

## Opinion

## Early adversity and the development of explore–exploit tradeoffs

Willem E. Frankenhuis<sup>1,2,4,\*</sup> and Alison Gopnik<sup>3,4</sup>

**Childhood adversity can have wide-ranging and long-lasting effects on later life. But what are the mechanisms that are responsible for these effects? This article brings together the cognitive science literature on explore–exploit tradeoffs, the empirical literature on early adversity, and the literature in evolutionary biology on ‘life history’ to explain how early experience influences later life. We propose one potential mechanism: early experiences influence ‘hyperparameters’ that determine the balance between exploration and exploitation. Adversity might accelerate a shift from exploration to exploitation, with broad and enduring effects on the adult brain and mind. These effects may be produced by life-history adaptations that use early experience to tailor development and learning to the likely future states of an organism and its environment.**

**Bridging life history and learning**

Early experiences can have broad and lasting effects on later health and thriving [1–3]. These effects are clearly not deterministic; many people are resilient, but early adversity increases the chance of later difficulties. The practical importance of these findings is clear. But they pose a crucial question: how and why does this happen? One possibility is that there is a cascade of causes: childhood adversity tends to lower school performance, which tends to lower incomes, which tends to lower health, and so on. Another possibility is that early adversity may damage the developing organism, leading to deficits that are difficult to overcome. Both these processes are likely to play a part in some cases. However, they do not seem to explain the coherent and wide-ranging effects of early experience, in particular, the consistent acceleration of various aspects of development (see later).

We hypothesize that early experiences influence **hyperparameters** (see [Glossary](#)), which control processes of **learning** and development. These parameters determine how learning takes place rather than the content of what is learned [4–7]. Hyperparameters have been employed to explain many types of learning and development, from **phenotypic plasticity** to **sensitive periods** to reinforcement learning to Bayesian causal learning. The values of hyperparameters may themselves be acquired or modified based on experience [8–12]. Here, we focus on a particular type of hyperparameter that controls the balance between exploration and exploitation. This hyperparameter appears to vary between individuals in ways that influence behavior across multiple different contexts [13], though more work on generalizability is needed.

The explore–exploit contrast has been discussed most often in the literature on reinforcement learning (e.g., [8]), but it applies much more widely. There are many cases, across diverse domains, where an agent must choose between searching widely versus within a narrower range of options, prioritizing information versus reward, or maintaining flexibility versus increasing specialization. These are all instances of a more general **tradeoff** between exploring the space of possible solutions and exploiting a particular solution: a tradeoff formalized in optimization theory

**Highlights**

This opinion article brings together research on the effects of early adversity, computational accounts of learning, and life-history theory.

These bodies of research have independently converged on the finding that exposures to chronic adversity may accelerate maturation, yet synergies between them have hardly been explored.

We connect these bodies of research by proposing the hypothesis that early experiences influence hyperparameters that determine the balance between exploration and exploitation.

Specifically, we argue that childhood adversity may accelerate a shift from exploration to exploitation, with wide-ranging effects on the adult brain and mind.

Understanding which types of adversity accelerate the explore–exploit shift, and formal theory exploring when such a response may be adaptive, are key directions for future work.

<sup>1</sup>Department of Psychology, Utrecht University, Utrecht, The Netherlands

<sup>2</sup>Max Planck Institute for the Study of Crime, Security and Law, Freiburg, Germany

<sup>3</sup>Department of Psychology and Berkeley Artificial Intelligence Research, University of California at Berkeley, CA, USA

<sup>4</sup>These authors contributed equally to this work.

\*Correspondence:  
[w.e.frankenhuis@uu.nl](mailto:w.e.frankenhuis@uu.nl)  
(W.E. Frankenhuis).



[14]. Exploration and exploitation are ends of a continuum – from broad to narrow, noisy to efficient, oriented to information or to reward – rather than a dichotomy.

Recent evidence suggests that there is a characteristic developmental shift from exploration to exploitation (see [4] for a review). We propose that early adversity may accelerate this shift. Moreover, we argue that acceleration is produced by evolved life-history adaptations that use early-life **cues** – observations that provide information about the state of the organism or the environment – to tailor the organism to anticipate future conditions. To make this argument, we bring together three different literatures: computational accounts of learning and development in cognitive science, empirical studies of early adversity, and **life-history theory** (Figure 1).

Bridging these research programs is also timely because these communities are already crossing this divide. First, neuroscientists are exploring the effects of adverse environments on the pace of brain development. This work is often guided by the hypothesis that, if energetic resources are adequate for growth [15,16], chronic adversity may accelerate the rate of brain maturation, in concert with life-history traits (e.g., puberty [17–20]), to prepare juveniles for early independence and reproduction [21–24]. If an explore–exploit shift is instantiated in processes of brain maturation, as seems plausible, these neural accelerations might reflect an acceleration in that shift, produce such an acceleration, or both. The proximate phenomenon of brain maturation acceleration might ultimately be explained by the evolutionary advantages of an earlier shift to exploitation, which would facilitate early independence.

Second, evolutionary biologists are studying how plasticity and learning develop in concert with classic life-history traits – such as longevity, the age at first birth, and grandparental investment

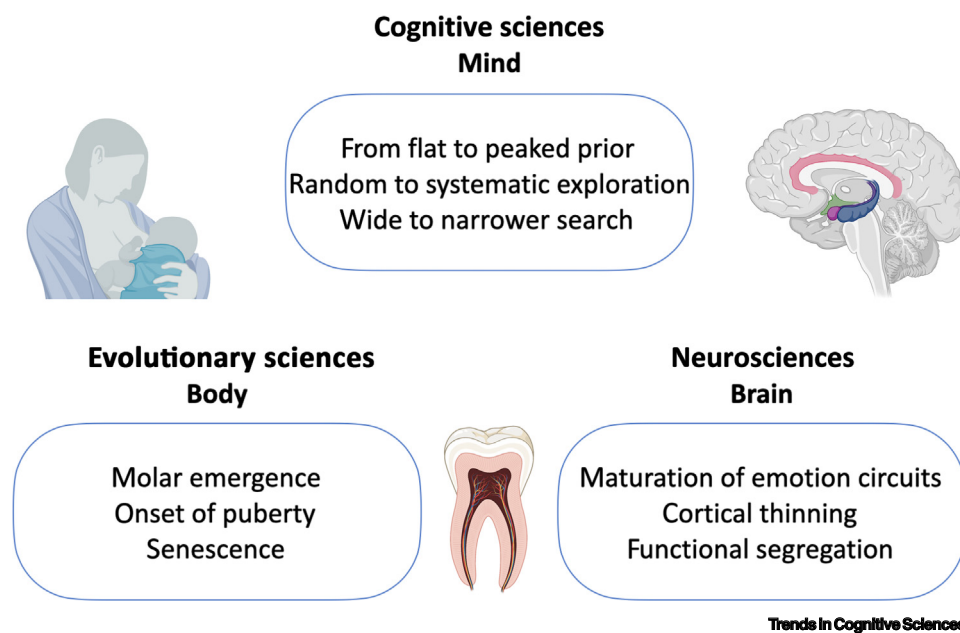


Figure 1. Developmental changes that may be accelerated by adversity. Over the past decade, three independent bodies of work have converged on findings suggesting that chronic adversity may accelerate particular developmental processes in the body, brain, and mind. We argue that these effects are produced by evolved adaptations, favored by natural selection, that use early experience to tailor individuals to anticipated future conditions. Although these bodies of work have operated largely independently, there is great scope for synergy between them. Figure created with [BioRender.com](https://www.biorender.com).

## Glossary

**Ancestral cue hypothesis:** the mind has evolved to respond to specific events (e.g., fights) that provide information about environmental conditions, current and future (e.g., mortality rate in the broader ecology).

**Cue:** an observation that provides information (i.e., reduces uncertainty) about the state of the organism or the environment.

**Directed exploration:** an agent focuses its search where it is most likely to gain new information.

**Hyperparameter:** a parameter that is set before learning begins, and whose value controls the learning process, rather than being estimated during learning.

**Learning:** the acquisition or modification of information, abilities, or responses as a result of experience.

**Life-history theory:** a framework for studying how organisms allocate limited resources (e.g., time, energy) between different activities (e.g., growth, reproduction) across the lifespan.

**Local optimum:** the best solution to a problem within a small neighborhood of possible solutions; all small changes will lead to worse options, but a large shift might lead to a substantially better solution (e.g., the global optimum).

**Multiarmed bandit task:** a classic reinforcement learning task modeled on a 'one-armed bandit' slot machine. Agents must choose between different options with different past histories of reward. An exploit choice is to always choose the arm with the greatest past history of reward. But exploration, by occasionally choosing the other arm, allows the agent to detect whether the environment has changed.

**Phenotypic plasticity:** the ability of an organism to change its phenotype in response to environmental conditions.

**Sensitive period:** a period or life stage in which experience shapes a phenotype to a larger extent than the same experience does in other periods or life stages.

**Simulated annealing:** a computational technique for searching through complex high-dimensional spaces to find accurate problem solutions, which typically begins with a broad 'high-temperature' search that visits a wider range of lower-probability solutions and then 'cools off' to settle on a narrower and less variable search.

[9,25,26] – and to formally model these relationships [27–30]. The time is ripe to connect these communities by bridging their contributions, delineating gaps, and outlining future directions [16].

The earlier life-history approaches have focused on developments (e.g., puberty onset or longevity) that have implications for reproductive success. We focus instead on hyperparameters that shape the way that organisms adapt to their environment during the lifespan, through both learning and development. We propose that early experiences shape these parameters in ways that on average increased survival and reproduction during human evolution, by facilitating earlier independence and exploitation in adverse environments, but allowing a longer period of exploration in more favorable ones. We do not claim that these responses are adaptive in contemporary societies.

We begin by explaining the explore–exploit tradeoff in cognitive science. Then we review empirical research suggesting that development does indeed involve a shift from exploration to exploitation. We then review evidence suggesting that different types of adversity accelerate various developmental processes. We hypothesize that adversity may have an accelerative effect on the explore–exploit shift, making it more likely that children move to a narrower, less exploratory, less plastic learning mode at an earlier age. Then we discuss the relationship of these ideas to formal models in life-history biology. We conclude by discussing future directions and policy implications.

### The explore–exploit tradeoff

Many types of learning and development can be understood as searches through a high-dimensional space of possible states, hypotheses, or policies [14,31]. An agent begins at a particular point in that space – with a particular set of states, hypotheses, or policies – and then moves through the space to discover new ones. There is an intrinsic tradeoff between searching widely, with the chance of finding optimal outcomes, and searching more narrowly for ‘good enough’ solutions that can be quickly implemented. A narrower, more exploitative search is more likely to lead to an effective solution quickly, but may leave the agent stuck in a **local optimum**. In local optima, all small changes will lead to worse options, but a large shift might lead to a substantially better solution.

Exploitation allows an agent to accumulate resources in the short run. Exploration is costly as it requires resources to learn about the environment. However, exploration may pay off in the longer run, particularly when the environment is complex and variable and time horizons are long. If the search space is large and has many dimensions, there is no simple way to resolve this tradeoff. However, different computational strategies ensure that exploration takes place despite its short-term cost [6,8,32,33] (Box 1).

A common approach is to begin with more exploration and gradually converge on exploitation (e.g., [34]). For example, in **simulated annealing**, systems begin with a higher-**temperature**, noisier, more random search, and gradually cool off to a more detailed and focused one [14]. Similarly, machine-learning systems typically start with a hyperparameter specifying a high learning rate – that is, the model changes more easily in response to new data leading to more exploration – and gradually decrease that rate, so that new data have less effect. This shift can take place on many different timescales. Some machine-learning techniques employ cycles of exploration followed by exploitation (e.g., [35]).

Biological organisms may employ a similar technique over their lifespan, beginning with a protected exploratory period, in which resources are provided by others, and gradually shifting to a more competent period of exploitation [4,36–38]. In this view, childhood itself could be construed as an evolved life-history solution to the explore–exploit tradeoff [4]. Consistent with this

**Statistical learning hypothesis:** the mind tracks the statistical structure of the environment by integrating experiences across ontogeny, including learning evolutionarily novel cues, updating cue weights, and prediction errors.

**Stress acceleration hypothesis:** early adversity leads to accelerated maturation of brain systems supporting early independence and reproduction.

**Temperature:** a feature of computational searches through a complex high-dimensional space of possibilities, where high-temperature searches are broader, more variable, more random, and more likely to consider low-probability possibilities than lower-temperature searches.

**Tradeoffs:** investments (e.g., time, energy) made in one activity (e.g., growth) reduce the resources available for investments in other activities (e.g., reproduction).

### Box 1. Computational strategies to ensure exploration

There are different specific computational strategies that instantiate the balance between exploration and exploitation. One way to ensure exploration may simply be to search more widely. In Bayesian learning, if some hypotheses have a much higher initial probability than others (a 'peaked' prior), they are more likely to be sampled or considered, and the search will be narrowed. If different hypotheses have more similar initial probabilities (a 'flat' prior), a wider range of hypotheses will be considered [52,53]. If agents start with flatter priors, which become more peaked as they learn more, this would naturally implement an explore-then-exploit sequence. The shift from flat to peaked priors characterizes human learning and might seem obvious and intuitive. But agents could also start with a highly peaked prior, with strong innate preferences for certain hypotheses or policies, and gradually flatten that distribution with experience. This does indeed seem to be the case for 'precocial' species, such as chickens, who mature quickly and rely more on innate structure than learning, and who may use general associative learning to modify that structure later.

Another group of strategies involves adding random noise to the search process. A common technique is to change the temperature of the search, using a 'hotter', noisier, more random search initially to ensure exploration and escape local optima and gradually 'cooling off' [8,14,33]. For example, in a reinforcement learning multiarmed bandit task, an agent must choose between a policy that has been more rewarding in the past and a less rewarding one. An agent might simply randomly decide to occasionally sample the less rewarding arm of the bandit: the explore rather than exploit choice. Human children do indeed seem to engage in more random exploration than adults [6,52].

Yet another strategy is to prioritize obtaining relevant information about the environment over immediate reward. There is considerable evidence that even infants are intrinsically motivated to seek information, and by early childhood children will consistently act to gain more information [87]. In **directed exploration** in reinforcement learning, agents prioritize actions that are most likely to produce relevant information rather than responding randomly. There is evidence that both adults and children produce directed as well as random exploration [33,51], but also that children are less strategic in their use of directed exploration than adults [88].

Finally, the temporal structure of exploration and exploitation may vary. Classically, exploration precedes exploitation, but there may be explore-exploit cycles. Exploration may also increase for particular domains in particular later periods: for example, adolescents may show an increase in exploration in the social domain [54], and a return to exploration may occur in adults in particular contexts [89]. Although these strategies differ in detail, they all help ensure that exploration takes place and so help resolve the broader explore-exploit tradeoff. An open question is whether and how early experience influences the development of these different forms of exploration.

idea is the correlation across species between the length of childhood and adult intelligence, flexibility, and learning, as well as brain size and neuron count ([18,39–41], but see [42]).

### Hyperparameters change across development: from exploration to exploitation

Recent theory and data suggest that there are indeed changes in hyperparameters of learning across ontogeny. There is increasing evidence for general shifts from exploration to exploitation across many types and levels of learning and development, ranging from phenotypic plasticity to abstract theory formation, as well as related changes in information processing, affect, and motivation [4,5,36–38,43].

Although the exploration-exploitation tradeoff is classically applied to learning problems, phenotypic plasticity – the ability of an organism to change its phenotype (including its body, mind, and brain) in response to environmental conditions – can also be understood in this way (Box 2). Organisms face a similar tradeoff when they search through the space of potential phenotypes and through the space of potential policies or hypotheses. More plastic species, or individuals in more plastic sensitive periods, explore a wider space of phenotypes than less plastic species, or individuals at less plastic stages. Plasticity comes with benefits – allowing adaptation to a wider range of environments – but also with costs: the more plastic immature organism is less well adapted to any environment.

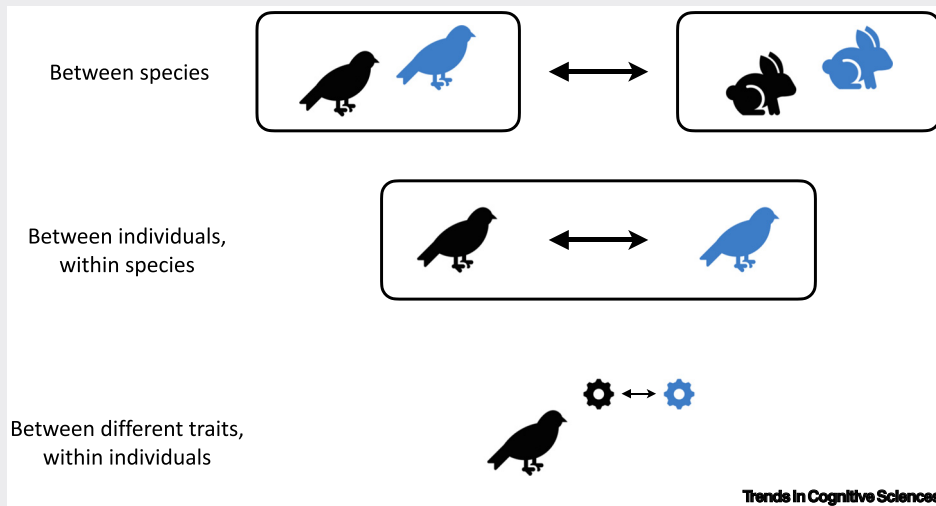
Across the tree of life, sensitive periods tend to occur early in ontogeny (Box 2). In human language learning, infants begin with the ability to learn all the phonetic contrasts in the world's

**Box 2. The evolution and development of sensitive periods**

Sensitive periods – in which the impact of experience on traits is larger than in other periods – are widespread in nature, but their evolution is not well understood [89]. Formal modelers have recently explored the conditions in which sensitive periods are favored by natural selection [73,90]. These models capture development as a process in which organisms learn about their environments and tailor their phenotypes to the environmental state. In some models, the environmental state changes very slowly relative to an organism’s lifespan, remaining nearly stable within generations [83,91–94]. In others, the environment fluctuates at a noticeable rate within generations as well [30,95]. Collectively, these models provide insight into key features of sensitive periods (e.g., timing, duration) at different levels of organization (Figure 1): between species (e.g., some birds learn new songs only early in life, others throughout their lifetimes), between individuals (e.g., children maturing at different rates following adversity), and between traits (e.g., cognitive and emotional systems recalibrating at different rates to changed environmental conditions).

Some findings are robust across models. If the environment is stable within generations, plasticity is often highest early in ontogeny and gradually declines while organisms reduce their uncertainty. This process might result in ‘critical periods’, after which plasticity drops to zero. If the environment fluctuates within generations, sensitive periods may occur at the onset, mid-way through, and even toward the end of ontogeny, but critical periods are not favored [30] unless changing phenotypes is inherently costly [95]. If the reliability of information increases across ontogeny, sensitive periods may occur at later developmental stages. Perhaps adolescents are particularly sensitive to peer feedback [96] because this feedback is more informative about social status or desirability in adulthood than feedback received in early childhood ([94], contextualized by Gee [97]).

In stable as well as fluctuating environments, the trajectory of plasticity depends on the degree to which uncertainty about environmental conditions persists across ontogeny. Individuals who have more consistent experiences lose plasticity earlier in ontogeny than individuals who have less consistent experiences [73]. This result fits evidence showing that noisy information or lack of information may prolong plasticity. For instance, bilingual infants show a more gradual decline in plasticity than monolingual infants [98]. In zebra finches, the absence of tutors extends the window for song learning [99]. In Japanese macaques, lack of face input prevents perceptual narrowing [100]. In rat pups, white noise delays auditory specialization [101]. Although the effects of noisy information or lack of information are complex, the rate at which the brain infers the statistical structure of the environment is likely one process that sets the parameters of learning mechanisms [102].



**Figure 1. Variation in plasticity across levels of organization.** The developmental trajectory of plasticity, such as the timing and duration of sensitive periods, varies: between species (top row), between individuals of the same species (middle row), and between different traits within a single individual (bottom row). A complete understanding of this variation requires proximate explanations, which focus on causation (neural, cognitive, and physiological processes) and ontogeny (developmental processes), and ultimate explanations, which focus on phylogeny (e.g., ancestral precursors) and function (e.g., adaptive value).

languages. As they gain more experience, they narrow the range of possibilities, until adults are unable to even hear distinctions in nonfamiliar languages [44,45]. In face perception, a broad non-specific system becomes gradually tuned for processing species-specific faces [46,47]. There is

also progressive specialization in motor development: behaviors are initially broadly tuned to their goal and become increasingly efficient during the first year of life [48]. In all these cases, an initial wider exploratory search converges on a narrower implementation of a solution.

The most extensive research on the explore–exploit tradeoff has taken place in reinforcement learning. In **multiarmed bandit tasks**, an agent must choose between implementing policies that have been rewarding in the past or experimenting with unknown or less rewarding options. Younger agents show more exploration in these tasks [13,49–51]. Moreover, recent work suggests a process akin to simulated annealing takes place across development, with children employing a more high-temperature search than adults [6,37].

The same pattern emerges in Bayesian hypothesis learning and theory formation, where agents ‘sample’ potential hypotheses and then compute their probabilities [52]. In causal learning tasks, children are better than adults at learning unusual hypotheses [53,54], suggesting a broader or ‘hotter’ search through the space of possible hypotheses (Box 1). Similarly, children are more likely than adults to search and find new hypotheses when the environment changes [7,32]. Younger children are also more willing to bear the costs of exploration than older children and adults, and this makes them better learners [55]. In a ‘creative foraging’ task, children find more variable and unique solutions than adults [56]. In an effort-based information sampling task, adolescents seek a higher evidence threshold before deciding compared with adults [57]. But not all studies have found a transition from more exploratory to less (e.g., [58,59]); future work may clarify boundary conditions.

There are related shifts in information processing. Children have a wider focus of attention than adults, a ‘lantern’ rather than a ‘spotlight’ [60], taking in information even when it is not immediately relevant to their goals [61]. Similarly, younger children are less efficient at remembering material that is relevant to their goals, but better at remembering unattended information [49]. There are associated shifts in motivation and affect. Curiosity and playfulness motivate exploration and are characteristic of younger animals, including children. Both human children and rat juveniles are less subject to avoidant fear conditioning than adults, and may actually approach cues to unpleasant stimuli, though only if a caregiver is present, prioritizing information over reward [62]. This pattern would also contribute to exploration.

This example also raises the question of the relationship between exploration and risk-taking. People living in adverse conditions tend to be oriented more toward immediate rewards. This may be necessary to meet their basic needs, and reasonable because future rewards are unlikely to materialize [63]. In some situations, this present-orientation may lead to a kind of risk-taking that involves less exploration (e.g., taking out a high-interest loan that will be difficult to repay, without considering alternatives). However, risk-taking can also be motivated by the desire for more exploration. In the fear conditioning case, for example, approaching the cue to the unpleasant stimuli is a form of exploration. It provides information – does the cue reliably predict the effect? – even at the cost of loss. This is also true in the balloon emotional learning task (see later), where the agent gets a reward each time they inflate a balloon, but loses all the reward if the balloon pops. The agent who continues to inflate learns the relation between action and outcome even as they experience the loss. Thus, in our view, people who experience adversity might show risk-prone behavior if this can lead to immediate reward, and risk-averse behavior if information search trades off with exploiting resources. So, our view differs from existing claims that people living in adverse conditions take more or less risk in general.

### Does adversity shorten childhood exploration?

It is plausible that human children have evolved capacities to produce adaptive responses to at least certain forms of adversity. Across human evolution, children have been exposed to higher levels of adversity (e.g., threat and deprivation) than is typical in industrialized societies [64]. Moreover, as these levels were highly variable over time and space, natural selection has likely favored the ability to develop adaptive responses to these adversities, notwithstanding their detrimental effects [64]. Shortening the period of childhood might be one such response.

The empirical record suggests an association between early adversity and accelerated life history [17–20], although there may be different patterns in different human populations [65]. For instance, both puberty and adult teeth emerge earlier in children with more adverse experiences [66]. Some types of neural development – changes in cortical thickness, functional connectivity, and amygdala connectivity – also appear to be accelerated in adverse conditions, and slowed down in enriched ones [24]. The **stress acceleration hypothesis** makes this same argument for emotion development, including fear conditioning and emotion regulation [21–24].

All of these accelerations may be related to explore–exploit shifts. However, much empirical work remains to be done to test the hypothesis that early adversity, in particular, accelerates the explore–exploit shifts in learning and development. An ideal test would concurrently document the development of life-history traits and explore–exploit behaviors in a longitudinal study.

There are some suggestive findings already. In adult rats, contemporaneous chronic stress and deprivation shift behavior from exploration to exploitation [67,68]. Similarly, human adults who experienced more stress in the past month exploited rather than explored in a virtual foraging task [69]. In another study, youth who had experienced early institutionalization explored less in the balloon emotional learning task [70]. Adults who experienced early economic adversity also showed less exploration on this task, and were averse to risk-taking in uncertain contexts [63]. Adolescents who experienced early maltreatment were less cognitively flexible; they were less likely to change their responses in the light of new evidence, analogous to a lower learning rate, and this response correlated with the amount of adversity [71]. All of these findings suggest less exploration.

Accelerated neural, emotional, and cognitive developmental changes may confer immediate survival benefits, although they come with a cost through less flexibility in the long run. This reflects the explore–exploit tradeoff. Settling quickly on good-enough solutions and focusing on immediate rewards might increase survival in an environment with scarce resources and short time horizons. Exploring possible solutions more widely and gathering more information is advantageous when the environment is richer and more complex and horizons are longer.

It may be adaptive to prioritize exploitation over exploration in response to early adversity for two reasons (Box 3). First, early adversity may be a cue that environmental conditions will be unfavorable in the future – limited or unpredictable resources, high levels of threat – favoring an earlier shift to exploitation (an external predictive adaptive response). Second, early adversity may directly cause somatic damage that shortens the anticipated lifespan favoring accelerated development, irrespective of future environmental conditions (an internal predictive adaptive response [72–74]). These theories are mutually compatible. Both would explain why early experience shapes the timing of the explore–exploit shift. To the extent that predictive adaptive responses continue to operate across the lifespan, both processes could also explain, at least in principle, why adverse conditions influence explore–exploit behavior in adulthood as well, although early experience appears to have a particularly large effect.

### Box 3. How do people assess environmental conditions to guide life histories?


There are numerous intriguing questions at the intersection of cognitive science and life history. Here, we ask: how do people assess environmental conditions to guide their life histories? We discuss two hypotheses – ancestral cues and statistical learning – based on the work of Young *et al.* [103].

According to the **'ancestral cue' hypothesis**, the human mind has evolved to respond to specific events (e.g., fights) that provide information about current and future environmental conditions (e.g., mortality rates in the broader ecology), which influences life-history development. There are examples in other animals. If female parasitic wasps detect changes in barometric pressure associated with thunderstorms, which might cause mortality, they increase reproductive effort by laying more eggs in low-quality hosts [104]. Similarly, if over human evolution exposure to violence predicted higher mortality, or moving into new territory predicted greater unpredictability, our 'stone-age minds' might use these cues to adjust life history (e.g., invest more in fertility and less in survival, but see Box 4).

According to the **'statistical learning' hypothesis**, humans estimate the statistical structure of the environment by integrating experiences across ontogeny [103,105], without necessarily relying on ancestral cues, and these estimates influence life-history development. Empirically, humans are able to learn evolutionarily novel cues (e.g., gang signs indicate danger), update the weights of cues [106], and use prediction errors (i.e., differences between anticipated and actual conditions) to estimate the unpredictability and controllability of the environment [107]. People might integrate across these sources of information (time series) to build models of their environments, which influence life histories. The 'ancestral cues' and 'statistical learning' hypotheses are mutually compatible: people may combine ancestral and learned cues in developing models of their world [103] (Figure 1). Empirical work suggests that people may vary in the extent to which they rely on each type of cue [108].

Ancestral cues and statistical learning have different pros and cons. Ancestral cues are efficient: a developmental response may be triggered by only limited exposure to the cue. There is no need to track, store, and use experiences to build models of the environment. However, this simplicity implies lower flexibility: people are unable to learn new cues, update cue reliabilities (e.g., if the environment changes), and extinguish the weight of ancestral cues [103].

Future theory can explore the conditions in which organisms should rely on ancestral cues or statistical learning, or both, using the framework of statistical decision theory [109]. Future empirical research can quantify the statistical structure of people's environments [103,105], and examine whether it predicts life-history outcomes as well, better, or worse, than their exposures to ancestral cues. Future work should also explore whether and how early experience has an influence on later development through ancestral or statistical cues or both.



**Figure 1. Assessing current and future environmental conditions.** The 'ancestral cue' hypothesis states that human development responds to specific events that provided information about broader ecological conditions across evolutionary time. The statistical learning hypothesis states that the mind uses a variety of learning strategies to develop estimates, including (but not limited to) learning novel cues, updating cue weights, and prediction errors. These hypotheses are not mutually exclusive: people might combine ancestral and learned cues in developing models of their world. Adapted from Young *et al.* [103].



### Modeling the evolution of development in life-history theory

Our hypothesis, then, is that adversity accelerates the explore–exploit shift, a plastic response produced by an evolved adaptation. The question of ‘how’ such an adaptation could evolve fits within a branch of evolutionary biology called life-history theory. Life-history theory is not one theory, but rather a framework for studying how organisms allocate their limited resources (e.g., time and energy) between different activities (e.g., growth and reproduction) across their lifespan [75,76]. It is a broad framework for studying questions about tradeoffs, not a particular set of predictions [20,77]. However, the framework allows us to build specific models, which do make predictions, very successfully in some cases.

In the past 15 years, the number of studies bridging psychology and life history has increased rapidly [78,79]. This trend has led to new developments. First, early work explored life-history traits that directly affect survival and reproduction, such as the age of first birth, the number and quality of offspring, reproductive lifespan, and aging. Recent research incorporates the idea that life-history strategies also include psychological traits, such as risk-taking and impulsivity [17–20,25,26,63]. We would add hyperparameters for exploration to this list.

Second, the early work focused mainly on variation between species or populations. Recent work often examines individual differences, and in particular, the hypothesis that there is a fast–slow

#### Box 4. Does extrinsic mortality favor ‘fast’ life histories?


A widespread claim is that higher extrinsic (i.e., unavoidable) mortality favors faster life histories; that is, organisms should invest more in fertility and less in survival [110]. The intuition is simple: if death looms large, organisms should reproduce early and often, before it is too late. However, the reality is more complex. Formal modeling shows that higher mortality might actually accelerate, decelerate, or have no impact on the pace of development [111]. Our hypothesis – that adversity accelerates the explore–exploit shift – fits with a fast–slow continuum of individual differences based on plastic responses [17–20] (for recent meta-analyses, see [112,113]). However, we do not subscribe to the widespread claim that life-history theory ‘predicts’ such a continuum. To clarify why, we need to share a basic insight from formal evolutionary theory. Our exposition follows the logic of a recent guide for the perplexed [80] (see also <https://www.danielnettle.org.uk/2022/02/18/live-fast-and-die-young-maybe/>).

Meet two strategies: fast and slow. Fast produces offspring only at time 1, then it dies. Slow produces offspring at times 1 and 2, if it survives to time 2. All else being equal, lower mortality means more slow individuals survive to time 2, increasing this strategy’s lifetime fitness relative to fast. This scenario is consistent with the widespread claim that lower mortality favors slower strategies (later reproduction). However, all else is not equal; when mortality goes down, the population begins to grow exponentially (Figure 1).

This places a premium on early reproduction: early-produced offspring contribute a greater share to future generations than late-produced offspring; after all, early-produced offspring contribute descendants as well [80]. This is analogous to a growing economy: the faster the growth, the better for those who invest early. In a fast-growing population, the population growth rate (which factors in the timing of reproductive events) is thus a better measure than lifetime reproductive success [114]. In contrast to the widespread claim, lower mortality does not change selection pressures for fast versus slow strategies. The premium on early reproduction cancels out exactly the benefit of the second reproductive attempt for the slow strategy in an exponentially growing population.

But no population can grow rapidly forever. It matters what makes population growth slow down: higher mortality, lower fertility, or migration. We focus on higher mortality occurring in a dense population (e.g., due to intense competition for resources, faster-spreading disease). The question then is: who suffers from high density, juveniles or adults? If juveniles, the outcome matches the widespread claim: more juveniles will die at high densities, while adults survive, and slow adults can reap the reproductive benefits of a robust body. But if adults suffer from high densities, it is pointless to invest in a robust body, and better to shift investment to reproduction. In this scenario, in the opposite of the classic prediction, decreased extrinsic mortality – which leads to higher densities that harm adults more than juveniles – selects for fast strategies. Finally, if juveniles and adults are affected equally, higher mortality has no impact on the pace of development [111].


Building a solid bridge between life history and learning requires embracing theoretical nuances and tailoring new models to the peculiar features of human life history, including social learning [27,28].



continuum, where faster individuals mature at a younger age, have more offspring, invest less in each offspring, and senesce and die earlier [17–20]. This broader focus dovetails with the new hypothesis that adversity promotes a ‘fast’ strategy, accelerating the pace of brain development [21–24]. Again, our hypothesis about the accelerated explore–exploit shift fits with this general approach. We emphasize, however, that both our and the other hypotheses about fast strategies in response to adversity are inferred from empirical observations, not deduced from formal theory [77,80,81]. In fact, formal theory ‘predicts’ accelerated maturation following adversity under a limited set of conditions, and not for the reasons people often think (Box 4).

There is a well-established formal theory exploring the conditions in which phenotypic plasticity is favored by natural selection over nonplastic development [73,82]. However, the majority of this work has not studied strategies for exploration. Rather, it has assumed a two-stage life history. In the first stage, organisms sample a cue, or they do not sample, relying instead on (epi)genetically inherited information. In the second stage, organisms develop a phenotype, instantaneously or after a time lag [83]. These models do not allow organisms to explore more or less, broadly or narrowly, or maintain flexibility or increase specialization for exploitation, while constructing their phenotypes incrementally in a series of small steps [83–85].

Recent work has modeled phenotypic plasticity as an explore–exploit tradeoff [73]. These models capture development as a sequential decision-making process, in which organisms learn about their environments while gradually tailoring phenotypes (Figure 2). This approach affords a



Trends in Cognitive Sciences

**Figure 2. Formal models of the evolution of development.** Organisms start out with a prior estimate of their environment. This estimate reflects the distribution of environmental states experienced by a species. Then organisms may sample cues that provide information (reduce uncertainty) about the current environmental state. The cue reliability determines the extent to which cues reduce uncertainty. Both the environment and the cue reliability may be stable or variable across ontogeny. If the environment is autocorrelated, cues predict future states of the environment. After sampling a cue, the organism updates its estimate (posterior) and may develop a phenotype. In models of incremental development, organisms gradually tailor their phenotypes to the environment, rather than instantaneously. Using optimization methods, we find the optimal policy for different evolutionary ecologies: combinations of priors, cue reliabilities, rates of environmental change, and other features of the environment, such as food availability, predator density, and pathogen risk. For each potential state of the organism, the optimal policy specifies the optimal decision: whether to sample and which phenotypes to develop. In response to experiences – cues and other events influencing the phenotype (e.g., lowered somatic quality following an accident or predation) – the policy produces developmental trajectories and mature phenotypes. Adapted, with permission, from Walasek (<https://osf.io/9tqs6/>).

wider range of developmental trajectories and outcomes, with organisms sampling more or less at different times in development. These new models have already shed light on species-typical sensitive periods, and the effects of postnatal experiences on their timing and duration (Box 2). But these models typically hold life history constant, exploring only how ecology (e.g., rates of change) and experiences (e.g., cue reliability) shape levels of plasticity across ontogeny. They have not allowed life history, plasticity, and learning to vary and influence each other in a coevolutionary dance.

We know of only two models exploring how variation in longevity, a life-history trait, affects plasticity across ontogeny. These models reach different conclusions. In one model, if adulthood is long, organisms rely not on childhood cues, but instead on inherited information about long-term environmental conditions [30]. In the other model [29], adults continue to be able to use cues to improve the phenotype–environment match, and thus longevity favors plasticity. In birds and mammals, high plasticity is associated with longer lifespans [4,42], but in certain species of fish, high plasticity is associated with shorter lifespans [86]. Future work could examine which factors moderate this pattern of associations.

Formal models have also explored how learning coevolves with life-history traits. One model has explored the extent to which social (vs. individual) learning is favored as a function of births and deaths in different age classes, and the degree of competition, in a population [27]. Its results are surprising: organisms with ‘fast’ life histories often evolve greater reliance on social learning than those with ‘slow’ life histories. It thus remains an open question why many primates – and humans, in particular – have slow life histories yet are brilliant social learners. Another model has explored the conditions that favor teaching as well as the optimal ages of teaching and learning [28]. In many conditions, this model favors youths tutoring each other, and elderly instructing youth, while peak-productive adults produce rather than teach. These predictions are supported by cross-cultural analyses of hunting [28].

### Concluding remarks

We envision exciting directions for future work (see [Outstanding questions](#)). There are different types of adversity, from neglect to threat to unpredictability, and different types of exploration and exploitation in different domains. A key task will be to examine how all of these are related. The developmental trajectory may also be more complex than simply exploration followed by exploitation; for instance, there may be explore–exploit cycles, and adults may recover their capacity for exploration in particular contexts. Moreover, the effects of adversity are clearly not deterministic; many people show resilience. Finally, there should be formal theory exploring in which environmental conditions it is adaptive for adverse experiences to accelerate the explore–exploit shift (Box 4).

The acceleration hypothesis also has practical implications. There are cognitive, emotional, and neural consequences of early adversity or enrichment. A common view is that early enrichments, like preschool, instill particular kinds of knowledge or skill that are useful later, and adversity makes these skills less available. Our hypothesis suggests a rather different yet complementary picture. Cues to early nurture may allow a longer period of more flexible learning rather than instilling particular skills or content. This may help explain the broad and enduring effects of early adversity, ranging from lower income and education to depression and anxiety. The nurture, care, and investment children receive may allow a world of wider possibilities.

### Acknowledgments

We thank Irene Godoy, Hanna Kokko, Daniel Nettle, Seth Pollak, Charlotte de Vries, Nicole Walasek, Esther Weijman, and Ethan Young for helpful feedback on previous drafts of this opinion article. The research of W.E.F. is supported by the Dutch

### Outstanding questions

Does the explore–exploit shift develop in concert with, or independently from, life-history traits (e.g., pubertal timing), physiological traits (e.g., metabolic rate), and other psychological traits (e.g., risk-taking, impulsivity)?

Do such correlations (e.g., between faster maturation and earlier explore–exploit shift) exist not only among individuals, but also across species?

What changes in affect and motivation accompany the shift from exploration to exploitation?

Is the explore–exploit shift specific to certain domains of development or general across them (e.g., cognitive, socioemotional, motor)? Do different domains have different developmental trajectories?

Does early experience influence the pace of development through ancestral or statistical cues or both?

Do explore–exploit shifts change reliance on individual learning versus social learning? Do such effects vary between different sources of information (e.g., parents or peers) and for different types of social learning (e.g., copying the majority or imitating prestigious individuals)?

What can formal modeling teach us about the effects of early adversity on exploration?

How do different types of adversity – such as threat, deprivation, or unpredictability – influence the development of exploration?

Are there different effects on different types of exploration (e.g., random vs. directed)?

Do evolved responses to anticipated future environmental conditions versus anticipated future somatic decline explain the developmental timing of the explore–exploit shift?

How are progressive specialization, developmental narrowing, and explore–exploit shifts related?

Does enrichment decelerate the shift from exploration to exploitation? If so,

Research Council (V1.Vidi.195.130) and the James S. McDonnell Foundation (<http://doi.org/10.37717/220020502>). The work of A.G. is supported by grants from DARPA (047498-002), DOD ONR MURI, the Templeton Foundation (61475), and the Templeton World Charity Foundation.

**Declaration of interests**

No interests are declared.

**References**

1. Gee, D.G. (2021) Early adversity and development: parsing heterogeneity and identifying pathways of risk and resilience. *Am. J. Psychiatry* 178, 998–1013
2. Nelson, C.A. et al. (2020) Adversity in childhood is linked to mental and physical health throughout life. *Br. Med. J.* 371, m3048
3. Repetti, R.L. et al. (2002) Risky families: family social environments and the mental and physical health of offspring. *Psychol. Bull.* 128, 330–366
4. Gopnik, A. (2020) Childhood as a solution to explore–exploit tensions. *Philos. Trans. R. Soc. B* 375, 20190502
5. Gualtieri, S. and Finn, A.S. (2022) The sweet spot: when children’s developing abilities, brains, and knowledge make them better learners than adults. *Perspect. Psychol. Sci.* 17, 1322–1338
6. Meder, B. et al. (2021) Development of directed and random exploration in children. *Dev. Sci.* 24, e13095
7. Sumner, E. et al. (2019) It’s not the treasure, it’s the hunt: children are more explorative on an explore/exploit task than adults. In *Proceedings of the 41st Annual Conference of the Cognitive Science Society* (Goel, A.K. et al., eds), pp. 2891–2897, Cognitive Science Society
8. Cohen, J.D. et al. (2007) Should I stay or should I go? How the human brain manages the tradeoff between exploitation and exploration. *Philos. Trans. R. Soc. B* 362, 933–942
9. Del Giudice, M. and Crespi, B.J. (2018) Basic functional tradeoffs in cognition: an integrative framework. *Cognition* 179, 56–70
10. Hills, T.T. et al. (2015) Exploration versus exploitation in space, mind, and society. *Trends Cogn. Sci.* 19, 46–54
11. Tenenbaum, J.B. et al. (2006) Theory-based Bayesian models of inductive learning and reasoning. *Trends Cogn. Sci.* 10, 309–318
12. Smith, L.B. et al. (2018) The developing infant creates a curriculum for statistical learning. *Trends Cogn. Sci.* 22, 325–336
13. Eckstein, M.K. et al. (2022) The interpretation of computational model parameters depends on the context. *eLife* 11, e75474
14. Kirkpatrick, S. et al. (1983) Optimization by simulated annealing. *Science* 220, 671–680
15. Coall, D.A. and Chisholm, J.S. (2003) Evolutionary perspectives on pregnancy: maternal age at menarche and infant birth weight. *Soc. Sci. Med.* 57, 1771–1781
16. Ellis, B.J. et al. (2022) Why and how does early adversity influence development? Toward an integrated model of dimensions of environmental experience. *Dev. Psychopathol.* 34, 447–471
17. Belsky, J. et al. (1991) Childhood experience, interpersonal development, and reproductive strategy. *Child Dev.* 62, 647–670
18. Ellis, B.J. et al. (2009) Fundamental dimensions of environmental risk: the impact of harsh versus unpredictable environments on the evolution and development of life history strategies. *Hum. Nat.* 20, 204–268
19. Del Giudice, M. (2020) Rethinking the fast-slow continuum of individual differences. *Evol. Hum. Behav.* 41, 536–549
20. Frankenhuis, W.E. and Nettle, D. (2020) Current debates in human life history research. *Evol. Hum. Behav.* 41, 469–473
21. Callaghan, B.L. and Tottenham, N. (2016) The stress acceleration hypothesis: effects of early-life adversity on emotion circuits and behavior. *Curr. Opin. Behav. Sci.* 7, 76–81
22. Gee, D.G. et al. (2013) Early developmental emergence of human amygdala–prefrontal connectivity after maternal deprivation. *Proc. Natl. Acad. Sci. U. S. A.* 110, 15638–15643
23. Gee, D.G. and Cohodes, E.M. (2021) Influences of caregiving on development: a sensitive period for biological embedding of predictability and safety cues. *Curr. Dir. Psychol. Sci.* 30, 376–383
24. Tooley, U.A. et al. (2021) Environmental influences on the pace of brain development. *Nat. Rev. Neurosci.* 22, 372–384
25. Réale, D. et al. (2010) Personality and the emergence of the pace-of-life syndrome concept at the population level. *Philos. Trans. R. Soc. B* 365, 4051–4063
26. Sih, A. and Del Giudice, M. (2012) Linking behavioural syndromes and cognition: a behavioural ecology perspective. *Philos. Trans. R. Soc. B* 367, 2762–2772
27. Deffner, D. and McElreath, R. (2020) The importance of life history and population regulation for the evolution of social learning. *Philos. Trans. R. Soc. B* 375, 20190492
28. Gurven, M.D. et al. (2020) The optimal timing of teaching and learning across the life course. *Philos. Trans. R. Soc. B* 375, 20190500
29. Ratikainen, I.I. and Kokko, H. (2019) The coevolution of lifespan and reversible plasticity. *Nat. Commun.* 10, 538
30. Walasek, N. et al. (2022) Sensitive periods, but not critical periods, evolve in a fluctuating environment: a model of incremental development. *Proc. R. Soc. B* 289, 20212623
31. Sutton, R.S. and Barto, A.G. (1998) *Reinforcement Learning: An Introduction*, MIT Press
32. Sumner, E. et al. (2019) The exploration advantage: children’s instinct to explore allows them to find information that adults miss. *PsyArXiv* Published online June 28, 2021. <http://doi.org/10.31234/osf.io/h437v>
33. Schulz, E. and Gershman, S.J. (2019) The algorithmic architecture of exploration in the human brain. *Curr. Opin. Neurobiol.* 55, 7–14
34. Whittle, P. (1982) *Optimization Over Time*, John Wiley & Sons
35. Smith, L.N. (2017) Cyclical learning rates for training neural networks. In *Proceedings of 2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pp. 464–472
36. Brown, R. et al. (2022) Prediction in the aging brain: merging cognitive, neurological, and evolutionary perspectives. *J. Gerontol. Psychol. Sci.* 77, 1580–1591
37. Giron, A.P. et al. (2023) Developmental changes in learning resemble stochastic optimization. *PsyArXiv* Published online February 17, 2023. <http://doi.org/10.31234/osf.io/9f4k3>
38. Spreng, R.N. and Turner, G.R. (2021) From exploration to exploitation: a shifting mental mode in late life development. *Trends Cogn. Sci.* 25, 1058–1071
39. Kaplan, H. et al. (2000) A theory of human life history evolution: diet, intelligence, and longevity. *Evol. Anthropol.* 9, 156–185
40. Reader, S.M. and Laland, K.N. (2002) Social intelligence, innovation, and enhanced brain size in primates. *Proc. Natl. Acad. Sci.* 99, 4436–4441
41. Uomini, N. et al. (2020) Extended parenting and the evolution of cognition. *Philos. Trans. R. Soc. B* 375, 20190495
42. Sol, D. et al. (2016) The life-history basis of behavioural innovations. *Philos. Trans. R. Soc. B* 371, 20150187
43. Chierchia, G. et al. (2022) Confirmatory reinforcement learning changes with age during adolescence. *Dev. Sci.* 26, e13330
44. Kuhl, P.K. (2004) Early language acquisition: cracking the speech code. *Nat. Rev. Neurosci.* 5, 831–843
45. Werker, J.F. et al. (2012) How do infants become experts at native-speech perception? *Curr. Dir. Psychol. Sci.* 21, 221–226
46. Nelson, C.A. (2001) The development and neural bases of face recognition. *Infant Child Dev.* 10, 3–18
47. Pascalis, O. et al. (2020) Development of face processing: are there critical or sensitive periods? *Curr. Opin. Behav. Sci.* 36, 7–12
48. D’Souza, H. et al. (2017) Specialization of the motor system in infancy: from broad tuning to selectively specialized purposeful actions. *Dev. Sci.* 20, e12409

does this effect depend on the type of enrichment (e.g., nutritional, psychosocial)?

Which factors other than adversity and enrichment influence the shift from exploration to exploitation?

What cues will reopen flexibility and exploration in adolescence and adulthood?

49. Blanco, N.J. and Sloutsky, V.M. (2021) Systematic exploration and uncertainty dominate young children's choices. *Dev. Sci.* 24, e13026
50. Nussenbaum, K. and Hartley, C.A. (2019) Reinforcement learning across development: what insights can we draw from a decade of research? *Dev. Cogn. Neurosci.* 40, 100733
51. Schulz, E. et al. (2019) Searching for rewards like a child means less generalization and more directed exploration. *Psychol. Sci.* 30, 1561–1572
52. Bonawitz, E. et al. (2014) Probabilistic models, learning algorithms, response variability: sampling in cognitive development. *Trends Cogn. Sci.* 18, 497–500
53. Lucas, C.G. et al. (2014) When children are better (or at least more open-minded) learners than adults: developmental differences in learning the forms of causal relationships. *Cognition* 131, 284–299
54. Gopnik, A. et al. (2017) Changes in cognitive flexibility and hypothesis search across human life history from childhood to adolescence to adulthood. *Proc. Natl. Acad. Sci. U. S. A.* 114, 7892–7899
55. Liquin, E.G. and Gopnik, A. (2022) Children are more exploratory and learn more than adults in an approach-avoid task. *Cognition* 218, 104940
56. Hart, Y. et al. (2022) The development of creative search strategies. *Cognition* 225, 105102
57. Niebaum, J.C. et al. (2022) Adolescents sample more information prior to decisions than adults when effort costs increase. *Dev. Psychol.* 58, 1974–1985
58. Neldner, K. et al. (2019) Creation across culture: children's tool innovation is influenced by cultural and developmental factors. *Dev. Psychol.* 55, 877–889
59. Peiz, M. and Kidd, C. (2020) The elaboration of exploratory play. *Philos. Trans. R. Soc. B* 375, 20190503
60. Gopnik, A. (2009) *The Philosophical Baby: What Children's Minds Tell Us about Truth, Love and the Meaning of Life*, Random House
61. Frank, S.M. et al. (2021) Fundamental differences in visual perceptual learning between children and adults. *Curr. Biol.* 33, 427–432
62. Tottenham, N. et al. (2019) Parental presence switches avoidance to attraction learning in children. *Nat. Hum. Behav.* 3, 1070–1077
63. Amir, D. et al. (2018) An uncertainty management perspective on long-run impacts of adversity: the influence of childhood socioeconomic status on risk, time, and social preferences. *J. Exp. Soc. Psychol.* 79, 217–226
64. Frankenhuus, W.E. and Amir, D. (2022) What is the expected human childhood? Insights from evolutionary anthropology. *Dev. Psychopathol.* 34, 473–497
65. Sear, E. (2020) Do human 'life history strategies' exist? *Evol. Hum. Behav.* 41, 513–526
66. McDermott, C.L. (2021) Early life stress is associated with earlier emergence of permanent molars. *Proc. Natl. Acad. Sci. U. S. A.* 118, e2105304118
67. Garcia-Marquez, C. and Armario, A. (1987) Chronic stress depresses exploratory activity and behavioral performance in the forced swimming test without altering ACTH response to a novel acute stressor. *Physiol. Behav.* 40, 33–38
68. Matisz, C.E. et al. (2021) Chronic unpredictable stress shifts rat behavior from exploration to exploitation. *Stress* 24, 635–644
69. Lenow, J.K. et al. (2017) Chronic and acute stress promote overexploitation in serial decision making. *J. Neurosci.* 37, 5681–5689
70. Humphreys, K.L. et al. (2015) Exploration–exploitation strategy is dependent on early experience. *Dev. Psychobiol.* 57, 313–321
71. Harms, M.B. et al. (2018) Instrumental learning and cognitive flexibility processes are impaired in children exposed to early life stress. *Dev. Sci.* 21, e12596
72. Nettle, D. et al. (2013) The evolution of predictive adaptive responses in human life history. *Proc. R. Soc. B* 280, 20131343
73. Frankenhuus, W.E. and Walasek, N. (2020) Modeling the evolution of sensitive periods. *Dev. Cogn. Neurosci.* 41, 100715
74. Geronimus, A.T. et al. (2006) 'Weathering' and age patterns of allostatic load scores among blacks and whites in the United States. *Am. J. Public Health* 96, 826–833
75. Roff, D.A. (2002) *Life History Evolution*, Sinauer
76. Stearns, S.C. (1992) *The Evolution of Life Histories*, Oxford University Press
77. Stearns, S.C. and Rodrigues, A.M. (2020) On the use of 'life history theory' in evolutionary psychology. *Evol. Hum. Behav.* 41, 474–485
78. Nettle, D. and Frankenhuus, W.E. (2019) The evolution of life history theory: a bibliometric analysis of an interdisciplinary research area. *Proc. R. Soc. B* 286, 20190040
79. Nettle, D. and Frankenhuus, W.E. (2020) Life history theory in psychology and evolutionary biology: one research programme or two? *Philos. Trans. R. Soc. B* 375, 20190490
80. de Vries, C. et al. (2023) Extrinsic mortality and senescence: a guide for the perplexed. *Peer Community J.* 3, e29
81. Mathot, K.J. and Frankenhuus, W.E. (2018) Models of pace-of-life syndromes (POLS): a systematic review. *Behav. Ecol. Sociobiol.* 72, 41
82. Snell-Rood, E.C. and Steck, M.K. (2019) Behaviour shapes environmental variation and selection on learning and plasticity: review of mechanisms and implications. *Anim. Behav.* 147, 147–156
83. Frankenhuus, W.E. and Panchanathan, K. (2011) Balancing sampling and specialization: an adaptationist model of incremental development. *Proc. R. Soc. B* 278, 3558–3565
84. Lew-Levy, S. et al. (2022) Socioecology shapes child and adolescent time allocation in twelve hunter-gatherer and mixed-subsistence forager societies. *Sci. Rep.* 12, 8054
85. Pretelli, I. et al. (2022) Foraging complexity and the evolution of childhood. *Sci. Adv.* 8, eabn9889
86. Sowersby, W. et al. (2021) Fast life-histories are associated with larger brain size in killifishes. *Evolution* 75, 2286–2298
87. Schulz, L. (2012) The origins of inquiry: inductive inference and exploration in early childhood. *Trends Cogn. Sci.* 16, 382–389
88. Somerville, L.H. et al. (2017) Charting the expansion of strategic exploratory behavior during adolescence. *J. Exp. Psychol.* 146, 155–164
89. Fawcett, T.W. and Frankenhuus, W.E. (2015) Adaptive explanations for sensitive windows in development. *Front. Zool.* 12, S3
90. Stamps, J.A. and Luttbeg, B. (2022) Sensitive period diversity: insights from evolutionary models. *Q. Rev. Biol.* 97, 243–295
91. English, S. et al. (2016) Adaptive use of information during growth can explain long-term effects of early life experiences. *Am. Nat.* 187, 620–632
92. Panchanathan, K. and Frankenhuus, W.E. (2016) The evolution of sensitive periods in a model of incremental development. *Proc. R. Soc. B* 283, 20152439
93. Stamps, J.A. and Krishnan, V.V. (2017) Age-dependent changes in behavioural plasticity: insights from Bayesian models of development. *Anim. Behav.* 126, 53–67
94. Walasek, N. et al. (2022) An evolutionary model of sensitive periods when the reliability of cues varies across ontogeny. *Behav. Ecol.* 33, 101–114
95. Fischer, B. et al. (2014) The evolution of age-dependent plasticity. *Am. Nat.* 183, 108–125
96. Fuhrmann, D. et al. (2015) Adolescence as a sensitive period of brain development. *Trends Cogn. Sci.* 19, 558–566
97. Gee, D.G. (2022) When do sensitive periods emerge later in development? *Trends Cogn. Sci.* 26, 97–98
98. D'Souza, D. and D'Souza, H. (2021) Bilingual adaptations in early development. *Trends Cogn. Sci.* 25, 727–729
99. Kelly, T.K. et al. (2018) Epigenetic regulation of transcriptional plasticity associated with developmental song learning. *Proc. R. Soc. B* 285, 20180160
100. Sugita, Y. (2008) Face perception in monkeys reared with no exposure to faces. *Proc. Natl. Acad. Sci. U. S. A.* 105, 394–398
101. Chang, E.F. and Merzenich, M.M. (2003) Environmental noise retards auditory cortical development. *Science* 300, 498e502
102. Moran, R.J. et al. (2014) The brain ages optimally to model its environment: evidence from sensory learning over the adult lifespan. *PLoS Comput. Biol.* 10, e1003422
103. Young, E.S. et al. (2020) Theory and measurement of environmental unpredictability. *Evol. Hum. Behav.* 41, 550–556
104. Roitberg, B. et al. (1993) Life expectancy and reproduction. *Nature* 364, 108

105. Frankenhuis, W.E. *et al.* (2019) A case for environmental statistics for early life effects. *Philos. Trans. R. Soc. B* 374, 20180110
106. Behrens, T.E.J. *et al.* (2007) Learning the value of information in an uncertain world. *Nat. Neurosci.* 10, 1214–1221
107. Moscarello, J.M. and Hartley, C.A. (2017) Agency and the calibration of motivated behavior. *Trends Cogn. Sci.* 21, 725–735
108. Li, Z. *et al.* (2023) Testing different sources of environmental unpredictability on adolescent functioning: Ancestral cue versus statistical learning and the role of temperament. *J. Child Psychol. Psychiatry* 64, 437–448
109. Dall, S.R. *et al.* (2005) Information and its use by animals in evolutionary ecology. *Trends Ecol. Evol.* 20, 187–193
110. Williams, G.C. (1957) Pleiotropy, natural selection, and the evolution of senescence. *Evolution* 11, 398–411
111. Moorad, J. *et al.* (2019) Evolutionary ecology of senescence and a reassessment of Williams' 'extrinsic mortality' hypothesis. *Trends Ecol. Evol.* 34, 519–530
112. Colich, N.L. *et al.* (2020) Accelerated biological aging following childhood experiences of threat and deprivation: a meta-analysis. *Psychol. Bull.* 146, 721–764
113. Sear, R. *et al.* (2019) Cross-cultural evidence does not support universal acceleration of puberty in father-absent households. *Philos. Trans. R. Soc. B* 374, 20180124
114. Caswell, H. (2001) *Matrix Population Models: Construction, Analysis and Interpretation* (2<sup>nd</sup> edn), Sinauer Associates