



PAPER

Bridging developmental systems theory and evolutionary psychology using dynamic optimization

Willem E. Frankenhuis,¹ Karthik Panchanathan² and H. Clark Barrett³

1. Department of Developmental Psychology, Behavioural Science Institute, Radboud University Nijmegen, The Netherlands and Cognitive Development Center, Central European University, Budapest, Hungary

2. Department of Anthropology, University of Missouri, USA

3. Center for Behavior, Evolution, and Culture, UCLA Department of Anthropology, and FPR-UCLA Center for Culture, Brain, and Development, University of California, Los Angeles, USA

Abstract

Interactions between evolutionary psychologists and developmental systems theorists have been largely antagonistic. This is unfortunate because potential synergies between the two approaches remain unexplored. This article presents a method that may help to bridge the divide, and that has proven fruitful in biology: dynamic optimization. Dynamic optimization integrates developmental systems theorists' focus on dynamics and contingency with the 'design stance' of evolutionary psychology. It provides a theoretical framework as well as a set of tools for exploring the properties of developmental systems that natural selection might favor, given particular evolutionary ecologies. We also discuss limitations of the approach.

Research highlights

- This paper introduces developmental psychologists to dynamic optimization, a method that has generated much interesting work in biology.
- Dynamic optimization can help bridge the current divide between developmental systems theory and evolutionary psychology.
- Dynamic optimization can be used to explore the properties of developmental systems that natural selection might favor depending on environment.
- Dynamic optimization provides an integrative theoretical framework as well as a set of tools.

Introduction

Evolution is the control of development by ecology. (Leigh van Valen, 1973)

Developmental systems theorists and evolutionary psychologists interact less than they should, and when

they do, these interactions tend to be antagonistic (Lickliter & Honeycutt, 2003a, 2003b, 2003c; and commentaries by Buss & Reeve, 2003; Crawford, 2003; Krebs, 2003; Tooby, Cosmides & Barrett, 2003). There have, of course, been both points of agreement (Badcock, 2012; Barrett, 2006, 2007; Bjorklund, 2003; Bjorklund, Ellis & Rosenberg, 2007; Ploeger, van der Maas & Raijmakers, 2008) and disagreement (Dennett, 2011; Sterelny & Griffiths, 1999). Where the sides seem to differ is in explanatory focus. Developmental systems theorists accuse evolutionary psychologists of underplaying the role of developmental causation in building phenotypes (i.e. genetic determinism). Evolutionary psychologists, on the other hand, claim that developmental systems theorists underplay the role of natural selection in organizing development (i.e. unconstrained holism), resulting in a lack of ability to predict species-typical cognition and behavior (Bjorklund *et al.*, 2007). Each 'camp' perceives the other's criticism to be unfair.

The current stalemate is unfortunate because potential synergies between the two fields remain unexplored. These synergies can be built on the assumption, shared by both approaches, that developmental systems are the central units of evolution (Barrett, 2006, 2007; Johnston & Turvey,

Address for correspondence: Willem E. Frankenhuis, Department of Developmental Psychology, Behavioural Science Institute, Radboud University Nijmegen, Montessorilaan 3, PO Box 9104, 6500 HE, Nijmegen, The Netherlands; e-mail: wfrankenhuis@gmail.com

1980; West & King, 1987; West-Eberhard, 2003). This assumption is justified by the fact that all phenotypes are products of developmental processes. Thus evolutionary processes, including natural selection, alter phenotypes by modifying developmental systems – the array of processes that construct organisms, including genes, gene regulatory systems, and other factors (Barrett, 2007; Finlay, 2007; Griffiths & Gray, 2005). This idea has always been central to developmental systems theory (Gottlieb, 1991; Lickliter, 2008; Oyama, Griffiths & Gray, 2001; Lickliter & Honeycutt 2003a, 2003c), and has more recently been incorporated by some evolutionary psychologists (Barrett, 2006, 2007; Bjorklund *et al.*, 2007; Frankenhuys & Panchanathan, 2011a, 2011b; see also Tooby & Cosmides, 1992, pp. 77–93). This point of agreement provides an important scaffold for the integration proposed here.

We hope to show that one approach, called dynamic optimization (Houston & McNamara, 1999; Hutchinson & McNamara, 2000; Mangel & Clark, 1988), has the potential to help forge a constructive bridge. This approach integrates developmental systems theorists' focus on dynamics and contingency with the 'design stance' of evolutionary psychology. In addition, it provides a set of tools for analyzing the properties of developmental systems that natural selection might favor, given particular evolutionary ecologies.¹

We recognize that developmental systems theory is not a single theory, but rather a theoretical perspective that encompasses a family of related ideas (Robert, Hall & Olson, 2001). For instance, developmental systems theorists agree that organisms inherit not only genes, but instead an entire developmental matrix; however, what elements comprise this matrix, and their relative importance (e.g. genes vs. other developmental resources), are debated (Rossiter, 1996; Sterelny, Smith & Dickison, 1996; Griesemer, Haber, Yamashita & Gannett, 2005). We distinguish between a 'hard' and a 'soft' version of developmental systems theory. The 'hard' version maintains that natural selection acts on the replication of entire organism-environment wholes, which comprise every repeatable influence that contributes to the construction of organisms anew each generation. These repeatable influences include not only factors encompassed within the organism's bodily envelope (e.g. genes, gene regulatory systems, etc.), but also factors outside it (e.g. social

and bio-geographical factors), which mainstream biology regards as part of the 'external environment'. Whether this 'hard' version can be integrated with dynamic optimization, and whether doing so would be productive, are open questions that we leave for a future study.

In this article, we consider a 'soft' version of developmental systems theory. Like the 'hard' version, the 'soft' version includes a focus on context, dynamics, contingency, and incremental development. Unlike the 'hard' version, the 'soft' version views organisms as separate from their environments and natural selection as maximizing the fitness of developing organisms (not of entire organism-environment wholes). What we mean by 'separate' is not that organisms are *independent* of their environments; indeed, organisms are always embedded in their environments. Instead, what we mean is that the properties of organisms and the properties of environments can be treated as distinct components in an evolutionary model, despite being causally *related*. Thus, in what follows, we treat 'developmental systems' as synonymous with 'developing organisms', dynamic entities comprising genetic, molecular, and cellular interactions at multiple levels, which are shaped by their external environments, but distinct from them. We acknowledge that some developmental systems theorists would reject this 'soft' version. We note, however, that our approach can incorporate bi-directional relations between organisms and their environments; i.e. developing organisms are not only shaped by their environments, but also reciprocally shape their own environments. We do not discuss such 'niche-construction' in detail (see Flynn, Laland, Kendal & Kendal, 2013; Laland & Sterelny, 2006; Laland, Sterelny, Odling-Smee, Hoppitt & Uller, 2011), but the dynamic optimization approach presented here can accommodate it.

Dynamic optimization is not a novel approach; it is commonly used in mathematics, economics, engineering, biology, and more recently in cognitive neuroscience (Busemeyer & Pleskac, 2009; Chhabra & Jacobs, 2006; Drugowitsch, Moreno-Bote, Churchland, Shadlen & Pouget, 2012), clinical psychology (Moodie, Richardson & Stephens, 2007; Murphy, 2003), and the study of sensorimotor control (Engelbrecht, Berthier & O'Sullivan, 2003; see also Körding & Wolpert, 2006; Todorov, 2004). However, the approach is rarely employed in developmental psychology. Here, we hope to show that dynamic optimization offers a fruitful bridge between evolutionary and developmental psychology.

Natural selection as optimization

In biology, the morphology, physiology, and behavior of plants and animals is often analyzed from an

¹ Evolutionary ecologies, in this context, are not a particular historical place or time, but rather the statistical composite of selection pressures that caused the design of an adaptation across generations (Tooby & Cosmides, 1992). Developmental ecologies exist on within-generation timescales (and are thus nested within evolutionary ecologies), and are those environments experienced by an adaptation between its conception and death.

optimization perspective (Fawcett, Hamblin & Giraldeau, 2013; Gardner, 2009; Grafen, 1984; Maynard Smith, 1978; Mayr, 1983). Optimization theories presume that there is a fit between the design of phenotypes and the adaptive problems they solve, where ‘optimality’ typically means optimality under constraint (i.e. natural selection does not produce perfect solutions, but selects among the best available ones). In this sense, for example, the wings of various bird species are optimized for the problems of flight that they face. Hummingbird wings are tailored for maneuverability, vultures’ for soaring, and gannets’ for diving at extremely high speeds (over 60 mph), enabling them to catch fish much deeper than most airborne birds. In each case, the process of natural selection has shaped these phenotypes to fulfill specific adaptive functions.

Developmental systems theorists object to viewing natural selection as a ‘designing force’ (Lickliter, 2008; Lickliter & Honeycutt, 2003a), because the process has no agency – no intentions or foresight. It is true, of course, that natural selection is a process without agency, as recognized by Darwin himself (indeed, the lack of agency was precisely what made his theory controversial). No modern-day biologist would attribute agency to natural selection. However, adaptations can have the ‘appearance of design’ (Williams, 1966), and their component parts can be understood, and in some cases predicted, with reference to their adaptive function (Gardner, 2009; see Grafen, 2007, for a formal justification of studying organisms as optimized by natural selection). For instance, functional considerations helped William Harvey discover that blood circulates through the body, rather than being consumed by organs, and that the heart functions as a pump (Mayr, 1983). We believe this particular disagreement is mostly about semantics, as long as both ‘camps’ agree on how natural selection actually works: it is the selective retention of inherited variations, with ‘variations’ referring not to genes alone, but to entire developmental systems (including genes, gene regulatory systems, etc.).

In ‘design’ language, we may suppose that natural selection favors developmental systems that tend to construct adult phenotypes that are successful, relative to other variants, at surviving and reproducing (Barrett, 2006, 2007; Geary, 2006; Geary & Bjorklund, 2000). In this sense, the properties of developmental systems will reflect, in terms of a form–function fit, the selection pressures that created them – in the same way that ever-faster running speed in gazelles reflects the enduring challenge of escaping mobile predators. Our approach, therefore, will be to analyze developmental systems through the lens of an engineer, focusing on the adaptive problems they solve for developing organisms.

There are many different ways to implement an engineering analysis (for examples, see de Beer, 2000; Frank, 1996; Osman, 2010). However, at a general level, what is required is a hypothesis regarding the problem that the system solves, the environments in which the system was designed to solve the problem, a specification of the features that the system brings to bear in solving the problem, a description of how the properties of the system and the properties of the environment interact to produce outcomes, and an analysis of how well the system performs in terms of its target solutions (Johnston & Turvey, 1980; Marr, 1982; Tooby & Cosmides, 1992). Here, we hope to show that dynamic optimization provides one (although not the only) suitable framework for modeling evolved developmental systems.

Dynamic optimization belongs to a broader class of optimization approaches (Maynard Smith, 1978). All these models share certain assumptions, but each has some unique features, which are useful for particular purposes. We focus on dynamic optimization because this approach is suitable for examining what kinds of dynamical systems natural selection might favor, given particular evolutionary ecologies. It allows us to ask such interesting questions as: When does natural selection favor developmental mechanisms that use properties of the external environment to guide development? How does this sensitivity depend on the reliability of these properties being present? And how does it depend on the system’s prior probability of developing within specific environmental conditions?

Our treatment will be primarily conceptual, but we also include brief discussions of how dynamic models can be implemented mathematically (Sections 3 and 4) and through computer simulation (Section 6). We provide references here (Houston, Clark, McNamara & Mangel, 1988; Mangel & Clark, 1988; Houston & McNamara, 1992, 1999; Mangel & Ludwig, 1992), and elsewhere in the paper, for readers seeking more information.

Static optimization

Before discussing dynamic optimization, we first discuss static optimization. Static optimization is useful when organisms make only a single ‘decision’, or multiple decisions that are independent of each other, so that previous decisions do not constrain future options. For example, an organism might make a single ‘decision’ to invest in growth or reproduction, to develop a particular level of vigilance, to stay or migrate, etc. Because static optimization does not model the changing state of the organism, it is not suitable for analyzing dynamical

aspects of developmental systems. Nonetheless, the method can generate insights that are useful to developmental psychologists. Let us consider an example.

Developmental plasticity refers to the ability of a developmental system to produce different phenotypes depending on the environment in which it develops (e.g. resulting from experience-dependent systems interacting with local ecologies; Greenough, Black & Wallace, 1987). Many psychologists are aware that capacities for developmental plasticity tend to evolve in environments that are variable across evolutionary time. However, it is less well known that different kinds of environmental variation may result in very different adaptations. In certain cases, static optimization can illuminate why.

One pivotal distinction is between spatial and temporal environmental variation.² In spatially varying environments, the environment is broken up spatially into different patches (like a mosaic), each with a particular state (e.g. in the simplest possible example, two states, safe or dangerous). Within a given generation, offspring may be born into either kind of patch; however, across generations, organisms always face the same spatial distribution. In spatially varying environments, parents' fitness can be computed by taking the arithmetic mean of the fitness of all their offspring;³ that is, the fraction of offspring in patch A multiplied by their fitness, plus the fraction of offspring in patch B multiplied by their fitness, and so forth. This implies that even if some offspring attain low fitness (e.g. they die), parents' fitness need not suffer greatly, as long as their other offspring attain high fitness.

Now consider temporal variation. In a temporally varying environment, all offspring within a generation are born into the same environmental state (e.g. safe); however, in the next generation, all individuals may be born in a different environmental state (e.g. dangerous). In a temporally varying environment, the entire population experiences variation across generations. In such environments, long-term fitness depends on the fitness of the lineage in generation 1, multiplied by the fitness in generation 2, and so forth. The average fitness of this series will not be their arithmetic mean, but their multiplicative, or geometric, mean – the n -th root of the product of n fitness values (Dempster, 1955). With this knowledge, we can develop a simple static optimization model showing that spatial and temporal variation may result in different adaptations.

Imagine two developmental systems (DS 1 and DS 2) that are exposed to two environmental states (E1 and E2), with equal probability (50%). Suppose that DS 1 constructs a similar phenotype in E1 and E2, which is highly fecund in E1, producing as many as nine offspring, but unproductive in E2, producing only one offspring. In contrast, DS 2 is plastic, producing different phenotypes in E1 and E2, each quite well matched to its environmental state. DS 2 performs worse than DS 1 in E1, but better than DS 1 in E2; say, four offspring in each environment. Which is favored when? In a spatially varying environment, DS 1 attains higher fitness than DS 2, because its arithmetic mean is greater: $(9 + 1)/2 > (4 + 4)/2 = 5 > 4$. However, in a temporally varying environment, the geometric mean of DS 2 is greater: $\sqrt[2]{(9 \times 1)} < \sqrt[2]{(4 \times 4)} = 3 < 4$.

Thus, spatial and temporal environmental variation may result in different adaptations, because they entail different fitness calculations. In particular, temporal variation more often selects for risk-averse strategies, because instances of low fitness have multiplicative (geometric) effects (Philippi & Seger, 1989). We just observed that parents might reduce temporal variance in fitness by producing plastic offspring – this is called 'conservative bet-hedging', because it avoids extremes. A second way in which parents can reduce fitness variance is by producing phenotypically variable offspring. This strategy of 'diversified bet-hedging' ensures that at least a fraction of offspring will be well adapted to the next generation, whatever its state (Childs, Metcalf & Rees, 2010). Indeed, Belsky and Pluess (2009) hypothesized that individual differences in plasticity in humans may be the result of parents hedging their bets.

Formal models (Philippi & Seger, 1989) and empirical research on plants and animals (Childs *et al.*, 2010) show that temporal variation leads to selection for bet hedging when fitness effects – i.e. the effects of the focal trait on the organism's overall fitness – are large. This is because greater variance in fitness implies lower geometric mean fitness, increasing the benefits of reducing this fitness variance by producing variable offspring. Static optimization models thus suggest novel avenues for testing the 'diversified bet-hedging' hypothesis of differential plasticity (Belsky & Pluess, 2009): we should expect large differences in fitness between low- and high-plasticity individuals in a given environment (Frankenhuis, Panchanathan & Belsky, in preparation).⁴ This prediction remains to be tested in humans.

² Some authors use the terms individual level and population level to distinguish these two kinds of environmental variation.

³ We assume that generations are non-overlapping: organisms are born; develop; reproduce; mature individuals die, and the cycle repeats. Parents and offspring do not coexist.

⁴ This prediction assumes that current environments have similar fitness-relevant properties to the ecologies (EEA, environments of evolutionary adaptedness) in which the relevant mechanisms evolved (i.e. no evolutionary disequilibrium).

Dynamic optimization

Development, by definition, involves change over time. These changes cannot be captured within the framework of static optimization, but they can be modeled using dynamic optimization. The goal of dynamic optimization is similar to that of dynamical systems theory: to understand how a system changes over time depending on interactions between its internal state and external influences (van der Maas, 1995; De Beer, 2000; Smith & Thelen, 2003). However, dynamic optimization specifically aims to find the best course of action for each possible state of the system (the optimal policy; Houston & McNamara, 1999; or in game theory, the optimal strategy; Maynard Smith, 1982). In our case, we want to determine optimal (i.e. fitness-maximizing) developmental ‘decisions’,⁵ which, combined, form optimal developmental trajectories.

In dynamic models, developing organisms are represented by a set of state variables. State variables are quantities that describe the state of the system and that are sufficient to predict its state, if only probabilistically, at future times. State variables can represent factors internal to the organism (e.g. energy, knowledge, skill) or external to it (e.g. the distribution of resources in the environment). Each state variable takes on a specific value (e.g. energy = 4, knowledge = 6; skill = 9), which might change over time, and the complete set of values describes the ‘state’ of the organism.

If we know the organism’s state and the policy it follows, then we can predict its decisions, but not necessarily its next state because decisions may be probabilistically (rather than deterministically) linked to outcomes. For instance, if an animal chooses to forage, it may or may not find a resource, and so its energy may be replenished or depleted in the next period. Thus, policies should be viewed as instructions (‘go forage’), not as outcomes (‘finding a resource’). Dynamical systems are called stochastic when the same decision may lead to more than one outcome, and deterministic when the state resulting from a decision is certain. If future states result stochastically from decisions, the payoff of these future states should be weighted by their probability (i.e. ‘expected payoffs’).

We will now provide an example of dynamic optimization. Although our example is simple, it captures an important tradeoff that many organisms (including humans) face, between growth and reproduction. All else being equal, organisms with higher fecundity should

be favored by natural selection. However, individuals who begin reproducing early in life may experience costs to fitness, including reduced growth and development (for evidence in humans, see Helle, 2008). Precisely how organisms should optimally allocate resources to growth and development, versus allocating them to reproduction, will depend on the nature of this tradeoff, which is in turn influenced by features of the environment.

One factor that has been linked to faster maturation and earlier reproduction in humans and other animals is the rate of extrinsic mortality (Belsky, Steinberg & Draper, 1991; Chisholm, 1993; Frankenhuis & Del Giudice, 2012) - the risk of dying due to factors outside of one’s own control, such as predation, accidents, disease, etc. (Ellis, Figueredo, Brumbach & Schlomer, 2009; Griskevicius, Delton, Robertson & Tybur, 2011). In humans, higher rates of extrinsic mortality predict earlier onset of menarche, earlier age of first sex, and earlier production of first offspring (Ellis, 2004; Walker, Gurven, Hill, Migliano, Chagnon, De Souza, Djurovic, Hames, Hurtado, Kaplan, Kramer, Oliver, Vaggia & Yamauchi, 2006).

We can formulate a dynamic model linking extrinsic mortality to the growth–reproduction tradeoff. Consider a female that lives for three years. In each year, she has two options: investing in growth or reproduction. During each year, she faces some fixed probability of surviving (s), which we assume to be independent of the developmental decision to grow or reproduce. Females who reproduce earlier sacrifice some of their own growth, hence have smaller bodies (Helle, 2008); as a result, they have fewer physical resources to invest in each offspring, and so these have lower viability (e.g. an increased risk of juvenile mortality) compared with offspring of mothers who developed a larger body. We capture this idea as follows: Each offspring yields $1-c^{1+n}$ to the parent – with c denoting the reduction in offspring viability (ranging between 0 and 1), and n denoting the number of time periods a parent has invested in her own growth. Thus, a parent who invested zero, one, and two periods in her own growth garners, respectively, $1-c$, $1-c^2$, $1-c^3$ for a given offspring. How should natural selection shape reproductive trajectories, depending on parents’ survival probability (s) and juvenile mortality (c) (see Figure 1, for a graphical representation of the decision tree)?

To calculate the payoffs of different developmental trajectories we use a technique called backwards induction. Backwards induction helps resolve problems of combinatorial explosion that can arise with an exhaustive search, if every decision (and so all developmental trajectories) were to be explored. This is not a concern in models with low dimensionality, like the current example, but it quickly becomes a major issue when more variables and/or time periods are included (Kokko, 2007; Roff,

⁵ We use the term ‘decisions’ metaphorically: ontogeny results from mechanistic processes, which need not be conscious or intentional.

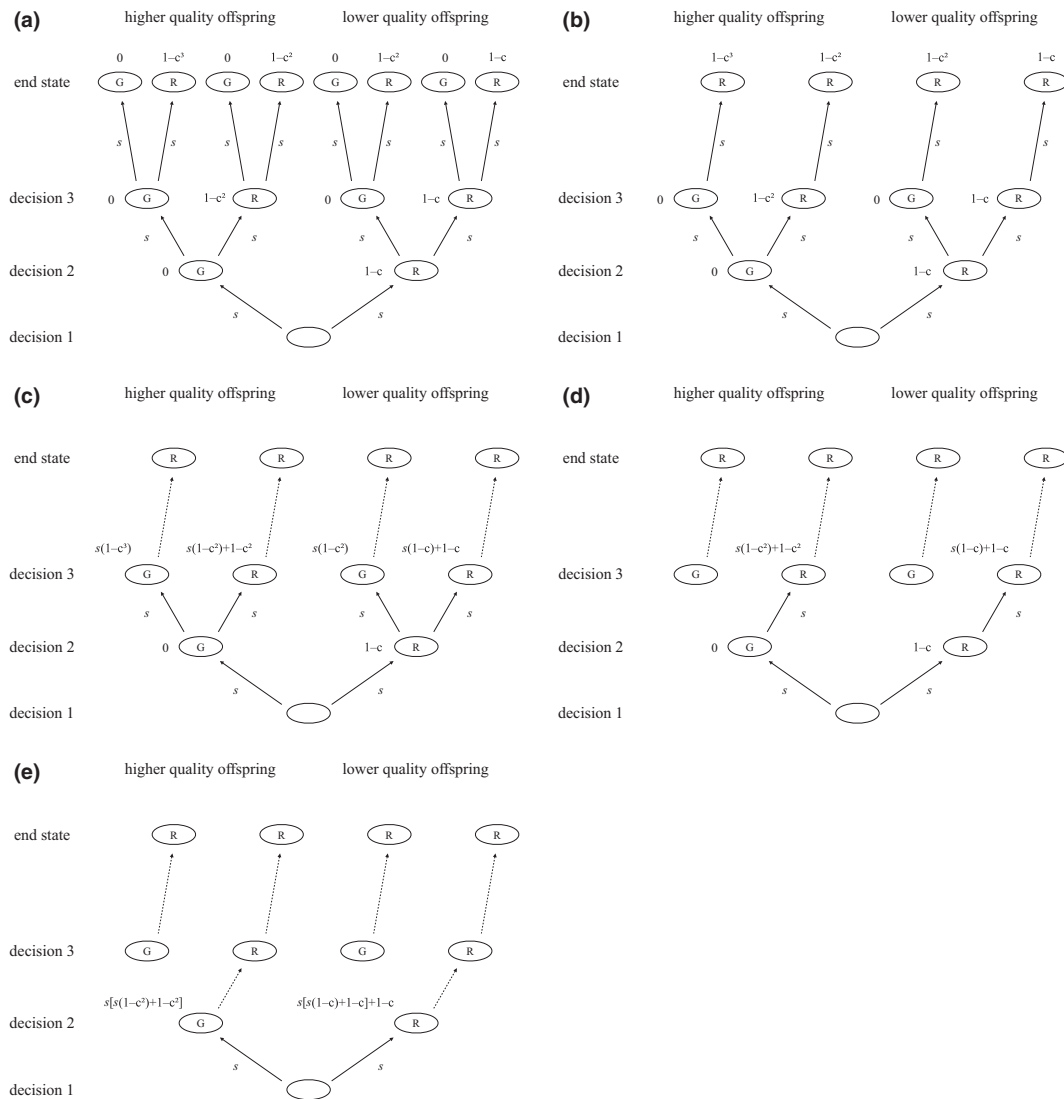


Figure 1 A schematic representation of a tradeoff developing organisms may face between investments in their own growth versus earlier reproduction. The model has eight end states (Figure 1a). In all four penultimate states, individuals should choose to reproduce. Hence, we eliminate the ‘growth’ trajectory emanating from each penultimate state (Figure 1b), and assign to it the expected optimal payoff associated with having reached that state, which equals the sum of the current payoff plus the expected gain given making the optimal decision (Figure 1c). In the antepenultimate period, there are two states, and in each of these, irrespective of the values of c and s , the optimal decision is to reproduce. Hence, we eliminate the ‘growth’ option emanating from them (Figure 1d), and assign to it the expected optimal payoff associated with having reached that state (Figure 1e). In the starting state, the optimal decision will depend on the values of c and s .

2010). Instead of moving forward through the state space, exploring every possibility for every state, this technique begins at the final time period and moves backwards, pruning the suboptimal branches (i.e. choices resulting in lower fitness), thus reducing the search space. Backwards induction does not imply that evolution works backwards; a complete forward search (e.g. via genetic algorithms) would yield the same answer. Nor does it imply that developmental processes can only be understood retro-

spectively. Rather, once we know the optimal policy (either through forward or backward search), we can make prospective predictions about how developmental systems will respond – both in terms of state change and behavioral response – to impinging influences.

We begin by determining the fitness associated with each of the end states. The current model has eight end states (Figure 1a). Four of these denote growing in the final period, and the other four reproducing. The payoff

associated with attaining an increment of growth in the final period is zero (i.e. no offspring are produced, and investments cannot be cashed out later); hence, in all four penultimate states, individuals should choose to reproduce ($1-c^{1+n} > 0$) (decision 3). Therefore, we can eliminate the 'growth' trajectory emanating from each penultimate state (Figure 1b), and assign to each penultimate state the expected optimal payoff associated with having reached that state. This expected payoff equals the sum of the current payoff plus the expected gain given making the optimal decision (Figure 1c). Now we move backwards.

In the antepenultimate period, there are two states ('G' and 'R'), and in each of these states the organism can, again, choose to grow or reproduce (decision 2). In contrast to the penultimate state, optimal decisions are not obvious in the antepenultimate period; we have to do some algebra to find the solution.⁶ It is possible to prove⁷ that in both antepenultimate states, irrespective of the values of c and s , the optimal decision is to reproduce. Hence, we eliminate the 'growth' option emanating from each antepenultimate state (Figure 1d), and assign to each of these states the expected optimal payoff associated with having reached that state (Figure 1e).

In the starting state, individuals again face a choice between growing and reproducing (decision 1). In this case, the optimal decision cannot be determined independently of specific parameter values; i.e. the optimal decision will depend on the values of c and s . Comparing the expected values associated with each decision, and simplifying, shows that organisms should invest in growth if $c(1-c)(s + s^2) > 1-c$. The right term represents the benefit of reproducing compared with growing in the first time period (i.e. only individuals who reproduce obtain this benefit). The left term represents the benefit of growing compared with reproducing in the first time period. We can decompose the left term as follows. Individuals who grow in the first period produce

offspring yielding $1-c^2$, whereas early reproducers garner $1-c$ for a given offspring. The marginal benefit of investing one period in growth is thus $1-c^2-(1-c)$, equaling $1-c^2-1+c$, equaling $-c^2+c$, equaling $c(1-c)$. Investing in growth in the first period yields this benefit in the second period (discounted by probability s) and in the third period (discounted by probability s^2).

We set out to analyze how natural selection should shape reproductive trajectories, depending on parents' survival probability (s) and juvenile mortality (c). What can we learn from the expression $c(1-c)(s + s^2) > 1-c$? One insight is that when expected lifespan is short (i.e. when extrinsic mortality is high), parents should invest less in their own growth in order to reproduce earlier. This result is consistent with the recurring finding that accelerated maturation correlates with high rates of extrinsic mortality in many animals, including humans (Ellis, 2004; Walker *et al.*, 2006). It is also consistent with research showing that individuals subjectively expecting a short lifespan discount the future more, focusing instead on immediate rewards. Based on qualitative interviews, longitudinal survey data, and existing demographic analyses, Brezina, Tekin and Topalli (2009) conclude: 'young people who perceive a high probability of early death may have little reason to delay gratification for the promise of future benefits, as the future itself is discounted. Consequently, these young people tend to pursue high-risk behaviors associated with immediate rewards, which include crime and violence' (p. 1091; for related experimental work, see Griskevicius, Tybur, Delton & Robertson, 2011).

Our expression $c(1-c)(s + s^2) > 1-c$ also yields a second insight: When juvenile mortality is high, natural selection more likely favors parents who invest in their own growth in order to increase offspring viability. This result becomes apparent when we divide both sides by $1-c$, yielding $c(s + s^2) > 1$. This finding is also intuitive: If juvenile mortality is high, natural selection favors parents who reduce this risk. However, even simple models like this one can raise interesting questions. For example, our modeling results show that organisms never reproduce before they grow: if an individual chooses to grow, she will always do this in the beginning. One might wonder whether this finding could be linked to the empirical observation that animals and plants do grow substantially before they begin reproducing. Of course, additional work would be needed to know whether this result generalizes, for instance, when more developmental periods are included.

We have shown that dynamic optimization can provide insights into evolved developmental systems. However, both our current example and our example of static optimization assume that developmental systems 'know'

⁶ A computer solving backwards induction problems would not determine optimal decision algebraically, as we do here; instead, our parameters (c and s) would simply be values, and the program would compute which decision in a given state results in the largest expected value, store this decision, and move backwards.

⁷ In the antepenultimate period, when in the 'growth' state, individuals should invest in growth if $s(1-c^3) > s(1-c^2) + 1-c^2$. To see why this never obtains, divide both sides by s ; deduct one from each side; add c^2 to both sides; multiply both sides by s ; taking a c out of the parentheses in the left term yields: $sc(c-c^2) > 1-c^2$. This is never true, because $c-c^2 > 1-c^2$ never holds (as $c < 1$). In the 'reproduction' state, individuals should invest in growth if $s(1-c^2) > s(1-c) + 1-c$. This never obtains either. To see why, distribute the s 's on both sides; deduct s from both sides; add sc to both sides; take sc out of parentheses on the left yields: $sc(1-c) > 1-c$, which never holds.

– i.e. embody in their design⁸ – all the parameter values relevant to their decisions. For instance, in the current example, individuals know the value of s when making their decisions; they do not estimate it – no new information is acquired during ontogeny. This limitation is lifted in adaptive dynamic models.

Adaptive dynamic optimization

Adaptive dynamic optimization should be of particular interest to developmental psychologists, because this approach can deal with state-dependent systems that change over time and that can acquire new information (Mangel & Ludwig, 1992; Houston & McNamara, 1992). For example, a developmental mechanism may be uncertain about the current state of the environment (e.g. safe or dangerous). However, it can sample cues to the environmental state in order to update its (genetically or epigenetically) inherited prior ‘belief’,⁹ resulting in a posterior belief. The changing knowledge state of the system, like other aspects of the phenotype, can be represented by state variables.¹⁰

We will now discuss an example of adaptive dynamic optimization. This example incorporates developmental dynamics, context sensitivity, and information acquisition, and also incremental construction, another key emphasis of developmental systems approaches (Oyama *et al.*, 2001). Our example will, hopefully, demonstrate how each approach benefits from integrating essential tenets of the other.

Consider a developing organism that faces the challenge of constructing an adaptive phenotype in an environment that remains stable across its lifetime, yet the parameters of which are not known to it, because

⁸ When we say that organisms ‘know’ the state of the world, we do not imply pre-specified representations in the developmental system. We mean that organisms adjust development in a way consistent with them ‘knowing’ the state of the world. For instance, if a lineage evolves in a predator-rich environment and, as a result, developmental systems are selected to construct protective armor, these systems are said to ‘assume’ that the world is predator-rich.

⁹ Again, we are using the language of belief formation metaphorically. The developmental system is merely predicted to behave ‘as if’ it holds these beliefs. Optimization models provide computational descriptions, not cognitive representation or neural implementation (Marr, 1982).

¹⁰ There is no limit to the number of state variables except the power and memory of one’s computer, and thus the method will allow for multiple domains of uncertainty and associated learning tracks. However, the ‘curse of dimensionality’ looms large; adding an extra dimension to a mathematical space will exponentially increase the volume of this space (Bellman, 1961). Therefore, in practice, dynamic optimization models will involve a limited set of variables.

over the course of evolution, its ancestors experienced different environmental states (Frankenhuis & Panchanathan, 2011a, 2011b; Frankenhuis & Del Giudice, 2012). As a result, the organism’s developmental systems reflect – in terms of a form–function fit – the prior probabilities of encountering these environmental states (McNamara, Green & Olssen, 2006). We also assume that environmental variation is spatial, not temporal; within each generation, some fraction of the lineage (say, 50%) develops and reproduces in a safe environment, and the rest in a dangerous environment. Finally, because across evolutionary time probabilistic cues indicating the environmental state were available, organisms evolved perceptual and cognitive abilities for using these cues to improve their estimates of the environment.

Many biological models have examined the conditions in which organisms evolve sensitivity to environmental cues in order to guide development (reviewed in Meyers & Bull, 2002; Schlichting & Pigliucci, 1998). These models typically assume a two-stage life history; organisms first sample a cue to the state of the world, and subsequently develop phenotypes, either instantaneously (e.g. Moran, 1992) or after a time lag (e.g. Padilla & Adolph, 1996). As noted by developmental systems theorists (Oyama *et al.*, 2001), models such as these tend to leave out that development is typically a constructive process in which phenotypes incrementally adapt to local ecologies (for notable exceptions, see Mangel & Clark, 1988; Houston & McNamara, 1999; Schlichting & Pigliucci, 1998).

The constructive nature of developmental processes is especially apparent in organisms that start building their traits during a life stage at which they are not needed yet (see Gluckman, Hanson, Spencer & Bateson, 2005). Such ‘anticipatory construction’ may be adaptive when it takes time to incrementally develop a phenotype, or when possessing the trait very early in life (e.g. at birth) is adaptive. For instance, mothers may program their offspring (e.g. inside the egg or the womb) to start developing protective armor that will help to reduce predation risk in the postnatal environment (e.g. Agrawal, Laforsch & Tollrian, 1999), fur coats for staying warm, coloration, wing shapes, and even adaptive behavior patterns that are adaptive in the postnatal environment (Gluckman *et al.*, 2005).

When the time allocated to building a phenotype correlates with the adaptive fit to the environment – such that the more time invested, the better the fit – organisms may benefit from starting to specialize earlier in ontogeny (Frankenhuis & Panchanathan, 2011a, 2011b). For instance, children growing up in hostile environments may tailor their perceptual and mental abilities to recognizing dangers, perceiving angry facial expressions

more accurately than other children (Pollak, 2005, 2008), developing better abilities for deception (Mealey, 1995), and specializing in behaviors that effectively avoid or impose harm, such as fighting ability. In contrast, children growing up in a safe environment may invest in growth and development (Helle, 2008), acquisition of skills and knowledge, and constructing long-term cooperative networks (Belsky *et al.*, 1991).

If, for a given trait, earlier specialization allows more time to develop an adaptive phenotype, should developing organisms start specializing as early as possible? Despite its benefits, earlier specialization may entail costs. For instance, earlier specialization implies less time for assessing the state of the environment, and therefore a greater risk of developing a phenotype not matching the environmental state (see Frankenhuis & Panchanathan, 2011a, 2011b, for further discussion of costs associated with earlier specialization). Developmental mismatch may impose major health costs in humans, including elevated risk of cardiovascular diseases, obesity, and diabetes (Barker, 1994; Monaghan, 2008; Gluckman *et al.*, 2005).

Thus, specializing early may lead to a more adaptive phenotype, but sampling more cues likely yields a more accurate estimate of the environmental state, reducing the risk of mismatch. How should a given developmental system optimally tradeoff sampling of environmental cues and phenotypic specialization, depending on environmental conditions? This question combines a focus on dynamics and contingency with the idea that system flexibility itself is the product of natural selection. In a recent paper (Frankenhuis & Panchanathan, 2011a), we provide a detailed analysis of this problem. We computed optimal policies using adaptive dynamic optimization for a range of evolutionary environments. These policies – which specify the optimal ‘decision’ for each state in each developmental period – provide insights into the developmental systems that natural selection might favor, given particular ecologies.

In our model, we explore such variables as the prior probabilities of each environmental state – for example, individuals may be equally likely to be born in either environmental state, or more likely in one than the other. We also varied the value of information – cues may be highly informative about the environmental state, or weakly informative. And we explored different mapping functions from specialization to fitness – linear fitness (the marginal increase in fitness is constant with each increment towards the environment-appropriate phenotype), diminishing fitness (the marginal increase decreases with each environment-appropriate increment), and increasing fitness (the marginal increase increases). Also, we computed distributions of mature

phenotypes (e.g. uniform, bimodal, *J*-shaped) produced by the developmental systems, which may provide insights into the patterning of individual differences in adulthood.

Our model thus links evolutionary ecologies to developmental systems, and developmental systems to phenotypic distributions. The model generates several novel insights (Frankenhuis & Panchanathan, 2011a, 2011b). Here, we mention just one: stochastic sampling can lead to individual differences in developmental plasticity itself. Some individuals obtain a homogenous sample of cues, resulting in a confident estimate about the environmental state – they specialize early. Other individuals obtain a heterogeneous, uninformative set of cues – they continue sampling. As a result, in a given developmental period, some individuals (the uncertain ones) may still be sampling while others (the confident ones) are specializing; the former will thus be more susceptible to environmental information – they are more plastic. Accordingly, inter-individual variation in sensitivity to environmental information may result from inter-individual variation in the consistency of earlier experiences.

Adaptive dynamic models have the potential to advance both theory and empirics. Theoretically, our model (and other ones) can inspire evolution-oriented psychologists to explicitly consider developmental dynamics, context sensitivity, and incremental developmental construction. Empirically, our specific result that stochastic sampling could lead to individual differences in plasticity itself will, hopefully, contribute to our understanding of why children differentially benefit or suffer from such experiences as nurturance or abuse (Belsky & Pluess, 2009).

Constructing a dynamic program

This section offers a brief conceptual outline of how dynamic programs may be implemented. Such a framework may be helpful to researchers who want to develop their own models as well as to scholars interested in teaching others. Unfortunately, to our knowledge, no instruction books about the approach exist for a psychological audience (providing another motivation for writing this paper). However, there are many useful sources describing dynamic programming techniques in biology (Mangel & Clark, 1988; Houston & McNamara, 1999; Kokko, 2007; Roff, 2010).

Models are by design simplified, idealized versions of reality. Among their various goals, they may strive to capture some essential components of a process or system. Though some models are built for prediction,

models may serve many other functions. In the case of dynamic optimization, models are often used as deductive aids, forcing us to make assumptions explicit, ensuring logical consistency in argumentation, and suggesting new ways of looking at the world, including normative hypotheses that can be subjected to empirical tests (Maynard Smith, 1978; Schlichting & Pigliucci, 1998; Staddon, Hinson & Mazur, 1983).

Developing a model first requires formulating a problem of interest (e.g. How should developing organisms allocate resources to growth vs. reproduction, or sampling vs. specializing, under tradeoffs?). We use the dimensions of interest to define the state space of the system; that is, all possible states of the system at all times. The developmental system is represented as a vector of variables, and any given state of the system is represented by a unique combination of values of the state variables. The state space can be conceptualized as a multidimensional matrix, with one axis representing time (e.g. the x -axis), and other axes representing the state variables – an N -dimensional matrix situated in time (with N representing the number of state variables).

We begin by determining the (expected) fitness associated with each and every possible end state of development (T_{end} , which can represent maturation, death, or any natural bound given the modeling question). How end states translate into payoffs depends on the fitness function. The fitness function will depend on the question of interest, and it is often worthwhile exploring multiple fitness functions. For instance, in our model of incremental development, we examined how the optimal balance between sampling and specializing varied as a function of the mapping of specialization to fitness (i.e. linear, diminishing, and increasing). In other cases, pre-existing empirical knowledge may determine the chosen fitness function. For instance, if it is known that having some of the trait (e.g. a fur coat) is much better than having none at all (e.g. no fur), and having some almost as good as having a lot of the trait (e.g. a thick fur coat), then the best fitness function will be one that captures marginally diminishing returns. In general, fitness functions include all the state variables, and output one payoff for each end state (e.g. represented as a column of end states, each with an associated payoff).

After computing the payoffs of end states, we move one step back in time (backwards induction) to the penultimate time period (here, $T_{end} - 1$). Although no decisions are made in the final time period – only payoffs calculated – developmental decisions are being made in time period $T_{end} - 1$. We do this by calculating, for each penultimate state, which of the available decisions results in the highest payoff (or ‘expected’ payoff, if outcomes are probabilistically linked to decisions). The next step

makes backwards induction computationally efficient: The (expected) payoff associated with the best decision for a state in time period $T_{end} - 1$ now becomes the (expected) payoff associated with that state in $T_{end} - 1$. This allows us to ignore T_{end} going forward, with all suboptimal branches being cut off. We repeat the same procedure for the previous developmental time period ($T_{end} - 2$), looking forward only to time period $T_{end} - 1$. We continue this process backwards to the beginning of development ($time = 1$), and having arrived there, we have a best decision for each state in the entire state space (i.e. except states in T_{end}). The collection of these decisions is the optimal policy, a manual for optimally navigating the state space.

In addition to optimal decisions, we might store second and third best decisions, and so forth. These suboptimal decisions will be informative if we want to estimate how strong natural selection needs to be – compared with other evolutionary processes such as drift and migration – for the optimal policy to evolve in the messy world of competing evolutionary forces. This is one kind of sensitivity analysis. A second kind of sensitivity analysis is one that explores a wider parameter range; for instance, if we model a system that samples cues to the environmental state, we can explore cues varying from highly to weakly informative. This allows us to estimate whether optimal developmental policy holds across a broad or narrow range of parameter values (a measure of robustness), which will have implications for the extent to which the results generalize across circumstances. Sensitivity analyses generally contribute to a better understanding of system dynamics.

Limitations of dynamic optimization

The goal of optimization models is often to explore how different variables affect each other, sharpening our intuitions about how, at a qualitative level, processes might work in nature – although sometimes, the goal may be to derive quantitative predictions in a specific ecological context (Maynard Smith, 1978). However, the models we presented are intended as ‘computational-level’ models, in Marr’s (1982) sense. They leave open many parameters of how the policies might be instantiated in the physical materials of which bodies and brains are made, constraining only the specific computational properties the models are intended to capture (Bayesians make a similar argument about their models of cognition; Tenenbaum, Griffiths & Kemp, 2006). Optimization models serve as a comparison benchmark for how systems can be expected to behave if they have been optimized (Staddon *et al.*, 1983).

Optimal policies can be thought of as very large contingency tables (or mathematical functions) which specify for each state of the system what decision the organism should make. It is unlikely, however, that natural systems use such look-up tables. Instead, the neural networks (or cognitive mechanisms) embodied in animals, including humans, might use coarse-grained noisy information to make fast judgments and decisions (Gigerenzer, Todd & the ABC Research Group, 1999). This point was nicely captured by McNamara and Houston (2009), two pioneers of state-dependent approaches in biology: ‘If this model really was a description of the world, then the optimal rule could be implemented by looking up the optimal action to adopt in each state. However, the world is not as simple as this or any other model. The number of situations is too vast to expect the optimal decision for every situation to evolve. Instead, it is likely that animals will evolve rules that perform well on average in their natural environment. These rules might be simple and might not be exactly optimal in any situation’ (pp. 670–671).

Like any approach, optimization methods have their limits and caveats (Maynard Smith, 1978). One limitation that applies to optimization models is that they employ what is called ‘the phenotypic gambit’ (Grafen, 1984), an assumption that optimal phenotypes can be produced, without considering developmental, genetic, and other constraints on phenotypic design. In reality, developmental systems may be intimately tied up with other systems, with each system affecting the dynamics of the others (e.g. Finlay, 2007; Finlay, Darlington & Nicastro, 2001; Maynard Smith, Burian, Kauffman, Alberch, Campbell, Goodwin, Lande, Raup & Wolpert, 1985; Ploeger *et al.*, 2008). And, because natural selection only acts on available variation, evolution cannot produce optimal solutions if those variants have never arisen. In cases that involve such constraints on optimality, developmental systems should not be modeled in isolation, but rather as one component of a cohesive package – i.e. a suite of inter-connected characters – in which each system might constrain the other ones (Gould & Lewontin, 1979). Developmental constraints can be integrated in state-dependent models. For instance, we can characterize the developmental system not only in terms of its internal state (e.g. energy levels) and the state of the external environment (e.g. resource distributions), but also in terms of its developmental effects on correlated traits (and vice versa).

There is a second way in which dynamic optimization can teach us about developmental constraints. In practice, it can be quite difficult to identify developmental constraints (Andrews, Gangestad & Matthews, 2002; Maynard Smith *et al.*, 1985). Counter-intuitively, opti-

mization models can facilitate discovery by describing how a system ought to behave, if unconstrained (Olson, 2012). When a system deviates from predictions, this could indicate a developmental constraint (or the model may be mistaken about what selection is maximizing). Moreover, knowing how a system deviates from optimality can provide insight into the factors likely constraining it (e.g. tradeoffs resulting from connections with other systems). Optimization models can thus be used as a comparison benchmark for the detection of constraints. As Nettle, Gibson, Lawson and Sear (2013) note: ‘the use of the term gambit is entirely apt; [it] is a way of opening the enquiry designed to gain some advantage in the quest to understand. It is not the end game (for more detailed discussion, see Fawcett *et al.*, 2013).

We modeled one form of developmental plasticity using dynamic programming, but not all forms of plasticity can be modeled using this approach. Dynamic programming models can explore a vast number of developmental trajectories, but each state existed in the state space before we ran the model, and is thus a predefined possibility. Since this approach optimizes only over the initial state space, there is no room for ‘novel environments’. And, since policies are always found within this pre-specified space, no ‘novel phenotypes’ can emerge.¹¹ This limitation is important, because it implies that dynamic programming may be unsuitable for capturing developmental dynamics as a source of phenotypic novelty, and resultant evolutionary change (Lickliter, 2008; Lickliter & Honeycutt, 2003a, 2003c; West-Eberhard, 2003).¹²

Finally, the method assumes that natural selection has explored the entire state space, and favored the optimal policy. This assumption is not always realistic; the extent to which selection has explored the state space will depend on a host of factors, including the available genetic variation, the time selection has had to find solutions, and the shape of the adaptive landscape (i.e. adaptive valleys might prevent the system from reaching an optimal peak).

Although the above features are genuine limitations, dynamic optimization models are well suited for analyz-

¹¹ Note: It is possible to compute an optimal policy in one environment and expose it to another environment (e.g. in a simulation) in order to observe how it fares.

¹² These limitations do not necessarily apply to models that use open-ended learning rules, including models of associative or reinforcement learning; such models might be more suitable for studying phenotypic development in evolutionarily novel environments, and the ontogeny of novel phenotypes (Montague & King-Casas, 2007; Sutton & Barto, 1998).

ing many important questions at the intersection between evolutionary and developmental science. Therefore, in our view, the method deserves to be incorporated into the toolkit of scholars working in this area.

Concluding remarks

This article has been written with two goals in mind. Our main goal is to introduce dynamic optimization into developmental psychology. Dynamic programming is increasingly being used in many fields of science, and may prove useful to developmental psychologists as well. Second, interactions between developmental systems and evolutionary psychological approaches have sometimes, unfortunately, been antagonistic. We have outlined one way in which the gap can be bridged.

Dynamic optimization offers one way of integrating developmental systems theory's focus on dynamics and contingency with the 'design stance' of evolutionary psychology. It provides tools for exploring the properties of developmental systems that natural selection might favor, given particular evolutionary ecologies. The approach allows for stochastic outcomes of decisions, and can incorporate uncertainty as well as learning about parameter values. Optimal policies may be found even in a large developmental state space using backwards induction, which involves (1) breaking the decision problem into temporal sub-problems, (2) finding optimal solutions for the last sub-problem and recording them, (3) taking one time step back, and (4) repeating the procedure to the onset of ontogeny. Because backwards induction allows pruning of suboptimal branches, it is computationally more efficient than an exhaustive forward search.

We have focused on adaptive problems spanning multiple years (e.g. growth–reproduction tradeoffs, sampling vs. specializing), but learning and development also occur on shorter timescales (e.g. minutes or days), and here dynamic optimization methods can be applied as well. For instance, infants face the challenge of learning which objects in the environment are social agents, and responsive to them (and to what extent), and which ones are not (Frankenhuis, Gergely & Watson, 2013). Research shows that infants actively contribute to this learning process by producing behaviors (e.g. smiling and cooing) that elicit responses from social agents around them (Movellan & Watson, 2002). This form of 'active learning' requires temporal spacing of behavioral units for other agents to be able to respond and to contrast this response with background noise. A recent analysis used adaptive dynamic optimization to examine whether, at different stages of development, infants' vocalization

patterns effectively enabled uncertainty reduction about the question whether a given entity is an agent (Butko & Movellan, 2010). This analysis highlights the exciting potential for integrating evolution and learning in the psychological sciences.

We hope our paper stimulates productive interactions between evolutionary psychologists and developmental systems theorists. Human development is highly flexible, depending on many state-dependent processes. Therefore, there should be great scope for profitable integration.

Acknowledgements

This research was supported by the Hungarian Science Foundation (OTKA NK 83997). We thank Rob Boyd, David Frankenhuis, Marco del Giudice, Ron Dotsch, Kotrina Kajokaite, Robert Lickliter, Daniel Nettle, Annemie Ploeger, Simone Schnall, and three anonymous reviewers and the editor Denis Mareschal, for their insightful comments.

References

- Agrawal, A.A., Laforsch, C., & Tollrian, R. (1999). Transgenerational induction of defences in animals and plants. *Nature*, **401**, 60–63.
- Andrews, P.W., Gangestad, S.W., & Matthews, D. (2002). Adaptationism – how to carry out an exaptationist program. *Behavioral and Brain Sciences*, **25**, 489–553.
- Badcock, P.B. (2012). Evolutionary systems theory: a unifying meta-theory of psychological science. *Review of General Psychology*, **16**, 10–23.
- Barker, D. (1994). *Mothers, babies, and disease in later life*. London: BMJ.
- Barrett, H.C. (2006). Modularity and design reincarnation. In P. Carruthers, S. Laurence & S. Stich (Eds.), *The innate mind: Culture and cognition* (pp. 199–217). New York: Oxford University Press.
- Barrett, H.C. (2007). Development as the target of evolution: a computational approach to developmental systems. In S. Gangestad & J. Simpson (Eds.), *The evolution of mind: Fundamental questions and controversies* (pp. 186–192). New York: Guilford.
- Bellman, R.E. (1961). *Adaptive control processes: A guided tour*. Princeton, NJ: Princeton University Press.
- Belsky, J., & Pluess, M. (2009). Beyond diathesis stress: differential susceptibility to environmental influences. *Psychological Bulletin*, **135**, 885–908.
- Belsky, J., Steinberg, L., & Draper, P. (1991). Childhood experience, interpersonal development, and reproductive strategy: an evolutionary theory of socialization. *Child Development*, **62**, 647–670.

- Bjorklund, D.F. (2003). Evolutionary psychology from a developmental systems perspective: comment on Lickliter and Honeycutt (2003). *Psychological Bulletin*, **129**, 836–841.
- Bjorklund, D.F., Ellis, B.J., & Rosenberg, J.S. (2007). Evolved probabilistic cognitive mechanisms: an evolutionary approach to gene \times environment \times development. In R.V. Kail (Ed.), *Advances in child development* (pp. 1–39). Oxford: Elsevier.
- Brezina, T., Tekin, E., & Topalli, V. (2009). ‘Might not be a tomorrow’: a multimethods approach to anticipated early death and youth crime. *Criminology*, **47**, 1091–1129.
- Busemeyer, J.R., & Pleskac, T. (2009). Theoretical tools for understanding and aiding dynamic decision making. *Journal of Mathematical Psychology*, **53**, 126–138.
- Buss, D.M., & Reeve, H.K. (2003). Evolutionary psychology and developmental dynamics: comment on Lickliter and Honeycutt (2003). *Psychological Bulletin*, **129**, 848–853.
- Butko, N.J., & Movellan, J.R. (2010). Detecting contingencies: an infomax approach. *Neural Networks*, **23**, 973–984.
- Chhabra, M., & Jacobs, R. (2006). Near-optimal human adaptive control across different noise environments. *Journal of Neuroscience*, **26**, 10883–10887.
- Childs, D.Z., Metcalf, C.J.E., & Rees, M. (2010). Evolutionary bet-hedging in the real world: empirical evidence and challenges revealed by plants. *Proceedings of the Royal Society B*, **277**, 3055–3064.
- Chisholm, J.S. (1993). Death, hope, and sex: life-history theory and the development of reproductive strategies. *Current Anthropology*, **34**, 1–24.
- Crawford, C.B. (2003). A prolegomenon for a viable evolutionary psychology: the myth and the reality. Comment on Lickliter and Honeycutt (2003). *Psychological Bulletin*, **129**, 854–857.
- De Beer, R.D. (2000). Dynamical approaches to cognitive science. *Trends in Cognitive Sciences*, **4**, 91–99.
- Dempster, E.R. (1955). Maintenance of genetic heterogeneity. *Cold Spring Harbor Symposium*, **20**, 25–32.
- Dennett, D.C. (2011). Shall we tango? No, but thanks for asking. *Journal of Consciousness Studies*, **18**, 23–34.
- Drugowitsch, J., Moreno-Bote, R., Churchland, A.K., Shadlen, M.N., & Pouget, A. (2012). The cost of accumulating evidence in perceptual decision making. *Journal of Neuroscience*, **32**, 3612–3628.
- Ellis, B.J. (2004). Timing of pubertal maturation in girls: an integrated life history approach. *Psychological Bulletin*, **130**, 920–958.
- Ellis, B.J., Figueredo, A.J., Brumbach, B.H., & Schlomer, G.L. (2009). Fundamental dimensions of environmental risk: the impact of harsh versus unpredictable environments on the evolution and development of life history strategies. *Human Nature*, **20**, 204–268.
- Engelbrecht, S.E., Berthier, N.E., & O’Sullivan, L.P. (2003). The undershoot bias: learning to act optimally under uncertainty. *Psychological Science*, **14**, 257–261.
- Fawcett, T.W., Hamblin, S., & Giraldeau, L.-A. (2013). Exposing the behavioral gambit: the evolution of learning and decision rules. *Behavioral Ecology*, **24**, 2–11.
- Finlay, B.L. (2007). Endless minds, most beautiful. *Developmental Science*, **10**, 30–34.
- Finlay, B.L., Darlington, R.D., & Nicastro, N. (2001). Developmental structure in brain evolution. *Behavioral and Brain Sciences*, **24**, 263–308.
- Flynn, E.G., Laland, K.N., Kendal, R.L., & Kendal, J.R. (2013). Developmental niche construction. *Developmental Science*, **16**, 296–313.
- Frank, S.A. (1996). The design of natural and artificial adaptive systems. In M.R. Rose & G.V. Lauder (Eds.), *Adaptation* (pp. 451–505). New York: Academic Press.
- Frankenhuys, W.E., & Del Giudice, M. (2012). When do adaptive developmental mechanisms yield maladaptive outcomes? *Developmental Psychology*, **48**, 628–642.
- Frankenhuys, W.E., Gergely, G., & Watson, J.S. (2013). Infants may use contingency analysis to estimate environmental states: an evolutionary, life-history perspective. *Child Development Perspectives*. doi:10.1111/cdep.12024.
- Frankenhuys, W.E., & Panchanathan, K. (2011a). Balancing sampling and specialization: an adaptationist model of incremental development. *Proceedings of the Royal Society B*, **278**, 3558–3565.
- Frankenhuys, W.E., & Panchanathan, K. (2011b). Individual differences in developmental plasticity may result from stochastic sampling. *Perspectives on Psychological Science*, **6**, 336–347.
- Frankenhuys, W.E., Panchanathan, K., & Belsky, J. (in preparation). Can variation in developmental plasticity evolve because parents hedge their bets? A mathematical model.
- Gardner, A. (2009). Adaptation as organism design. *Biology Letters*, **5**, 861–864.
- Geary, D.C. (2006). Evolutionary developmental psychology: current status and future directions. *Developmental Review*, **26**, 113–119.
- Geary, D.C., & Bjorklund, D.F. (2000). Evolutionary developmental psychology. *Child Development*, **71**, 57–65.
- Gigerenzer, G., & Todd, P.M., & the ABC Research Group (Eds.) (1999). *Simple heuristics that make us smart*. New York: Oxford University Press.
- 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. *Proceedings of the Royal Society B*, **272**, 671–677.
- Gottlieb, G. (1991). Experiential canalization of behavioral development: theory. *Developmental Psychology*, **27**, 4–13.
- Gould, S.J., & Lewontin, R.C. (1979). The spandrels of San Marco and the Panglossian paradigm: a critique of the adaptationist programme. *Proceedings of the Royal Society London B*, **205**, 581–598.
- Grafen, A. (1984). Natural selection, kin selection and group selection. In J.R. Krebs & N.B. Davies (Eds.), *Behavioural ecology: An evolutionary approach* (2nd edn., pp. 62–84). Oxford: Blackwell Scientific Publications.
- Grafen, A. (2007). The formal Darwinism project: a mid-term report. *Journal of Evolutionary Biology*, **20**, 1243–1254.

- Greenough, W.T., Black, J.E., & Wallace, C.S. (1987). Experience and brain development. *Child Development*, **58**, 539–559.
- Griesemer, J., Haber, M.H., Yamashita, G., & Gannett, L. (2005). Critical notice: Cycles of contingency: developmental systems and evolution. *Biology and Philosophy*, **20**, 517–544.
- Griffiths, P.E., & Gray, R.D. (2005). Discussion: Three ways to misunderstand developmental systems theory. *Biology and Philosophy*, **20**, 417–425.
- Griskevicius, V., Delton, A.W., Robertson, T.E., & Tybur, J.M. (2011a). The environmental contingency of life history strategies: influences of mortality and socioeconomic status on reproductive timing. *Journal of Personality and Social Psychology*, **100**, 241–254.
- Griskevicius, V., Tybur, J.M., Delton, A.W., & Robertson, T.E. (2011b). The influence of mortality and socioeconomic status on risk and delayed rewards: a life history theory approach. *Journal of Personality and Social Psychology*, **100**, 1015–1026.
- Helle, S. (2008). A tradeoff between reproduction and growth in contemporary Finnish women. *Evolution and Human Behavior*, **29**, 189–195.
- Houston, A., Clark, C., McNamara, J., & Mangel, M. (1988). Dynamic models in behavioural and evolutionary ecology. *Nature*, **332**, 29–34.
- Houston, A.I., & McNamara, J.M. (1992). Phenotypic plasticity as a state-dependent life-history decision. *Evolutionary Ecology*, **6**, 243–253.
- Houston, A.I., & McNamara, J.M. (1999). *Models of adaptive behavior: An approach based on state*. Cambridge: Cambridge University Press.
- Hutchinson, J.M.C., & McNamara, J.M. (2000). Ways to test stochastic dynamic programming models empirically. *Animal Behaviour*, **59**, 665–676.
- Johnston, T.D., & Turvey, M.T. (1980). A sketch of an ecological metatheory for theories of learning. In G.H. Bower (Ed.), *The psychology of learning and memory* (pp. 148–207). New York: Academic Press.
- Kokko, H. (2007). *Modelling for field biologists and other interesting people*. Cambridge: Cambridge University Press.
- Körding, K.P., & Wolpert, D.M. (2006). Bayesian decision theory in sensorimotor control. *Trends in Cognitive Sciences*, **10**, 319–326.
- Krebs, D.L. (2003). Fictions and facts about evolutionary approaches to human behavior: comment on Lickliter and Honeycutt (2003). *Psychological Bulletin*, **129**, 842–847.
- Laland, K.N., & Sterelny, K. (2006). Seven reasons (not) to neglect niche construction. *Evolution*, **60**, 1751–1762.
- Laland, K.N., Sterelny, K., Odling-Smee, J., Hoppitt, W., & Uller, T. (2011). Cause and effect in biology revisited: is Mayr's proximate–ultimate dichotomy still useful? *Science*, **334**, 1512–1516.
- Lickliter, R. (2008). The growth of developmental thought: implications for a new evolutionary psychology. *New Ideas in Psychology*, **26**, 353–369.
- Lickliter, R., & Honeycutt, H. (2003a). Developmental dynamics: towards a biologically plausible evolutionary psychology. *Psychological Bulletin*, **129**, 819–835.
- Lickliter, R., & Honeycutt, H. (2003b). Developmental dynamics and contemporary evolutionary dynamics: status quo or irreconcilable views? Response to Bjorklund (2003), Krebs (2003), Buss and Reeve (2003), Crawford (2003), and Tooby et al. (2003). *Psychological Bulletin*, **129**, 866–872.
- Lickliter, R., & Honeycutt, H. (2003c). Evolutionary approaches to cognitive development: status and strategy. *Journal of Cognition and Development*, **4**, 459–473.
- McNamara, J.M., Green, R.F., & Olssen, O. (2006). Bayes' theorem and its applications in animal behavior. *Oikos*, **112**, 243–251.
- McNamara, J.M., & Houston, A.J. (2009). Integrating function and mechanism. *Trends in Ecology and Evolution*, **24**, 670–675.
- Mangel, M., & Clark, C.W. (1988). *Dynamic modeling in behavioral ecology*. Princeton, NJ: Princeton University Press.
- Mangel, M., & Ludwig, D. (1992). Definition and evaluation of the fitness of behavioral and developmental programs. *Annual Review of Ecology, Evolution, and Systematics*, **23**, 507–536.
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco, CA: Freeman.
- Maynard Smith, J.T. (1978). Optimization theory in evolution. *Annual Review of Ecology and Systematics*, **9**, 31–56.
- Maynard Smith, J.T. (1982). *Evolution and the theory of games*. Cambridge: Cambridge University Press.
- Maynard Smith, J., Burian, R., Kauffman, S., Alberch, P., Campbell, J., Goodwin, B., Lande, R., Raup, D., & Wolpert, L. (1985). Developmental constraints and evolution: a perspective from the Mountain Lake Conference on Development and Evolution. *Quarterly Review of Biology*, **60**, 265–287.
- Mayr, E. (1983). How to carry out the adaptationist program? *American Naturalist*, **121**, 324–334.
- Mealey, L. (1995). The sociobiology of sociopathy: an integrated evolutionary model. *Behavioral and Brain Sciences*, **18**, 523–541.
- Meyers, L.A., & Bull, J.J. (2002). Fighting change with change: adaptive variation in an uncertain world. *Trends in Ecology and Evolution*, **17**, 551–557.
- Monaghan, P. (2008). Early growth conditions, phenotypic development and environmental change. *Proceedings of the Royal Society B*, **363**, 1635–1645.
- Montague, P.R., & King-Casas, B. (2007). Efficient statistics, common currencies and the problem of reward-harvesting. *Trends in Cognitive Sciences*, **11**, 514–519.
- Moodie, E.E.E.M., Richardson, T.S., & Stephens, D.A. (2007). Demystifying optimal dynamic treatment regimes. *Biometrics*, **63**, 447–455.
- Moran, N.A. (1992). The evolutionary maintenance of alternative phenotypes. *American Naturalist*, **139**, 971–989.
- Movellan, J.R., & Watson, J.S. (2002). The development of gaze following as a Bayesian systems identification problem. *Proceedings of the 2nd international conference on development and learning* (pp. 34–40). Los Alamitos, CA: IEEE Computer Society.

- Murphy, S. (2003). Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society B*, **65**, 331–366.
- Nettle, D., Gibson, M.A., Lawson, D.W., & Sear, R. (2013). Human behavioral ecology: current research and future prospects. *Behavioral Ecology*. doi:10.1093/beheco/ars222
- Olson, M.E. (2012). The developmental renaissance in adaptationism. *Trends in Ecology and Evolution*, **27**, 278–287.
- Osman, M. (2010). Controlling uncertainty: a review of human behavior in complex dynamic environments. *Psychological Bulletin*, **136**, 65–86.
- Oyama, S., Griffiths, P.E., & Gray, R.D. (Eds.) (2001). *Cycles of contingency: Developmental systems and evolution*. Cambridge, MA: MIT Press.
- Padilla, D.K., & Adolph, S.C. (1996). Plastic inducible morphologies are not always adaptive: the importance of time delays in a stochastic environment. *Evolutionary Ecology*, **10**, 105–117.
- Philippi, T., & Seger, J. (1989). Hedging one's evolutionary bets, revisited. *Trends in Ecology and Evolution*, **4**, 41–44.
- 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. *Psychological Inquiry*, **19**, 1–18.
- Pollak, S.D. (2005). Early adversity and mechanisms of plasticity: integrating affective neuroscience with developmental approaches to psychopathology. *Development and Psychopathology*, **17**, 735–752.
- Pollak, S.D. (2008). Mechanisms linking early experience and the emergence of emotions. *Current Directions in Psychological Science*, **17**, 370–375.
- Robert, J.S., Hall, B.K., & Olson, W.M. (2001). Bridging the gap between developmental systems theory and evolutionary developmental biology. *BioEssays*, **23**, 954–962.
- Roff, D.A. (2010). *Modeling evolution: An introduction to numerical methods*. Oxford: Oxford University Press.
- Rossiter, M. (1996). Incidence and consequences of inherited environmental effects. *Annual Review of Ecology and Systematics*, **27**, 451–476.
- Schlichting, C.D., & Pigliucci, M. (1998). *Phenotypic evolution: A reaction norm perspective*. Sunderland: Sinauer Associates.
- Smith, L.B., & Thelen, L. (2003). Development as a dynamic system. *Trends in Cognitive Sciences*, **7**, 343–348.
- Staddon, J.E.R., Hinson, J.M., & Mazur, J.E. (1983). Optimization: a result or a mechanism? *Science*, **221**, 976–977.
- Sterelny, K., & Griffiths, P.E. (1999). *Sex and death: An introduction of philosophy of biology*. Chicago, IL: The University of Chicago Press.
- Sterelny, K., Smith, K., & Dickison, M. (1996). The extended replicator. *Biology and Philosophy*, **11**, 377–403.
- Sutton, R.S., & Barto, A.G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- Tenenbaum, J.B., Griffiths, T.L., & Kemp, C. (2006). Theory-based Bayesian models of inductive learning and reasoning. *Trends in Cognitive Sciences*, **10**, 309–318.
- Todorov, E. (2004). Optimality principles in sensorimotor control. *Nature Neuroscience*, **7**, 907–915.
- Tooby, J., & Cosmides, L. (1992). The psychological foundations of culture. In J. Barkow, L. Cosmides & J. Tooby (Eds.), *The adapted mind: Evolutionary psychology and the generation of culture* (pp. 19–136). New York: Oxford University Press.
- Tooby, J., Cosmides, L., & Barrett, H.C. (2003). The second law of thermodynamics is the first law of psychology: evolutionary developmental psychology and the theory of tandem, coordinated inheritances: comment on Lickliter and Honeycutt (2003). *Psychological Bulletin*, **129**, 858–865.
- van der Maas, H.L.J. (1995). Beyond the metaphor? *Cognitive Development*, **10**, 621–642.
- Van Valen, L. (1973). Festschrift. *Science*, **180**, 488.
- Walker, R.S., Gurven, M., Hill, K., Migliano, A., Chagnon, N., De Souza, R., Djurovic, G., Hames, R., Hurtado, A.M., Kaplan, H., Kramer, K., Oliver, W.J., Vallengia, C., & Yamauchi, T. (2006). Growth rates and life histories in twenty-two small-scale societies. *American Journal of Human Biology*, **18**, 295–311.
- West, M.J., & King, A.P. (1987). Settling nature and nurture into an ontogenetic niche. *Developmental Psychobiology*, **20**, 549–562.
- West-Eberhard, M.J. (2003). *Developmental plasticity and evolution*. New York: Oxford University Press.
- Williams, G.C. (1966). *Adaptation and natural selection*. Princeton NJ: Princeton University Press.

Received: 17 July 2012

Accepted: 8 January 2013

Copyright of Developmental Science is the property of Wiley-Blackwell and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.