision.

is because adaptation costs can be lower even with less memory in a stable environment. With the savings, the organism either can reduce total brain size, or it can invest in a larger behavioral repertoire capable of monopolizing a larger number of environmental states.

Discussion

The many facets of adaptation constitute a major area of research across fields. In an overarching taxonomy, approaches connect (I) bias-stability, (II) bias-precision, and (III) stability-precision indicated in Table 1, whereas work that has implicitly combined all of these factors is noted under "integrated" (IV). Studies in category I focus on rules that apply in variable environments (bias) where environmental distributions are prone to rapid change (stability). Studies in category II focus on rules that apply when environments are likely to change (stability) and where making the correct decision depends on sensitivity to signals (precision or also "accuracy" in some literature). Studies in category III focus on rules that apply in variable environments (bias) and where making the correct decision depends on power of sensors to detect signal (precision). Studies in category IV include elements of I–III and apply to variable environments prone to rapid change, where sensory precision varies. Feedback with the environment, where agent inference modifies the input statistics and timescales are formally coupled, remains little considered despite being a central premise of research on niche construction and stigmergy. An important contribution of our work is to distill the properties that contribute to adaptation and to map them onto separate quantitative variables. With the clarified representation, we integrate them with an information theoretic and game dynamical modeling framework.

Regardless of how adaptation time is influenced by cognitive factors like computational precision and ecological factors like degree of environmental bias and stability, we find it varies sublinearly with the environmental timescale. Sublinear scaling of adaptation and environment might describe a wide range of systems in nature despite tremendous diversity in biological systems in how memory is implemented, ranging from biased walks in chemotaxis (70) to pheromone stigmergy (25, 71, 72) to distributed collective social action (73). How different mechanisms like selective forgetting might tune the balance may be a unifying question to explore (21, 74–76).

We consider information costs of adaptation in relation to metabolic and other costs for the evolution of active modification and outsourcing of memory to the environment. Our assumption that brain mass grows linearly with environmental time scales leads to a superlinear scaling of brain mass supporting memory with memory duration, or τ_{m}^2. In contrast, metabolic scaling predicts a substantially faster scaling of total brain mass as τ_{m}^6, which is supported by allometric measurements (40). Beyond the metabolic economy-of-scale implied, we anticipate that stable environments foster memory outsourcing when organisms come to strategically leverage a stable environment such as a slowly changing power structure as we discuss below. In turn, we propose that the savings in neural tissue can either lead to a proportional reduction in brain size or can be dedicated to encoding greater environmental complexity: tests of this prediction, for example, suggest a new direction for precision experiments in insect neural development that have yet to consider environmental outsourcing (43, 45). Savings in neural tissue motivate the cost of memory complexity, along with stabilization and sensory costs, when we consider the emergence of stabilization strategies.

When we focus on these information costs in game dynamics, the implications are consistent with prior work on social dynamics showing long-term social structure can serve as a form of collectively constructed societal memory. In pigtailed macaque society, a social power distribution arises from the exchange of subordination signals and provides a temporary but predictable social background against which individuals can adapt (10, 11, 77). In this form of social niche construction, individuals collectively stabilize perceived asymmetries about group members' capacity to successfully use force, thereby reducing uncertainty about outcome and cost of social interactions. Memory of who is perceived to be capable of winning becomes encoded in the slowly changing power structure, reducing the demands on individual memory because they no longer need to recall a varied history of fight outcomes.

This is consistent with our finding that a stable environment, the collectively computed power distribution, could permit individuals to invest in other facets of cognitive complexity such as encoding a greater complexity or number of social states (78). We then predict that stabilizing niche construction and outsourcing are more likely to evolve when costs of long memory trade-off with the costs of complex memory and the social environment is rich, meaning when it is valuable to be able to track diverse strategies and states. Hence active modification initially increases determinism and mutual information with the environment, but we anticipate that it contributes to the evolution of social complexity over the long run by freeing up cognitive resources.

Such predictions present a rich range of intriguing questions and extensions of our model to ongoing research in bacteria, animals, and social institutions using our general framework to conceptualize and to analyze adaptive niche construction.

Materials and Methods

The code used to generate these results will be made available on GitHub at https://github.com/eltrompetero/adaptation.

Numerical solution to model. Given Eqs 1–4 defining the binary agent, we calculate agent behavior in two ways. The first method is with

agent-based simulation (ABS). We generate a long time series either letting the environment fluctuate independently and training the agent at each moment in time or coupling environmental fluctuations at each time step with the state of the agent. By sampling over many such iterations, we compute the distribution over agent bias given environmental bias, $q(h|h_E)$, which converges to a stationary form.

This principle of stationarity motivates our second solution of the model using an eigenfunction method. If the distribution is stationary, then we expect that under time evolution that the conditional agent distribution map onto itself

$$q(h|h_E) = \mathcal{T}[q(h|h_E)]. \tag{19}$$

If the time-evolution operator $\mathcal{T}$ evolves the distribution over a single time step, the external field can either stay the same with probability $1 - 1/\tau_E$ or reverse with probability $1/\tau_E$.

For either for these two possible alternatives over a single time step, we must convolve the distribution with the distribution of noise for the field η_{τ_c}. The distribution of noise derives from agent perceptual errors ϵ_{τ_c} on the estimated probabilistic bias of the environment (Eq 2). Hence, the corresponding error distribution for the bias η_{τ_c} originates from the binomial distribution through a transformation of variables. We can simplify this because in the limit of large sensory cell sample size τ_c the binomial distribution converges to a Gaussian and a concise representation of the distribution of η_{τ_c} becomes accurate. Using Eq 1, we find that the distribution of perceptual errors in the bias yields

$$\begin{aligned}\rho(\eta_{\tau_c}, t) = (8\pi\sigma^2)^{-1/2} \exp\Big\{ -[\tanh h_E(t) - \\ \tanh(h_E(t) + \eta_{\tau_c})]^2 / 8\sigma^2 \Big\} \operatorname{sech}^2 (h_E(t) + \eta_{\tau_c}) .\end{aligned} \tag{20}$$

Here, the agent's perceptual estimate of the environment includes finite-sample noise determined by the sensory cell precision $1/\tau_c$. At finite τ_c, there is the possibility that the agent measure a sample from the environment of all identical states. In our formulation, the fields then diverge as do the fields averaged over many separate measurements. We do not permit such a "zero-temperature" agent that freezes in a single configuration in our simulation just as thermodynamic noise imposes a fundamental limit on invariability in nature. Our agents inhabit an *in silico* world, where the corresponding limit is fixed by the numerical precision of the computer substrate, so we limit the average of the bits sampled from the environment to be within the interval $[-1 + 10^{-15}, 1 - 10^{-15}]$. This is one amongst variations of this idea that inference is constrained by regularization, Bayesian priors, Laplace counting (in the frequentist setting), etc. Regardless of the particular approach with which finite bounds might be established, they are only important in the small τ_c limit. See SI Appendix A.

Given the Gaussian approximation to precision error, we propagate the conditional distribution over a single time step, defining a self-consistent equation that can be solved by iterated application. To make this calculation more efficient, we only solve for abscissa of the Chebyshev basis in the domain $\beta \in [0, 1]$, fixing both the endpoints of the interval including the exact value for $\beta = 1$ from Eq 24 (79) (more details in SI Appendices A and B). In Figure S7, we show that our two methods align for a wide range of agent memory τ_m. Importantly, the eigenfunction approach is much faster than ABS for large τ_c because the latter can require a large number of time steps to converge. On the other hand, ABS is relatively fast for small τ_c. Thus, these two approaches present complementary methods for checking our calculation of agent adaptation.

Divergence curves. To measure how well agent behavior is aligned with the environment, we compare environment $p_E(s, t)$ and agent $p(s, t)$ with the KL divergence at each time step to obtain the agent's typical loss in Eq 9. Equivalently, we can average over the stationary distribution of fields conditional on environment

$$\bar{D} = \frac{1}{N_E} \sum_E \int_{-\infty}^{\infty} dh\, q(h|h_E) D_{\mathrm{KL}}[p_E(h_E)||p(h)], \tag{21}$$

where we sum over all possible environments E and weight them inversely with the number of total environments N_E. For the binary case, $N_E = 2$. We furthermore simplify this for the binary case as

$$\bar{D} = \int_{-\infty}^{\infty} dh\, q(h|h_E = h_0) D_{\mathrm{KL}}[p_E(h_E = h_0)||p(h)]. \tag{22}$$

In Eq 22, we have combined the two equal terms that arise from both positive $h_E = h_0$ and negative $h_E = -h_0$ biases of the environment.

In Figure 2A and B, we show divergence as a function of agent memory over a variety of environments of varying correlation time $\bar{D}(\tau_m, \tau_E)$. When the agent has no memory, its behavior is given solely by the properties of the sensory cells as is determined by the integration time τ_c. Then, we only need account for the probability that the environment is in either of the two symmetric configurations and how well the memoryless agent does in both situations. Since the configurations are symmetric, the divergence at zero memory is

$$\begin{aligned}\bar{D}(\tau_m = 0) = \int_{-\infty}^{\infty} d\eta_{\tau_c}\, \rho(\eta_{\tau_c}|h_E = h_0) \times \\ \sum_{s\in\{-1,1\}} p_E(s|h_E = h_0) \log_2 \left(\frac{p_E(s|h_E = h_0)}{p(s)} \right),\end{aligned} \tag{23}$$

where the biased distribution of environmental state p_E and the error distribution ρ from Eq 20 are calculated with environmental bias set to $h_E = h_0$. Note that this is simply Eq 22 explicitly written out for this case.

At the limit of infinite agent memory, as in the right hand side of Figure 2A, passive agents have perfect memory and behavior does not budge from its initial state. Assuming that we start with an unbiased agent such that $q(h) = \delta(h)$, the Dirac delta function, the agent's field is forever fixed at $h = 0$. Then, divergence reduces to

$$\bar{D}(\tau_m = \infty) = 1 - S[p_E], \tag{24}$$

where the conditional entropy $S[p_E] = -p_E(s|h = h_0) \log_2 p_E(s|h = h_0) - [1 - p_E(s|h = h_0)] \log_2[1 - p_E(s|h = h_0)]$.

Scaling argument for optimal memory. As is summarized by Eq 10, the value of optimal memory can be thought of as a trade-off between the costs of mismatch with environment during the transient adaptation phase and gain from remembering the past during stable episodes. In order to apply this argument to the scaling of divergence, we consider the limit where the environment decay time τ_E is very long and agent memory τ_m is long though not as long as the environment's. In other words, we are interested in the double limit $\tau_m \to \infty$ and $\tau_m/\tau_E \to 0$. Then, it is appropriate to expand divergence in terms of the error in estimating the bias

$$\begin{aligned}\bar{D} = \Bigg\langle \sum_{s\in\{-1,1\}} p_E(s,t) \log p_E(s,t) - \\ p_E(s,t) \log[p_E(s,t) + \epsilon_{\tau_c}(t)] \Bigg\rangle,\end{aligned} \tag{25}$$

where the average is taken over time. Considering only the second term and simplifying notation by replacing $\epsilon_{\tau_c}(t)$ with ϵ,

$$\begin{aligned}&\langle p_E(s,t) \log p_E(s,t) + \log[1 + \epsilon/p_E(s,t)] \rangle \\ &\approx \left\langle p_E(s,t) \log p_E(s,t) + \frac{\epsilon}{p_E(s,t)} - \frac{1}{2} \frac{\epsilon^2}{p_E(s,t)^2} \right\rangle,\end{aligned} \tag{26}$$

where the average error $\langle \epsilon \rangle = 0$ and assuming that the third order correction $\mathcal{O}\left(\left\langle \epsilon^3 \right\rangle\right)$ is negligible. Plugging this back into Eq 25,

$$\bar{D} \approx \sum_{s\in\{-1,1\}} \frac{\tau_E - \tau_m}{\tau_E} \frac{\left\langle \epsilon^2 \right\rangle}{p_E(s)^2} + \frac{\tau_m}{\tau_E} \left\langle \frac{\epsilon^2}{p_E(s,t)^2} \right\rangle. \tag{27}$$

The first term in Eq 27 relies on the fact that when environmental timescales are much longer than agent memory, the errors become independent of the state of the environment. Thus, we can average over the errors separately, and the environment configuration average can be treated independently of time $p_E(s,t) \to p_E(s)$. The

second term, however, encases the transient dynamics that follow immediately after a switch in environmental bias while the agent remembers the previous bias. It is in the limit $\tau_{\mathrm{m}}/\tau_E \to 0$ that we can completely ignore this term and the scaling for optimal memory $\tau_{\mathrm{m}}^* \sim \tau_E^{1/2}$ from Eq 11 is the relevant limit that we consider here.

Since the errors with which the agent's matching of environmental bias is given by a Gaussian distribution of errors, the precision increases with the number of samples taken of the environment: it should increase with both sensory cell measurement time τ_{c} as well as the typical number of time steps in the past considered, $\tau_{\mathrm{m}} = -1/\log\beta$. Thus, we expect the scaling of divergence at optimal memory to be

$$\bar{D}^* \sim \frac{1}{\tau_{\mathrm{m}}^* \tau_{\mathrm{c}}}, \qquad [28]$$

which with Eq 11 leads to the scaling of optimal memory with environment decay time Eq 12. Though the scaling with precision timescale τ_{c} in Eq 28 is at $\tau_{\mathrm{m}} = \tau_{\mathrm{m}}^*$, it is clear that a similar scaling with τ_{c} holds at $\tau_{\mathrm{m}} = 0$, where only precision determines divergence. However, such a scaling does not generally hold for any fixed τ_{m}, the trivial case being at $\tau_{\mathrm{m}} = \infty$, where divergence must go to a constant determined by environmental bias.

ACKNOWLEDGMENTS. E.D.L. was supported by the Omega Miller Program at the Santa Fe Institute. D.C.K. and J.F. are grateful for support from the James S. McDonnell Foundation 21st Century Science Initiative-Understanding Dynamic and Multi-scale Systems.

A. Agent-based simulation

To complement the eigenfunction solution described in Appendix B, we present a simple agent-based simulation.

After having specified the environmental bias $h_E(t)$, we generate a sample of τ_c binary digits from the distribution $p_E(s,t)$. From this sample, we calculate the mean of the environment $\langle s \rangle$ which is bounded in the interval $[-1+10^{-15}, 1-10^{-15}]$. These bounds are necessary to prevent the measured field $\hat{h}(t)$ from diverging and reflects the fact that *in silico* agents have a finite bound in the values they can represent, mirroring finite cognitive resources for biological or social agents as discussed in Materials and Methods. We combine this estimated field $\hat{h}(t)$ with the one from the aggregator having set the initial value condition $H(0) = 0$. Given the estimate of the field $h(t)$, we compute the Kullback-Leibler (KL) divergence between the agent distribution $p(s)$ and the environment $p_E(s)$.

When we calculate the divergence landscape across a range of different agent memories, we randomly generate the environment using the same seed for the random number generator. Though this introduces bias in the pseudorandom variation between divergence for agents of different types, it makes clearer the form of the divergence landscape by eliminating different offsets between the points. Our comparison of this approach with the eigenfunction solution in Appendix B provides evidence that such bias is small with sufficiently long simulations. For the examples shown in the main text, we find that total time $T = 10^7$ or $T = 10^8$ are sufficient for convergence to the stationary distribution after ignoring the first $t = 10^4$ time steps.

B. Eigenfunction solution

We present more details on top of those in Materials and Methods on the iterative, eigenfunction solution to the divergence of an agent relying on the fact that the distribution of agent bias $q(h)$ becomes stationary at long times.

Let us first consider the case of the passive agent. After sufficiently long time, the distribution of agent behavior $q(h)$ and the distribution conditioned on the two states of the environment $q(h|h_E = h_0)$ and $q(h|h_E = -h_0)$ converge to stationary forms. Assuming that the distributions have converged, we evolve the distribution a single time step. If the external field $h_E(t) = h_0$, then it either stays fixed with probability $1 - 1/\tau_E$ or it switches to the mirrored configuration with probability $1/\tau_E$.

Considering now the evolution of the conditional probability $q(h|h_E = h_0)$, we note that the state of the agent will be either be convolved by the distribution of sampling error at the next time step or lose probability density from a switching field. Since we are considering a symmetric configuration, however, the mirrored conditional density will reflect the same probability density back such as in Eq S1. Thus, Eq S1 is satisfied by the conditional density of agent bias that is solved by the eigenfunction for $q(h|h_E)$ with eigenvalue 1. By the Perron-Frobenius theorem when considering normalized eigenvectors, this is the unique and largest eigenvalue that returns the stationary solution.

To extend this formulation to active agents, we must also account for the dependence of the rate of switching on the distance between agent and environmental bias. This additional complication only requires a modification of Eq S1 to include such dependence in the rate coefficients. Thus, all types of agents can be captured by this eigenfunction solution and solved by iteration til convergence.

Eq S1 is only independent of time when agent memory $\tau_m = 0$. When there is finite memory, or $\beta > 0$, the distribution $q(h,t)$ "remembers" the previous state of the environment such that we must iterate Eq S1 again. Over many iterations, we will converge to the solution, but the convergence slows with agent memory which introduces ever slower decaying eigenfunctions. An additional difficult arises because the narrowing in the peak of the agent's estimate of the environment, like the peaks shown in Figure S7, require increased numerical precision. As a result, increasing memory and computational costs make it infeasible to calculate the eigenfunction with high precision for β close to 1.

Instead of calculating the full functional form directly below but not at the limit $\beta \to 1$, we use the output of the iterative eigenfunction procedure as input for an interpolation procedure using Chebyshev polynomials. We iterate Eq S1 for β equal to the Gauss-Lobatto abscissa of the Chebyshev polynomial of degree d, mapping the interval $\beta \in [0,1]$ to the domain $x \in [-1,1]$ for the set of Chebyshev polynomials (79). The Gauss-Lobatto points include the endpoints $\beta = 0$ and $\beta = 1$, the first of which is trivial numerically and the latter for which we have an exact solution given in Eq 24. Then, we exclude calculated values for large β that show large iteration error $\epsilon > 10^{-4}$. This threshold, however, leaves the coefficients of the Chebyshev polynomial undetermined. We instead interpolate these remaining $N - k$ points by by fitting a Chebyshev polynomial of degree $N - k - 1$ with least-squares on the logarithm of the divergence. A similar procedure can be run for the stabilization cost from Eq 16 to obtain Figure S6B. We find that typically $N = 30$ or $N = 40$ starting abscissa with a maximum of 10^3 iterations are sufficient to obtain close agreement with the agent-based simulation (ABS) from Appendix A (Figure S8). This interpolation procedure does not work well with ABS because small stochastic errors can lead to high-frequency modes in interpolation (and thus large oscillations), errors that can be essentially driven to zero exponentially fast for the eigenfunction method.

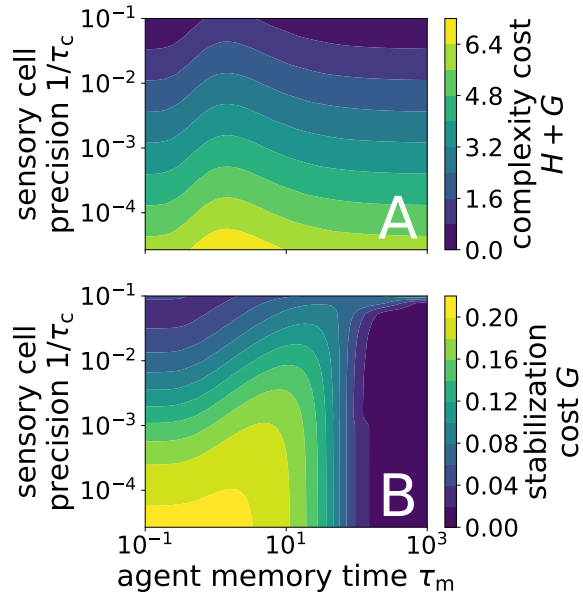

Fig. S6. Landscape of the two costs we consider, (A) agent complexity and (B) stabilization. (A) The values have been offset to ensure that the costs are positive over the shown landscape calculated from memory (Eq 15) and sensory cost (Eq 17). (B) Isocontours defined as sum of costs. As memory $\tau_m \to \infty$, stabilization cost converges to a finite value that can be calculated exactly from noting that agent behavior has probability density fixed at its starting point, $q(h) = \delta(h)$. A kink in the contours at $1/\tau_c = 10^{-3}$ arises from numerical precision errors where we matched up ABS and eigenfunction methods.

C. Algorithmic costs truncate scaling

Agents must expend energy and time to preserve or erase memory (21, 80, 81), to measure the environment (82, 83), and to modify the environment (84–86). Here, we account for agent complexity in terms of the information costs they incur, providing a general way to account for agent design and actions discussed further in Appendix C (44).

We separately account for costs of agent memory, sensory cell precision, and environmental stabilization. For memory, we consider the resulting complexity of agent behavior. We might expect infinite memory to be cheap because it is a minimally demanding strategy: no history is required because the agent adopts a preset configuration that does not change. This is captured by memory complexity, $H = \log(1 + 1/\tau_m)$, plotted in Figure S10 as a dot-dashed line (87). In neural circuits, this cost might be measured as the metabolic cost of enhancing or degrading new memories (21, 74, 75). Sensory complexity means that higher precision implies higher expenditure to

$$q(h,t|h_E=h_0)=\left(1-\frac{1}{\tau_E}\right)\int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\rho(\eta_{\tau_c}|h_E=h_0)q(h,t-1|h_E=h_0)\delta(h-h_0-\eta_{\tau_c})\,d\eta_{\tau_c}\,dh+ \frac{1}{\tau_E}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\rho(\eta_{\tau_c}|h_E=-h_0)q(h,t-1|h_E=-h_0)\delta(h+h_0-\eta_{\tau_c})\,d\eta_{\tau_c}\,dh. \quad [S1]$$

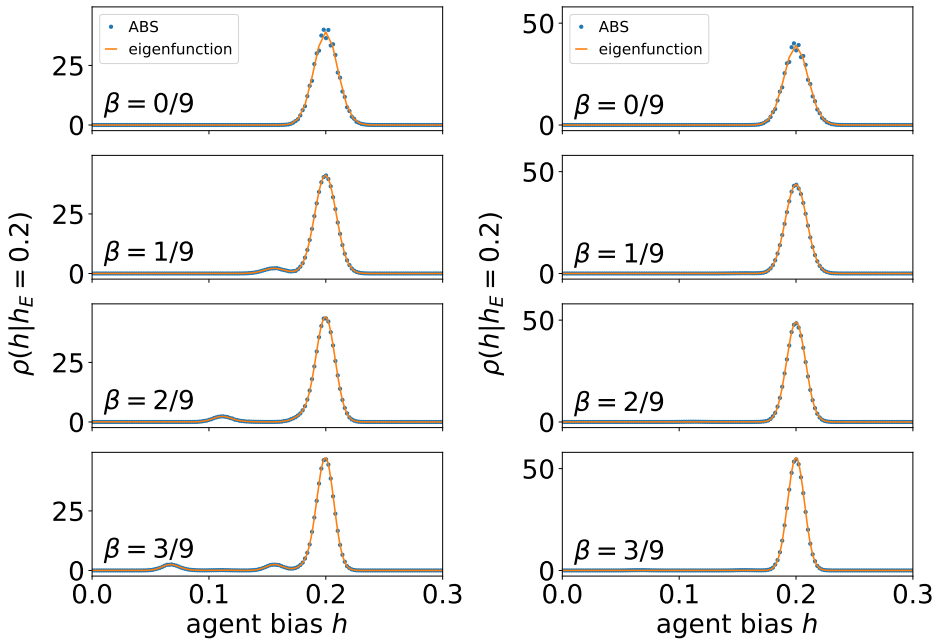

Fig. S7. Comparison of agent-based simulation and eigenfunction solution for the conditional probability distribution of agent bias $q(h|h_0)$ for (left) a passive agent and (right) stabilizer. Agent-based simulation returns a normalized histogram that aligns closely with the eigenfunction solution. Environment timescale $\tau_E = 20$ and bias $h_0 = 0.2$. Spacing of discrete domain in eigenfunction solution determined in proportion with typical width of the peak around $h = h_0$, which scales as in Eq 20 and inversely with the square root of agent memory τ_m.

obtain such precision, given by the KL divergence between environment configuration and agent behavior, $C \sim -\log_2(\sigma^2)$ leaving out constants. This depends on the variance of agent measurement noise $\sigma^2 = p_E(s,t)[1-p_E(s,t)]/\tau_c$. Infinitely precise sensory cells lead to diverging cost, whereas imprecise cells are cheap (Figure S6A). We also consider the cost G that stabilizers incur modifying the environment, here captured by stability, as the time-averaged KL divergence between the environmental decay rate with and without stabilizing effects (Eq 16). As we show in Figure S10, the cost is largest near optimal memory and decays to a constant for long agent memory τ_m that can be calculated exactly (Figure S6B). We account for these costs using information-theoretic measurements that present a generalizable formulation that could be extended to empirical measurements.

In Figure S10, we show each the divergence of a stabilizer without such costs in blue, each of these costs separately in black, and their sum in orange to generate the total divergence in Eq 18.

D. Evolution of reduced complexity

We consider a population of passive agents, or an agent with stabilization parameter $\alpha = 0$, precision timescale τ_c, and optimal memory τ_m^*, the variables that determine agent fitness. Assuming that the canonical equation for evolution applies (i.e. mutations only change phenotype and fitness slightly, the population dynamics move much faster than the evolutionary landscape such that we can assume a single phenotype dominates), the rate at which the population evolves across the phenotypic landscape is proportional to the fitness gradient. In addition to this assumption, we will assume that the population is always poised at optimal memory, an assumption that will be made clear below.

We recall that the total divergence consists of the time-averaged divergence $\bar{D}$, statistical complexity cost H, stabilization cost G, and precision cost C

$$\mathcal{D} = \bar{D} + \mu H(\tau_m) + \chi G(\tau_E, \tilde{\tau}_E) + \beta C(\tau_c) \quad [S2]$$

with semi-positive weights μ, χ, and β. In order to find the local dynamics of evolution, we must calculate the gradient $(\partial_{\tau_m}\mathcal{D}, \partial_\alpha\mathcal{D}, \partial_{\tau_c}\mathcal{D})$ determining the evolution in the properties of the agent. We calculate these term by term and then put them together at the end.

We assume that agent memory τ_m is at the minimum of the combination of time-averaged divergence $\bar{D}$ and statistical complexity cost μH (stabilization is zero for passive agents). Since divergence has a unique minimum and complexity monotonically approaches $H(\tau_m = \infty) = 0$, the addition of complexity only shifts optimal memory to a larger value. Without the complexity cost, we have that small deviations about optimal memory can be represented by a quadratic function for some positive constant a,

$$\bar{D} = D^* + a(\tau_m - \tau_m^*)^2, \quad [S3]$$

where we write

$$D^* = D_0(\tau_m^*)^{1/2} \quad [S4]$$

for some positive constant D_0. Once we have accounted for a perturbative addition from memory complexity, however, we have a shifted optimal memory

$$\tau_m^{**} = \tau_m^* + \frac{\mu}{2a\tau_m^*(\tau_m^*+1)} + \mathcal{O}(\mu^2) \quad [S5]$$

obtained from $\partial_{\tau_m}[\bar{D} + \mu H] = 0$. Then, the optimal divergence

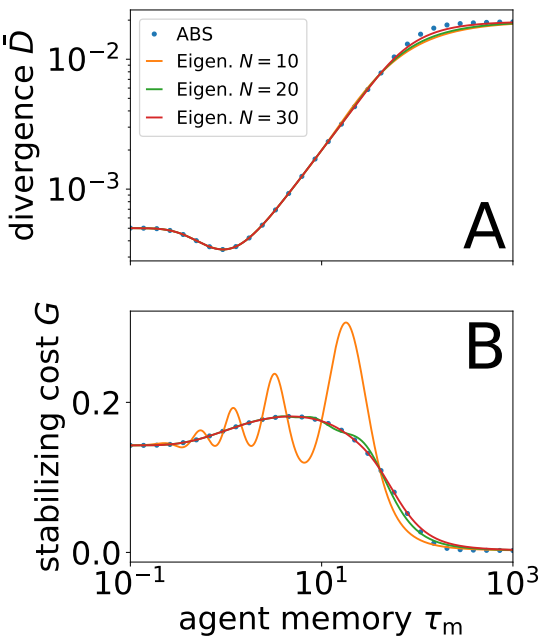

Fig. S8. Example of convergence of least-squares fit of Chebyshev polynomial with increasing number of abscissa N with the eigenfunction solution. (top) divergence D as calculated from the agent-based simulation (ABS) is closely modeled. The fit is close even with a relatively small number of points fit to a 9th-degree Chebyshev polynomial. Both methods are especially effective when the environmental timescale is small as is here, where $\tau_E = 10$. The bias $h_0 = 0.2$. (bottom) Stabilizing cost is similarly interpolated, but it is slower to converge with visible oscillations disappearing by $N = 30$. For $N = 20$ and $N = 30$, not all the points fell within the convergence criterion and only 19 and 28 points were fit, respectively. For both plots, the Chebyshev polynomial approximation is, unsurprisingly, slowest to converge near the sharp bends at large τ_m. ABS is run for 10^7 time steps.

becomes

$$\bar{D}'(\tau_m = \tau_m^{**}) = D^* + a\frac{\mu^2}{4(\tau_m^*)^2(\tau_m^*+1)^2} + \mathcal{O}(\mu^3). \qquad \text{[S6]}$$

Again, perturbations about the local optimum lead to

$$\bar{D}'(\tau_m) \approx D^* + a\frac{\mu^2}{4(\tau_m^*)^2(\tau_m^*+1)^2} + b(\tau_m - \tau_m^{**})^2 \qquad \text{[S7]}$$

for some positive constant b, which implicitly depends on the complexity cost. Eq S7 expresses local convexity about shifted optimal memory τ_m^{**} according to the corresponding shifted divergence $\bar{D}'$. This indicates how the population is poised along the ridge of optimal memory given a perturbative cost of memory complexity.

Then, time-averaged divergence will grow because optimal memory changes. Assuming the population is at optimal memory

$$\begin{aligned}\partial_\alpha \bar{D}' &= \partial_\alpha\left[D_0(\tau_m^*)^{-1/2} + a\frac{\mu^2}{4(\tau_m^*)^2(\tau_m^*+1)^2}\right] \\ &= \left[\frac{D_0}{2}(\tau_m^*)^{-3/2} + a\frac{\mu^2(2\tau_m^*+1)}{2(\tau_m^*)^3(\tau_m^*+1)^3}\right]\left|\frac{\partial\tau_m^*}{\partial\alpha}\right|,\end{aligned} \qquad \text{[S8]}$$

where we have used the fact that optimal memory must increase with stronger stabilizer, or that $\partial_\alpha \tau_m^* < 0$, to explicitly pull out a negative sign. This confirms that in Eq S8 that divergence at optimal memory decreases as α approaches -1 from above as expected. We have also assumed that we can discard terms higher order than linear in μ.

Niche-constructing stabilization changes the environmental timescale through feedback. We start by considering over a long period of time the average over many environmental switches

$$\begin{aligned}\langle 1/\tilde{\tau}_E\rangle &= 1/\tau_E + \alpha\left\langle\frac{v^2}{v^2+(h-h_E)^2}\right\rangle, \\ &= 1/\tau_E + \alpha f(\tau_m).\end{aligned} \qquad \text{[S9]}$$

Since we do not know the exact form of the second term on the right hand side, we represent it as some function f that represents an average over time. For notational simplicity, we only make explicit f's dependence on τ_m, but it depends on agent properties and environmental timescale. Now, a change in α also indirectly affects τ_m^* because the environmental timescale will change, reducing or increasing the agents ability to track the new environment. For example, with the passive agent, an increase in α introduces environmental stabilization, driving the effective environmental timescale slower and moving the optimal memory timescale up. Accounting for these derivatives means that

$$d_\alpha\langle 1/\tilde{\tau}_E\rangle = f(\tau_m) + \alpha\partial_{\tau_m} f(\tau_m)\partial_\alpha\tau_m. \qquad \text{[S10]}$$

Now, we will again make use of the assumption that τ_m is close τ_m^* such that we can make the linear approximation $f(\tau_m) \approx f(\tau_m^*) + (\tau_m - \tau_m^*)f'(\tau_m^*)$. Putting this in, we find

$$\begin{aligned}d_\alpha\langle 1/\tilde{\tau}_E\rangle(\tau_m) = f(\tau_m^*) + (\tau_m-\tau_m^*)f'(\tau_m^*) + \\ \alpha\partial_{\tau_m}[f(\tau_m^*) + (\tau_m-\tau_m^*)f'(\tau_m^*)]\partial_\alpha\tau_m^*\end{aligned} \qquad \text{[S11]}$$

For a passive agent, this simplifies because we know $\alpha = 0$. Furthermore, we know that $f'(\tau_m^*) = 0$ because we have assumed that the agent is at optimal memory so any deviation from optimal memory must generally increase the typical distance between environmental and agent bias $(h - h_E)^2$. Then,

$$d_\alpha\langle 1/\tilde{\tau}_E\rangle(\tau_m) = f(\tau_m^*). \qquad \text{[S12]}$$

Eq S12 is already clear from Eq S9 given the assumptions we have made, but these steps take us through the general problem (when not situated exactly at optimal memory and when $\alpha \neq 0$ are more complicated). In other words, decreasing α for the weak stabilizer will reduce the probability that the environment switches by the term in Eq S12 because $f > 0$ and f' — the change in probability is not just dependent on the rate effect f but also its derivative.

Under such a change, the new environmental timescale will deviate from τ_E and so the stabilization cost can be expanded as

$$\begin{aligned}G(\tau_E,\tilde{\tau}_E) &= \frac{1}{\tau_E}\log\left(\frac{1/\tau_E}{\langle 1/\tilde{\tau}_E\rangle}\right) + \left(\frac{1}{\tau_E}\right)\log\left(\frac{1-1/\tau_E}{1-\langle 1/\tilde{\tau}_E\rangle}\right) \\ &\approx \frac{1}{2\tau_E}\left[\langle 1/\tilde{\tau}_E\rangle - 1/\tau_E\right]^2 + \\ &\quad \frac{1}{2}\left(1-\frac{1}{\tau_E}\right)\left[\langle 1/\tilde{\tau}_E\rangle - 1/\tau_E\right]^2 \\ &= \frac{1}{2}\left[\langle 1/\tilde{\tau}_E\rangle - 1/\tau_E\right]^2,\end{aligned} \qquad \text{[S13]}$$

a cost that increases quadratically with the change in the averaged switch probability $\langle 1/\tilde{\tau}_E\rangle$ away from $1/\tau_E$. For a passive agent, this direction is 0 unless we allow for α to vary, which leads to the relation

$$G(\tau_E,\tilde{\tau}_E) = \frac{\alpha^2}{2}f(\tau_m^*)^2. \qquad \text{[S14]}$$

Eq S14 tells us that if we vary α, we must pay a stabilization cost that, at least locally, grows quadratically with the strength of stabilization with zero gradient.

The simplest contribution is with respect to the change in the precision timescale τ_c. Divergence, as derived in Materials & Methods, is proportional to $1/\tau_c$. On the other hand, precision cost is $C = \log\tau_c$. Since optimal memory timescale does not depend on τ_c, the change of the total divergence is

$$\partial_{\tau_c}\left[D_1/\tau_c + \beta\log_2\tau_c\right] = -D_1/{\tau_c}^2 + \beta/\tau_c, \qquad \text{[S15]}$$

where we take $D^* = D_1/\tau_c$ to encapsulate the terms in the divergence apart from the scaling with precision timescale. If this has a minimum at positive τ_c, the value of τ_c at which the minimum is reached is ${\tau_c}^* = D_1/\beta$.

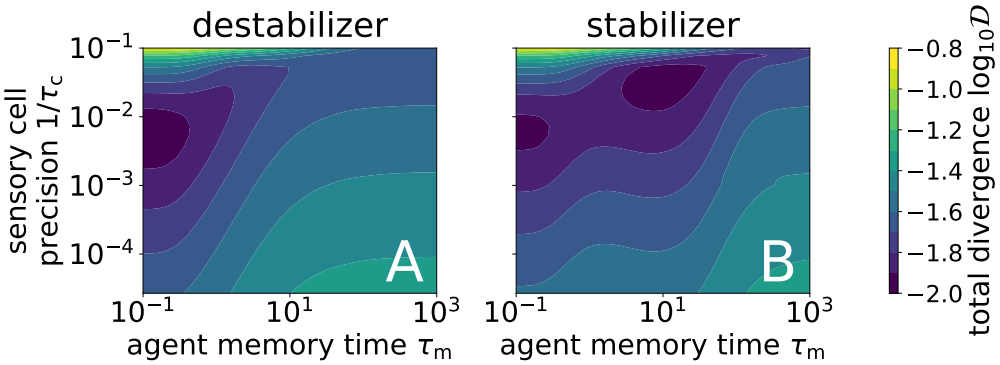

Fig. S9. Divergence for destabilizers and stabilizers when accounting for information costs. (B) Stabilizers, in this example, have two degenerate minima. Environment timescale $\tau_E = 10$, environmental bias $h_0 = 0.2$, and cost weights $\chi = 0.004$ and $\mu = \beta = 0.01$.

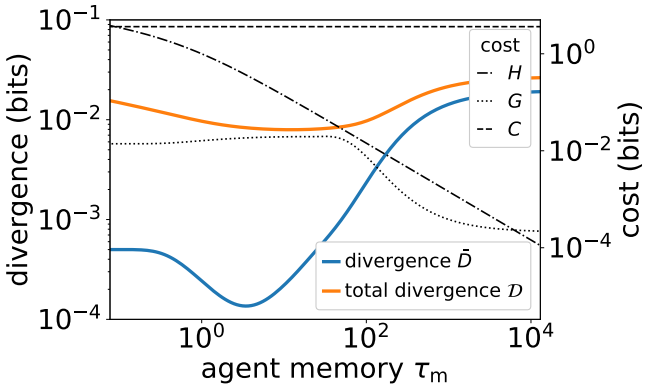

Fig. S10. Example of cost functions for stabilizers with varying memory but fixed sensory precision. (blue) Without costs, divergence profile shows only a single global minimum. (orange) With costs, we obtain degenerate minima at memory values around $\tau_m = 0$ and $\tau_m = 20$. Eigenfunction solution parameters specified in Materials & Methods code.

Putting all of these together, we have the terms in the gradient

$$\begin{aligned}
\partial_{\tau_m}\mathcal{D} &= 2a(\tau_m - \tau_m^*) \\
\partial_\alpha \mathcal{D} &= \left[\frac{D_0}{2}(\tau_m^*)^{-3/2} + a\frac{\mu^2(2\tau_m^*+1)}{2(\tau_m^*)^3(\tau_m^*+1)^3}\right]\left|\frac{\partial\tau_m^*}{\partial\alpha}\right| + \\
&\quad \chi\left[2\alpha f(\tau_m^*)^2 + \alpha^2 f(\tau_m^*)\partial_\alpha f(\tau_m^*)\partial_\alpha\tau_m^*\right] \qquad \text{[S16]} \\
\partial_{\tau_c}\mathcal{D} &= \beta/\tau_c - D_1/{\tau_c}^2
\end{aligned}$$

When the cost gradient $\partial_\alpha\mathcal{D} < 0$, a population of passive agents is driven towards niche construction and when $\partial_{\tau_c}\mathcal{D} < 0$ towards precision reduction.

Note that a similar derivation can be made for stabilizers, or agents with $\alpha < 0$. However, this requires us to use the full derivative with respect to α instead of assuming it to be zero. Furthermore, the stabilization cost is no longer at a minimum and will instead contribute to the gradient. The change in the environmental timescale is difficult to calculate analytically, but it is clear that the qualitative results will be the same because increasing α will drive effective environmental timescale up, but the exact rate at which it increases will depend on the local curvature of the function. Thus, the conditions that lead to reduction in agent complexity by increasing memory, enhancing stabilization, and lowering precision are captured by these gradients.

1. Feldman MW, Laland KN (1996) Gene-culture coevolutionary theory. *Trends in Ecology & Evolution* 11(11):453–457.
2. Gerbault P, et al. (2011) Evolution of lactase persistence: an example of human niche construction. *Philosophical Transactions of the Royal Society B: Biological Sciences* 366(1566):863–877.
3. Odling-Smee FJ, Laland KN, Feldman MW (1996) Niche Construction. *The American Naturalist* 147(4):641–648. Publisher: The University of Chicago Press.
4. Laland KN, O'Brien MJ (2011) Cultural Niche Construction: An Introduction. *Biological Theory* 6(3):191–202.
5. Clark AD, Deffner D, Laland K, Odling-Smee J, Endler J (2020) Niche Construction Affects the Variability and Strength of Natural Selection. *The American Naturalist* 195(1):16–30.
6. Buser CC, Newcomb RD, Gaskett AC, Goddard MR (2014) Niche construction initiates the evolution of mutualistic interactions. *Ecology Letters* 17(10):1257–1264. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/ele.12331.
7. Bowles S (2006) *Microeconomics*.
8. Poon P, Flack JC, Krakauer DC (2022) Institutional dynamics and learning networks. *PLOS ONE* 17(5):e0267688. Publisher: Public Library of Science.
9. Ofosu EK, Chambers MK, Chen JM, Hehman E (2019) Same-sex marriage legalization associated with reduced implicit and explicit antigay bias. *Proceedings of the National Academy of Sciences* 116(18):8846–8851. Publisher: Proceedings of the National Academy of Sciences.
10. Flack JC (2017) Coarse-graining as a downward causation mechanism. *Phil. Trans. R. Soc. A* 375(2109):20160338.
11. Flack J (2017) Life's information hierarchy in *From Matter to Life: Information and Causality*, eds. Ellis GFR, Davies PCW, Walker SI. (Cambridge University Press, Cambridge), pp. 283–302.
12. Mukherjee S, Bassler BL (2019) Bacterial quorum sensing in complex and dynamically changing environments. *Nature Reviews Microbiology* 17(6):371–382. Number: 6 Publisher: Nature Publishing Group.
13. Merton RK (1948) The Self-Fulfilling Prophecy. *The Antioch Review* 8(2):193–210. Publisher: Antioch Review, Inc.
14. Strathern M (1997) Improving ratings: audit in the British University system. *European Review* 5(3). Publisher: Cambridge University Press.
15. Soros G (2013) Fallibility, reflexivity, and the human uncertainty principle. *Journal of Economic Methodology* 20(4):309–329. Publisher: Routledge _eprint: https://doi.org/10.1080/1350178X.2013.859415.
16. Manheim D, Garrabrant S (2019) Categorizing Variants of Goodhart's Law. *arXiv:1803.04585 [cs, q-fin, stat]*. arXiv: 1803.04585.
17. Jakab S (2022) *The Revolution That Wasn't: GameStop, Reddit, and the Fleecing of Small Investors*. (Portfolio, New York, NY).
18. Kalman RE (1960) A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering* pp. 35–45.
19. Brenner N, Bialek W, de Ruyter van Steveninck R (2000) Adaptive Rescaling Maximizes Information Transmission. *Neuron* 26(3):695–702.
20. Gershman SJ, Radulescu A, Norman KA, Niv Y (2014) Statistical Computations Underlying the Dynamics of Memory Updating. *PLoS Comput Biol* 10(11):e1003939.
21. Davis RL, Zhong Y (2017) The Biology of Forgetting—A Perspective. *Neuron* 95(3):490–503.
22. Ratcliff R (1978) A Theory of Memory Retrieval. *Psychological Review* 85(2):59–108.
23. Chowdhury SM, Kovenock D, Sheremeta RM (2013) An experimental investigation of Colonel Blotto games. *Economic Theory* 52(3):833–861.
24. Krakauer D, Bertschinger N, Olbrich E, Flack JC, Ay N (2020) The information theory of individuality. *Theory in Biosciences* 139(2):209–223.
25. Theraulaz G, Bonabeau E (1999) A Brief History of Stigmergy. *Artificial Life* 5(2):97–116.
26. Brush ER, Krakauer DC, Flack JC (2018) Conflicts of interest improve collective computation of adaptive social structures. *Science Advances* 4(1):e1603311. Publisher: American Association for the Advancement of Science.
27. Ramos-Fernandez G, Smith Aguilar SE, Krakauer DC, Flack JC (2020) Collective Computation in Animal Fission-Fusion Dynamics. *Frontiers in Robotics and AI* 7.
28. Couzin ID, Krause J, Franks NR, Levin SA (2005) Effective leadership and decision-making in animal groups on the move. *Nature* 433(7025):513–516.
29. Franks NR, et al. (2007) Reconnaissance and latent learning in ants. *Proc. R. Soc. B.* 274(1617):1505–1509.
30. Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD (2006) The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol Rev* 113(4):700–765.
31. Brunton BW, Botvinick MM, Brody CD (2013) Rats and Humans Can Optimally Accumulate Evidence for Decision-Making. *Science* 340(6128):95–98.

32. McNamara JM, Houston AI (1987) Memory and the efficient use of information. *Journal of Theoretical Biology* 125(4):385–395.
33. Donaldson-Matasci MC, Bergstrom CT, Lachmann M (2010) The fitness value of information. *Oikos* 119(2):219–230.
34. Krakauer DC (2011) Darwinian demons, evolutionary complexity, and information maximization. *Chaos: An Interdisciplinary Journal of Nonlinear Science* 21(3):037110.
35. Kelly JL (1956) A New Interpretation of Information Rate. *the bell system technical journal* p. 10.
36. Cover TM, Thomas JA (2006) *Elements of Information Theory*. (John Wiley & Sons, Hoboken), Second edition.
37. Conti D, Mora T (2020) Non-equilibrium dynamics of adaptation in sensory systems. *arXiv:2011.09958 [nlin, q-bio]*.
38. West GB (1999) The Fourth Dimension of Life: Fractal Geometry and Allometric Scaling of Organisms. *Science* 284(5420):1677–1679.
39. White CR, Seymour RS (2003) Mammalian basal metabolic rate is proportional to body mass2/3. *Proceedings of the National Academy of Sciences* 100(7):4046–4049.
40. Burger JR, George MA, Leadbetter C, Shaikh F (2019) The allometry of brain size in mammals. *Journal of Mammalogy* 100(2):276–283.
41. West GB, Woodruff WH, Brown JH (2002) Allometric scaling of metabolic rate from molecules and mitochondria to cells and mammals. *Proceedings of the National Academy of Sciences* 99(Supplement 1):2473–2478.
42. Savage VM, Deeds EJ, Fontana W (2008) Sizing Up Allometric Scaling Theory. *PLoS Comput. Biol.* 4(9):e1000171.
43. Snell-Rood EC, Papaj DR, Gronenberg W (2009) Brain Size: A Global or Induced Cost of Learning? *Brain Behav Evol* 73(2):111–128.
44. Liefting M, Rohmann JL, Le Lann C, Ellers J (2019) What are the costs of learning? Modest trade-offs and constitutive costs do not set the price of fast associative learning ability in a parasitoid wasp. *Anim Cogn* 22(5):851–861.
45. Woude E, Groothuis J, Smid HM (2019) No gains for bigger brains: Functional and neuroanatomical consequences of relative brain size in a parasitic wasp. *J Evol Biol* p. jeb.13450.
46. Klyubin AS, Polani D, Nehaniv CL (2004) Tracking Information Flow through the Environment: Simple Cases of Stigmergy in *Artificial Life IX: Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems*, ed. Pollack J. (MIT Press).
47. Krakauer DC, Page KM, Erwin DH (2009) Diversity, Dilemmas, and Monopolies of Niche Construction. *The American Naturalist* 173(1):26–40.
48. Wen XL, Wen P, Dahlsjö CAL, Sillam-Dussès D, Šobotník J (2017) Breaking the cipher: Ant eavesdropping on the variational trail pheromone of its termite prey. *Proc. R. Soc. B.* 284(1853):20170121.
49. Smith CC, Reichman OJ (1984) The Evolution of Food Caching by Birds and Mammals. *Ann. Rev. Ecol. Syst.* 15:329–351.
50. Hall K, et al. (2017) Chimpanzee uses manipulative gaze cues to conceal and reveal information to foraging competitor. *Am J Primatol* 79(3):e22622.
51. Brännström Å, Johansson J, von Festenberg N (2013) The Hitchhiker's Guide to Adaptive Dynamics. *Games* 4(3):304–328.
52. Dieckmann U (year?) Ph.D. thesis.
53. Dieckmann U, Law R (1996) The dynamical theory of coevolution: A derivation from stochastic ecological processes. *J. Math. Biology* 34(5-6):579–612.
54. Koehl MAR, Cooper T (2015) Swimming in an Unsteady World. *Integr. Comp. Biol.* 55(4):683–697.
55. Evans MEK, Dennehy JJ (2005) Germ Banking: Bet-Hedging and Variable Release from Egg and Seed Dormancy. *The Quarterly Review of Biology* 80(4):431–451.
56. Slivkins A (2019) *Introduction to Multi-Armed Bandits*. (Now Publishers).
57. Kaspi H, Mandelbaum A (1995) Levy Bandits: Multi-Armed Bandits Driven by Levy Processes. *Ann. Appl. Probab.* 5(2):541–565.
58. Schofield N (2008) Divergence in the spatial stochastic model of voting in *Power, Freedom, and Voting*, eds. Braham M, Steffen F. (Springer Berlin Heidelberg, Berlin, Heidelberg), pp. 259–287.
59. Yuan DL, Chen Q (2010) Particle swarm optimisation algorithm with forgetting character. *International Journal of Bio-Inspired Computation* 2(1):59.
60. Tindall MJ, Gaffney EA, Maini PK, Armitage JP (2012) Theoretical insights into bacterial chemotaxis: Theoretical bacterial chemotaxis. *WIREs Syst Biol Med* 4(3):247–259.
61. Bifet A, Gavaldà R (2007) Learning from Time-Changing Data with Adaptive Windowing in *Proceedings of the 2007 SIAM International Conference on Data Mining*. (Society for Industrial and Applied Mathematics), pp. 443–448.
62. Kosina P, Gama J (2012) Handling time changing data with adaptive very fast decision rules in *Machine Learning and Knowledge Discovery in Databases*, eds. Flach PA, De Bie T, Cristianini N. (Springer Berlin Heidelberg, Berlin, Heidelberg), pp. 827–842.
63. Rolls ET, Deco G (2015) Stochastic cortical neurodynamics underlying the memory and cognitive changes in aging. *Neurobiology of Learning and Memory* 118:150–161.
64. Ratcliff R, Rouder JN (1998) Modeling Response Times for Two-Choice Decisions. *Psychol Sci* 9(5):347–356.
65. Miletic S, Boag R, Mathiopoulou V, Forstmann B (2019) Speed and accuracy in learning: A combined Q-learning diffusion decision model analysis in *2019 Conference on Cognitive Computational Neuroscience*. (Cognitive Computational Neuroscience, Berlin, Germany).
66. Kar S, Moura J (2009) Distributed Consensus Algorithms in Sensor Networks With Imperfect Communication: Link Failures and Channel Noise. *IEEE Trans. Signal Process.* 57(1):355–369.
67. Schweitzer ME, Cachon GP (2000) Decision Bias in the Newsvendor Problem with a Known Demand Distribution: Experimental Evidence. *Management Science* 46(3):404–420.
68. Tregenza T (1995) Building on the Ideal Free Distribution in *Advances in Ecological Research*. (Elsevier) Vol. 26, pp. 253–307.
69. Musa HH, Noureldien A (2018) Comparing the ranking performance of page rank algorithm and weighted page rank algorithm. *Advanced Science Letters* 24(1):750–753.
70. Kirkegaard JB, Goldstein RE (2018) The role of tumbling frequency and persistence in optimal run-and-tumble chemotaxis. *IMA Journal of Applied Mathematics* 83(4):700–719.
71. Bonabeau E, Theraulaz G, Deneubourg JL, Aron S, Camazine S (1997) Self-organization in social insects. *TREE* 12(5):188–193.
72. Dorigo M, Bonabeau E, Theraulaz G (2000) Ant algorithms and stigmergy. *Future Generation Computer Systems* 16(8):851–871.
73. Lee ED, Daniels BC, Krakauer DC, Flack JC (2017) Collective memory in primate conflict implied by temporal scaling collapse. *J. R. Soc. Interface* 14(134):20170223.
74. Hardt O, Nader K, Nadel L (2013) Decay happens: The role of active forgetting in memory. *Trends in Cognitive Sciences* 17(3):111–120.
75. Wixted JT (2004) The Psychology and Neuroscience of Forgetting. *Annu. Rev. Psychol.* 55(1):235–269.
76. Tello-Ramos MC, Branch CL, Kozlovsky DY, Pitera AM, Pravosudov VV (2019) Spatial memory and cognitive flexibility trade-offs: To be or not to be flexible, that is the question. *Animal Behaviour* 147:129–136.
77. Brush ER, Krakauer DC, Flack JC (2013) A Family of Algorithms for Computing Consensus about Node State from Network Data. *PLoS Comput Biol* 9(7):e1003109.
78. Daniels BC, Krakauer DC, Flack JC (2012) Sparse code of conflict in a primate society. *Proc. Natl. Acad. Sci. U.S.A.* 109(35):14259–14264.
79. Press WH, Teukolsky SA, Vetterling WT, Flannery BP (2007) *Numerical Recipes: The Art of Scientific Computing*. (Cambridge University Press, New York), 3rd edition.
80. Szilard L (1929) On Entropy Reduction in a Thermodynamic System by Interference by Intelligent Subjects. *Zeitschrift für Physik* 53:840–856.
81. Bennett CH (1987) Demons, Engines and the Second Law. *Sci Am* 257(5):108–116.
82. Alatalo RV, Carlson A, Lundberg A (1988) The search cost in mate choice of the pied flycatcher. *Animal Behaviour* 36(1):289–291.
83. Franks NR, Dornhaus A, Fitzsimmons JP, Stevens M (2003) Speed versus accuracy in collective decision making. *Proc. R. Soc. Lond. B* 270(1532):2457–2463.
84. Kylafis G, Loreau M (2008) Ecological and evolutionary consequences of niche construction for its agent. *Ecology Letters* 11(10):1072–1081.
85. Menge DNL, Levin SA, Hedin LO (2008) Evolutionary tradeoffs can select against nitrogen fixation and thereby maintain nitrogen limitation. *Proceedings of the National Academy of Sciences* 105(5):1573–1578.
86. Mohlenhoff KA, Codding BF (2017) When does it pay to invest in a patch? The evolution of intentional niche construction. *Evol. Anthropol.* 26(5):218–227.
87. Shannon CE (1948) A Mathematical Theory of Communication. *Bell Syst. Tech. J.* 27:379–423, 623–656.