Balancing sampling and specialization: an adaptationist model of incremental development

Willem E. Frankenhuis and Karthik Panchanathan

Supplementary data	"Data Supplement" http://rspb.royalsocietypublishing.org/content/suppl/2011/04/05/rspb.2011.0055.DC1.html
References	**This article cites 32 articles, 8 of which can be accessed free** http://rspb.royalsocietypublishing.org/content/278/1724/3558.full.html#ref-list-1
Email alerting service	Receive free email alerts when new articles cite this article - sign up in the box at the top right-hand corner of the article or click **here**

To subscribe to *Proc. R. Soc. B* go to: **http://rspb.royalsocietypublishing.org/subscriptions**

Balancing sampling and specialization: an adaptationist model of incremental development

Willem E. Frankenhuis[*,†] **and Karthik Panchanathan**[†]

UCLA Department of Anthropology, 341 Haines Hall, Box 951553, Los Angeles, CA 90095-1553, USA

Development is typically a constructive process, in which phenotypes incrementally adapt to local ecologies. Here, we present a novel model in which natural selection shapes developmental systems based on the evolutionary ecology, and these systems adaptively guide phenotypic development. We assume that phenotypic construction is incremental and trades off with sampling cues to the environmental state. We computed the optimal developmental programmes across a range of evolutionary ecological conditions. Using these programmes, we simulated distributions of mature phenotypes. Our results show that organisms sample the environment most extensively when cues are moderately, not highly, informative. When the developmental programme relies heavily on sampling, individuals transition from sampling to specialization at different times in ontogeny, depending on the consistency of their sampled cue set; this finding suggests that stochastic sampling may result in individual differences in plasticity itself. In addition, we find that different selection pressures may favour similar developmental mechanisms, and that organisms may incorrectly calibrate development despite stable ontogenetic environments. We hope our model will stimulate adaptationist research on the constructive processes guiding development.

Keywords: evolutionary ecology; development; adaptation; differential plasticity; life history; stochastic dynamic programming

1. INTRODUCTION

Adaptive phenotypic plasticity, the ability of organisms to tailor development to environmental conditions, is widespread in nature [1–4], and can evolve in response to spatially or temporally varying environments, when cues indicate the state of the world [5–9]. Evolutionary models of adaptive phenotypic plasticity typically assume a two-stage life history: organisms first sample a cue to the state of the environment, and then develop phenotypes, either instantaneously (e.g. [7,10]) or after a time lag (e.g. [11–13]). These models ignore the fact that development is often a constructive process, in which phenotypes incrementally adapt to local ecologies ([14,15]; but see [2,3]).

In this paper, we present an evolutionary model of incremental phenotypic development. Viewing development as gradual reveals novel trade-offs. Here, we consider one such trade-off: between learning (sampling cues to the environmental state) and specializing (tailoring the phenotype to the local ecology). We make two assumptions. First, there is a positive correlation between fitness and the time invested in developing an environment-specific phenotype. This implies that an organism may benefit from specializing earlier in ontogeny, because it allows more time for achieving an adaptive 'fit' with the environment. Second, organisms are uncertain about the state of the environment, and can reduce this uncertainty through sampling environmental cues. The more time an organism spends sampling, the more it can reduce uncertainty. Combining these assumptions reveals a trade-off: earlier specialization may result in a more adaptive phenotype, however, it implies less time for sampling, increasing the risk of miscalibration.

These two assumptions are plausible. Correlations between the time invested in phenotypic development and the adaptive fit to the environment have been documented in a variety of species. Water fleas (*Daphnia*) that begin tailoring their phenotype prenatally towards a predator-rich environment develop more effective protective helmets than individuals who begin specializing only after birth [16]. Social insects progressively become more efficient at performing a variety of tasks, including identification of food sources, food handling techniques, nest repair, nest-mate recognition, comb building and nest climate control [17]. Many fish species can increase their foraging efficiency by incrementally tailoring morphology and behaviour for capturing locally abundant prey [18]. In hunter–gatherer societies, becoming a skilled hunter requires a substantial developmental commitment: it may take men more than 20 years before they produce more calories than they consume ([19]; for a related hypothesis in the domain of human reproductive strategies, see [20]).

Despite benefits, earlier specialization may entail costs. Specializing earlier implies less time spent learning. If cues probabilistically indicate the environmental state, less learning increases the risk of miscalibrating [21]: incorrectly inferring the environmental state and so developing an inappropriate phenotype. Miscalibration may be costly for several reasons [22]. First, *if* phenotypic development can be reversed in light of new experiences,

* Author for correspondence (wfrankenhuis@gmail.com).
[†] These authors contributed equally to the study.

there may be costs to shifting from one phenotype to another. These costs may result from the reworking of tissues that have in part committed to other functions. Second, development is a path-dependent process; what happens at earlier stages might narrow future options. For instance, developing a larger body now may preclude faster speed in the future, because developing the required musculature for speed would be too expensive. Third, developmental specialization may be irreversible, in which case organisms are 'stuck' their entire lifetime with a maladapted phenotype (see [4] for examples).

In this paper, we use state-dependent life-history theory, implemented by stochastic dynamic programming [23,24], to model how natural selection should shape developmental systems to optimally balance sampling environmental cues and phenotypic specialization.

2. MODEL
We assume that the environment varies spatially (not temporally), being in one of two states, *environment 1* or *environment 0* (e.g. predator-rich or predator-poor). There is one optimal phenotype for each environmental state, *phenotype 1* or *phenotype 0* (e.g. armoured or fast moving). Furthermore, fitness increases with the time an organism invests in constructing the environment-appropriate phenotype. We assume the following life history: organisms are born; fully disperse and settle into different patches; develop; reproduce proportional to the adaptive fit to their current patch; mature individuals die and the cycle repeats. (An equivalent life history: organisms are born and develop in their natal patch; fully disperse; reproduce proportional to the adaptive fit to their natal patch; mature individuals die and the cycle repeats.).

Because the distribution of patches is fixed, natural selection can shape developmental programmes with expectations about the distribution of environmental states. For example, if 70 per cent of patches are in one state and the remaining 30 per cent in the other, neonates might 'assume' that they are 70 per cent likely to be in one patch type and 30 per cent likely to be in the other. The distribution of patches constitutes an evolutionarily supplied 'prior' for developing organisms [7,25–27].

Ontogeny consists of 20 discrete, non-overlapping time periods. In each period, organisms choose to either sample an environmental cue or develop by one increment towards one of two phenotypic targets. These choices are mutually exclusive, implying that each sampled cue means one lost increment of specialization. Cues reflect the environmental state with some fixed probability. The cue validity is the probability of receiving a particular cue conditioned on being in the corresponding environmental state. The cue validities for each environmental state are equal: the probability of receiving *cue 1* in environment 1 is equal to the probability of receiving *cue 0* in environment 0. The probability of receiving an incorrect cue is one minus the cue validity. We assume that these cue validities were recurrent features of the evolutionary environment, and as a result are embodied in the developmental programmes of organisms [7,24–26].

Sampling a cue updates the organism's estimate of being in one environmental state or the other according to Bayes' theorem. Natural selection is an optimizing process, and so it might favour information use consistent with Bayesian learning—the optimal way of information updating [26,27]. This assumption does not imply that animals cognitively represent information in Bayesian terms [27,28]. Natural selection is the Bayesian (not organisms), producing optimal developmental programmes. These programmes could be instantiated in organisms as nothing more than a series of simple *if-then* rules (e.g. sample one cue and specialize accordingly).

Phenotypic specialization can translate into fitness in many ways. Here, we consider three: linear returns (the marginal increase in fitness is constant with each increment towards the appropriate phenotype), diminishing returns (the marginal increase decreases with each appropriate increment) and increasing returns (the marginal increase increases). In our model, developmental decisions do not affect fitness during ontogeny; instead, they translate into fitness at some later life stage. Models of predictive adaptive response [22] and time lag (e.g. [11–13]) often share this assumption. However, in our model, the adaptive fit later in life depends on the time invested in progressively constructing a phenotype.

Using a stochastic dynamic programme, we computed optimal developmental policies, specifying the choice that maximizes expected fitness for every possible state of the system (see the electronic supplementary material, appendix S1 for dynamic programming equations). Stochastic dynamic programming is suitable for problems that involve both uncertainty (e.g. about the state of the environment) and interdependent decisions across time (i.e. current decisions affect future options). We explore a range of evolutionary-ecological conditions, varying the fitness function (diminishing, linear and increasing), cue validity (ranging from 0.55 to 0.95) and prior probability (ranging from 0.50 to 0.95).

3. RESULTS
We present four kinds of results, summarized in four figures. At a high level, we compare the reliance on sampling favoured by natural selection across a range of conditions (figure 1). This measures the expected fraction of ontogeny that organisms spend sampling cues, as opposed to specializing. Next, we compare the optimal developmental policies favoured by natural selection across a range of conditions (figure 2). Owing to the stochastic nature of cues, the same developmental programme may generate a range of mature phenotypes (figure 3). Finally, we compare the optimal policies, which may include plasticity, with two developmentally fixed types, who never sample (figure 4).

We first consider a population-level statistic describing the expected fraction of ontogeny organisms spend sampling (figure 1). Consistent with existing models that do not incorporate incremental development (for reviews, see [8,9]), adaptive plasticity is more likely to evolve when cues are moderately to highly informative. However, in our model, the greatest reliance on sampling evolves when cue validities are intermediate (figure 1*a,b*). When cues are highly informative, small samples yield accurate estimates; the opportunity cost of foregone expertise disfavours continued learning. When cues are weakly informative, organisms need many samples for an accurate estimate; however, exclusively sampling garners no fitness (for a related finding, see [29]).
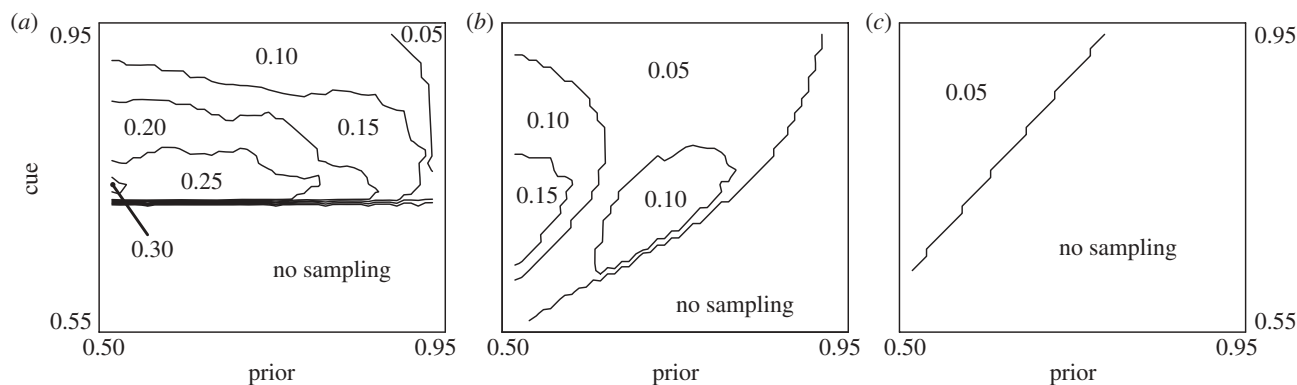
Figure 1. Reliance on sampling. The horizontal axis depicts the prior probability distribution, reflecting the likelihood of being born in one state of the world or the other. The vertical axis depicts the cue validity, the probability of sampling a particular cue, given a particular state of the world. Fitness correlates with the degree of specialization in a (*a*) diminishing, (*b*) linear or (*c*) increasing fashion. The contour lines show the expected fraction of ontogeny spent sampling cues (as opposed to specializing), following the optimal developmental programme. These expectations are taken across the two environmental states, and across individuals.

The reliance on sampling decreases as the prior probability of one environmental state increases relative to the other (figure 1*a*−*c*). As this asymmetry increases, more informative cues are required for sampling strategies to evolve (figure 1*b*,*c*). Selection favours the most reliance on sampling when the gain in fitness returns diminishes and the least when returns increase. With diminishing returns, most of the fitness gains require little expertise; by contrast, with increasing rewards, additional expertise garners more and more fitness.

Despite identical developmental programmes, stochastic sampling may lead individuals to obtain different cue sets and divergent estimates about the environmental state. As a result, organisms may follow different developmental trajectories (figure 2; here, the prior is 0.5; see the electronic supplementary material, appendix S2 for other priors). Organisms may transition from sampling to specialization at different points during ontogeny (figure 2*b*,*c*,*e*). Some individuals obtain a homogeneous sample, resulting in a confident estimate about the environmental state, and specialize early. Others obtain a heterogeneous, uninformative set, and so continue sampling. This leads to the novel hypothesis that inter-individual variation in plasticity may result from inter-individual variation in the consistency of earlier experiences ([30]; see also [31,32]). This hypothesis generates two predictions. First, plasticity should diminish faster across the lifespan in individuals who sample more consistent cues compared with individuals who sample less consistent cues. Second, variation in the timing of switching from sampling to specialization will be largest when cues are moderately informative. With intermediate cues, organisms sample the most extensively, resulting in the most variety of cue sets.

Different evolutionary ecologies may result in similar developmental dynamics (figure 2). For example, when fitness returns increase linearly, minimal sampling evolves with both high and low cue validities (figure 2*d*,*f*). As is well-known, such convergences in outcomes suggest caution when inferring selection histories from phenotypic design. Comparing diminishing and increasing returns, no sampling evolves when cues are poor (figure 2*a*,*g*). In this case, however, the developmental dynamics are different.

With diminishing returns and poor cues, selection favours a (risk-averse) 'generalist' strategy, developing specializations towards both phenotypic targets (figure 2*a*). This strategy maximizes the expected mean fitness *within* a single generation, in contrast to 'conservative bet-hedging' to reduce variance in fitness *across* generations [33]. With increasing returns and poor cues, selection favours (risk-prone) pure 'specialists', who commit at birth to fully specializing towards either one of the phenotypic targets (figure 2*g*). Since the expected fitness of each specialist type is identical, selection is ambivalent about the frequency of the different types. For this reason, this result is not an instance of developmental 'coin-flipping' to reduce variance in fitness across generations [34].

Different developmental programmes generate different distributions of mature phenotypes (figure 3; here, the prior is 0.5; see the electronic supplementary material, appendix S3 for other priors). If the developmental programme calls for no sampling, phenotypic distributions are composed either of 'generalists' (figure 3*a*) or pure specialists (figure 3*g*). When all organisms sample the same number of times, individuals transition from sampling to specialization around the same time, and two distinct types emerge (figure 3*d*,*f*,*h*,*i*). By contrast, when organisms vary in the onset of specialization (i.e. differential plasticity), J-shaped phenotypic distributions result (figure 3*b*,*c*,*e*). Because environmental cues are stochastic, at any instant, an approximately fixed fraction of the population will reach their desired level of certainty and transition from sampling to specialization. This leads to a prediction: when sampling trades off with specialization, the loss of plasticity at the population level should follow a geometric decay-like distribution. Specifically, when organisms rely on minimal sampling (e.g. figure 2*c*), a large fraction transitions from sampling to specialization at any moment, leading to a steep decline: most individuals specialize early in ontogeny, followed by a smaller and smaller number in subsequent time periods. By contrast, when organisms sample more extensively (e.g. figure 2*b*,*e*), the fraction that transitions from sampling to specialization at any moment is smaller, leading to a less steep decline.

Even though selection maximizes the expected fitness of the developmental programmes, not all individuals
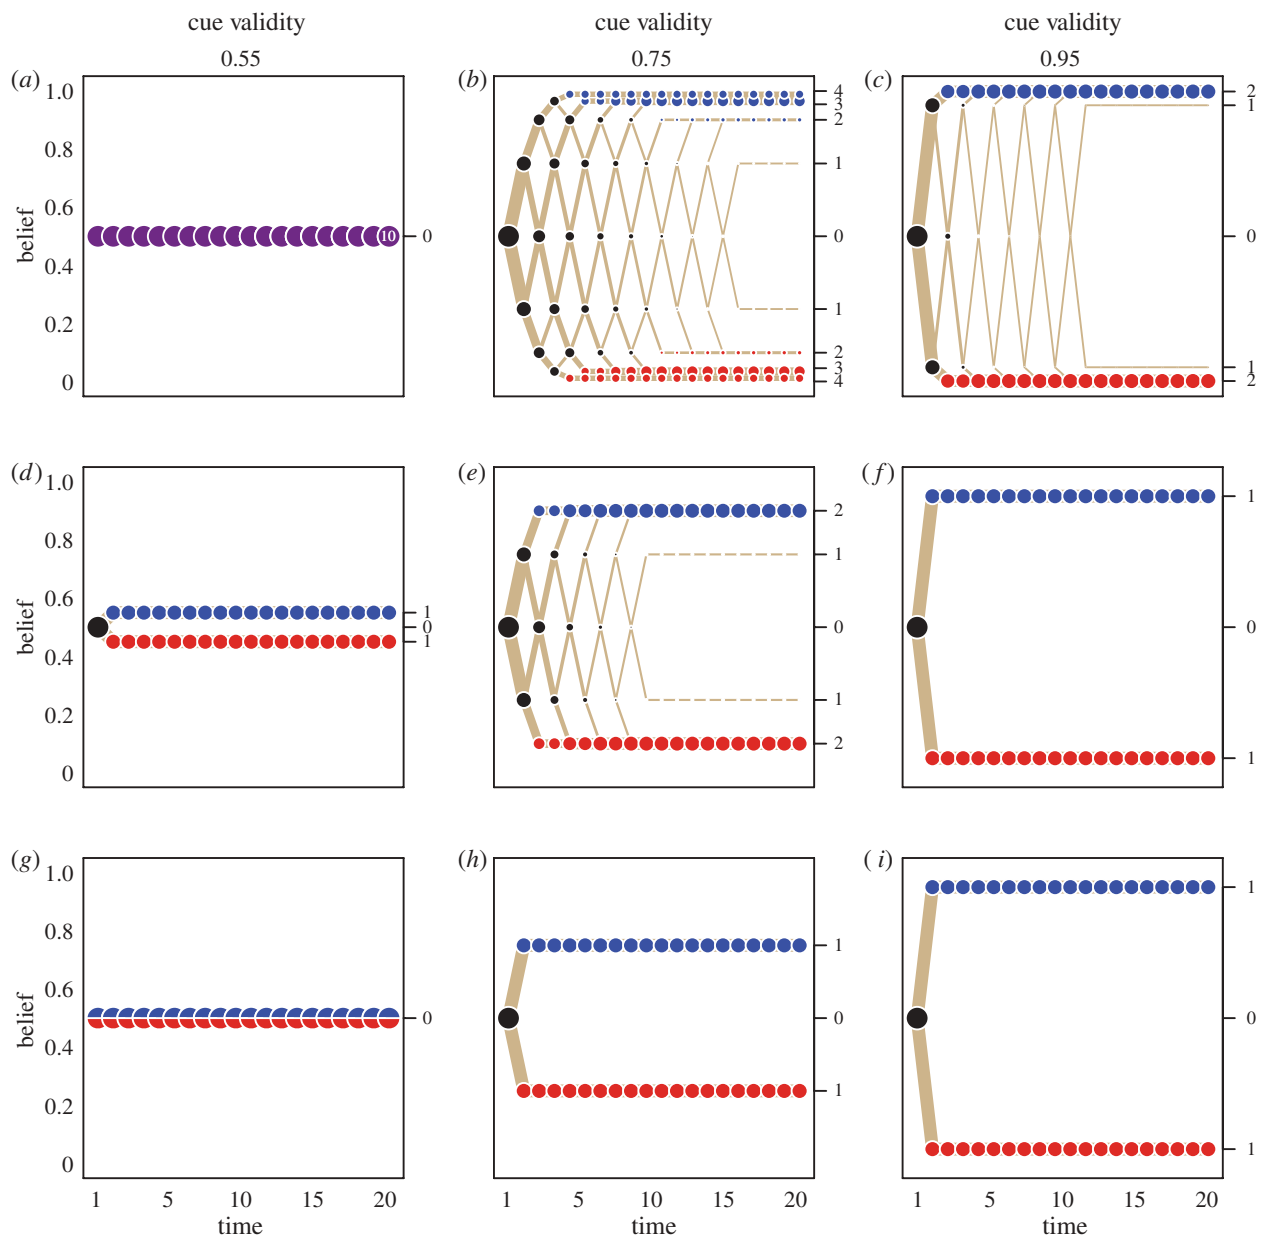
Figure 2. Optimal developmental programmes. The horizontal axis depicts ontogeny. The left, vertical axis depicts the degree of belief of being in World 1. The right, vertical axis depicts the difference in the numbers of cues sampled to each state of the world. Fitness correlates with the degree of specialization in a (a–c) diminishing, (d–f) linear or (g–i) increasing fashion. Sampled cues (a,d,g) weakly, (b,e,h) intermediately, or (c,f,i) strongly indicate the environmental state. Organisms begin ontogeny with a prior probability distribution, reflecting the evolutionary distribution of environments (here, the prior is 0.5; see the electronic supplementary material, appendix S2 for other priors). In each time period, the organism makes one decision: black represents sampling, blue specializing towards phenotype 1 and red specializing towards phenotype 0. The area of a circle is proportional to the probability of reaching a particular state. Within a time period, these areas sum to one. The beige lines represent developmental pathways. The thickness of a line is proportional to the probability of having reached the source state from which the line emanates multiplied by the probability of reaching the destination state.

correctly calibrate. Because sampling trades off with phenotypic development, selection may favour the onset of specialization long before posterior beliefs are strong. When cues are poor, and fitness returns are either linear or increasing, a substantial fraction of organisms miscalibrate development (figure 3d,g). Miscalibration is also common with increasing returns and moderately informative cues (figure 3h), because organisms rely on minimal sampling.

Besides an opportunity cost to learning, there may be additional costs to plasticity, such as developing and maintaining learning machinery [35,36]. Selection favours

learning strategies only when the benefits are sufficiently large to compensate for the constitutive costs of plasticity. Although the optimal policy always achieves (equal or) higher expected fitness than any other strategy (by definition), it is informative to observe *how much* better it does than alternatives. Here, we benchmark the optimal developmental programme, which may include sampling, against two non-plastic strategies (figure 4): a pure 'specialist', which adapts from birth towards the environmental state with the higher prior probability and a pure 'generalist', which specializes halfway towards each phenotypic target. With linear returns, the optimal strategy
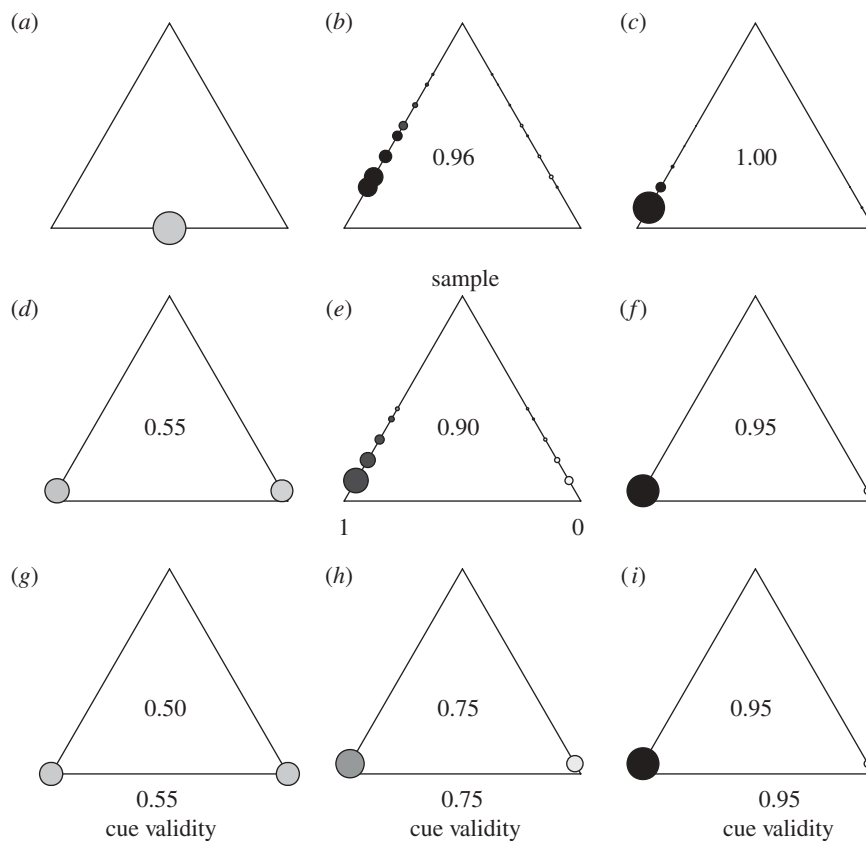
Figure 3. Distributions of mature phenotypes. The triangular plots show phenotypic distributions at the end of development, obtained by simulating 10 000 individuals who follow the optimal developmental programme. To see what a field researcher might observe in a particular ecological setting, we fixed the environment to state 1. Each column of plots depicts a specific cue validity, (a,d,g) 0.55, (b,e,h) 0.75 and (c,f,i) 0.95. Fitness correlates with the degree of specialization in a (a−c) diminishing, (d−f) linear or (g−i) increasing fashion. The evolutionary prior is 0.5, meaning that organisms are equally likely to develop in either environmental state (see electronic supplementary material, appendix S3 for other priors). Phenotypes are characterized by three numbers: the number of sampling bouts (top corner), the degrees of specialization for environment 1 (lower left corner) and environment 0 (lower right corner). The shading of a circle is proportional to the posterior belief: dark represents certainty of being in World 1, white of being in World 0 and grey uncertainty. The areas of circles are proportional to fractions of individuals. Individuals on the right side of the triangle are miscalibrated (i.e. specialized for environment 0 despite developing in environment 1). The number inside the triangle represents the fraction of correctly calibrated specialists: the number of correctly calibrated individuals (left side of the triangle) divided by all the specialists (left and right sides). Generalists, organisms partially specializing for both environments (bottom side), are left out of this fraction.

outperforms the specialist when cues are informative and priors uninformative (figure 4e). The optimal strategy outcompetes the generalist when priors are informative, when cues are informative, or both (figure 4f). With diminishing returns, the optimal strategy outcompetes the specialist when priors are uninformative, regardless of cue validity (figure 4b). The optimal strategy slightly outperforms the generalist when priors are informative, when cues are informative or both (figure 4c). With increasing returns, the optimal strategy outperforms the specialist when priors are uninformative combined with high cue validities (figure 4h), and outcompetes the generalist across the entire parameter space (figure 4i). Across the full range of fitness functions, sampling strategies are favoured over both specialist and generalist strategies when priors are uniformly distributed and cues are informative (figure 4).

4. DISCUSSION
Phenotypic plasticity is ubiquitous in nature [1–4], and theoretical models provide an understanding of when plasticity should evolve [5,7–10,25]. However,

as developmental system theorists have noted, adaptationist models rarely consider ontogeny as a constructive process in which phenotypes incrementally develop ([14,15]; but see [2,3]). In our model, natural selection shapes developmental systems based on the evolutionary ecology, and these systems guide ontogeny such that adaptive phenotypes result through incremental construction.

Our approach shows one way in which evolutionary and developmental processes may jointly construct adaptation, with one causal process (development) nested in the other (evolution). We investigated a trade-off between sampling environmental cues and incremental phenotypic construction. Consistent with existing models, which do not include incremental development [8,9], we find that adaptive plasticity is more likely to evolve when priors are less informative, and when cues are moderately to highly informative; similar developmental programmes can emerge from different selective pressures; with poor cues, generalists emerge when fitness returns are diminishing, whereas specialists emerge with increasing returns; organisms may miscalibrate development when cues are stochastic, despite stable ontogenetic environments.
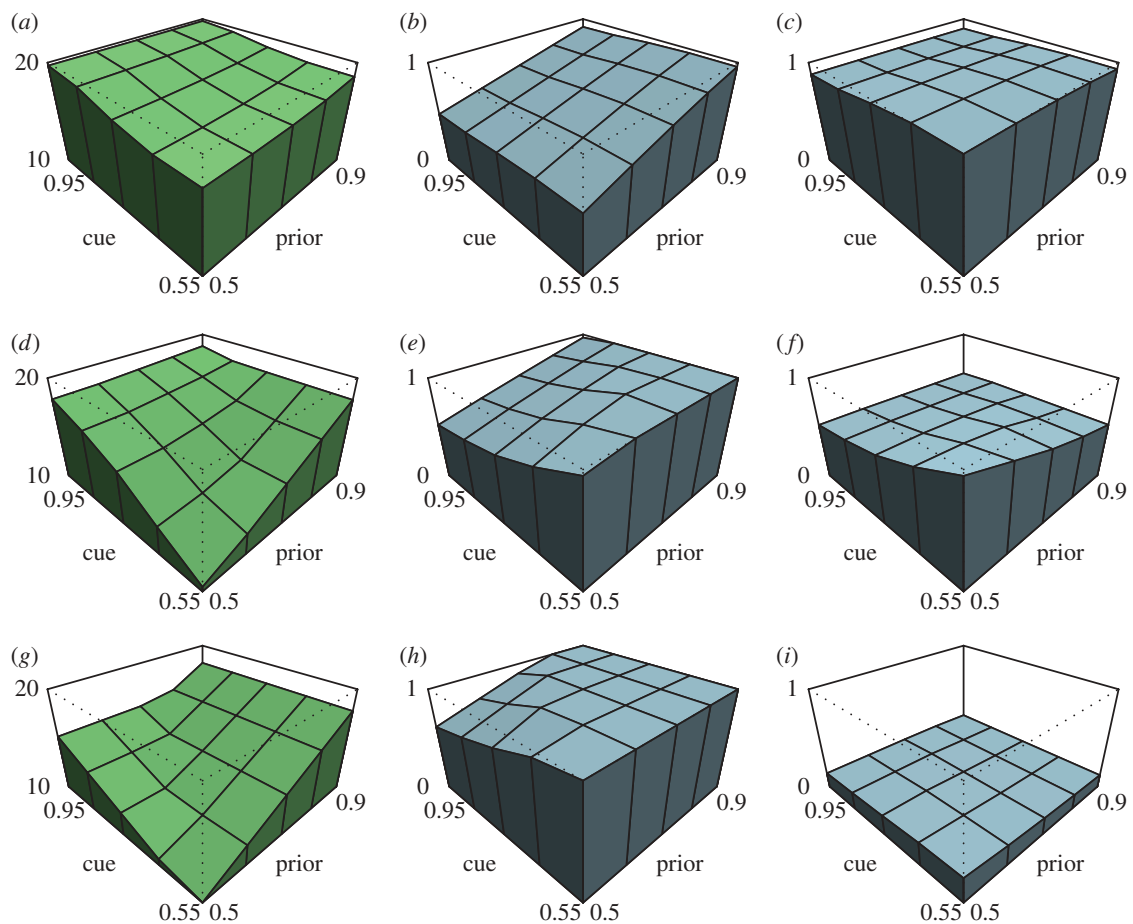
Figure 4. Fitness of the optimal policy. We compare the optimal developmental programme against a 'specialist', which adapts from birth towards the more likely prior environmental state, and a 'generalist', which specializes halfway towards each phenotypic target. For each subplot, the lower right axis depicts a range of priors the lower left axis depicts a range of cue validities, and the vertical axis fitness. The (*a,d,g*) left column of plots shows the expected fitnesses of the optimal developmental programmes. The (*b,e,h*) centre and (*c,f,i*) right columns show the fitnesses of the (non-sampling) specialist and generalist strategies as fractions of the optimal policies' fitnesses. When this fraction equals one, the optimal policy is equivalent to the non-sampling strategy. The top row of plots depicts (*a–c*) diminishing fitness returns, the middle row (*d–f*) linear returns, and the bottom row (*g–i*) increasing returns.

Our model also offers new insights. We focus on the emergence of 'critical windows' and inter-individual differences in plasticity. When cues are either weakly or strongly informative, selection favours all individuals to sample minimally or not at all. In such conditions, the duration of critical windows is developmentally inflexible—that is, not a function of the cues an individual samples—and so species-typical, stage-wise developmental patterns emerge [37]. When cues are moderately informative, individuals may sample extensively, and, because cues are stochastic, some individuals achieve their desired level of certainty early in ontogeny, others late. This process can result in inter-individual differences in durations of critical windows for learning: the critical window closes at different times for different individuals, and so individual differences in plasticity emerge ([30]; see also [31,32]). Specifically, at any time, an approximately fixed fraction of the remaining plastic individuals closes their learning window. At the population level, this process gives rise to a geometric decay-like distribution (J-shaped) of critical windows. As the reliance on sampling increases, the steepness of this distribution decreases.

We believe that our approach could be generative, leading to additional theory exploring how natural selection might shape the dynamics of developmental systems. We

mention one possibility. So far, we have assumed a strict trade-off between sampling and specialization; in any given time period, these choices are mutually exclusive. However, one might imagine that organisms passively acquire environmental cues *throughout* development while constructing their phenotype, with the possibility of switching the direction of phenotypic specialization. If so, individuals might onset specialization in the direction of their strongest prior belief, and possibly switch the direction of development if their beliefs change. Or, individuals might delay specialization, favouring an initial period of passive sampling before beginning to construct a phenotype. Which strategy selection favours could be a function of the costs of switching the direction of phenotypic specialization. This switch cost may be fixed, or it could be a function of the degree of specialization already attained. A model like this addresses the evolution of reversible and irreversible developmental trajectories [9,11], incorporating incremental phenotypic development.

Although we have conceptualized the environmental condition as external to the organism, it is possible to interpret these conditions as stable attributes of the organism (e.g. fighting ability). Organisms might benefit from spending some fraction of ontogeny assessing their own aptitudes (e.g. relative strength), and consequently

adjusting their developmental trajectory. For instance, physically larger individuals may benefit more from an aggressive strategy than smaller individuals, because for them the costs of conflict might be smaller [38]. Recent work in humans suggests that physically stronger men endorse violence more, are more prone to anger, and feel more entitled to better treatment [39]. In addition, stronger men are more extraverted—they more actively approach others, compete for social attention and openly pursue status [40]. For our model to apply to state-dependent personality development [41,42], aptitude assessment must trade-off with specialization, and payoffs should be a function of incrementally developed expertise (e.g. the more time spent fighting, the better the fighter). Our model may yield insights into the conditions in which natural selection favours organisms to learn about their own qualities versus adherence to a fixed strategy, and how much reliance on learning would be optimal.

We presented an adaptationist model of incremental development. We assumed one kind of developmental trade-off, between sampling cues to the environmental state and phenotypic specialization. However, viewing development as gradual and constructive suggests other trade-offs yet to be explored. We hope our model will stimulate this kind of research, drawing together adaptationist approaches and developmental systems theory.

REFERENCES

1 Bradshaw, A. D. 1965 Evolutionary significance of phenotypic plasticity in plants. *Adv. Genet.* **13**, 115–155.

2 DeWitt, T. J. & Scheiner, S. M. (eds) 2004 *Phenotypic plasticity: functional and conceptual approaches*. New York, NY: Oxford University Press.

3 Schlichting, C. D. & Pigliucci, M. 1998 *Phenotypic evolution: a reaction norm perspective*. Sunderland, MA: Sinauer Associates.

4 West-Eberhard, M. J. 2003 *Developmental plasticity and evolution*. New York, NY: Oxford University Press.

5 Dall, S. R. X., Giraldeau, L.-A., Ollson, O., McNamara, J. M. & Stephens, D. W. 2005 Information and its use by animals in evolutionary ecology. *Trends Ecol. Evol.* **20**, 187–193. (doi:10.1016/j.tree.2005.01.010)

6 Donaldson-Matasci, M. C., Bergstrom, C. T. & Lachmann, M. 2010 The fitness value of information. *Oikos* **119**, 219–230. (doi:10.1111/j.1600-0706.2009.17781.x)

7 Levins, R. 1968 *Evolution in changing environments*. Princeton, NJ: Princeton University Press.

8 Meyers, L. A. & Bull, J. J. 2002 Fighting change with change: adaptive variation in an uncertain world. *Trends Ecol. Evol.* **17**, 551–557. (doi:10.1016/S0169-5347(02)02633-2)

9 Piersma, T. & Drent, J. 2003 Phenotypic flexibility and the evolution of organismal design. *Trends Ecol. Evol.* **18**, 228–233. (doi:10.1016/S0169-5347(03)00036-3)

10 Moran, N. A. 1992 The evolutionary maintenance of alternative phenotypes. *Am. Nat.* **139**, 971–989. (doi:10.1086/285369)

11 Gabriel, W. 2006 Selective advantage of irreversible and reversible phenotypic plasticity. *Arch. Hydrobiol.* **167**, 1–20. (doi:10.1127/0003-9136/2006/0167-0001)

12 Jablonka, E., Oborny, B., Molnár, I., Kisdi, E., Hofbauer, J. & Czárán, T. 1995 The adaptive advantage of phenotypic memory. *Phil. Trans. R. Soc. Lond. B* **350**, 133–141. (doi:10.1098/rstb.1995.0147)

13 Padilla, D. K. & Adolph, S. C. 1996 Plastic inducible morphologies are not always adaptive: the importance of time delays in a stochastic environment. *Evol. Ecol.* **10**, 105–117. (doi:10.1007/BF01239351)

14 Lickliter, R. & Honeycutt, H. 2003 Developmental dynamics: towards a biologically plausible evolutionary psychology. *Psychol. Bull.* **129**, 819–835. (doi:10.1037/0033-2909.129.6.819)

15 Oyama, S., Griffiths, P. E. & Gray, R. D. (eds) 2001 *Cycles of contingency: developmental systems and evolution*. Cambridge, MA: MIT Press.

16 Agrawal, A. A., Laforsch, C. & Tollrian, R. 1999 Trans-generational induction of defences in animals and plants. *Nature* **401**, 60–63. (doi:10.1038/43425)

17 Chittka, L. & Muller, H. 2009 Learning, specialization, efficiency and task allocation in social insects. *Commun. Integr. Biol.* **2**, 151–154.

18 Dill, L. M. 1983 Adaptive flexibility in the foraging behavior of fishes. *Can. J. Fish. Aquat. Sci.* **40**, 398–408. (doi:10.1139/f83-058)

19 Kaplan, H., Hill, K., Lancaster, J. & Hurtado, A. A. 2000 A theory of human life history evolution: diet, intelligence, and longevity. *Evol. Anthropol.* **9**, 156–185. (doi:10.1002/1520-6505(2000))

20 Draper, P. & Harpending, H. 1982 Father absence and reproductive strategy: an evolutionary perspective. *J. Anthropol. Res.* **38**, 255–273.

21 Nepomnaschy, P. & Flinn, M. 2009 Early life influences on the ontogeny of the neuroendocrine stress response in the human child. In *The endocrinology of social relationships* (eds P. Gray & P. Ellison), pp. 364–384. Cambridge, MA: Harvard University Press.

22 Gluckman, P. D., Hanson, M. A., Spencer, H. G. & Bateson, P. 2005 Environmental influences during development and their later consequences for health and disease: implications for the interpretation of empirical studies. *Proc. R. Soc. B* **272**, 671–677. (doi:10.1098/rspb.2004.3001)

23 Houston, A. I. & McNamara, J. M. 1999 *Models of adaptive behavior: an approach based on state*. Cambridge, UK: Cambridge University Press.

24 Mangel, M. & Clark, C. W. 1988 *Dynamic modeling in behavioral ecology*. Princeton, NJ: Princeton University Press.

25 Figueredo, A. J., Hammond, K. R. & McKiernan, E. C. 2006 A Brunswikian evolutionary developmental theory of preparedness and plasticity. *Intelligence* **34**, 211–227. (doi:10.1016/j.intell.2005.03.006)

26 McNamara, J. & Houston, A. I. 1980 The application of statistical decision theory to animal behaviour. *J. Theor. Biol.* **85**, 673–690. (doi:10.1016/0022-5193(80)90265-9)

27 McNamara, J. M., Green, R. F. & Olssen, O. 2006 Bayes' theorem and its applications in animal behavior. *Oikos* **112**, 243–251. (doi:10.1111/j.0030-1299.2006.14228.x)

28 Tenenbaum, J. B., Griffiths, T. L. & Kemp, C. 2006 Theory-based Bayesian models of inductive learning and reasoning. *Trends Cogn. Sci.* **10**, 309–318. (doi:10.1016/j.tics.2006.05.009)

29 Todd, P. M. & Miller, G. F. 1991 Exploring adaptive agency II: simulating the evolution of associative learning. In *From Animals to Animats: Proc. first Int. Conf. on simulation of Adaptive Behavior* (eds J.-A. Meyer & S. W. Wilson), pp. 306–315. Cambridge, MA: MIT Press/Bradford Books.

30 Frankenhuis, W. E. & Panchanathan, K. In press. Individual differences in developmental plasticity may result from stochastic sampling. *Perspect. Psychol. Sci.*

31 Belsky, J. & Pluess, M. 2009 The nature (and nurture?) of plasticity in early human development. *Perspect. Psychol. Sci.* **4**, 345–351. (doi:10.1111/j.1745-6924.2009.01136.x)

32 Ellis, B. J. & Boyce, W. T. 2008 Biological sensitivity to context. *Curr. Dir. Psychol. Sci.* **17**, 183–187. (doi:10.1111/j.1467-8721.2008.00571.x)

33 Philippi, T. & Seger, J. 1989 Hedging one's evolutionary bets, revisited. *Trends Ecol. Evol.* **4**, 41–44. (doi:10.1016/0169-5347(89)90138-9)

34 Cooper, W. S. & Kaplan, R. H. 1982 Adaptive 'coin-flipping': a decision-theoretic examination of natural selection for random individual variation. *J. Theor. Biol.* **94**, 135–151. (doi:10.1016/0022-5193(82)90336-8)

35 Auld, J. R., Agrawal, A. A. & Relyea, R. A. 2010 Re-evaluating the costs and limits of adaptive phenotypic plasticity. *Proc. R. Soc. B* **277**, 503–511. (doi:10.1098/rspb.2009.1355)

36 DeWitt, T. J., Sih, A. & Wilson, D. S. 1998 Costs and limits of plasticity. *Trends Ecol. Evol.* **13**, 77–81. (doi:10.1016/S0169-5347(97)01274-3)

37 Ploeger, A., Van der Maas, H. L. J. & Raijmakers, M. E. J. 2008 Is evolutionary psychology a metatheory for psychology? A discussion of four major issues in psychology from an evolutionary developmental perspective. *Psychol. Inq.* **19**, 1–18. (doi:10.1080/10478400701774006)

38 Tooby, J. & Cosmides, L. 1990 On the universality of human nature and the uniqueness of the individual: the role of genetics and adaptation. *J. Personal.* **58**, 17–67. (doi:10.1111/j.1467-6494.1990.tb00907.x)

39 Sell, A., Tooby, J. & Cosmides, L. 2009 Formidability and the logic of human anger. *Proc. Natl Acad. Sci. USA* **106**, 15 073–15 078. (doi:10.1073/pnas.0904312106)

40 Lukaszewski, A. W. & Roney, J. R. 2011 The origins of extraversion: joint effects of facultative calibration and genetic polymorphism. *Pers. Soc. Psychol. Bull.* **37**, 409–421. (doi:10.1177/0146167210397209)

41 Dall, S. R. X., Houston, A. I. & McNamara, J. M. 2004 The behavioural ecology of personality: consistent individual differences from an adaptive perspective. *Ecol. Lett.* **7**, 734–739. (doi:10.1111/j.1461-0248.2004.00618.x)

42 Wolf, M. & Weissing, F. J. 2010 An explanatory framework for adaptive personality differences. *Phil. Trans. R. Soc. B* **365**, 3959–3968. (doi:10.1098/rstb.2010.0215)